

vRealize Automation 8.3 Load Balancing Guide

04 February 2021

vRealize Automation 8.3

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

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	vRealize Automation and vRealize Orchestrator Load Balancing	5
2	Load Balancing Concepts	6
	SSL Pass-Through	6
	Load Balancer Notifications	6
	One-Arm and Multi-Arm Topologies	7
3	Prerequisites for Configuring Load Balancers for vRealize Automation	8
	Complete the vRealize Automation/ vRealize Orchestrator Initial Installation	9
4	Configuring NSX-V	10
	Configure Global Settings	10
	Configure Application Profiles	12
	Configure Service Monitoring	13
	Configure Server Pools	14
	Configure Virtual Servers	16
5	Configuring NSX-T	18
	Configure NSX-T Application Profiles	18
	Configure NSX-T Active Health Monitor	19
	Configure NSX-T Server Pools	22
	Configure NSX-T Virtual Servers	23
	Configure Load Balancer	24
	Add Virtual Servers to Load Balancer	25
6	Configuring F5 Big-IP LTM	27
	Configure Monitors	27
	Configure F5 Server Pools	29
	Configure F5 Virtual Servers	30
7	Configuring Citrix ADC (NetScaler ADC)	33
	Configure Citrix Monitors	33
	Configure Citrix Service Groups	36
	Configure Citrix Virtual Servers	37
8	Configuring AVI Load Balancer	39
	Create Pool	39
	Create an Active Monitor	40

[Configure Virtual Service](#) 42

9 Troubleshooting 45

[Errors during vRealize Automation installation when using NSX-V as a load-balancer for Workspace ONE](#) 45

[Provisioning Failures When Using OneConnect with F5 BIG-IP](#) 46

[F5 BIG-IP License Limits Network Bandwidth](#) 46

vRealize Automation and vRealize Orchestrator Load Balancing

1

This document describes the load balancing configuration of vRealize Automation and vRealize Orchestrator in a distributed and highly available cluster deployment using VMware NSX, F5 Networks BIG-IP (F5), and Citrix NetScaler technologies.

This document is not an installation guide, but rather a configuration guide that supplements the vRealize Automation and vRealize Orchestrator installation and configuration documentation available in the [VMware vRealize Automation product documentation](#) and [VMware vRealize Orchestrator product documentation](#).

This information is for the following products and versions.

Table 1-1.

Product	Version
NSX-T	2.4, 2.5, 3.0
NSX-V	6.2.x, 6.3.x, 6.4.x
F5 BIG-IP LTM	11.x, 12.x, 13.x, 14.x, 15.x
Citrix NetScaler ADC	10.5, 11.x, 12.x, 13.x
vRealize Automation	8.0, 8.1, 8.2
vRealize Orchestrator	8.0, 8.1

Refer to the [VMware Product Interoperability Matrices](#) for more details.

Load Balancing Concepts

2

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

Follow your organization's policy for backing up load balancers, keeping in mind the preservation of network topology and VMware products backup planning.

This chapter includes the following topics:

- [SSL Pass-Through](#)
- [Load Balancer Notifications](#)
- [One-Arm and Multi-Arm Topologies](#)

SSL Pass-Through

SSL pass-through is used with the load balancing configurations.

SSL pass-through is used for these reasons:

- Ease of deployment
 - Not having to deploy the vRealize Automation, or vRealize Orchestrator certificates to the load balancer simplifies deployment and reduces complexity.
- No operational overhead
 - At the time of certificate renewal, no configuration changes are required to the load balancer.
- Ease of communication
 - The individual host names of the load-balanced components are the subject alternate name field of the certificates, so the client can easily communicate with the load balanced nodes.

Load Balancer Notifications

It is a recommended practice to enable notifications any time a vRealize Automation or vRealize Orchestrator node in a server pool goes down.

VMware NSX Data Center supports enabling notifications when an alert is raised in vRealize Operations Manager and vRealize Network Insight. Refer to the vRealize Operations Manager and vRealize Network Insight documentation.

For NetScaler, configure specific SNMP traps and an SNMP manager to send alerts. Consult the NetScaler documentation for information on SNMP configuration.

You can set up email notification with F5 using these methods:

- [Configuring the BIG-IP system to deliver locally generated email messages](#)
- [Configuring custom SNMP traps](#)
- [Configuring alerts to send email notifications](#)

One-Arm and Multi-Arm Topologies

One-arm and multi-arm deployments route load balancer traffic differently.

In one-arm deployment, the load balancer is not physically in line of the traffic, which means that the load balancer's ingress and egress traffic goes through the same network interface. Traffic from the client through the load balancer is network address translated (NAT) with the load balancer as its source address. The nodes send their return traffic to the load balancer before being passed back to the client. Without this reverse packet flow, return traffic would try to reach the client directly, causing connections to fail.

In a multi-arm configuration, the traffic is routed through the load balancer. The end devices typically have the load balancer as their default gateway.

The most common deployment is a one-arm configuration. The same principles apply to multi-arm deployments, and they both work with F5 and NetScaler.

For this document, the vRealize Automation and vRealize Orchestrator components are deployed in a one-arm configuration. Multi-arm deployments are also supported, and their configuration are generally similar to the one-arm configuration.

One-Arm Configuration:



Prerequisites for Configuring Load Balancers for vRealize Automation

3

Before configuring load balancers, perform these prerequisites.

- NSX-V/T - Before you can start a high-availability implementation of vRealize Automation or vRealize Orchestrator using NSX-V/T as a load balancer, ensure that your NSX-V/T topology is configured and that your version of NSX-V/T is supported. This document covers the load balancing aspect of an NSX-V/T configuration and assumes that NSX-V/T is configured and validated to work properly on the target environment and networks. To verify that your version is supported, see the product [interoperability matrix](#).
- F5 BIG-IP LTM - Before you can start a high-availability implementation of vRealize Automation or vRealize Orchestrator using F5 LTM load balancer, ensure that the load balancer is installed and licensed and that the DNS server configuration is complete.
- NetScaler - Before you can start a high-availability implementation of vRealize Automation or vRealize Orchestrator using the NetScaler load balancer, ensure that NetScaler is installed and has at least a Standard Edition license.
- Certificates - Request Certificate Authority (CA) signed certificate containing the load-balancer fully qualified domain name and the hostnames of the cluster nodes in the SubjectAltNames section. This configuration enables the load balancer to serve traffic without SSLerrors.
- Identity provider - Starting with vRealize Automation 8.0, the Identity Provider is Workspace ONE Access, which is deployed external to the vRealize Automation appliances and cluster.

