

Installing vRealize Network Insight

VMware vRealize Network Insight 5.1

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

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

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

Copyright © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About vRealize Network Insight Installation Guide	5
1 Preparing for Installation	6
System Recommendations and Requirements	6
Privileges	9
System Ports	10
Network Communication Ports	16
Supported Products and Versions	18
2 Installing vRealize Network Insight	21
Installation Workflow	21
Deploying vRealize Network Insight Platform OVA	23
Deployment using vSphere Web Client	23
Deployment Using vSphere Windows Native Client	25
Activating the License	26
Generating Shared Secret	27
Setting up Network Insight Collector (OVA)	27
Deployment Using vSphere Web Client	27
Deployment using vSphere Windows Native Client	28
Setting up Network Insight Collector (AMI) in AWS for VMware SD-WAN	30
Deploy Additional Collector to an Existing Setup	31
3 Accessing vRealize Network Insight by using the Evaluation License	32
Add vCenter Server	32
Analyze Traffic Flows	33
Generate a Report	34
4 Planning to Scale up your Deployment	35
Planning to Scale up the Platform Cluster	35
Planning to Scale up the Collector	36
Increase the Brick Size of your Setup	37
5 Upgrading vRealize Network Insight	38
Online Upgrade	39
Single-Click Offline Upgrade	41
CLI Upgrade	44
6 Uninstall vRealize Network Insight	46

Remove Collector IP When Netflow is Enabled in vCenter	46
Remove Collector IP When Netflow is Enabled in NSX	47

About vRealize Network Insight Installation Guide

The *vRealize Network Insight Installation Guide* is intended for administrators or specialists responsible for installing vRealize Network Insight.

Intended Audience

This information is intended for administrators or specialists responsible for installing vRealize Network Insight. The information is written for experienced virtual machine administrators who are familiar with enterprise management applications and datacenter operations.

Preparing for Installation

1

Before you install vRealize Network Insight, prepare the deployment environment to meet the system requirements.

This chapter includes the following topics:

- [System Recommendations and Requirements](#)
- [Supported Products and Versions](#)

System Recommendations and Requirements

For optimum performance, you must match the minimum recommendations for the deployment.

Recommendations for the Platform Deployment

Table 1-1. Specifications for Platform Brick Size

Brick Size	Cores required for 2.1 GHz CPU	Cores required for 2.3 GHz CPU	Cores required for 2.6 GHz CPU	RAM	Disk
Medium	10	9	8	32 GB	1 TB
Large	15	14	12	48 GB	1 TB
Extra Large	20	18	16	64 GB	2 TB

Note

- The reservation for the CPU speed and RAM for each node must be 100% of the value specified above.
- To match your setup to all the specifications, you might have to add the resources (RAM, Disk, CPU). See <https://kb.vmware.com/s/article/53550> and [Increase the Brick Size of your Setup](#).

Table 1-2. Non-Cluster Deployment - Maximum Capacity

Brick Size	Number of VMs	Flows per Day	Total Flows	Flow Planning
Medium	4K	1M	4M	2M
Large	6K	2M	8M	4M

Table 1-3. Non-Cluster Deployment - Maximum Capacity for VMware SD-WAN

Brick Size	Number of Edges	Flows per Day	Total Flows
Medium	2K	1M	4M
Large	2K	2M	8M

Note

- The count of VMs includes the templates on the vCenter as well.
- Total Flows is the maximum count of flows the system can store for the retention period.
- Flow Planning is the total flows for which the system can perform security planning.

Table 1-4. Cluster Deployment - Maximum Capacity

Brick Size	Cluster Size	Number of VMs	Flows per Day	Total Flows	Flow Planning
Large	3	10K	2M	8M	4M
Extra Large	3	18K	6M	24M	4M
Extra Large	5	30K	10M	40M	4M
Extra Large	10	100K	10M	40M	4M

Table 1-5. Cluster Deployment - Maximum Capacity for VMware SD-WAN

Brick Size	Cluster Size	Number of Edges	Flows per Day	Total Flows
Large	3	4K	2M	8M
Extra Large	3	6K	6M	24M
Extra Large	5	10K	10M	40M

Note

- The count of VMs includes the templates on the vCenter as well.
- Cluster size is the total number of nodes in the cluster.
- Total Flows is the count of flows in the system for the retention period.
- The query to determine the Total Flows is count of flows in last 31 days, assuming the retention period as 31 days.
- Flow Planning is the total flows for which the system can perform security planning.

Recommendation for the Collector Deployment

Table 1-6. Specifications for Collector Brick Size

Brick Size	Cores required for 2.1 GHz CPU	Cores required for 2.3 GHz CPU	Cores required for 2.6 GHz CPU	RAM	Disk
Medium	5	5	4	12 GB	200 GB
Large	10	9	8	16 GB	200 GB
Extra Large	10	9	8	24 GB	200 GB

Note The reservation for the CPU speed and RAM for each node must be 100% of the value specified above.

Table 1-7. Collector Deployment - Maximum Capacity

Collector Size	Number of VMs	Flows per Day	Flow count in 4 days
Medium	4K	2.5M	3.25M
Large	10K	5M	6.5M
Extra Large *	20K	10M	13M

* For Extra Large deployment, you must contact VMware Customer Support.

Table 1-8. Collector Deployment - Maximum Capacity for VMware SD-WAN

Collector Size	Number of Edges	Flows per Day	Flow count in 4 days
Medium	4K	2.5M	3.25M
Large	6K	5M	6.5M
Extra Large *	10K	10M	13M

* For Extra Large deployment, you must contact VMware Customer Support.

Note

- The count of VMs includes the templates on the vCenter as well.
- For a single deployment with more than one collector, the limitation on the total flows across collectors is based on the capacity of the platform.

Other Requirements and Considerations

- The maximum time skew between the platform nodes has to be lesser than 30 seconds.
- The availability of the NTP service is critical to system operations. Ensure that you do not reboot the platform node or the collector node when the NTP service is not available.
- When the existing compute resources are completely used by the other processes on the platform, vRealize Network Insight crashes and does not recover automatically. If the services fail to recover, reboot the platform node.

- If the network latency between platform node and upgrade server is greater than 500ms, the vRealize Network Insight upgrade might encounter an error. So, the network latency must be less than 500ms.
- The recommended disk latency for optimal performance is up to 5ms. If the disk latency is greater than 5ms, the system performance degrades.
- The recommended disk IOPS is 7500.

Supported Web Browser

- Google Chrome: The latest two versions.
- Mozilla Firefox: The latest two versions.

Recommendations to Support High Availability

You can customize vSphere HA options to enable vSphere high availability.

- **Host Failure** - Restart VMs
- **Host Isolation**- Disabled
- **Guest not heartbeating**- Disabled

Privileges

Privileges Required for Data Sources

- Privileges required to configure and use IPFIX
 - vCenter Server Credentials with privileges:
 - Distributed Switch: Modify
 - dvPort group: Modify
 - The predefined roles in the vCenter server must have the following privileges assigned at root level that need to be propagated to the children roles:
 - System.Anonymous
 - System.Read
 - System.View
 - global.settings

To know more about roles in vCenter, see Using Roles to Assign Privileges section in *vSphere Security* guide.
- Privileges required for NSX Manager Data Provider
 - NSX Manager Data Provider requires the **Enterprise** role.

- If Central CLI is enabled, then the system admin credentials are required for NSX Manager Data Provider.
- User privileges required on Cisco switches for metrics collection
 - vRealize Network Insight is capable of collecting metric data via SNMP as well as configuration via SSH from Cisco Switches. Cisco Switches UCS platform requires the use of both SSH and API for collection.

Table 1-9.

Type of data	User Privileges
Configuration Data	Read-Only
Metric Data	SNMP read-only
	SNMPv2 read-only SNMP community
	SNMPv3 read-only

System Ports

Following is the list of ports required for the vRealize Network Insight inbound communication:

Ports for the Platform Cluster Setup

Table 1-10.

