



# vRealize Automation Load Balancing

Configuration Guide  
Version 7.1.x

TECHNICAL WHITE PAPER

AUGUST 2016

VERSION 2.4

## Table of Contents

Introduction.....	4
Load Balancing Concepts.....	4
SSL Pass-Through.....	4
Session Persistence.....	4
Source IP Address Hash (NSX).....	5
Email notifications on Load Balancer.....	5
One-arm or Multi-arm Topologies.....	5
Prerequisites for configuring Load Balancers with vRealize Automation.....	6
Completing the vRealize Automation Initial Installation.....	7
Configuring F5 Big IP.....	8
Configure Custom Persistence Profile.....	8
Configure Monitors.....	8
Configure Server Pools.....	10
Configure Virtual Servers.....	11
Configuring NSX 6.2.x.....	14
Configure Global Settings.....	14
Add Application Profiles.....	15
Add Service Monitoring.....	16
Add Pools.....	17
Add Virtual Servers.....	17
Configuring Citrix NetScaler.....	18
Configure Monitors.....	18
Configure Service Groups.....	21
Configure Virtual Servers.....	22
Configure Persistence Group.....	24

**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"><li>▪ Updated timeout to 10 seconds for Configure Monitors and Add Service Monitoring in F5 and NSX sections respectively</li><li>▪ Added source IP persistence and timeout of 1800 seconds for Add Application Profiles section</li><li>▪ Updated all the screenshots to match the content</li><li>▪ Updated NSX load balancing method to be round robin</li></ul>
August 2016	2.4	<ul style="list-style-type: none"><li>▪ Added configuration for Citrix NetScaler</li><li>▪ Updated for NSX 6.2</li></ul>

# Introduction

This document describes the configuration of the load balancing modules of F5 Networks BIG-IP software (F5) and NSX load balancers for vRealize Automation 7.0.x in a distributed and high availability deployment. This document is not an installation guide, but a load-balancing configuration guide that supplements the vRealize Automation installation and configuration documentation available in the *vRealize Automation Installation and Configuration* guide in the VMware vRealize Automation 7.0 Documentation Center.

This information is for the following products and versions.

PRODUCT	VERSION
F5 BIG IP	Tested with 11.6
NSX	6.2.3, 6.1.4, 6.1.3 (versions earlier than 6.1.3 are not supported)
Citrix NetScaler	Tested with 10.5
vRealize Automation	7.0.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 recommended in this document is the result of extensive testing and represents the best balance between stability, performance, and scalability.

#### *Destination Address (F5)*

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 & NSX)*

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. 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 (30 minutes) to match the vRealize Automation GUI timeout.

### Source IP Address Hash (NSX)

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 does not support email notification for such a scenario.

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 components to be load balanced and the load balancers' virtual IP are on the same network. 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 traffic, return traffic goes directly back to the client and connections 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. For the purpose of this document, the vRealize Automation components are deployed as a one-arm configuration as shown in [Figure 1](#).

