



vRealize Automation Load Balancing

Configuration Guide
Version 7.5

TECHNICAL WHITE PAPER

APRIL 2019

VERSION 2.9.2

Table of Contents

Introduction	5
Load Balancing Concepts	5
SSL Pass-Through.....	5
Session Persistence.....	5
Destination Address (F5 and NetScaler)	6
Source (IP) Address (F5, NetScaler, and NSX-V)	6
Source IP Address Hash (NSX-V)	6
Email Notifications on Load Balancer	6
One-Arm or Multi-Arm Topologies.....	6
Prerequisites for Configuring Load Balancers with vRealize Automation.....	7
Complete the vRealize Automation Initial Installation	8
Configure vRealize Automation Identity Connectors for Kerberos (Optional).....	10
Configuring F5 Big-IP LTM.....	11
Configure Custom Persistence Profile.....	11
Configure Monitors	11
Configure Server Pools.....	13
Configure Virtual Servers	17
Configuring NSX-V	20
Configure Global Settings	20
Add Application Profiles	22
Add Service Monitoring	23
Add Pools	24
Add Virtual Servers	25
Configuring NSX-T	27
Add Application Profiles	27
Add Persistence Profile	28
Add Active Health Monitor	28
Configure Server Pools.....	32
Configure Virtual Servers	33
Configure Load Balancer.....	35
Add Virtual Servers to Load Balancer.....	35
Configuring Citrix ADC (NetScaler ADC).....	36
Configure Monitors	36
Configure Service Groups	38
Configure Virtual Servers	40

Configure Persistency Group.....	42
Troubleshooting	43
Provisioning failures when using OneConnect with F5 BIG-IP for a virtual server with SSL pass-through	43
F5 BIG-IP license limits network bandwidth	43
Proxy agent ping failure	43
Connection reset errors in the catalina.out log file	44
Proxy Agents cannot connect to load balanced Manger Service endpoint	44

Revision History

DATE	VERSION	DESCRIPTION
August 2015	1.0	Initial version
December 2015	1.1	Minor updates
December 2015	2.0	Updates for vRealize Automation 7.0
January 2016	2.1	Minor updates
May 2016	2.2	Updates for vRealize Automation 7.0.x
June 2016	2.3	<ul style="list-style-type: none"> • Updated timeout to 10 seconds for Configure Monitors and Add Service Monitoring in F5 and NSX-V sections respectively • Added source IP persistence and timeout of 1800 seconds for Add Application Profiles section • Updated all the screenshots to match the content • Updated NSX-V load balancing method to be round robin
August 2016	2.4	<ul style="list-style-type: none"> • Added configuration for Citrix NetScaler • Updated for NSX-V 6.2
November 2016	2.5	<ul style="list-style-type: none"> • Updated interval to 5 seconds for Configure Monitors in Citrix NetScaler section • Updated timeout to 4 seconds for Configure Monitors in Citrix NetScaler section
May 2017	2.6	Minor updates
May 2017	2.7	<ul style="list-style-type: none"> • Added monitor and pool configurations for vRealize Orchestrator Control Center. • Added troubleshooting section.
May 2018	2.8	<ul style="list-style-type: none"> • Updated version to 7.4 • Added information about expected result for load balancer installation • Added troubleshooting topic for increasing connection time • Revised SSL pass-through information. • Revised Configure Persistence Group section
November 2018	2.9	<ul style="list-style-type: none"> • Updated version to 7.5 • Updated load balancer software versions from the latest testing • Added support for three virtual appliances • Added support for Manager Service automatic failover • Added pools, virtual IPs for port 80 • General refresh of the document (grammar, style, and formatting)
December 2018	2.9.1	<ul style="list-style-type: none"> • Corrected the health monitor receive string to include a backslash.

April 2019	2.9.2	<ul style="list-style-type: none"> • Add supporting information for NSX -T • Clarified mentions of NSX to reference NSX-V • Added NSX-T 2.4 to the support matrix
------------	-------	--

Introduction

This document describes the configuration of the load balancing modules of F5 Networks BIG-IP software (F5), Citrix NetScaler, and NSX load balancers for vRealize Automation 7.x in a distributed and highly available deployment. This document is not an installation guide, but a load-balancing configuration guide that supplements the vRealize Automation installation and configuration documentation available at [VMware vRealize Automation product documentation](#).

This information is for the following products and versions.

PRODUCT	VERSION
F5 BIG-IP	11.x, 12.x, 13.x, 14.x
NSX-V	6.2.x, 6.3.x, 6.4.x (please refer to the VMware Product Interoperability Matrices for more details)
NSX-T	2.4
Citrix NetScaler	10.5, 11.x, 12.x
vRealize Automation	7.x

Load Balancing Concepts

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

SSL Pass-Through

SSL pass-through is used with the load balancing configurations for the following reasons:

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

Session Persistence

The persistence option overrides any load balancing algorithm option, for example setting dest_addr overrides, setting round robin, and so on. Different components in the vRealize Automation architecture benefit from different persistence methods. The configuration described in this document is the result of extensive testing and represents the best balance between stability, performance, and scalability. SSL pass-through is a required implementation for vRealize Automation load balancing.

Destination Address (F5 and NetScaler)

Destination address affinity persistence, also known as sticky persistence, supports TCP and UDP protocols, and directs session requests to the same server based on the destination IP address of a packet.

Source (IP) Address (F5, NetScaler, and NSX-V)

The default source IP address persistence option persists traffic based on the source IP address of the client for the life of that session and until the persistence entry timeout expires. The default for this persistence is 180 seconds (30 minutes). The next time a persistent session from that same client is initiated, it might be persisted to a different member of the pool. This decision is made by the load balancing algorithm and is non-deterministic.

NOTE: Set the persistence entry timeout to 1800 seconds to match the vRealize Automation GUI timeout.

Source IP Address Hash (NSX-V)

The source IP address is hashed and divided by the total weight of the running servers to designate which server receives the request. This process ensures that the same client IP address always reaches the same server if no server fails or starts. For more information on IP Hash load balancing, see VMware knowledge base article [KB 2006129](#).

Email Notifications on Load Balancer

It is a good practice to set up an email notification on the load balancer that sends emails to the system administrator every time a vRealize Automation or vRealize Orchestrator node goes down. Currently, NSX-V does not support email notification for such a scenario.

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 an email notification with F5 by following methods:

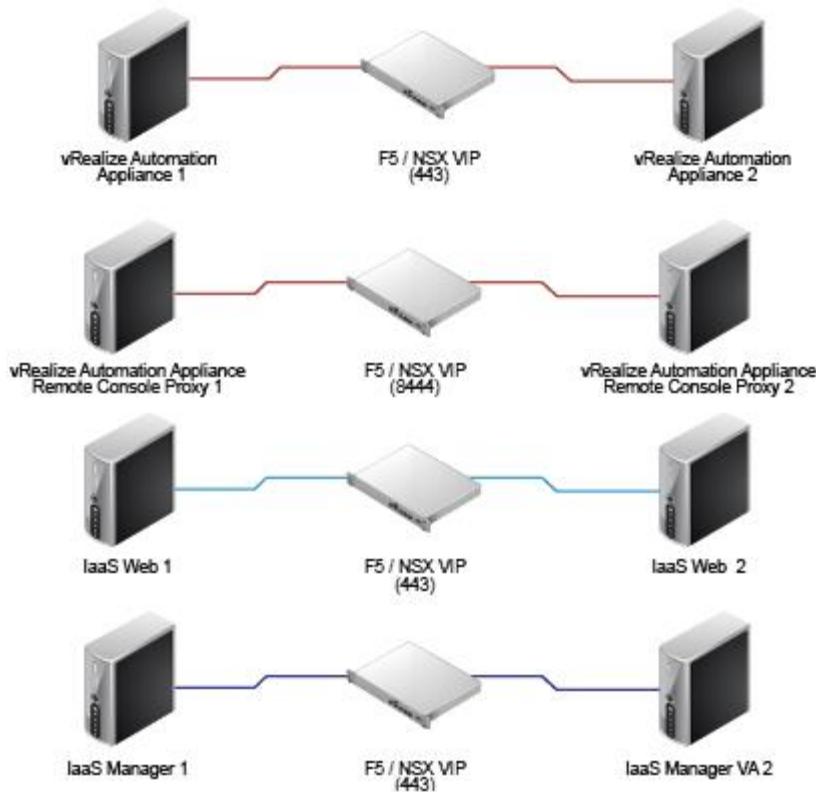
- [Configuring the BIG-IP system to deliver locally generated email messages](#)
- [Configuring custom SNMP traps](#)
- [Configuring alerts to send email notifications](#)

One-Arm or Multi-Arm Topologies

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 components are deployed as a one-arm configuration as shown in [Figure 1](#).