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
SSH client	Platform	22	SSH	CLI or host access	No	Yes	User/Password or SSH key-based authentication
Client Web-Browser and vRNI Collector	Platform	443	HTTPS	UI/API access and communication with vRNI Collector	Yes	Yes	SSL channel encrypted with 2048b RSA key based SHA2 cert (or User configured custom cert). Collector to Platform messages on this channel also encrypted further with HMAC.

Table 1-10. (continued)

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
Platform	Platform	2181	HTTP	Communication between zookeeper servers on other nodes (in case of cluster). And stores metadata information (znode data)	No	No	
Platform	Platform	2888	HTTP	Used to connect to zookeeper leader	No	No	
Platform	Platform	3000	HTTP	Used for email notifications	Yes	No	
Platform	Platform	3888	HTTP	Used for zookeeper leader election	Yes	No	
Platform	Platform	5432	jdbc	Storing VM configuration data and infra meta data	Yes	No	
Platform	Platform	8020	TCP/RPC	Communication between other name node(s) and data nodes	Yes	No	
Platform	Platform	8025	HTTP	Node managers use this port to connect to resource manager	No	No	
Platform	Platform	8030	HTTP	Used by resource manager to schedule the tasks	No	No	

Table 1-10. (continued)

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
Platform	Platform	8032	HTTP	The address of the applications manager interface in the RM	No	No	
Platform	Platform	8033	HTTP	The address of the RM admin interface	No	No	
Platform	Platform	8042	HTTP	Node manager web app address	No	No	
Platform	Platform	8080	HTTP	Serves UI requests	Yes	No	
Platform	Platform	8088	HTTP	The HTTP address of the Resource Manager web application	No	No	
Platform	Platform	8480	TCP/RPC	JournalNode HTTP server	No	No	
Platform	Platform	8485	TCP/RPC	HDFS shared edits data dir	No	No	
Platform	Platform	9090	HTTP	Serves requests from collector and sends commands to collector	Yes	Yes (protected via nginx)	
Platform	Platform	9092	Binary over TCP	Port on which other brokers communicate	Yes	No	

Table 1-10. (continued)

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
Platform	Platform	9200-9300	HTTP	Serves search requests. ES uses range of ports to listen, if 9200 is by it uses next port available.	Yes	No	
Platform	Platform	9300	HTTP	Serves search requests. ES uses range of ports to listen, if 9200 is by it uses next port available.	Yes	No	
Platform	Platform	30000:65535	TCP	Ephemeral ports range used by various processes to make the TCP connection with the other processes	No	No	
Platform	Platform	60000	IPC	Used for communication between other hbase primary and region servers	Yes	No	
Platform	Platform	60010	HTTP	Used for hbase web UI	No	No	

Table 1-10. (continued)

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
Platform	Platform	60020	IPC	Communication between hbase primary and region server	Yes	No	
Platform	Platform	4500-4510	TCP	Communication between Foundation DB servers running on different platforms	Yes	No	

Ports for the Single Platform Setup

Table 1-11.

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
SSH client	Platform	22	SSH	CLI or host access	No	Yes	User/Password or SSH key-based authentication
Client Web-Browser and vRNI Collector	Platform	443	HTTPS	UI/API access and communication with vRNI Collector	Yes	Yes	SSL channel encrypted with 2048b RSA key based SHA2 cert (or User configured custom cert). Collector to Platform messages on this channel also encrypted further with HMAC.

Ports for the Collector Server

Table 1-12.

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
SSH client	Collector	22	SSH	CLI or host access	No	Yes	User/Password or SSH key-based authentication
vRNI Collector	Platform	443	HTTPS	Primary communication channel with Platform	Yes	Yes	SSL channel encrypted with 2048b RSA key based SHA2 cert (or User configured custom cert). Collector to Platform messages on this channel also encrypted further with HMAC.
Flow Forwarder	Collector	UDP 2055	NetFlow/IPFIX	Flows from target are pushed to this port	Yes	No	
Flow Forwarder	Collector	UDP 6343	sFlow	Flows from target are pushed to this port	Yes	No	

Table 1-12. (continued)

Source	Target	Port	Protocol	Purpose	Sensitive	SSL	Authentication
ESXi Host	Collector	1991	TCP	Collecting latency measurement of virtual infrastructure, for example: latency between vNIC to pNIC, VTEP to VTEP, TEP to TEP, and so on.	No	No	
Dell OS10	Collector	50000	GRPC	Receiving buffer stats telemetry information from Dell OS10 devices	No	No	

Network Communication Ports

The following table lists the ports and the protocols that are used for the network communication in vRealize Network Insight.

You can also see the list of ports at <https://ports.vmware.com/home/vRealize-Network-Insight>.

Table 1-13.

Purpose	From	To	Port	Protocol
Communication between the VMs of vRealize Network Insight	Collector	Platform Note The port must be enabled for all platforms.	443	HTTPS
Services that require Internet access	Platform and Collector	svc.ni.vmware.com support2.ni.vmware.com reg.ni.vmware.com	443	HTTPS
Communication for miscellaneous services configured	Platform	LDAP server	389, 636	LDAP and LDAPS
		SNMP server	Configurable	SNMP
	Platform and Collector	DNS server	53	UDP
		Syslog server	Configurable	

Table 1-13. (continued)

Purpose	From	To	Port	Protocol
	ESXi Hosts	Collector	2055	
	ESXi Hosts	Collector	1991	TCP
Communication with AWS as a data source	Collector	AWS(*.amazonaws.com)	443	HTTPS
Communicate with Telemetry service	Browser	Telemetry URL https://vcsa.vmware.com	433	HTTPS
Communication with other data sources within the data center	Collector	Arista switches	161 and 22	SNMP and SSH
		Azure	443	HTTPS
		Brocade switches	161 and 22	SNMP and SSH
		Check Point firewall	443	HTTPS
		Cisco Nexus	161 and 22	SNMP and SSH
		Cisco UCS (Unified Computing System)	161, 22, and 443	SNMP, SSH, and HTTPS
		Cisco Catalyst switches	161 and 22	SNMP and SSH
		Cisco ACI Switches	161	SNMP
		Cisco APIC Controller	161 and 443	HTTPS and SNMP
		Dell switches	161 and 22	SNMP and SSH
		Dell OS10	50000	TCP
		VeloCloud	443, 2055	HTTPS
		HP	22	SSH
		Juniper Switches	161 and 22	SNMP and SSH
		Palo Alto Networks	443	HTTPS
		VMware vSphere	443	HTTPS
		VMware NSX - V (All Component)	22 and 443	SSH and HTTPS
		NSX-T Manager	443	TCP
		VMware PKS API Server	8443 and 9021	TCP
		Kubernetes API Server	8443	TCP
		vRealize Log Insight	443	HTTPS
		Fortinet FortiManager	443	HTTPS

Supported Products and Versions

vRealize Network Insight supports several products and versions.

Data Source	Version/Model	Connection Protocol	Permissions/Privileges
Amazon Web Services (Enterprise License Only)	Not Applicable	HTTPS	See the Add Datasources section in the User Guide.
Arista switches	7050TX, 7250QX, 7050QX-32S, 7280SE-72	SSH, SNMP	See the Add Datasources section in the User Guide.
Azure Subscription	Not Applicable	HTTPS	See the Add Datasources section in the User Guide.
Brocade Switches	VDX 6740, VDX 6940, MLX, MLXe	SSH, SNMP	See the Add Datasources section in the User Guide.
Check Point Firewall	Check Point R80 , R80.10	HTTPS, SSH	See the Add Datasources section in the User Guide.
Cisco ACI	3.2	HTTPS (to APIC controller) SNMP (to APIC controller and ACI switches)	See the Add Datasources section in the User Guide.
Cisco ASA	X Series with OS 9.4	SSH, SNMP	See the Add Datasources section in the User Guide.
Cisco Catalyst	3000, 3750, 4500, 6000, 6500	SSH, SNMP	See the Add Datasources section in the User Guide.
Cisco Nexus	3000, 5000, 6000, 7000, 9000	SSH, SNMP	Read only user Read only SNMP user
Cisco UCS (Unified Computing System)	Series B blade servers, Series C rack servers, Chassis, Fabric interconnect	UCS Manager: HTTPS UCS Fabric: SSH, SNMP	Read only user Read only SNMP user
Dell switches	FORCE10 MXL 10, FORCE10 S6000, S4048, Z9100, S4810, PowerConnect 8024, Dell OS10	SSH, SNMP	Read only user Read only SNMP user
Fortinet FortiManager	6.0.1	HTTPS	The user must have: <ul style="list-style-type: none"> ■ at least the Restricted User role with access to all ADOMs and policy packages. ■ the rpc-permit read access enabled from Command Line Interface (CLI).
F5 BIG - IP	12.1.2 and later	HTTPS, SSH, SNMP	The user must have at least the guest role. Also, TMSH must be enabled and must have access to all partitions. F5 BIG-IP supports both routing and load balancing.