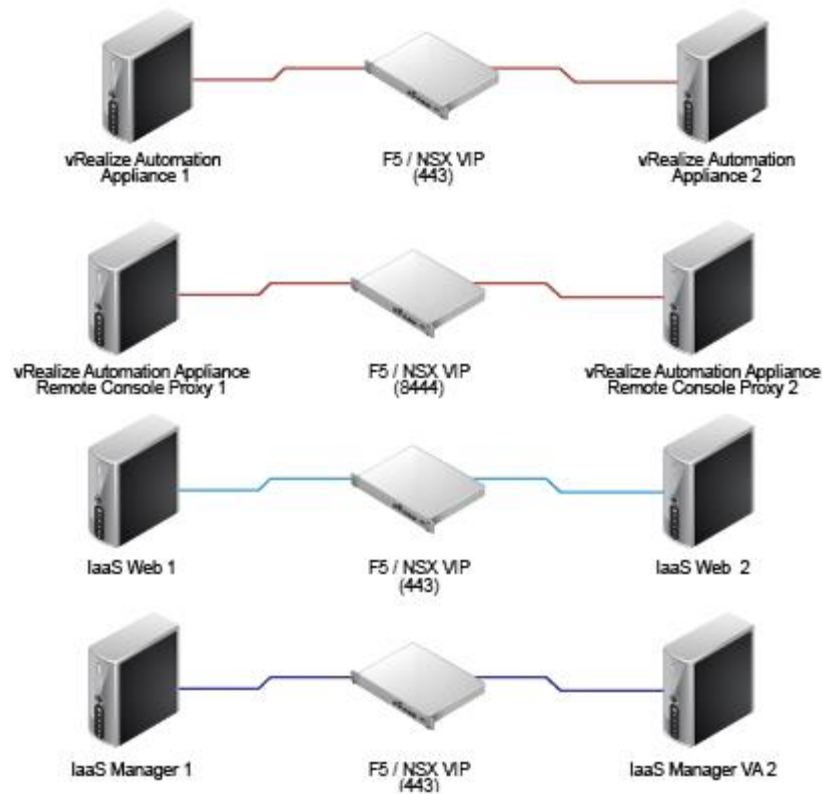


FIGURE 1. ONE-ARM CONFIGURATION

## Prerequisites for Configuring Load Balancers with vRealize Automation

- **F5** - Before you start the 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.
- **NSX** - Before you start the HA implementation of vRealize Automation using NSX as a load balancer, ensure that your NSX topology is configured and that your version of NSX is supported. This document covers the load balancing aspect of an NSX configuration, and assumes that NSX 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** - Create signed or self-signed certificates to contain the vRealize Automation virtual IP and the hostnames 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** - 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 an Microsoft SQL Server instance.

For more information on installation and configuration see [vRealize Automation 7.0 product documentation](#).

If required, external 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 Orchestrator cluster.

## Complete the vRealize Automation Initial Installation

During the initial setup process, the load balancer with all nodes enabled routes half of the traffic to the secondary nodes, which are not yet installed, and the installation fails. To avoid these failures and to complete the initial installation of a vRealize Automation, you must perform the following tasks.

1. Configure the F5, NSX, or NetScaler load balancer. See [Configuring F5 Big IP](#), [Configuring NSX 6.2.x](#), 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 of the components are installed, enable all nodes on the load balancer.
6. Configure the load balancer with all of the 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 available **connectors** are enabled and configured for authentication for each virtual appliance node.

The screenshot shows the vRealize Automation Administration console. The left sidebar contains a navigation menu with options: Administration, Directories, Policies, Identity Providers (selected), Connectors, User Attributes, Network Ranges, and Password Recovery. The main content area is titled 'WorkspaceIDP\_\_1' and shows the configuration for an Identity Provider. The configuration includes fields for Identity Provider Name (WorkspaceIDP\_\_1), Users (sqa), Network (ALL RANGES), Authentication Methods (Password), Connector(s) (vra-va-node1-host, vra-va-node2-host), and IdP Hostname (vra-va-load-balancer-vip). A 'Disable IdP' button is visible below the main configuration area. At the bottom right, there are 'Save' and 'Cancel' buttons.

## Configuring F5 Big IP

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 and can be deployed in 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 earlier 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 see [HTTP health checks may fail even though the node is responding correctly](#).

### Configure Custom Persistence Profile

You can configure 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 can configure required monitors for your F5 load balancer.

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	TIMEOUT	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_mgr	HTTPS	3	10	GET /VMPSProvision\r\n	ProvisionService	



### Example

The completed 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	Services on the vRealize Automation Virtual Appliance
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	
Cipher List	DEFAULT:+SHA:+3DES:+kEDH
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	443 HTTPS ▼

Cancel Repeat Finished

## Configure Server Pools

You can configure server pools for your F5 load balancer.

1. Log in to the F5 load balancer and select **Local Traffic > Pools**.
2. Click **Create** and provide the required information. Leave the default when nothing is specified.
3. Enter each pool member as a **New Node** and add it to the **New Members**.
4. Repeat steps 1, 2, and 3 for each row of information in Table 2.
5. To check the network map for an overall view of the server pools, select **LTM > Network Map**.

