



vRealize Operations Manager Load Balancing

Configuration Guide
Version 7.5

TECHNICAL WHITE PAPER
JULY 2019
VERSION 1.9.1

Table of Contents

Introduction	5
Load Balancing Concepts	5
Selecting a Load Balancer	5
How to Handle SSL UI Certificates with a Load Balancer	6
vRealize Operations Manager Overview	6
vRealize Operations Manager Architecture	7
Configuring End Point Operations Agents	9
HAProxy Installation and Configuration	10
Install Single-Node HAProxy	10
Configure Logging for HAProxy	11
Configure HAProxy	11
Configure HAProxy for vRealize Operations Manager Analytics	11
Configure EPOps HAProxy	13
Verify HAProxy Configuration	15
Advanced Configuration: HAProxy with Keepalived	15
Configure HAProxy with Keepalived	17
F5 Big IP Installation & Configuration	21
Configure Custom Persistence Profile	21
Configure Health Monitors	23
Configure Server Pools	25
Configure Virtual Servers	26
Verify Component and Pool Status	28
Citrix NetScaler Installation & Configuration	29
Configure Health Monitors	29
Configure Service Groups	32
Configure Virtual Servers	33
Configure Persistence Group	33
NSX-V Installation & Configuration	35
Install and Configure Edge for Load Balancing	35
Configure Application Profiles	36
Add Service Monitoring	37
Add Pools	39
Add Virtual Servers	40
Configure Auto Redirect from HTTP to HTTPS	42
Configure Application Profile for HTTPS Redirect	42

Configure the Virtual Server for HTTPS Redirect	43
Verify Component and Pool Status	44
NSX-T Installation & Configuration	46
NSX-T Version 2.2, 2.3	47
Configure Application Profiles	47
Configure Persistence Profile	48
Add Active Health Monitor	49
Configure Server Pools	53
Configure Virtual Servers	57
Configure Load Balancer	62
Verify Components, Pool and Virtual Server Status	64
NSX-T Version 2.4	66
Configure Load Balancer	66
Configure Application Profiles	67
Configure Persistence Profile	68
Add Active Health Monitor	68
Configure Server Pools	72
Configure Virtual Servers	73

Revision History

DATE	VERSION	DESCRIPTION
July 2019	1.9.1	Update to NSX-V prerequisites. Update F5 healthcheck requirements
April 2019	1.9	Update to NSX-V and NSX-T configuration. Minor updates to include vRealize Operations Manager version 7.5
September 2018	1.7	Addition of NSX-T and updates to HAProxy, F5 BIG-IP Minor updates to include vRealize Operations Manager version 7.0
April 2018	1.6	Update to NSX-V and F5 configurations
April 2017	1.5	Update to include new values for interval/timeout health checks and lower the potential downtime. Minor updates to include vRealize Operations Manager 6.5
January 2017	1.4	Updates to include newer versions of load balancing software.
November 2016	1.3	Minor updates to include vRealize Operations Manager version 6.4
August 2016	1.2	Minor updates to include vRealize Operations Manager version 6.3
February 2016	1.1	Minor updates to include vRealize Operations Manager version 6.2
December 2015	1.0	Initial version.

Introduction

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, 6.7, 7.0, 7.5	https://www.vmware.com/support/pubs/vmware-vrops-suite-pubs.html
F5 BIG-IP	11.5, 11.6, 12.1, 13.0, 14x	https://support.f5.com/kb/en-us.html
Citrix NetScaler	10.5*, 11.0*, 11.x, 12.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, 2.3, 2.4	https://docs.vmware.com/en/VMware-NSX-T/index.html
HA Proxy	v1.5.x	http://www.haproxy.org/
RHEL	v7.x	https://access.redhat.com/documentation/en-US/index.html
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.

Load Balancing Concepts

Load balancers distribute connections among servers in high availability (HA) deployments. The system administrator backs up the load balancers on a regular basis at the same time as other components.

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

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.

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. You can use a load balancer 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 through 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:

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.

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 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. This 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

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
- CPU: 2 vCPU
- Memory: 4GB
- Disk space: 50GB
- HAProxy 1.5.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

Install Single-Node HAProxy

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

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, see [Configure HAProxy Analytics](#) and to configure EPOps balancer, see [Configure 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
```

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.

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
```

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.

Configure HAProxy for vRealize Operations Manager Analytics

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  400
    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\ HTTP/1.0\r\n\r\n
tcp-check expect rstring ONLINE

server node1 <Insert node1 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
server node2 <Insert node2 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
server node3 <Insert node3 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
server node4 <Insert node4 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6

```

Configure EPOps HAProxy

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 60s check-ssl maxconn 140 fall 6 rise 6
server node2 <Insert node2 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
server node3 <Insert node3 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
server node4 <Insert node4 ip address here>:443 check inter 60s check-ssl maxconn 140 fall 6 rise 6
```

NOTE: The line “listen stats :9090” configures the status listener of HAProxy.

Verify HAProxy Configuration

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.

