NetScaler SD-WAN 9.3

Aug 17, 2017

The NetScaler SD-WAN product was formerly called "CloudBridge". Refer to the links below to access CloudBridge documentation.

CloudBridge

CloudBridge 9.0 CloudBridge 8.1 CloudBridge 8.0 CloudBridge 7.4 CloudBridge 7.4 CloudBridge 7.3

Similar to the previous product portfolio, NetScaler SD-WAN is available in three different editions, allowing you to deploy the features you need at each location with easy upgrades, configuration, and monitoring.

Note

All references to the term "CloudBridge" is applicable to the new product term "NetScaler SD-WAN".

For more information about the NetScaler SD-WAN platform editions, see the product datasheet and product portfolio at: https://www.citrix.com/content/dam/citrix/en_us/documents/data-sheet/netscaler-sd-wan-datasheet.pdf and https://www.citrix.com/products/netscaler-sd-wan/platforms.html.

Important

For information about NetScaler SD-WAN WANOP 9.1, 9.2, and 9.3 installation, deployment, and feature configuration, please refer to the CloudBridge 7.4 documentation. The features and procedures for the NetScaler SD-WAN WANOP 9.1, 9.2, and 9.3 are similar to the procedures documented in CloudBridge 7.4 release.

Command Center End of Life Notification:

End of Life process for **Citrix Command Center** tool was initiated on 15-May-2017.

It is recommended that you migrate to the new management tool **NetScaler MAS** for your WAN Optimization deployments at the earliest.

Please refer to the articles below for the Command Center EOL calendar and associated details.

- CTX223806 - Notice of Status Change Announcement for Citrix Command Center Software Version 5.2
- CTX223786 - FAQ: Citrix Command Center - End Of Life

Citrix Command Center tool for NetScaler SD-WAN WANOP edition is supported only till the NetScaler SD-WAN 9.2 appliance software release.

Starting with the NetScaler SD-WAN 9.3 software release, NetScaler MAS will be the management tool for SD-WAN WANOP edition appliances.

There are three NetScaler SD-WAN Editions each with a different set or subset of NetScaler SD-WAN features.

- **NetScaler SD-WAN Enterprise Edition (EE)** – This Edition includes both Standard Edition and WAN Optimization features. Enterprise Edition Integrates WAN virtualization with WAN optimization capabilities to optimize branch and mobile user experience and to achieve fully resilient applications regardless of network quality For release 9.1, Enterprise Edition is available for the 1000-VW and 2000-VW branch hardware appliances, only.
- **NetScaler SD-WAN Standard Edition (VW/SE)** – This Edition includes Standard Edition Virtual WAN features, only. It supports software-defined WAN capability to create a highly reliable network from multiple network links and to ensure that each application takes the best path to achieve the highest application performance.
- **NetScaler SD-WAN Optimization Edition (WANOP)** – This Edition includes WAN Optimization features, only. It supports application acceleration, data reduction and protocol control to optimize applications across the WAN. Optionally, it can include virtual Windows Server to simplify branch infrastructure and mobile PC plug-in capability.

> ## Note
>
> The WANOP Edition and Standard/Enterprise Edition are separate hardware platforms running different software.
>
> In a data center, administrators can deploy one Standard Edition and one WANOP Edition to achieve Enterprise Edition capabilities. In the branch office, administrators can choose to deploy either a Standard Edition or WANOP Edition. Alternatively, the benefits of both Standard and WANOP Editions can be accomplished by deploying a single Enterprise Edition at the branch office.

See the Licensing section, for more information about the license options available for using NetScaler SD-WAN platform editions.

Release Notes

Oct 05, 2017

This release notes describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

For information about the previous release versions, see the NetScaler SD-WAN 9.2 and NetScaler SD-WAN 9.1 documentation.

Fixed Issues

SD-WAN 4100-SE

Issue ID 675715: On a NetScaler SD-WAN 4100-SE appliance, changing Interface settings for 1G interface does not work and causes link to become inactive. For example; changing the speed to 100MB does not work. The interface settings change option is disabled for all 1G ports similar to the 10G ports as it is not supported on the 4100-SE appliance.

SD-WAN 4000-WANOP and 4000-SE

Issue ID 680778: A configuration audit error occurs in two-box mode deployment when a NetScaler SD-WAN 4000-SE appliance with two interface groups is configured with first interface group having bridged pair with two Ethernet interfaces selected, and second interface group is connected to the WANOP appliance. The error occurs when the first interface group is enabled with WCCP listener indicating that multiple Ethernet interfaces cannot be enabled with WCCP. When you revert configuration by disabling WCCP on the first interface group and enabling it on the second interface group, the same configuration audit error is displayed even though only one Ethernet interface is enabled on the interface group.

Issue ID 679121: While upgrading SD-WAN 4000 appliance from old releases to 9.2 release, the SD-WAN GUI appears before the upgrade process is completed. The old image is listed in the GUI.

Issue ID 680825: On a NetScaler SD-WAN 4000 appliance with release version 9.2, the HTTP service does not work for one of the SD-WAN instances and fails to start or restart the https service.

TCP Fragmented traffic

Issue ID 681472: Virtual WAN drops TCP Fragmented traffic when firewall connection tracking is enabled.

NTP Server Time Settings

Issue ID 680987: On NetScaler SD-WAN 2000 appliances, when you change the NTP server settings, the Enterprise Edition appliance time settings sync up with the new NTP server time settings and the correct time zone format is displayed. However, the new NTP server time settings on a WANOP appliance are not synchronized with the new NTP server time settings.

Diagnostic tool

Issue ID 680251: In a NetScaler SD-WAN VPX appliance setup, multiple IPREF client TCP sessions are initiated while server session is still on causing the server to display additional entries even when the client has stopped sending any further traffic.

Rules Group Tab

Issue ID 681562: The Rule group tab in SD-WAN Center report page does not show any data for the configured applications.

DPI- No audit error on disabling DPI

Issue ID 681175: If an application object created with DPI application is associated to a firewall policy template, and is used in firewall and then if the DPI is disabled, there is no audit error message displayed indicating that there are rules still associated with firewall as the firewall is still functional.

SSL Profile Name

Issue ID 681482: In a NetScaler SD-WAN VPX appliance setup, when you create an SSL profile and try to edit the profile and save it, the following error message is displayed: "*No object with profile name exists".*

SSL Profile page

Issue ID 681443: When creating or editing an SSL profile, the settings are saved but the application does not get redirected to the SSL Profile home page.

GUI

Issue ID 681649: Unable to enable DHCP Server and Relay for management from the UI. On selecting Enbale DHCP Server, the fields Lease Time, Domain Name, Start IP Address and End IP Addresss should be editable, but these fields are not editable.

Security Vulnerability

Issue ID 690709: Unauthenticated remote code execution on the NetScaler SD-WAN Enterprise Edition and Standard Edition appliances. This security hotfix addresses the vulnerabilities as described in the CTX security bulletin article (CTX225990).

SD-WAN WANOP

Issue ID 675452: NetScaler SD-WAN WANOP client info displays OS version as Windows 8 even when plugin is installed in Windows 10 OS.

Simplified Configuration

Issue ID 678342: In the SD-WAN configuration editor, secondary level confirmation is not provided when deleting a WAN Link, Interface Group, or Static Route from the **Basics** view.

Ethernet Interfaces Configuration

Issue ID 680585: In a NetScaler SD-WAN Standard Edition appliance web GUI, the Basic View under Configuration Editor allows you to create Interface without selecting Ethernet interfaces. The created interface is displayed in the Advanced View as VLAN 0 instead of displaying in the Basic View.

DPI - Traffic classified as unknown when the traffic flows through EE appliances

Issue ID 677504: Applications are classified as Unknown protocol when the traffic flows through EE appliances,

because the compressed traffic is not classified. Therefore, the Firewall rules do not work on EE appliance with DPI enabled when rules are configured with Application, Application Family or Application Object firewall policies. This issue occurs only when a WANOP Service Class Compression policy is configured on a Standard Edition/Enterprise Edition or Standard Edition/Standard Edition appliance with a WANOP deployment mode.

DPI – Any application traffic sent via GRE Tunnel is reported as GRE in SD-WAN Center

Issue ID 680994: Ideally, any application traffic (example HTTP) sent through the GRE tunnel should be classified by DPI reported as both GRE and the real application traffic (example HTTP) in the Application section of **Reporting** page in SD-WAN Center. Due to this bug, the real application (example HTTP) is also reported as GRE traffic.

This bug is only a reporting issue and the real classification has no issues in the site level DPI. Both the classification and firewall actions after DPI will have no impact in any site.

SD-WAN GUI

Issue ID 683520: In the SD-WAN GUI, changing the interface settings for interface under **Configuration** > **Appliance Settings** > **Network adapters** > **Ethernet** does not work for the SD-WAN 1000-EE, 2000-EE and 400-SE platforms.

XS 6.5 Upgrade Support

Issue ID 662041: Once your appliance is upgrade to NetScaler SD-WAN 9.3 software, you can also upgrade XenServer to version 6.5 in case if you are using 6.0 or 6.2 version currently.

Networking

Issue ID 668835: WCCP does not work with GRE redirection when loopback IP is configured on the router. If you have loopback IP as the WCCP router IP configured on SD-WAN-WANOP on any previous software release version other than 9.3.x, upgrading to software release version 9.3.x will not resolve the issue.

Workaround: You need to reconfigure the cache IP addresses and service group by performing **Change Mode** in the SD-WAN WANOP GUI by navigating to **Configuration** >**Advance deployments** >**WCCP**.

Known Issues

SD-WAN 4100 and 5100 WANOP

Issue ID 688095 and 687990: In NetScaler SD-WAN 4100 and 5100 WANOP appliances, when the time zone or date is changed, the NetScaler instance reboots. While rebooting the CB broker is unable to communicate with the NetScaler instance and hence displays the IP address of the NetScaler instance as 0.0.0.0. The corresponding IP address of the NetScaler instance is displayed after the reboot.

Issue ID 691656: When provisioning SD-WAN 4100 and 5100 WANOP appliances with NetScaler SD-WAN 9.3 build, the provisioning fails and a message that the NetScaler instance is down appears.

SD-WAN Appliances

Issue ID 677856: SD-WAN appliance will not honor drop or reject firewall filter rules for any traffic when the appliance

goes to Fail-to-wire (FTW) mode.

SD-WAN VPX Appliances Software Downgrade

Issue ID 670142: SD-WAN software downgrade for SD-WAN VPX appliances from release 9.1.1 to version 9.1.0 does not work in XenServer, ESXi, and AWS environments.

SD-WAN 4000 WANOP and 4000 SE

Issue ID 681550: On a NetScaler SD-WAN 4000 WANOP appliance, uploading DER encoded certificate for the SSL profile is ignored and no error message is displayed in the web GUI. Only PEM encoded certificates are accepted.

Two Box Mode

Issue ID 681680: After a factory reset on the SD-WAN SE appliance in a two box mode, configuration sync between SD-WAN WANOP and SD-WAN SE appliances fails due to stale SSL certificates.

Workaround: Disable and re-enable two box mode on the SD-WAN WANOP appliance.

SD-WAN 1000 / 2000

Issue ID 681663: When you upgrade SD-WAN 1000 / 2000 appliance from release build version 9.1.2.26 to 9.2.x, a warning is displayed in the browser.

Workaround: Perform the upgrade in an in-cognito mode window of the Google Chrome browser.

SD-WAN WANOP 4000, 4100, and 5100 – NetScaler and WANOP instance information unavailable in the GUI after TACACS readonly user login

Issue ID 688948: On SD-WAN WANOP 4000 platform editions with SVM running software release version earlier than 9.3.0, TACACS user with readonly viewer privilege setting did not require executable command for user profile in the remote TACACS server.

In release 9.3.x, readonly command needs to be configured for read only user in the remote TACACS server.

Sample user configuration in remote TACACS server:

```
user = tac_super1 {

        login = cleartext tac_super1_pwd

        cmd = superuser { permit .* }

}

user = tac_ro1 {

        login = cleartext tac_ro1_pwd

        cmd = readonly { permit .* } // earlier this was not required

}
```

WAN GRE Tunnel

Issue ID 681171: Fragmented GRE tunnel packets are not reassembled properly by a NetScaler SD-WAN appliance.

IPSec Tunnel Configuration

Issue ID 681121: On a NetScaler SD-WAN VPX appliance, a web GUI error is displayed and configuration fails when you try to add and configure IPSec tunnel through the SD-WAN configuration editor.

Workaround: Configure IKE and IPsec parameters except protected networks and save the configuration. Edit the configuration to add protected networks.

Enterprise Edition as MCN – SSL Profile

Issue ID 680199: On a factory shipped Enterprise Edition appliance when you create an SSL profile and associate a Service Class to the profile with unidirectional setting, the SSL profile is not checked/enabled in the SSL Profile page of the SD-WAN EE web GUI. Also, the service class is not associated to the SSL profile.

Workaround: Create a new SSL profile and associate unidirectional service classes.

Configuration and Reporting

Issue ID 683882: Audit errors are reported when you create more than one Service Class on an SD-WAN appliance with override options. This issue occurs only when you perform override for service class and create more than one service class. It is not observed when you create only one Service Class under the default section.

Upgrade Failure

Issue ID 689362: In CB VPX, upgrade from 7.2.2 to 9.3.0.74 image failed. The following error message is displayed- "Unable to upload patch. Patch size may be too large!". There is a limit for patch size in older builds.

Workaround: Upgrade to any intermediate build (7.4.3 build or any build post 7.4.3 release) and then from the intermediate build upgrade to 9.3.0 build.

Transparent proxy support for TLS 1.2

Issue ID 691900: In NetScaler SD-WAN WANOP 9.3.0, for SSL compression the SSL profile has to be configured in split mode only as transparent proxy mode is not supported.

SD-WAN GUI Audit Error

Issue ID 687693: In the NetScaler SD-WAN GUI, when you navigate to **Basic** view > **Add Service Provider** with maximum number in the **Physical Rate** field, the generated audit error is misleading - *Integer must be less than or equal to XXX.XXX (decimal number*).

Issue ID 687701: In the NetScaler SD-WAN GUI, Service Provider Queue rate value percentage is represented incorrectly when max value is added in the physical rate field.

Change Management (Single Step Upgrade) SD-WAN GUI

Issue ID 691080: The single step upgrade procedure fails on an SD-WAN MCN appliance in a high-availability mode. When you attempt to perform change management procedure using the *.zip* single step upgrade file, the non-Virtual WAN software components, such as the WANOP package transfer is initiated only when manual toggle of the appliance happens for the Primary appliance being Active. This results in version mismatch between the WANOP appliances and the single step upgrade process is not successful.

Issue ID 691359: You can download LCM package by clicking on the active/staged hyperlink under **Download Package** when using the *tar.gz* files to perform Change Management. This will download only the Virtual WAN package as in the previous software release versions. If you use the *. zip* file of single step upgrade procedure to perform Change Management, the staged/active hyperlink under **Download LCM Package** downloads the single step upgrade package.

Issue ID 691571: On low-end platform editions, such as the SD-WAN 400, 100, 2000, or VPX appliances with 4 GB or smaller memory assigned, if concurrent local change management package downloads are initiated the appliance runs out of memory and becomes unresponsive.

Workaround: Download local change management package one at a time, this reduces the load on the appliance.

Issue ID 691953: During software upgrade on an appliance using an SE license a WAN optimization related waring message appears. After the scheduled upgrade and after the WAN optimization, SVM and XenServer hotfixes are installed the warning message is cleared.

Workaround: Clear the warning messages manually or open the SD-WAN web UI in an incognito window.

Issue ID 691746: In SD-WAN 1000 and 2000 appliances, when the software is upgraded from software release version 8.1.0 to version 9.3.0 and the appliance license is changed from SE to EE, the WAN Optimization node is not displayed in the **Configuration** and **Monitoring** tabs.