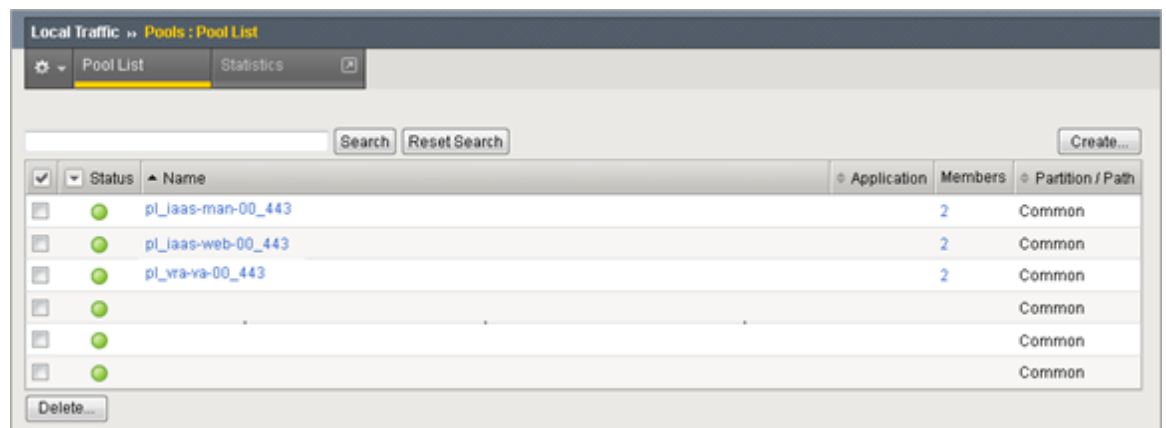
**TABLE 2 – CONFIGURE SERVER POOLS**

NAME	HEALTH MONITORS	LOAD BALANCING METHOD	NODE NAME	ADDRESS	SERVICE PORT
pl_vra-va-00_443	vra_https_va_web	Round Robin	ra-vra-va-01	10.26.38.44	443
			ra-vra-va-02	10.26.38.45	443
pl_iaas-web-00_443	vra_https_iaas_web	Round Robin	ra-web-01	10.26.38.49	443
			ra-web-02	10.26.38.50	443
pl_iaas-man-00_443	vra_https_iaas_mgr	Round Robin	ra-man-01	10.26.38.46	443
			ra-man-02	10.26.38.59	443
*pl_vra-va-00_8444	vra_https_va_web	Round Robin	ra-vra-va-01	10.26.38.44	8444
			ra-vra-va-02	10.26.38.45	8444

\*Port 8444 is optional – it is used for the remote-console connectivity

### Example

The completed configuration should look similar to the following screen.



## Configure Virtual Servers

You can configure virtual servers for your F5 load balancer.

1. Log in to the F5 load balancer and select **Local Traffic > Virtual Servers**.
2. Click **Create** and provide the required information. Leave the default when nothing is specified.
3. Repeat steps 1 and 2 for each entry in Table 3.
4. To check the network map for an overall view 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_443	Performance (Layer 4)	10.26.38.40	443	Auto Map	pl_vra-va-00_443	source_addr_vra
vs_web-00_443	Performance (Layer 4)	10.26.38.41	443	Auto Map	pl_iaas-web-00_443	source_addr_vra
vs_man-00_443	Performance (Layer 4)	10.26.38.42	443	Auto Map	pl_iaas-man-00_443	None
vs_vra-va-00_8444	Performance (Layer 4)	10.26.38.40	8444	Auto Map	pl_vra-va-00_8444	source_addr_vra

*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	
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network Address: 10.23.38.40
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 ▼
SMTSPS Profile	None ▼
VLAN and Tunnel Traffic	All VLANs and Tunnels ▼
Source Address Translation	Auto Map ▼

