vRealize Operations Manager Load Balancing

vRealize Operations 8.4

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

Contents

Introduction

<div style="text-align: right; font-size: 2em;">1</div>

This document describes the configuration of the load balancing modules of F5 Networks BIG-IP software (F5), Citrix NetScaler, HAProxy and NSX load balancers for vRealize Operations Manager.

This document is not an installation guide, but a load-balancing configuration guide that supplements the vRealize Operations Manager installation and configuration documentation available in the vRealize Operations Manager Documentation Center.

This information is for the following products and versions.

Product	Version	Documentation
vRealize Operations Manager	6.6.1, 6.7, 7.x, 8.x	https://docs.vmware.com/en/vRealize-Operations-Manager/index.html
F5 BIG-IPLTM	11.5, 11.6, 12.1, 13.0, 14.x, 15.x, 16.x	https://support.f5.com/csp/knowledge-center/software/BIG-IP?module=BIG-IP%20LTM
F5 BIG-IP GTM**	15.x, 16.x	https://support.f5.com/csp/knowledge-center/software/BIG-IP?module=BIG-IP%20GTM
Citrix NetScaler	10.5*, 11.0*, 11.x, 12.x, 13.x	https://www.citrix.com/products/netscaler-adc/
NSX-V	6.1.3, 6.2.x, 6.3.x, 6.4.x	https://pubs.vmware.com/NSX-6/index.jsp#Welcome/welcome.html
NSX-T	2.2.x, 2.3.x, 2.4.x, 2.5.x, 3.x.x	https://docs.vmware.com/en/VMware-NSX-T/index.html
HA Proxy	v1.5.x, 1.8.x	http://www.haproxy.org/
RHEL	v7.x, v8.x	https://access.redhat.com/documentation/
Keepalived	v1.3.x	http://www.keepalived.org/

* Citrix NetScaler VPX versions prior to 11.0 65.35 have a bug which prevents them from using TLS 1.1/1.2. For more information, please refer to the NetScaler section of this document.

** F5 BIG-IP GTM is supported only for use with vRealize Operations Continuous Availability feature and could not be considered as a replacement for BIG-IP LTM

This chapter includes the following topics:

■ Load Balancing Concepts

- ■ [vRealize Operations Manager Overview](#)

Load Balancing Concepts

Loadbalancers distribute connections among servers in high availability (HA) deployments.

Following are the advantages of using a load balancer in front of the vRealize Operations Manager cluster:

- ■ Utilizing a load balancer ensures that the deployed cluster is properly balanced for performance of UI traffic.

- ■ Allows all nodes in the cluster to equally participate in the handling of UI sessions and traffic.

- ■ Provides high availability if any admin or data node fails, by directing UI traffic only to serving nodes in the cluster.

- ■ Provides simpler access for the users. Instead of accessing each node individually the user only needs one URL to access the entire cluster and not be concerned with which node is available.

- ■ Provides load balancing, high availability and ease of configuration for the End Point Operations (EPOps) agents.

The system administrator backs up the loadbalancers on a regular basis at the same time as other components.

Follow your site policy for backing up loadbalancers, keeping in mind the preservation of network topology and vRealize Operations Manager backup planning.

Selecting a Load Balancer

There are no specific requirements for selecting a load balancer platform for vRealize Operations Manager. Majority of Load Balancers available today support complex web servers and SSL.

Load balancer can be used in front of a vRealize Operations Manager cluster if certain parameters and configuration variables are followed. HAProxy was chosen for this example due to its ease of deployment, open source availability, stability, capability handling SSL sessions, and performance. Following are some of the parameters that should be considered for configuring other brands of load balancers:

- ■ You must use TCP Mode. HTTP mode is not supported.

- ■ It is not recommended to use round-robin balancing mode

- ■ Cookie persistence does not work

- ■ SSL pass-through is used, SSL termination is not supported

- ■ IP Hash type balancing is recommended to ensure that the same client IP address always reaches the same node, if the node is available

- ■ Health checks should be performed with public API provided by vRealize Operations Manager.

How to Handle SSL UI Certificates with a Load Balancer

In all the default installations of vRealize Operations Manager nodes a default self-signed VMware certificate is included. You can implement your own SSL certificate from an internal Certificate Authority or external Certificate Authority.

For more information on the certificate installation procedures, see Requirements for Custom vRealize Operations Manager SSL Certificates.

In addition to these configuration variables it is important to understand how SSL certificates are distributed in a cluster. If you upload a certificate to a node in the cluster, for example: the master node, the certificate will then be pushed to all nodes in the cluster. To handle UI sessions by all the nodes in the cluster you must upload an SSL certificate that contains all the DNS names (optional: IP addresses and DNS names) in the **Subject Alternative Name** field of the uploaded certificate. The common name should be the Load Balancer DNS name. The subject alternative names are used to support access to the admin UI page.

When the certificate is uploaded trough admin UI page it is pushed to all the nodes in the cluster. Currently, when you use a load balancer with vRealize Operations Manager, the only supported method is SSL pass-through, which means the SSL certificate cannot be terminated on the load balancer.

To change SSL certificate on a cluster deployment:

Procedure

1 Log in to the master node by using the following link: https://<ipaddress>/admin.

2 On the top right side, click the certificate button to change the certificate.

3 Click on Install New Certificate

4 Click on Browse button and choose PEM certificate file.

5 After certificate verification click Install.

Results

When you view the certificate on the node that you are accessing, you will see all nodes in the cluster listed in the certificate SAN.

vRealize Operations Manager Overview

The vRealize Operations Manager clusters consist of a master node, an optional replica node for high availability, optional data nodes, and optional remote collector nodes.

You can access and interact with the product by using the product UI available on the master and data nodes. The remote collector nodes do not contain a product UI and are used for data collection only. The product UI is powered by a Tomcat instance that resides across each node but is not load balanced out of the box. You can scale up vRealize Operations Manager environment by adding nodes when the environment grows larger.

vRealize Operations Manager supports high availability by enabling a replica node for the vRealize Operations Manager master node. A high availability replica node can take over the functions that a master node provides. When a problem occurs with the master node, fail-over to the replica node is automatic and requires only 2 to 3 minutes of vRealize Operations Manager downtime. Data stored on the master node is always backed up on the replica node. In addition, with high availability enabled, the cluster can survive the loss of a data node without losing any data.

Node Role	Functions
Master Node	It is the initial, required node in the cluster. All other nodes are managed by the master node. It contains the product UI.
	In a single-node installation, the master node performs data collection and analysis as it is the only node where vRealize Operations Manager adapters are installed.
Data Node	In larger deployments, only data nodes have adapters installed to perform collection and analysis. It contains the product UI.
Replica Node	To enable high availability, the cluster requires that you convert a data node in to a replica of the master node. It does not contain product UI.

vRealize Operations Manager Architecture

Information about vRealize Operations Manager maximum supported nodes

Information about vRealize Operations Manager maximum supported nodes in analytics cluster as well as other information related to High Availability can be found in the sizing guideline document.

Remote collectors are not considered part of the analytics clusters as they do not participate in any type of data calculations or processing. EPOps traffic is load balanced to the same cluster.

Note The load balancer cannot decrypt the traffic, hence cannot differentiate between EPOps and analytics traffic.

Following is a basic architecture overview of a vRealize Operations Manager 8-node cluster with high availability enabled.
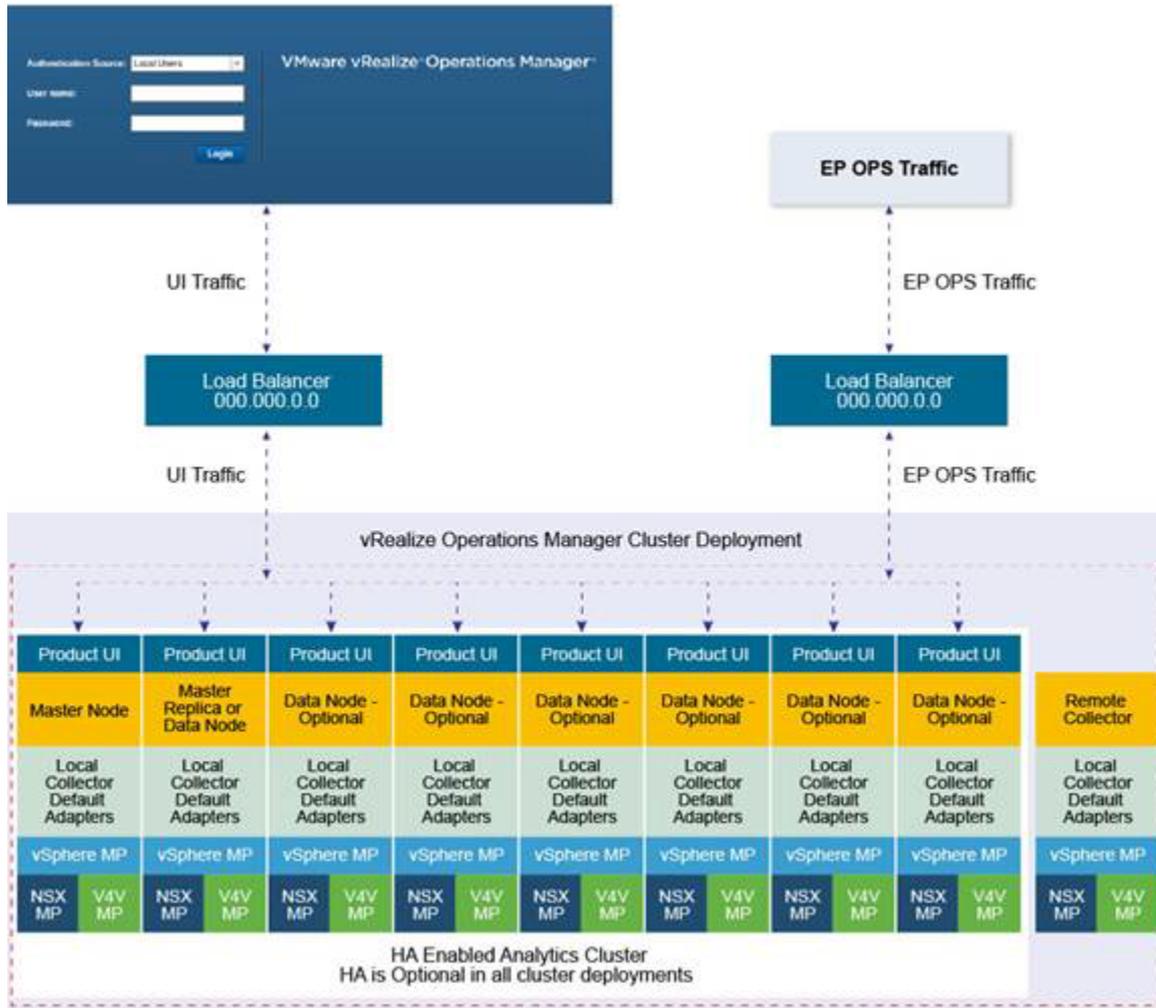
Figure 1. vRealize Operations Manager 8-Nodes Cluster with High Availability

Configuring End Point Operations Agents

End Point Operations agents are used to gather operating system metrics to monitor availability of remote platforms and applications.

End Point Operations agents collected metrics are sent to the vRealize Operations Manager server. You can configure additional load balancer profile or dedicated load balancer to separate analytics traffic from EPOps traffic.

The steps to configure EPOps load balancer are described as required throughout this document.

You must shut down that the load balancer while upgrading or shutting down vRealize Operations Manager cluster. The load balancer should be restarted after the cluster is upgraded.

In the case of EPOps balancing, the overall latency between agent, load balancer, and cluster should be lower than 20 milliseconds. If the latency is higher, you must install a remote collector and direct the agents directly to it.

HAProxy Installation and Configuration

2

HAProxy offers high availability, load balancing, and proxying for TCP and HTTP-based applications. Both multi-arm and one-arm configurations are tested and supported.

Prerequisites

Following are the prerequisites to ensure a functional load balancer configuration and deployment.

- OS: Red Hat Enterprise Linux (RHEL) v7.x or v8.x

- CPU: 2 vCPU

- Memory: 4GB

- Disk space: 50GB

- HAProxy 1.5.x for RHEL 7.x or v1.8.x for RHEL 8.x

- Fully functioning DNS with both forward and reverse lookups

- All nodes in the vRealize Operations Manager cluster operating correctly

- HAProxy deployed in same datacenter and preferably on the same cluster as vRealize Operations Manager

- HAProxy not deployed on the same ESX hosts as vRealize Operations Manager cluster to ensure availability

- Minimum 2-node deployment of vRealize Operations Manager cluster

- Deployment does not require high availability to be enabled, but it is recommended that you enable high availability

- One master node and at least one data node is required for using a load balancer beneficially

This chapter includes the following topics:

- Install Single-Node HAProxy

- Configure Logging for HAProxy

- Configure HAProxy

- Advanced Configuration: HAProxy with Keepalived

Install Single-Node HAProxy

HAProxy installation guide.

HAProxy installation is supported and tested on Red Hat Enterprise Linux (RHEL) 7.x or 8.x and can be obtained from the official Red Hat repository. You can install HAProxy on RHEL 7.x or 8.x by using yum package manager. To configure HAProxy as a load-balancer for vRealize Operations Manager please follow the steps below:

Procedure

1   Perform a package update on system to ensure all packages are up-to-date:

```
yum update
```

2   Install HAProxy:

```
yum -y install haproxy
```

3   Copy original HAProxy configuration to backup file:

```
cp /etc/haproxy/haproxy.cfg /etc/haproxy/haproxy.cfg.bak
```

4   Configure HAProxy configuration. To configure analytics balancer, seeConfigure HAProxy Analyticsand to configure EPOps balancer, seeConfigure EPOps HAProxy.

5   Allow firewall traffic through for the ports needed for HAProxy to function:

```
firewall-cmd --permanent --zone=public --add-port=80/tcp
firewall-cmd --permanent --zone=public --add-port=9090/tcp
firewall-cmd --permanent --zone=public --add-port=443/tcp
```

6   Reload the firewall configuration:

```
systemctl reload firewalld
```

7   Enable HAProxy to connect to any interface:

```
setsebool -P haproxy_connect_any 1
```

8   Enable HAProxy service:

```
systemctl enable haproxy
```

Results

Configure Logging for HAProxy

An administrator might want to configure logging of the HAProxy service to aid in monitoring and troubleshooting an environment.

The HAProxy logger allows for the use rsyslog internally on the Linux installation to log to a local file. You can also utilize vRealize Log Insight integration to send this log to a vRealize Log Insight deployment by utilizing the new Log Insight Linux agent to greatly simplify the configuration and logging of Linux platforms. To configure basic applications logging using rsyslog locally on the server perform the following steps.

Procedure

1 Configure the rsyslog configuration file to accept UDP syslog reception:

```
vi /etc/rsyslog.conf
```

2 Uncomment the following lines:

```
# Provides UDP syslog reception
$ModLoad imudp
$UDPServerAddress 127.0.0.1
$UDPServerRun 514
```

3 Save the file:

```
wq!
```

4 Create the HAProxy logging configuration file for specific application parameters

```
vi /etc/rsyslog.d/haproxy.conf
```

5 Add the following line:

```
if ($programname == 'haproxy') then -/var/log/haproxy.log
```

6 Save the file:

```
wq!
```

7 Create HAProxy Log file and set proper permissions:

```
touch /var/log/haproxy.log
chmod 755 /var/log/haproxy.log
```

8 Restart the rsyslog service:

```
Service rsyslog restart
```

Results

Configure HAProxy

The HAProxy configuration has been tested against an 8-node vRealize Operations Manager cluster.