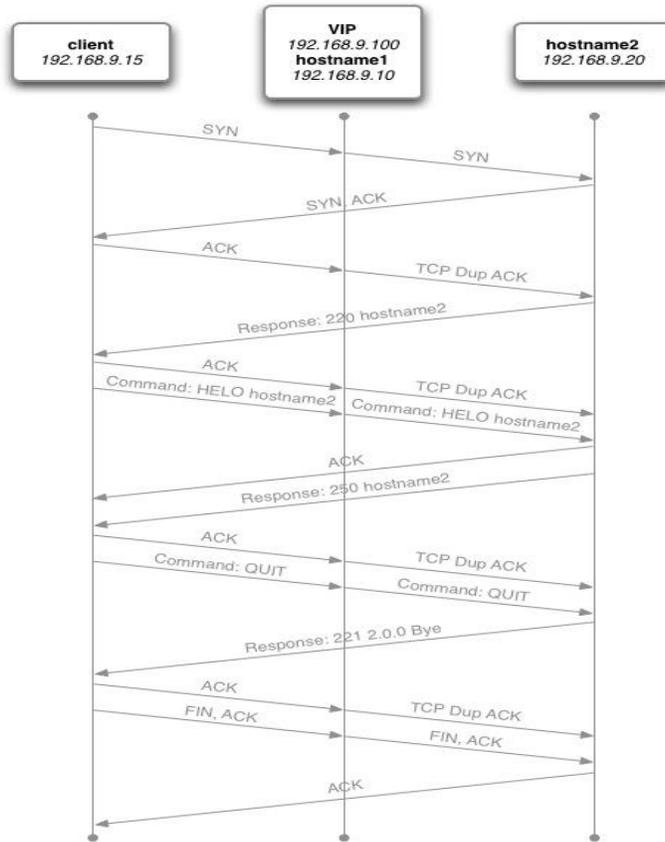


FIGURE 2. HAPROXY WITH KEEPALIVED

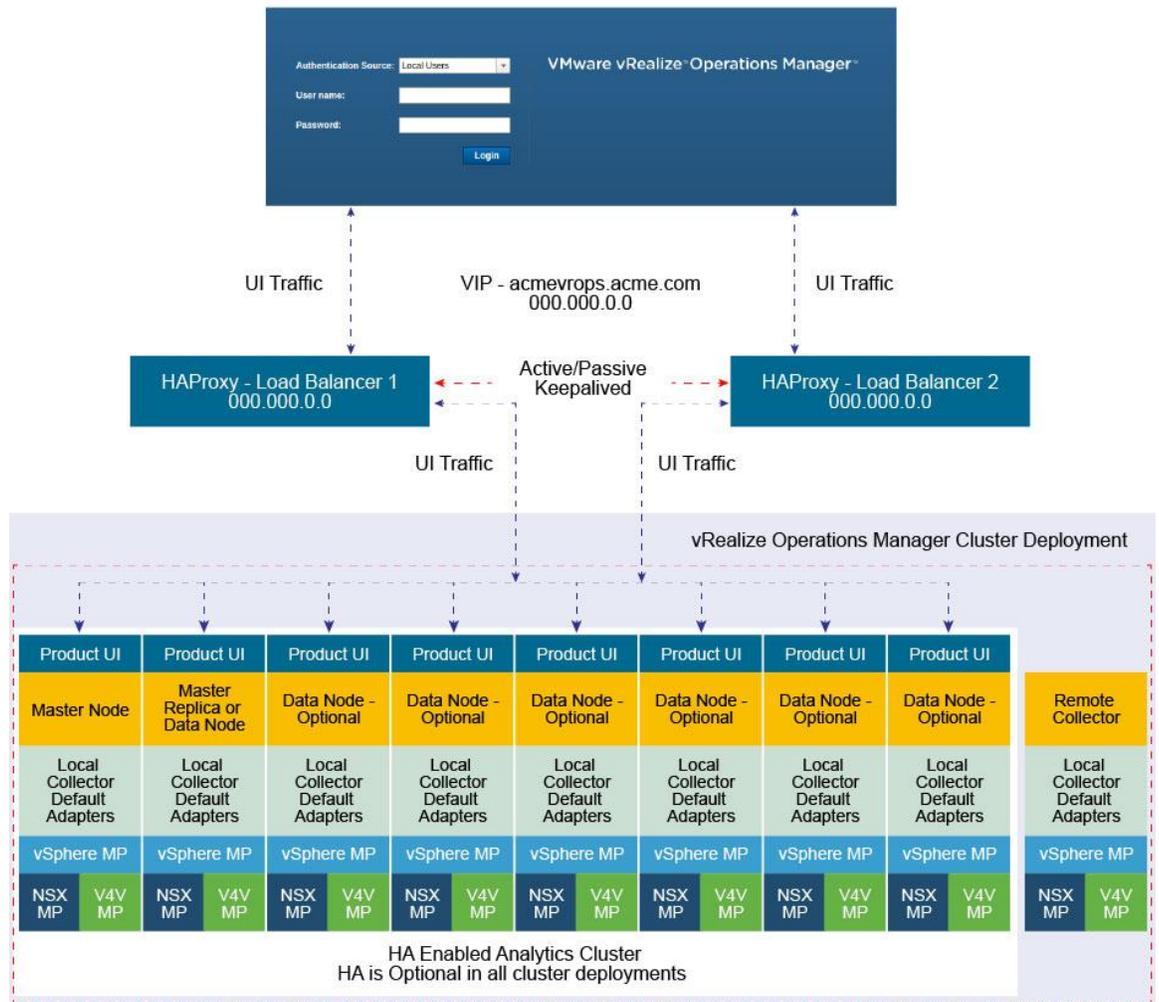


FIGURE 3. vREALIZE OPERATIONS MANAGER 8-NODES CLUSTER USING HAProxy WITH KEEPALIVED

Configure HAProxy with Keepalived

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 between 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  
service haproxy 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 Installation & Configuration

The F5 Big IP load balancer configuration is similar to the HAProxy configuration. The F5 uses the SSL pass-through in the same manner as the HAProxy configuration. The F5 configuration has been tested in both one-arm and multi-arm topologies.

Prerequisites

The following are the prerequisites for a functional F5 configuration in front of a vRealize Operations Manager cluster:

- This document assumes that an F5 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 F5 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.

Configure Custom Persistence Profile

There are multiple possible profiles provided out of box in most F5 deployments and creating a custom persistence profile using source addresses affinity. You must create a customer persistence profile by using the following steps:

1. Log in to the F5 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**.

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 F5 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	<input type="checkbox"/>
Match Across Virtual Servers	<input type="checkbox"/>
Match Across Pools	<input type="checkbox"/>
Hash Algorithm	Default ▼
Timeout	Specify... ▼ 1800 seconds
Prefix Length	None ▼
Map Proxies	<input checked="" type="checkbox"/> Enabled
Override Connection Limit	<input type="checkbox"/>

Example for EPOps configuration:

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

Configuration	
Match Across Services	<input type="checkbox"/>
Match Across Virtual Servers	<input type="checkbox"/>
Match Across Pools	<input type="checkbox"/>
Hash Algorithm	Default ▼
Timeout	Specify... ▼ 1800 seconds
Prefix Length	None ▼
Map Proxies	<input checked="" type="checkbox"/> Enabled
Override Connection Limit	<input type="checkbox"/>

Configure Health Monitors

Health monitors are required to ensure the F5 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.