**FIGURE 1. ONE-ARM CONFIGURATION**

Prerequisites for Configuring Load Balancers with vRealize Automation

- **F5** – Before you start an HA implementation of vRealize Automation using an F5 load balancer, ensure that F5 is installed and licensed and that the DNS server configuration is complete.
- **NetScaler** – Before you start an HA implementation of vRealize Automation by using the NetScaler load balancer, ensure that NetScaler is installed and has installed at least a Standard Edition license.
- **NSX-V** – Before you start an HA implementation of vRealize Automation using NSX-V as a load balancer, ensure that your NSX-V topology is configured and that your version of NSX-V is supported. This document covers the load balancing aspect of an NSX-V configuration and assumes that NSX-V is configured and validated to work properly on the target environment and networks.
To verify that your version is supported, see the [vRealize Automation Support Matrix](#) for the current release.
- **Certificates** – Request Certificate Authority (CA) signed or create self-signed certificates containing the vRealize Automation virtual IP and the host names of the vRealize Automation nodes in the SubjectAltNames section. This configuration enables the load balancer to serve traffic without SSL errors. If you need to replace the self-signed certificates with CA signed certificates, see the VMware knowledge base article [KB 2107816](#). For more information about certificate troubleshooting and supportability, see the VMware knowledge base article [KB 2106583](#).
- **Identity provider** – Starting with vRealize Automation 7.0, the preferred Identity Provider is [VMware Identity Manager](#), which is embedded in the vRealize Automation Appliance.
- **Database** – Verify that supported database servers are available for vRealize Infrastructure-as-a-Service (IaaS) nodes. IaaS components require a Microsoft SQL Server instance.

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

If required, 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.

Complete the vRealize Automation Initial Installation

During the installation process of vRealize Automation, a load balancer 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, you must perform the following tasks.

1. Configure the F5, NSX, or NetScaler load balancer. See [Configuring F5 BIG-IP](#), [Configuring NSX](#), and [Configuring Citrix NetScaler](#).
2. Turn off the health monitors or change them temporarily to default TCP, and ensure traffic is still forwarding to your primary nodes.
3. Disable all secondary nodes (VA and IaaS) from the load balancer pools.
4. Install and configure all the system components as detailed in vRealize Automation 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](#).

After you have configured a directory for at least one tenant, ensure that the **IdP Hostname** is set to the load balancer virtual IP for the vRealize Automation virtual appliances and all available **connectors** are enabled and configured for authentication for each virtual appliance node.

The screenshot shows the vRealize Automation interface with the 'Administration' tab selected. On the left, a sidebar lists 'Identity Providers' under 'Administration'. A single entry, 'WorkspaceIDP__1', is shown with its details: Type: AUTOMATIC, Status: Enabled, and a 'Disable IdP' button. To the right, the main panel shows the configuration for this IdP. It includes fields for 'Identity Provider Name' (set to 'WorkspaceIDP__1'), 'Users' (with a checkbox for 'sqa'), 'Network' (set to 'ALL RANGES'), 'Authentication Methods' (set to 'Password'), 'Connector(s)' (listing 'vra-va-node1-host' and 'vra-va-node2-host' with checkboxes checked), and 'IdP Hostname' (set to 'vra-va-load-balancer-vip'). A note at the bottom right says: 'You can deploy external connectors and add them to this IdP for high availability. Create that connector for this IdP.' At the bottom are 'Save' and 'Cancel' buttons.

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

All Web nodes should be available, running, and enabled.
 - c. The pool, service groups, and virtual server of the Manager Service nodes should be available and running.
 - The active Manager Service node should appear as online, and enabled in the load balancer
 - Each passive Manager Service node should appear as offline, but enabled in the load balancer

Note: Starting with vRealize Automation 7.3.x, the Manager Service uses active-passive configuration and supports automatic failover. It is normal to have the passive nodes detected as offline, however; they should be enabled so the load balancer can start forwarding traffic to them automatically once a passive node is promoted to active one. For more information, see [About Automatic Manager Service Failover](#) in *Installing vRealize Automation*.

Configure vRealize Automation Identity Connectors for Kerberos (Optional)

If you require your vRealize Automation users to be authenticated by using Kerberos and they log in to the administration console by using Windows Single Sign-on, then you must configure the identity connectors.

1. Connect vRealize Automation to an Active Directory by using the IWA method.
2. Configure policies and methods for Kerberos authentication.
3. Open the administration console and navigate to **Administration > Directories Management > Connectors**. See [Using Directories Management to Create an Active Directory](#) in *Configuring vRealize Automation*.
4. In the Worker column, select a worker to view the connector details and navigate to the Auth Adapters page.
5. Click **KerberosIdpAdapter**.
You will be redirected to the vIDM Console.
6. Select **Enable Windows Authentication**.
7. Select **Enable Redirect**.
8. In Redirect Host Name, enter the **FQDN** of the appliance you are configuring.
9. Configure the KerberosIdPAdapter on all the connectors in your cluster. Ensure the configuration of the adapter is identical on all the connectors, except for the Redirect Host Name value, which should be specific to each connector.

The Enable Redirect option directs HTTPS access for all client computers to a vRealize Automation node on port 443. The user's browser is redirected away from the load balancing vRealize Automation endpoint to the appliance's FQDN to perform authentication. After authentication, the user's browser is redirected to the load balancing vRealize Automation FQDN.

Configuring the same HOST SPN (the load balancing FQDN) for each computer object in AD is not supported because Kerberos authentication relies on SPNs configured in Active Directory, which should be unique for every host. The only way to enable single sign-on for all members of a vRealize Automation cluster is to tell vIDM to redirect from the load balancing FQDN to the individual node, verify the ticket, and then redirect to the FQDN.

If you do not want to expose individual vRealize Automation appliances directly to the user and you still require Kerberos authentication, then your best option is to federate the vRealize Automation solution with a third-party SAML identity provider such as Active Directory Federation Services.

Configuring F5 Big-IP LTM

This document assumes that the F5 device is already deployed in the environment and can access vRealize Automation components over a network.

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

If you are using an F5 version before 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, see [HTTP health checks may fail even though the node is responding correctly](#).

Configure Custom Persistence Profile

You can configure the persistence profile for your F5 load balancer.

1. Log in to the F5 and select **Local Traffic > Profiles > Persistence**.
2. Click **Create**.
3. Enter the name **source_addr_vra** 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**.

Configure Monitors

You need to add the following monitors for vRealize Automation.

1. Log in to the F5 load balancer and select **Local Traffic > Monitors**.
2. Click **Create** and provide the required information.
Leave the default value when nothing is specified.
3. Repeat steps 1 and 2 for each row of information in [Table 1](#).
4. To check the network map for an overall view of the monitors, select **LTM > Network Map**.

TABLE 1 - CONFIGURE MONITORS

NAME	TYPE	INTERVAL	TIME OUT	SEND STRING	RECEIVE STRING	ALIAS SERVICE PORT
vra_https_va_web	HTTPS	3	10	GET/vcac/services/api/health\r\n	HTTP/1.(0 1) (200 204)	443
vra_https_iaas_web	HTTPS	3	10	GET /wapi/api/status/web\r\n	REGISTERED	
vra_https_iaas_mngr	HTTPS	3	10	GET /VMPSProvision\r\n	ProvisionService	
vro_https_8283	HTTPS	3	10	GET /veo-controlcenter/docs/	HTTP/1.(0 1) (200)	8283

Example

The configuration for a VA monitor should look similar to the following screen:

Local Traffic > Monitors > New Monitor...