Clusters with fewer nodes up to a maximum of 16 analytics nodes are also supported and require the same configuration. Every time the cluster is expanded, and a new node is deployed you must edit the HAProxy configuration and add the IP address of the new node. After editing the configuration file, the HAProxy service should always be restarted so the configuration is reloaded. We recommended to set HAProxy global max. connections parameter (2000) and node max. connections parameter (140) which covers most of the cases. However, we strongly suggested to check the sizing of your environment and adjust those settings based on vROps load.

Configure HAProxy for vRealize Operations Manager Analytics

HAProxy for vRealize Operations Manager analytics confuguration guide

You can configure the HAProxy for vRealize Operations Manager analytics as follows:

```
# Configuration file to balance both web and epops
#global parameters global
    log         127.0.0.1 local2
    chroot      /var/lib/haproxy
    pidfile     /var/run/haproxy.pid
    maxconn     2000
    user        haproxy
    group       haproxy
    daemon
    stats socket /var/lib/haproxy/stats
    ssl-server-verify none
#default parameters unless otherwise specified defaults
    log global
    mode http
    option httplog
    option tcplog
    option dontlognull
    timeout connect 5000ms
    timeout client 50000ms
    timeout server 50000ms
#listener settings for stats webpage can be optional but highly recommended  listen stats
:9090
    balance
    mode http
    stats enable
    stats auth admin:admin
    stats uri /
    stats realm Haproxy\ Statistics
#automatic redirect for http to https connections
frontend vrops_unsecured_redirect *:80
        redirect location https://<insert_fqdn_address_here>
#front settings in this case we bind to all addresses on system or specify an interface
    frontend vrops_frontend_secure
        bind <web dedicated ip>:443
        mode tcp
        option tcplog
        default_backend vrops_backend_secure
#backend configuration of receiving servers containing tcp-checks health checks and hashing
#needed for a proper configuration and page sessions
```

```
#adjust the server parameters to your environment
    backend vrops_backend_secure
        mode tcp
        option tcplog
    balance source
    hash-type consistent
    option tcp-check
    tcp-check connect port 443 ssl
    tcp-check send GET\ /suite-api/api/deployment/node/status?
services=api&services=adminui&services=ui\ HTTP/1.0\r\n\r\n
## For older versions of vROPS from 6.6.1 to 7.5 please use the following "tcp-check"
# tcp-check send GET\ /suite-api/api/deployment/node/status\ HTTP/1.0\r\n\r\n
tcp-check expect rstring ONLINE
server node1 <Insert node1 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node2 <Insert node2 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node3 <Insert node3 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node4 <Insert node4 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
```

Note Please make sure to use proper tcp-check call in above instruction. Starting from vROps 8.0 status API enhanced to track separate services status. Old "tcp-check" call provided above in comments.

Configure EPOps HAProxy

EPOps via HAProxy configuration guide

You can configure EPOps HAProxy as follows:

```
# EPOPS Load Balancer configuration.
#global parameters
global
    log         127.0.0.1 local2
    chroot      /var/lib/haproxy
    pidfile     /var/run/haproxy.pid
    maxconn     2000
    user        haproxy
    group       haproxy
    daemon
    stats socket /var/lib/haproxy/stats
    ssl-server-verify none
#default parameters unless otherwise specified
defaults
    log global
    mode http
    option httplog
    option tcplog
    option dontlognull
    timeout connect 5000ms
    timeout client 50000ms
    timeout server 50000ms
```

```
#listener settings for stats webpage can be optional but highly recommended
    listen stats :9090
    balance
    mode http
    stats enable
    stats auth admin:admin
    stats uri /
    stats realm Haproxy\ Statistics
#automatic redirect for http to https connections
    frontend vrops_unsecured_redirect *:80
    redirect location <Insert https fqdn here >
    frontend epops_frontend_secure
    bind <epops dedicated ip>:443
    mode tcp
    option tcplog
    use_backend epops_backend_secure
    #adjust the server parameters to your environment
    backend epops_backend_secure
    mode tcp
    option tcplog
    balance source
    hash-type consistent
    option tcp-check
    timeout queue 20s
    tcp-check connect port 443 ssl
    tcp-check send GET\ /epops-webapp/health-check\ HTTP/1.0\r\n
    tcp-check send \r\n
    tcp-check expect string ONLINE
server node1 <Insert node1 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node2 <Insert node2 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node3 <Insert node3 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
server node4 <Insert node4 ip address here>:443 check inter 15s check-ssl maxconn 140 fall 3
rise 3
```

Note   The line "listen stats :9090" configures the status listener of HAProxy.

Verify HAProxy Configuration

HA Configuration verification

Procedure

1   When the configuration is completed, connect to http://haproxy_ip_address:9090 by using the username and password used to configure HAProxy. In the above example, username: admin and password: admin.

2   Verify that all the nodes rows are shown in green.

Advanced Configuration: HAProxy with Keepalived

In some circumstances and deployments, dual highly available HAProxy is required. In a single-node deployment HAProxy becomes the single point of failure in the deployment and adds potential reliability concerns.

Also, if the HAProxy needs patches, updates, or other maintenance, the HAProxy becomes a single point of downtime. To remediate this concern, deployment of two HAProxys and Keepalived is used to ensure one node is always available. The configuration of the HAProxy can be exactly same across nodes, simply adjusting for local node IP addresses. In most cases the first deployed HAProxy virtual machine can simply be cloned and used as the secondary node.

Failover of a failed HAProxy node by using Keepalived has been tested to occur in less than 5 seconds depending on the network variables. The failover period was rarely noticed by the user or effecting the UI session, during the limited testing. Keepalived uses Linux Virtual Router Redundancy Protocol (VRRP) and multicast advertisements from the master node. If the master node stops sending advertisements the backup proceeds to send a gratuitous ARP to the network and taking ownership of the VIP address and owns the hardware address that master previously owned. The master and the backup monitor each other with multicast events at a rate of once per second.

Figure 2-1. HAProxy with Keepalived
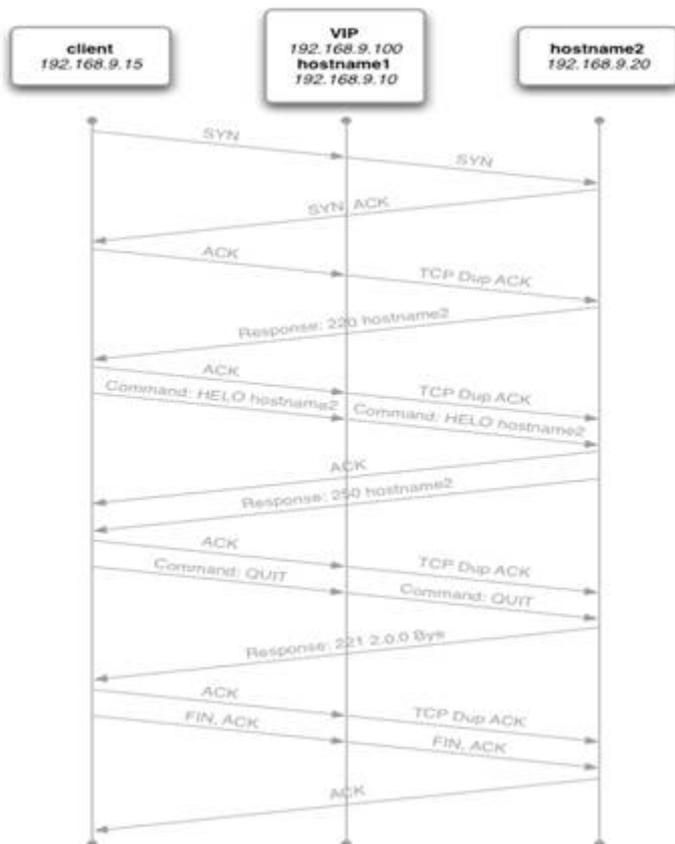
Figure 2-2. vRealize Operations Manager 8-Nodes Cluster using HAProxy with Keepalived
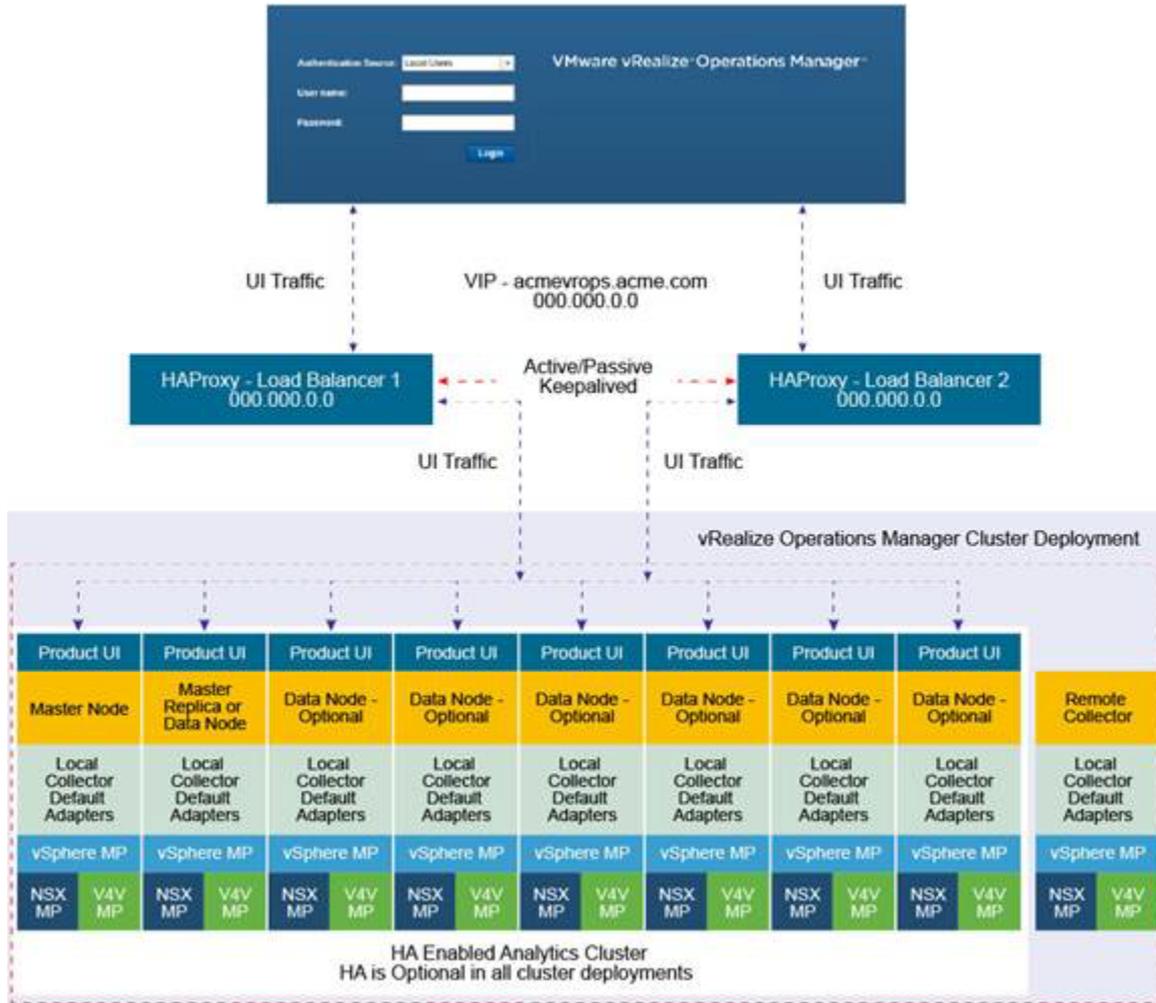
Configure HAProxy with Keepalived

HAProxy with Keepalived configuration guide

1   Clone the HAProxy VM or install a new VM with the same configuration as the first deployed HAProxy.

2   Change Hostname and IP Address

3   Create VIP and point to main DNS record for vRealize Operations Manager cluster. For example: acmevrops6.acme.com / 192.168.1.5)

   You will now have 2x HAProxy load balancers running. For example: LB1/192.168.1.6 and LB2/192.168.1.7.

4   Verify HAProxy configuration is located on both the load balancers. You should be able to access either one and access vRealize Operations Manager cluster successfully.

When both the HAProxies are confirmed working and contain identical configurations, you should configure the Keepalived to ensure that you have availability between the two load balancers.

5   SSH to LB1 which we will consider is the MASTER election.

```
yum install keepalived
```

6   You should configure the kernel to use a VIP to bind to vi /etc/sysctl.conf. Add the following line to the file

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

7   For the kernel to pick up the new changes without rebooting, run the following command:

```
sysctl -p
```

8   Delete the file:

```
/etc/keepalived/keepalived.conf
```

9   Create a new file:

```
/etc/keepalived/keepalived.conf
```

10   In the new keepalived.conf file add the following

```
Master Node
global_defs {
  router_id haproxy2 # The hostname of this host.
}
vrrp_script haproxy {
  script "killall -0 haproxy"
  interval 2
  weight 2
}
vrrp_instance 50 {
  virtual_router_id 50
  advert_int 1
  priority 50
  state MASTER
  interface eth0
  virtual_ipaddress {
     Virtual_IPaddress dev eth0 # The virtual IP address that will be shared between
MASTER and BACKUP
  }
  track_script {
     haproxy
  }
}
```

11   Verify that above the Router_ID is the HOSTNAME of the local load balancer that you are setting up.

12 Verify that you have set up the correct network device, check if you are using eth0.

13 Verify that above the Virtual_IPaddress is the VIP address, and not the local IP address of the LB1 node.

14 Set the priority in increments of 50. In this example, the node has the highest priority, so it is set to 100. Verify that the node is set as the master node.

15 Save the configuration file and restart the services.

16 You must enable the Keepalived service:

```
systemctl enable keepalived
```

17 Run the commands:

```
service keepalived restart
service haproxy restart
```

18 To display if the node has the active load balancer IP, run:

```
ip a | grep eth0
```

19 If the system you are on displays the primary IP address of the load balancer, then this is the active system processing traffic. Verify that only one system displays the primary IP address of the load balancer.

20 If the address is present on both the machines, the configuration is incorrect, and both the machines might not be able to communicate with each other.

21 To configure the second LB2 Keepalived service perform the same steps as above and configure Keepalived service on LB2.

22 In the new keepalived.conf file add the following for the slave node:

```
global_defs {
  router_id haproxy4 # The hostname of this host!
}
vrrp_script haproxy {
  script "killall -0 haproxy"
  interval 2
  weight 2
}
vrrp_instance 50 {
  virtual_router_id 50
  advert_int 1
  priority 50
  state BACKUP
  interface eth0
  virtual_ipaddress {
      Virtual_IPaddress dev eth0 # The virtual IP address that will be shared betwee MASTER
and BACKUP.
  }
```

```
  track_script {
    haproxy
  }
}
```

23 Verify that the Router_ID is the HOSTNAME of the local load balancer that you are setting up.

24 Verify that above the Virtual_IPaddress is the VIP address and not the local IP address of the LB1 node.

25 Set the priority in increments of 50. In this example, the node has the highest priority, so it is set to 100. Verify that the node is set as the backup.