1. Log in to the F5 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.

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\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-check HTTP/1.0\r\n	ONLINE	Heartbeat page used to monitor the epops health

Example for vRealize Operations Manager analytics configuration:

Local Traffic » Monitors » wrops_https

⚙️ Properties Instances

General Properties

Name	wrops_https
Partition / Path	Common
Description	Default HTTPS monitor to ensure the HTTPS page is accessible
Type	HTTPS
Parent Monitor	https

Configuration: Basic ▾

Interval	20 seconds
Timeout	61 seconds
Send String	GET /suite-api/api/deployment/node/status\r\n
Receive String	ONLINE
Receive Disable String	
Cipher List	DEFAULT:+SHA:+3DES:+kEDH
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	* All Ports
Adaptive	<input type="checkbox"/> Enabled

Update Delete

Example for EPOps configuration:

Local Traffic » Monitors » vrops_epops	
	Properties
	Instances
General Properties	
Name	vrops_epops
Partition / Path	Common
Description	Heartbeat page used to monitor the epops health
Type	HTTPS
Parent Monitor	https
Configuration: Basic	
Interval	20 seconds
Timeout	61 seconds
Send String	GET /epops-webapp/health-check HTTP/1.0\r\n
Receive String	ONLINE
Receive Disable String	
Cipher List	DEFAULT:+SHA:+3DES:+kEDH
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	* All Ports
Adaptive	<input type="checkbox"/> Enabled
<input type="button" value="Update"/> <input type="button" value="Delete"/>	

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.

1. Log in to the F5 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.

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

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:

The screenshot displays the vRealize Operations Manager interface. The left sidebar shows navigation options like Main, Help, About, Statistics, Mapps, Local Traffic, Acceleration, Device Management, Network, and System. The main content area is titled 'Local Traffic > Pools: Pool List > ha-vrops-pool'. Under the 'Members' tab, the 'Load Balancing' section shows 'Load Balancing Method' set to 'Least Connections (node)' and 'Priority Group Activation' set to 'Disabled'. Below this is an 'Update' button. The 'Current Members' section features a table with columns for Status, Member, Address, Service Port, FQDN, Ephemeral, Ratio, Priority Group, Connection Limit, and Partition / Path. Two members are listed: 'ra-vrops-n1.443' and 'ra-vrops-n2.443', both with a status of 'Active', a ratio of 1, and a priority group of 0. Below the table are buttons for 'Enable', 'Disable', 'Force Offline', and 'Remove'.

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.

1. Log in to the F5 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.

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	ha-vrops-profile	None
epops-vip	Performance (Layer 4)	<ipaddress>	443	None	Auto Map	ha-epops-prod	ha-vrops-profile	None

Example:

The screenshot displays the F5 configuration interface for a virtual server. The left sidebar shows the navigation menu with 'Local Traffic' selected. The main panel shows the configuration for the virtual server 'ra-vrops-vip-http'.

General Properties:

- Name: ra-vrops-vip-http
- Partition / Path: Common
- Description:
- Type: Performance (Layer 4)
- Source Address: 0.0.0.0/0
- Destination Address:
- Service Port: 80 HTTP
- Notify Status to Virtual Address:
- Availability: Available (Enabled) - The virtual server is available
- Synccookie Status: Off
- State: Enabled

Configuration: Basic

- Protocol: TCP
- Protocol Profile (Client): fastL4
- HTTP Profile: http
- VLAN and Tunnel Traffic: Enabled on
- VLANs and Tunnels: Selected (Common SiteA-Res1), Available (Common SiteA-Mgmt, Common SiteA-Res2, http-tunnel, socks-tunnel)
- Source Address Translation: Auto Map

Acceleration: Basic

- Rate Class: None
- SPOY Profile: None

Buttons: Update, Delete

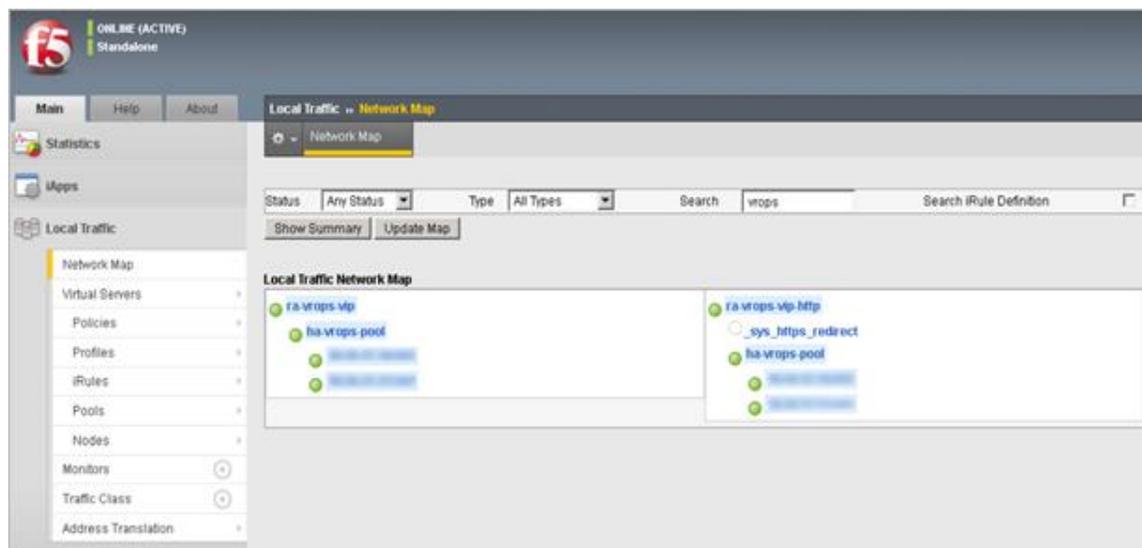
Verify Component and Pool Status

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.

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:

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.



Citrix NetScaler Installation & Configuration

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 article <http://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. Please refer to the articles below for more information.

[Configure a certificate for use with vRealize Operations Manager](#)

[FAQ: Key Sizes/Certificates Supported by NetScaler](#)

Configure Health Monitors

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.

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	ONLINE	443	yes
vrops_epops	HTTP-EVC	16 sec.	15 sec.	3	GET /epops-webapp/health-check	ONLINE	443	yes

Example:

Configure Monitor

Name
wrops_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 Monitor

Name
vrops_https

Type
HTTP-ECV

Standard Parameters Special Parameters

Send String
GET /suite-api/api/deploym

Receive String
ONLINE

Custom Header

OK Close

Configure Service Groups

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.

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:

Load Balancing Service Group

Basic Settings ✎

Name ha-vrops-prod_443 Protocol SSL_BRIDGE State ENABLED Effective State ● Up Traffic Domain 0	Cache Type SERVER Cacheable NO Health Monitoring YES AppFlow Logging ENABLED Number of Active Connections 0 AutoScale Mode -
--	---

Service Group Members

4 Service Group Members >

Settings ✎ ✕

SureConnect OFF Surge Protection OFF Use Proxy Port YES Down State Flush ENABLED	Use Client IP NO Client Keep-alive NO TCP Buffering YES Client IP DISABLED Header AutoScale Mode -
---	--

Monitors ✕

1 Service Group to Monitor Binding >

Configure Virtual Servers

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.