**Acceleration**

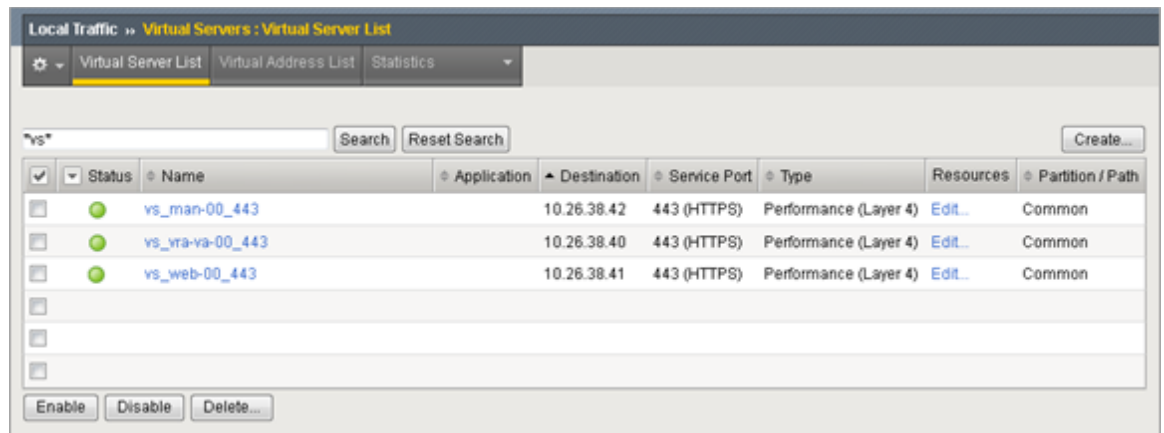
Rate Class	None ▼
SPDY Profile	None ▼

**Resources**

iRules	Enabled	Available
	<div> <div></div> <div>&lt;&lt;</div> <div>&gt;&gt;</div> <div>Up Down</div> </div>	<div> Common _sys_APM_ExchangeSupport_OA_BasicAuth _sys_APM_ExchangeSupport_OA_NtlmAuth _sys_APM_ExchangeSupport_helper _sys_APM_ExchangeSupport_main </div>
Default Pool	+ pl_vra-va-00_443 ▼	
Default Persistence Profile	source_addr_vra ▼	
Fallback Persistence Profile	None ▼	

Cancel Repeat Finished

The completed configuration should look similar to the following screen.



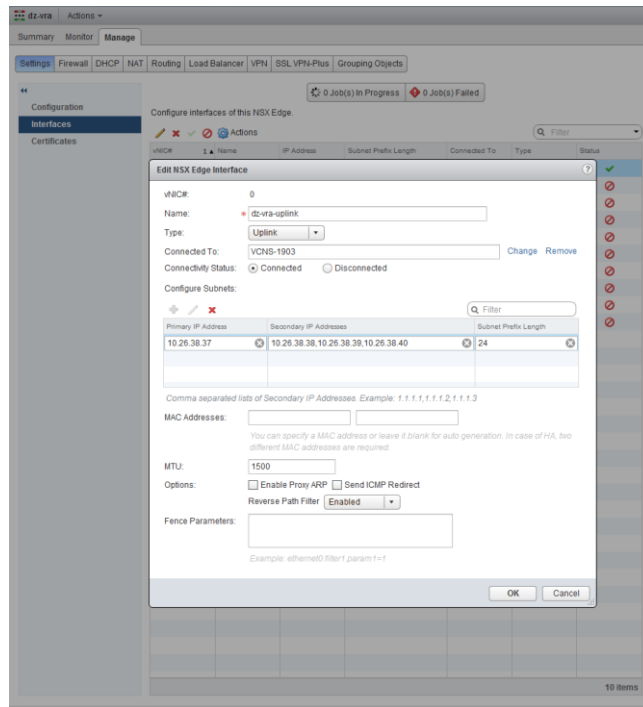
## Configuring NSX 6.2.x

You can deploy a new NSX Edge Services Gateway or use an existing one. It must have network connectivity to and from the vRealize Automation components being load balanced.

### Configure Global Settings

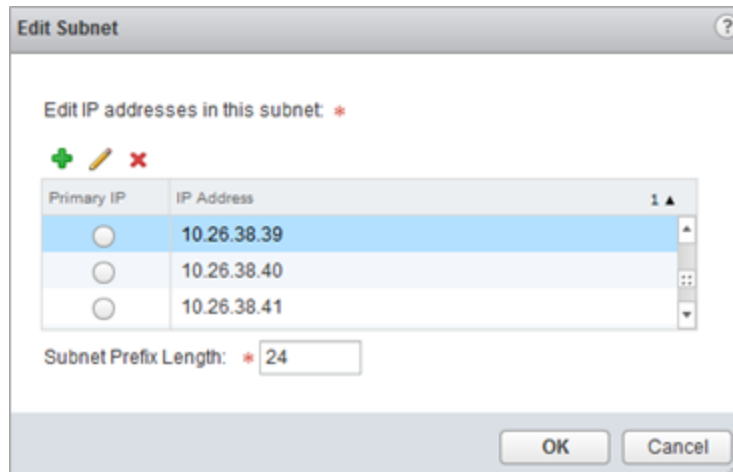
You can configure the global settings by using the following steps.