General Properties

Name	vra_https_va_web
Description	
Type	HTTPS
Parent Monitor	https

Configuration: Basic

Interval	3 seconds
Timeout	10 seconds
Send String	GET /vcac/services/api/health\r\n
Receive String	HTTP/1\.(0 1) (200 204)
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	443 HTTPS
Adaptive	<input type="checkbox"/> Enabled

Cancel Repeat Finished

Configure Server Pools

You must configure the following server pools for vRealize Automation.

1. Log in to the F5 load balancer and select **Local Traffic > Pools**.
2. Click **Create** and provide the required information. Leave the default value 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 in

5. Table 2.
6. **Note:** The environment depicted in

Table 2 is an example. Your environment might contain two or three vRealize Automation virtual appliance nodes and two or more nodes per IaaS role.

7. Select **pl_iaas-man-00_443** on Local Traffic -> Pools. Click **Advanced** in the configuration section and Select **Drop for Action on Service Down**. Click **OK** and click **Finished**.
8. To check the network map for an overall view of the server pools, select **LTM > Network Map**.

Table 2 – Configure Server Pools

NAME	HEALTH MONITORS	LOAD BALANCING METHOD	NODE NAME	ADDRESS	SERVICE PORT
pl_vra_va-00_80	http	Least connections (member)	ra-vra-va-01	IP Address	80
			ra-vra-va-02	IP Address	80
			ra-vra-va-03	IP Address	80
pl_vra-va-00_443	vra_https_va_web	Least connections (member)	ra-vra-va-01	IP Address	443
			ra-vra-va-02	IP Address	443
			ra-vra-va-03	IP Address	443
*pl_vra-va-00_8444	vra_https_va_web	Least connections (member)	ra-vra-va-01	IP Address	8444
			ra-vra-va-02	IP Address	8444
			ra-vra-va-03	IP Address	8444
pl_vro-cc-00_8283	vro_https_8283	Least connections (member)	ra-vra-va-01	IP Address	8283
			ra-vra-va-02	IP Address	8283
			ra-vra-va-03	IP Address	8283
pl_iaas-web-00_443	vra_https_iaas_web	Least connections (member)	ra-web-01	IP Address	443
			ra-web-02	IP Address	443
pl_iaas-man-00_443	vra_https_iaas_mngr	**Least connections (member)	ra-man-01	IP Address	443
			ra-man-02	IP Address	443

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

** The Manager Service uses [active-passive type of configuration](#) hence the load balancer will always send the traffic to the current active node regardless of the load-balancing method.

Example

Your pool configuration should look similar to the following screen.

The screenshot shows the F5 Local Traffic > Pools : Pool List interface. The search bar contains 'pl'. The table lists six pools:

	Name	Members	Partition / Path
<input type="checkbox"/>	pl_iaas-man-00_443	2	Common
<input type="checkbox"/>	pl_iaas-web-00_443	2	Common
<input type="checkbox"/>	pl_vra-va-00_443	3	Common
<input type="checkbox"/>	pl_vra-va-00_8444	3	Common
<input type="checkbox"/>	pl_vra_va-00_80	3	Common
<input type="checkbox"/>	pl_vro-cc-00_8283	3	Common

Configure Virtual Servers

You must configure the following virtual servers for vRealize Automation.

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 value when nothing is specified.
3. Repeat steps 1 and 2 for each entry in Table 3.
4. For an overall view and status of the virtual servers, select **LTM > Network Map**.

TABLE 3 – CONFIGURE VIRTUAL SERVERS

NAME	TYPE	DESTINATION ADDRESS	SERVICE PORT	SOURCE ADDRESS TRANSLATION	DEFAULT POOL	DEFAULT PERSISTENCE PROFILE
vs_vra-va-00_80	Performance (HTTP)	IP Address	80	Auto Map	pl_vra-va-00_80	None
vs_vra-va-00_443	Performance (Layer 4)	IP Address	443	Auto Map	pl_vra-va-00_443	source_addr_vra
vs_vra-va-00_8444	Performance (Layer 4)	IP Address	8444	Auto Map	pl_vra-va-00_8444	source_addr_vra
vs_vro-00_8283	Performance (Layer 4)	IP Address	8283	Auto Map	pl_vro-cc-00_8283	source_addr_vra
vs_web-00_443	Performance (Layer 4)	IP Address	443	Auto Map	pl_iaas-web-00_443	source_addr_vra
vs_man-00_443	Performance (Layer 4)	IP Address	443	Auto Map	pl_iaas-man-00_443	None

Example

Local Traffic > Virtual Servers : Virtual Server List > New Virtual Server...

General Properties	
Name	vs_vra-va-00_443
Description	
Type	Performance (Layer 4)
Source Address	
Destination Address/Mask	10.23.89.44
Service Port	443 HTTPS
Notify Status to Virtual Address	<input checked="" type="checkbox"/>
State	Enabled

Configuration: Basic

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

Acceleration

Rate Class	None
------------	------

Resources

iRules	<table border="1"> <thead> <tr> <th>Enabled</th> <th>Available</th> </tr> </thead> <tbody> <tr> <td></td> <td> /Common Rule_Demo_test_ADO_Model _sys_APM_ExchangeSupport_OA_BasicAuth _sys_APM_ExchangeSupport_OA_NtlmAuth _sys_APM_ExchangeSupport_helper </td> </tr> </tbody> </table>	Enabled	Available		/Common Rule_Demo_test_ADO_Model _sys_APM_ExchangeSupport_OA_BasicAuth _sys_APM_ExchangeSupport_OA_NtlmAuth _sys_APM_ExchangeSupport_helper
Enabled	Available				
	/Common Rule_Demo_test_ADO_Model _sys_APM_ExchangeSupport_OA_BasicAuth _sys_APM_ExchangeSupport_OA_NtlmAuth _sys_APM_ExchangeSupport_helper				
Default Pool	+ pl_vra-va-00_443				
Default Persistence Profile	source_addr-30min				
Fallback Persistence Profile	None				

The completed configuration should look similar to the following screen:

The screenshot shows the 'Local Traffic > Virtual Servers : Virtual Server List' interface. The top navigation bar includes tabs for 'Virtual Server List' (which is selected and highlighted in yellow), 'Virtual Address List', and 'Statistics'. Below the navigation is a search bar with fields for 'vs_...' and buttons for 'Search' and 'Reset Search'. A 'Create...' button is located in the top right corner of the search area. The main content is a table titled 'Virtual Server List' with the following columns: Status, Name, Description, Application, Destination, Service Port, Type, Resources, and Partition / Path. The table contains six rows, each representing a virtual server with a green status icon and a checkbox column:

Status	Name	Description	Application	Destination	Service Port	Type	Resources	Partition / Path
<input type="checkbox"/>	vs_web-00_443			10.23.89.45	443 (HTTPS)	Performance (Layer 4)	Edit...	Common
<input type="checkbox"/>	vs_vro-00_8283			10.23.89.44	8283	Performance (Layer 4)	Edit...	Common
<input type="checkbox"/>	vs_vra-va-00_8444			10.23.89.44	8444	Performance (Layer 4)	Edit...	Common
<input type="checkbox"/>	vs_vra-va-00_80			10.23.89.44	80 (HTTP)	Performance (HTTP)	Edit...	Common
<input type="checkbox"/>	vs_vra-va-00_443			10.23.89.44	443 (HTTPS)	Performance (Layer 4)	Edit...	Common
<input type="checkbox"/>	vs_man-00_443			10.23.89.46	443 (HTTPS)	Performance (Layer 4)	Edit...	Common

At the bottom left of the table are buttons for 'Enable', 'Disable', and 'Delete...'. The entire interface is contained within a dark blue-bordered window.