Data Source	Version/Model	Connection Protocol	Permissions/Privileges
HP	HP Virtual Connect Manager 4.41, HP OneView 3.0	HP OneView 3.0: HTTPS HP Virtual Connect Manager 4.41: SSH	Read only user
Huawei Cloud Engine	6800, 7800, 8800	SSH, SNMP	Read only user Read only SNMP user
Infoblox	Infoblox NIOS version 8.0, 8.1, 8.2	HTTPS	Read only user with API Interface access Read-only permissions for DNS object types as follows: <ul style="list-style-type: none"> ■ Permission Type - DNS ■ Resource - A Records, DNS Zones, DNS Views
Juniper Switches	EX3300, QFX 51xx Series (JunOS v12 & v15, without QFabric)	Netconf, SSH, SNMP	Read only user Read only SNMP user
Kubernetes	<ul style="list-style-type: none"> ■ 1.12 on NSX-T 2.3.1 ■ 1.12 on NSX-T 2.3.2 ■ 1.13 on NSX-T 2.3.2 	HTTPS	User must have cluster admin role with read permissions.
Palo Alto Networks	Panorama 7.0.x, 7.1, 8.x, 9.0	HTTPS	User must have admin role with XML API access. For details, see the Palo Alto Networks section in the <i>vRealize Network Insight User Guide</i> .
ServiceNow	London	HTTPS	User must have admin role
VMware SD-WAN	VeloCloud Orchestrator and Edge Version 3.3.1 and later	HTTPS	User must have Account Role with any of the following permission: <ul style="list-style-type: none"> ■ Superuser ■ Standard Admin ■ Customer Support
VMC on AWS - vCenter	M5P2 and above Note Only NSX-T based VMware Cloud on AWS SDDCs are supported.	HTTPS	User must have the following permission: <ul style="list-style-type: none"> ■ Cloud Administrator: To add data source and enable IPFIX.

Data Source	Version/Model	Connection Protocol	Permissions/Privileges
VMC on AWS - NSX Manager	M5P2 and above Note Only NSX-T based VMware Cloud on AWS SDDCs are supported.	HTTPS	User must have any of the following permission: ■ Org Member.Administrator: To add data source and enable IPFIX. ■ Org Member.Cloud Admin: To add data source and enable IPFIX. ■ Org Member.VMware Cloud on AWS (all role): To add data source and enable IPFIX. ■ Org Member.Cloud Auditor: To add data source.
VMware Identity Manager	3.3 and later	HTTPS	User must have admin role.
VMware PKS	PKS 1.3.2 on NSX-T 2.3.1 PKS 1.3.2 on NSX-T 2.3.2 PKS 1.5.0 on NSX-T 2.5.0 PKS 1.6.0 on NSX-T 2.5.1		User must have cluster admin role with read permissions.
VMware NSX Manager (VMware NSX-V)	Supported Versions	SSH, HTTPS	See the Edge Data Collection section in the <i>vRealize Network Insight User Guide</i> .
VMware NSX-T Manager	2.4. For additional supported version, see Supported Versions	HTTPS	Read only user
VMware vRealize Log Insight	Supported Versions	HTTPS	API user with permissions to install, configure, and manage the content pack
VMware vSphere	Supported Versions For IPFIX, VMware ESXi version needed: ■ 5.5 Update 2 (Build 2068190) and above ■ 6.0 Update 1b (Build 3380124) and above ■ VMware VDS 5.5 and above Note VMware tools should be installed on all the VMs in the data center to identify the VM to VM path.	HTTPS	Read only user Privileges required to configure and use IPFIX vCenter Server Credentials with privileges: Distributed Switch: Modify dvPort group: Modify The predefined roles in the vCenter server must have the following privileges assigned at root level that need to be propagated to the children roles: System.Anonymous System.Read System.View global.settings

Installing vRealize Network Insight

2

You can deploy vRealize Network Insight using vSphere Web client or vSphere Windows native client.

Note After you successfully deploy vRealize Network Insight Platform OVA, verify whether the given static IP is set on vCenter Server.

To automate installation, configuration, upgrade, patch, configuration management, drift remediation and health from within a single pane of glass, you can use vRealize Suite Lifecycle Manager. If you are a new user, click here to install [vRealize Suite Lifecycle Manager](#). This provides the IT Managers of Cloud admin resources to focus on business-critical initiatives, while improving time to value (TTV), reliability and consistency.

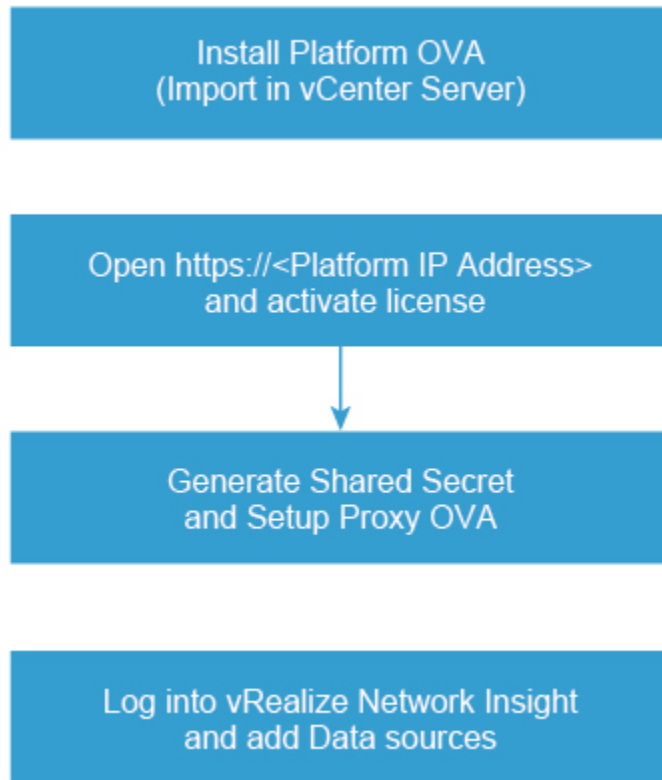
You can also install and upgrade vRealize Network Insight by using vRealize Suite Lifecycle Manager. For more information, see the [vRealize Suite Lifecycle Manager Installation, Upgrade, and Management Guide](#).

This chapter includes the following topics:

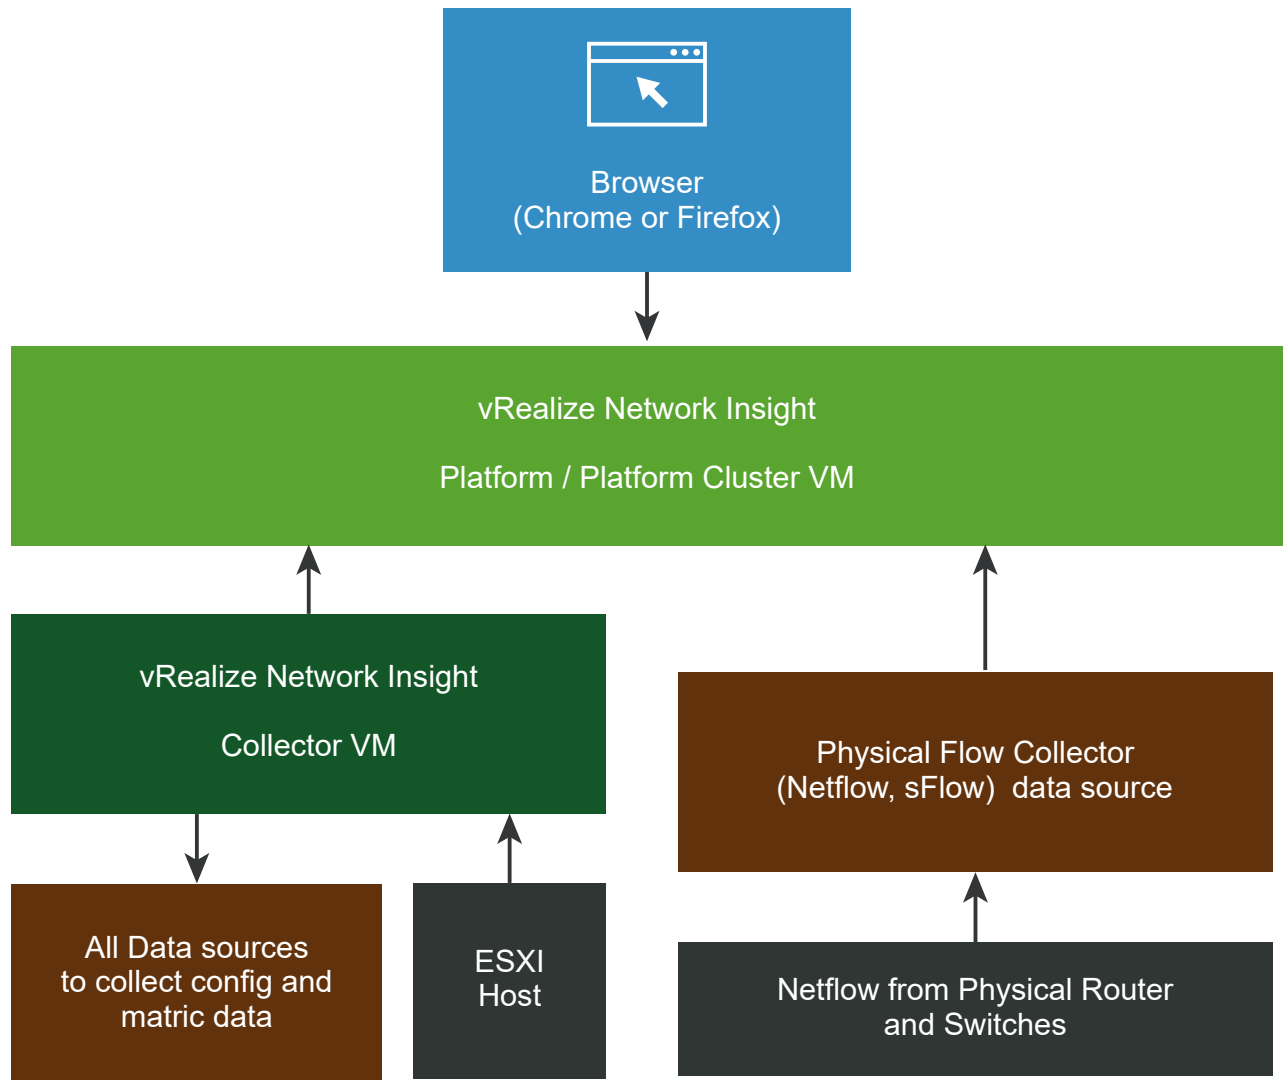
- [Installation Workflow](#)
- [Deploying vRealize Network Insight Platform OVA](#)
- [Activating the License](#)
- [Generating Shared Secret](#)
- [Setting up Network Insight Collector \(OVA\)](#)
- [Setting up Network Insight Collector \(AMI\) in AWS for VMware SD-WAN](#)
- [Deploy Additional Collector to an Existing Setup](#)

Installation Workflow

To install vRealize Network Insight, you install the platform OVA, activate the license, generate a shared secret, and setup collector OVA.



A simplified deployment diagram of vRealize Network Insight is as follows:



Deploying vRealize Network Insight Platform OVA

You can import the vRealize Network Insight Platform OVA to your vCenter Server.

Note Deployment of vRealize Network Insight Platform OVA on the VMC SDDC is not supported.

Deployment using vSphere Web Client

You can deploy vRealize Network Insight using vSphere Web Client.

Procedure

- 1 Right-click the **Datacenter** where you want to install the appliance and select **Deploy OVF Template**.

- 2 Enter the URL to download and install the OVA package or browse to select the source location of the OVA package.
- 3 Enter the OVA name. Select the destination folder for deployment.
- 4 Select a host or a cluster or a resource pool where you want to run the deployed template.
- 5 Verify the OVF template details.
- 6 Read the End User License Agreement and click **Accept**.
- 7 Select a deployment configuration. Click **Next**.
- 8 Select the location to store the files for the deployed template. Select **Thin Provision** as the Virtual Disk format. Select the datastore or the datastore clusters where you want to store the files. Click **Next**.
- 9 Select the network that the deployed VM will use.

The selected network should allow the appliance to reach out to Internet for support and upgrade.
- 10 To customize the template for the deployment, you will have to manually configure the appliance using the VM console. Click **Next**.
- 11 Verify the configuration details and click **Finish**.
- 12 [Increase the Brick Size of your Setup](#) to match the [System Recommendations and Requirements](#).
- 13 Once the platform is installed, start the VM and launch the console.
- 14 Log in with the console credential that you see on the screen and run the setup command.
- 15 Create the password for the *support* login and change the password for the *consoleuser*.

Note

- Your password must contain a minimum of 6 characters. A single quote (') is not allowed.
 - You must change the *support* and *consoleuser* password periodically to comply with your organization policy.
-

- 16 Enter the following details to configure the network:
 - a **IPv4 Address**: Second reserved static IP address
 - b **Netmask**: Subnet mask for the above static IP
 - c **Default Gateway**: Default gateway of your network
 - d **DNS** : DNS server of your environment

Note For multiple DNS servers, ensure that they are separated by space.

- e **Domain Search List** : The domain that needs to be appended for dns lookups
- f Enter y to save the configuration.

- 17 Enter the NTP Server and ensure that it can be reached from the VM. The services will fail to start if NTP time is out of sync.

Note For multiple NTP servers, ensure that they are separated by commas.

- 18 (Optional) To configure Web Proxy, enter y.
- 19 All the services are verified.
- 20 Add additional disk space based on your setup requirement. See <https://kb.vmware.com/s/article/53550>.

Deployment Using vSphere Windows Native Client

You can deploy vRealize Network Insight using vSphere Windows native client.

Procedure

- 1 Click **File > Deploy OVF Template**.
- 2 Enter the URL to download and install the OVA package from the internet or browse to select the source location of the OVA package on your computer.
- 3 Click **Next** and verify the OVF template details.
- 4 Read the End-User License Agreement and click **Accept**.
- 5 Provide a name and specify the location for the deployed template. Click **Next**.
- 6 Select the **Deployment Configuration**.
- 7 Select a **Host/Cluster** where you want to run the deployed template.
- 8 Select the **Resource Pool** in which you want to deploy this template.
- 9 Select a destination storage for the VM files. Click **Next**.
- 10 Specify the format in which you want to store the virtual disks. Select **Thin Provision** as the virtual disk format. Click **Next**.
- 11 Specify the network that the deployed template should use. Map the network from OVA to your inventory.
- 12 Customize the template for the deployment. Provide the shared secret that was generated on the onboarding page. You will have to manually configure the appliance using the VM console. Click **Next**.
- 13 Verify all the configuration data. Check **Power on after deployment**. Click **Finish**.
- 14 [Increase the Brick Size of your Setup](#) to match the [System Recommendations and Requirements](#).
- 15 Once the Collector OVA is installed, start the VM and launch the console.
- 16 Log in with the console credential that you see on the screen and run the setup command.

- 17 Create the password for the *support* login and change the password for the *consoleuser*.

Note

- Your password must contain a minimum of 6 characters. A single quote (') is not allowed.
 - You must change the *support* and *consoleuser* password periodically to comply with your organization policy.
-

- 18 Enter the following details to configure the network:

- a **IPv4 Address:** Second reserved static IP address
- b **Netmask:** Subnet mask for the above static IP
- c **Default Gateway:** Default gateway of your network
- d **DNS :** DNS server of your environment

Note For multiple DNS servers, ensure that they are separated by space.