Issue ID 695430 : If an MCN is activated using single step upgrade and you are pushing a configuration to the branch appliances that are configured using ZTD, the change management fails.

Workaround: Upload the SD-WAN appliance package using change management, after single step upgrade.

HDX CGP over SSL

Issue ID 690794: HDX ICA/CGP over SSL sessions behavior In Virtual WAN Standard Edition:

- HDX sessions are not being negotiated as multi stream sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under interactive_very_low class. This may cause issues in QoS, bandwidth allocation and so on as application QoS will not be triggered because the traffic is not classified as HDX sessions.

Issue ID 690805:HDX ICA/CGP over SSL behavior In Virtual WAN Enterprise Edition:

- HDX sessions are negotiates as multi stream session.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under HDX_priority_tag_1 class. But, this traffic is not reported in Application QoS reports and in HDX reports in SD-WAN Center. However, in WANOP reports display CGP over SSL sessions as HDX session.

DPI- ICMP Functionality

Issue ID 677356: A firewall policy for blocking ICMP as an application blocks only pings (echo requests). All other ICMP types are allowed to pass through.

Workaround: Instead of blocking ICMP as an application, block IP-protocol > ICMP.

DPI – Dual- mode IPERF test identifies traffic only from one node

Issue ID 678131: When dual-mode IPERF test is performed between two appliances, the traffic in NetScaler SD-WAN web management interface under Monitoring > Firewall > Connections with DPI identifies traffic flow only from one of the connections.

DPI –Traffic for Top App Family as "Standard" and Top App as "Unknown Virtual protocol" for a Standard Edition appliance

Issue IDs 678373, 678339, 678545, 675063, 676017: On a NetScaler SD-WAN Standard Edition appliance, enable EDT policy for MSI+MP for Win7 and Win2K12 XD 7.12 VDAs on ports 2598, 2599, 2600, 2601 and subsequently disable Session Reliability policy for Win7 VDA.

Start sending internet traffic and check the monitoring flows in the Standard-Edition web management interface for Classes, Rule groups – ICAUDP and ICACGPUDP, and Firewall. Check the Dashboard and Reporting page in SD-WAN Center web management interface. The results display Top Application Family as Standard and Top Applications as Unknown Virtual Protocol.

SD-WAN Center

Issue ID 683419: In the SD-WAN Center dashboard, read-only user login access generates the following GUI error: *Error in retrieving top applications*.

Issue ID 692484: In the SD-WAN Center network map dashboard, sites with manually added Static or Dynamic Virtual paths are not accounted symmetrically for both the sites. When visualizing sites in the network map, only one of the sites constituting Static or Dynamic Virtual path is displayed.

Issue ID 692486: In SD-WAN Center, intermittent 550 site information for all sites on the dashboard is displayed in yellow tile. These sites are considered as BAD sites. However, data for the sites gets auto corrected and displays correct information for all sites.

Issue ID 692487: In the SD-WAN Center dashboard, configuration setup for monitoring 400 or more sites can take approximately 4 minutes to load.

Issue ID 692500: The SD-WAN Center dashboard does not work on Internet Explorer browser, all other pages of SD-WAN Center web interface works fine on Internet Explorer browser.

Workaround: Use other browsers like Firefox or Google Chrome.

Limitations

- The number of users is equal to the total number of HDX sessions. The number of users is not based on distinct user names. That is, two sessions started by a single user on two different machines or the same machine is counted as two users.
- HDX sessions are not being negotiated as Multi Stream Sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HTML5 receiver and ICA over SSL are not supported.

- SD-WAN web GUI Diagnostic tool will not be supported on UNTRUSTED links and Dynamic Virtual Paths.
- In the SD-WAN Center Reporting page, the Application name, Application Family, and Site filter do not contain scrollable search drop-down menu.

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. VPX-SE on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the VPX-SE and the LAN subnet of the Client/Server Host have to be different. This will require additional routing configuration to be done in two places.
  – User Defined Routes (UDR) have to be added in Azure directing all Virtual WAN Data traffic from the Client/Server LAN Subnet to the LAN interface of the VPX-SE in Azure.

– Routes have to be added in the Virtual WAN Configuration File directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.

- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management Subnet unreachability.

NetScaler SD-WAN 9.3.5 Release Notes

May 23, 2018

This release notes describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3.5 for the SD-WAN Standard, WANOP, and Enterprise Edition appliances.

The NetScaler SD-WAN release 9.3 version 5 introduces the 210 Standard Edition LTE – R1 and R 2 platform. See the 210-SE LTE topic for more information about the 210-SE LTE R1/R2 platform.

Issue ID 706340: In a SD-WAN deployment with serial HA, the secondary appliance does not perform fail-to-wire and bridging traffic to LAN. If secondary appliance is turned off, then fail-to-wire occurs, and traffic is bridged to LAN.

Issue ID 706721: The GUI session stops working when you try to navigate to the configuration tab if multiple tech-support files remain in the two-box (Virtual WAN, WANOP) solution.

Issue ID 707003: During peak traffic or load condition, STS generation may cause system out-of-memory issues and creates incomplete diagnostics file.

Issue ID 707397: Virtual Paths reset because of HA failover during configuration Activation.

Issue ID 707561: Under rare condition, enabling dynamic virtual paths causes system crash.

Issue ID 707569: In a multi-routing domain environment, system crashes when it receives out of order packets.

Issue ID 709267: System crash related to dynamic virtual path.

Issue ID 705855: After configuration and activation, the SD-WAN service might crash when traffic flow is moved from one WAN service to another WAN service where NATing is enabled.

Issue ID 702711: Backing up of configuration files from SD-WAN WANOP 7.4 or pre 9.3 releases and restoring it on SD-WAN WANOP 9.3 and above releases is not supported.

Issue ID 706510: The SD-WAN Standard Edition and Enterprise Edition appliances report incorrect NetFlow data. The output interface number in the NetFlow record is inaccurate. You can ignore NetFlow data related to the output interface.

Issue ID 708986: If you switch the LTE modem firmware or restart the LTE modem, the LTE modem might crash or become inactive. This can cause the 210-SE LTE appliance to become unresponsive.

Workaround: Power cycle the 210-SE LTE appliance to recover it.

Issue ID 708995: Virtual paths become inactive when the LTE modem is rebooted after SIM card reset or firmware change.

Workaround: Reboot the 210-SE LTE appliance.

NetScaler SD-WAN 9.3.4 Release Notes

Mar 11, 2018

This release note describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3.4 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

Fixed Issues

Issue ID 704221: In NetScaler SD-WAN release version 9.2.0, the Virtual Paths become inactive during fragmentation because of improper fragment processing packets (Data/Control packets).

Issue ID 704195: In NetScaler SD-WAN release version 9.3.0, the MCN appliance encounters an exception because of incorrect appliance rollback during configuration mismatch query leading to the appliance crash.

Issue ID 703987: In NetScaler SD-WAN release version 9.3.0, the appliance might crash because of multiple events such as MTU changes, operation stats collection, or database cleanup are in progress, when the MySQL daemon is at 100% of CPU.

Issue ID 703777: In NetScaler SD-WAN release version 9.3.2, MCN failover occurs with coredump because of incorrect sequence of events that causes the route update process to stall.

Issue ID 697517: In NetScaler SD-WAN release version 9.3.0, the appliance goes out of memory triggering virtual paths to become inactive because of DPI library allocation from the dynamic memory.

Issue ID 703072: In NetScaler SD-WAN release version 9.3.2, the SD-WAN MCN appliance running as a DHCP Server might not assign IP addresses to some branch appliances if the DHCP directories are missing.

Known Issues

Issue ID 702711: When you attempt to restore configuration files from previous SD-WAN release versions such as 9.2.x to the release version 9.3.x, the *Backup file parsing* error is encountered.

NetScaler SD-WAN 9.3.3 Release Notes

Jan 23, 2018

This release notes describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3.3 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

For information about the previous release versions, see the NetScaler SD-WAN documentation on docs.citrix.com.

Fixed Issues

Issue ID 700046: NetScaler SD-WAN appliance crashes when you try to generate STS while processing high traffic volume.

Issue ID 693737: In a HA deployment with NetScaler SD-WAN VPX appliance on VMware ESXi platform, the virtual path service becomes inactive.

Issue ID 697906: NetScaler SD-WAN service is disabled when MPLS QOS queues with all tagged queues are configured through a WAN Link Template, and on receiving a packet without DSCP tag.

Issue ID 699982: When GRE Routes with Gateway eligibility are enabled, ICMP Packets to check gateway eligibility are not guaranteed to be transmitted through the GRE Tunnel.
If the tunnel is down or the gateway address does not route to the tunnel, the packet uses standard IP routing. This leads to a GRE route being eligible inappropriately.

Issue ID 7000247: When DPI multi-threading is enabled, it can cause conn_mgr/other threads to wait for connection lock on platforms. This occurs when the release conn->lock is removed before returning from firewall.

Issue ID 695993: Packets sent to the SD-WAN (IP host) Virtual IP address from a trusted WAN interface is dropped when the flow table is exhausted or unable to allocate flows. In this case, proper ICMP reply was expected from the SD-WAN appliance.

Issue ID 699665: When you configure Virtual Path IPsec between MCN appliance and two of the branch appliances, and when you attempt to send traffic from Branch1 appliance to Branch2, WAN-to-WAN forwarding packets with large size are dropped by IPsec due to buffer overflow.

Issue ID 700183: When GRE tunnel is transmitted through an untrusted interface, for example; Internet Service, the ping requests are responded, but the IP host will not forward replies/messages back to the GRE tunnel.

Issue ID 700285: Do not update VLAN ID in the packet descriptor after DPI processing is complete.

Issue ID 700585: Disabling service on SD-WAN appliances configured with BGP peering for more BGP holdtime duration, results in the BGP session becoming disabled after enabling the service.

Issue ID 701855: DHCP is enabled on LOM by default on some factory shipped 410-SE appliances. Assign an unreachable IP address to the LOM.

Issue ID 701845: SD-WAN Center encounters an internal error when trying to export configuration. The config file should be saved with a new name and then exported to change management to prevent internal error.

Issue ID 700181: The ability to reconfigure or disable two box mode is not possible when caches are configured with any other subnets other than /24.

Issue ID 698803: As part of change management procedure during SD-WAN appliance staging phase, configuration fails when you change MTU on the intermediate router to 600.

Issue ID 702520: NetScaler Management and Analytics System is not supported in SD-WAN 410 Standard Edition Platform.

Known Issues

Platform

SD-WAN Appliances

Issue ID 686768: When you configure Primary MCN and Secondary MCN for the first time, the Primary MCN is incorrectly configured as secondary MCN.

Workaround: Configure the site, which must be the Secondary MCN, in client mode. Once the configuration is active, reconfigure the site as Secondary MCN and activate the new configuration.

SD-WAN VPX Appliances

Issue ID 694837: For High Availability in Amazon Web Services (AWS) environment, Virtual WAN service is disabled on a NetScaler SD-WAN VPX Primary (active) appliance citing duplicate IP address when the HA interface on the primary appliance goes down.

Issue ID 701427: On ESXi platform, converting SD-WAN VPX to VPX-L platform requires software re-image. You need to perform the first boot on a 16GB VM.

SD-WAN 4000 WANOP and 4000 SE

Issue ID 681550: On a NetScaler SD-WAN 4000 WANOP appliance, uploading DER encoded certificate for the SSL profile is ignored and no error message is displayed in the web GUI. Only PEM encoded certificates are accepted.

Two Box Mode

Issue ID 681680: After a factory reset on the SD-WAN SE appliance in a two box mode, configuration sync between SD-WAN WANOP and SD-WAN SE appliances fails due to stale SSL certificates.

Workaround: Disable and re-enable two-box mode on the SD-WAN WANOP appliance.

SD-WAN 1000 / 2000

Issue ID 681663: When you upgrade SD-WAN 1000 / 2000 appliance from release build version 9.1.2.26 to 9.2.x, a warning is displayed in the browser.

Workaround: Perform the upgrade in an in-cognito mode window of the Google Chrome browser.

HDX CGP over SSL

Issue ID 690794: HDX ICA/CGP over SSL session's behavior In Virtual WAN Standard Edition:

- HDX sessions are not being negotiated as multi stream sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under **interactive_very_low** class. This may cause issues in QoS, bandwidth allocation and so on as application QoS will not be triggered because the traffic is not classified as HDX sessions.

Configuration

WAN GRE Tunnel

Issue ID 681171: a NetScaler SD-WAN appliance does not reassemble fragmented GRE tunnel packets properly.

IPsec Tunnel Configuration

Issue ID 681121: On a NetScaler SD-WAN VPX appliance, a web GUI error is displayed and configuration fails when you try to add and configure IPsec tunnel through the SD-WAN configuration editor.

Workaround: Configure IKE and IPsec parameters except protected networks and save the configuration. Edit the configuration to add protected networks.

Enterprise Edition as MCN – SSL Profile

Issue ID 680199: On a factory shipped Enterprise Edition appliance when you create an SSL profile and associate a Service Class to the profile with unidirectional setting, the SSL profile is not checked/enabled in the SSL Profile page of the SD-WAN EE web GUI. In addition, the service class is not associated to the SSL profile.

Workaround: Create a new SSL profile and associate unidirectional service classes.

Configuration and Reporting

Issue ID 683882: Audit errors are reported when you create more than one Service Class on an SD-WAN appliance with override options. This issue occurs only when you perform override for service class and create more than one service class. This behavior is not observed when you create only one Service Class under the default section.

Transparent proxy support for TLS 1.2

Issue ID 691900: In NetScaler SD-WAN WANOP 9.3.0, for SSL compression the SSL profile has to be configured in split mode only as transparent proxy mode is not supported.

Change Management (Single Step Upgrade) SD-WAN GUI

Issue ID 691571: On low-end platform editions, such as the SD-WAN 400, 100, 2000, or VPX appliances with 4 GB or smaller memory assigned, if concurrent local change management package downloads are initiated the appliance runs out of memory and becomes unresponsive.

Workaround: Download local change management package one at a time, this reduces the load on the appliance.

Issue ID 691953: During software upgrade on an appliance using a Standard Edition license, a WAN optimization related waring message appears. After the scheduled upgrade and after the WAN optimization, SVM and XenServer hotfixes are installed the warning message is cleared.

Workaround: Clear the warning messages manually or open the SD-WAN web UI in an incognito browser window.

Secure Peering Certificate and Keys

Issue ID 695363: In the SD-WAN GUI, on the Secure Peering Certificate and Keys page, the CA certificate contents are displayed when the private CA radio button is selected after setting the KeyStore password on a new appliance.

Workaround: You need to switch between the radio buttons of the 'Private CA' and 'CA Certificate' once to get the correct contents displayed under 'Private CA' and 'CA Certificate' for Secure Peering Certificate and Keys.

Multicast Traffic

Issue ID 694894: When you configure Application QoS rule with match type as "Application" to match 'icmp' and change the class to Real-time, and mode to load balance which overrides the default rule, the multicast traffic is not processed.

DPI Functionality

DPI- ICMP Functionality

Issue ID 677356: A firewall policy for blocking ICMP as an application blocks only pings (echo requests). All other ICMP types are allowed to pass through.

Workaround: Instead of blocking ICMP as an application, block **IP-protocol** > **ICMP**.

DPI – Dual- mode IPERF test identifies traffic only from one node

Issue ID 678131: When dual-mode IPERF test is performed between two appliances, the traffic in NetScaler SD-WAN web management interface under **Monitoring** > **Firewall** > **Connections** with DPI identifies traffic flow only from one of the connections.

DPI –Traffic for Top App Family as "Standard" and Top App as "Unknown Virtual protocol" for a Standard Edition appliance

Issue IDs 678373, 678339, 678545, 675063, 676017: On a NetScaler SD-WAN Standard Edition appliance, enable EDT policy for MSI+MP for Win7 and Win2K12 XD 7.12 VDAs on ports 2598, 2599, 2600, 2601 and subsequently disable Session Reliability policy for Win7 VDA.