NAME	PROTOCOL	DESTINATION ADDRESS	PORT	LOAD BALANCING METHOD	SERVICE GROUP BINDING
ha-vrops-prod-VIP_80	HTTP	10.23.90.18	80	Leastconnection	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:

Load Balancing Virtual Server | [Export as a Template](#)

Basic Settings

Name	ha-vrops-prod-VIP_443	Listen Priority	-
Protocol	SSL_BRIDGE	Listen Policy Expression	NONE
State	Up	Range	1
IP Address	10.23.90.18	Redirection Mode	IP
Port	443	RHI State	PASSIVE
Traffic Domain	0	AppFlow Logging	ENABLED

Services and Service Groups

- No Load Balancing Virtual Server Service Binding >
- 1 Load Balancing Virtual Server ServiceGroup Binding >

Method

Load Balancing Method	LEASTCONNECTION	New Service Startup Request Rate	0
Backup LB Method	ROUNDROBIN	New Service Request unit	PER_SECOND
		Increment Interval	-

Done

Configure Persistence Group

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.

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:

Configure Persistency Group

Group Name

Persistence*

IPv4 Netmask

IPv6 Mask Length

Time-out

Backup Persistence*

Virtual Server Name*

Configured (2) Remove All

ha-vrops-prod-VIP_80	-	+ Add
ha-vrops-prod-VIP_443	-	

NSX-V Installation & Configuration

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.

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. Accelerated mode uses L4 and LVS and non-accelerated mode uses L7

Install and Configure Edge for Load Balancing

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

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 from NSX-V 6.2.0:

Edit Load balancer global configuration
?

Enable Load Balancer

Enable Acceleration

Logging

Log Level: Info ▼

Enable Service Insertion

Service Definition: ▼

Service Configuration: ▼

Deployment Specification: ▼

Runtime NICs
Attributes
Typed Attributes

+
✎
✕
⇅
⇄

Name	Connected To	ConnectivityType	IP Address	Subnet Mask	Gateway Address

OK
Cancel

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.

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 (+) 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: HTTPS**
9. Select **Enable SSL Passthrough**.
10. Select **Persistence** as **Source IP**.

11. Enter **1800** for **Expires in (seconds)**.
12. Select **Ignore** for **Client Authentication**.
13. Click **OK** to save the Profile

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:

The screenshot shows the 'Edit Profile' dialog box with the following configuration:

- Name: vrops_https
- Type: HTTPS
- Enable SSL Passthrough:
- HTTP Redirect URL: (empty)
- Persistence: Source IP
- Cookie Name: (empty)
- Mode: (empty)
- Expires in (Seconds): 1800
- Insert X-Forwarded-For HTTP header:
- Enable Pool Side SSL:
- Virtual Server Certificates: Pool Certificates
- Service Certificates: CA Certificates, CRL
- Table with columns: Common Name, Issuer, Validity
- Cipher: (empty)
- Client Authentication: Ignore
- Buttons: OK, Cancel

Add Service Monitoring

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.

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 (+) 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.

NAME	INTERVAL	TIMEOUT	RETRIES	TYPE	METHOD	URL	RECEIVE :
vROps_Monitor	5	16	3	HTTPS	GET	/suite-api/api/deployment/node/status	ONLINE (upper case)
EPPOS_Monitor	5	16	3	HTTPS	GET	/epops-webapp/health-check	ONLINE (upper case)

Example:

Edit Service Monitor

Name: * vROPs_Monitor

Interval: 5 (seconds)

Timeout: 16 (seconds)

Max Retries: 3

Type: HTTPS

Expected:

Method: GET

URL: /suite-api/api/deployment/node/status

Send:

Receive: ONLINE

Extension:

OK Cancel

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.

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. Enter a name for the load balancer pool. For example: vROps_Pool.
7. (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. Click the Add (+) icon to add your member servers and the required information:
 - a. Name
 - b. IP Address
 - c. Weight: 1
 - d. Monitor Port: 443

- e. Port: 443
- f. Max Connections: *Set the limit based on the NSX-V LB sizing*
- g. Min Connections: 8

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	8	8
EPOps_Pool	LEASTCONN	EPOps_Monitor	EPOps_Node1	x.x.x.x	1	443	443	8	8

Example:

Edit Pool

Name: * vROps_Pool

Description:

Algorithm: LEASTCONN

Algorithm Parameters:

Monitors: vROps_Monitor

Members:

Enabled	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connection.
✓	ra-vrops...		1	443	443	8	8
✓	ra-vrops...		1	443	443	8	8

Transparent

OK Cancel

Add Virtual Servers

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

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 (+) 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 **HTTPS**.
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.

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: erops_http and EPOPS_Pool.

Example:

The screenshot shows the 'Edit Virtual Server' dialog box with the following configuration:

- Enable Virtual Server:**
- Enable Acceleration:**
- Application Profile:** vrops_https
- Name:** * vrops_https
- Description:** (empty)
- IP Address:** * (redacted) [Select IP Address](#)
- Protocol:** HTTPS
- Port / Port Range:** * 443
- Default Pool:** vROPS_POOL
- Connection Limit:** 0
- Connection Rate Limit:** 0 (CPS)

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

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 () 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>/vcops-web-ent/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:

Edit Profile

Name:

Type:

Enable SSL Passthrough

HTTP Redirect URL:

Persistence:

Cookie Name:

Mode:

Expires in (Seconds):

Insert X-Forwarded-For HTTP header

Enable Pool Side SSL

Virtual Server Certifica... Pool Certificates

Service Certificates CA Certificates CRL

Common Name	Issuer	Validity

Cipher:

Client Authentication:

OK Cancel

Configure the Virtual Server for HTTPS Redirect

You can configure the virtual server for HTTPS redirect.