- e **Domain Search List :** The domain that needs to be appended for dns lookup.
- f Enter y to save the configuration.

- 19 Enter the NTP Sever and ensure that it can reached from the VM. The services will fail to start if NTP time is out of sync.

Note For multiple NTP servers, ensure that they are separated by commas.

- 20 (Optional) To configure Web Proxy, enter y.

- 21 All the services are verified.

- 22 Add additional disk space based on your setup requirement. See <https://kb.vmware.com/s/article/53550>.

Activating the License

After installing the vRealize Network Insight Platform OVA, open *https://<vRealize Network Insight Platform IP address>* in the Chrome Web browser.

Procedure

- 1 Enter the license key received in the welcome email.
- 2 For UI admin (admin@local) user name, set the password.

Note Your password must be alphanumeric, with a minimum of 8 characters and a maximum of 100 characters. Space between the characters is not allowed.

- 3 Click **Activate**.
- 4 Add the vRealize Network Insight Collector after activating the license.

Generating Shared Secret

You can generate and import the vRealize Network Insight collector virtual appliance.

Generate a shared secret and import the vRealize Network Insight collector virtual appliance:

Procedure

- 1 Generate a shared secret after activating the license on the **Setup Proxy Virtual Appliance** page.
- 2 Copy the shared secret.

You will require this during the deployment of vRealize Network Insight Collector OVA.

Setting up Network Insight Collector (OVA)

You can set up vRealize Network Insight collector by importing OVA to your vCenter server.

Follow the steps below to import the vRealize Network Insight collector OVA to your vCenter Server.

Deployment Using vSphere Web Client

You can import the vRealize Network Insight Collector OVA using vSphere Web Client.

Procedure

- 1 Right-click the **Datacenter** where you want to install the appliance and select **Deploy OVF Template**.
- 2 Enter the URL to download and install the OVA package from the internet or browse to select the source location of OVA from your computer.
- 3 Provide a name and specify the location for the deployed template. Click **Next**.
- 4 Select a resource (host or a cluster) where you want to run the deployed template. Click **Next**.
- 5 Verify all the details of the template. Click **Next**.
- 6 Read the End-User License Agreement and click **Accept**. Click **Next**.
- 7 Select a deployment configuration. Click **Next**.
- 8 Select the location where you want to store the files for the deployed template. Specify the format in which you want to store the virtual disks. Select **Thin Provision** as the virtual disk format. Select the Datastore in which you want to install the files. Click **Next**.
- 9 Specify the destination network for the source network. Click **Next**.
- 10 Customize the template for the deployment. Provide the shared secret that was generated from the UI. You will have to manually configure the appliance using the VM console. Click **Next**.

- 11 Verify all the configuration data. Click **Finish**.
- 12 Once the Collector OVA is installed, start the VM and launch the console.
- 13 Log in with the console credential that you see on the screen and run the `setup` command.
- 14 Create the password for the *support* login and change the password for the *consoleuser*.

Note

- Your password must contain a minimum of 6 characters. A single quote (') is not allowed.
 - You must change the *support* and *consoleuser* password periodically to comply with your organization policy.
-

- 15 Enter the following details to configure the network:

- a **IPv4 Address:** Second reserved static IP address
- b **Netmask:** Subnet mask for the above static IP
- c **Default Gateway:** Default gateway of your network
- d **DNS :** DNS server of your environment

Note For multiple DNS servers, ensure that they are separated by space.

- e **Domain Search List :** The domain that needs to be appended for dns lookups
- f Enter `y` to save the configuration.

- 16 Enter the NTP Sever and ensure that it can reached from the VM. The services will fail to start if NTP time is out of sync.

Note For multiple NTP servers, ensure that they are separated by commas.

- 17 (Optional) To configure Web Proxy, enter `y`.
- 18 A check is made to see if the shared secret key has been configured. The collector is paired with the corresponding platform. This may take few minutes.
- 19 All the services are verified.
- 20 Click **Finish**, once **Proxy Detected!** message is displayed on the onboarding page. It will redirect to the Login Page.

Deployment using vSphere Windows Native Client

You can import the vRealize Network Insight Collector OVA using vSphere Windows native client.

Procedure

- 1 Click **File > Deploy OVF Template**.
- 2 Enter the URL to download and install the OVA package from the internet or browse to select the source location of the OVA package on your computer.

- 3 Verify the OVF template details. Click **Next**.
- 4 Read the End-User License Agreement and click **Accept**. Click **Next**.
- 5 Provide a name and specify the location for the deployed template. Click **Next**.
- 6 Select a **Deployment Configuration**. Click **Next**.
- 7 Select a **Host/Cluster** where you want to run the deployed template. Click **Next**.
- 8 Select the **Resource Pool** in which you want to deploy this template. Click **Next**.
- 9 Select a destination storage for the VM files. Click **Next**.
- 10 Specify the format in which you want to store the virtual disks. Select **Thin Provision** as the virtual disk format. Click **Next**.
- 11 Specify the network that the deployed template should use. Map the network from OVA to your inventory.
- 12 Customize the template for the deployment. Provide the shared secret that was generated on the onboarding page. You will have to manually configure the appliance using the VM console. Click **Next**.
- 13 Verify all the configuration data. Check **Power on after deployment**. Click **Finish**.
- 14 Once the Collector OVA is installed, start the VM and launch the console.
- 15 Log in with the given console credentials. Run the setup command.
- 16 Create the password for the support login. Change the password for the consoleuser.
- 17 Enter the following details to configure the network:
 - a **IPv4 Address**: Second reserved static IP address
 - b **Netmask**: Subnet mask for the above static IP
 - c **Default Gateway**: Default gateway of your network
 - d **DNS** : DNS server of your environment

Note For multiple DNS servers, ensure that they are separated by space.

 - e **Domain Search List** : The domain that needs to be appended for dns lookup.
 - f Enter y to save the configuration.
- 18 Enter the NTP Sever and ensure that it can reached from the VM. The services will fail to start if NTP time is out of sync.

Note For multiple NTP servers, ensure that they are separated by commas.

- 19 (Optional) To configure Web Proxy, enter y.
- 20 A check is made to see if the shared secret key has been configured. The collector is paired with the corresponding platform. This may take few minutes.

21 All the services are verified.

22 Click **Finish**, once **Proxy Detected!** message is displayed on the onboarding page. It will redirect to the Login Page.

Setting up Network Insight Collector (AMI) in AWS for VMware SD-WAN

You can set up vRealize Network Insight collector for AWS by importing Amazon Machine Image (AMI) to your AWS environment.

If your environment does not have a vCenter server, and you want to deploy your collector in a cloud environment then you can deploy your collector in AWS.

Note Currently, vRealize Network Insight supports the collector deployment in AWS using AMI only for VMware SD-WAN.

The procedure and task related to EC2 instances are documented in <https://docs.aws.amazon.com/efs/index.html>.

Procedure

- 1 Launch your EC2 instance using the VMware provided AMI in the Amazon EC2 console. For procedure details, see Create Your EC2 Resources and Launch Your EC2 Instance topic in the *Amazon Elastic File System* documentation.

Note When you Launch your EC2 instance in AWS, you must select the following:

Option	Action
Instance type	m4.xlarge (MEDIUM BRICK)
Network	Select an appropriate network and subnet.
Storage	Default Storage.
Tags	As per customer Policies.
Security Group	Allow Outbound to 0.0.0.0/0 for port 443 (or for restricted rules, allow outbound for NI SaaS Prod FQDN for port 443).
Key	Select appropriate Key (SSH Login is enabled for the AMI).

- 2 When your EC2 instance is in the running state, log in to your EC2 instance.
- 3 Log in with the given console credentials. Run the setup command.

- 4 Create the password for the support login. Change the password for the consoleuser.

Note After you change the password, the network options will be skipped during setup CLI.

Proxy AMI does not support the following:

- IP change
- IPv6
- Web Proxy Configuration.

- 5 Enter the NTP Server and ensure that it can be reached from the VM. The services fail to start if the NTP time is out of sync.