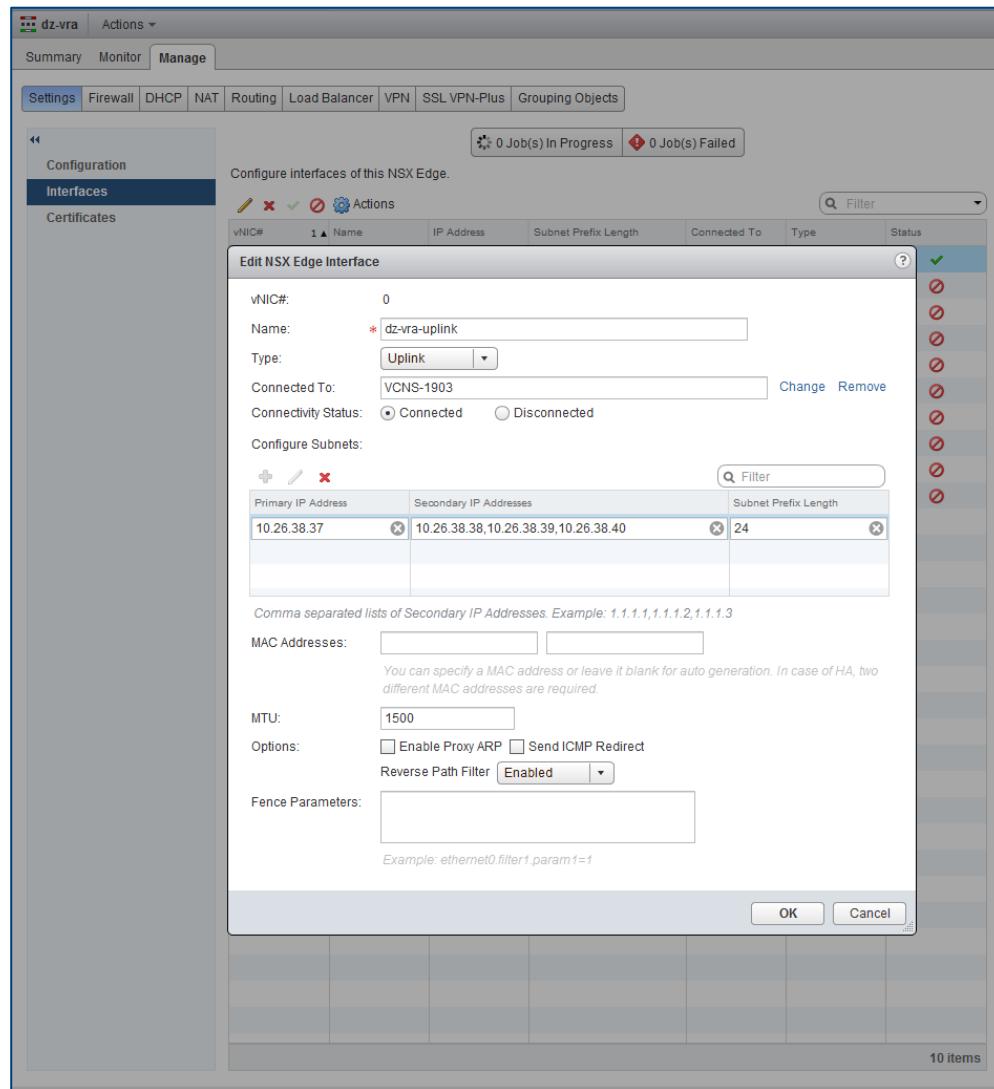
Configuring NSX-V

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 Automation components being load balanced.

Configure Global Settings

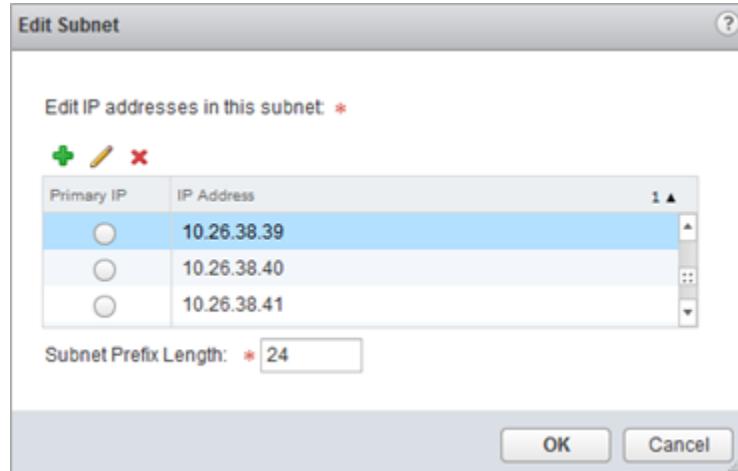
You can configure the global settings by using the following steps.

1. Log in to the NSX-V, click the **Manage** tab, click **Settings**, and select **Interfaces**.
2. Double-click on your Edge device in 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.

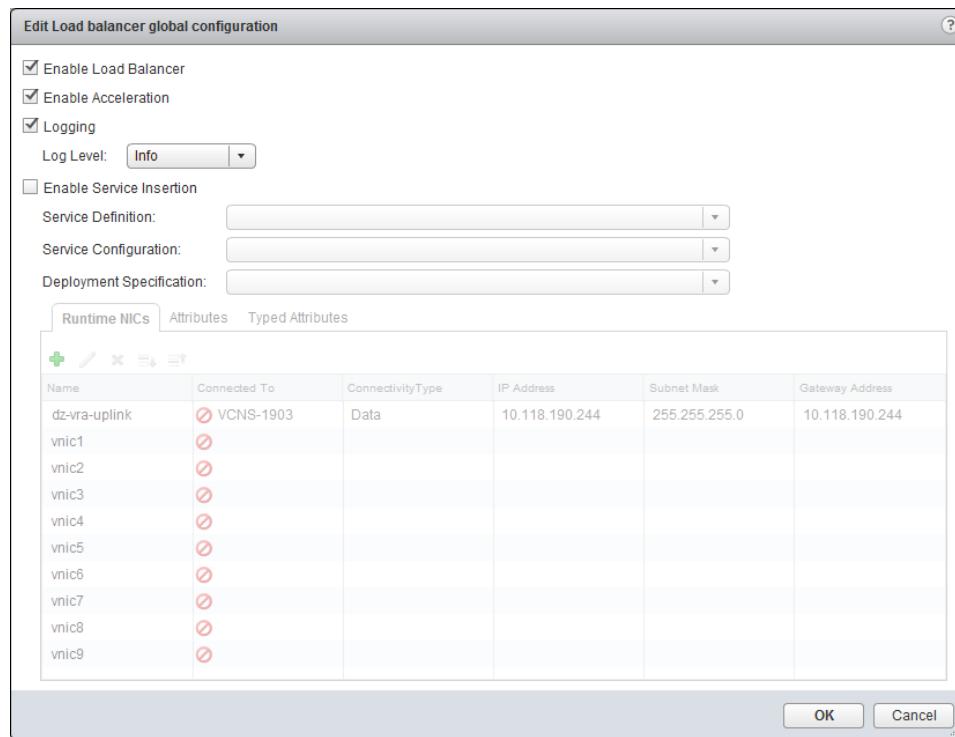


* This interface might look slightly different in NSX-V 6.1.x and earlier.

5. Add the IP addresses assigned to the virtual IPs and click **OK**.
6. Click **OK** to exit the interface configuration page.



7. Go to the **Load Balancer** tab and click the **Edit** icon.
8. Select **Enable Load Balancer**, **Enable Acceleration**, and **Logging**, if required, and click **OK**.



* This interface might look slightly different in NSX-V 6.1.x and earlier.

Add Application Profiles

You must add application profiles for the different components of vRealize Automation.

1. Click **Application Profiles** in the pane on the left.
2. Click the **Add** icon to create the application profiles required for vRealize Automation by using the information in Table 4. Leave the default value when nothing is specified.