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 (**+**) 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:

The screenshot shows the 'Edit Virtual Server' dialog box with the following configuration:

- Enable Virtual Server
- Enable Acceleration
- Application Profile: vrops_redirect
- Name: * vrops_redirect
- Description: (empty)
- IP Address: * (redacted) [Select IP Address](#)
- Protocol: HTTP
- Port / Port Range: * 80
- Default Pool: NONE
- Connection Limit: 0
- Connection Rate Limit: 0 (CPS)

Verify Component and Pool Status

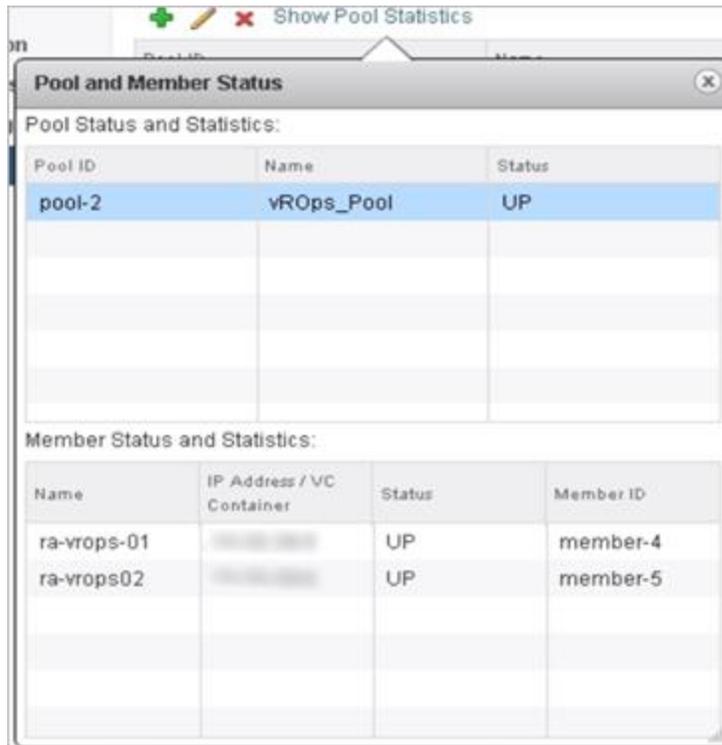
You can verify the status of the components running on the load balancer and you can check the status of the pools from inside the UI of the vSphere Web Client.

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. Select the pool you want to verify. For example: vROps_Pool.
7. Click **Show Pool Statistics**. A **Pool and Member Status** pop-up window appears.

8. Select a pool ID. For example: vROps_Pool.

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

Example:



The screenshot shows a dialog box titled "Pool and Member Status" with a close button (X) in the top right corner. The dialog is divided into two sections: "Pool Status and Statistics:" and "Member Status and Statistics:".

Pool Status and Statistics:

Pool ID	Name	Status
pool-2	vROps_Pool	UP

Member Status and Statistics:

Name	IP Address / VC Container	Status	Member ID
ra-vrops-01	10.10.10.10	UP	member-4
ra-vrops02	10.10.10.10	UP	member-5

NSX-T Installation & Configuration

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

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:

1. Redirect HTTP to HTTPS
2. 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)

Edit HTTP Profile
 

Name *

Description

Redirection

None
 HTTP Redirect
 HTTP to HTTPS Redirect

Profile Configuration

X-Forwarded-For

Advanced Properties

Connection Idle Timeout (sec)

Request Header Size (bytes)

Request Body Size (bytes)

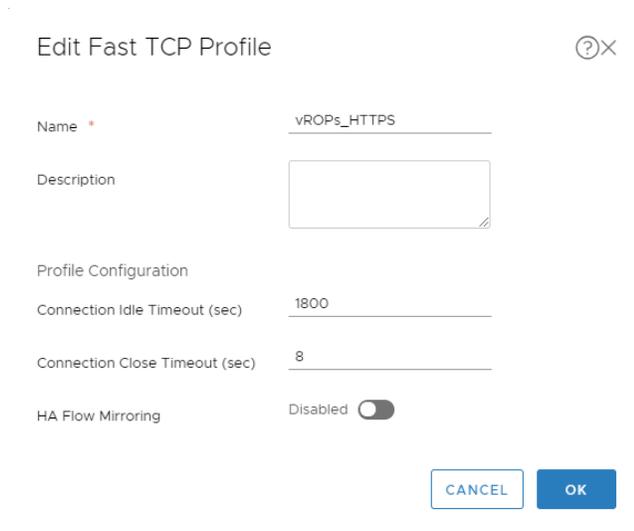
If not specified, it is unlimited

NTLM Authentication Disabled

CANCEL
OK

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)



Edit Fast TCP Profile ? ×

Name *

Description

Profile Configuration

Connection Idle Timeout (sec)

Connection Close Timeout (sec)

HA Flow Mirroring Disabled

Configure Persistence Profile

- 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 New Source IP Persistence Profile (?) X

Name *

Description

Profile Configuration

Share Persistence Disabled

Persistence Entry Timeout (seconds)

HA Persistence Mirroring Disabled

Purge Entries when Full Enabled

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 with the following values:

1. Request Method
GET
2. Request URL
/suite-api/api/deployment/node/status
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	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:

Edit Active Health Monitor

1 Monitor Properties

2 Health Check Parameters

Monitor Properties

Name * vROPS_MONITOR

Description vROPS_MONITOR

Health Check Protocol * LbHttpsMonitor

Monitoring Port

Monitoring interval (sec) * 5

Fall Count * 3

Rise Count * 3

Timeout Period (sec) * 16

CANCEL NEXT

Edit Active Health Monitor

1 Monitor Properties

2 Health Check Parameters

SSL and HTTP Health Check Parameters

Configure the SSL Connection sent before the HTTP Request

SSL Protocols

Available(3)	Selected(2)
<input type="checkbox"/> TLS_V1 <input type="checkbox"/> TLS_V1_1 <input type="checkbox"/> TLS_V1_2	<input type="checkbox"/> TLS_V1_1 <input type="checkbox"/> TLS_V1_2

SSL Ciphers

Available(31)	Selected(30)
<input type="checkbox"/> TLS_ECDHE_RSA_WITH_AES_128_GC... <input type="checkbox"/> TLS_ECDHE_RSA_WITH_AES_256_GC...	<input type="checkbox"/> TLS_ECDHE_RSA_WITH_AES_128... <input type="checkbox"/> TLS_ECDHE_RSA_WITH_AES_256...