For more information on installation and configuration, see vRealize Automation documentation on [docs.vmware.com](#).

If necessary, an external vRealize Orchestrator cluster can be configured to work with the vRealize Automation system. This can be done after the vRealize Automation system is up and running. However, a vRealize Automation Highly Available setup already includes an embedded vRealize Orchestrator cluster.

This chapter includes the following topics:

- [Complete the vRealize Automation/ vRealize Orchestrator Initial Installation](#)

Complete the vRealize Automation/ vRealize Orchestrator Initial Installation

You must configure your load balancer before completing the initial installation of vRealize Automation, vRealize Orchestrator.

During the installation process of vRealize Automation or vRealize Orchestrator, a load balancer typically will route half of the traffic to the secondary nodes, which will not yet be configured, causing the installation to fail. To avoid these failures and to complete the initial installation of vRealize Automation or vRealize Orchestrator, you must perform these steps.

Procedure

- 1 Configure the F5, NSX, or NetScaler load balancer. See [Chapter 6 Configuring F5 Big-IP LTM](#), [Chapter 5 Configuring NSX-T](#), and [Chapter 7 Configuring Citrix ADC \(NetScaler ADC\)](#).
- 2 Turn off the health monitors or change them temporarily to default ICMP, and ensure traffic is still forwarding to your primary nodes.
- 3 Disable all secondary nodes from the load balancer pools.
- 4 Install and configure all system components as detailed in vRealize Automation / vRealize Orchestrator Installation and Configuration documentation.
- 5 When all components are installed, enable all non-primary nodes on the load balancer.
- 6 Configure the load balancer with all monitors (health checks) enabled.

After you complete this procedure, update the monitor that you created in [Configure Monitors](#).

- 7 Ensure that all nodes are in the expected state with the health monitor enabled in the load balancer after installation. The pool, service groups, and virtual server of the virtual appliance nodes should be available and running. All virtual appliance nodes should be available, running, and enabled.

Configuring NSX-V

4

You can deploy a new NSX-V Edge Services Gateway or reuse an existing one. However, it must have network connectivity to and from the vRealize components being load balanced.

Note Refer to the [VMware Workspace One](#) load-balancing documentation in order to configure highly-available identity provider for vRealize Automation.

This chapter includes the following topics:

- [Configure Global Settings](#)
- [Configure Application Profiles](#)
- [Configure Service Monitoring](#)
- [Configure Server Pools](#)
- [Configure Virtual Servers](#)

Configure Global Settings

Configure global settings using these steps.

Procedure

- 1** Log in to the NSX-V, click **Manager > Settings** and select **Interfaces**.
- 2** Select your Edge device from the list.
- 3** Click **vNIC#** for the external interface that hosts the virtual IP addresses and click the **Edit** icon.

- 4 Select the appropriate network range for the NSX-V Edge and click the **Edit** icon.

Edit Interface | nic0

Basic Advanced

vNIC# 0

Name ^{*} nic0

Type ☐ Internal ☒ Uplink ☐ Trunk

Connected To ^{*} Prod-01

Connectivity Status ☒ Connected

Configure Subnets

+ ADD DELETE

<input type="checkbox"/>	Primary IP Address	Secondary IP Addresses	Subnet Prefix Length
<input type="checkbox"/>	192.168.208.102		24

1 items

CANCEL SAVE

- 5 Add the IP addresses assigned to the virtual IPs and click **Save**.
- 6 Click **Ok** to exit the interface configuration page.
- 7 Navigate to the **Load Balancer** tab and click the **Edit** icon.
- 8 Select **Enable Load Balancer** and **Logging**, if necessary, and click **Save**.

Edit Load Balancer Global Configuration

Load Balancer ☒ Enable

Acceleration ☐ Disable

Logging ☒ Enable

Log Level

CANCEL SAVE

Configure Application Profiles

It is required to add application profiles for vRealize Automation and for an external vRealize Orchestrator (optional).

Procedure

- 1 Click **Application Profiles** in the left pane.
- 2 Click the **Add** icon to create the application profiles required for the specific product as outlined in this table. Use the default value if nothing is specified.

Table 4-1. Application Profiles

Name	Type	Persistence	Expires In
vRealize Automation	SSL Passthrough	None	None
vRealize Orchestrator	SSL Passthrough	None	None
Note Use only for external vRealize Orchestrator instances.			

Results

The completed configuration should look similar to this screen:

New Application Profile

Application Profile Type

SSL Passthrough

General

Client SSL

Server SSL

Name *

vRealize Automation / vRealize Orchestrator VA Web

HTTP Redirect URL

Persistence

None

Cookie Name

Mode

Expires in

(Seconds)

Insert X-Forwarded-For HTTP header

Disable

CANCEL

ADD

Configure Service Monitoring

It is required to add service monitors for vRealize Automation and for an external vRealize Orchestrator (optional).

Procedure

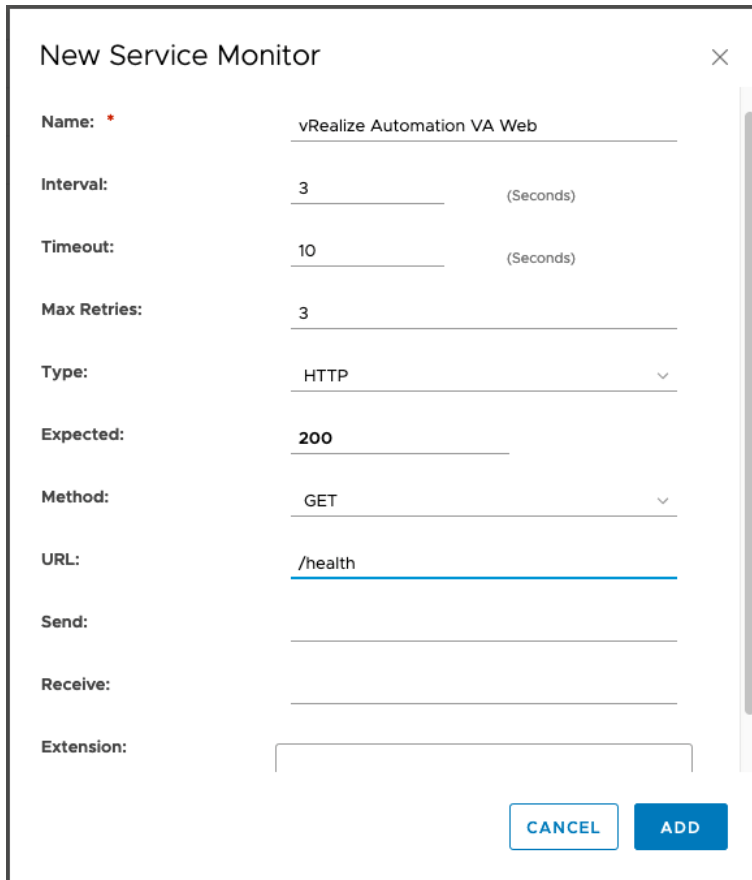
- 1 Click **Service Monitoring** in the left pane.
- 2 Click the **Add** icon to create the service monitors required for the specific product as outlined in this table. Use the default value if nothing is specified.