1. Log in to the NSX, click the **Manage** tab, click **Settings**, and select **Interfaces**.
2. Double-click to 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 Edge and click the **Edit** icon.



This interface might look different in NSX 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.



**Edit Subnet**

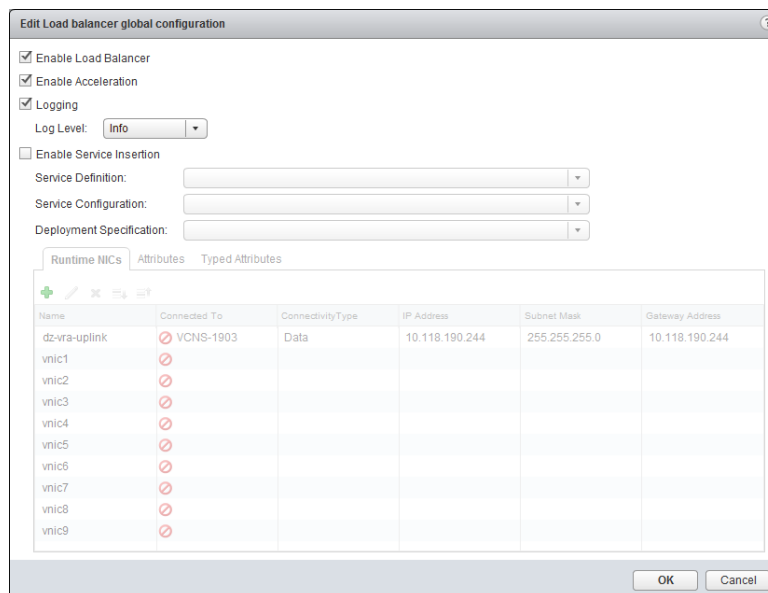
Edit IP addresses in this subnet: \*

Primary IP	IP Address
<input type="radio"/>	10.26.38.39
<input type="radio"/>	10.26.38.40
<input type="radio"/>	10.26.38.41

Subnet Prefix Length: \* 24

OK Cancel

- On the **Load Balancer** tab, click the **Edit** icon.
- Select **Enable Load Balancer**, **Enable Acceleration**, and **Logging**, if required, and click **OK**.



**Edit Load balancer global configuration**

☒ Enable Load Balancer  
☒ Enable Acceleration  
☒ Logging  
 Log Level: Info

☐ Enable Service Insertion  
 Service Definition:   
 Service Configuration:   
 Deployment Specification:

Runtime NICs Attributes Typed Attributes

Name	Connected To	Connectivity Type	IP Address	Subnet Mask	Gateway Address
dz-vra-uplink	✓ VCNS-1903	Data	10.118.190.244	255.255.255.0	10.118.190.244
vnic1	✗				
vnic2	✗				
vnic3	✗				
vnic4	✗				
vnic5	✗				
vnic6	✗				
vnic7	✗				
vnic8	✗				
vnic9	✗				

OK Cancel

This interface might look different in NSX 6.1.x and earlier.

## Add Application Profiles

You can add application profiles for different components of vRealize Automation.

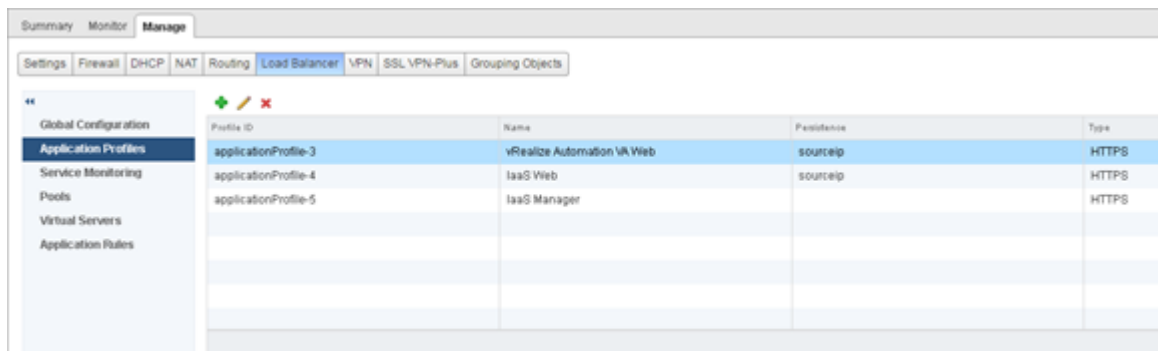
- Click **Application Profiles** in the pane on the left.
- Click the **Add** icon to create the application profiles required for vRealize Automation by using information in Table 4. Leave the default when nothing is specified.