CANCEL BACK FINISH

- Example for vROPS_Monitor

Edit Active Health Monitor

1 Monitor Properties

2 Health Check Parameters

SSL and HTTP Health Check Parameters

HTTP Request URL

HTTP Request Version

HTTP Request Headers

+ ADD DELETE

Header Name	Header Value
	

HTTP Request Body

HTTP Response Configuration

HTTP Response Code

Specify response codes separated by comma (support up to 64 codes)

HTTP Response Body

Regular Expression is not allowed

CANCEL BACK FINISH

- Example for EPOPS_Monitor

Edit Active Health Monitor

1 Monitor Properties

2 Health Check Parameters

SSL and HTTP Health Check Parameters

HTTP Request URL

HTTP Request Version

HTTP Request Headers

+ ADD DELETE

Header Name	Header Value
	

HTTP Request Body

HTTP Response Configuration

HTTP Response Code

Specify response codes separated by comma (support up to 64 codes)

HTTP Response Body

Regular Expression is not allowed

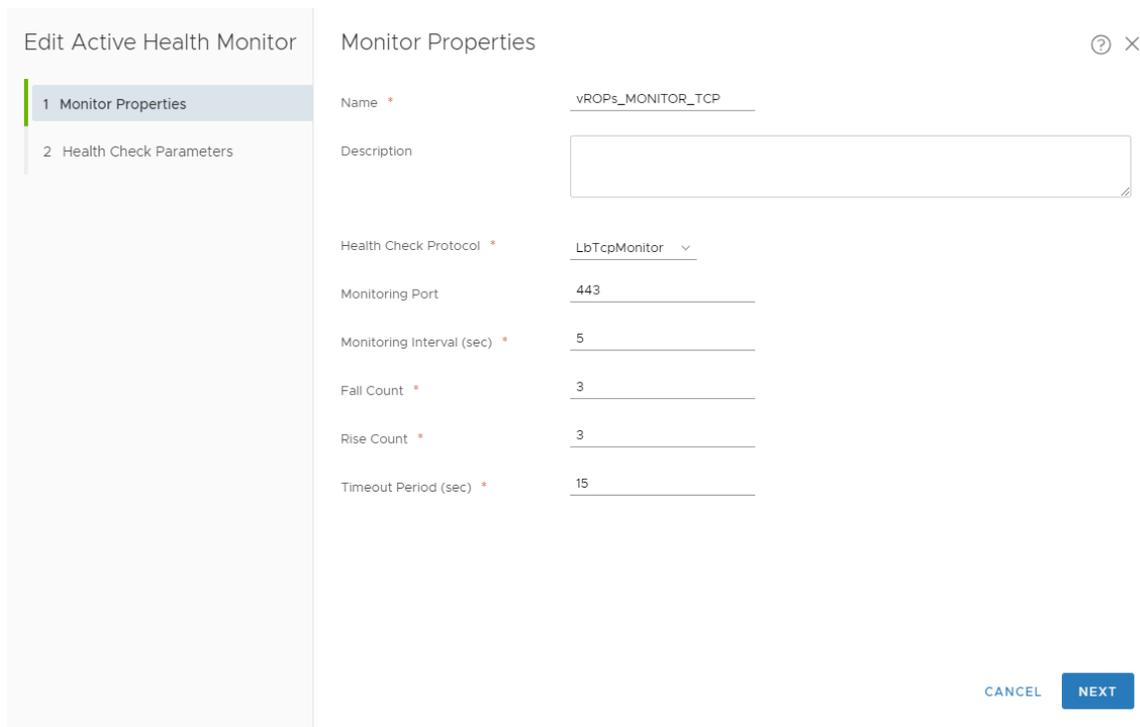
CANCEL BACK FINISH

- **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



Edit Active Health Monitor

Monitor Properties

1 Monitor Properties

2 Health Check Parameters

Name * vROPs_MONITOR_TCP

Description

Health Check Protocol * LbTcpMonitor

Monitoring Port 443

Monitoring Interval (sec) * 5

Fall Count * 3

Rise Count * 3

Timeout Period (sec) * 15

CANCEL **NEXT**

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
 - f. 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
EOPS_POOL	LEASTCONN	EOPS_Monitor	EOPS_NODE1	x.x.x.x	1	443	ENABLED

Edit Server Pool

- 1 General Properties
- 2 SNAT Translation
- 3 Pool Members
- 4 Health Monitors

General Properties

Name *

Description

Load Balancing Algorithm

Advanced Properties

TCP Multiplexing Disabled

Maximum Multiplexing Connections

?

×

CANCEL

Edit Server Pool

- 1 General Properties
- 2 SNAT Translation**
- 3 Pool Members
- 4 Health Monitors

SNAT Translation ? ×

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

[CANCEL](#) [BACK](#) [NEXT](#)

Edit Server Pool

- 1 General Properties
- 2 SNAT Translation
- 3 Pool Members
- 4 Health Monitors

Pool Members ? X

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

	Name	IP	Port	Weight	State	Backup Member	Max. Concurrent Connection
<input type="radio"/>	Master	[REDACTED]	443	1	ENABLED	<input type="checkbox"/>	
<input type="radio"/>	Replica	[REDACTED]	443	1	ENABLED	<input type="checkbox"/>	
<input type="radio"/>	Data1	[REDACTED]	443	1	ENABLED	<input type="checkbox"/>	
<input type="radio"/>	Data2	[REDACTED]	443	1	ENABLED	<input type="checkbox"/>	

4 Pool Members

Edit Server Pool

- 1 General Properties
- 2 SNAT Translation
- 3 Pool Members
- 4 Health Monitors

Health Monitors ? X

Minimum Active Members

Active Health Monitor [Create A New Active Monitor](#)

Passive Health Monitor [Create A New Passive Monitor](#)

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

Edit Virtual Server

- 1 General Properties
- 2 Virtual Server Identifiers
- 3 Server Pool and Rules
- 4 Load Balancing Profiles
 - A Persistence Profiles
 - B Client Side SSL
 - C Server Side SSL

General Properties



Name * VROPS-NSXT2-HTTP

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 v

Application Profile * VROPS_HTTP_to_HTTPS v

Access Log Disabled

CANCEL NEXT

Edit Virtual Server

- 1 General Properties
- 2 Virtual Server Identifiers
- 3 Server Pool and Rules
- 4 Load Balancing Profiles
 - A Persistence Profiles
 - B Client Side SSL
 - C Server Side SSL

Virtual Server Identifiers



IP Address * 192.168.207.200

Port * 80

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 80

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 and Rules
4 Load Balancing Profiles

A Persistence Profiles
B Client Side SSL
C Server Side SSL

Persistence Profiles ? ×

Persistence Profiles

Persistence Enabled

Source Ip Persistence VROPS_PERSISTENCE
[Create A New Source Persistence Profile](#)

Cookie Persistence
[Create A New Cookie Persistence Profile](#)

CANCEL BACK NEXT

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
- Virtual Server Identifiers
- Server Pool
- Load Balancing Profiles

General Properties

Name *

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

Application Profile *

Access Log Enabled

CANCEL NEXT

Edit Virtual Server

- General Properties
- Virtual Server Identifiers
- Server Pool
- Load Balancing Profiles

Virtual Server Identifiers

IP Address *

Port *

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

Protocol

Advanced Properties

Maximum Concurrent Connection

Maximum New Connection Rate

Default Pool Member Port

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

You need to specify a load-balancer configuration parameter and 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)

Name *

Description

Load Balancer Size *

Select from one of the three available choices of size for the Load Balancer

 Warning: Changing the Load Balancer Size will disrupt the active traffic on the Load Balancer. Service Disruption is to be expected.