Start sending internet traffic and check the monitoring flows in the Standard-Edition web management interface for Classes, Rule groups – ICAUDP and ICACGPUDP, and Firewall. Check the Dashboard and Reporting page in SD-WAN Center web management interface. The results display **Top Application Family** as **Standard** and **Top Applications** as **Unknown Virtual Protocol**.

SD-WAN Center

Issue ID 692487: In the SD-WAN Center dashboard, configuration setup for monitoring 400 or more sites can take approximately 4 minutes to load.

Issue ID 693436: The clear connections/flows clears SD WAN connection table entries and subsequently all the ICA sessions The SD-WAN Center dashboard shows incorrect results for HDX TCP and EDT classification sessions and reports it as “Not Classified”.

Issue ID 693026: For HDX configuration, only UDP ICA sessions are classified by ICA classifier. The FrameHawk ICA

session are ignored. The SD-WAN DPI fails to classify the FrameHawk sessions.

Issue ID 694541: NetScaler SD-WAN Center dashboard reports an MSI + MP session as single stream instead of multi stream. When you configure any of the default ICA/CGP ports, 1494 or 2598 as part of the DPI ICA Port range under the Global Application Settings, the ports will not be honored.

Workaround: Do not use that port for the port range.

Issue ID 692487: In a network setup with 550 sites, if you perform any network level changes in the setup, it takes approximately ~5-8 mins for the information to be displayed in the SD- WAN center Dashboard.

Limitations

- The number of users is equal to the total number of HDX sessions. The number of users is not based on distinct user names. That is, two sessions started by a single user on two different machines or the same machine is counted as two users.
- HDX sessions are not being negotiated as Multi Stream Sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HTML5 receiver and ICA over SSL are not supported.

- The audio over UDP ports; 16500 to 16509 for Real Time Transport traffic Classification and Reporting is not supported in release 9.3.1.x.

- SD-WAN web GUI Diagnostic tool will not be supported on UNTRUSTED links and Dynamic Virtual Paths.
- In the SD-WAN Center Reporting page, the Application name, Application Family, and Site filter do not contain scrollable search drop-down menu.
- For SD-WAN Standard Edition, a connection with high latency is displayed as **Poor Connection** in the Citrix Quality Indicator (CQI) Tool. While, in SD-WAN Center the Network HDX QOE is still displayed as **Good**.

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. VPX-SE on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the VPX-SE and the LAN subnet of the Client/Server Host have to be different. This will require additional routing configuration to be done in two places.
  – User Defined Routes (UDR) have to be added in Azure directing all Virtual WAN Data traffic from the Client/Server LAN Subnet to the LAN interface of the VPX-SE in Azure.
  – Routes have to be added in the Virtual WAN Configuration File directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.
- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management

Subnet unreachability.

NetScaler SD-WAN 9.3.2 Release Notes

Dec 05, 2017

This release notes describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3.2 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

For information about the previous release versions, see the NetScaler SD-WAN documentation on docs.citrix.com.

Fixed Issues

Issue ID 699095: NetScaler SD-WAN service restarts when a Dynamic virtual path is removed while traffic is traversing through it.

Issue ID 699115: NetScaler SDWAN appliance crashes and service gets disabled while handling route failovers for existing flows processed in the system.

Issue ID 698346: In the SD-WAN GUI, under Firewall Static NAT Policies, a new option "**Proxy ARP**" is added which is disabled by default. If "**Proxy ARP**" option is enabled, SD-WAN appliance will respond to the ARP requests received for the outside IP address based on the Static Inbound/Outbound NAT rule configured.

Issue ID 698202: In NetScaler SD-WAN the "Virtual Path" service type for Firewall Static NAT policies is added. You can now configure the Static Inbound/Outbound NAT for Virtual Path service by specifying the Source and Destination Zones, Inside, and Outside IP address.

Issue ID 697539: NetScaler SD-WAN might set faulty MTU's due to the loss of MTU probes in a packet loss network. This could cause bad paths causing performance issues or impacting change management.

Issue ID 698659: When you upgrade SD-WAN release version 9.2 to 9.3 using the single step upgrade procedure, the upgrade fails due to exceeded file size limit.

Issue ID 677856: SD-WAN appliance will not honor drop or reject firewall filter rules for any traffic when the appliance goes to Fail-to-wire (FTW) mode.

Issue ID 697804: NetScaler SD-WAN Enterprise edition crashes when an optimized traffic flow becomes unoptimized.

Issue ID 692486: In the SD-WAN Center, intermittent 550 site information for all sites on the dashboard is displayed in yellow tile. These sites are considered as BAD sites. However, data for the sites gets auto corrected and displays correct information for all sites.

Issue ID 694613/ 697736: When SD-WAN appliance in HA mode is peered with neighbor router through OSPF, changes to static routes from external (type 5) to type1 may result in neighboring router retaining the routes. This is more likely to happen when standby appliance becomes active.

Issue ID 698353: On NetScaler SD-WAN 4100/5100 platforms, packets may be erroneously dropped during application classification when DPI is enabled.

Issue ID 699375/699402: When a packet is lost in SureFT, it is supposed to be retransmitted. The retransmits occur, but with zero length data (observable though internal packet capture utility). Additionally, when the connection to the MCN is lost, it will continue to make requests to the MCN for all outstanding file blocks which may use as much as 100kbps for

each transaction. The transactions never terminate leaving the branch continuing to request forever.

Issue ID 697294: NetScaler SD-WAN network might experience a service restart in scale environments when Paths or Virtual Paths are changing states frequently.

Issue ID 695506: In a HA mode, on a NetScaler SD-WAN 3000-SE secondary appliance, you cannot edit and configure LDAP authentication for user administration.

Issue ID 694594: On NetScaler SD-WAN Enterprise Edition appliances, configuration update causing WCCP router to reboot multiple times results in the WAN OP functionality to be disabled.

Issue ID 696499: On a NetScaler SD-WAN 5100-SE appliance, an unexpected service restart occurs due to an IPMI failure.

Issue ID 697548: When an MCN is deployed with an HA peer, an unexpected service restart may cause Change Management on both appliances to try to manage the network causing. When this occurs, failures in Change Management staging may be observed.

Issue ID 698350: In NetScaler SD-WAN, for Firewall Static NAT Inbound Policies, traffic received on the outside IP address configured is forwarded to the originator.

Issue ID 696748: When HA is configured and active for an MCN, change management may repeatedly fail for some appliances if a HA switch is triggered unexpectedly during a change management procedure. The failures may persist even after change management has been completed for the other appliances in the network.

Known Issues

Platform

SD-WAN Appliances

Issue ID 686768: When you configure Primary MCN and Secondary MCN for the first time, the Primary MCN is incorrectly configured as secondary MCN.

SD-WAN VPX Appliances

Issue ID 694837: For High Availability in Amazon Web Services (AWS) environment, Virtual WAN service is disabled on a NetScaler SD-WAN VPX Primary (active) appliance citing duplicate IP address when the HA interface on the primary appliance goes down.

SD-WAN 4000 WANOP and 4000 SE

Issue ID 681550: On a NetScaler SD-WAN 4000 WANOP appliance, uploading DER encoded certificate for the SSL profile is ignored and no error message is displayed in the web GUI. Only PEM encoded certificates are accepted.

Two Box Mode

Issue ID 681680: After a factory reset on the SD-WAN SE appliance in a two box mode, configuration sync between SD-WAN WANOP and SD-WAN SE appliances fails due to stale SSL certificates.

Workaround: Disable and re-enable two-box mode on the SD-WAN WANOP appliance.

SD-WAN 1000 / 2000

Issue ID 681663: When you upgrade SD-WAN 1000 / 2000 appliance from release build version 9.1.2.26 to 9.2.x, a warning is displayed in the browser.

Workaround: Perform the upgrade in an in-cognito mode window of the Google Chrome browser.

HDX CGP over SSL

Issue ID 690794: HDX ICA/CGP over SSL session's behavior In Virtual WAN Standard Edition:

- HDX sessions are not being negotiated as multi stream sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under **interactive_very_low** class. This may cause issues in QoS, bandwidth allocation and so on as application QoS will not be triggered because the traffic is not classified as HDX sessions.

Configuration

WAN GRE Tunnel

Issue ID 681171: a NetScaler SD-WAN appliance does not reassemble fragmented GRE tunnel packets properly.

IPsec Tunnel Configuration

Issue ID 681121: On a NetScaler SD-WAN VPX appliance, a web GUI error is displayed and configuration fails when you try to add and configure IPsec tunnel through the SD-WAN configuration editor.

Workaround: Configure IKE and IPsec parameters except protected networks and save the configuration. Edit the configuration to add protected networks.

Enterprise Edition as MCN – SSL Profile

Issue ID 680199: On a factory shipped Enterprise Edition appliance when you create an SSL profile and associate a Service Class to the profile with unidirectional setting, the SSL profile is not checked/enabled in the SSL Profile page of the SD-WAN EE web GUI. In addition, the service class is not associated to the SSL profile.

Workaround: Create a new SSL profile and associate unidirectional service classes.

Configuration and Reporting

Issue ID 683882: Audit errors are reported when you create more than one Service Class on an SD-WAN appliance with override options. This issue occurs only when you perform override for service class and create more than one service class. This behavior is not observed when you create only one Service Class under the default section.

Transparent proxy support for TLS 1.2

Issue ID 691900: In NetScaler SD-WAN WANOP 9.3.0, for SSL compression the SSL profile has to be configured in split mode only as transparent proxy mode is not supported.

Change Management SD-WAN GUI

Issue ID 698803: As part of change management procedure during SD-WAN appliance staging phase, configuration fails when you change MTU on the intermediate router to 600.

Change Management (Single Step Upgrade) SD-WAN GUI

Issue ID 691571: On low-end platform editions, such as the SD-WAN 400, 100, 2000, or VPX appliances with 4 GB or smaller memory assigned, if concurrent local change management package downloads are initiated the appliance runs out of memory and becomes unresponsive.

Workaround: Download local change management package one at a time, this reduces the load on the appliance.

Issue ID 691953: During software upgrade on an appliance using a Standard Edition license, a WAN optimization related waring message appears. After the scheduled upgrade and after the WAN optimization, SVM and XenServer hotfixes are installed the warning message is cleared.

Workaround: Clear the warning messages manually or open the SD-WAN web UI in an incognito browser window.

Issue ID 691080: The single step upgrade procedure fails on an SD-WAN MCN appliance in a high availability mode. Single Step Upgrade involves transmission of WAN Opt packages by MCN once Virtual WAN package Activation is done. If the MCN site is in HA, there is a possibility that on Activation of Virtual WAN packages the HA role might toggle. Due to this the Secondary appliance becomes Active which never had the WAN Opt packages uploaded hence would not be able to transmit it to other sites.

Workaround: If the HA toggle happens post Virtual WAN package activation, then manually perform HA toggle. Navigate to **Configuration** > **Appliance Settings** > **Administrator Interface** > **Miscellaneous** and click Switch **HA Mode**. This will make the Primary appliance as Active MCN and hence the WAN Opt package transfer will resume.

Secure Peering Certificate and Keys

Issue ID 695363: In the SD-WAN GUI, on the Secure Peering Certificate and Keys page, the CA certificate contents are displayed when the private CA radio button is selected after setting the KeyStore password on a new appliance.

Workaround: You need to switch between the radio buttons of the 'Private CA' and 'CA Certificate' once to get the correct contents displayed under 'Private CA' and 'CA Certificate' for Secure Peering Certificate and Keys.

Multicast Traffic

Issue ID 694894: When you configure Application QoS rule with match type as "Application" to match 'icmp' and change the class to Real-time, and mode to load balance which overrides the default rule, the multicast traffic is not processed.

DPI Functionality

DPI- ICMP Functionality

Issue ID 677356: A firewall policy for blocking ICMP as an application blocks only pings (echo requests). All other ICMP types are allowed to pass through.

Workaround: Instead of blocking ICMP as an application, block **IP-protocol** > **ICMP**.

DPI – Dual- mode IPERF test identifies traffic only from one node

Issue ID 678131: When dual-mode IPERF test is performed between two appliances, the traffic in NetScaler SD-WAN

web management interface under **Monitoring** > **Firewall** > **Connections** with DPI identifies traffic flow only from one of the connections.

DPI – Traffic for Top App Family as "Standard" and Top App as "Unknown Virtual protocol" for a Standard Edition appliance

Issue IDs 678373, 678339, 678545, 675063, 676017: On a NetScaler SD-WAN Standard Edition appliance, enable EDT policy for MSI+MP for Win7 and Win2K12 XD 7.12 VDAs on ports 2598, 2599, 2600, 2601 and subsequently disable Session Reliability policy for Win7 VDA.

Start sending internet traffic and check the monitoring flows in the Standard-Edition web management interface for Classes, Rule groups – ICAUDP and ICACGPUDP, and Firewall. Check the Dashboard and Reporting page in SD-WAN Center web management interface. The results display **Top Application Family** as **Standard** and **Top Applications** as **Unknown Virtual Protocol**.

SD-WAN Center

Issue ID 692487: In the SD-WAN Center dashboard, configuration setup for monitoring 400 or more sites can take approximately 4 minutes to load.

Issue ID 693436: The clear connections/flows clears SD WAN connection table entries and subsequently all the ICA sessions The SD-WAN Center dashboard shows incorrect results for HDX TCP and EDT classification sessions and reports it as “Not Classified”.

Issue ID 693026: For HDX configuration, only UDP ICA sessions are classified by ICA classifier. The FrameHawk ICA session are ignored. The SD-WAN DPI fails to classify the FrameHawk sessions.

Issue ID 694541: NetScaler SD-WAN Center dashboard reports an MSI + MP session as single stream instead of multi stream. When you configure any of the default ICA/CGP ports, 1494 or 2598 as part of the DPI ICA Port range under the Global Application Settings, the ports will not be honored.

Workaround: Do not use that port for the port range.

Issue ID 692487: In a network setup with 550 sites, if you perform any network level changes in the setup, it takes approximately ~5-8 mins for the information to be displayed in the SD- WAN center Dashboard.

Limitations

- The number of users is equal to the total number of HDX sessions. The number of users is not based on distinct user names. That is, two sessions started by a single user on two different machines or the same machine is counted as two users.
- HDX sessions are not being negotiated as Multi Stream Sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HTML5 receiver and ICA over SSL are not supported.

- The audio over UDP ports; 16500 to 16509 for Real Time Transport traffic Classification and Reporting is not supported

in release 9.3.1.x.

- SD-WAN web GUI Diagnostic tool will not be supported on UNTRUSTED links and Dynamic Virtual Paths.
- In the SD-WAN Center Reporting page, the Application name, Application Family, and Site filter do not contain scrollable search drop-down menu.
- For SD-WAN Standard Edition, a connection with high latency is displayed as **Poor Connection** in the Citrix Quality Indicator (CQI) Tool. While, in SD-WAN Center the Network HDX QOE is still displayed as **Good**.

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. VPX-SE on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the VPX-SE and the LAN subnet of the Client/Server Host have to be different. This will require additional routing configuration to be done in two places.
  – User Defined Routes (UDR) have to be added in Azure directing all Virtual WAN Data traffic from the Client/Server LAN Subnet to the LAN interface of the VPX-SE in Azure.
  – Routes have to be added in the Virtual WAN Configuration File directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.
- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management Subnet unreachability.

NetScaler SD-WAN 9.3.1 Release Notes

Oct 05, 2017

This release notes describes known issues, and fixed issues applicable to Citrix NetScaler SD-WAN software release 9.3.1 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

For information about the previous release versions, see the NetScaler SD-WAN documentation on docs.citrix.com.