**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

*Example*

The completed configuration should look similar to the following screen.

**Add Service Monitoring**

You can add service monitoring for different components of vRealize Automation.

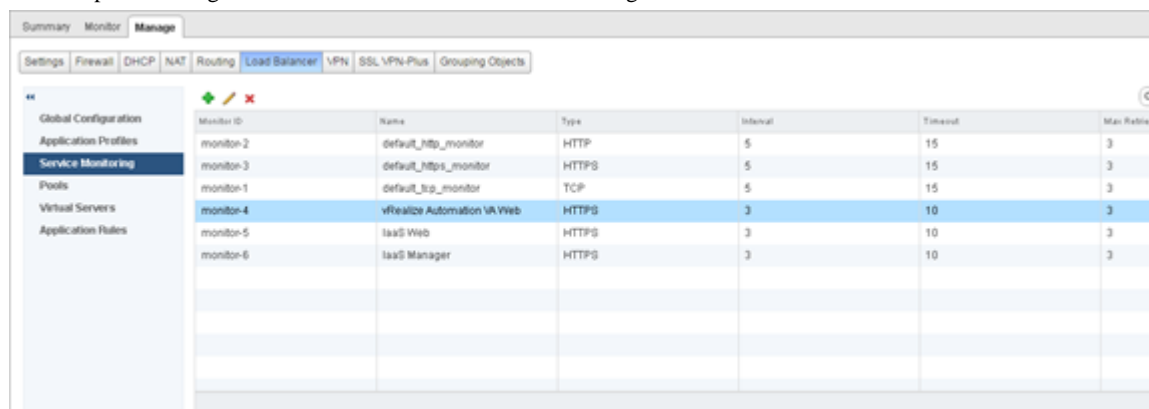
1. Click **Service Monitoring** in the left pane.
2. Click the **Add** icon to create the service monitors required for vRealize Automation using information in Table 5. Leave the default when nothing is specified.

**TABLE 5 – ADD SERVICE MONITORING**

NAME	INTERVAL	TIME OUT	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	/VMPSProvision	ProvisionService	



The completed configuration should look similar to the following screen.



## Add Pools

You can create pools.

1. Click **Pools** in the left pane.
2. Click the **Add** icon to create the Pools required for vRealize Automation using information in Table 6. Leave the default when nothing is specified.

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_443	Round Robin	vRA VA Web	vRA VA1	10.26.38.44	443	
			vRA VA2	10.26.38.45	443	
pool_iaas-web_443	Round Robin	IaaS Web	IaaS Web1	10.26.38.49	443	
			IaaS Web2	10.26.38.50	443	
pool_iaas-manager_443	Round Robin	IaaS Manager	IaaS Man1	10.26.38.49	443	
			IaaS Man2	10.26.38.50	443	
*pool_vra-rconsole_8444	Round Robin	vRA VA Web	vRA VA1	10.26.38.44	8444	443
			vRA VA2	10.26.38.45	8444	443

\*Only needed if remote-console access is used

## Add Virtual Servers

You can add virtual servers.

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 when nothing is specified.

**TABLE 7 - ADD VIRTUAL SERVERS**

NAME	IP ADDRESS	PROTOCOL	PORT	DEFAULT POOL	APPLICATION PROFILE	APPLICATION RULE
vs_vra-va-web_443	10.26.38.40	HTTPS	443	pool_vra-va-web_443	vRA VA	
vs_iaas-web_443	10.26.38.41	HTTPS	443	pool_iaas-web_443	IaaS Web	
vs_iaas-manager_443	10.26.38.42	HTTPS	443	pool_iaas-manager_443	IaaS Manager	
*vs_vra-va-rconsole_8444	10.26.38.40	HTTPS	8444	pool_vra-rconsole_8444	vRA VA	

\*Only needed if remote-console access is used

The completed configuration should look similar to the following screen.