Table 4-2. Service Monitoring

Name	Interval	Timeout	Retries	Type	Method	URL	Receive	Expected
vRealize Automation	3	10	3	HTTP	GET	/health		200
vRealize Orchestrator	3	10	3	HTTP	GET	/health		200
Note Use only for external vRealize Orchestrator instances.								

Results

The completed configuration should look similar to this screen:



The screenshot shows a 'New Service Monitor' dialog box with the following fields and values:

Field	Value
Name *	vRealize Automation VA Web
Interval	3 (Seconds)
Timeout	10 (Seconds)
Max Retries	3
Type	HTTP
Expected	200
Method	GET
URL	/health
Send	
Receive	
Extension	

At the bottom right, there are two buttons: 'CANCEL' and 'ADD'.

Configure Server Pools

It is required to create server pools for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Click **Pools** in the left pane.

- 2 Click the **Add** icon to create the pools required for the specific product as outlined in this table.

Table 4-3. Server Pools

Pool Name	Algorithm	Monitors	Member Name	IP Address/ vCenter Container	Port	Monitor Port
vRealize Automation	Least connections	vRealize Automation	VA1 VA2 VA	IP Address	443	8008
vRealize Orchestrator	Least connections	vRealize Orchestrator	VA1 VA2 VA3	IP Address	443	8008

Note Use only for external vRealize Orchestrator instances.

Results

The completed configuration should look similar to this screen:

New Pool ×

General

Members

+ ADD

EDIT

DELETE

	Name	IP Address / VC Container	Weight	Monitor Port	Port	Max Connections	Min Connections
<input type="radio"/>	vRA_VA_1	10.10.10.10	1	8008	443		
<input type="radio"/>	vRA_VA_3	10.10.10.12	1	8008	443		
<input type="radio"/>	vRA_VA_2	10.10.10.11	1	8008	443		

1 - 3 of 3 items

CANCEL

ADD

Configure Virtual Servers

It is required to configure virtual servers for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Click **Virtual Servers** in the left pane.
- 2 Click the **Add** icon to create the virtual servers required for the different product as outlined in this table. Use default values if nothing is specified.

Table 4-4. Virtual Servers

Name	Acceleration	IP Address	Protocol	Port	Default Pool	Application Profile		
vRealize Automation	Disabled	IP Address	HTTPS	443	vRealize Automation	vRealize Automation		
vRealize Orchestrator	Disabled	IP Address	HTTPS	443	vRealize Orchestrator	vRealize Orchestrator		
Note Use only for external vRealize Orchestrator instances.								

Results

The completed configuration should look similar to this screen.

New Virtual Server ×

Virtual Server *

☒ Enable

Acceleration *

☐ Disable

Application Profile:

vRealize Automation VA Web ▼

Name: *

vs_vra-va-web_443

Description:

IP Address: *

10.10.10.8 [Select IP Address](#)

Protocol:

HTTPS ▼

Port / Port Range: *

443
e.g.: 9000,9010-9020

Default Pool:

pool_vra-va-web_443 ▼

CANCEL

ADD

Configuring NSX-T

5

Before configuring, the NSX-T must be deployed in the environment and the Tier-1 gateway with the load balancer must have access to the vRealize components over a network.

Note Refer to the [VMware Workspace One](#) load-balancing documentation in order to configure highly-available identity provider for vRealize Automation.

Note NSX-T version 2.3 does not support the HTTPS monitor for the FAST TCP virtual server pool. The HTTPS monitor is supported for NSX-T versions 2.4 and later.

This chapter includes the following topics:

- [Configure NSX-T Application Profiles](#)
- [Configure NSX-T Active Health Monitor](#)
- [Configure NSX-T Server Pools](#)
- [Configure NSX-T Virtual Servers](#)
- [Configure Load Balancer](#)
- [Add Virtual Servers to Load Balancer](#)

Configure NSX-T Application Profiles

You can add an application profile in NSX-T for HTTPS requests.

Procedure

- 1 Navigate to **Networking > Load Balancing > Profiles**.
- 2 Select **Application** as the profile type.
- 3 Click **Add Application Profile** and select **Fast TCP Profile**.
- 4 Enter a name for the profile.

Results

The completed application profile for the HTTPS request should look similar to this screen:

LOAD BALANCERS VIRTUAL SERVERS SERVER POOLS **PROFILES** MONITORS • About

Select Profile Type APPLICATION ▾

ADD APPLICATION PROFILE ▾

Name	Type	Idle Timeout (sec)	HA Flow Mirroring
vRA_HTTPS *	Fast TCP	1800	<input type="checkbox"/> Disabled

Description

Tags ✓
Maximum 30 tags are allowed.

SAVE CANCEL

Connection Close Timeout

Configure NSX-T Active Health Monitor

To configure an active health monitor for NSX-T follow these steps.

Procedure

- 1 Navigate to **Networking > Load Balancing > Monitors**.
- 2 Click **Add Active Monitor** and select **HTTP**.
- 3 Enter a name for the health monitor.

4 Configure the health monitor as outlined in this table:

Table 5-1. Configure Health Monitor

Name	Monitoring Port	Interval	Timeout	Fall Count	Type	Method	URL	Response Code	Response Body
vRealize Automation	8008	3	10	3	HTTP	GET	/health	200	None
vRealize Orchestrator	8008	3	10	3	HTTP	GET	/health	200	None

Note
Use only for external vRealize Orchestrator instance s.

Results

The completed configuration should look similar to these screens.

LOAD BALANCERS VIRTUAL SERVERS SERVER POOLS PROFILES **MONITORS** [About](#)

Select Monitor Type **ACTIVE** ▾

[ADD ACTIVE MONITOR](#) ▾ [COLLAPSE ALL](#)

Name	Protocol	Monitoring Port	Monitoring Interval	Timeout Period (sec)	Server Pools
vRealize Automation VA *	HTTP	8008	3	10	

Description

Fall Count

Tags ☒

Maximum 30 tags are allowed.