SD-WAN 4100 and 5100 WANOP

Issue ID 688095 and 687990: In NetScaler SD-WAN 4100 and 5100 WANOP appliances, when the time zone or date is changed, the NetScaler instance reboots. While rebooting the CB broker is unable to communicate with the NetScaler instance and hence displays the IP address of the NetScaler instance as 0.0.0.0. The corresponding IP address of the NetScaler instance is displayed after the reboot.

Change Management (Single Step Upgrade) SD-WAN GUI

Issue ID 691359: You can download LCM package by clicking on the active/staged hyperlink under **Download Package** when using the tar.gz files to perform change Management. This will download only the Virtual WAN package as in the previous release versions. If you use the .zip file of single step upgrade procedure to perform change management, the staged/active hyperlink under **Download LCM Package** downloads the single step upgrade package.

Issue ID 691746: In SD-WAN 1000 and 2000 appliances, when the software is upgraded from 8.1.0 to 9.3.0 and the appliance license is changed from SE to EE, the WAN Optimization node is not displayed in the **Configuration** and **Monitoring** tabs

SD-WAN Center

Issue ID 683419: In the SD-WAN Center, read-only user login access generates the following GUI error:

Error in retrieving top applications.

Issue ID 692484: In the SD-WAN Center network map dashboard, sites with manually added Static or Dynamic Virtual paths are not accounted symmetrically for both the sites. When visualizing sites in the network map, only one of the sites constituting Static or Dynamic Virtual path is displayed.

Issue ID 692500: The SD-WAN Center dashboard does not work on Internet Explorer browser, all other pages of SD-WAN Center web interface works fine on Internet Explorer browser.

Platform

SD-WAN Appliances

Issue ID 686768: When you configure Primary MCN and Secondary MCN for the first time, the Primary MCN is incorrectly configured as secondary MCN.

Issue ID 677856: SD-WAN appliance will not honor drop or reject firewall filter rules for any traffic when the appliance goes to Fail-to-wire (FTW) mode.

SD-WAN VPX Appliances

Issue ID 694837: For High Availability in Amazon Web Services (AWS) environment, Virtual WAN service is disabled on a NetScaler SD-WAN VPX Primary (active) appliance citing duplicate IP address when the HA interface on the primary appliance goes down.

SD-WAN VPX Appliances Software Downgrade

Issue ID 670142: SD-WAN software downgrade for SD-WAN VPX appliances from release 9.1.1 to version 9.1.0 does not work in XenServer, ESXi, and AWS environments.

SD-WAN 4000 WANOP and 4000 SE

Issue ID 681550: On a NetScaler SD-WAN 4000 WANOP appliance, uploading DER encoded certificate for the SSL profile is ignored and no error message is displayed in the web GUI. Only PEM encoded certificates are accepted.

SD-WAN 4100 and 5100 WANOP

Issue ID 691656: When provisioning SD-WAN 4100 and 5100 WANOP appliances with NetScaler SD-WAN 9.3 build, the provisioning fails and a message that the NetScaler SD-WAN instance is down appears. Retrying the operation succeeds.

Two Box Mode

Issue ID 681680: After a factory reset on the SD-WAN SE appliance in a two box mode, configuration sync between SD-WAN WANOP and SD-WAN SE appliances fails due to stale SSL certificates.

Workaround: Disable and re-enable two-box mode on the SD-WAN WANOP appliance.

SD-WAN 1000 / 2000

Issue ID 681663: When you upgrade SD-WAN 1000 / 2000 appliance from release build version 9.1.2.26 to 9.2.x, a warning is displayed in the browser.

Workaround: Perform the upgrade in an in-cognito mode window of the Google Chrome browser.

SD-WAN WANOP 4000, 4100, and 5100

Issue ID 688948: On SD-WAN WANOP 4000 platform editions with SVM running software release version earlier than 9.3.0, TACACS user with read-only viewer privilege setting did not require executable command for user profile in the remote TACACS server.

In release 9.3.x, read-only command needs to be configured for read only user in the remote TACACS server.

Sample user configuration in remote TACACS server:

```
user = tac_super1 {

        login = cleartext tac_super1_pwd

        cmd = superuser { permit .* }

}

user = tac_ro1 {

        login = cleartext tac_ro1_pwd

        cmd = read-only { permit .* } // earlier this was not required

}
```

Upgrade Failure

Issue ID 689362: In CB VPX, upgrade from 7.2.2 to 9.3.0.74 image failed. The following error message is displayed- "Unable to upload patch. Patch size may be too large!" There is a limit for patch size in older builds.

SD-WAN GUI Audit Error

Issue ID 687693: In the NetScaler SD-WAN GUI, when you navigate to Basic view > Add Service Provider with maximum number in the Physical Rate field, the generated audit error is misleading - *Integer must be less than or equal to XXX.XXX (decimal number*).

Issue ID 687701: In the NetScaler SD-WAN GUI, Service Provider Queue rate value percentage is represented incorrectly when max value is added in the physical rate field.

MSI behavior on Standard Edition and Enterprise Edition Appliances

Issue ID 690826: On NetScaler SD-WAN Standard Edition appliance although there is no optimization feature, the new HDX implementation in 9.3.0_x prevents HDX sessions from being negotiated as Multistream or Multistream + Multiport even if policies are Enabled on XD Studio.

SD-WAN release 9.2.1_X did not exhibit this behavior. The SD MSI and MSI+MP policies were honored and sessions would get negotiated as multistream if these policies were set. This behavior is deprecated in 9.3.0_x. This behavior is applicable to Enterprise Edition appliance as well.

HDX CGP over SSL

Issue ID 690794: HDX ICA/CGP over SSL session's behavior In Virtual WAN Standard Edition:

- HDX sessions are not being negotiated as multi stream sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under **interactive_very_low** class. This may cause issues in QoS, bandwidth allocation and so on as application QoS will not be triggered because the traffic is not classified as HDX sessions.

Issue ID 690805: HDX ICA/CGP over SSL behavior In Virtual WAN Enterprise Edition:

- HDX sessions are negotiates as multi stream session.
- HDX traffic is classified as belonging to Hyper Text Transfer Protocol Secure (https) application and web family.
- HDX traffic falls under **HDX_priority_tag_1** class. However, this traffic is not reported in Application QoS reports and in HDX reports in SD-WAN Center. However, in WANOP reports display CGP over SSL sessions as HDX session.

Configuration

WAN GRE Tunnel

Issue ID 681171: a NetScaler SD-WAN appliance does not reassemble fragmented GRE tunnel packets properly.

IPsec Tunnel Configuration

Issue ID 681121: On a NetScaler SD-WAN VPX appliance, a web GUI error is displayed and configuration fails when you try to add and configure IPsec tunnel through the SD-WAN configuration editor.

Workaround: Configure IKE and IPsec parameters except protected networks and save the configuration. Edit the configuration to add protected networks.

Enterprise Edition as MCN – SSL Profile

Issue ID 680199: On a factory shipped Enterprise Edition appliance when you create an SSL profile and associate a Service Class to the profile with unidirectional setting, the SSL profile is not checked/enabled in the SSL Profile page of the SD-WAN EE web GUI. In addition, the service class is not associated to the SSL profile.

Workaround: Create a new SSL profile and associate unidirectional service classes.

Configuration and Reporting

Issue ID 683882: Audit errors are reported when you create more than one Service Class on an SD-WAN appliance with override options. This issue occurs only when you perform override for service class and create more than one service class. This behavior is not observed when you create only one Service Class under the default section.

Transparent proxy support for TLS 1.2

Issue ID 691900: In NetScaler SD-WAN WANOP 9.3.0, for SSL compression the SSL profile has to be configured in split mode only as transparent proxy mode is not supported

Change Management (Single Step Upgrade) SD-WAN GUI

Issue ID 691571: On low-end platform editions, such as the SD-WAN 400, 100, 2000, or VPX appliances with 4 GB or smaller memory assigned, if concurrent local change management package downloads are initiated the appliance runs out of memory and becomes unresponsive.

Workaround: Download local change management package one at a time, this reduces the load on the appliance.

Issue ID 691953: During software upgrade on an appliance using a Standard Edition license, a WAN optimization related waring message appears. After the scheduled upgrade and after the WAN optimization, SVM and XenServer hotfixes are installed the warning message is cleared.

Workaround: Clear the warning messages manually or open the SD-WAN web UI in an incognito browser window.

Issue ID 691080: The single step upgrade procedure fails on an SD-WAN MCN appliance in a high availability mode. Single Step Upgrade involves transmission of WAN Opt packages by MCN once Virtual WAN package Activation is done. If the MCN site is in HA, there is a possibility that on Activation of Virtual WAN packages the HA role might toggle. Due to this the Secondary appliance becomes Active which never had the WAN Opt packages uploaded hence would not be able to transmit it to other sites.

Workaround: If the HA toggle happens post Virtual WAN package activation, then manually perform HA toggle. Navigate to **Configuration** > **Appliance Settings** > **Administrator Interface** > **Miscellaneous** and click **Switch HA Mode**. This will make the Primary appliance as Active MCN and hence the WAN Opt package transfer will resume.

Secure Peering Certificate and Keys

Issue ID 695363: In the SD-WAN GUI, on the Secure Peering Certificate and Keys page, the CA certificate contents are displayed when the private CA radio button is selected after setting the KeyStore password on a new appliance.

Workaround: You need to switch between the radio buttons of the 'Private CA' and 'CA Certificate' once to get the correct contents displayed under 'Private CA' and 'CA Certificate' for Secure Peering Certificate and Keys.

Multicast Traffic

Issue ID 694894: When you configure Application QoS rule with match type as "Application" to match 'icmp' and change the class to Realtime, and mode to loadbalance which overrides the default rule, the multicast traffic is not processed.

DPI- ICMP Functionality

Issue ID 677356: A firewall policy for blocking ICMP as an application blocks only pings (echo requests). All other ICMP types are allowed to pass through.

Workaround: Instead of blocking ICMP as an application, block **IP-protocol** > **ICMP**.

DPI – Dual- mode IPERF test identifies traffic only from one node

Issue ID 678131: When dual-mode IPERF test is performed between two appliances, the traffic in NetScaler SD-WAN web management interface under **Monitoring** > **Firewall** > **Connections** with DPI identifies traffic flow only from one of the connections.

DPI –Traffic for Top App Family as "Standard" and Top App as "Unknown Virtual protocol" for a Standard Edition appliance

Issue IDs 678373, 678339, 678545, 675063, 676017: On a NetScaler SD-WAN Standard Edition appliance, enable EDT policy for MSI+MP for Win7 and Win2K12 XD 7.12 VDAs on ports 2598, 2599, 2600, 2601 and subsequently disable Session Reliability policy for Win7 VDA.

Start sending internet traffic and check the monitoring flows in the Standard-Edition web management interface for Classes, Rule groups – ICAUDP and ICACGPUDP, and Firewall. Check the Dashboard and Reporting page in SD-WAN Center web management interface. The results display **Top Application Family** as **Standard** and **Top Applications** as **Unknown Virtual Protocol**.

SD-WAN Center

Issue ID 692486: In the SD-WAN Center, intermittent 550 site information for all sites on the dashboard is displayed in yellow tile. These sites are considered as BAD sites. However, data for the sites gets auto corrected and displays correct information for all sites.

Issue ID 692487: In the SD-WAN Center dashboard, configuration setup for monitoring 400 or more sites can take approximately 4 minutes to load.

Issue ID 693436: The clear connections/flows clears SD WAN connection table entries and subsequently all the ICA sessions The SD-WAN Center dashboard shows incorrect results for HDX TCP and EDT classification sessions and reports it as “Not Classified”.

Issue ID 693026: For HDX configuration, only UDP ICA sessions are classified by ICA classifier. The FrameHawk ICA session are ignored. The SD-WAN DPI fails to classify the FrameHawk sessions.

Issue ID 694541: NetScaler SD-WAN Center dashboard reports an MSI + MP session as single stream instead of multi stream. When you configure any of the default ICA/CGP ports, 1494 or 2598 as part of the DPI ICA Port range under the Global Application Settings, the ports will not be honored.

Workaround: Do not use that port for the port range.

Limitations

- The number of users is equal to the total number of HDX sessions. The number of users is not based on distinct user names. That is, two sessions started by a single user on two different machines or the same machine is counted as two users.
- HDX sessions are not being negotiated as Multi Stream Sessions even though MSI is enabled on the appliance and MSI+MP policies are set on incoming ICA traffic.
- HTML5 receiver and ICA over SSL are not supported.

- SD-WAN web GUI Diagnostic tool will not be supported on UNTRUSTED links and Dynamic Virtual Paths.
- In the SD-WAN Center Reporting page, the Application name, Application Family, and Site filter do not contain scrollable search drop-down menu.
- For SD-WAN Standard Edition, a connection with high latency is displayed as **Poor Connection** in the Citrix Quality Indicator (CQI) Tool. While, in SD-WAN Center the Network HDX QOE is still displayed as **Good**.

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. VPX-SE on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the VPX-SE and the LAN subnet of the Client/Server Host have to be different. This will require additional routing configuration to be done in two places.
  – User Defined Routes (UDR) have to be added in Azure directing all Virtual WAN Data traffic from the Client/Server LAN Subnet to the LAN interface of the VPX-SE in Azure.
  – Routes have to be added in the Virtual WAN Configuration File directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.
- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management Subnet unreachability.

NetScaler SD-WAN 9.2.1 Release Notes

Oct 05, 2017

This release notes describes the fixed issues, known issues, and limitations applicable to Citrix NetScaler SD-WAN software release 9.2.1 for the SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

> **Note**
>
> - This release note document does not include security related fixes. For a list of security related fixes and advisories, see the Citrix security bulletin.
> - The [# XXXXXX] labels for issue descriptions are internal tracking IDs used by the SD-WAN support team.

What's New in Release 9.2.1 (build 1002)

The following new features and enhancements were introduced in NetScaler SD-WAN Release 9.2.1 build 1002:

NetScaler SD-WAN release 9.2.1, build 1002 has new images with security fix for CVE-2017-14602.

This vulnerability is only present when the above versions are used on the following appliance models:

- Citrix NetScaler SD-WAN model 5100 WAN Optimization appliances
- Citrix NetScaler SD-WAN (CloudBridge) model 5000 WAN Optimization appliances
- Citrix NetScaler SD-WAN model 4100 WAN Optimization appliances
- Citrix NetScaler SD-WAN (CloudBridge) model 4000 WAN Optimization appliances

For additional information related to this security fix, impacted editions, and platforms, refer to the security bulletin posted at https://support.citrix.com/article/CTX228091.
.Best practices for use of WAN Optimization products are now available at: Read more

Fixed Issues

SD-WAN 4000 WANOP and 4000 SE

- Issue ID 680778: A configuration audit error occurs in two-box mode deployment when a NetScaler SD-WAN 4000 SE appliance with two interface groups is configured with first interface group having bridged pair with two ethernet interfaces selected, and second interface group is connected to the WANOP appliance. The error occurs when the first interface group is enabled with WCCP listener indicating that multiple ethernet interfaces cannot be enabled with WCCP. When you revert configuration by disabling WCCP on the first interface group and enabling it on the second interface group, the same configuration audit error is displayed even though only one ethernet interface is enabled on the interface group.
- Issue ID 680825: On a NetScaler SD-WAN 4000 appliance with release version 9.2, the HTTP service does not work for one of the SD-WAN instances and fails to start or restart the HTTPS service.
- Issue ID 679121: While upgrading SD-WAN 4000 appliance from old releases to 9.2 release, the SD-WAN GUI appears before the upgrade process is completed. The old image is listed in the GUI.

SD-WAN 4100 SE

- Issue ID 675715: On a NetScaler SD-WAN 4100 SE appliance, changing Interface settings for 1G interface does not work and causes link to become inactive. For example; changing the speed to 100MB does not work. The interface settings change option is disabled for all 1G ports similar to the 10G ports as it is not supported on the 4100-SE bare metal platform.