TABLE 4 – ADD APPLICATION PROFILES

NAME	TYPE	ENABLE SSL PASS-THROUGH	TIMEOUT	PERSISTENCE
IaaS Manager	HTTPS	Checked	-	None
IaaS Web	HTTPS	Checked	1800 seconds	Source IP
vRealize Automation VA Web	HTTPS	Checked	1800 seconds	Source IP
vRealize Orchestrator Control Center	HTTPS	Checked		Source IP

Example

The completed configuration should look similar to the following screen:

The screenshot shows the 'Manage' tab selected in the top navigation bar. Below it, the 'Load Balancer' tab is also selected in a secondary navigation bar. The main content area displays a table titled 'Application Profiles'. The table has columns for 'Profile ID', 'Name', 'Persistence', and 'Type'. There are four entries in the table:

Profile ID	Name	Persistence	Type
applicationPrc	IaaS Manager		HTTPS
applicationPrc	IaaS Web	sourceip	HTTPS
applicationPrc	vRealize Orchestrator Control Center	sourceip	HTTPS
applicationPrc	vRealize Automation VA Web	sourceip	HTTPS

Below the table, a section titled 'Profile vRealize Automation VA Web Details:' contains fields for 'Cipher:' and 'Client Authentication:'. At the bottom, there are tabs for 'Virtual Server Certificates' and 'Pool Certificates', with 'Service Certificates' currently selected. A small table for certificate details is partially visible.

Add Service Monitoring

You must add service monitors for the different components of vRealize Automation.

1. Click **Service Monitoring** in the left pane.
2. Click the **Add** icon to create the service monitors required for vRealize Automation using information in Table 5. Leave the default value when nothing is specified.

TABLE 5 – ADD SERVICE MONITORING

NAME	INTERVAL	TIMEOUT	RETRIES	TYPE	METHOD	URL	RECEIVE	EXPECTED
vRealize Automation VA Web	3	10	3	HTTPS	GET	/vcac/services/api/health		200, 204 (for 7.0) 204 (for 7.0.1 and later)
IaaS Web	3	10	3	HTTPS	GET	/wapi/api/status/web	REGISTERED	
IaaS Manager	3	10	3	HTTPS	GET	/VMPSP provision	ProvisionService	
vRealize Orchestrator Control Center	3	10	3	HTTPS	GET	/vco-controlcenter/docs/		200

The completed configuration should look similar to the following screen:

The screenshot shows the 'Manage' tab selected in the top navigation bar. Under the 'Load Balancer' tab, the 'Service Monitoring' section is active. On the left, a sidebar lists 'Global Configuration', 'Application Profiles', 'Service Monitoring' (which is highlighted in blue), 'Pools', 'Virtual Servers', and 'Application Rules'. The main area displays a table of service monitors with the following data:

Monitor ID	Name	Type	Interval	Timeout	Max Retries
monitor-1	default_tcp_monitor	TCP	5	15	3
monitor-2	default_http_monitor	HTTP	5	15	3
monitor-3	default_https_monitor	HTTPS	5	15	3
monitor-4	vRealize Automation VA Web	HTTPS	3	10	3
monitor-5	IaaS Web	HTTPS	3	10	3
monitor-6	IaaS Manager	HTTPS	3	10	3
monitor-7	vRealize Orchestrator Control Center	HTTPS	3	10	3

Below the table, a summary for 'Service Monitor vRealize Orchestrator Control Center Details' is shown:

Name	vRealize Orchestrator Control	Type	https
Interval	3	Expected	200
Timeout	10	URL	/vco-controlcenter/docs/
Max Retries	3	Send	
Receive		Method	GET
Extension			

Add Pools

You must create the following pools for vRealize Automation.

1. Click **Pools** in the left pane.
2. Click the **Add** icon to create the pools required for vRealize Automation using the information in Table 6.
 - The environment depicted in Table 6 is an example. Your environment might contain two or three vRealize Automation virtual appliance nodes and two or more nodes per IaaS role.
 - You can either use the IP address of the pool members or select them as a Virtual Center Container.

TABLE 6 - ADD POOLS

POOL NAME	ALGORITHM	MONITORS	MEMBER NAME	EXAMPLE IP ADDRESS / VCENTER CONTAINER	PORT	MONITOR PORT
pool_vra-va-web_80	Least connections	default_http_monitor	vRA VA1	IP Address	80	
			vRA VA2	IP Address	80	
			vRA VA3	IP Address	80	
pool_vra-va-web_443	Least connections	vRA VA Web	vRA VA1	IP Address	443	
			vRA VA2	IP Address	443	
			vRA VA3	IP Address	443	
*pool_vra-rconsole_8444	Least connections	vRA VA Web	vRA VA1	IP Address	8444	443
			vRA VA2	IP Address	8444	443
			vRA VA3	IP Address	8444	443
pool_vro-cc_8283	Least connections	vRealize Orchestrator Control Center	vRA VA1	IP Address	8283	
			vRA VA2	IP Address	8283	
			vRA VA3	IP Address	8283	
pool_iaas-web_443	Least connections	IaaS Web	IaaS Web1	IP Address	443	
			IaaS Web2	IP Address	443	
pool_iaas-manager_443	**Least connections	IaaS Manager	IaaS Man1	IP Address	443	
			IaaS Man2	IP Address	443	

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

** The Manager Service uses [active-passive type of configuration](#) hence the load balancer will always send the traffic to the current active node regardless of the load balancing method.

Add Virtual Servers

You must to add the following Virtual Servers for vRealize Automation.

1. Click **Virtual Servers** on the left pane.
2. Click the **Add** icon to create the virtual servers required for vRealize Automation using the information in Table 7. Leave the default value when nothing is specified.

TABLE 7 - ADD VIRTUAL SERVERS

NAME	IP ADDRESS	PROTOCOL	PORT	DEFAULT POOL	APPLICATION PROFILE	APPLICATION RULE
vs_vra-va-web_80	IP Address	HTTP	80	pool_vra-va-web_80	vRA VA	
vs_vra-va-web_443	IP Address	HTTPS	443	pool_vra-va-web_443	vRA VA	
vs_iaas-web_443	IP Address	HTTPS	443	pool_iaas-web_443	IaaS Web	
vs_iaas-manager_443	IP Address	HTTPS	443	pool_iaas-manager_443	IaaS Manager	
*vs_vra-va-rconsole_8444	IP Address	HTTPS	8444	pool_vra-rconsole_8444	vRA VA	
vs_vro-cc_8283	IP Address	HTTPS	8283	pool_vro-cc_8283	vRealize Orchestrator Control Center	

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

The completed configuration should look similar to the following screen.

The screenshot shows the 'Manage' tab selected in the top navigation bar. Under the 'Virtual Servers' section, a table lists six virtual servers:

Virtual Server ID	Name	Default Pool	IP Address	Protocol	Port
virtualServer-2	vs_iaas-web_443	pool-3	10.145.153.104	https	443
virtualServer-3	vs_iaas-manager_443	pool-4	10.145.153.105	https	443
virtualServer-4	vs_vra-va-rconsole_8444	pool-5	10.145.153.103	https	8444
virtualServer-5	vs_vro-cc_8283	pool-6	10.145.153.103	https	8283
virtualServer-1	vs_vra-va-web_443	pool-2	10.145.153.103	https	443
virtualServer-6	vs_vra-va-web_80	pool-9	10.145.153.103	http	80

A tooltip indicates there are 6 items. Below the table, the details for 'Virtual Server vs_vra-va-web_443' are shown:

Description:		Application Profile:	applicationProfile-3
Connection Limit:	0	Connection Rate Limit:	0
Service Insertion Status:	Disabled	Acceleration Status:	Disabled
		Service Profile Status:	

At the bottom, a table lists application rules:

Rule Id	Name	Script
---------	------	--------

Configuring NSX-T

This document assumes that the NSX-T is already deployed in the environment and the Tier-1 gateway with the load balancer can access vRealize Automation components over a network.

Note: NSX-T 2.3 has a known issue, HTTPS monitor is not supported for FAST TCP virtual server's pool, which is fixed in 2.4.

Add Application Profiles

You need to create two application profiles.

Add the Application Profile for HTTP requests

1. Go to **Networking** → **Load Balancing** → **PROFILES**
2. Select Profile Type **APPLICATION**
3. Click the **ADD APPLICATION PROFILE** and select **HTTP**
4. Choose a **Name** for the profile

Example

The completed configuration for an application profile for HTTP request should look similar to the following screen:

The screenshot shows the 'ADD APPLICATION PROFILE' dialog for an HTTP profile. The profile is named 'vRA_HTTP_to_HTTPS'. The 'Type' is set to 'HTTP' and 'Idle Timeout (sec)' is set to '15'. Other settings include 'Request Header Size' (1024), 'Request Body Size' (empty), and 'NTLM Authentication' (disabled). The 'Tags' section shows 'Tag (Required)' and 'Scope (Optional)'. At the bottom are 'SAVE' and 'CANCEL' buttons.