Additional Properties ▾

HTTP Request [Configure](#) HTTP Response [Configure](#)

[SAVE](#) [CANCEL](#)


HTTP Request and Response Configuration ×

Active Health Monitor -

HTTP Request Configuration

HTTP Response Configuration

HTTP Method Get ▼HTTP Request URL /healthHTTP Request Version 1.1 ▼ADD

Header Name	Header Value
 Request Header not found	

HTTP Request Body

CANCELAPPLYHTTP Request and Response Configuration ×

Active Health Monitor -

HTTP Request Configuration

HTTP Response ConfigurationHTTP Response Code 200 ×

1 or more response codes

HTTP Response Body

Configure NSX-T Server Pools

You must configure server pools for vRealize Automation, and an external vRealize Orchestrator (optional).

Procedure

- 1 Navigate to **Networking > Load Balancing > Server Pools**.
- 2 Click **Add Server Pool**.
- 3 Enter a name for the pool.
- 4 Configure the pool as outlined in this table:

Table 5-2. Configure Server Pools

Pool Name	Algorithm	Active Monitor	Name	IP	Port
vRealize Automation	Least Connections	vRealize Automation	VA1 VA2 VA3	IP	443
vRealize Orchestrator	Least Connections	vRealize Orchestrator	VA1 VA2 VA3	IP	443
Note Use only for external vRealize Orchestrator instances.					

Results

The completed configuration should look similar to these screens.

The first screenshot shows the 'SERVER POOLS' tab in the vRealize Automation interface. It includes a navigation bar with 'LOAD BALANCERS', 'VIRTUAL SERVERS', 'SERVER POOLS' (selected), 'PROFILES', 'MONITORS', and 'About'. Below the navigation bar is a table with columns: Name, Algorithm, Members/Group, and Virtual Servers. The 'Name' column contains 'pool_vra-va-web_443'. The 'Algorithm' column has a dropdown menu set to 'Least Contr'. The 'Members/Group' column has a 'Select Members' button. The 'Virtual Servers' column has a dropdown menu set to 'vra_htt'. Below the table is a form for configuring the server pool. It includes a 'Description' field with the placeholder 'Enter Description', an 'Active Monitor' dropdown menu, and a 'SNAT Translation Mode' dropdown menu set to 'Automap'. There is also an 'Additional Properties' section with a right-pointing arrow. At the bottom of the form are 'SAVE' and 'CANCEL' buttons.

The second screenshot shows the 'Configure Server Pool Members' dialog box. It has a title bar with a close button. Below the title bar is a subtitle 'Server Pool - pool_iaas-manager_443'. There are two radio buttons: 'Enter individual members' (selected) and 'Select a group'. Below the radio buttons is an 'ADD MEMBER' button and a search bar with the placeholder 'Search'. Below the search bar is a table with columns: Name, IP, Port, Weight, State, Backup Member, and Max Concurrent Connections. The table has two rows of data. The first row has 'Name' as a text input, 'IP' as a text input, 'Port' as '443', 'Weight' as '1', 'State' as 'Enabled', 'Backup Member' as a radio button set to 'Disabled', and 'Max Concurrent Connections' as a text input. The second row has 'Name' as a text input, 'IP' as a text input, 'Port' as '443', 'Weight' as '1', 'State' as 'Enabled', 'Backup Member' as a radio button set to 'Disabled', and 'Max Concurrent Connections' as a text input. Below the table are 'CANCEL' and 'APPLY' buttons.

Configure NSX-T Virtual Servers

It is required to configure virtual servers for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Navigate to **Networking > Load Balancing > Virtual Servers**.
- 2 Click **Add virtual server** and select **Layer**.

3 Configure the virtual servers as outlined in this table:

Table 5-3. Configure Virtual Servers

Name	Type	Application Profile	IP Address	Port	Server Pool	Persistence Profile
vRealize Automation	L4 TCP	vRealize Automation	IP	443	vRealize Automation	None
vRealize Orchestrator	L4 TCP	vRealize Orchestrator	IP	443	vRealize Orchestrator	None

Note Use only for external vRealize Orchestrator instances.

Results

The completed configuration should look similar to this screen.

The screenshot displays the 'VIRTUAL SERVERS' configuration page in the vRealize Automation interface. A table at the top lists virtual servers, with the first entry 'vs_vra-va-web_443' selected. Below the table, the configuration details for this server are shown in a form. The form includes fields for Description, Persistence (set to Disabled), Additional Properties (Max Concurrent Connections: Unlimited, Max New Connection Rate: Unlimited, Default Pool Member Ports: 443), Admin State (Enabled), and Tags. The 'Application Profile' is set to 'vRA_HTTP'. The 'Load Balancer' is set to 'r34r3r4'. The 'Server Pool' is set to 'pool_...'. The 'Access Log' is disabled. The 'Tag (Required)' field is empty, and the 'Scope (Optional)' field is also empty. The 'Maximum 30 tags are allowed.' message is displayed below the tags section. The 'SAVE' and 'CANCEL' buttons are at the bottom.

Configure Load Balancer

Specify a load balancer for each vRealize Automation, and for an external vRealize Orchestrator (optional) instance.

Procedure

- 1 Navigate to **Networking > Load Balancing > Load Balancers**.
- 2 Click **Add Load Balancer**.
- 3 Enter a name and select the appropriate **Load Balancer Size** (depends on vRealize Automation cluster size).
- 4 Select the **Tier 1 Logical Router**.

Note In NSX-T version 2.4, the monitor health checks are performed using the IP address of Tiers-1 uplink (or first service port for Tiers-1 standalone SR) for all load balancer server pools. Ensure that server pools are accessible from this IP address.

Results

The configuration should look similar to this screen:

The screenshot displays the 'LOAD BALANCERS' configuration page in vRealize Automation. The top navigation bar includes 'LOAD BALANCERS', 'VIRTUAL SERVERS', 'SERVER POOLS', 'PROFILES', 'MONITORS', and 'About'. A blue button labeled 'ADD LOAD BALANCER' is visible. Below the navigation bar, a table lists the configured load balancers. The first entry is 'vra75_lb' with a size of 'Small' and a 'Tier-1 Gateway' of 'vRA-LB-Tier-1-Router'. Below the table, there are input fields for 'Description', 'Tags' (with a note 'Maximum 30 tags are allowed'), 'Error Log Level', and 'Admin State' (a toggle switch). A section for 'VIRTUAL SERVERS' is collapsed. At the bottom, there are 'SAVE' and 'CANCEL' buttons.