26 Save the configuration file and restart the services.

27 You must enable the Keepalived service:

```
systemctl enable keepalived
```

28 Run the commands:

```
service keepalived restart
```

29 To display if the node has the active load balancer IP, run:

```
ip a | grep eth0
```

30 If the system you are on displays the primary IP address of the load balancer, then this is the active system processing traffic

F5 BIG-IP LTM Installation & Configuration

3

The F5 BIG-IP **L**ocal **T**raffic **M**anager load balancer configuration is similar to the HAProxy configuration.

The LTM uses SSL pass-through in the same manner as with the HAProxy configuration. The LTM configuration has been tested in both one-arm and multi-arm topologies.

Prerequisites

The following are the prerequisites for a functional LTM configuration in front of a vRealize Operations Manager cluster:

- This document assumes that an LTM device is already deployed in the environment and is configured with network connectivity to the deployed environment where the load balancer instance would be used and run from.

- The LTM can be either physical or virtual and can be deployed in one-arm or multi-arm topologies

- The Local Traffic Module (LTM) must be configured and licensed as Nominal, Minimum, or Dedicated. You can configure LTM on System > Resource Provisioning page.

- A vRealize Operations Manager cluster has been deployed in the environment and is fully functional and all nodes in the cluster are accepting UI traffic. This cluster might have high availability enabled but it is not a requirement.

- An additional VIP/Virtual Server IP address for vRealize Operations Manager analytics.

- An additional VIP/Virtual Server IP address for EPOps in case you are configuring separate load balancers for analytics and EPOps.

This chapter includes the following topics:

- Configure Custom Persistence Profile

- Configure Health Monitors

- Configure Server Pools

- Configure Virtual Servers

- Verify Component and Pool Status

Configure Custom Persistence Profile

There are multiple possible profiles provided out of box in most LTM deployments and creating a custom persistence profile using source addresses affinity.

You must create a customer persistence profile by using the following steps:

Procedure

1   Log in to the LTM and select **Local Traffic** > **Profiles** > **Persistence**.

2   Click **Create**.

3   Enter the name **source_addr_vrops** and select **Source Address Affinity** from the drop-down menu.

4   Enable **Custom** mode.

5   Set the **Timeout** to **1800 seconds** (**30 minutes**).

6   Click **Finished**.

Results

Note   The timeout of the vRealize Operations Manager user sessions, configured through the Global Settings page is 30 minutes is, consistent with vRealize Operations Manager configuration. If the timeout value is updated for vRealize Operations Manager, it should be updated for LTM too.

Example for vRealize Operations Manager analytics configuration:

General Properties	
Name	source_addr_vrops
Partition / Path	Common
Persistence Type	Source Address Affinity
Parent Profile	source_addr

Configuration	
Match Across Services	☐
Match Across Virtual Servers	☐
Match Across Pools	☐
Hash Algorithm	Default
Timeout	Specify... 1800 seconds
Prefix Length	None
Map Proxies	☑ Enabled
Override Connection Limit	☐

Example for EPOps configuration:

Configure Health Monitors

Health monitors are required to ensure the LTM has the proper endpoints on the vRealize Operations Manager node to test to make sure the node is available and functioning for clients to access the node.

In this case, create a few Health Monitors to ensure all URLs are checked properly for availability.

Procedure

1   Log in to the LTM and from the main menu select **Local Traffic** > **Monitors**.

2   Click **Create** and provide the required information as shown in the following tables.Leave the default when nothing is specified.

Results

vRealize Operations Manager Analytics configuration:

Name	type	interval	timeout	send string	Receive string	Description
vrops_https	https	20	61	GET /suite-api/api/deployment/node/status? services=api&services=adminui&services=ui\r\n ---------------------------- Note: For older versions of vROPS from 6.6.1 to 7.5 please use the following URL call, as starting from vROps 8.0 status API enhanced to track separate services status. GET /suite-api/api/deployment/node/status \r\n	ONLINE	Default HTTPS monitor to ensure the HTTPS page is accessible

EPOPS configuration:

name	type	interval	timeout	send string	receive string	Description
vrops_epops	https	20	61	GET /epops-webapp/ health-checkHTTP/1.0\r\n	ONLINE	Heartbeat page used to monitor the epops health

Example for vRealize Operations Manager analytics configuration:

Example for EPOps configuration:

Configure Server Pools

Server Pools are used to contain the pools of members or nodes that will be receiving traffic.

You will only need to create a single pool for a vRealize Operations Manager cluster with all nodes participating in the UI traffic as members. In most cases, you will add each node in the cluster except for the remote collectors.

Procedure

1 Log in to the LTM load balancer and select **Local Traffic** > **Pools**.

2 Click **Create** and provide the required information.Leave the default when nothing is specified.

3 Enter each pool member as a **New Node** and add it to the **New Members**.

4 Repeat steps 1, 2, and 3 for each row of information in the following table.

5 On the **Members** page, select the **Load Balancing Method** as the **Least Connections (node)** and **Priority Group Activation** as **Disabled**.

Results

vRealize Operations Manager Analytics configuration:

Name	Description	Health Monitors	Load Balancing Method	Node Name
ha-vrops-prod	vRealize Operations Manager Pool	vrops_https	Least Connections	vrops_node1:<ipaddress> vrops_node2:<ipaddress> vrops_node3:<ipaddress>

EPOps configuration:

Name	Description	Health Monitors	Load Balancing Method	Node Name
ha-epops-prod	vRealize Operations Manager Pool	vrops_epops	Least Connections	vrops_node1:<ipaddress> vrops_node2:<ipaddress> vrops_node3:<ipaddress>

Note Ensure that you are using the correct service port: 443 for SSL.

Example

Example:
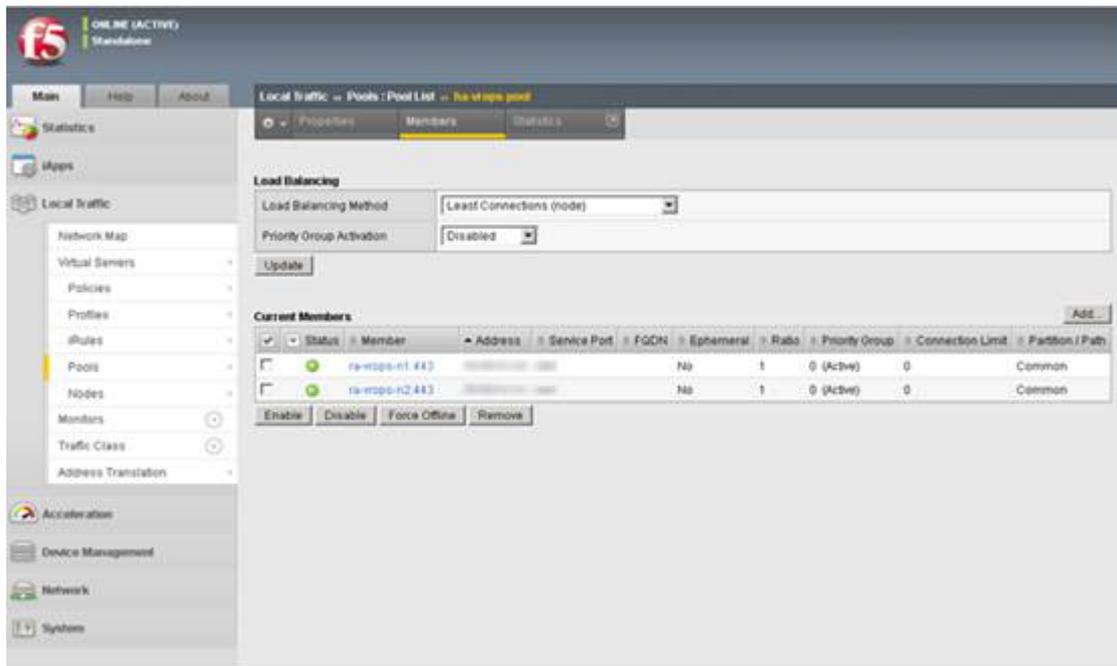
Configure Virtual Servers

Virtual servers contain the virtual IP address (VIP) for the pools of nodes that will be accessed.

In this case, there are two separate VIP's created with the same IP address. One virtual server will be for insecure traffic which will leverage a custom iRule to ensure the traffic gets redirected properly to the HTTPS session. The second virtual server will be used for secure traffic to ensure traffic will be sent directly to the secure HTTPS web page normally.

Procedure

1   Log in to the LTM load balancer and select **Local Traffic** > **Virtual Servers**.

2   Click **Create** and provide the required information. Leave the default when nothing is specified.

3   When all the settings are configured, click **Update** to create the first virtual server.

4   Repeat the steps to configure the second virtual server by using the settings in the table below.

Results

Name	Type	Destination Address	Service Port	HTTP Profile	Service Address Translation	Default Pool	Default Persistence Profile	iRules
ra-vrops-vip-http	Standard	<ipaddress>	80	HTTP	Auto Map	None	None	_sys_https_redirect
ra-vrops-vip	Performance (Layer 4)	<ipaddress>	443	None	Auto Map	ha-vrops-prod	source_addr_vrops	None
epops-vip	Performance (Layer 4)	<ipaddress>	443	None	Auto Map	ha-epops-prod	source_addr_vrops	None

Example

Example:
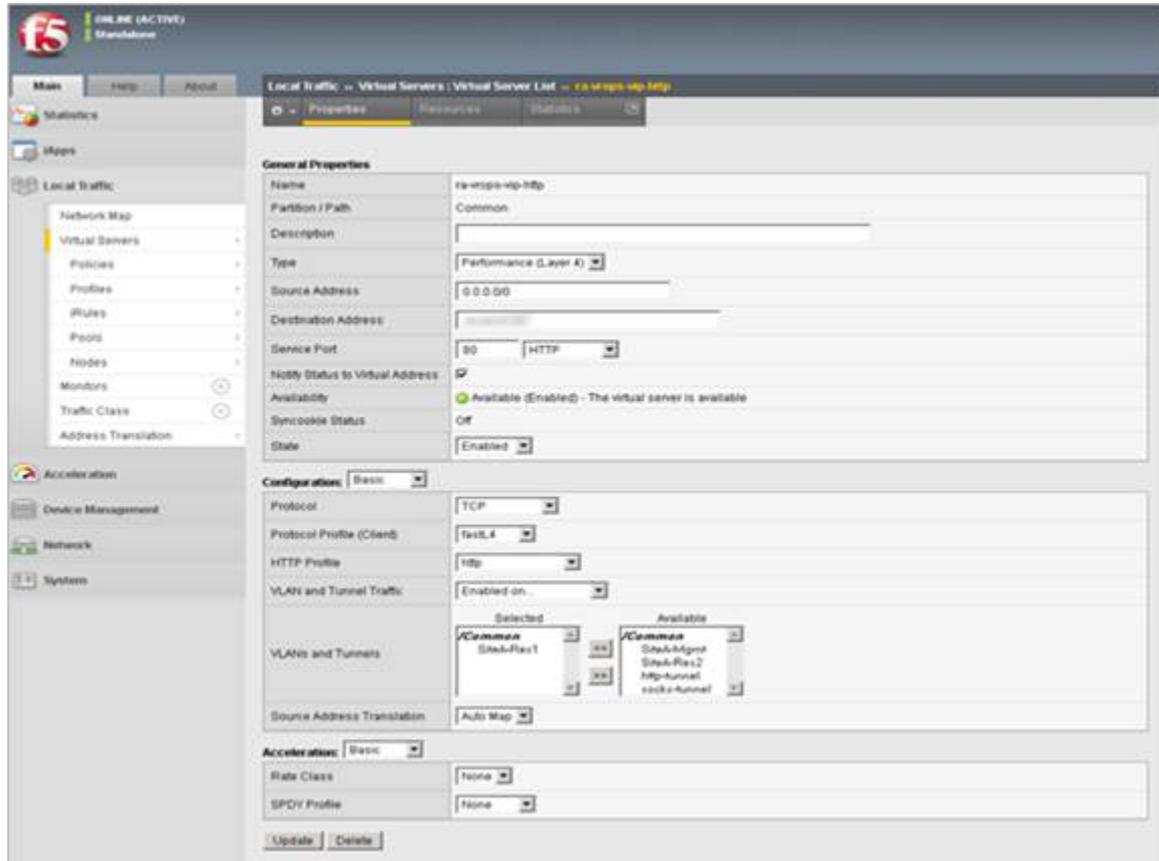
Verify Component and Pool Status

After completing configuration for health monitors, server pools, and virtual servers, verify the status of the configured environment and filter to the specific deployment that was just configured to get an overall view of the nodes, pools, and virtual servers.

Verification steps:

Procedure

1   To check the network map for an overall view of the server pools, select **LTM > Network Map**.

2   Filter the **Network Map** by using the search box to enter the name of the virtual server name used in the configuration.

3   Each status indicator represents the status of the node, the pool, and virtual server or assigned VIP.

Example

Example:

In the following example, you can see both the ra-vrops-vip and the ra-vrops-vip-http VIP are functioning normally. When one of the nodes fail, the indicator will turn red and the indicator for the pool turns yellow to represent a failure in the pool.
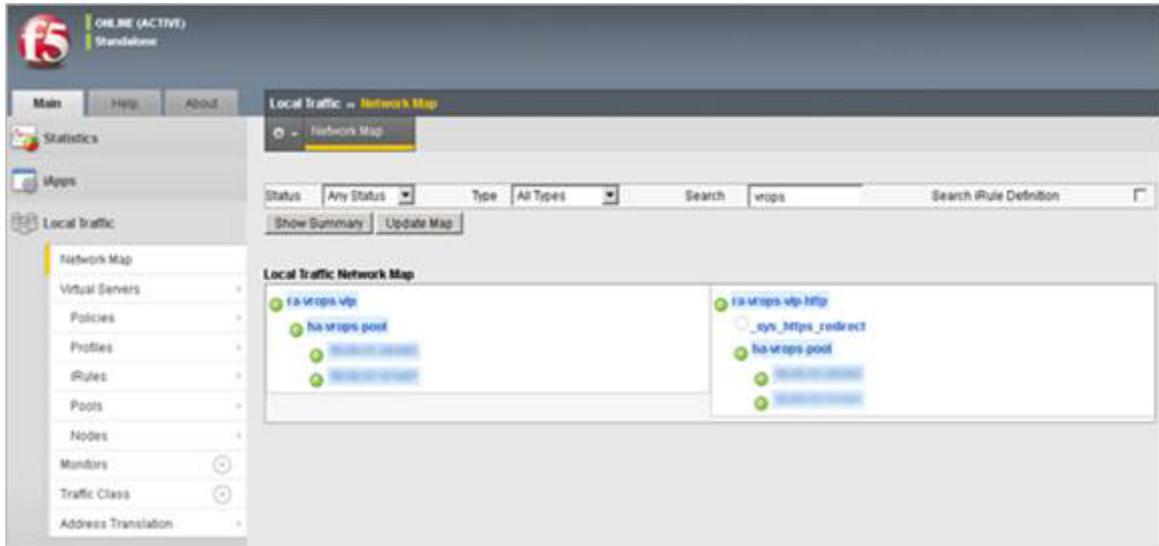
F5 BIG-IP GTM Installation & Configuration

4

The F5 BIG-IP **G**lobal **T**raffic **M**anager DNS based load balancer is designed to be used together with F5's **L**ocal **T**raffic **M**anager for delivering globally distributed applications.

vRealize Operations supports the use of GTM only with the Continuous Availability feature enabled and only for cross datacenter load-balancing between different Fault Domains.