Add the Application Profile for HTTPS requests

1. Go to **Networking** → **Load Balancing** → **PROFILES**
2. Select Profile Type **APPLICATION**
3. Click the **ADD APPLICATION PROFILE** and select **Fast TCP Profile**
4. Choose a **Name** for the profile

Example

The completed configuration for an application profile for HTTPS request should look similar to the following screen:

The screenshot shows the 'PROFILES' tab selected in the navigation bar. A sub-menu 'Select Profile Type' is open, showing 'APPLICATION' is selected. Below this, a table lists a single profile:

Name	Type	Idle Timeout (sec)	HA Flow Mirroring	Virtual Servers
vRA_HTTPS	Fast TCP	1800	Disabled	

Below the table are fields for 'Description' (Enter Description), 'Tags' (Tag (Required) and Scope (Optional)), and 'Connection Close Timeout' (set to 8). At the bottom are 'SAVE' and 'CANCEL' buttons.

Add Persistence Profile

1. Go to **Networking** → **Load Balancing** → **PROFILES**
2. Select Profile Type **PERSISTENCE**
3. Click the **ADD PERSISTENCE PROFILE** and select **Source IP**
4. Choose a Name for the profile

Example

The completed configuration for a persistence profile should look similar to the following screen:

The screenshot shows the 'PROFILES' tab selected in the navigation bar. A sub-menu 'Select Profile Type' is open, showing 'PERSISTENCE' is selected. Below this, a table lists a single profile:

Name	Type	Share Persistence	Virtual Servers
source_addr_vra		Disabled	

Below the table are fields for 'Description' (Enter Description), 'Tags' (Tag (Required) and Scope (Optional)), and 'Persistence Entry Timeout' (set to 300). Under 'Advanced Settings', 'Purge Entries when Full' is set to 'Enabled' and 'HA Persistence Mirroring' is set to 'Disabled'. At the bottom are 'SAVE' and 'CANCEL' buttons.

Add Active Health Monitor

1. Go to **Networking** → **Load Balancing** → **MONITORS**
2. Click the **Add ACTIVE MONITOR**, select **HTTPS**
3. Choose a Name for the Health Monitor. Set **Monitoring Port**, **Monitoring Interval**, **Timeout Period**, **Fall Count** and **Rise Count** (please refer to the table and example below)
4. Click **HTTP Request Configure** (please refer to the table and example below)
5. Click **HTTP Response Configure** (please refer to the table and example below)
6. Click **SSL Configuration Configure**
 - Server SSL **Enabled**
 - Server SSL Profile: **default-balance-server-ssl-profile**

TABLE 8 – CONFIGURE HEALTH MONITORS

NAME	PORT	INTERVAL	TIMEOUT	METHOD	REQUEST URL	RESPONSE CODE	RESPONSE BODY
vra_http_80	80	3	10	Get	/	302	
vra_https_va_web	443	3	10	Get	/vcac/services/api/health	200,204	
vra_https_iaas_web	443	3	10	Get	/wapi/api/status/web		REGISTERED
vra_https_iaas_mgr	443	3	10	Get	/VMPSProvision		ProvisionService
vro_https_8283	8283	3	10	Get	/vco-controlcenter/docs/	200	

Example

The completed configuration for a health monitor should look similar to the following screens:

The screenshot shows the 'MONITORS' tab selected in the navigation bar. A new monitor named 'vra_https_va_web' is being configured. The configuration fields include:

- Name:** vra_https_va_web
- Protocol:** HTTPS
- Monitoring Port:** 443
- Monitoring Interval:** 3
- Timeout Period (sec):** 10
- Description:** Enter Description
- Fall Count:** 3
- Tags:** Tag (Required) Scope (Optional)
- Rise Count:** 3
- Additional Properties:** HTTP Request, SSL Configuration, HTTP Response, Configure

At the bottom, there are 'SAVE' and 'CANCEL' buttons.

HTTP Request and Response Configuration X

Active Health Monitor -

HTTP Request Configuration HTTP Response Configuration

HTTP Method: Get ▼

HTTP Request URL: /vcac/services/api/heal

HTTP Request Version: 1.1 ▼

ADD

Header Name	Header Value
(empty)	Request Header not found

HTTP Request Body

CANCEL **APPLY**

HTTP Request and Response Configuration

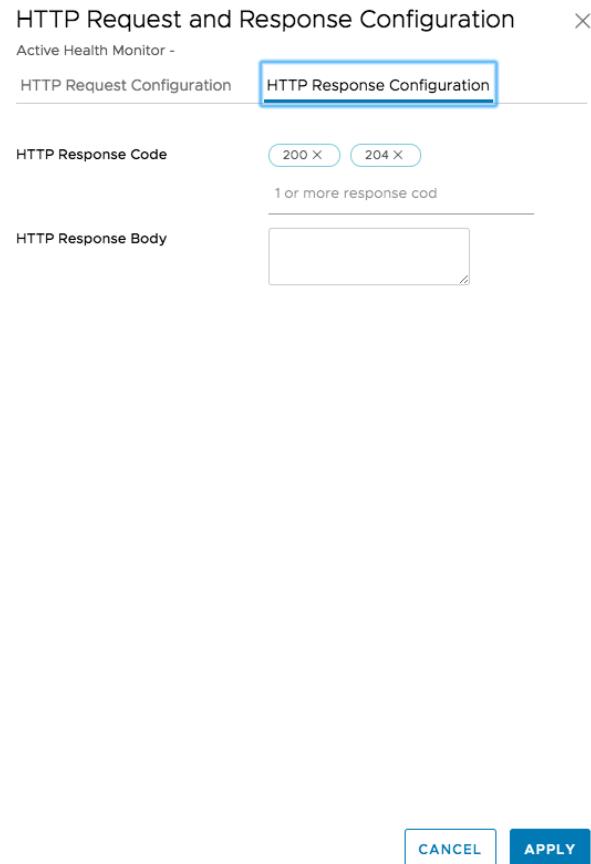
Active Health Monitor -

HTTP Request Configuration **HTTP Response Configuration**

HTTP Response Code **200 X 204 X**
1 or more response cod

HTTP Response Body

CANCEL **APPLY**



Edit SSL Configuration

Active Monitor vra_https_va...

Server SSL **Enabled**

Client Certificate Select Certificate ▾
default-balanced-server-ssl-pr (X) ▾

Server SSL Profile default-balanced-server-ssl-pr ▾

Advanced Properties

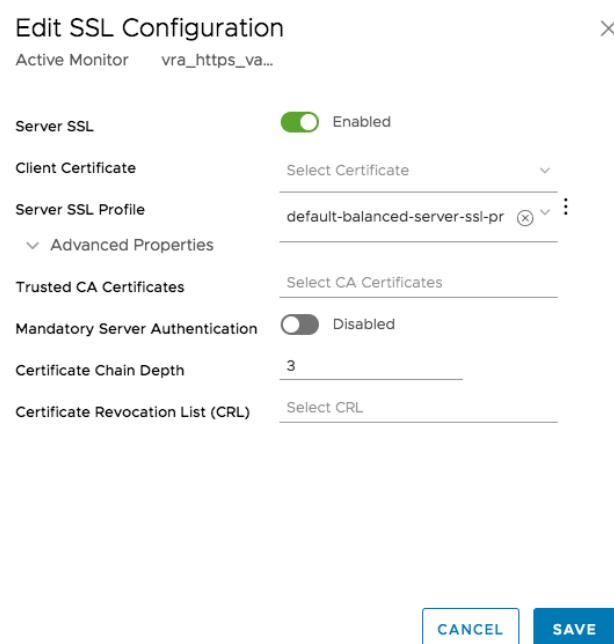
Trusted CA Certificates Select CA Certificates