Add Virtual Servers to Load Balancer

Once you've configured the load balancer, you can add virtual servers.

Procedure

- 1 Navigate to **Networking > Load Balancing > Virtual Servers**.
- 2 Edit the configured virtual servers.
- 3 Assign the previously configured load balancer as the **Load Balancer**.

Results

The configuration should look similar to this screen:

Name	IP Address	Ports	Type	Load Balancer	Server
vs_vra-va-web_443 *	192.168.205.10 * e.g. 10.10.10.10	443 x Enter Ports or Port Rang	L4 TCP	vRA_LB (x) v	p
Description		Enter Description		Application Profile *	vRA_HTTPS
Persistence		Disabled v			
> Additional Properties					
<div>SAVE CANCEL</div>					

Configuring F5 Big-IP LTM

6

Before configuring your F5 device, it must be deployed in the environment with access to vRealize components over a network.

Note Refer to the [Workspace One](#) load-balancing documentation in order to configure highly-available identity provider for vRealize Automation.

For configuration, the F5 device must meet these requirements:

- The F5 device can be either physical or virtual.
- The F5 Local Traffic module (LTM) load balancer can be deployed in either one-arm or multi-arm topologies.
- The LTM must be configured and licensed as either Nominal, Minimum, or Dedicated. You can configure the LTM by navigating to **System > Resource Provisioning**.

If you are using an F5 LTM version older than 11.x, you might need to change your health monitor settings related to the Send string. For more information about how to set up your health monitor send string for the different versions of F5 LTM, see [HTTP health checks may fail even though the node is responding correctly](#).

This chapter includes the following topics:

- [Configure Monitors](#)
- [Configure F5 Server Pools](#)
- [Configure F5 Virtual Servers](#)

Configure Monitors

It is required to add monitors for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Log in to the F5 load balancer and navigate to **Local Traffic > Monitor**.

- 2 Click **Create** and configure the monitor as outlined in this table. Use the default value if nothing is specified.

Table 6-1. Configure Monitors

Name	Type	Interval	Timeout	Send String.	Receive String.	Alias Service Port
vRealize Automation	HTTP	3	10	GET /health HTTP/1.0\r\n\r\n	HTTP/1\.(0 1) (200)	8008
vRealize Orchestrator	HTTP	3	10	GET /health HTTP/1.0\r\n\r\n	HTTP/1\.(0 1) (200)	8008

Note Use only for external vRealize Orchestrator instances.

Results

The configuration should look similar to this screen.

The screenshot shows the 'New Monitor...' configuration window in the F5 Load Balancer. The breadcrumb trail at the top is 'Local Traffic > Monitors > New Monitor...'. The 'General Properties' section includes fields for Name (vra_http_va_web), Description, Type (HTTP), and Parent Monitor (http). The 'Configuration' section is set to 'Basic' and includes fields for Interval (3 seconds), Timeout (10 seconds), Send String (GET /health HTTP/1.0\r\n\r\n), Receive String (HTTP/1\.(0|1) (200)), Receive Disable String, User Name, Password, Reverse (No), Transparent (No), Alias Address (* All Addresses), Alias Service Port (8008), and Adaptive (Enabled). At the bottom are buttons for Cancel, Repeat, and Finished.

General Properties	
Name	vra_http_va_web
Description	
Type	HTTP
Parent Monitor	http

Configuration: Basic

Interval	3 seconds
Timeout	10 seconds
Send String	GET /health HTTP/1.0\r\n\r\n
Receive String	HTTP/1\.(0 1) (200)
Receive Disable String	
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	8008 Other:
Adaptive	<input type="checkbox"/> Enabled

Cancel Repeat Finished

Configure F5 Server Pools

It is required to configure service pools for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Log in to the F5 load balancer and navigate to **Local Traffic > Pools**.

- 2 Click **Create** and configure the pool as outlined in this table. Use the default value if nothing is specified.

Table 6-2. Configure Server Pools

Name	Health Monitors	Load Balancing Method	Node Name	Address	Service Port
vRealize Automation	vRealize Automation	Least Connections (member)	VA1 VA2 VA3	IP Address	443
vRealize Orchestrator	vRealize Orchestrator	Least Connections (member)	VA1 VA2 VA3	IP Address	443

Note Use only for external vRealize Orchestrator instances.

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

Results

The configuration should look similar to this screen.

Local Traffic » Pools : Pool List » **pl_vra-va-00_443**

⚙ Properties **Members** Statistics

Load Balancing

Load Balancing Method: Least Connections (member)

Priority Group Activation: Disabled

Update

Current Members

<input checked="" type="checkbox"/>	<input type="checkbox"/> Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio	Priority Group
<input type="checkbox"/>		dz-vra8-node1.sof-mbu.eng.vmware.com:443	192.168.10.30	443		No	1	0 (Active)
<input type="checkbox"/>		dz-vra8-node2.sof-mbu.eng.vmware.com:443	192.168.10.31	443		No	1	0 (Active)
<input type="checkbox"/>		dz-vra8-node3.sof-mbu.eng.vmware.com:443	192.168.10.32	443		No	1	0 (Active)

Enable Disable Force Offline Remove

Configure F5 Virtual Servers

It is required to configure virtual servers for vRealize Automation, and for an external vRealize Orchestrator (optional).

Procedure

- 1 Log in to the F5 load balancer and navigate to **Local Traffic > Virtual Servers**.
- 2 Click **Create** and configure the virtual server as outlined in this table. Use the default value if nothing is specified.

Table 6-3. Configure Virtual Servers

Name	Type	Destination Address	Service Port	Source Address Translation	Default Pool	Default Persistence Profile
vRealize Automation	Performance (Layer 4)	IP Address	443	Auto Map	vRealize Automation	None
vRealize Orchestrator	Performance (Layer 4)	IP Address	443	Auto Map	vRealize Orchestrator	None
Note Use only for external vRealize Orchestrator instances.						

- 3 For an overall view and the status of the virtual servers, select **Local Traffic > Virtual Servers**.

Results

The configuration should look similar to these screens.