Virtual Server ID	Name	Default Pool	IP Address	Protocol	Port
virtualServer-16	vs_iaas-manager_443	pool-4	10.26.38.42	https	443
virtualServer-15	vs_iaas-web_443	pool-3	10.26.38.41	https	443
virtualServer-14	vs_vra-va-web_443	pool-2	10.26.38.40	https	443

## Configuring Citrix NetScaler

Before beginning 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 in single or clustered configuration.
- Enable the Load Balancer and SSL modules. You can do so from the **NetScaler > System > Settings > Configure Basic Features** page.

### Configure Monitors

- Log in to the NetScaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Monitors**
- Click **Add** and provide the required information for each row of information in Table 8. **Error! Reference source not found.**  
Leave the default value when nothing is specified.

**TABLE 8 – CONFIGURE MONITORS**

NAME	TYPE	INTERVAL	TIMEOUT	SEND STRING	RECEIVE STRING	DEST. PORT	SECURE
vra_https_va_web	HTTP	3	10	GET /vcac/services/api/health	HTTP/1\.(0 1) (200 204)	443	yes
vra_https_iaas_web	HTTP-ECV	3	10	GET /wapi/api/status/web	REGISTERED	443	yes
vra_https_iaas_mgr	HTTP-ECV	3	10	GET /VMPSProvision	ProvisionService	443	yes

## Example

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

Configure Monitor

Name

vra\_https\_va\_web

Type

HTTP

Standard Parameters

Special Parameters

Interval

5

Second

Destination IP

.

.

.

IPv6

Response Time-out

4

Second

Destination Port

Bound Service

Down Time

30

Second

Dynamic Time-out

0

Deviation

0

Second

Dynamic Interval

0

Retries

3

Resp Time-out Threshold

0

SNMP Alert Retries

0

Action

Success Retries

1

Failure Retries

0

Net Profile

TOS

TOS ID

Enabled

Reverse

Transparent

LRTM (Least Response Time using Monitoring)

Secure

IP Tunnel

OK

Close

**Configure Monitor**

Name  
vra\_https\_va\_web

Type  
HTTP

Standard Parameters    Special Parameters

HTTP Request  
GET /vrao/services/api/hea

Response Codes  
200  
204

Custom Header

OK Close

## Configure Service Groups

1. Log in to the NetScaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Service Groups**.
2. Click **Add** and provide the required information for each row in Table 9. Add details of each pool member as a **Member** and add it to the **New Members** type **Server Based**. Leave the default when nothing is specified.

**TABLE 9 – CONFIGURE SERVICE GROUPS**

NAME	HEALTH MONITORS	PROTOCOL	SG MEMBERS	ADDRESS	PORT
pl_vra-va-00_443	vra_https_va_web	SSL Bridge	ra-vra-va-01	10.23.90.102	443
			ra-vra-va-02	10.23.90.103	443
pl_iaas-web-00_443	vra_https_iaas_web	SSL Bridge	ra-web-01	10.23.90.106	443
			ra-web-02	10.23.90.107	443
pl_iaas-man-00_443	vra_https_iaas_mgr	SSL Bridge	ra-man-01	10.23.90.109	443
			ra-man-02	10.23.90.110	443
*pl_vra-va-00_8444	vra_https_va_web	SSL Bridge	ra-vra-va-01	10.23.90.102	8444
			ra-vra-va-02	10.23.90.103	8444

\*port 8444 is optional – it is used for the remote-console connectivity

## Example

Complete service group for IaaS Managment.