TCP Fragmented traffic

- Issue ID 681472: Virtual WAN drops TCP Fragmented traffic when firewall connection tracking is enabled.

NTP Server Time Settings

- Issue ID 680987: On NetScaler SD-WAN 2000 appliances, when you change the NTP server settings, the Enterprise Edition appliance time settings sync up with the new NTP server time settings and the correct time zone format is displayed. However, the new NTP server time settings on a WANOP appliance are not synchronized with the new NTP server time settings.

Diagnostic tool

- Issue ID 680251: In a NetScaler SD-WAN VPX appliance setup, multiple IPREF client TCP sessions are initiated while server session is still on causing the server to display additional entries even when the client has stopped sending any further traffic.

Rules Group Tab

- Issue ID 681562: The Rule group tab in SD-WAN Center report page does not show any data for the configured applications.

DPI- No audit error on disabling DPI

- Issue ID 681175: If an application object created with DPI application is associated to a firewall policy template, and is used in firewall and then if the DPI is disabled, there is no audit error message displayed indicating that there are rules still associated with firewall as the firewall is still functional.

SSL Profile Name

- Issue ID 681482: In a NetScaler SD-WAN VPX appliance setup, when you create an SSL profile and try to edit the profile and save it, the following error message is displayed: “No object with profile name exists”.

SSL Profile Page

- Issue ID 681443: When creating or editing an SSL profile, the settings are saved but the application does not redirected to the SSL Profile home page.

GUI

- Issue ID 681649: Unable to enable DHCP Server and Relay for management from the UI. On selecting **Enbale DHCP Server**, the fields **Lease Time**, **Domain Name**, **Start IP Address** and **End IP Addresss** should be editable, but these fields are not editable.

Security Vulnerability

- Issue ID 690709: Unauthenticated remote code execution on NetScaler SD-WAN. This security hotfix addresses the

vulnerabilities as described in the security bulletin article (CTX225990).

Known Issues

SD-WAN 4000 WANOP and 4000 SE

- Issue ID 681550: On a NetScaler SD-WAN 4000 WANOP appliance, uploading DER encoded certificate for the SSL profile is ignored and no error message is displayed in the web GUI. Only PEM encoded certificates are accepted.

Two Box Mode

- Issue ID 681680: After a factory reset on the SD-WAN SE appliance in a two box mode, configuration sync between SD-WAN WANOP and SD-WAN SE appliances fails due to stale SSL certificates.

  Workaround: Disable and re-enable two box mode on the SD-WAN WANOP appliance.

SD-WAN 1000 / 2000

- Issue ID 681663: When you upgrade SD-WAN 1000 / 2000 appliance from release build version 9.1.2.26 to 9.2.x, a warning is displayed in the browser.

  Workaround: Perform the upgrade in an in-cognito mode window of the Google Chrome browser.

SD-WAN WANOP

- Issue ID 675452: NetScaler SD-WAN WANOP client info displays OS version as Windows 8 even when plugin is installed in Windows 10 OS.

SD-WAN GUI

- Issue ID 683520: In the SD-WAN GUI, changing the interface settings for interface under **Configuration** > **Appliance Settings** > **Network adapters** > **Ethernet** does not work for the SD-WAN 1000-EE, 2000-EE and 400-SE platforms.

WAN GRE Tunnel

- Issue ID 681171: Fragmented GRE tunnel packets are not reassembled properly by a NetScaler SD-WAN appliance.

IPSec Tunnel Configuration

- Issue ID 681121: On a NetScaler SD-WAN VPX appliance, a web GUI error is displayed and configuration fails when you try to add and configure IPSec tunnel through the SD-WAN configuration editor.

  Workaround: Configure IKE and IPsec parameters except protected networks and save the configuration. Edit the configuration to add protected networks.

Enterprise Edition as MCN – SSL Profile

- Issue ID 680199: On a factory shipped Enterprise Edition appliance when you create an SSL profile and associate a Service Class to the profile with unidirectional setting, the SSL profile is not checked/enabled in the SSL Profile page of

the SD-WAN EE web GUI. Also, the service class is not associated to the SSL profile.

Workaround: Create a new SSL profile and associate unidirectional service class (es).

Simplified Configuration

- Issue ID 678342: In the SD-WAN configuration editor, secondary level confirmation is not provided when deleting a WAN Link, Interface Group, or Static Route from the Basics view.

Ethernet Interfaces Configuration

- Issue ID 680585: In a NetScaler SD-WAN Standard Edition appliance web GUI, the Basic View under Configuration Editor allows you to create Interface without selecting ethernet interfaces. The created interface is displayed in the Advanced View as VLAN 0 instead of displaying in the Basic View.

Configuration and Reporting

- Issue ID 683882: Audit errors are reported when you create more than one Service Class on an SD-WAN appliance with override options. This issue occurs only when you perform override for service class and create more than one service class. It is not observed when you create only one Service Class under the default section.

DPI- ICMP Functionality

- Issue ID 677356: A firewall policy for blocking ICMP as an application blocks only pings (echo requests). All other ICMP types are allowed to pass through.

Workaround: Instead of blocking ICMP as an application, block IP-protocol > ICMP.

DPI – Dual- mode IPERF test identifies traffic only from one node

- Issue ID 678131: When dual-mode IPERF test is performed between two appliances, the traffic in NetScaler SD-WAN web management interface under Monitoring > Firewall > Connections with DPI identifies traffic flow only from one of the connections.

DPI - Traffic classified as unknown when the traffic flows through EE appliances

- Issue ID 677504: Applications are classified as Unknown protocol when the traffic flows through EE appliances, because the compressed traffic is not classified. Therefore, the Firewall rules do not work on EE appliance with DPI enabled when rules are configured with Application, Application Family or Application Object firewall policies. This issue occurs only when a WANOP Service Class Compression policy is configured on a Standard Edition/Enterprise Edition or Standard Edition/Standard Edition appliance with a WANOP deployment mode.

DPI – Any application traffic sent via GRE Tunnel is reported as GRE in SD-WAN Center

- Issue ID 680994: Ideally, any application traffic (example HTTP) sent through the GRE tunnel should be classified by DPI reported as both GRE and the real application traffic (example HTTP) in the Application section of Reporting page in SD-WAN Center. Due to this bug, the real application (example HTTP) is also reported as GRE traffic. This bug is only a reporting issue and the real classification has no issues in the site level DPI. Both the classification and firewall actions after DPI will have no impact in any site.

DPI –Traffic for Top App Family as "Standard" and Top App as "Unknown Virtual protocol" for a Standard

Edition appliance

- Issue IDs 678373, 678339, 678545, 675063, 676017: On a NetScaler SD-WAN Standard Edition appliance, enable EDT policy for MSI+MP for Win7 and Win2K12 XD 7.12 VDAs on ports 2598, 2599, 2600, 2601 and subsequently disable Session Reliability policy for Win7 VDA.

Start sending internet traffic and check the monitoring flows in the Standard-Edition web management interface for Classes, Rule groups – ICAUDP and ICACGPUDP, and Firewall. Check the Dashboard and Reporting page in SD-WAN Center web management interface. The results display Top Application Family as Standard and Top Applications as Unknown Virtual Protocol.

SD-WAN Center – GUI Error

- Issue ID 683419: In the SD-WAN Center, read-only user login access generates the following GUI error: Error in retrieving top applications.

Limitations

SD-WAN Center and Diagnostic Tool

- SD-WAN web GUI Diagnostic tool will not be supported on UNTRUSTED links and Dynamic Virtual Paths.
- In the SD-WAN Center Reporting page, the Application name, Application Family, and Site filter do not contain scrollable search drop-down menu.

Microsoft Azure

- A VM in Azure can have Public IP on only one interface. This VM needs to be on the WAN link to establish Virtual Path. Management is accessed over Private network. While configuring SD-WAN SE-VPX, network interfaces have to added in following order:

    a) WAN interface (Private IP, Public IP)

    b) LAN interface (Private IP)

    c) Management interface (Private IP)

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. SE-VPX on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the SE-VPX and the LAN subnet of the Client/Server Host have to be different. This will require additional routing configuration to be done in two places.

    – User Defined Routes (UDR) have to be added in Azure directing all Virtual WAN Data traffic from the Client/Server LAN Subnet to the LAN interface of the SD-WAN SE-VPX in Azure.

    – Routes have to be added in the Virtual WAN Configuration File directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.

- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management

Subnet unreachability.

What's New

Jan 21, 2018

What's New in Release 9.3.3

The SD-WAN 9.3.3 release introduces the 210 Standard Edition appliance.

For more information, see; NetScaler SD-WAN 210-SE

What’s New in Release 9.3.0 (build 1000)

The following new features and enhancements were introduced in NetScaler SD-WAN Release 9.3.0 build 1000:

NetScaler SD-WAN release 9.3.0, build 1000 has new images with security fix for CVE-2017-14602.

This vulnerability is only present when the above versions are used on the following appliance models:

- Citrix NetScaler SD-WAN model 5100 WAN Optimization appliances
- Citrix NetScaler SD-WAN (CloudBridge) model 5000 WAN Optimization appliances
- Citrix NetScaler SD-WAN model 4100 WAN Optimization appliances
- Citrix NetScaler SD-WAN (CloudBridge) model 4000 WAN Optimization appliances

For additional information related to this security fix, impacted editions, and platforms, refer to the security bulletin posted at https://support.citrix.com/article/CTX228091.

Best practices for use of WAN Optimization products are now available at: Read more

What’s New in Release 9.3.0 (build 161)

The following new software, hardware features and enhancements were introduced in NetScaler SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances for Release 9.3.

HDX Traffic Identification and MSI support

- You can enable DPI for Citrix ICA applications to classify Citrix ICA HDX sessions over any transport protocol such as TCP, UDP, CGP or HTTP.

- You can enable ICA multi stream to allow multiple ICA streams in a session.

Application QoS Rules

- You can now take advantage of the application DPI Engine to filter traffic flows based on application, application family, or application object match-types and apply application QoS rules to them.

Application QoE

- You can now assess the user experience of ICA / HDX applications using the Quality of Experience (QoE) parameter. The QoE is a numeric value between 0–100, the higher the value the better the user experience. QoE is enabled by default for all ICA / HDX applications.

Alarm Enhancements

- In NetScaler SD-WAN 9.3 the events notification system is enhanced by including a feature that allows you to configure alarms. You can now configure your SD-WAN appliance to identify alarm conditions based on your network priorities, generate alerts, and receive notifications via email, syslog or SNMP trap.

DPI support on SD-WAN Enterprise Edition Appliance

- NetScaler SD-WAN 9.3 extends support for Deep Packet Inspection to all Enterprise Edition Appliances.

Destination NAT - Integration with Forcepoint and Zscaler for Firewall Traffic Redirection

- In NetScaler SD-WAN 9.3, you can redirect internet (http and https) traffic from an SD-WAN appliance at the enterprise edge to the Forcepoint cloud-hosted security module through the Firewall redirect (transparent proxy by Destination NAT) feature. You can redirect HTTP traffic from port 80 to port 8081, and HTTPS traffic from port 443 to port 8443 of the nearest Forcepoint cloud proxy server. For SD-WAN 9.3, only firewall redirect feature has been implemented.

- NetScaler SD-WAN appliances can connect to Zscaler cloud network through GRE tunnels and IPsec tunnels at the customer's site. When implementing Zscaler using SD-WAN appliances, the following functionality is supported:

- GRE traffic forwarding mode only to Zscaler, enabling direct Internet breakout.
- IPsec Tunnel traffic forwarding to Zscaler.
- Support for direct internet access (DIA) using Zscaler on a per customer site basis.
- On some sites, you may want to provide DIA with on-premises security equipment and not use Zscaler.
- On some sites, you may choose to backhaul all traffic another customer site and provide internet access.
- Virtual Routing and Forwarding deployment.
- One WAN link as part of internet services.

FIPS Compliance Mode Using NetScaler SD-WAN GUI

In NetScaler SD-WAN 9.3, FIPS mode enforces configuring FIPS compliant settings for IPsec Tunnels and IPsec settings for Virtual Paths.

- Displays the FIPS compliant IKE Mode.
- Displays FIPS Compliant IKE DH Group for users to select the required parameters to use when configuring the appliance in FIPS compliant mode (2,5,14 – 21).
- Displays the FIPS compliant IPsec Tunnel Type in IPsec settings for Virtual Paths

    a. IKE Hash and (IKEv2) Integrity mode, IPsec auth mode.

    b. Performs audit errors for FIPS based Lifetime Settings.

Email Authentication Support

- In SD-WAN Center 9.3 release, along with configuring the email settings you can also configure SMTP Authentication.

Event SNMP/Syslog Support

- You can now configure notification settings to receive event alerts by email, SNMP traps or Syslog messages on SD-WAN Center.

SD-WAN Center Dashboard Improvements

- The SD-WAN Center dashboard is updated to include HDX visibility, the following widgets are included:

- Network HDX: Quality Summary
- Network HDX: Users and Sessions
- Network HDX: Bottom 5 Poor Sites
- Site HDX: Users
- Site HDX: Sessions
- Site HDX: QoE

SD-WAN Center HDX Insight

- You can view the Quality of Experience (QoE) of HDX applications at each site along with other HDX statistics as a report in SD-WAN Center.

Configuring Zero Touch Deployment in SD-WAN Center Using Proxy Settings

- You can configure proxy settings for Zero Touch Deployment to function properly in SD-WAN Center, if it is connected to the internet through a proxy server.

Metering and Standby WAN Links

NetScaler SD-WAN supports enabling metered links, which can be configured such that user traffic is only transmitted on a specific Internet WAN Link when all other available WAN Links are disabled.

Metered links conserve bandwidth on links that are billed based on usage. With the metered links you can configure the links as the Last Resort link, which disallows the usage of the link until all other non-metered links are down or degraded. Set Last Resort is typically enabled when there are three WAN Links to a site (i.e. MPLS, Broadband Internet, 4G/LTE) and one of the WAN links is 4G/LTE and may be too costly for a business to allow usage unless it is absolutely necessary. Metering is not enabled by default and can be enabled on a WAN link of any access type (Public Internet / Private MPLS / Private Intranet). If metering is enabled, you can optionally configure the following:

- data cap.
- billing frequency (weekly/monthly).
- start date of the billing cycle.
- active heartbeat interval
  - interval at which a heartbeat message is sent by an appliance to its peer on the other end of the virtual path when there has been no traffic (user/control) on the path for at least a heartbeat interval.
  - configurable values: default 50ms/1s/2s/3s/4s/5s/6s/7s/8s/9s/10s.

Simplified Configuration

- In NetScaler SD-WAN 9.3 release, the configuration editor is further simplified. The simplified basic configuration mode has two views Global and Site.

Using the **Global** tab, you can:

- Set the global virtual WAN network encryption settings.
- Create multiple WAN Link Templates and map it to Service Providers.
- Create WAN Link Template for MPLS links.
- Configure the WAN Link speeds in Mbps or Kbps.
- Set up MPLS Queues using % or kbps.

The new updates in the **Sites** tab are:

- Enable site as intermediate node or enable dynamic virtual path.
- Clone Sites

Change Management

- The change management UI is updated to make it easier and faster for the user to perform a change management operation.

- Once touch start feature is introduced in SD-WAN 9.3 release, this allows you to easily and quickly configure your SD-WAN appliance as a Client on first time start up.

Configuration Rollback

- The Configuration Rollback feature allows the Change Management system to detect and recover from certain software / configuration errors by reverting to the previously active software/configuration.. This feature can detect network outage and appliance crash.

Single Step Upgrade

- In release 9.3, a single step upgrade package using the SD-WAN GUI change management option to upgrade non-SD-WAN components in the network for all applicable platform editions has been introduced. The MCN distributes all necessary software components to the sites (Branch) in the network. When the branch site receives the upgrade component files, these can be installed at scheduled time intervals. If the scheduled time is not specified, a default time which is set by MCN for all branches is used for installation.