General Properties

Name	vs_vra-va-00_443
Description	
Type	Performance (Layer 4)
Source Address	<input checked="" type="radio"/> Host <input type="radio"/> Address List
Destination Address/Mask	<input checked="" type="radio"/> Host <input type="radio"/> Address List 192.168.10.33
Service Port	<input checked="" type="radio"/> Port <input type="radio"/> Port List 443 HTTPS
Notify Status to Virtual Address	<input checked="" type="checkbox"/>
State	Enabled

Configuration: Basic

Protocol	TCP
Protocol Profile (Client)	fastL4
HTTP Profile (Client)	None
HTTP Profile (Server)	(Use Client Profile)
HTTP Proxy Connect Profile	None
VLAN and Tunnel Traffic	All VLANs and Tunnels
Source Address Translation	Auto Map

Acceleration: Basic

iSession Profile	None
Rate Class	None

Resources

iRules	Enabled	Available
		/Common _sys_APM_ExchangeSupport_OA_BasicAuth _sys_APM_ExchangeSupport_OA_NtimAuth _sys_APM_ExchangeSupport_helper _sys_APM_ExchangeSupport_main
Default Pool	+ pl_vra-va-00_443	
Default Persistence Profile	None	
Fallback Persistence Profile	None	

Cancel Repeat Finished

● vs_vra-va-00_443

STATS DIAGRAM

☐ List other virtual servers that share these pools ☐ List other pools that use these nodes

Virtual Server

Pools

Pool Members

● vs_vra-va-00_443
192.168.10.33:443

● pl_vra-va-00_443

● dz-vra8-node1.sof-mbu.er
192.168.10.30
 ● dz-vra8-node2.sof-mbu.er
192.168.10.31
 ● dz-vra8-node3.sof-mbu.er
192.168.10.32

Configuring Citrix ADC (NetScaler ADC)

7

Before you configure Citrix ADC, ensure the NetScaler device is deployed in the environment with access to the vRealize Components.

For configuration, the Citrix ADC must meet these requirements:

- You can use either a virtual or physical NetScaler.
- The Citrix load balancer can be deployed in either a one-arm or multi-arm topologies.
- Enable the load balancer and SSL modules by navigating to **NetScaler > System > Settings > Configure > Basic Features**.

This chapter includes the following topics:

- [Configure Citrix Monitors](#)
- [Configure Citrix Service Groups](#)
- [Configure Citrix Virtual Servers](#)

Configure Citrix Monitors

You can configure a Citrix monitor by performing these steps.

Procedure

- 1 Log in to the NetScaler Load Balancer and navigate to **NetScaler > Traffic Management > Load Balancing > Monitors**.

- 2 Click **Add** and configure the monitor as outlined in this table. Use the default value if nothing is specified.

Table 7-1. Configure Citrix Monitors

Name	Type	Interval	Timeout	Retries	Success Retries	HTTP Request/Send String	Response Codes	Receive String	Dest. Port	Secure
vRealize Automation	HTTP	5	4	3	1	GET / health	200	None	8008	No
vRealize Orchestrator	HTTP	5	4	3	1	GET / health	200	None	8008	No
Note Use only for external vRealize Orchestrator instances.										

Results

The configuration should look similar to this screen.

← Create Monitor

Name*

vra_https_va_web

i

Type*

HTTP

>

i

Basic Parameters

Interval

5

Second

▼

Response Time-out

4

Second

▼

i

Response Codes

+

200

×

Custom Header

HTTP Request

GET /health

i

☐ Secure

Advanced Parameters

Destination IP

Destination Port

8008

i

Down Time

30

Second

▼

TROFS Code

TROFS String

Dynamic Time-out

i

Deviation

Second

▼

Dynamic Interval

Retries

3

i

Configure Citrix Service Groups

You can configure service groups by performing these steps.

Procedure

- 1 Log in to the NetScaler load balancer and navigate to **NetScaler > Traffic Management > Load Balancing > Service Groups**.
- 2 Click **Add** and configure the service groups as outlined in this table.

Table 7-2. Configure Service Groups

Name	Health Monitors	Protocol	SG Members	Address	Port
vRealize Automation	vRealize Automation	SSL Bridge	VA1 VA2 VA3	IP Address	443
vRealize Orchestrator	vRealize Orchestrator	SSL Bridge	VA1 VA2 VA3	IP Address	443

Note Use only for external vRealize Orchestrator instances.

Results

The configuration should look similar to this screen:

← Load Balancing Service Group

Basic Settings

Name	pl_vra-va-00_443	Cache Type	SERVER
Protocol	SSL_BRIDGE	Cacheable	NO
State	ENABLED	Health Monitoring	YES
Effective State	● UP	AppFlow Logging	ENABLED
Traffic Domain	0	Monitoring Connection Close Bit	NONE
Comment		Number of Active Connections	0
		AutoScale Mode	DISABLED

Service Group Members

3 Service Group Members >

Settings

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

Monitors

1 Service Group to Monitor Binding >

Done

Configure Citrix Virtual Servers

You can configure virtual servers by performing these steps.

Procedure

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

- 2 Click **Add** and configure the virtual server as outlined in this table. Use the default value if nothing is specified.

Table 7-3. Configure Virtual Servers

Name	Protocol	Destination Address	Port	Load Balancing Method	Service Group Binding
vRealize Automation	SSL Bridge	IP Address	443	Least Connections	vRealize Automation
vRealize Orchestrator	SSL Bridge	IP Address	443	Least Connections	vRealize Orchestrator

Note Use only for external vRealize Orchestrator instances.

Results

The configuration should look similar to this screen:

Load Balancing Virtual Server

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

Basic Settings

Name	vs_vra-va-00_443	Listen Priority	-
Protocol	SSL_BRIDGE	Listen Policy Expression	NONE
State	● UP	Redirection Mode	IP
IP Address	10.71.226.23	Range	1
Port	443	IPset	-
Traffic Domain	0	RHI State	PASSIVE
		AppFlow Logging	ENABLED
		Retain Connections on Cluster	NO

Services and Service Groups

No Load Balancing Virtual Server Service Binding >

1 Load Balancing Virtual Server ServiceGroup Binding >

Traffic Settings

Health Threshold	0	Priority Queuing	
Client Idle Time-out	180	Sure Connect	
Minimum Autoscale Members	0	Down State Flush	ENABLED
Maximum Autoscale Members	0	Layer 2 Parameters	OFF
ICMP Virtual Server Response	PASSIVE	Trofs Persistence	ENABLED

Done

Configuring AVI Load Balancer

8

You can configure an AVI load balancer by performing these steps.

Ensure that you have deployed a Service Engine in the vCenter where the vRealize Automation instance is located and that the Service Engine interface is configured in the same network as the vRealize Automation.