<input checked="" type="radio"/> SMALL	<input type="radio"/> MEDIUM	<input type="radio"/> LARGE
Virtual Servers  10	Virtual Servers  100	Virtual Servers  1000
Pool Members  30	Pool Members  300	Pool Members  3000
CPU 2 Memory 4GB	CPU 4 Memory 8GB	CPU 12 Memory 16GB

Error Log Level *

The screenshot shows the vRealize Operations Manager interface. On the left, a navigation sidebar lists various tools and services, with 'Load Balancing' selected and 'Load Balancers' highlighted. The main content area displays a list of load balancers, each with a checkbox and a status icon. The 'VROPS' load balancer is selected, indicated by a blue checkmark and a blue highlight. To the right of the list, a configuration panel for the selected load balancer is visible, showing sections for 'Summary', 'Attachment', and 'Manage Tags'. The 'Attachment' section includes fields for 'Tier-1 Router', 'Active Edge Node', and 'Standby Edge Node'.

Attach to a Logical Router ✕

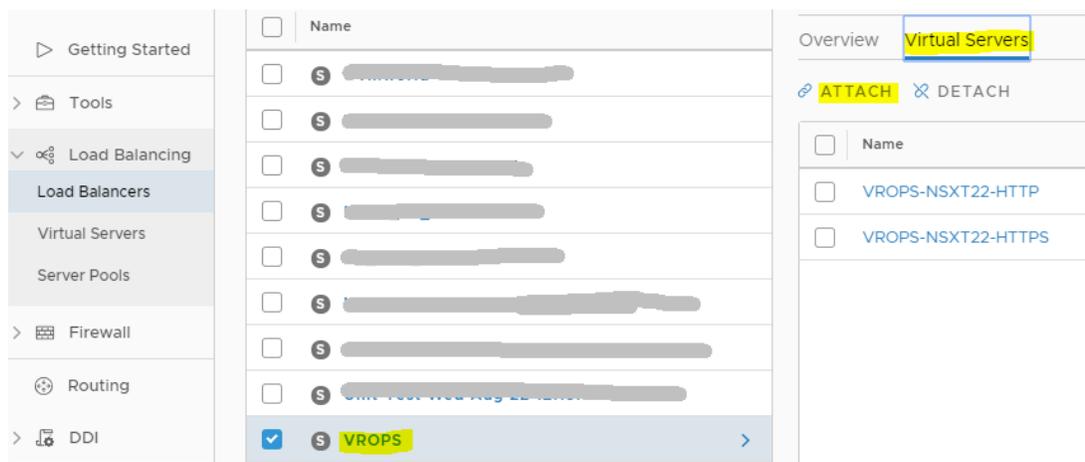
Select the Router to which the Load Balancer VROPS-NSXT22 is to be attached. Only Tier-1 Routers in 'Active Standby' are currently supported.
 Note: The Load Balancer can only be Enabled if it had a Virtual Server associated with it.

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)

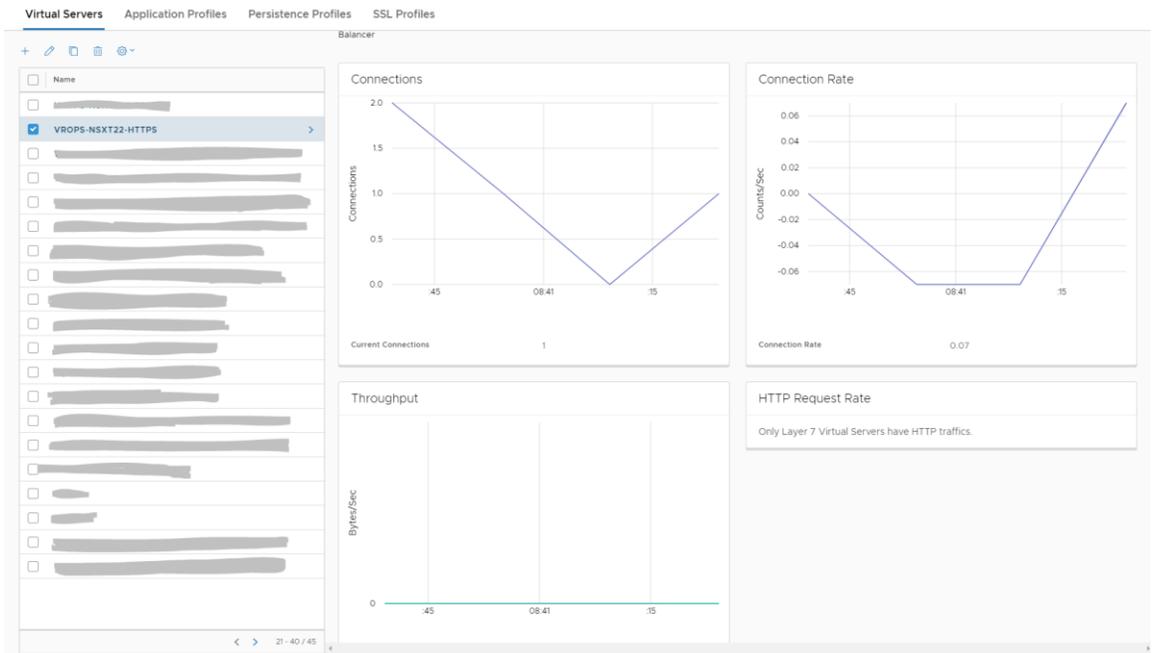
vROPS-POOL

Overview Virtual Servers Pool Members **Pool Member Statistics**

Display Statistics from Load Balancer VROPS

IP:Port	Status	Current Sessions	Max Sessions	Bytes in	Bytes out	Http Request Rate
10.10.10.443	↑ UP	0	19	0	0	0
10.10.10.443	↑ UP	0	18	0	0	0
10.10.10.443	↑ UP	0	0	0	0	0
10.10.10.443	↑ UP	0	13	0	0	0

- 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.