API Reference - NITRO API

NITRO APIs (REST APIs) have been introduced in NetScaler SD-WAN release 9.3. NITRO APIs can be used for third-party software integration. NITRO APIs are introduced for Change Management, Local Change Management, and few more functionalities. Detailed API documentation is available in the product installation. APIs can be downloaded from the NetScaler SD-WAN GUI by navigating to **Configuration** > **Appliance Settings** > **NITRO API** and click **Download Nitro API Doc**.

TLS 1.2 support

- In NetScaler SD-WAN WANOP 9.3, to enable secure access with SSL tunnel, the latest SSL protocol TLS 1.2 is used in SSL proxy. You can choose to use TLS1.2 protocol only or use TLS1.0, TLS1.1 and TLS1.2 protocols.

- SSL protocols SSL v3 and SSL v2 are no longer supported.

In release 9.3, two new platform editions are introduced.

- SD-WAN WANOP 4100
- SD-WAN WANOP 5100

Deploying SD-WAN VPX-SE appliance in AWS and Microsoft Azure

- NetScaler SD-WAN 9.3 release supports deploying high availability Standard Edition VPX appliances in AWS and Microsoft Azure environments.

Deploying High Availability SD-WAN VPX Appliances in Linux-KVM Hypervisor

- NetScaler SD-WAN 9.3 release supports deploying high availability VPX appliances in Linux KVM Hypervisor environments. It also supports deploying high availability VPX appliance instances on the same host.

Support for deploying SD-WAN VPX appliances in Microsoft Azure, Hyper-V 2012 R2, AWS, VMWare ESXi, and XenServer

- NetScaler SD-WAN 9.3 release extends support for deploying SD-WAN VPX appliances in Microsoft Azure, Hyper-V 2012, AWS, VMWare ESXi, and XenServer environments.

Scalability - Routing Domain Support

- In NetScaler SD-WAN 9.3, support for more number of routing domains and associated attributes has been added. See the list below for the supported parameters:

- Maximum Routing Domains: 255
- Maximum Access Interfaces per WAN Link: 64
- Maximum BGP neighbours per site: 255
- Maximum OSPF area per site: 255
- Maximum Virtual Interfaces per OSPF area: 255
- Maximum Route Learning import filters per site: 512
- Maximum Route Learning export filters per site: 512
- Maximum BGP routing policies: 255
- Maximum BGP community string objects: 255

XenServer 6.5 Upgrade Support - NetScaler SD-WAN Appliances

- XenServer versions on most appliances shipped with SD-WAN software version older than 9.1.1 are supported with XS 6.0

or 6.2. For SD-WAN 9.3 software release, you need to upgrade XenServer from 6.0 or 6.2 to 6.5. All SD-WAN appliances support this upgrade. To do the XS 6.5 upgrade, you will need to run SD-WAN software version 9.0 or newer version if you are running an older version than 9.0. After an upgrade to 9.0, it is recommended to upgrade the software to version 9.3 or latest version.

Note: Any appliance shipped with version 9.1 and later will have XS 6.5 support already.

For more information about each of these supported features, see the topics listed on the left navigation panel.

FAQs

May 16, 2018

> **Note**
>
> The WANOP, SVM, and XenServer supplementals / hot fixes are referred to as OS components.

Should I use *.tar.gz,* or single step upgrade *.zip* package to upgrade to 9.3.x from my current version (8.1.x, 9.1.x, 9.2.x)?

Use the *.tar.gz* files of the concerned platforms to upgrade the SD-WAN software to 9.3.x. After the SD-WAN software is upgraded to 9.3.x version, perform change management using the *.zip* package to transfer/stage OS component software packages. After activation, the MCN transfers/stages OS components for all the relevant branches.

After upgrading to 9.3.0 using single step upgrade package (.zip file) do, I need to perform *.upg* upgrade on each appliance?

No, OS software update/upgrade will be taken care by the single step upgrade *.zip* package and it is installed as per the scheduling details provided by you in the Change Management Settings of the respective sites.

Why should I use *.tar.gz* followed by *.zip* package to upgrade from pre-9.3 to 9.3.x, and why not directly use *.zip* package of 9.3.x?

Single Step upgrade package is supported from release 9.3.0 build 161 onwards. On earlier release versions (prior to 9.3) this package is not recognized. When the single step upgrade *.zip* package is uploaded into the Change Management inbox, the system throws an error stating that the package is not recognized. Hence, first upgrade the SD-WAN software to release 9.3 or above version and then perform Change Management using the *.zip* package.

How will the OS Components be installed through single step upgrade, if *.upg* upgrade is not performed?

The MCN will transfer/stage OS components software packages based on the appliance model, after the Change Management is completed using single step upgrade *.zip* package. After activation, the MCN starts transferring/staging the OS components software packages for the branches that need them for the scheduled update/upgrade.

How do I install OS components, without scheduling for later installations?

Set the **Maintenance Window** value to '**0**' for instant installation of the OS components.

> **Note**
>
> The installation starts only when the appliance has received all the package that is needed for the site, even when **Maintenance Window** value is set to '**0**'.

What is the use of scheduling installation? Can I use schedule instructions to upgrade SD-WAN alone?

Scheduled installation was introduced in SD-WAN 9.3 release and is applicable for OS components only and not for SD-WAN software upgrade. With single step upgrade, you need not log into each appliance to perform OS components upgrade and the scheduling option allows you to schedule the OS components installation at a different time other than SD-WAN software version upgrade.

Why does the scheduling information in Change Management Settings page displays past schedule date by default and what does it mean?

The **Change Management Settings** page displays the default scheduling information that is, *'"start": "2016-05-21 21:20:00," "window": 1, "repeat": 1, "unit": "days"'*. If the date is a past date it means that, the scheduled installation is based on the time and other parameters like maintenance window, repeat window, and unit and not the date.

What is default schedule installation date/time set to, is it generic or local appliance dependent?

By default the scheduling details is set as *'2016-05-21 at 21:20:00 (Maintenance window of 1 hour and repeated every 1 day)'*. This detail is local appliance site dependent.

How can I install OS Components immediately without waiting for the maintenance / scheduled window?

Set the **Maintenance Window** value to '**0**' in **Change Management Setting** page, this overrides the scheduled installation time.

Which package should I use for upgrade when current software version is 9.3.x or above?

Use single step upgrade *.zip* package to upgrade to any higher versions when the current software version is 9.3.x or above.

When does the OS Components files get transferred/staged to the branches?

The OS components files are transferred/staged to relevant branches after the activation is completed when Change Management is done using single step upgrade *.zip* package to upgrade the system.

Which appliances receive OS Components files? Is it platform dependent or do all branches receive it.

Appliances that are hypervisor based such as, **SD-WAN – 400, 800, 1000, 2000 SD-WAN** and Bare metal **SD-WAN - 2100** running on EE license will receive OS components to upgrade.

How does scheduling work?

By default the scheduling details is set as '*2016-05-21 at 21:20:00 (Maintenance window of 1 hour and repeated every 1 day)*' and it implies that the system will check if new software is available for installation every day as repeat value is set to '**1 day**' and will have maintenance window of '**1 hour**' and the installation will get triggered/attempted (if new software is available) at **21:20:00** (local appliance time) effective from '**2016-05-21**'.

How do I get to know if the OS Components have been upgraded?

In the **Status** column, you can see a green tick mark. On hovering over it, you see '***Upgrade is Successful***' message.

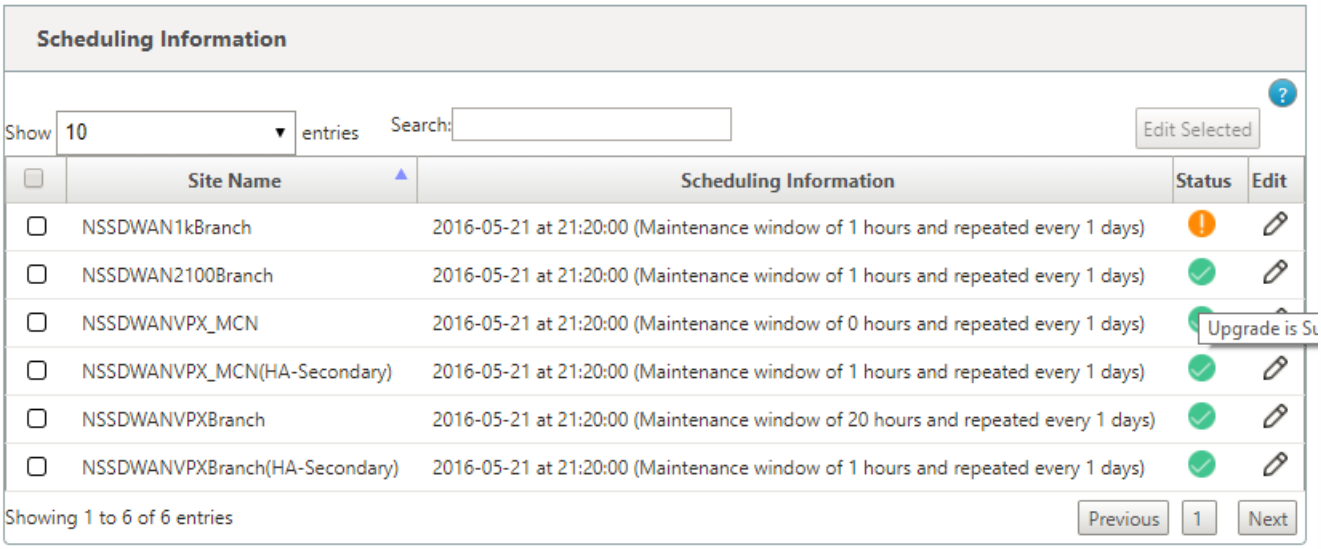

Can I use *tar.gz* file to upgrade to next release, when single step upgrade was used for previous software upgrade?

You can use *tar.gz* file to upgrade, but it is not recommended because you would need to perform *.upg* upload to upgrade OS component software by logging into each applicable appliance. From release 9.3 version 1 onwards the '***Update Operating System Software***' page is depreciated, hence you would need to perform change management with *.zip* package to upgrade the OS components.

How can we validate the current running versions of OS Components?

Currently you cannot validate the current running versions of OS components from the SD-WAN GUI. You need to login from each console or get STS to view this information.

What difference it would make if I have bare metal appliances in my network? Does scheduling impact bare metal / Virtual appliances?

Bare Metal appliances such as, **SD-WAN – 410,2100,4100,5100 SD-WAN** run only SD-WAN software and they do not need any OS components packages. These platforms are treated on par with SD-WAN VPX in terms of software need. The MCN will not transfer OS components packages to these appliances. Setting scheduling information will not take effect for these appliances, as they do not have any OS components that need upgrade.

How does Single Step Upgrade work in HA environment / deployment?

In HA deployment at MCN, we have a limitation, where the active MCN switch's/toggles the role of primary MCN during Change Management and Standby/Secondary MCN takes over. In this case, you need to perform Change Management once again with the *.zip* package on the active MCN for the packages or you can switch back to primary MCN by toggling the role of active MCN so that original primary MCN can take up the role for the OS components packages to be staged to other branches.

How does single step upgrade work in HA environment / deployment?

While performing single step upgrade in HA deployment, the role of the primary MCN and the Standby MCN is toggled. This is a limitation. If this happens, perform Change Management again with the *.zip* package on the active MCN. Alternatively, you can switch back to the primary MCN by toggling the role of the active MCN so that the original primary MCN can stage OS components packages to the branches.

Is single step upgrade support for ZTD to boot strap the appliances?

Yes, it can be used.

Can I use single step upgrade to upgrade my standalone WAN opt?

No.

Can I use single step upgrade to upgrade standalone WANOP appliance deployed in 2 Box mode?

No. Only SD-WAN appliance part of two Box mode would be upgraded and not the WANOP standalone appliance.

What is the 210-SE LTE platform edition?

210-SE LTE platform is part of the NetScaler SD-WAN family supporting Standard Edition with Integrated LTE. It provides the capability to utilize the LTE radio on the appliance as another WAN link. It includes all the Standard Edition features such as, Path selection with quick adaptation, full link bonding, Routing, Firewall, Application centric policies, and centralized management.

How many LTE modems does it support?

The 210-SE LTE platforms support one integrated LTE modem.

What are the top use cases?

1. Single Circuit Branches: 210-SE LTE is a great fit to start deploying at sites with only one existing circuit. Once deployed, the LTE could be utilized as a secondary or another active circuit for site with ease.
2. Avoid External Modems: In existing networks with an external LTE modem, 210-SE LTE can consolidate both the SD-WAN and LTE modem into one device reducing the footprint.
3. Use LTE as a backup link: An existing branch could augment the LTE connection provided by 210-SE LTE with existing circuits and act as a backup or standby link.
4. Sites with no wired circuits available: In remote locations where you cannot get a wired WAN connection, you can quickly bring the site online and make it part of your rest of your network.
5. Any remote site with the need for a wireless connection.

What are the supported bands?

210-SE LTE platforms are available in two variants to support the full global range of bands.

- R1 (Region / Radio interface 1) supports LTE-A: B1-B5, B7, B8, B12, B13, B20, B25, B26, B29, B30, B41
- R2 (Region / Radio interface 2) supports LTE-A: B1, B3, B5, B7, B8, B18, B19, B21, B28, B38-B41

When is platform be releasing?

The software release on the platform is released 9.3 version 5. It is also supported on future releases 10.0 version 2 and beyond.

What are countries that these models are qualified for at FCS?

Please check the 210-SE LTE-Selling-Guidelines on the SalesIQ for this information.

Which carriers are supported?

Please check the 210-SE LTE-Selling-Guidelines on the SalesIQ for this information.

How is this different compared to other LTE routers in the market?

The differentiation comes from the SD-WAN path selection capabilities. By employing packet-by-packet decisions, Citrix SD-WAN solution can fully bond the WAN links and enable a single flow to utilize all available bandwidth. The software also continuously monitors the network conditions for degradations and moves the traffic to utilize the best available link in real-time. When used as a primary link, the software can conserve the amount of bandwidth used by reducing overheads. When used as a backup, SD-WAN software ensures that the link is healthy and brings it into service when required based on business policies.

How is it different than the USB modems?

The USB modems used are primarily consumer grade with low reliability, minimal visibility into health of the link, and no long term supportability. In contrast, the 210-SE LTE utilizes an enterprise grade modem with high reliability, detailed visibility into network health, and long term software support.

Are dual LTE modems supported?

Yes, there is a variant of the platform to support dual modems in the future.

Is the new SD-WAN Orchestrator supported on the 210-SE LTE platform?

Yes. SD-WAN Orchestrator supports this platform when it is released.

Do these 210-SE LTE appliances have room to grow?

Yes. With upcoming software updates, these appliances can support up to 100 Mbps full duplex bandwidth. You can upgrade appliance with a software license.

Unable to log into Citrix Workspace Cloud after clicking the login tab in the Zero Touch Landing page of SD-WAN Center GUI

1. Turn on **Allow pop-ups** on the browser.
2. Check if DNS is configured in SD-WAN Center and it has internet connectivity.

Is it mandatory to discover MCN and have working MCN before proceeding with ZTD service?

Yes. It is mandatory to have working MCN before proceeding with ZTD service

Is it mandatory to Import/create a new configuration file using Config Editor to proceed with ZTD Deployment?

Yes, it is mandatory to import or create a configuration file, if using software release version earlier than 9.3.4.

What if configuration file is empty in the Prepare New Site tab?

Check that MCN discovery is successful, and verify that SD-WAN Center and MCN certificates are synced. MCN management IP address is unreachable or changed after the discovery.

After Deploying a new Site using ZTD, where is the status displayed?