This chapter includes the following topics:

- [Create Pool](#)
- [Create an Active Monitor](#)
- [Configure Virtual Service](#)

Create Pool

You can create pools for an AVI load balancer by performing the following steps.

To create a pool:

- 1 Navigate to the **Menu** and click **Applications**.
- 2 Click the **Pool** tab and enter these details.

Appliance Name	Default Server Port	Lookup Server by Name	Real time metrics	Enable SSL	SSL Profile
vRealize Automation	443	Enabled	Enabled	Enabled	System Standard
vRealize Orchestrator	443	Enabled	Enabled	Enabled	System Standard
Note Use only for external vRealize Orchestrator instances.					

- 3 Click **Next** and add servers to the pool.

Edit Pool: vRA Cluster 1-pool

Settings Servers Advanced

Name Enabled ☒ AutoScale Policy

Default Server Port AutoScale Launch Config

Graceful Disable Timeout Minutes Persistence

Load Balance Analytics Profile

Health Monitors ☐ Passive Health Monitor ☒ Lookup Server by Name ☒ Rewrite Host Header to Server Name ☐ Enable real time metrics ☒

Min. Health Monitors to consider server 'up'

[+ Add Active Monitor](#)

• SSL to Backend Servers •

☒ Enable SSL

Cancel Save

New Pool: vRA Cluster 1-pool

Step 1: Settings Step 2: Servers Step 3: Advanced Step 4: Review

• Add Servers •

Select Servers

Server IP Address [Add Server](#)

• Servers •

☐ Enable HTTP2

Displaying 2 items

<input type="checkbox"/>	Status	Server Name	Resolve by DNS	IP Address	Port	Ratio	Description	Network	Header ...	Rewrite ...
<input type="checkbox"/>	Enabled	<input type="text"/>	<input type="checkbox"/>	10.71.224.161	<input type="text" value="443"/>	1	<input type="text"/>		<input type="text" value="Header"/>	<input type="checkbox"/>
<input type="checkbox"/>	Enabled	<input type="text"/>	<input type="checkbox"/>	10.71.224.162	<input type="text" value="443"/>	1	<input type="text"/>		<input type="text" value="Header"/>	<input type="checkbox"/>

Cancel [Previous](#) [Next](#)

4 Follow the remaining steps in the wizard to finish creating the pool.

Create an Active Monitor

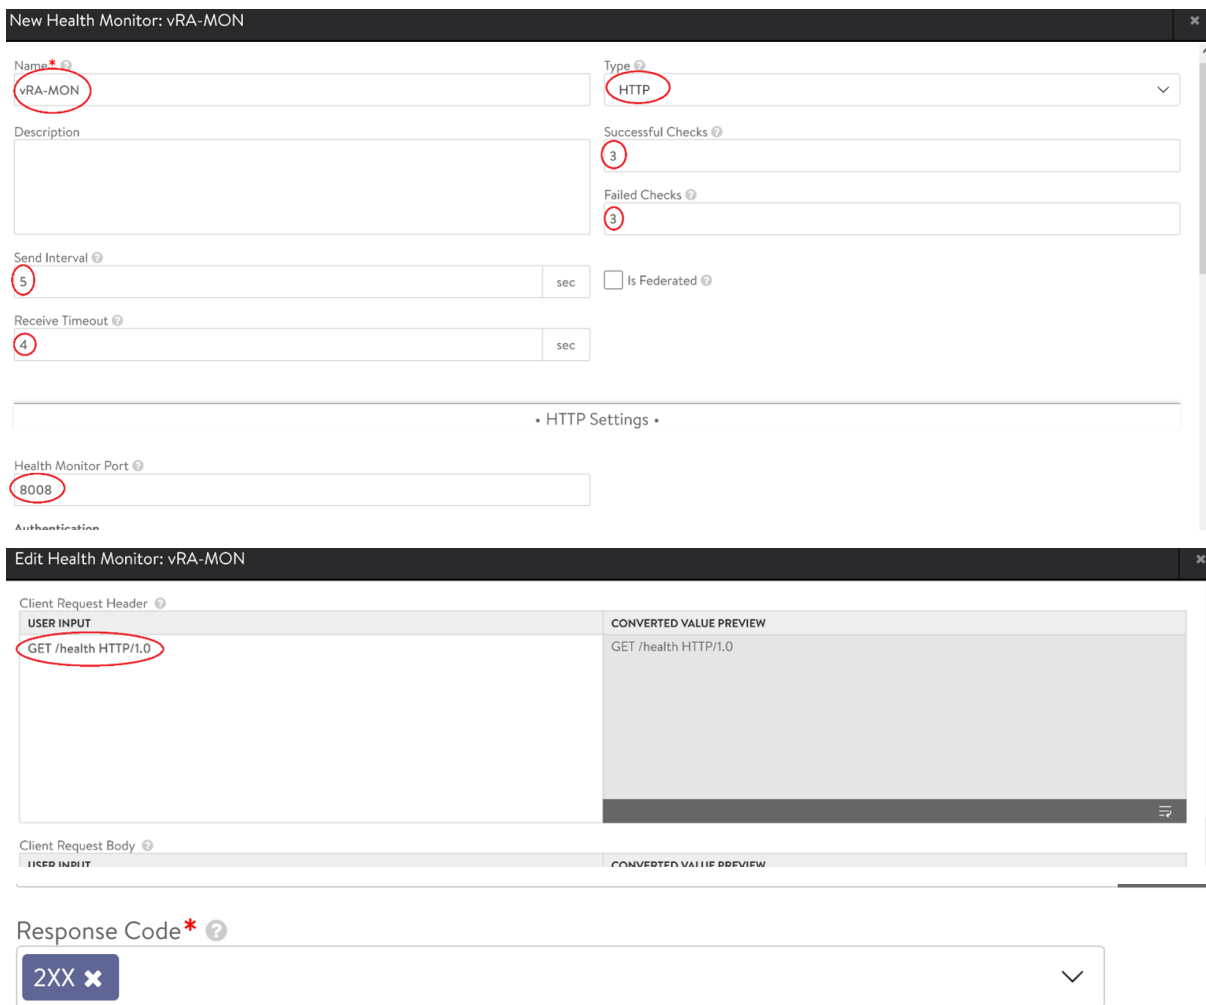
You can create an active monitor by following these steps.

To create an active monitor you must edit the pool configuration.

- 1 From the pool, click the **Edit** icon. to open the context window in the Settings tab.
- 2 Click **Add Active Monitor** and then click the down arrow.