This chapter includes the following topics:

- Terminology
- Architecture
- Prerequisites
- Configure Health Monitors
- Configure GSLB Pools
- Configure Wide-IP

Terminology

F5 Load Balancer terminology

GTM – Global Traffic Manager – DNS based load-balancer, used for cross-DC traffic routing

LTM – Local Traffic Manager – TCP/UDP based load-balancer, typically used in a single DC for multi-server load balancing

CA – Continuous Availability – A vRealize Operations feature which enables you to stretch a cluster across two DCs

FD – Fault Domain - A group of vRealize Operations nodes residing in a single DC. CA supports up to 2 DCs or 2 FDs

Architecture

Typical deployment for vRealize Operations in CA mode includes 2, 4, 6, or 8 nodes based on the appropriate sizing requirements.

Those nodes should be deployed equally into two independent datacenters. One additional witness node should be deployed in a third independent datacenter. Each datacenter is then grouped in a Fault Domain e.g. FD #1 and FD #2. To distribute the traffic between nodes in a Fault Domain, we also need to configure a LTM appliance for each FD (two in total) by following the instructions in this guide.

Since a GTM device primarily handles dynamic DNS record updates, we need to plan the DNS naming before the deployment of the Fault Domains. We also need to ensure all of the DNS records are included into the vRealize Operations SSL certificate – at this point, the installer will not include the address of the LTM VIPs or GTM Wide-IPs; therefore, it will be required to issue and sign (either with external trusted CA or internal one) a new certificate.

In the example below, there are 4 data nodes per Fault Domain, 2 LTM VIPs and 1 GTM Wide-IP. The idea behind this structure is to allow access to the GTM Wide-IP which is globally distributed hence it will point to either FD #1 or FD #2 depending or the current availability (you can also choose to use latency based traffic redirection so a user will be sent to the closest available FD) or access a given FD directly by its LTM VIP for debugging purposes or as a last resort fail-safe.

Name	Type	ADDRESS
vrops-node1.dc1.example.com	A	IP
vrops-node2.dc1.example.com	A	IP
vrops-node3.dc1.example.com	A	IP
vrops-node4.dc1.example.com	A	IP
vrops-node5.dc2.example.com	A	IP
vrops-node6.dc2.example.com	A	IP
vrops-node7.dc2.example.com	A	IP
vrops-node8.dc2.example.com	A	IP
vrops-fd1.dc1.example.com	A	LTM VIP
vrops-fd2.dc2.example.com	A	LTM VIP
vrops.example.com	Wide-IP/A	To be configured later in this chapter

The architecture should look similar to the diagram below:
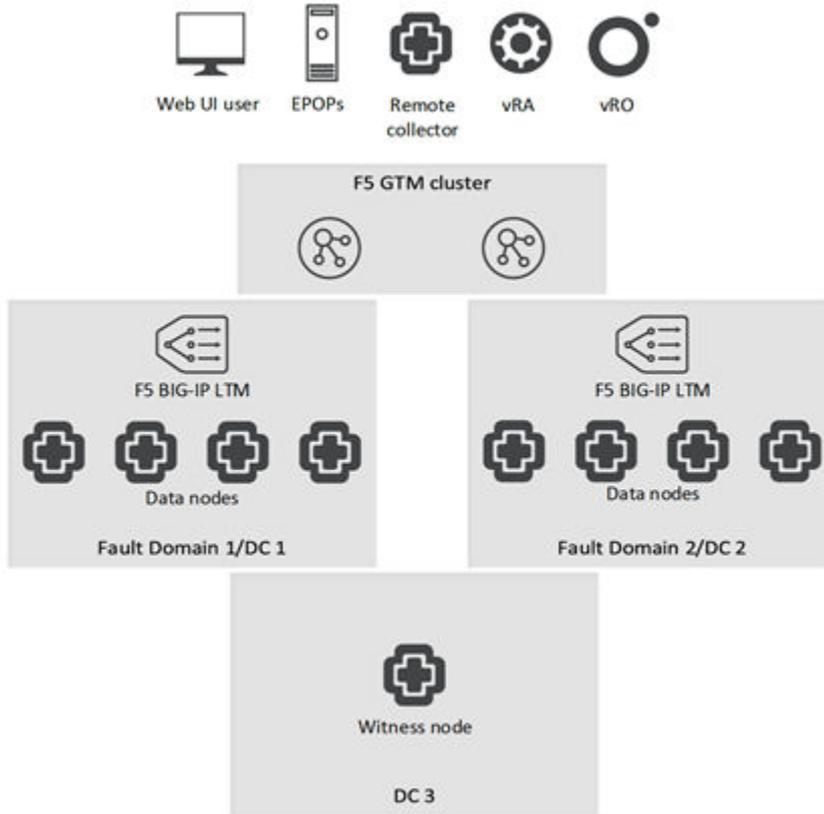
After deploying nodes in each FD and configuring the respective LTM load-balancers, we can procced with the configuration of the GTM nodes. The GTM cluster itself can be deployed in any architecture supported by F5. For our testing, we have used a GTM + LTM combined virtual appliances deployed in each datacenter. We have also clustered only the GTM module since there is no need for clustering on the LTM level. Having separate GTM and LTM appliances or physical systems is supported.

A fully configured and deployed solution during normal operation:
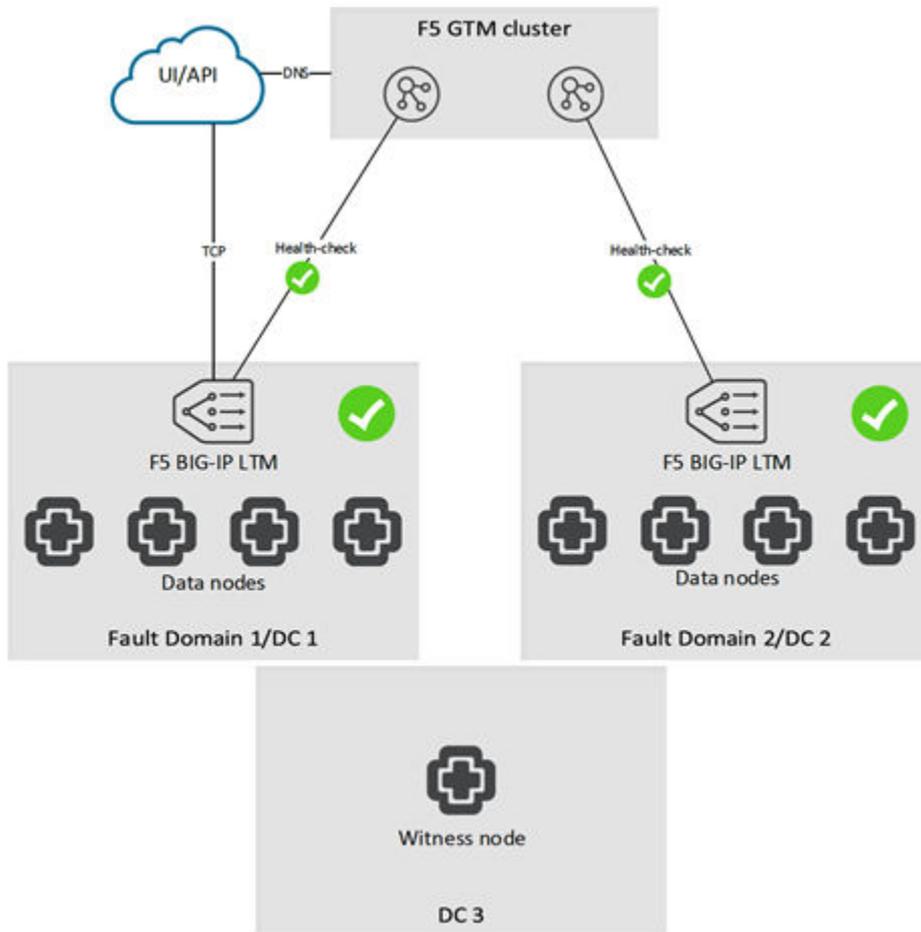
Having the LTMs monitoring each individual vRealize Operations nodes and the GTMs monitoring the accessibility of the entire Fault Domain, ensures the maximum possible fault protection with the least possible overhead.

- In case there is only a single node failure in a Fault Domain, the local LTM will prevent any traffic hitting the affected node while the entire Fault Domain will continue to remain functional

- In case we experience an outage in the entire datacenter, the GTMs will re-route the traffic to a healthy datacenter

- Failover and recovery are automatic in both scenarios

Figure 4-1. Failover scenario #1 – single node failure
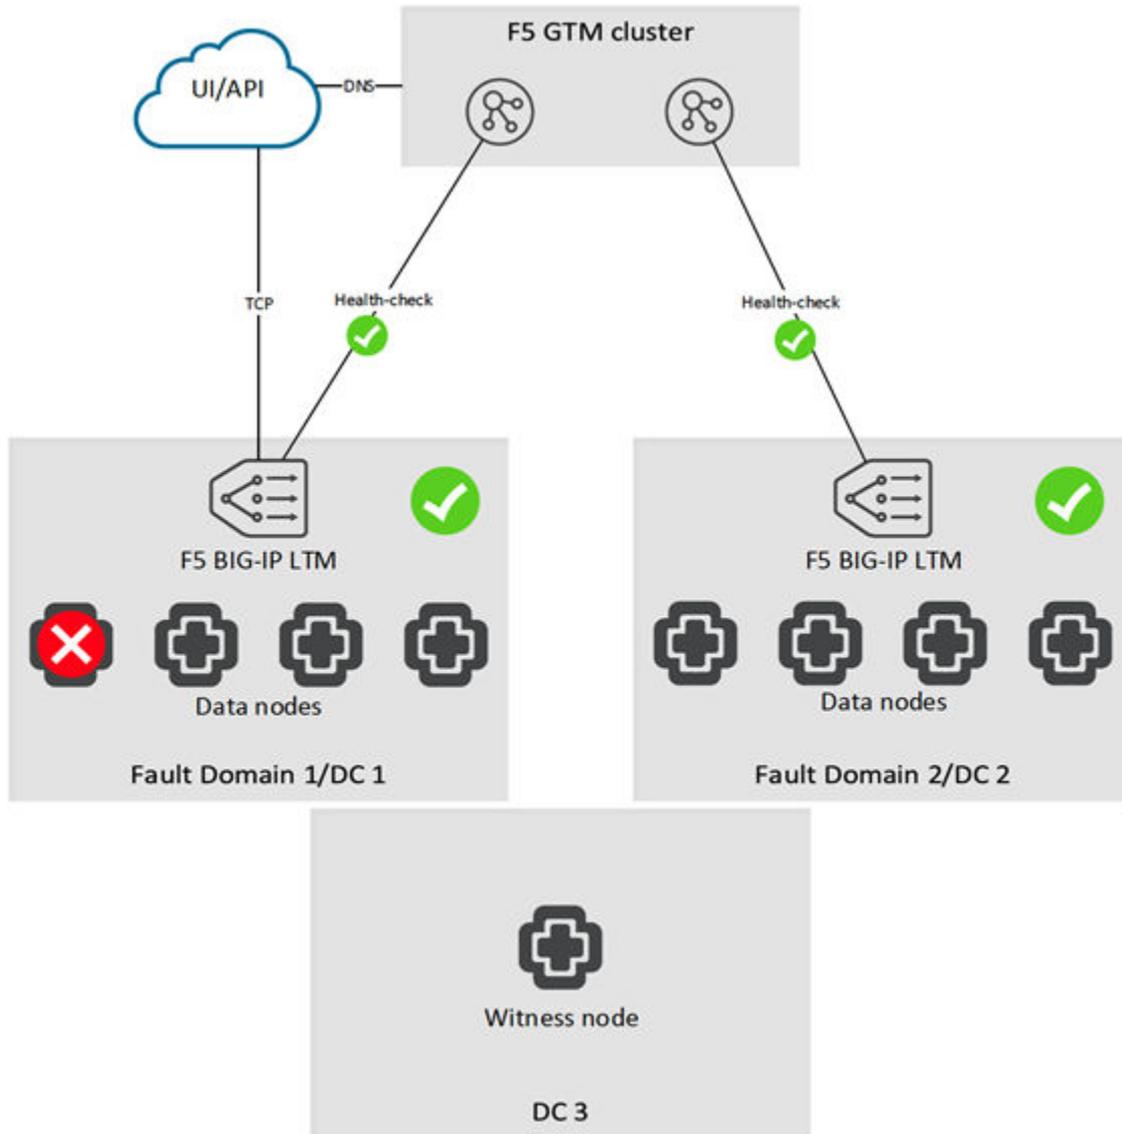
Figure 4-2. Failover scenario #2 – full datacenter outage
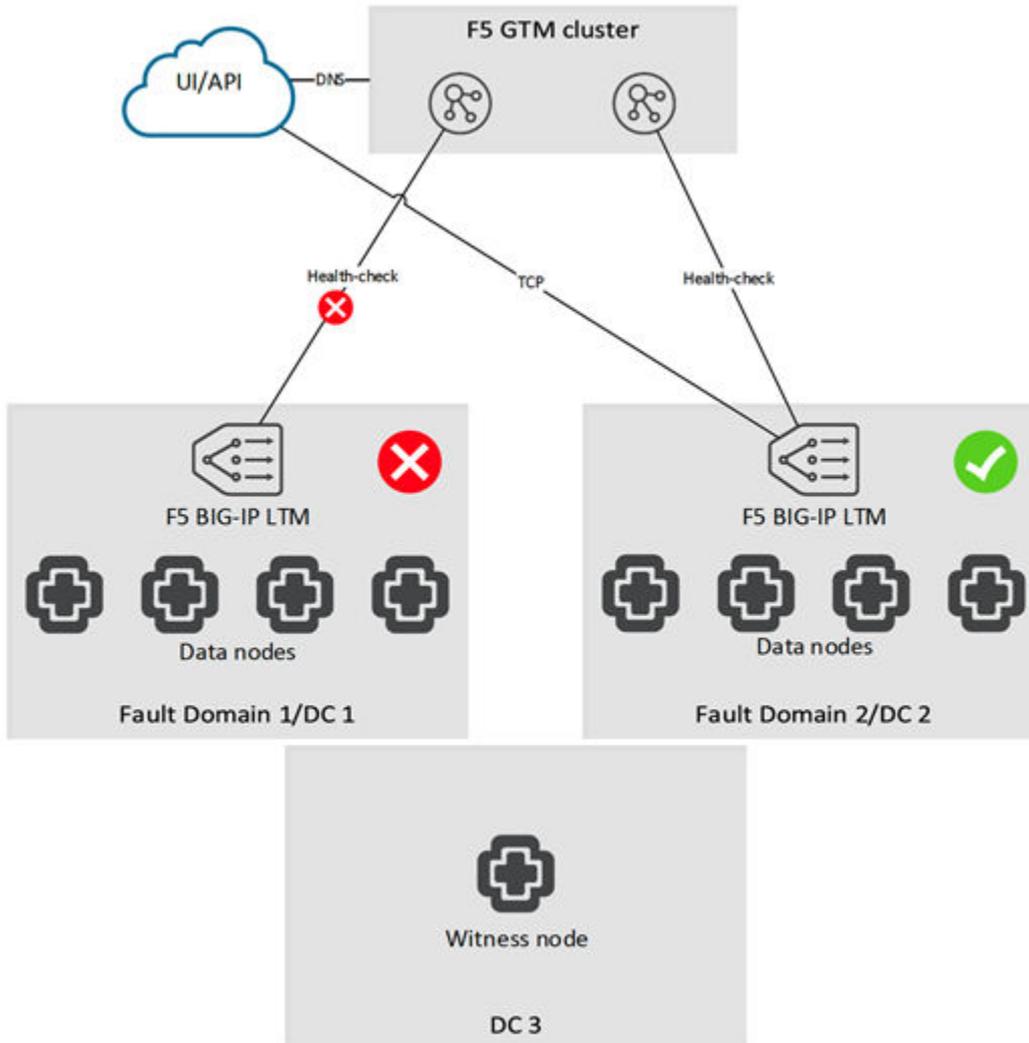
Prerequisites

The following are the prerequisites for a functional GTM configuration managing a vRealize Operations Manager CA enabled cluster

- GTM appliances have to be more than 1 and hosted in more than 1 independent datacenter

- GTM appliances can be deployed in any datacenter globally as long as they are in the same cluster

- LTM appliances have to be in the same datacenter as the respective Fault Domain which they serve

- GTM and LTM appliances have to be paired and trust must be established between them. This is required so the GTM appliances can retrieve the health-check status from the LTM appliances by utilizing the big3d agent.