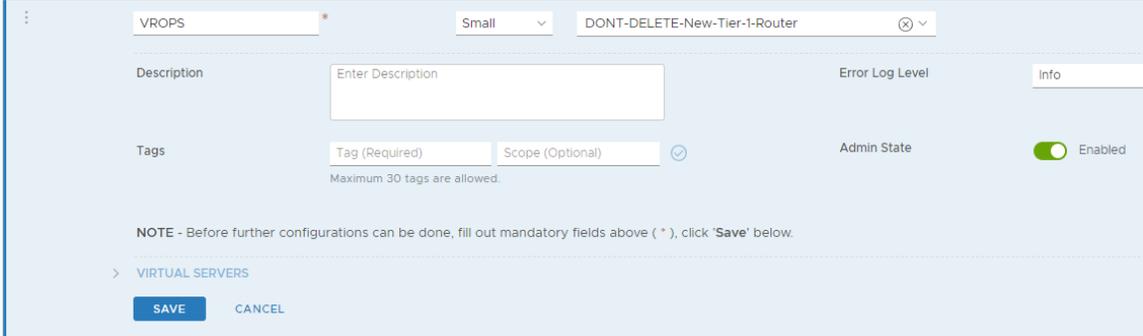


NSX-T Version 2.4

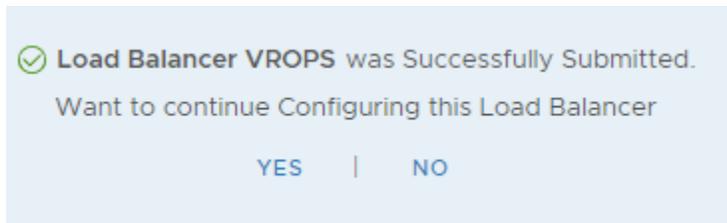
Configure Load Balancer

You need to specify a load-balancer configuration parameter and 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:

3. Redirect HTTP to HTTPS
4. 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 **APPLICATION** ▾
- Click the **Add** (**ADD APPLICATION PROFILE** ▾) icon and choose **HTTP Profile**.
- Choose a name for the profile and enter parameters (please refer to the example below)

The screenshot shows the configuration form for an application profile. The profile name is 'vROps_HTTP_to_HTTPS' and the type is 'HTTP'. The 'Request Header Size' is set to '15'. The 'Description' field contains 'Enter Description'. The 'X-Forwarded-For' dropdown is set to 'Insert'. The 'Redirection' dropdown is set to 'HTTP to HTTPS Redirect'. The 'NTLM Authentication' toggle is 'Disabled'. The 'Tags' section has two input fields: 'Tag (Required)' and 'Scope (Optional)', with a checkmark icon and the note 'Maximum 30 tags are allowed.' At the bottom, there are 'SAVE' and 'CANCEL' buttons.

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

- 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 with the following values:
 3. HTTP Method
GET
 4. HTTP Request URL
/suite-api/api/deployment/node/status (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	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:

The screenshot shows a configuration form for a health check. The form is titled "vROPs_MONITOR" and "HTTPS". The "Interval" is set to 5 and "Timeout" is set to 16. The "Description" field is empty. The "Tags" field has a "Tag (Required)" and "Scope (Optional)" dropdown, with a checkmark icon and the text "Maximum 30 tags are allowed." The "Fall Count" is set to 3 and "Rise Count" is set to 3. There are links to "Configure" for "HTTP Request", "HTTP Response", and "SSL Configuration". At the bottom, there are "SAVE" and "CANCEL" buttons.

- 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 X 204 X 301 X

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

HTTP Request URL

HTTP Request Version

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

1 or more response code

HTTP Response Body

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)

The screenshot shows the configuration page for a server pool named 'vROPs-POOL'. The algorithm is 'Least Conn'. The SNAT Translation Mode is 'Automap'. The Active Monitor is 'vROPs_MONITOR'. The TCP Multiplexing is disabled. The Max Multiplexing Connections is 6. The Min Active Members is 1. There are fields for Tag (Required) and Scope (Optional). The page has 'SAVE' and 'CANCEL' buttons at the bottom.

- Add the pool members via **Select Members** (vRealize Operations Manager data nodes IP addresses and Port)
 - Name
 - IP Address
 - Weight: 1
 - Port: 443
 - State: ENABLED

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 Server Pool Members



Server Pool - vROPS-POOL

 Enter individual members Select a group[ADD MEMBER](#)

Name	IP	Port	Weight	State	Backup Member	Max Concurrent Connections
: DATA3	[REDACTED]	443	1	Enabled	<input type="radio"/> Disabled	
: DATA2	[REDACTED]	443	1	Enabled	<input type="radio"/> Disabled	
: DATA1	[REDACTED]	443	1	Enabled	<input type="radio"/> Disabled	

[CANCEL](#)[APPLY](#)

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

Name	IP Address	Ports	Type	Load Balancer	Server Pool
VROPS-NSXT22-HTTPS *	192.168.207.10 * e.g. 10.10.10.10	443 x * Enter Ports or Port	L4 TCP	VROPS ⊗ v	VROPS-POOL ⊗ v
Description	Enter Description		Application Profile *	vROPS_HTTPS	
Persistence	Source IP				
Source IP *	VROPS_PERSISTENCE ⊗ v				
Additional Properties					
Max Concurrent Connections	Unlimited		Max New Connection Rate	Unlimited	
Sorry Server Pool	Select Server Pool		Default Pool Member Ports	Enter Ports or Port Ranges (e.g. 8080, 80-90, 443)	
Admin State	<input checked="" type="checkbox"/> Enabled		Access Log	<input type="checkbox"/> Disabled	
Tags	Tag (Required) Scope (Optional) ✓ Maximum 30 tags are allowed.				
<input type="button" value="SAVE"/> <input type="button" value="CANCEL"/>					

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

Name	IP Address	Ports	Type	Load Balancer	Server Pool
VROPS-NSXT22-HTTP *	192.168.207.10 * e.g. 10.10.10.10	80 *	L7 HTTP	VROPS ⊗ v	Select Server P...
Description	Enter Description		Application Profile *	vROPS_HTTP_to_HTTPS	
Persistence	Disabled		SSL Configuration	Configure	
Load Balancer Rules					
Additional Properties					
<input type="button" value="SAVE"/> <input type="button" value="CANCEL"/>					