Note For multiple NTP servers, ensure that they are separated by commas.

- 6 A check is made to see if the shared secret key has been configured. The collector is paired with the corresponding platform. This process can take few minutes.
- 7 All the services are verified.

What to do next

Enable the flow collection from Edges to the collector you deployed in AWS. To enable the flow collection, do the following:

- Make the collector you deployed in AWS as a Non-VeloCloud Site. For details, contact VMware support.

Deploy Additional Collector to an Existing Setup

You can add additional vRealize Network Insight collector to an existing setup.

Procedure

- 1 Log into the vRealize Network Insight UI. Navigate to **Settings > Install and Support**.
- 2 Click **Add Proxy VM**.
- 3 Copy the shared secret from the dialog that is displayed.
- 4 Follow the steps in section [Setting up Network Insight Collector \(OVA\)](#) in step 3.

Accessing vRealize Network Insight by using the Evaluation License

vRealize Network Insight starts in the NSX assessment mode when you use the evaluation license.

You can add a data source to vRealize Network Insight, analyze traffic flow, and generate reports.

Note To switch to the Full Product mode, click Switch to Full Product Evaluation located in the bottom right corner.

This chapter includes the following topics:

- [Add vCenter Server](#)
- [Analyze Traffic Flows](#)
- [Generate a Report](#)

Add vCenter Server

You can add vCenter Servers as data source to vRealize Network Insight.

Multiple vCenter Servers can be added to vRealize Network Insight to start monitoring data.

Prerequisites

- The predefined roles in the vCenter server must have the following privileges assigned at root level that need to be propagated to the children roles:
 - **System.Anonymous**
 - **System.Read**
 - **System.View**
 - **Global.Settings**
- Following vCenter Server privileges are required to configure and use IPFIX:
 - **Distributed Switch: Modify**
 - **dvPort group: Modify**

To know more about roles in vCenter, see Using Roles to Assign Privileges section in vSphere Security guide.

Procedure

- 1 Click **Add vCenter**.
- 2 Click **Add new source** and customize the options.

Option	Action
Collector VM	Select a collector VM from the drop-down menu.
IP Address/FQDN	Enter the IP address or fully qualified domain name of the vCenter Server.
Username	Enter the user name with the following privileges: <ul style="list-style-type: none"> ■ Distributed Switch: Modify ■ dvPort group: Modify
Password	Enter the password for vRealize Network Insight software to access the vCenter Server system.

- 3 Click **Validate**.

If the number of VMs discovered exceeds the capacity of the platform or a collector node or both, the validation fails. You will not be allowed to add a data source until you increase the brick size of the platform or create a cluster.

The specified capacity for each brick size with and without flows is as follows:

Brick Size	VMs	State of Flows
Large	6k	Enabled
Large	10k	Disabled
Medium	3k	Enabled
Medium	6k	Disabled

- 4 Select **Enable Netflow (IPFIX) on this vCenter** to enable IPFIX.

For more information on IPFIX, see the Enabling IPFIX configuration on VDS and DVPG section.

- 5 Add advanced data collection sources to your vCenter Server system.
- 6 Click **Submit** to add the vCenter Server system. The vCenter Server systems appear on the homepage.

Analyze Traffic Flows

You can use vRealize Network Insight to analyze flows in your datacenter.

Prerequisites

At least two hours of data collection must occur before starting the flow analysis.

Procedure

- 1 Specify the scope of the analysis. For example, if you are interested in flows of all virtual machines in a **Cluster**, select Cluster from the dropdown menu. You can alternately select all virtual machines connected to a VLAN or VXLAN.
- 2 Select the entity name for which you want to analyze the flows.
- 3 Select the duration and click **Analyze**.

Generate a Report

You can generate a report of the flow assessment.

Prerequisites

Analyze traffic flows in the datacenter. For comprehensive reports, collect 24 hours of data before the analysis.

Procedure

- 1 In the **EVAL NSX Assessment Mode**, click **Generate Report** in the Analyze Flows page.
- 2 In the **Non EVAL Mode**, on the **Microsegmentation** page, click **Traffic Distribution > More Options > Assessment Report**.

Planning to Scale up your Deployment

4

If the VM count or the number of active flows in your setup are high or expected to grow, you can increase the size of the platform or collector.

This chapter includes the following topics:

- [Planning to Scale up the Platform Cluster](#)
- [Planning to Scale up the Collector](#)
- [Increase the Brick Size of your Setup](#)

Planning to Scale up the Platform Cluster

You can scale up the platform cluster to meet the increasing load. Based on the load, you can either scale up by increasing the brick size or creating or expanding a platform cluster. Three LARGE platform bricks can be connected together to form a platform cluster. If a platform is of LARGE or EXTRA LARGE brick size, then you have to scale up by creating a platform cluster.

To decide platform brick size and number of platform bricks, see [System Recommendations and Requirements](#).

Note The platform cluster does not support the high availability configuration. All the platform nodes need to be up and running for the cluster to work at optimal performance levels.

Scaling up Scenarios for the Platform Cluster

- Scenario 1 : Your platform is running 5000 VMs and 1.5 million active flows
Convert your platform MEDIUM to LARGE. See [Increase the Brick Size of your Setup](#).
- Scenario 2 : Your platform is running a single LARGE node with 9000 VMs and 2 million active flows
Add two more LARGE brick nodes to convert into 3-node LARGE brick cluster. See *Expand Clusters in the vRealize Network Insight User Guide*.
- Scenario 3 : Your platform is running a 3-node LARGE cluster with one or more collectors, 15000 VMs and 4 million active flows.

Convert your existing platform nodes from LARGE to EXTRA-LARGE. See [Increase the Brick Size of your Setup](#).

- Scenario 4 : Your platform is running a 3-node EXTRA-LARGE cluster with one or more collectors, 25000 VMs and 8 million active flows.

Add two more EXTRA-LARGE brick nodes to convert into 5-node Extra-LARGE cluster. See *Expand Clusters in the vRealize Network Insight User Guide*.

Planning to Scale up the Collector

The collector capacity is based on the brick size. The datasource that you can add to a collector is depended on the capacity of the collector (VMs and flows).

See [Table 1-7. Collector Deployment - Maximum Capacity](#). After a collector is of LARGE brick size, you have to add more collectors. You can scale up each collector to EXTRA-LARGE size.

You can add multiple datasources to a collector based on the supported collector capacity. However, you cannot add same datasource to multiple collectors.

Scaling up Scenarios for the Collectors

- Scenario1: 2000 VMs in a vCenter.

Install one medium collector VM. Add the vCenter to this collector. See [Add vCenter Server](#).

- Scenario 2: 1000 VMs in vCenter1 and 2000 VMs in vCenter2 (all of them are in one data center)

Install one medium collector VM. Add both vCenters to this collector. See [Add vCenter Server](#).

- Scenario 3: 1000 VMs in vCenter1 (data center1) and 2000 VMs in vCenter2 (data center2)

Install one medium collector VM in each data center. Add vCenter1 to a collector VM in same data center and Add vCenter2 to a collector VM in its data center. See [Add vCenter Server](#).

- Scenario 4: VM count exceeds 4000, active flows exceeds 2.5 Million.

Convert your collector VM from MEDIUM to LARGE. See [Increase the Brick Size of your Setup](#).

- Scenario 5: 9,000 VMs in vCenter1 without flows (data center1).

Install one large collector VM. Add this vCenter to the collector. See [Add vCenter Server](#).

- Scenario 6: VM count is less than or equal to 10000, but the active flow exceeds 5 million.

Convert your collector VM from LARGE to EXTRA-LARGE. See [Increase the Brick Size of your Setup](#).

- Scenario 8: Two vCenters, vCenter1 has 10000 VMs and 9 million active flows, and vCenter2 has 10000 VMs and 4 million active flows.

Install one EXTRA-LARGE and one LARGE proxies. Add vCenter1 to EXTRA-LARGE proxy and add vCenter2 to LARGE proxy.

- Scenario 9: One vCenter that runs 10000 VMs and 9 million active flows.

Install one EXTRA-LARGE proxy and add the vCenter to the proxy.