- GTM and LTM solutions can be either virtual machines or physical systems

- GTM and LTM solutions can be on the same systems or deployed separately

- This document assumes that the LTM and GTM devices are already deployed in the environment and network connectivity is configured. Generic configuration of LTM and GTM devices is not covered in this document, please review the F5's official documentation on how to configure Prober Pools, DNS Listeners and Zones, and how to pair the devices and group them into Datacenters

- vRealize Operations must be deployed and the Continuous Availability feature needs to be enabled

- Configure static DNS records for all vRealize Operations nodes and Fault Domains

Example:

Name	Type	ADDRESS
vrops-node1.dc1.example.com	A	IP
vrops-node2.dc1.example.com	A	IP
vrops-node3.dc1.example.com	A	IP
vrops-node4.dc1.example.com	A	IP
vrops-node5.dc2.example.com	A	IP
vrops-node6.dc2.example.com	A	IP
vrops-node7.dc2.example.com	A	IP
vrops-node8.dc2.example.com	A	IP
vrops-fd1.dc1.example.com	A	LTM VIP
vrops-fd2.dc2.example.com	A	LTM VIP
vrops.example.com	Wide-IP/A	To be configured later in this chapter

- Issue and sign an SSL certificate containing all related DNS records

Configure Health Monitors

GTM health monitors are used to determine the current status of an LTM Virtual IP and redirect the traffic accordingly.

In case of Fault Domain failure our monitors will notice that and send the traffic to the remaining Fault Domain.More about health monitors.

Procedure

1 Log in to the GTM web UI and select **DNS > GSLB > Monitors**.

2 Click **Create** and provide the required information. Leave the default when nothing is specified.

3 Repeat steps 1 and 2 for each row of information in the table below.

Results

vRealize Operations Manager Analytics configuration:

Name	Type	Interval	Timeout	p. Timeout	Send String	Receive String
vrops_https	HTTPS	30 sec.	120 sec.	5 sec.	GET /suite-api/api/deployment/node/status?service=api&service=admin&service=ui\r\n	ONLINE

Configure GSLB Pools

Global Server Load Balancing (GSLB) pools are an GTM objects that group collection of LTM Virtual IPs in order to provide load-balancing and global availability between them

In our architecture it works together with the GTM health monitors and the big3d agents in order to establish the best available datacenter to send user traffic to. More about GSLB Pools.

Procedure

1 Log in to the GTM web UI and select **DNS > GSLB > Pools**.

2 Click **Create** and provide the required information. Leave the default when nothing is specified.

3 Repeat steps 1 and 2 for each row of information in the table below.

Results

Name	Type	Health Monitors	TTL	Maximum answers returned	Load Balancing Method	Members
vrops_pool	A	vrops_https	30 sec.	1	Preferred: Global Availability Alternate: Ration Fallback: Fallback IP Fallback IP: The IP address of your master node	Select the Virtual IPs which resides on each linked LTM and set their desired ratio
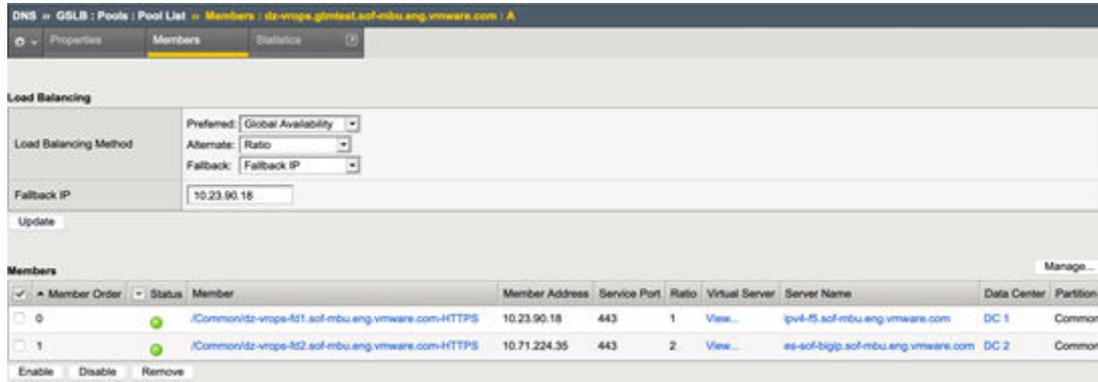
Configure Wide-IP

A wide IP maps a fully-qualified domain name (FQDN) to one or more pools of virtual servers that host the content of a domain.

When an LDNS issues a DNS name resolution for a wide IP, the configuration of the wide IP indicates which pools of virtual servers are eligible to respond to the request, and which load balancing methods BIG-IP GTM uses to select the pool.More about Wide IPs.

Procedure

1   Log in to the GTM web UI and select **DNS > GSLB > Wide IPs**.

2   Click **Create** and provide the required information. Leave the default when nothing is specified.

3   Repeat steps 1 and 2 for each row of information in the table below.

Results

Name	Type	Load-balancing method	Peristance	Last resort Pool	Pools
vrops.example.com	A	Global Availability	Disabled	vrops_pool	vrops_pool

Citrix NetScaler Installation & Configuration

5

Citrix NetScaler configuration guide

Before starting with this configuration make sure that the Netscaler device is deployed in the environment and has access to the vRealize Operations components.

- You can use either virtual or physical Netscaler in single or clustered configuration.

- Enable the **Load Balancer (LB)** and **SSL** modules. You can do so from the **NetScaler > System > Settings > Configure Basic Features** page.

- In case you experience SSL timeout issues with the virtual edition of NetScaler please update the appliance to version 11.0 65.35 or disable TLS 1.1/1.2 as per articlehttp://support.citrix.com/article/CTX205578.

   This is a known NetScaler bug – reference ID: 600155.

- You can use either multi-arm or one-arm configuration. Our tests were done in multi-arm configuration.

- VPX versions of Netscaler doesn't support certificates larger than 2048bits on the back-end servers.

   If you are planning to use VPX you will need to change the vRealize Operations certificate.

   Configure a certificate for use with vRealize Operations Manager

   FAQ: Key Sizes/Certificates Supported by NetScaler

This chapter includes the following topics:

- Configure Health Monitors
- Configure Service Groups
- Configure Virtual Servers
- Configure Persistence Group

Configure Health Monitors

NetScaler health monitors are used to determine the current status of an Virtual IP and redirect the traffic accordingly.

Procedure

1 Log in to the Netscaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Monitors**.

2 Click **Add** and provide the required information. Leave the default when nothing is specified.

3 Repeat steps 1 and 2 for each row of information in the table below.

Results

vRealize Operations Manager Analytics configuration:

Name	Type	Interval	Timeout	RETRIES	Send String	Receive String	DEST. PORT	secure
vrops_http	HTTP	16 sec.	15 sec.	3	GET /	(200| 204| 301)	80	no
vrops_https	HTTP-EVC	16 sec.	15 sec.	3	GET /suite-api/api/ deployment/node/status? services=api&services=adminui&services=ui ---------------------- Note: For older versions of vROPS from 6.6.1 to 7.5 please use the following URL call, as starting from vROps 8.0 status API enhanced to track separate services status. GET /suite-api/api/deployment/node/status \r\n	ONLINE	443	yes
vrops_epops	HTTP-EVC	16 sec.	15 sec.	3	GET /epops-webapp/health-check	ONLINE	443	yes

Example

Example:

Configure Monitor

Name
vrops_https

Type
HTTP-ECV

Standard Parameters | Special Parameters

Interval
16 | Second

Destination IP
. . . | ☐ IPv6

Response Time-out
15 | Second

Destination Port
443

Down Time
30 | Second

TROFS Code
0

TROFS String

Dynamic Time-out
0

Deviation
0 | Second

Dynamic Interval
0

Retries
3

Resp Time-out Threshold
0

SNMP Alert Retries
0

Action

Success Retries
1

Failure Retries
0

Net Profile

☐ TOS
TOS ID

☑ Enabled
☐ Reverse
☐ Transparent
☐ LRTM (Least Response Time using Monitoring)
☑ Secure
☐ IP Tunnel

OK Close

Configure Service Groups

Server Groups are used to contain the pools of members or nodes that will be receiving traffic.

Procedure

1   Log in to the Netscaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Service Groups**.

2   Click **Add** and provide the required information.Leave the default when nothing is specified.

3   Enter each pool member as a **Member** and add it to the **New Members** type **Server Based**.

4   Repeat steps 1, 2, and 3 for each row of information in the table below.

Results

Name	Health Monitors	Protocol	SG MEMBERS	address	Port
ha-vrops-prod_80	vrops_http	HTTP	vrops_node1 vrops_node2 vrops_node3	vrops_node1:<ipaddress> vrops_node2:<ipaddress> vrops_node3:<ipaddress>	80
ha-vrops-prod_443	vrops_https	SSL Bridge	vrops_node1 vrops_node2 vrops_node3	vrops_node1:<ipaddress> vrops_node2:<ipaddress> vrops_node3:<ipaddress>	443
ha-epops-prod_443	vrops_epops	SSL Bridge	vrops_node1 vrops_node2 vrops_node3	vrops_node1:<ipaddress> vrops_node2:<ipaddress> vrops_node3:<ipaddress>	443

Example:

Configure Virtual Servers

Virtual servers contain the virtual IP address (VIP) for the pools of nodes that will be accessed.

Procedure

1   Log in to the Netscaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Virtual Servers**.

2   Click **Add** and provide the required information. Leave the default when nothing is specified.

3 Repeat steps 1 and 2 for each entry in the table below.

Results

Name	Protocol	Destination address	port	LOAD BALANCING METHOD	SERVICE GROUP BINDING
ha-vrops-prod-VIP_80	HTTP	10.23.90.18	80	Leastection	ha-vrops-prod_80
ha-vrops-prod-VIP_443	SSL Bridge	10.23.90.18	443	Leastconnection	ha-vrops-prod_443
ha-vrops-epops-VIP_443	SSL Bridge	10.23.90.19	443	Leastconnection	ha-epops-prod_443

Example

Example:

Configure Persistence Group

Persistence profile using source addresses affinity.

Prerequisites

You must create a customer persistence profile by using the following steps:

Procedure

1 Log in to the Netscaler and select **NetScaler > Traffic Management > Load Balancing > Persistency Groups**.

2 Click Add and provide the required information. Leave the default when nothing is specified.

3 Repeat steps 1 and 2 for each entry in the table below.

Results

group Name	Persistence	timeout	Virtual Server Name
source_addr_vrops	SOURCEIP	30 min.	ha-vrops-prod-VIP_80 ha-vrops-prod-VIP_443
source_addr_epops	SOURCEIP	30 min.	ha-vrops-epops-VIP_443

Note The timeout of the vRealize Operations Manager user sessions, configured through the Global Settings page is 30 minutes is, consistent with vRealize Operations Manager configuration. If the timeout value is updated for vRealize Operations Manager, it should be updated for Netscaler too.

Example

Example:

NSX-V Installation & Configuration

6

The NSX-V virtual networking solution includes the capability of deploying an Edge gateway as a load balancer.

Currently, the NSX-V load balancer has basic load balancing functionality and it should not be considered a full-fledged load balancer with advanced configuration like F5 LTM.

Note Use NSX-V version 6.1.3 and higher for all deployments as many issues with the load balancers have been resolved in this release.

Prerequisites

The following are the prerequisites for a functional NSX-V load balancer in front of a vRealize Operations Manager cluster:

- This document assumes that NSX-V deployment is already deployed in the environment and is fully functional.

- The NSX-V deployment is of version 6.1.3 or higher.

- NSX-V Edge is deployed and has access to the network on which vRealize Operations Manager cluster is deployed.

- Edge can be enabled for high availability, however it is not a requirement

- Currently, there are 2 types of modes the load balancer can be used: Accelerated and Non-Accelerated. Difference between Acceleration enabled/disabled is the LB will passthrough TCP connection (enabled) or terminate the TCP connection (disabled), and then send once the TCP connection is done, it will do open a TCP connection to the pool member.

This chapter includes the following topics:

- Install and Configure Edge for Load Balancing

- Configure Application Profiles

- Add Service Monitoring

- Add Pools

- Add Virtual Servers

- Configure Auto Redirect from HTTP to HTTPS

- Verify Component and Pool Status

Install and Configure Edge for Load Balancing

Enable Load Balancing service

You can specify global load balancer configuration parameters and configure the NSX-V Edge for load balancing by enabling the load balancer service.

Procedure

1   Log in to the vSphere Web Client.

2   Click **Networking & Security** and then click **NSX Edges**.

3   Double-click an NSX-V Edge.

4   Click **Manage** and then click the **Load Balancer** tab.

5   Click **Edit** and select **Enable Load Balancer** and **Enable Acceleration**

6   Click **OK** to save changes and enable the service on the Edge.

Example

Example from NSX-V 6.4.x:
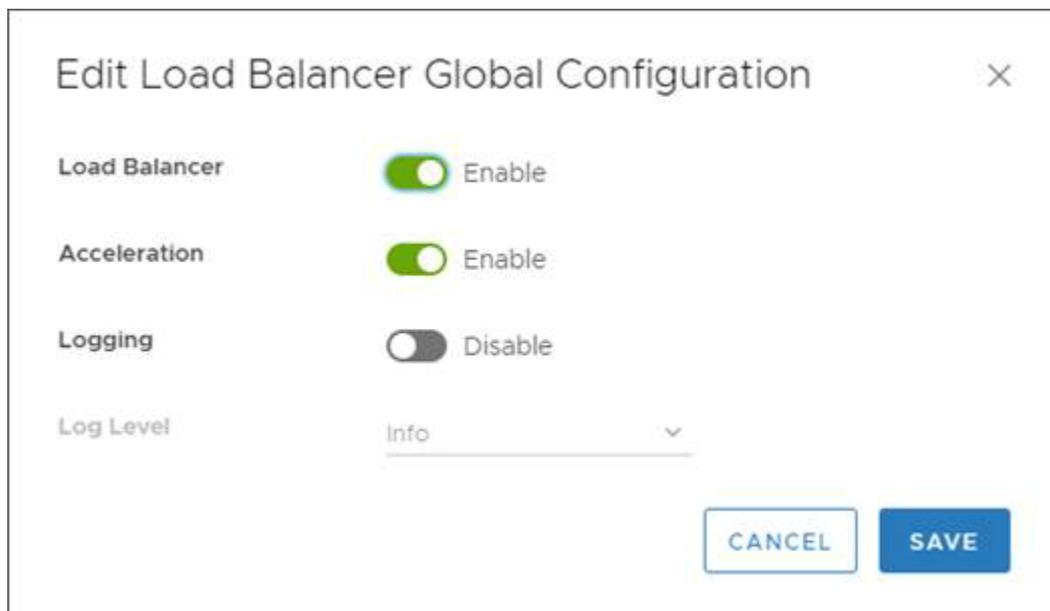
Configure Application Profiles

You must create an application profile to define the behavior of a particular type of network traffic.