1. You can navigate to the **Pending Activation** page to see the status for newly deployed site.
2. After deploying a new site, you can go through the status by opening a link provided in email sent from the ZTD Cloud Services team.

Why does the **Activation** page hang in waiting for installer state after deploying a new site?

1. This occurs when a Serial number for Device A but connected Device B is provided.
2. This can also occur when agent is not installed or able to communicate with the Cloud Service.
3. Log in to admin console and run ztd_diagnostics command for debugging.

Where can I access ZTD specific log files?

ZTD specific log files are available under the */home/sdwan/agent_logs* directory

What needs to be done if you want to reuse the same appliance on a different Customer network or RMA deployment time?

1. Perform Factory Reset.
2. Delete the appliance from **Activation History** in SD-WAN Center.
3. Activate the appliance again (with serial number and new configuration).

From ZTD perspective is there any difference in Factory shipped and RMA appliance?

No, there is no difference between factory shipped and RMA appliance.

Can a non-admin SD-WAN Center user perform ZTD?

Yes, if you are using a software release version higher than release 9.3.4.

Receiving Configuration version Mismatch error for a ZTD Deployed site

This is a known issue with SD-WAN software release versions earlier than 9.3.4. You need to upgrade MCN to release 9.3.4 and above. Then, perform ZTD deployment after upgrade is complete.

ZTD is supported on which platforms?

ZTD is supported on 210, 210 LTE, 410, 1000, 2000, and 2100 standard edition platforms.

How does ZTD work on 210-LTE Platform, when Management interface is used to reach the Cloud Service?

The workflow is similar to the existing platforms that are supported for ZTD.

If Management Port was already connected and LTE Port needs to be used for Internet connectivity, what procedures are to be performed?

1. If Management interface is configured with DHCP IP address (default option), unplug the Management port and restart the appliance.
2. If Management interface is configured with Static IP address, modify the Management port to have DHCP IP address, apply the configuration, unplug the Management port, and restart the appliance.

What is the pre-requisite that needs to be followed when LTE Ports are used for Internet connectivity in

210-SE LTE appliance?

1. Management Port should not be connected.
2. For the Branch Site, Internet Service using LTE Interface should be additionally configured in Virtual WAN Configuration.
3. SIM card with data connectivity should be available. After inserting the SIM card, 210-SE LTE appliance needs to be restarted.
4. LTE Signal Coverage - LTE IP address assignment or status can be verified using SD-WAN CLI -> LTE - Help / Status.

Licensing

Jan 04, 2018

There are three NetScaler SD-WAN Editions each with a different set or subset of NetScaler SD-WAN features. The type of license you install determines the NetScaler SD-WAN Standard Edition, WANOP, and Enterprise Edition appliances.

> **Note**
>
> When installing and applying a license, make sure that your specific appliance supports the NetScaler SD-WAN Edition you want to enable, and that you have the correct software version in place.

The following table illustrates which NetScaler SD-WAN platforms are supported for each of the available NetScaler SD-WAN software versions.

Version	WAN Optimization Edition	Standard Edition	Enterprise Edition
Release .7.X	Yes	—	—
Release .8.X	—	Yes	—
Release .9.0	—	Yes	Yes
Release 9.1	Yes	Yes	Yes
Release 9.2	Yes	Yes	Yes
Release 9.3	Yes	Yes	Yes

NetScaler SD-WAN 9.3 introduced a new set of licenses specific to the SD-WAN solution. Earlier version of licenses, including those compatible with release 7.x, are not supported with the NetScaler SD-WAN release. The existing process to obtain NetScaler SD-WAN licenses remains consistent with the CloudBridge 8.0.x, and 9.0.x releases. Once obtained, the licenses can be activated through the appliance's management web interface.

The following table lists all the appliance models supported in NetScaler SD-WAN 9.3 release:

Platform Edition	License Model
Standard Edition VPX	VPX-020-SE, 050-SE, 100-SE, 200-SE, 500-SE, 1000-SE
Standard Edition 410	410-020-SE, 410-050-SE, 410-100-SE, 410-150-SE
Standard Edition 1000	1000-020-SE, 1000-050-SE, 1000-100-SE
Standard Edition 2000	2000-100-SE, 2000-200-SE, 2000-300-SE
Standard Edition 2100	2100-200-SE, 2100-300-SE, 2100-500-SE, 2100-1000-SE, 2100-1500-SE
Standard Edition 4100	4100-1000-SE, 4100-2000-SE, 4100-3000-SE
Standard Edition 5100	5100-1000-SE, 5100-2000-SE, 5100-3000-SE, 5100-4000-SE, 5100-5000-SE
WANOP Edition VPX	VPX-2, 6, 10, 20, 50, 100, 200
WANOP Edition 800	800-002, 800-006, 800-010
WANOP Edition 1000 Windows Server	1000WS-006, 1000WS-010, 1000WS-020
WANOP Edition 1000	1000-006, 1000-010, 1000-020
WANOP Edition 2000	2000-010, 2000-020, 2000-050
WANOP Edition 3000	3000-050, 3000-100, 3000-155
WANOP Edition 4000	4000-310, 4000-500, 4000-1000
WANOP Edition 4100	4100-310-WANOP, 4100-500-WANOP, 4100-1000-WANOP
WANOP Edition 5000	5000-1500, 5000-2000
WANOP Edition 5100	5100-1500-WANOP, 5100-2000-WANOP
Enterprise Edition 1000	1000-010-EE, 1000-020-EE, 1000-050-EE, 1000-100-EE
Enterprise Edition 2000	2000-100-EE, 2000-200-EE, 2000-250-EE
Standard Edition 210	210-SE-020, 210-SE-050

VPX models allow 2, 6, 10, 20, 50, 100, and 200 Mbps bandwidth licenses. At least two 2.1 GHZ CPUs are required in order to support the VPX instances.

Before you can download the software, you must obtain and register a NetScaler SD-WAN software license. For instructions on obtaining a NetScaler SD-WAN software license, contact Citrix NetScaler SD-WAN Customer Support. Instructions for uploading and installing the license file on your appliances are provided in the section, Uploading and Installing the SD-WAN Software License File. However, before installing the license, you must first setup the appliance

hardware, and set the date and time for the appliance.

To return or reallocate a license, you must use the Citrix NetScaler SD-WAN Licensing Portal. You also have the option to use the Licensing Portal for license allocation. For instructions, see the Knowledge Base article entitled, "My Account All Licensing Tools User Guide," at this location:

http://support.citrix.com/article/ctx131110

Provisioning Licensing

Aug 09, 2017

The license procedure for provisioning licensing for SD-WAN platform editions covers the following topics:

- Supported SD-WAN license model for 9.0, 9.1, 9.2, and 9.3
- Remote License Server support for SD-WAN VPX-SE appliances
- Pre-requisites for using Remote License Server
- Use Cases
  - Deployment scenarios supported for 9.1, 9.2, and 9.3
    - Remote License server reachable in Management network (without using data/aPA Ports)
    - Remote License server in the Branch network
    - SD-WAN VPX-SE - PBR Deployment in the Branch Office
    - SD-WAN VPX-SE - Microsoft Azure
    - Two Box Deployment
    - Hairpin Mode
  - Deployment scenarios supported only on 9.3

    - SD-WAN VPX-SE - AWS

    - SD-WAN VPX-SE - Hyper V 2012 R2 and 2016, Linux-KVM Hypervisor

- Import SD-WAN VPX-SE license on XenServer/ESXi 9.3
  - Local License
  - Remote License
- Deployment scenarios not supported for 9.2 and 9.3
  - Remote License server deployed in Data Center (data/apA Ports)

License Procedure

Aug 09, 2017

Pre-requisites for using Remote License Server for SD-WAN appliances.

- NTP should be configured for both License server and SD-WAN (date and time should be in-sync)
- Remote License Server version should be 11.13.1 or earlier.

It is recommended that you use the latest License Server version:

- Release 9.2: 11.13.1 L.S
- Release 9.1: 11.13.1 L.S
- Release 9.0: 11.13.1 L.S
- Release 8.1: 11.12.1 L.S

1. Remote license server reachable through the management network without using data/apA Ports.

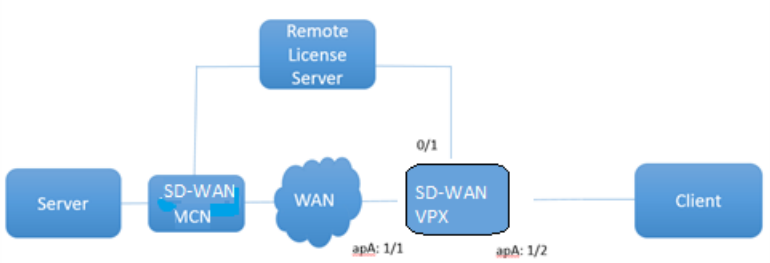

2. Remote license server in the Branch network.

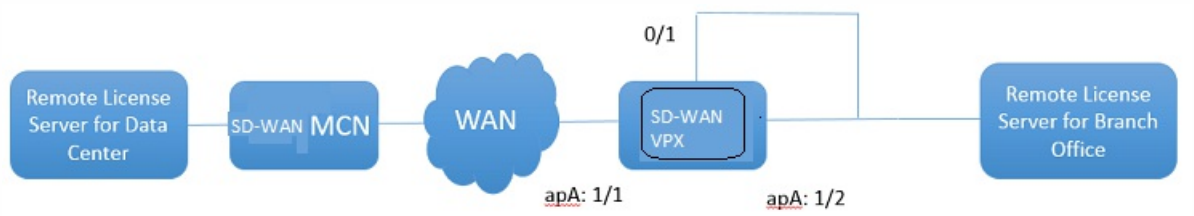

3. SD-WAN VPX-SE - PBR deployment in the Branch office.

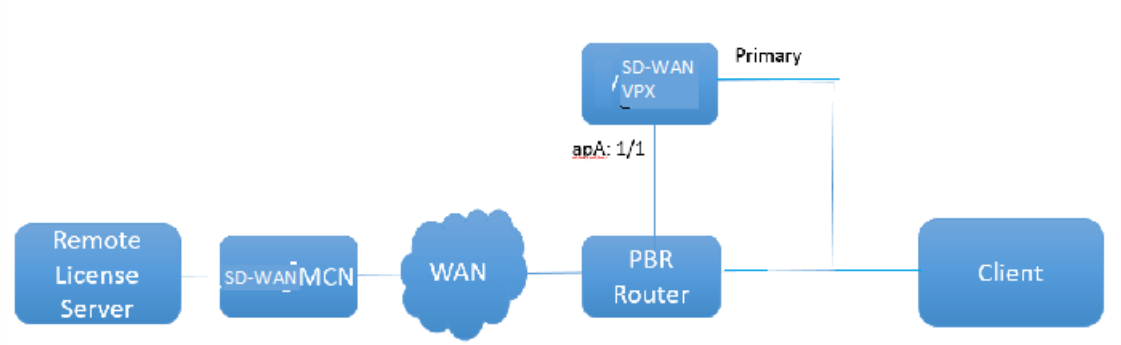

Deployment scenarios not supported for 9.2

4. Remote license server deployed in Data Center reachable through the data/apA Ports.

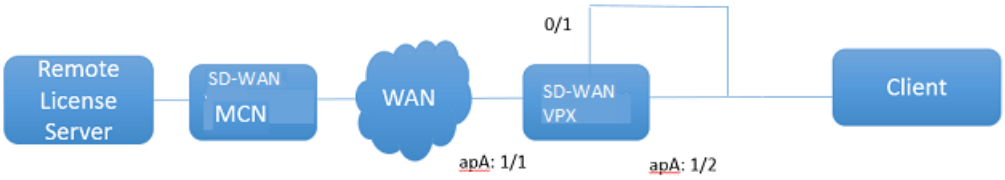

Importing SD-WAN VPX-SE license deployed on XenServer/ESXi:

1. In the SD-WAN web management interface, navigate to **Configuration > Appliance Settings > Licensing**.
2. Select **Local** and upload the License. Click **Upload and Install**.
3. Save your changes by clicking **Apply Settings**.

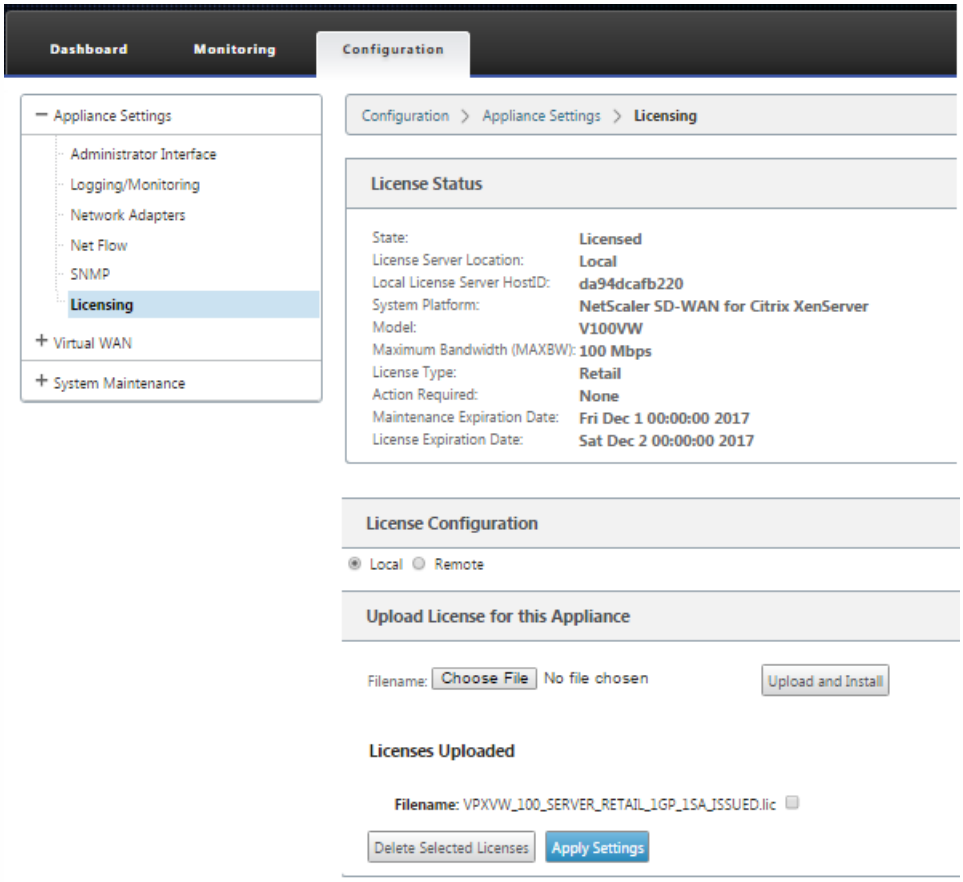

1. In the SD-WAN web management interface, navigate to **Configuration > Appliance Settings > Licensing**.
2. Select **Remote** and enter the Remote Server-IP address details.