Increase the Brick Size of your Setup

To match your requirements, you can change the brick size of your platform or the collector appliance from MEDIUM to LARGE or LARGE to EXTRA-LARGE.

Procedure

- ◆ Perform the steps that are relevant to your setup.

Option	Description
For a single node platform or fresh independent OVA	<ol style="list-style-type: none"> Log in to vCenter. Shutdown the platform VM. Increase the disk, RAM, total vCPU and corresponding reservation of the VM to match the target brick size. For more information, see the System Recommendations and Requirements page. Restart Platform VM.
For a cluster platform	<ol style="list-style-type: none"> Log in to vCenter. Shutdown the platform VM in the reverse chronological order. For example: Shut down from Node 3 to Node 1. Increase the disk, RAM, total vCPU and corresponding reservation. For more information, see the System Recommendations and Requirements. Restart Platform VMs in the chronological order. For example: Restart from Node 1 to Node 3.
For a collector	<ol style="list-style-type: none"> Log in to vCenter. Shutdown the collector VM. Increase the disk, RAM, total vCPU and corresponding reservation of the VM to match the target brick size. For more information, see the System Recommendations and Requirements page. Restart the collector VM.

Upgrading vRealize Network Insight

5

You can upgrade your current vRealize Network Insight environment to the latest version.

Important points to consider before upgrade:

- After upgrade, vRealize Network Insight takes around 12 to 24 hours to process the data that were in pipeline during the upgrade operation and reflect on the UI.
- vRealize Network Insight does not support rollback or product downgrade. You must take a backup before you proceed to upgrade. For more information about the back up and restore process, see the <https://kb.vmware.com/s/article/55829> KB article.
- In a cluster environment, you must perform the upgrade operation only on the platform 1 node.
- After upgrading to vRealize Network Insight 5.1, some of the firewall rule IDs might change to the new IDs returned by the VMware Cloud on AWS 1.9 API. If any VMware Cloud on AWS 1.8 firewall rules that are attached to the flows exists:
 - the correct or respective VMware Cloud on AWS 1.9 firewall rules would be attached immediately after upgrade for all active flows.
 - the firewall rules would refer to non-existent rules for the flows whose period of inactivity is greater than 24 hours prior to upgrading 1.8 to 1.9 version.

Note If issues such as upload failure or UI failure come up while performing the centralized upgrade, contact VMware support.

Migration to Foundation DB

To distribute configuration data across datastores in the cluster, vRealize Network Insight 5.1 replaces PostgreSQL to Foundation DB for storing the configuration data. This enables vRealize Network Insight to:

- reduce the load on platform 1 node
- avoid single point of failure
- improve resiliency
- enhance performance

- share the disk uniformly across cluster nodes

The migration process automatically,

- shuts down all services
- starts table to table migration from PostgreSQL to Foundation DB
- displays the dynamic migration progress information on the platform 1 UI

The migration time to move data from PostgreSQL to Foundation DB depends on the disk speed and the node count (more nodes provides more Foundation DB write throughput)

The time taken to complete the migration process depends on the size of the database.

Setup Size	Data Size	Node Count	Typical Migration Time
Small	20GB to 40GB	1 node	1 to 2 hours
Medium	60GB to 100GB	3 nodes	7 to 10 hours
Large One cloud setups	500GB	10 nodes cluster	15 to 20 hours
XL (Megatron)	1 TB	10 node cluster	35 to 40 hours

Note that the migration happens as a part of vRealize Network Insight upgrade process. So, the upgrade time might be longer, which you see on the screen during the process.

vRealize Network Insight provides the different modes of upgrade.

This chapter includes the following topics:

- [Online Upgrade](#)
- [Single-Click Offline Upgrade](#)
- [CLI Upgrade](#)

Online Upgrade

Whenever there is a new version of vRealize Network Insight available, you receive a notification.

Prerequisites

- The upgrade steps might fail if there is insufficient space in the /tmp directory. Verify you meet the following disk space requirements for platform and collector server:
 - /tmp - 6 GB
 - /home - 2 GB
- Verify you meet the following disk space requirement for platform server:
 - /- 6 GB (Only for the Platform1 node)
 - /var - 40 GB

- Verify you have the minimum bandwidth requirement of 500 KB/s to download the upgrade bundle from the server. The **Install and the Support** page throws an error, if the download bandwidth is not sufficient.
- Ensure that all the nodes are online. If any node is inactive, you will not be allowed to trigger the upgrade.
- Take the snapshots of the VMs.
- Note the following values to verify after migration:
 - Count of VMs
 - VM where Snapshot Count > 0
 - Count of Firewall Rule
 - Count of Security Group
 - Count of NSX Firewall

Procedure

- 1 When an update is available, you see **Update available** message notification.

Note

- If the update notification is not available, verify that both vRealize Network Insight platform and collector VMs have connectivity to `svc.ni.vmware.com` on port 443 and `reg.ni.vmware.com` on port 443 by running the `show-connectivity-status` command. If this connectivity requires `http proxy`, configure it on each VM using the `set-web-proxy` command. Ensure that the output contains upgrade connectivity status as `Passed`.
 - File a support ticket and provide the service tag from the product UI. The service tag is shown under **Settings > About**.
 - Log into the appliance and run the `show-connectivity-status` command. Provide a screenshot of the command output from each vRealize Network Insight platform and collector VMs.
-

- 2 In the Update Available message notification, click **View details** to view details of update. vRealize Network Insight Upgrade screen appears.
- 3 Read the **Before you proceed** instructions and click **Continue**.
- 4 Wait for the pre-checks to complete, which verifies:
 - the disk space, including the space required for migration
 - the version
 - the NTP sync status
 - the bandwidth

You can see the approximate time required to complete the upgrade process (including the migration duration) on your setup.

5 Click **Install Now**.

6 Once the upgrade process begins, the vRealize Network Insight Upgrade screen provides the status of the upgrade process.

Note

- If a node becomes inactive, the upgrade process does not continue. The upgrade will not resume until the node becomes active again.
 - The platform1 becomes the upgrade server. If platform1 is offline, then no other node is upgraded.
 - Once the platforms are upgraded, you can resume your normal vRealize Network Insight operations even though the collector upgrade happens in parallel. Until the upgrade process is completely over, the `Node Version Mismatch detected` message is shown in the Install and Support page.
-

- After the services are upgraded, Nginx restarts to display the migration process. So, you might not be able to access the UI for a short period (one to two minutes).
- vRealize Network Insight starts migrating data to Foundation database. On the Data Migration Status screen, you see:
 - the overall status
 - the elapsed time
 - the table by table status
 - the number of records migrated

For any issues, you can use the **Export Migration logs** option to share with the VMware support team.

- The PostgreSQL data on the collectors are also migrated to Foundation DB as part of upgrade process. However, the collector migration status is not shown on the UI.

7 Upon the completion of upgrade process, you see the confirmation message.

All platforms and the collectors nodes are upgraded.

What to do next

- Log in to vRealize Network Insight and perform your tasks.
- After two or three days, delete the snapshots to save the disk space.

Single-Click Offline Upgrade

vRealize Network Insight supports the single-click offline upgrade of the product from Release 3.7 and later.

Prerequisites

- The upgrade steps might fail if there is insufficient space in the /tmp directory. Verify you meet the following disk space requirements for platform and collector server:
 - /tmp - 6 GB
 - /home - 2 GB
- Verify you meet the following disk space requirement for platform server:
 - /- 12 GB (Only for the Platform1 node)
 - /var - 40 GB

Note The bundle upload and the subsequent upgrade steps may fail if there is insufficient space in the /tmp directory.

- To avoid the UI session timeout, go to **Settings > System Configuration > User Session Timeout** and increase the **User Session Timeout** to at least 2 hours. After you change the session timeout duration, you must log in again to the system.
- Ensure that all the nodes are online. If any node is inactive, you will not be allowed to trigger the upgrade.
- Take the snapshots of the VMs.
- Note the following values to verify after migration:
 - Count of VMs
 - VM where Snapshot Count > 0
 - Count of Firewall Rule
 - Count of Security Group
 - Count of NSX Firewall

Procedure

- 1 Download the required upgrade bundle file from [My VMware](#) and save the update package in your local disk.
- 2 Check and ensure that the MD5SUM value of the downloaded bundle matches the MD5SUM value specified in the VMware website.
- 3 On the **Install and Support** page, under **Software version**, select **Click here**.