After configuring a profile, you should associate the profile with a virtual server. The virtual server then processes traffic according to the values specified in the profile. Using profiles enhances your control over managing network traffic and makes traffic-management tasks easier and more efficient.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Application Profiles**.

6 Click the Add (+ ADD) icon.

7 Enter a name for the profile and select the traffic type for which you are creating the profile. For example: vrops_https.

8 Select the **Type**: **TCP**

9 Select **Persistence** as **Source IP**.

10 Enter **1800** for **Expires in (seconds)**.

11 Select **Ignore** for **Client Authentication**.

12 Click **OK** to save the Profile

Results

Note   When the encrypted traffic is balanced, the load balancer cannot differentiate between the traffic for vRealize Operations Manager analytics and EPOps.If you plan to use two load balancers, one for vRealize Operations Manager analytics and one for EPOps, you could use the same profile as both the profiles are identical. If you create two different profiles, only the name of the profiles is different, but the configurations for both the profiles are identical.

Example

Example:

Add Service Monitoring

Health monitors are required to ensure the NSX-V has the proper endpoints on the vRealize Operations Manager node to test to make sure the node is available and functioning for clients to access the node

Configuring service monitoring is similar to creating health checks on other platforms. In NSX-V 6.1, there is a limitation on how many health checks can be performed against a single node. Currently, you can only have a single health check run against a node to ensure availability.

When you associate a service monitor with a pool, the pool members are monitored according to the service monitor parameters. To configure a Service Monitor, perform the following steps.

Procedure

1   Log in to the vSphere Web Client

2   Click **Networking & Security** and then click **NSX Edges**.

3   Double-click an NSX Edge.

4   Click **Manage** and then click the **Load Balancer** tab.

5   In the left navigation panel, click **Service Monitoring**.

6   Click the Add (+ ADD) icon.

7 Enter a name for the service monitor. For example: vROps_Monitor

8 Enter an **Interval** at which a server is to be pinged.

9 Enter a **Timeout** in seconds, maximum time within which a response from the server must be received.

10 Enter the number of times the server must be pinged before it is declared down.

11 Select the **Method** in which you want to send the health check request to the server. For example: GET.

12 Insert the health check URL as shown in the following table.

13 Enter the **Receive** data string needed for a successful health check response. For example: ONLINE.

14 Click **OK** to save the new Service Monitor.

Results

Name	Interval	Timeout	Retries	Type	Method	URL	Receive:
vROps_Monitor	5	16	3	HTTPS	GET	/suite-api/api/deployment/node/status? services=api&services=adminui&services=ui ------------------------ Note: For older versions of vROPS from 6.6.1 to 7.5 please use the following URL call, as starting from vROps 8.0 status API enhanced to track separate services status. /suite-api/api/deployment/node/status \r\n	ONLINE (upper case)
EPPOS_Monitor	5	16	3	HTTPS	GET	/epops-webapp/health-check	ONLINE (upper case)

Example

Example:

Add Pools

You can add a server pool to manage and share backend servers, flexibly and efficiently.

A pool manages load balancer distribution methods and has a service monitor attached to it for health check parameters.

**Procedure**

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Pools**.

6 Click the Add (+ ADD) to add Pool

7 Enter a name for the load balancer pool. For example: vROps_Pool. (Optional) Enter a description.

8 Select an **Algorithm** from the drop-down list. For example: **LEASTCONN**.

9   Select the **Monitors** from the drop-down list. For example: vROps_Monitor.

10   **Transparent**can be Enabled for 2-arm setup and should be disabled for 1-arm setup. (1-arm when Virtual Server IP and vROPS IP addresses are from same subnet, 2-arm – from different)

11   At the top navigation panel menu click **Members**.

12   Click the Add (＋ ADD) icon to add your member servers and the required information:

   a   Name

   b   IP Addressc.

   c   Weight: 1d.

   d   Monitor Port: 443e.

   e   Port: 443f.

   f   Max Connections: 0g.

   g   Min Connections: 0

Pool Name	Algorithm	Monitors	Member Name	IP Address/ vCenter Container	Weight	Port	Monitor Port	Max Conns	Min Conns
vROps_Pool	LEASTCONN	vROps_Monitor	vROps_Node1	x.x.x.x	1	443	443	0	0
EPOps_Pool	LEASTCONN	EPOps_Monitor	EPOps_Node1	x.x.x.x	1	443	443	0	0

Example

Example:

Add Virtual Servers

Virtual servers contain the virtual IP address (VIP) for the pools of nodes that will be accessed

You can add an NSX Edge internal or uplink interface as a virtual server.

Procedure

1 Log in to the vSphere Web Client.

2 Click **Networking & Security** and then click **NSX Edges**.

3 Double-click an NSX Edge.

4 Click **Manage** and then click the **Load Balancer** tab.

5 In the left navigation panel, click **Virtual Servers**.

6 Click the Add (+ ADD) icon.

7 Enter a name for the virtual server. For example: vROps_Virtual_Server

8 Select **Enable Virtual Server** and **Enable Acceleration**.

9 Select the **Application Profile** name from the drop-down list. For example: Exp: vrops_https

10 Enter a **Name** for the virtual server.

11 (Optional) Enter a description.

12 Enter the IP Address to be used for the VIP.

13 From the drop-down list for **Protocol**, select **TCP**.

14 Enter the **Port** value as 443.

15 From the drop-down list for **Default Pool**, select the default pool that you have configured. For example: vROps_Pool

16 For **Connection Limit** and **Connection Rate Limit**, leave the default as 0.

Results

Note If you are using separate load balancers for vRealize Operations Manager and EPOps, the above steps need to be repeated for EPOps virtual server. Use different names for EPOps profile and respective pool. For example: epops_http and EPOPS_Pool.

Example

Example:

Edit Virtual Server | vrops_tcp

General   Advanced

Virtual Server *	Enable	
Acceleration *	Enable	
Application Profile:	vrops_tcp	
Name: *	vrops_tcp	
Description:		
IP Address: *		Select IP Address
Protocol:	TCP	
Port / Port Range: *	443	
	e.g.: 9000,9010-9020	

CANCEL   SAVE

Configure Auto Redirect from HTTP to HTTPS

When using the NSX-V load balancer in front of the vRealize Operations Manager cluster you may want the URL to automatically redirect to the HTTPS login page.

If you do not configure this the user will need to insert the https field in front of the URL/IP Address. Similar setting is also required in a HAProxy configuration to ensure the redirect works properly. You must configure application profiles and virtual servers for HTTPS redirect.

Note   Ensure that you are using the HTTPS URLs in a correct manner.

Configure Application Profile for HTTPS Redirect

When using the NSX-V load balancer in front of the vRealize Operations Manager cluster you may want the URL to automatically redirect to the HTTPS login page. If you do not configure this the user will need to insert the https field in front of the URL/IP Address

You must configure application profiles and virtual servers for HTTPS redirect.

NOTE: Ensure that you are using the HTTPS URLs in a correct manner.

Procedure

1   Log in to the vSphere Web Client.

2   Click **Networking & Security** and then click **NSX Edges**.

3   Double-click an NSX Edge.

4   Click **Manage** and then click the **Load Balancer** tab.

5   In the left navigation panel, click **Application Profiles**.

6   Click the Add (+ ADD) icon.

7   Enter a name for the Application Profile. For example: vROps_Redirect

8   From the drop-down list for **Type**, select **HTTP**.

9   For **HTTP Redirect URL**, enter https://<ip_address_of_vip>/ui/login.action.

10 From the drop-down list for **Persistence**, select **Source IP**.

11 Enter **1800** for **Expires in (seconds)**.

12 Click **OK** to save.

Example

Example:

Configure the Virtual Server for HTTPS Redirect

Virtual server fort HTTPS redirect

You can configure the virtual server for HTTPS redirect.

Procedure

1   Log in to the vSphere Web Client.

2   Click **Networking & Security** and then click **NSX Edges**.

3   Double-click an NSX Edge.

4   Click **Manage** and then click the **Load Balancer** tab.

5   In the left navigation panel, click **Virtual Servers**.

6   Click the Add (+ ADD) icon.

7   Select **Enable Virtual Server**.

8   Select an **Application Profile** from the drop-down list that you have created. For example: vrops_redirect

9   Enter a **Name** for the virtual server.

10   (Optional) Enter a **Description**.

11   Enter IP Address for the VIP.

12   From the drop-down list for **Protocol**, select **HTTP**.

13   Enter the **Port** value as 80.

14   From the drop-down list for **Default Pool**, select **None**. For NSX-V versions 6.2.7 and 6.3.0, create an empty pool and assign it as the default pool.

15   For **Connection Limit** and **Connection Rate Limit**, leave the default as 0.

Example

Example:

Verify Component and Pool Status

After completing configuration for health monitors, server pools, and virtual servers, verify the status of the configured environment and filter to the specific deployment that was just configured to get an overall view of the nodes, pools, and virtual servers.

You can verify the status of the components running on the load balancer and you can check the status of the pools from inside the UI of the vSphere Web Client.

Procedure

1   Log in to the vSphere Web Client.

2   Click **Networking & Security** and then click **NSX Edges**.

3   Double-click an NSX Edge.

4   Click **Manage** and then click the **Load Balancer** tab.

5   In the left navigation panel, click **Pools**.

6   Click **Show Pool Status (** **)**. A **Pool and Member Status** pop-up window appears.

7   Select a pool ID. For example: vROps_Pool.

Results

The member ID and status of the selected pool are displayed. The status can be **UP** or **DOWN**.

Example

Example:

NSX-T Installation & Configuration 7

The NSX-T virtual networking solution includes the capability of deploying an Edge gateway as a load-balancer.

It offers high availability and load balancing for TCP and HTTP-based applications.

Note Please use NSX-T version 2.2 or higher if you like to handle SSL Certificates within the load-balancer.

Prerequisites

The following are the prerequisites for a functional NSX-T load balancer in front of a vRealize Operations Manager cluster:

- This document assumes that NSX-T is already deployed in the environment and is fully functional

- The NSX-T deployment is version 2.2 or higher

- NSX-T Edge has access to the network on which the vRealize Operations Manager cluster is deployed

- NSX-T Tier-1 edge for load balancing is configured

- A vRealize Operations Manager cluster has been deployed in the environment and is fully functional with all nodes in the cluster accepting traffic. The cluster might have high availability enabled, but it is not a requirement

- 1 Virtual Server IP address for vRealize Operations Manager analytics

An additional VIP/Virtual Server IP address for EPOps traffic, in case of separate load balancers is used for analytics and EPOps

This chapter includes the following topics:

- NSX-T Version 2.2, 2.3

- Configure Application Profiles

- Configure Persistence Profile

- Add Active Health Monitor

- Configure Server Pools

- Configure Virtual Servers

- Configure Load Balancer

- Verify Components, Pool and Virtual Server Status

- NSX-T Version 2.4, 2.5.X, 3.X.X

- Configure Load Balancer

- Configure Application Profiles

- Configure Persistence Profile

- Add Active Health Monitor

- Configure Server Pools

- Configure Virtual Servers

NSX-T Version 2.2, 2.3

Configuration guide for NSX-T version 2.2, 2.3

Configure Application Profiles

Application profile must be created to define the behavior of a particular type of network traffic.

For NSX-T, two application profiles need to be created to:

- Redirect HTTP to HTTPS

- Handle HTTPS traffic

After the configuration of an application profile, the same should be associated with a virtual server. The virtual server then processes traffic according to the options specified in the application profile.

Log in to the NSX-T UI:

Configure the Application Profile for HTTP requests:

- Go to **Load Balancing -> Virtual Servers -> Application Profiles**

- Click the **Add** (

) icon and choose **HTTP Profile**.

- Choose a name for the profile and enter parameters (please refer to the example below)

Configure the Application Profile for HTTPS requests:

■   Go to **Load Balancing -> Virtual Servers -> Application Profiles**

■   Click the **Add** (

) icon and choose **Fast TCP Profile**.

■   Choose a name for the profile and enter parameters (please refer to the example below)

Configure Persistence Profile

Persistence profile using source addresses affinity

- Go to **Load Balancing** → **Virtual Servers** → **Persistent Profiles**

- Click the **Add** (

) icon and select **Source IP Persistence**

- Choose a name for the profile and enter parameters (please refer to the example below)

Add Active Health Monitor

Configuring active health monitoring is similar to creating health checks on other load-balancers.