Mandatory Server Authentication **Disabled**

Certificate Chain Depth 3

Certificate Revocation List (CRL) Select CRL

CANCEL **SAVE**



Configure Server Pools

You need to configure the following server pools for vRealize Automation

1. Go to **Networking → Load Balancing → SERVER POOLS**
2. Click the **ADD SERVER POOL**
3. Choose a **Name** for the pool. Set **Algorithm** as LEAST_CONNECTION
4. Configure **SNAT Translation** as Auto Map
5. Click **Select Members** and **ADD MEMBER** (please refer to the table and example below)
 - **Name**
 - **IP**
 - **Weight: 1**
 - **Port**
 - **State: ENABLED**

TABLE 9 – CONFIGURE SERVER POOLS

POOL NAME	ALGORITHM	ACTIVE MONITOR	NAME	IP	PORT
pool_vra-va-web_80	Least connections	vra_http_80	vra_va1	IP	80
			vra_va2	IP	80
pool_vra-va-web_443	Least connections	vra_https_va_web	vra_va1	IP	443
			vra_va2	IP	443
*pool_vra-rconsole_8444	Least connections	vra_https_va_web	vra_va1	IP	8444
			vra_va2	IP	8444
pool_vro-cc_8283	Least connections	vro_https_8283	vra_va1	IP	8283
			vra_va2	IP	8283
pool_iaas-web_443	Least connections	vra_https_iaas_web	vra_web1	IP	443
			vra_web2	IP	443
pool_iaas-manager_443	**Least connections	vra_https_iaas_mngr	vra_ms1	IP	443
			vra_ms1	IP	443

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

** The Manager Service uses active-passive type of configuration hence the load balancer will always send the traffic to the current active node regardless of the load balancing method.

Example

The completed configuration for a server pool should look similar to the following screen:

The screenshot shows two overlapping configuration screens. The top screen is titled "ADD SERVER POOL" and displays settings for a pool named "pool_vra-va-web_443". It includes fields for Description (Enter Description), SNAT Translation Mode (Automap), and Active Monitor (vra_https_va_web). The bottom screen is titled "Configure Server Pool Members" and shows the configuration of members for the pool "pool_iias-manager_443". It has two tabs: "Enter individual members" (selected) and "Select a group". The table lists two members, both assigned port 443 with weight 1 and state Enabled, and both marked as Disabled as backup members. At the bottom right of the member configuration screen are "CANCEL" and "APPLY" buttons.

Configure Virtual Servers

You need to add the following Virtual Servers for vRealize Automation.

1. Go to **Networking → Load Balancing → VIRTUAL SERVERS**
2. Click the **ADD VIRTUAL SERVER**, select Layer 4/7 (please refer to the table below)
3. Choose a **Name** for Virtual Server
4. Assign **IP Address** (Virtual IP) and **Port** (please refer to the table below)
5. Choose the **Server Pool** previously configured
6. Choose the **Application Profile** previously configured
7. Set **Persistence** if required (please refer to the table below)
8. Set the **Default Pool Member Ports** (please refer to the table below)

TABLE 10 – CONFIGURE VIRTUAL SERVERS

NAME	TYPE	APPLICATION PROFILE	IP ADDR	PORT	SERVER POOL	PERSISTENCE PROFILE
vs_vra-va-web_80	L7 HTTP	vRA_HTTP_to_HTTPS	IP	80	pool_vra-va-web_80	None
vs_vra-va-web_443	L4 TCP	vRA_HTTPS	IP	443	pool_vra-va-web_443	source_addr_vra
vs_iaas-web_443	L4 TCP	vRA_HTTPS	IP	443	pool_iaas-web_443	source_addr_vra
vs_iaas-manager_443	L4 TCP	vRA_HTTPS	IP	443	pool_iaas-manager_443	None
*vs_vra-va-rconsole_8444	L4 TCP	vRA_HTTPS	IP	8444	pool_vra-rconsole_8444	source_addr_vra
vs_vro-cc_8283	L4 TCP	vRA_HTTPS	IP	8283	pool_vro-cc_8283	source_addr_vra

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

Example

The completed configuration for a virtual server should look similar to the following screen:

The screenshot shows the 'LOAD BALANCERS' interface with the 'VIRTUAL SERVERS' tab selected. A new virtual server is being configured with the following details:

- Name:** vs_vra-va-web_443
- IP Address:** e.g. 10.10.10.10
- Ports:** 443
- Type:** L4 TCP
- Load Balancer:** Select Load Balancer
- Server Pool:** pool_vra-va-w
- Application Profile:** vRA_HTTPS
- Persistence:** Source IP
- Source IP:** source_addr_vra
- Additional Properties:**
 - Max Concurrent Connections: Unlimited
 - Default Pool Member Ports: 443 (Enter Ports or Port Ranges: e.g. 8080, 80-90, 443)
 - Sorry Server Pool: Select Server Pool
 - Admin State: Enabled
 - Access Log: Disabled
 - Tags: Tag (Required) Scope (Optional)

At the bottom, there are 'SAVE' and 'CANCEL' buttons.

Configure Load Balancer

You need to specify a load-balancer configuration parameter for vRealize Automation.

1. Go to **Networking → Load Balancing → LOAD BALANCERS**
2. Click the **ADD LOAD BALANCER**
3. Choose a **Name**, select appropriate **Load Balancer Size** (depends on vRA cluster size)
4. Choose the pre-created **Tier 1 Logical Router**

Note: In 2.4, the monitor health checks are done using the IP address of Tiers-1 uplink (or first service port for Tiers-1 standalone SR) for all server pools of the load-balancer. Please ensure that server pools are reachable from this IP address.

Example

The completed configuration for a load balancer should look similar to the following screen:

Name	Size	Tier-1 Gateway	Virtual Servers	Status
vra75_lb	Small	vRA-LB-Tier-1-Router		

Description: Enter Description

Error Log Level: Info

Tags: Tag (Required), Scope (Optional)

Admin State: Enabled

VIRTUAL SERVERS

SAVE **CANCEL**

Add Virtual Servers to Load Balancer

1. Go to **Networking → Load Balancing → VIRTUAL SERVERS**
2. Edit configured Virtual Servers
3. Assign **Load Balancer** as the previously configured Load Balancer

Example

The completed configuration for a virtual server should look similar to the following screen:

vs_vra-va-web_443	192.168.205.10	443	L4 TCP	vra75_lb	pool_vra-va-w
-------------------	----------------	-----	--------	----------	---------------

Description: Enter Description

Persistence: Source IP

Source IP: source_addr_vra

SAVE **CANCEL**

Configuring Citrix ADC (NetScaler ADC)

Before starting this configuration, ensure that the NetScaler device is deployed in the environment and has access to the vRealize Automation components.

- You can use either virtual or physical NetScaler
- The Citrix load balancer can be deployed in either one-arm or multi-arm topologies
- Enable the Load Balancer and SSL modules. You can do so from **NetScaler > System > Settings > Configure Basic Features** page.

Configure 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 for each row in Table . Leave the default value when nothing is specified.

TABLE 11 – CONFIGURE MONITORS

NAME	TYPE	INTERVAL	TIMEOUT	SEND STRING	RECEIVE STRING	DEST. PORT	SECURE
vra_https_va_web	HTTP	5	4	GET /vcac/services/api/health	HTTP/1\.(0 1)(200 204)	443	yes
vra_https_iaas_web	HTTP-ECV	5	4	GET /wapi/api/status/web	REGISTERED	443	yes
vra_https_iaas_mngr	HTTP-ECV	5	4	GET /VMPSProvision	ProvisionService	443	yes
vro_https_8283	HTTP	5	4	GET /vco-controlcenter/docs/	HTTP/1\.(0 1)(200)	8283	yes

Example

The completed configuration for a virtual appliance monitor should look similar to the following screen:

[Create Monitor](#)

Name*
vra_https_va_web

Type*
HTTP

Type of monitor that you want to create.

Basic Parameters

Interval
5 Second

Response Time-out
4 Second

Response Codes

204	x	?
200	x	?