4 Click **Browse** to select the file and click **Upload**.

When the upload is complete, vRealize Network Insight show the **Bundle Upload Complete** message notification within 2-3 minutes and the bundle processing happens in the background.

Note

- Until the upload of the package happens, ensure that the session is not closed. If the session ends, you have to restart the upload process.
 - Do not refresh the page after bundle upload, until you see the **Update Available** message notification.
-

5 In the **Update Available** message notification, click **View details**.

vRealize Network Insight Upgrade screen appears.

6 Read the **Before you proceed** instruction and click **Continue**.

7 Wait for the pre-checks to complete, which verifies:

- the disk space, including the space required for migration
- the version
- the NTP sync status
- the bundle

8 Click **Install Now**.

You can see the approximate time required to complete the upgrade process on your setup.

9 Once the upgrade process begins, the vRealize Network Insight Upgrade screen provides the status of the upgrade process.

Note

- If a node becomes inactive, the upgrade process does not continue. The upgrade will not resume until the node becomes active again.
 - The platform1 becomes the upgrade server. If platform1 is offline, then no other node is upgraded.
 - Once the platforms are upgraded, you can resume your normal vRealize Network Insight operations even though the collector upgrade happens in parallel. Until the upgrade process is completely over, the **Node Version Mismatch detected** message is shown in the **Install and Support** page.
-
- After the services are upgraded, Nginx restarts to display the migration process. So, you might not be able to access the UI for a short period (one to two minutes).

- vRealize Network Insight starts migrating data to Foundation database. On the Data Migration Status screen, you see:
 - the overall status
 - the elapsed time
 - the table by table status
 - the number of records migrated

For any issues, you can use the **Export Migration logs** option to share with the VMware support team.

- The PostgreSQL data on the collectors are also migrated to Foundation DB as part of upgrade process. However, the collector migration status is not shown on the UI.

10 Upon the completion of upgrade process, you see the confirmation message.

All platforms and the collectors nodes are upgraded.

What to do next

- Log in to vRealize Network Insight and perform your tasks.
- After two or three days, delete the snapshots to save the disk space.

CLI Upgrade

Consider the CLI upgrade only if both online upgrade or single-click offline upgrade does not work. You must upgrade Platform VMs before Collector VMs. However, you must contact VMware Support before initiating Offline Upgrade using CLI.

In a cluster environment, you must perform the upgrade operation only from Platform 1 (P1) node and the other Platform nodes in the cluster get upgraded automatically. But you must upgrade each Collector individually.

Prerequisites

- The upgrade steps might fail if there is insufficient space in the /tmp directory. Verify you meet the following disk space requirements for platform and collector server:
 - /tmp - 6 GB
 - /home - 2 GB
 - /var - 40 GB
- Ensure that all the nodes are online. If any node is inactive, you will not be allowed to trigger the upgrade.
- Take the snapshots of the VMs.
- Note the following values to verify after migration:
 - Count of VMs

- VM where Snapshot Count > 0
- Count of Firewall Rule
- Count of Security Group
- Count of NSX Firewall

Procedure

- 1 Download the required upgrade bundle file from [My VMware](#).
- 2 Check and ensure that the MD5SUM value of the downloaded bundle matches the MD5SUM value specified in the VMware website.
- 3 Copy the upgrade bundle to vRealize Network Insight Platform 1 VM and all Collector VMs.
 - To copy the file from Linux VM to vRealize Network Insight VM, run command `scp <filename>.upgrade.bundle consoleuser@<IP_Address_vRNI_VM>:~/.`
 - To copy the file from Windows VM to vRealize Network Insight VM, run command `pscp -scp <SOURCE_PATH>\<filename>.upgrade.bundle consoleuser@<IP_Address_vRNI_VM>:~/.`

Note Use the pscp utility from <https://the.earth.li/~sgtatham/putty/latest/w64/pscp.exe>.

- 4 Log in to the vRealize Network Insight Platform 1 through CLI using `consoleuser` and run the following commands:
 - `package-installer copy --host localhost --user consoleuser --path /home/consoleuser/<filename>.upgrade.bundle`
 - `package-installer upgrade --name <filename>.upgrade.bundle`

Note You must perform the Platform upgrade first then start the Collector update.

- 5 Run the `package-installer upgrade` command again after the setup is rebooted as part of OS upgrade.

Important If you get an SSH session timeout error, you must check `/var/log/arkin/centralized_upgrade.log` to know if the reboot has already happened. If the reboot is successful, you must run the `package-installer upgrade` command again.

- 6 Log in to each Collector node through CLI and perform the upgrade using the same commands used for platform upgrade.

Note You can upgrade all the Collectors simultaneously.

- 7 Verify the upgraded version using the `show-version` command.

Uninstall vRealize Network Insight

6

You must uninstall vRealize Network Insight through vSphere Web Client.

Procedure

- 1 If you can access the vRealize Network Insight web portal, do the following:
 - a Log in to the vRealize Network Insight web portal.
 - b Go to **Settings > Accounts and Datasources**.
 - c Turn off and delete all datasources.

Deletion of the vCenter datasource removes IPFIX settings (if configured) on VDS.
Similarly deletion of the NSX Manager datasource removes IPFIX settings from NSX Flow Monitor.
- 2 If you are unable to access the vRealize Network Insight web portal, do the following:
 - a If Netflow (IPFIX) is enabled on vCenter, remove vRealize Network Insight collector IP from VDS/DVPG IPFIX settings. See [Remove Collector IP When Netflow is Enabled in vCenter](#).
 - b If IPFIX is enabled on NSX, remove vRealize Network Insight collector IP Flow Monitoring settings. See [Remove Collector IP When Netflow is Enabled in NSX](#).
 - c If Netflow is configured on physical switches to send Netflow to vRealize Network Insight Netflow Collector, modify the configuration at switches to stop sending NetFlow information.
- 3 If any specific firewall or routing rules are created to allow or route traffic to and from vRealize Network Insight VMs, remove those firewall/routing rules.
- 4 For security reasons, clean up access credentials used to configure data sources in vRealize Network Insight.
- 5 Shutdown and delete all vRealize Network Insight Collectors and Platform VMs.

Remove Collector IP When Netflow is Enabled in vCenter

If Netflow (IPFIX) is enabled in vCenter, use this procedure to remove vRealize Network Insight Collector IP from Virtual Dedicated Server (VDS)/Distributed Virtual Port Group (DVPG) IPFIX settings.

Procedure

- 1 Log in to vSphere Web Client.
- 2 Go to **Home > Networking**.
- 3 In the left pane, select the **VDS** and click **Configure > Edit**.
- 4 In the **Collector IP address** field, remove vRealize Network Insight Collector IP details.
- 5 In the **Collector Port** field, remove the port details.
- 6 Click **OK**.

You must wait around two minutes before you move to the next step.

- 7 Select the DVPG of this VDS and click **Configure > Policies > Edit**.
- 8 In the **Netflow** field, select **Disable** from the drop-down.
- 9 Verify your settings and click **Apply**.

What to do next

Perform the steps again for each VDS and its DVPGs for which IPFIX is enabled to remove vRealize Network Insight Collector IP.

Remove Collector IP When Netflow is Enabled in NSX

If Netflow (IPFIX) is enabled in NSX, use this procedure to remove vRealize Network Insight (vRealize Network Insight) Collector IP flow monitoring settings.

Procedure

- 1 Log in to vSphere Web Client.
- 2 Click **Home > Networking & Security > Tools > Flow Monitoring > Configuration**.
- 3 In the **Global Flow Collection Status**, click **Disable**.
- 4 To disable the flow connection, click **IPFIX**.
- 5 In the **IPFIX** tab, select the **Collector IP** and click **Delete**.
- 6 If there are no more IPs left, then click **Edit** and clear **Enable IPFIX Configuration** check-box.
- 7 Click **Save**.