When you associate an active health monitor with a pool, the pool members are monitored according to the active health monitor parameters. To configure an **Active Health Monitor**, perform the following steps:

- Go to **Load Balancing → Server Pools → Active Health Monitors**

- Click the **Add** (   ) icon

- Choose a name for the active health monitor and enter **Monitor Properties** (please refer to the example below)

  Note LbHttpsMonitor is pre-configured monitor for HTTPS protocol and it can be used for this Active Health Monitor

- Configure Health check parameters withthe following values:

```
1.    Request Method
         GET
2.    Request URL
         /suite-api/api/deployment/node/status?services=api&services=adminui&services=ui
3.    Request Version
         HTTP_VERSION_1_1
4.    Response Status Codes
         200, 204, 301
5.    Response Body
         ONLINE (upper case)
6.    Ciphers
         TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,
         TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,
         TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA, TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,
         TLS_ECDH_RSA_WITH_AES_256_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA,
         TLS_RSA_WITH_AES_128_CBC_SHA, TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,
         TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,
         TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384, TLS_RSA_WITH_AES_128_CBC_SHA256,
         TLS_RSA_WITH_AES_128_GCM_SHA256, TLS_RSA_WITH_AES_256_CBC_SHA256,
         TLS_RSA_WITH_AES_256_GCM_SHA384, TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,
         TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,
         TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,
         TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,
         TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,
         TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,
         TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256,
         TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256,
         TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384,
         TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384, TLS_ECDH_RSA_WITH_AES_128_CBC_SHA,
         TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256, TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256,
         TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384, TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384

         Note: Ciphers selection can be vary based on security requirements.
```

```
7.    Protocols
        TLS_V1_1, TLS_V1_2
8.    Server Auth
        IGNORE
9.    Certificate Chain Depth
        3
```

Name	Interval	Timeout	Retries	Type	Method	URL	Receive:
vROPs_MONITOR	5	16	3	HTTPS	GET	/suite-api/api/deployment/node/status?services=api&services=adminui&services=ui ------------------------------------ Note: For older versions of vROPS from 6.6.1 to 7.5 please use the following URL call, as starting from vROps 8.0 status API enhanced to track separate services status: /suite-api/api/deployment/node/status \r\n	ONLINE (upper case)
EPOPS_Monitor	5	16	3	HTTPS	GET	/epops-webapp/health-check	ONLINE (upper case)

- Here is an example of how the configuration should look like:

- Example for vROPS_Monitor

- Example for EPOPS_Monitor
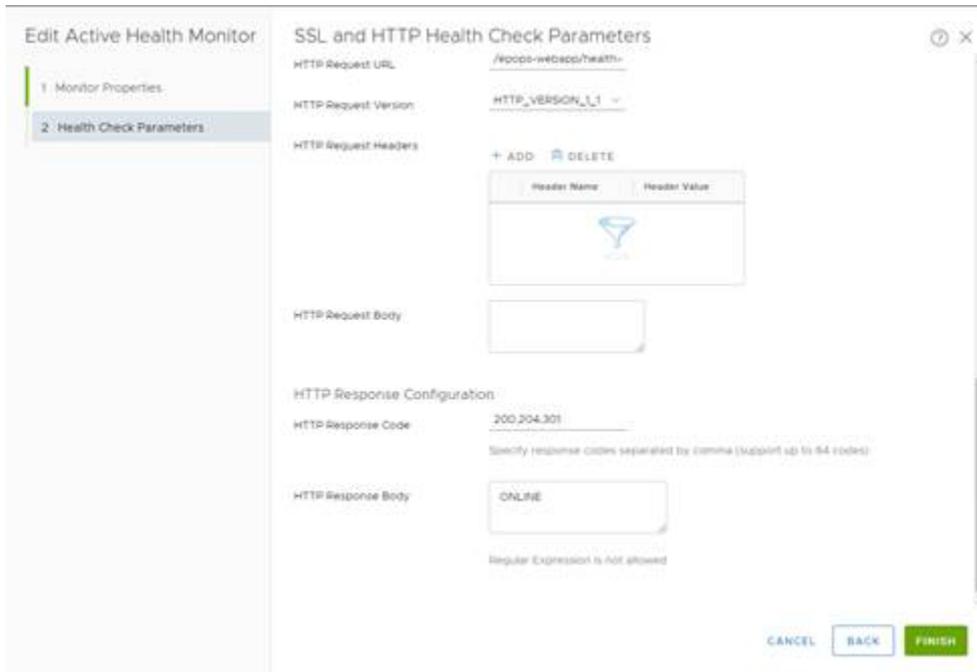
- **Note:**There is an issue with active health monitor in version 2.3.0.1. For this version Active Health Monitor should be configured by the following way in order to avoid unexpected Virtual Servers down (Upgrade to NSX-T Version 2.4 is the permanent recommendation)

- Click the **Add** ( ) icon

- Choose a name for the active health monitor and enter **Monitor Properties** (please refer to the example below)

```
1. Health Check Protocol
   LbTcpMonitor
2. Monitoring Port
   443
```

Note LbTcpMonitor is pre-configured monitor for TCP protocol and it can be used for this Active Health Monitor

Configure Server Pools

NSX-T Server Pools are used to contain the nodes that are receiving traffic.

You will need to create a single pool per vRealize Operations Manager cluster with all the data nodes participating in the cluster as members. Remote collectors should not be added into this pool.

Configure a Server Pool:

- Go to **Load Balancing** → **Server Pools** → **Server Pools**

- Click the **Add** ( ) icon

- Choose a **Name** for the pool. For example: vROPs-POOL

- Set **Load Balancing Algorithm** as LEAST_CONNECTION

- Configure **SNAT Translation** as **Auto Map**

- Add the pool members (vRealize Operations Manager data nodes IP addresses and Port)

  a   Name

  b   IP Address

  c   Weight: 1

  d   Port: 443

  e   State: ENABLED

▪ Attach an **Active Health Monitor** to the pool (please refer to the example below)

Pool Name	Algorithm	Monitors	Member Name	IP Address	Weight	Port	STATE
vROPS-POOL	LEASTCONN	vROPS_MONITOR	vROPS_NODE1	x.x.x.x	1	443	ENABLED
EPOPS_POOL	LEASTCONN	EPOPS_Monitor	EPOPS_NODE1	x.x.x.x	1	443	ENABLED

Edit Server Pool

SNAT Translation ⑦ ✕

1 General Properties

2 SNAT Translation

3 Pool Members

4 Health Monitors

Three Modes based on the topology are supported. In case of inline deployment of Load Balancer, use Transparent (NO_SNAT) to preserve original Client IP and Port. Auto Map mode uses LB interface IP and ephemeral port. In scenarios where both Clients and Pool Members are attached to the same Logical Router, SNAT (Auto Map or IP List) must be used.

Translation Mode *   ○ Transparent ● Auto Map ○ IP List

Port Overload   Enabled ●

Overload Factor   2 ▾

CANCEL   BACK   NEXT

Edit Server Pool

Pool Members ⑦ ✕

1 General Properties

2 SNAT Translation

3 Pool Members

4 Health Monitors

Pool Members can either be Static members that allows you to add IPs and Ports of individual servers or Dynamic Members as defined by NSGroup Membership Criteria. The admin state in case of the Dynamic Members can be set after Server Pool creation in the Members section of the Server Pool. Currently only IPv4 addressing is supported.

Membership Type   ● Static ○ Dynamic

Static Membership

+ ADD   ☐ CLONE   ☐ DELETE

	Name	IP	Port	Weight	State	Backup Member	Max. Concurrent Connection
○	Master		443	1	ENABLED	◯	
○	Replica		443	1	ENABLED	◯	
○	Data1		443	1	ENABLED	◯	
○	Data2		443	1	ENABLED	◯	

▦ COLUMNS   4 Pool Members

CANCEL   BACK   NEXT
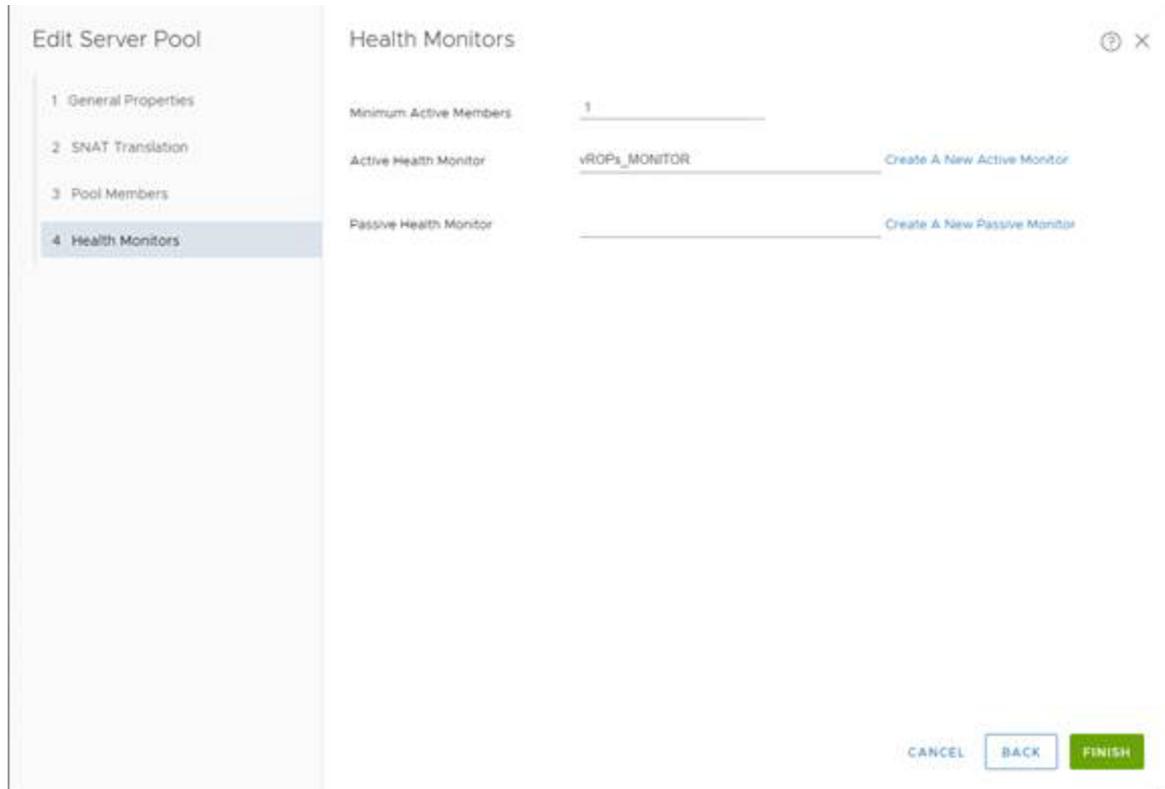
Configure Virtual Servers

NSX-T Virtual Servers contain the Virtual IP address (VIP) for the pools of nodes that will be accessed.

In this case, there are two separate VIPs created with the same IP address. One virtual server is used for redirecting insecure HTTP (port 80) traffic to a secure-channel connection – HTTPS (port 443). The second virtual server is used for handling and forwarding secure-channel traffic (HTTPS) to the backend systems.

Configure the Virtual Servers for HTTP requests:

- Go to **Load Balancing** → **Virtual Servers** → **Virtual Servers**

- Click the **Add** ( ) icon

- Choose a name for Virtual Server

- Configure **Application Type** as **Layer 7**

- Assign appropriate **Application Profile** (please refer to the example below)

- Assign VIP (Virtual IP) and port 80 to handle HTTP requests

- Add **Default Pool Member Port** 80

■ Assign appropriate **Persistent Profile** (please refer to the example below)

Note There is no need to configure any Server Pool for this Virtual Server

Configure the Virtual Servers for HTTPS requests:

- Go to **Load Balancing** → **Virtual Servers** → **Virtual Servers**

- Click the **Add** ( ) icon

- Choose a name for the Virtual Server

- Configure **Application Type** as **Layer 4**

- Assign appropriate **Application Profile** (please refer to the example below)

- Assign a VIP (Virtual IP) and port 443 to handle HTTPS requests

- Add **Default Member Port** 443.

- Assign appropriate **Server Pool** (please refer to the example below)

- Assign appropriate **Load Balancing Profile** (please refer to the example below)

Edit Virtual Server

General Properties

1 General Properties
2 Virtual Server Identifiers
3 Server Pool
4 Load Balancing Profiles

Name * VROPS-NSXT22-HTTPS

Description

Load Balancer Application Profile

Load Balancer Application Profile defines the application protocol characteristics of the Virtual Server. The current release supports three types of App Profiles: Fast TCP Profile, Fast UDP Profile and HTTP Profile. For HTTP and HTTPS applications (Layer-7 load balancing), a HTTP Profile must be chosen as the Application Profile. For Non-HTTP application you may select a Fast TCP or Fast UDP Application Profiles.

Application Type *   ○ Layer 7   ● Layer 4   TCP ∨

Application Profile *   vROPs_HTTPS

Access Log   Enabled ●

CANCEL   NEXT

Edit Virtual Server

Virtual Server Identifiers

1 General Properties
2 Virtual Server Identifiers
3 Server Pool
4 Load Balancing Profiles

IP Address *   192.168.207.200

Port *   443

Specify port (e.g. 8080) or port range (e.g. 80-90) or both separated by comma (e.g. 8080, 80-90, 20)

Protocol   TCP

Advanced Properties

Maximum Concurrent Connection

Maximum New Connection Rate

Default Pool Member Port   443

Specify port (e.g. 8080) or port range (e.g. 80-90) or both separated by comma (e.g. 8080, 80-90, 20)

CANCEL   BACK   NEXT
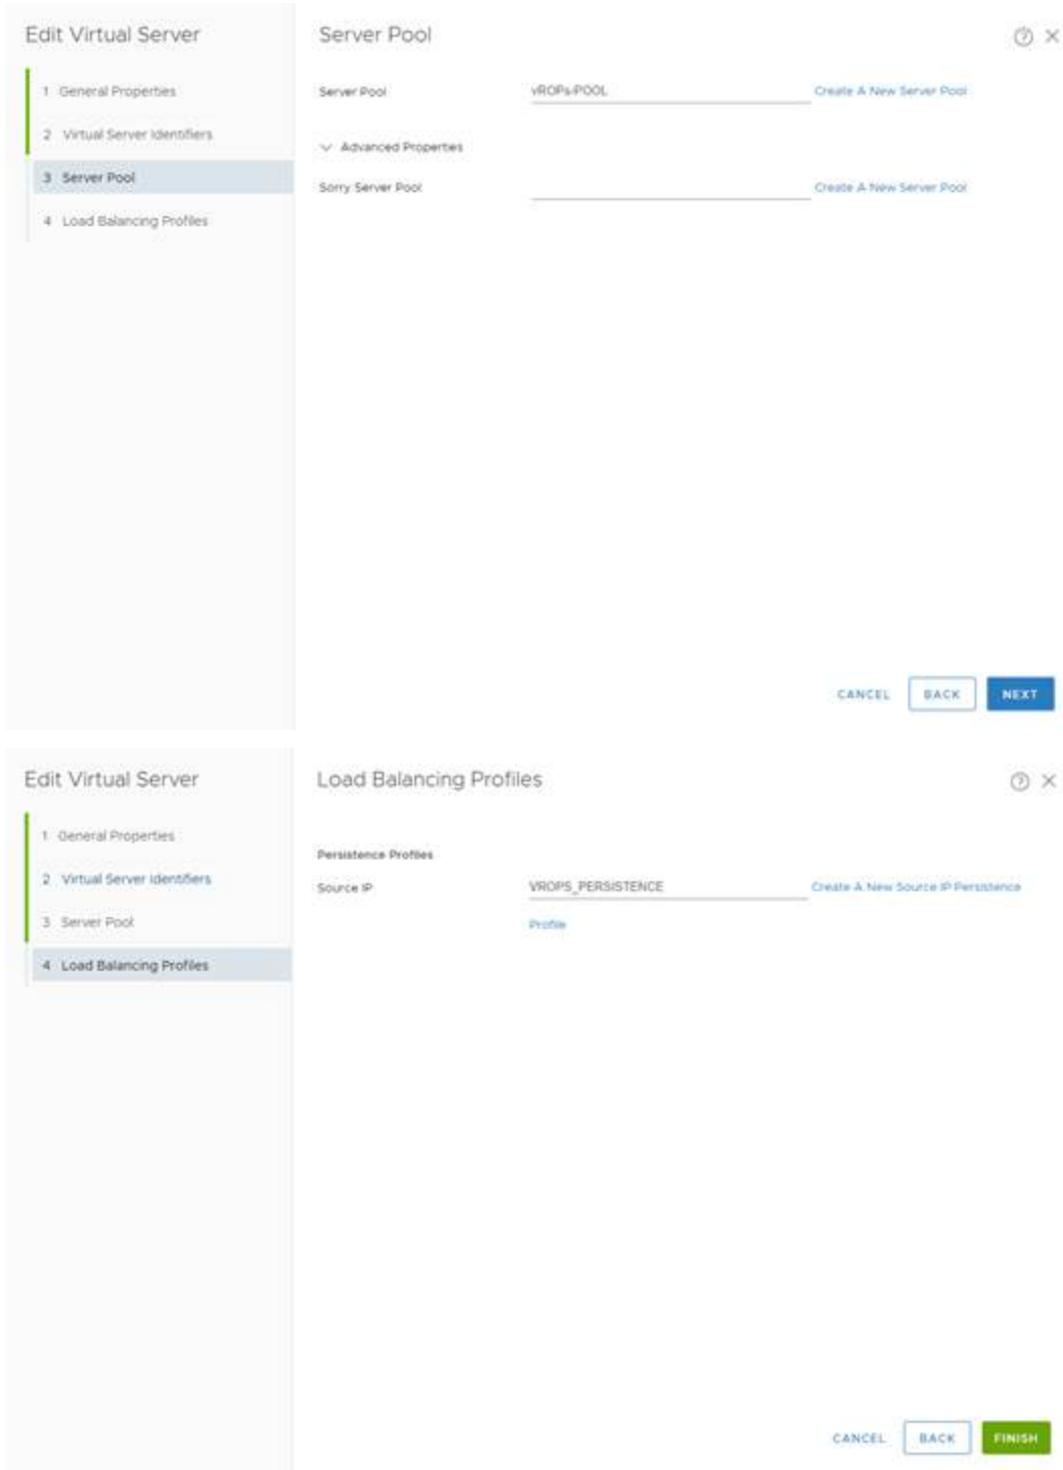
Edit Virtual Server

1 General Properties

2 Virtual Server Identifiers

3 Server Pool

4 Load Balancing Profiles

Server Pool ⓘ ✕

Server Pool   vROPs-POOL   Create A New Server Pool

∨ Advanced Properties

Sorry Server Pool   Create A New Server Pool

CANCEL BACK NEXT

Edit Virtual Server

1 General Properties

2 Virtual Server Identifiers

3 Server Pool

4 Load Balancing Profiles

Load Balancing Profiles ⓘ ✕

Persistence Profiles

Source IP   VROPS_PERSISTENCE   Create A New Source IP Persistence

Profile

CANCEL BACK FINISH

Configure Load Balancer

Specify a load-balancer configuration parameter

You need to configure the NSX-T appliance for load balancing by creating the respective service.