Custom Header

HTTP Request
GET /vcac/services/api/health

Secure

SSL Profile

Bind Delete

Certificate Name

No items

Advanced Parameters

Create Close

This screenshot shows the 'Create Monitor' dialog box. The 'Name*' field is set to 'vra_https_va_web'. The 'Type*' field is set to 'HTTP', with a tooltip explaining it's the type of monitor being created. Under 'Basic Parameters', the 'Interval' is set to '5' with a unit of 'Second'. The 'Response Time-out' is set to '4' with a unit of 'Second'. In the 'Response Codes' section, two codes are listed: '204' and '200', each with a delete button ('x') and a help icon ('?'). The 'HTTP Request' field contains the URL 'GET /vcac/services/api/health'. The 'Secure' checkbox is checked. There is a 'SSL Profile' dropdown with a plus sign (+) and a edit icon. Below the request fields are 'Bind' and 'Delete' buttons. A 'Certificate Name' section shows 'No items'. At the bottom are 'Create' and 'Close' buttons.

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 for each row in Table .

Note: The environment depicted in Table is an example. Your environment might contain two or three vRealize Automation virtual appliance nodes and two or more nodes per IaaS role.

TABLE 12 – CONFIGURE SERVICE GROUPS

NAME	HEALTH MONITORS	PROTOCOL	SG MEMBERS	ADDRESS	PORT
pl_vra-va-00_80	http-ecv	HTTP	ra-vra-va-01	IP Address	80
			ra-vra-va-02	IP Address	80
			ra-vra-va-03	IP Address	80
pl_vra-va-00_443	vra_https_va_web	SSL Bridge	ra-vra-va-01	IP Address	443
			ra-vra-va-02	IP Address	443
			ra-vra-va-03	IP Address	443
*pl_vra-va-00_8444	vra_https_va_web	SSL Bridge	ra-vra-va-01	IP Address	8444
			ra-vra-va-02	IP Address	8444
			ra-vra-va-03	IP Address	8444
pl_vro-cc-00_8283	vro_https_8283	SSL Bridge	ra-vra-va-01	IP Address	8283
			ra-vra-va-02	IP Address	8283
			ra-vra-va-03	IP Address	8283
pl_iaas-web-00_443	vra_https_iaas_web	SSL Bridge	ra-web-01	IP Address	443
			ra-web-02	IP Address	443
pl_iaas-man-00_443	vra_https_iaas_mgr	SSL Bridge	ra-man-01	IP Address	443
			ra-man-02	IP Address	443

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

Example

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	OFF	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 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 for each entry in Table . Leave the default value when nothing is specified.

TABLE 13 – CONFIGURE VIRTUAL SERVERS

NAME	PROTOCOL	DESTINATION ADDRESS	PORT	LOAD BALANCING METHOD	SERVICE GROUP BINDING
vs_vra-va-00_80	HTTP	IP Address	80	Least connections	pl_vra-va-00_80
vs_vra-va-00_443	SSL Bridge	IP Address	443	Least connections	pl_vra-va-00_443
vs_web-00_443	SSL Bridge	IP Address	443	Least connections	pl_iaas-web-00_443
vs_man-00_443	SSL Bridge	IP Address	443	**Least connections	pl_iaas-man-00_443
*vs_vra-va-00_8444	SSL Bridge	IP Address	8444	Least connections	pl_vra-va-00_8444
vs_vro-cc-00_8283	SSL Bridge	IP Address	8283	Least connections	pl_vro-cc-00_8283

* Port 8444 is optional – it is required only if you want to use remote console from vRealize Automation.

** The Manager Service uses an [active-passive type of configuration](#) hence the load balancer will always send the traffic to the current active node regardless of the load-balancing method.

Example

Load Balancing Virtual Server

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

Basic Settings			
Name	SSL_BRIDGE	Listen Priority	-
Protocol	SSL_BRIDGE	Listen Policy Expression	NONE
State	UP	Range	1
IP Address	10.23.90.21	Redirection Mode	IP
Port	443	RHI State	PASSIVE
Traffic Domain	0	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	>	

Persistence			
Persistence	SOURCEIP	Time-out (mins)	30
		IPv4 Netmask	255.255.255.255
		IPv6 Mask Length	128

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

Configure Persistence Group

1. Log in to the NetScaler and select **NetScaler > Traffic Management > Load Balancing > Persistence Groups**.
2. Click **Add**.
3. Enter the name **source_addr_vra** and select **Persistence > SOURCEIP** from the drop-down menu.
4. Set the **Timeout** to **30 minutes**.
5. Add all Virtual Servers related to vRealize Automation.
 - vra_80
 - vra_443
 - vro_443
 - web_443
 - controlcenter_8283
6. Click **OK**.

Example

The completed configuration should look like the following screen.

Create Persistence Group

Group Name*
source_addr_vra

Persistence*
SOURCEIP

IPv4 Netmask
255 . 255 . 255 . 255

IPv6 Mask Length
128

Time-out
30

Backup Persistence*
NONE

Use vServer Persistence*

Virtual Server Name*

Available (12)	Select All
nnrvava_443	
nnrvaweb	
nnrvava_8283	
nnrvava_80	
New	

Configured (0)	Remove All
No items	

Create **Close**

Troubleshooting

Provisioning failures when using OneConnect with F5 BIG-IP for a virtual server with SSL pass-through

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 select **Local Traffic > Virtual Servers > Virtual Server List**.
2. Click the name of the virtual server to modify.
3. Choose **None** for the **OneConnect Profile** option in the **Acceleration** section and click **Finish**.

F5 BIG-IP license limits network bandwidth

If you experience provisioning failures or issues loading vRealize Automation console pages, especially during periods of a high utilization, network traffic to and from the load balancer might exceed what the F5 BIG-IP license allows.

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

Proxy agent ping failure

After starting the Manager Service on a second manager server, the proxy agent is unable to reconnect. This happens because the F5 appliance is still maintaining an SSL session with the agent by sending keepalives while the agent is trying to establish a new session.

To configure the load balancer to drop all packets and prevent it from sending keepalives to resolve this issue:

1. Log in to the F5 load balancer and select **Local Traffic > Pools**.
2. Select the **Manager Service** pool.
3. Click **Advanced** in the **Configuration** section.
4. Select **Drop** for the **Action On Service Down** option.
5. Click **OK** and click **Finished**.

Connection reset errors in the catalina.out log file

When the system is under a heavy load due to many simultaneously requested provisions through the IaaS components, you might see connection reset errors in the catalina.out log file on the vRealize Automation appliances. This can happen when a session between the appliances and the Web servers expires. You can work around this problem by increasing the timeout setting for your load balancer.

F5

Use the AskF5 procedure K7166: “Changing the idle timeout for a protocol profile” at <https://support.f5.com/csp/article/K7166> to change the Idle Timeout for a virtual server. Perform this procedure on the Web Service virtual load balancer server and set the timeout initially to 600 seconds. A best practice is to gradually increase the timeout until there are no connection reset errors.

NetScaler

Use the Citrix procedure in “Setting a Time-out Value for Idle Client Connections” at <https://docs.citrix.com/en-us/netscaler/11/traffic-management/load-balancing/load-balancing-manage-clienttraffic/client-idle-timeout-value.html> to change the Idle Timeout for a virtual server. Perform this procedure on the Web Service virtual load balancer server and set the timeout initially to 600 seconds. A best practice is to gradually increase the timeout until there are no connection reset errors.

NSX-V

Use the procedure in the [knowledge base article 2147156](#) to change the Idle Timeout for a virtual server. Perform this procedure on the Web Service virtual load balancer server and set the timeout initially to 600 seconds. A best practice is to gradually increase the timeout until there are no connection reset errors.

Proxy Agents cannot connect to load balanced Manager Service endpoint

With NSX as a load balancing solution, after a Manager Service failover the proxy agents cannot connect to the load balanced Manager Service.

This issue occurs when the IIS role is installed on the servers that are running the Manager Service. When a Manager Service is stopped, the monitors configured on NSX-V are flagged as DOWN. Because IIS is running and accepting connections on port 443, NSX-V keeps the established sessions, and the Proxy Agent service keeps trying to reuse the session.

To resolve this issue, remove the IIS role from the servers that are running the Manager Service. The Manager Service is a self-hosted service that does not require IIS.



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2015-2018 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.