3 Select **Create Health Monitor** and enter the following details.

Appliance Name	Type	Interval	Timeout	Successful checks	Failed checks	Health Monitor Port	Client request header	Response code
vRealize Automation	HTTP	5	4	3	3	8008	GET / health HTTP/1.0	2XX
vRealize Orchestrator	HTTP	5	4	3	3	8008	GET / health HTTP/1.0	2XX
Note Use for external vRealize Orchestrator instances only.								



New Health Monitor: vRA-MON

Name ^{*} [?] vRA-MON Type [?] HTTP

Description

Successful Checks [?] 3

Failed Checks [?] 3

Send Interval [?] 5 sec ☐ Is Federated [?]

Receive Timeout [?] 4 sec

• HTTP Settings •

Health Monitor Port [?] 8008

Authentication

Edit Health Monitor: vRA-MON

Client Request Header [?]

USER INPUT

GET /health HTTP/1.0

CONVERTED VALUE PREVIEW

GET /health HTTP/1.0

Client Request Body [?]

USER INPUT

CONVERTED VALUE PREVIEW

Response Code ^{*} [?]

2XX ✕

Configure Virtual Service

You can configure virtual service for an AVI load balancer by following these steps.

To configure virtual service:

- 1 From the menu, click **Applications**.
- 2 Click the **Virtual Services** tab, and then click **Create Virtual Service**.
- 3 Enter these configuration details:

Appliance Name	FQDN or IP Address	TCP/UDP Profile	Application Profile	Services	Pool
vRealize Automation	VIP Address or FQDN	System-TCP-Proxy	System-L4-Application	443	vRealize Automation
vRealize Orchestrator	VIP Address or FQDN	System-TCP-Proxy	System-L4-Application	443	vRealize Orchestrator
Note Use for external vRealize Orchestrator instances only.					

The screenshot shows the 'New Virtual Service: vRA Cluster 1' configuration window. The 'Step 1: Settings' tab is active. The 'Name' field is 'vRA Cluster 1'. The 'VIP Address' section shows 'FQDN or IPv4 Address' as '10.71.224.165' and 'Service Port' as '443'. The 'Profiles' section shows 'TCP/UDP Profile' as 'System-TCP-Proxy' and 'Application Profile' as 'System-L4-Application'. The 'Pool' section shows 'Pool' selected and 'vRA Cluster 1-pool' as the pool name. The 'Next' button is highlighted in green.

- 4 Click **Next** to navigate to the **Advanced** tab and enter the following information.

Appliance Name	Placement Network	IPv4 Subnet	Server Network Profile	SE Group	Use VIP, as SNAT
vRealize Automation	Network where VIP is	Network and netmask	System-TCP-Proxy	SE Group where the appropriate SE is located	Enabled
vRealize Orchestrator	Network where VIP is	Network and netmask	System-TCP-Proxy	SE Group where the appropriate SE is located	Enabled
Note Use for external vRealize Orchestrator instances only.					

Edit Virtual Service: vRA Cluster 1 Help ×

Settings Policies Analytics Advanced

• Quality of Service •

Weight ⓘ
1

Fairness ⓘ
Throughput And Delay Fairness Throughput Fairness

• Virtual IP Placement Settings •

Virtual IP
10.71.224.165

Placement Network ⓘ
10.71.224 (vlan1224) (Static) - 10.71.224.0/24, 10.71.0.0/16

IPv4 Subnet ⓘ
10.71.224.0/24

IPv6 Subnet ⓘ
2001::1/24

+ Add Placement Network

• Other Settings •

Server Network Profile ⓘ
System-TCP-Proxy

SE Group ⓘ
Default-Group

☒ Auto Gateway ⓘ ☒ Use VIP as SNAT ⓘ

☐ Advertise VIP via BGP ⓘ ☐ Advertise SNAT via BGP ⓘ

Cancel Save

Troubleshooting

9

This chapter includes the following topics:

- [Errors during vRealize Automation installation when using NSX-V as a load-balancer for Workspace ONE](#)
- [Provisioning Failures When Using OneConnect with F5 BIG-IP](#)
- [F5 BIG-IP License Limits Network Bandwidth](#)

Errors during vRealize Automation installation when using NSX-V as a load-balancer for Workspace ONE

If you see errors when installing vRealize Automation while using Workspace ONE as load-balancer, follow these troubleshooting steps.

When using NSX-V as a load-balancer for VMware Workspace ONE there might be specific network limitations which will result in errors and timeouts during the installation of vRealize Automation similar to:

```
2020-06-30 09:10:08.751+0000 INFO 16 --- [or-http-epoll-3]
com.vmware.identity.rest.RestClient : POST https://default-49-29.sqa.local/SAAS/API/1.0/oauth2/token?
grant_type=client_credentials
2020-06-30 09:10:08.755+0000 WARN 16 --- [or-http-epoll-3]
r.netty.http.client.HttpClientConnect : [id: 0x754860c7, L:/10.244.0.206:48686 !
R:default-49-29.sqa.local/10.198.49.29:443] The connection observed an error
reactor.netty.http.client.PrematureCloseException: Connection prematurely closed BEFORE response
```

You can mitigate those errors by extending the NSX-V idle connection close time to 5 minutes instead of the default of 1 second.

This can be achieved with an application rule containing the following:

```
timeout http-keep-alive 300s
```

Provisioning Failures When Using OneConnect with F5 BIG-IP

When you use the OneConnect feature with F5 BIG-IP for a virtual server, provisioning tasks sometimes fail.

OneConnect ensures connections from the load balancer to the back-end servers are multiplexed and reused. This lowers the load on the servers and makes them more resilient.

Using OneConnect with a virtual server that has SSL pass-through is not recommended by F5 and might result in failed provisioning attempts. This happens because the load balancer attempts to establish a new SSL session over an existing session while the back-end servers expect the client to either close or renegotiate the existing session, which results in a dropped connection. Disable OneConnect to resolve this issue.

- 1 Log in to the F5 load balancer and navigate to **Local Traffic > Virtual Servers > Virtual Servers List**.
- 2 Click the name of the virtual server you want to modify.
- 3 In the **Acceleration** section, select **None** for the **OneConnect Profile**.
- 4 Click **Finish**.

F5 BIG-IP License Limits Network Bandwidth

You might experience provisioning failures or problems loading vRealize Automation console pages due to load balancer network traffic exceeding the F5 BIG-IP license limit.

To check if the BIG-IP platform is experiencing this problem, see [How the BIG-IP VE system enforces the licensed throughput rate](#).