- Go to **Load Balancing → Load Balancers**

- Click the **Add** (

) icon

- Choose a name, select appropriate sizing (depends on vROPS cluster size) and error log level and press OK

- Attach the previously created during installation and configuration "Tier 1 Logical Router" to the newly created Load Balancer (**Overview** → **Attachment** → **EDIT**)

- Attach the previously created Virtual Servers for HTTP and HTTPS to the Load Balancer (Virtual Servers → ATTACH)

Attach to a Logical Router

Select the Router to which the Load Balancer VROPS-NSXT22 is to be attached. Only Tier-1 Routers in 'Active Standby' are currently supported. Note: The Load Balancer can only be Enabled if it had a Virtual Server associated with it.

Tier-1 Logical Router   *   DONT-DELETE-VROPS-Tier-1-Router

CANCEL   OK

Verify Components, Pool and Virtual Server Status

After completion of configuration, status of components running on the load balance can be verified.

To get an overall view of the nodes, pools and virtual servers need to use steps described below:

- Go to **Load Balancing → Server Pools → Server Pools**

- Select the pool that you want to verify. For example: vROPS-POOL

- Click on **Pool Member Statistics.** The member IP:Port and status of the selected pool are displayed. The status should be UP. (can be UP or DOWN)


- Go to **Load Balancing → Virtual Servers → Virtual Servers**

- Select the virtual server that you want to verify. For example: VROPS-NSXT22-HTTPS

- Click on **Statistics. Connections, Connection Rate** and **Throughput** should be displayed. If configuration is mentioned metrics should display status graphs.
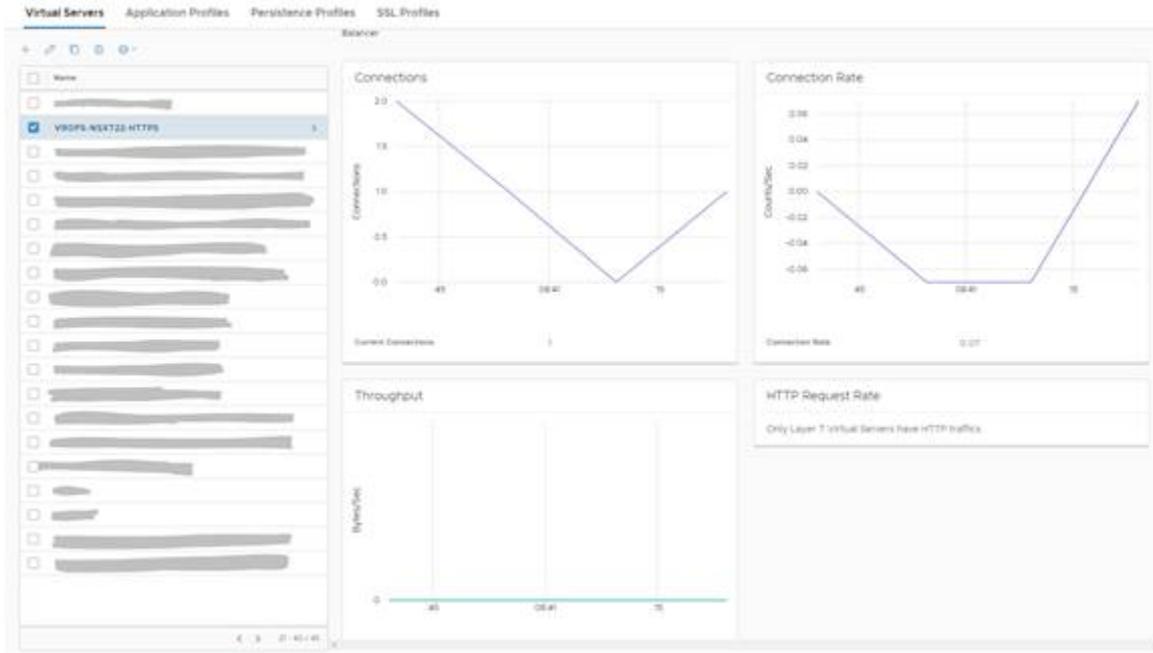
NSX-T Version 2.4, 2.5.X, 3.X.X

Configuration guide for NSX-T version 2.4, 2.5.X, 3.X.X

Configure Load Balancer

Specify a load-balancer configuration parameter

You need to configure the NSX-T appliance for load balancing by creating the respective service.

- Go to **Networking**→Load Balancing → Load Balancers

- Click the **Add** (
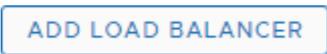
) icon

- Choose a name, select appropriate sizing (depends on vROPS cluster size), error log level, previously created during installation and configuration "Tier 1 Logical Router" and press OK

- For the following dialog select NO
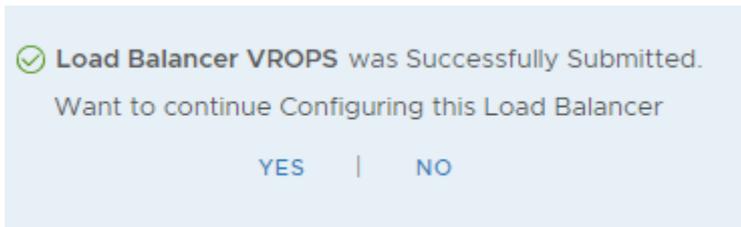
Configure Application Profiles

Application profile must be created to define the behavior of a particular type of network traffic

For NSX-T, two application profiles need to be created to:

- Redirect HTTP to HTTPS

- Handle HTTPS traffic

After the configuration of an application profile, the same should be associated with a virtual server. The virtual server then processes traffic according to the options specified in the application profile.

Configure the Application Profile for HTTP requests:

- Go to **Networking -> Load balancing -> Profiles**

- Select Profile Type

- Click the **Add** (

) icon and choose **HTTP Profile**.

- Choose a name for the profile and enter parameters (please refer to the example below)

-

Configure the Application Profile for HTTPS requests:

■ Go to **Networking -> Load balancing -> Profiles**

■ Select profile type

APPLICATION ˅

■ Click the **Add** (

ADD APPLICATION PROFILE ˅

) icon and choose **Fast TCP Profile**.

■ Choose a name for the profile and enter parameters (please refer to the example below)

Configure Persistence Profile

Persistence profile using source addresses affinity

■ Go to **Networking -> Load balancing -> Profiles**

■ Select profile type

PERSISTENCE ˅

■ Click the **Add** (

ADD PERSISTENCE PROFILE ˅

) icon and select **Source IP**

■ Choose a name for the profile and enter parameters (please refer to the example below)

Add Active Health Monitor

Configuring active health monitoring is similar to creating health checks on other load-balancers.

When you associate an active health monitor with a pool, the pool members are monitored according to the active health monitor parameters. To configure an **Active Health Monitor**, perform the following steps:

- Go to **Networking -> Load balancing -> Monitors**

- Select monitor type **ACTIVE ⌄**

- Click the **Add** (**ADD ACTIVE MONITOR ⌄**) icon and select **HTTPS**

- Choose a name for the active monitor and enter **Monitor Properties** (please refer to the example below)

- Configure Health check parameters withthe following values:

```
3. HTTP Method
   GET
4. HTTP Request URL
   /suite-api/api/deployment/node/status?services=api&services=adminui&services=ui (or /
epops-webapp/health-check for EPOPS)
5. HTTP Request Version
   1.1
6. HTTP Response Code
   200, 204, 301
7. HTTP Response Body
   ONLINE (upper case)
```

Name	Interval	Timeout	Retries	Type	Method	URL	Receive:
vROPs_MONITOR	5	16	3	HTTPS	GET	/suite-api/api/deployment/node/status?services=api&services=adminui&services=ui ----------------------------- Note: For older versions of vROPS from 6.6.1 to 7.5 please use the following URL call, as starting from vROps 8.0 status API enhanced to track separate services status: /suite-api/api/deployment/node/status \r\n	ONLINE (upper case)
EPOPS_Monitor	5	16	3	HTTPS	GET	/epops-webapp/health-check	ONLINE (upper case)

- Here is an example of how the configuration should look like:

- Example for vROPS_Monitor

HTTP Request and Response Configuration ✕

Active Health Monitor - vROPs_MONITOR

HTTP Request Configuration   HTTP Response Configuration

HTTP Method	Get ⌄
HTTP Request URL	/suite-api/api/deployme
HTTP Request Version	1.1 ⌄

[ADD]

Header Name	Header Value

Request Header not found

HTTP Request Body

HTTP Request and Response Configuration ✕

Active Health Monitor - vROPs_MONITOR

HTTP Request Configuration   **HTTP Response Configuration**

HTTP Response Code	(200 ✕) (204 ✕) (301 ✕)
	1 or more response code
HTTP Response Body	ONLINE

▪ Example for EPOPS_Monitor

HTTP Request and Response Configuration ✕

Active Health Monitor - vROPs_MONITOR

HTTP Request Configuration   HTTP Response Configuration

HTTP Method	Get ∨
HTTP Request URL	/epops-webapp/health~
HTTP Request Version	1.1 ∨

ADD

Header Name	Header Value

Request Header not found

HTTP Request Body

HTTP Request and Response Configuration ✕

Active Health Monitor - vROPs_MONITOR

HTTP Request Configuration   **HTTP Response Configuration**

HTTP Response Code	200 ✕ 204 ✕ 301 ✕
	1 or more response code
HTTP Response Body	ONLINE

Configure Server Pools

NSX-T Server Pools are used to contain the nodes that are receiving traffic.

You will need to create a single pool per vRealize Operations Manager cluster with all the data nodes participating in the cluster as members. Remote collectors should not be added into this pool.

Configure a Server Pool:

- Go to **Networking** →**Load Balancing** → **Server Pools**

- Click the **Add** (ADD SERVER POOL) icon

- Choose a **Name** for the pool. For example: vROPs-POOL

- Set **Algorithm** as **LEAST CONNECTION**

- Configure **SNAT Translation Mode** as **Automap**

- Attach an **Active Monitor** to the pool (please refer to the example below)
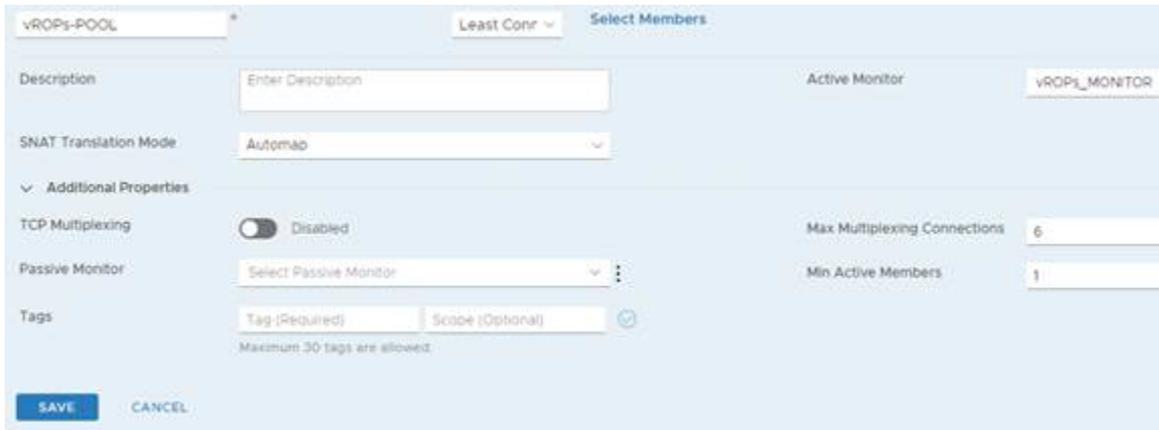
- Add the pool members via Select Members (vRealize Operations Manager data nodes IP addresses and Port)

  a   Name

  b   IP Address

  c   Weight: 1

  d   Port: 443

  e   State: ENABLED f.

Pool Name	Algorithm	Monitors	Member Name	IP Address	Weight	Port	STATE
vROPS-POOL	LEASTCONN	vROPS_MONITOR	vROPS_NODE1	x.x.x.x	1	443	ENABLED
EPOPS_POOL	LEASTCONN	EPOPS_Monitor	EPOPS_NODE1	x.x.x.x	1	443	ENABLED

Configure Virtual Servers

NSX-T Virtual Servers contain the Virtual IP address (VIP) for the pools of nodes that will be accessed.

In this case, there are two separate VIPs created with the same IP address. One virtual server is used for redirecting insecure HTTP (port 80) traffic to a secure-channel connection – HTTPS (port 443). The second virtual server is used for handling and forwarding secure-channel traffic (HTTPS) to the backend systems.

Configure the Virtual Servers for HTTPS requests:

- Go to **Networking → Load Balancing → Virtual Servers**

- Click the **Add** (ADD VIRTUAL SERVER ⌄) icon and select **L4 TCP**

- Choose a name for the Virtual Server

- Assign appropriate **Application Profile** (please refer to the example below)

- Assign appropriate **Load Balancer** (please refer to the example below)

- Assign appropriate **Server Pool** (please refer to the example below)

- Select **Persistence** as **Source IP** (please refer to the example below)

- Assign appropriate **Source IP** profile (please refer to the example below)

- Assign a VIP (Virtual IP) and port 443 to handle HTTPS requests

Configure the Virtual Servers for HTTP requests:

■ Go to **Networking → Load Balancing → Virtual Servers**

■ Click the **Add** (ADD VIRTUAL SERVER ∨) icon and select **L7 HTTP**

■ Assign appropriate **Application Profile** (please refer to the example below)

■ Assign VIP (Virtual IP) and port 80 to handle HTTP requests

Note There is no need to configure any Server Pool for this Virtual Server