Load Balancing Service Group

**Basic Settings**

Name  
Protocol  
State  
Effective State  
Traffic Domain

pl\_iaaS-man-00\_443  
SSL\_BRIDGE  
**ENABLED**  
Up  
0

Cache Type  
Cacheable  
Health Monitoring  
AppFlow Logging  
Number of Active Connections  
AutoScale Mode

**SERVER**  
**NO**  
**YES**  
**ENABLED**  
**0**  
-

**Service Group Members**

1 Service Group Member

**Monitors**

1 Service Group to Monitor Binding

Done

NetScaler > Traffic Management > Load Balancing > Servers

Name	State	IPAddress / Domain	Traffic Domain
ab-vra-va01	Enabled	10.23.90.102	0
ab-vra-va02	Enabled	10.23.90.103	0
ab-vra-web01	Enabled	10.23.90.106	0
ab-vra-web02	Enabled	10.23.90.107	0
ab-vra-mgr01	Enabled	10.23.90.109	0
ab-vra-mgr02	Disabled	10.23.90.110	0

After you have created all of your service groups, your service groups should look like the following screen.

NetScaler > Traffic Management > Load Balancing > Service Groups

Service Group Name	State	Effective State	Protocol	Max Clients	Max Requests	Maximum Bandwidth (Kbps)	Monitor Threshold	Traffic Domain
pl_vra-va-00_443	ENABLED	UP	SSL_BRIDGE	0	0	0	0	0
pl_iaaS-web-00_443	ENABLED	UP	SSL_BRIDGE	0	0	0	0	0
pl_iaaS-man-00_443	ENABLED	UP	SSL_BRIDGE	0	0	0	0	0

## Configure Virtual Servers

- Log in to the NetScaler load balancer and select **NetScaler > Traffic Management > Load Balancing > Virtual Servers**.
- Click **Add** and provide the required information for each entry in Table 10. Leave the default when nothing is specified.

**TABLE 10 – CONFIGURE VIRTUAL SERVERS**

NAME	PROTOCOL	DESTINATION ADDRESS	PORT	LOAD BALANCING METHOD	SERVICE GROUP BINDING
vs_vra-va-00_443	SSL Bridge	10.23.89.98	443	Roundrobin	pl_vra-va-00_443
vs_web-00_443	SSL Bridge	10.23.89.99	443	Roundrobin	pl_iaas-web-00_443
vs_man-00_443	SSL Bridge	10.23.89.100	443	Roundrobin	pl_iaas-man-00_443
vs_vra-va-00_8444	SSL Bridge	10.23.89.98	8444	Roundrobin	pl_vra-va-00_8444

### Example

Load Balancing Virtual Server

Basic Settings

Name

vs\_man-00\_443

Protocol

SSL\_BRIDGE

State

Down

IP Address

10.23.89.100

Port

443

Traffic Domain

0

Listen Priority

-

Listen Policy Expression

None

Range

1

Redirection Mode

IP

RHI State

PASSIVE

AppFlow Logging

ENABLED

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

Method

Load Balancing Method

ROUNDROBIN

New Service Startup Request Rate

-

New Service Request unit

PER\_SECOND

Increment Interval

-

Done

Help

Advanced

Policies

Protection

Profiles

Traffic Settings

The completed configuration should look similar to the following screen.

NetScaler > Traffic Management > Load Balancing > Virtual Servers										<a href="#">Refresh</a> <a href="#">Help</a> <a href="#">Save</a>
<a href="#">Add</a> <a href="#">Edit</a> <a href="#">Delete</a> <a href="#">Enable</a> <a href="#">Disable</a> <a href="#">Statistics</a> <a href="#">Action</a>										
Name	State	Effective State	IP Address	Port	Protocol	Method	Persistence	% Health	Traffic Domain	
▶ vs_vra-va-00_443	Up	Up	10.23.89.98	443	SSL_BRIDGE	ROUNDROBIN	SOURCEIP	100.00% 2 UP/0 DOWN	0	
▶ vs_web-00_443	Up	Up	10.23.89.99	443	SSL_BRIDGE	ROUNDROBIN	SOURCEIP	100.00% 2 UP/0 DOWN	0	
▶ vs_man-00_443	Up	Up	10.23.89.100	443	SSL_BRIDGE	ROUNDROBIN	SOURCEIP	100.00% 1 UP/0 DOWN	0	

## 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 related Virtual Servers.
6. Click **OK**.

### Example

The completed configuration should look similar to 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

Virtual Server Name\*

Available (0)

Select All

No items

New

Configured (3)

Remove All

vs\_man-00\_443

vs\_web-00\_443

vs\_vra-va-00\_443

Create

Close



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 [www.vmware.com](http://www.vmware.com)

Copyright © 2015 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.