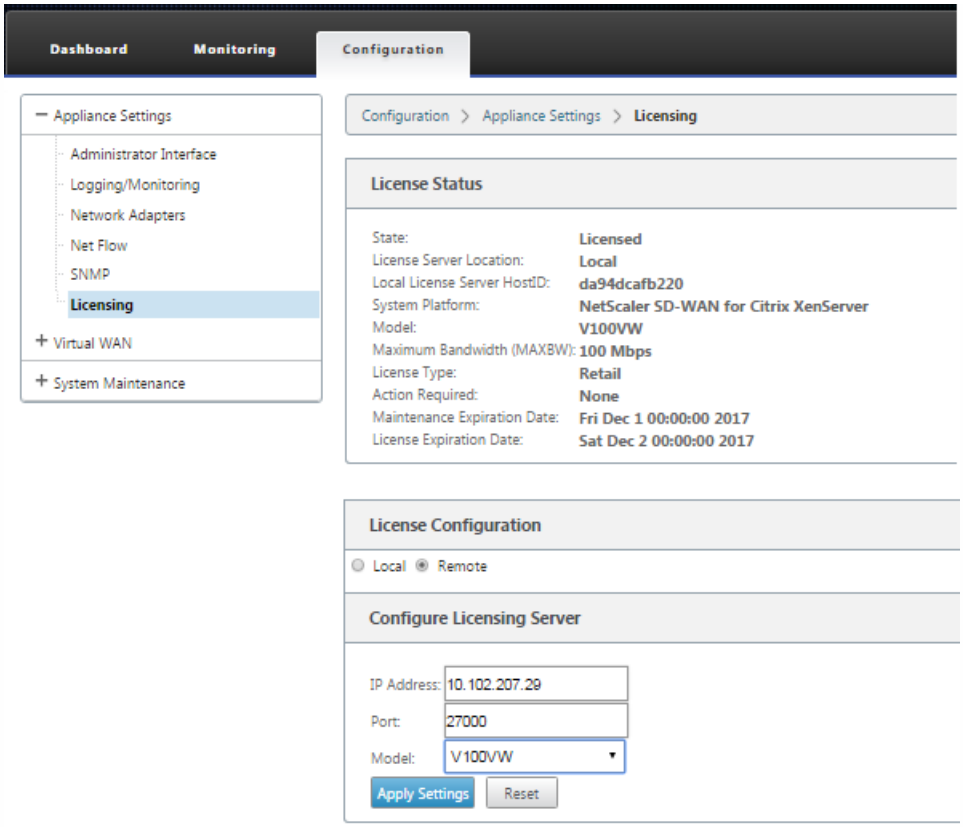

Updating and Upgrading to NetScaler SD-WAN 9.3

Nov 22, 2017

There are two main upgrade scenarios:

1. Upgrade appliances with working Virtual WAN configuration from a previous release version, 8.1, 9.0, 9.1, and 9.2 to the current version 9.3.x.

2. Upgrade appliances to release 9.3.x without existing Virtual WAN configuration.

Important

Appliances shipped with 8.0.x image are not supported to upgrade to Enterprise Edition.

Note

Upgrading to release 9.3 is a multi-step process. SD-WAN software is upgraded centrally from the MCN appliance. Operating system software can be upgraded centrally from the MCN appliance by performing Change Management using the single step upgrade .zip software package, and the same can be scheduled for later installation.

Upgrade to 9.3 With Working Virtual WAN Configuration

Jul 26, 2018

> **Note**
>
> Upgrading to 9.3 release is a multi-step process. Virtual WAN software is upgraded centrally from the MCN appliance using *tar.gz* files. Refer to the upgrade instructions in the following sections.

a. Targeted appliances for upgrade to Enterprise Edition (1000-EE or 2000-EE) are required to have:
- factory image of 9.0.0.x RTM build, if your appliance is WANOP edition which has been converted to Enterprise Edition using USB. See, Convert SD-WAN 1000 or 2000 WANOP to Enterprise Edition with USB.
- factory image of 8.1.0.x RTM build and higher.

b. Have a valid SD-WAN license.

c. Have a working Virtual WAN configuration running 8.1.x, 9.0.x, 9.1.x, or 9.2.x build with virtual paths established from MCN to the branch sites.

1. On the MCN appliance, navigate to **Configuration** > **Virtual WAN** > **Change Management**.
2. Obtain applicable *cb-vw_<APPLIANCE-MODEL>_9.3.X.tar.gz* file for all sites in the Virtual WAN network from Citrix download page for NetScaler SD-WAN Release 9.3 at: https://www.citrix.com/downloads/netscaler-sd-wan.html
3. Upload the *cb-vw-<ApplianceModel>-9.3.x.tar.gz* file for the branches defined in the configuration file for which upgrade needs to be performed. Perform **Change Management** in SD-WAN web interface for the MCN appliance and complete the change management process.

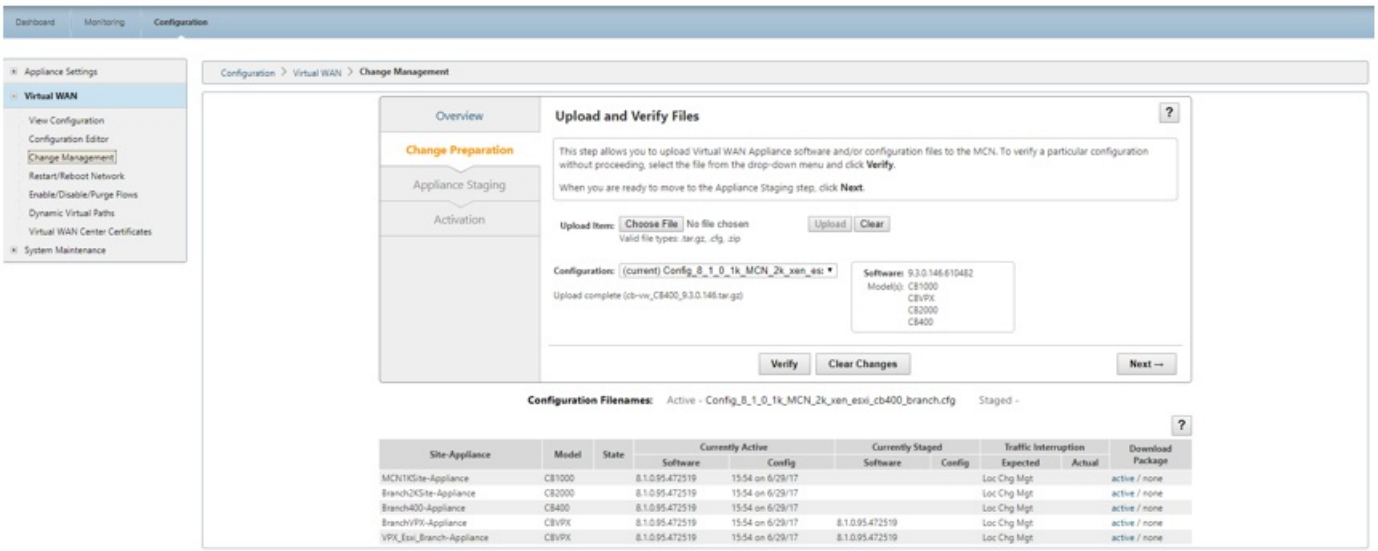

4. Click **Next** to proceed further.

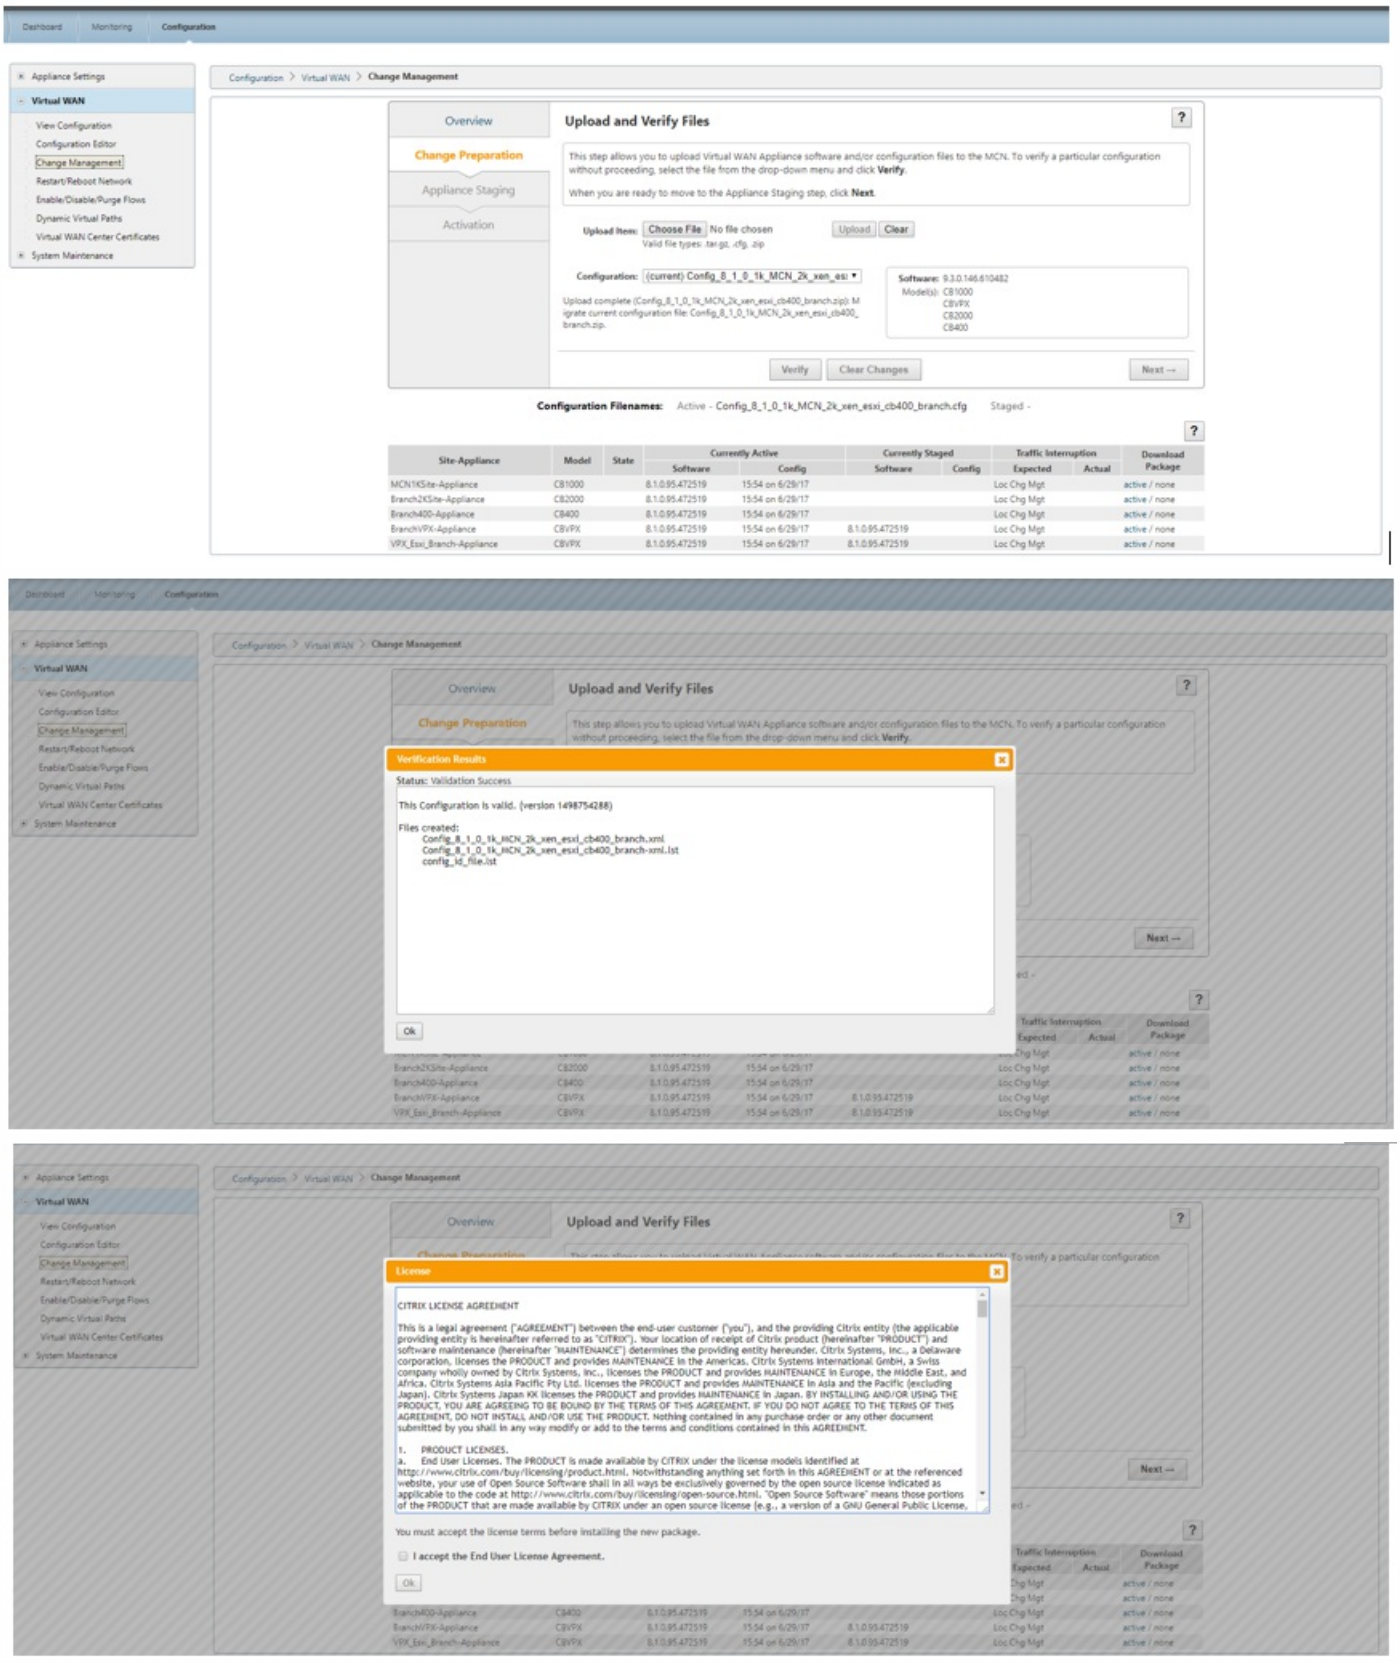

5. After accepting license agreement, you are navigated to **Appliance Staging** where appliances can be staged by clicking on **Stage Appliances**.

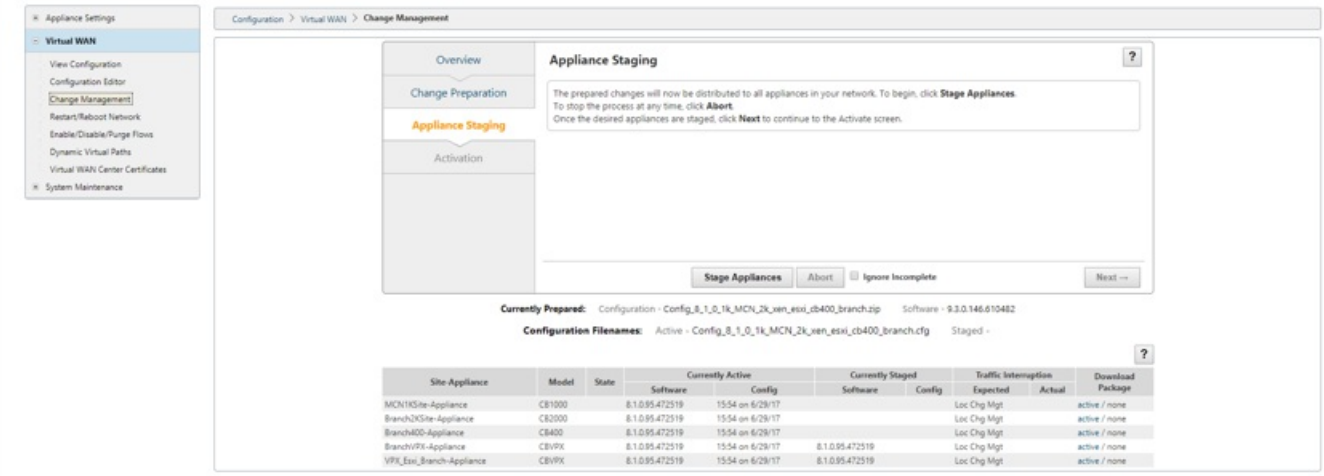

6. Transfer Progress status is displayed as part of preparing and staging the software packages to the appliances.

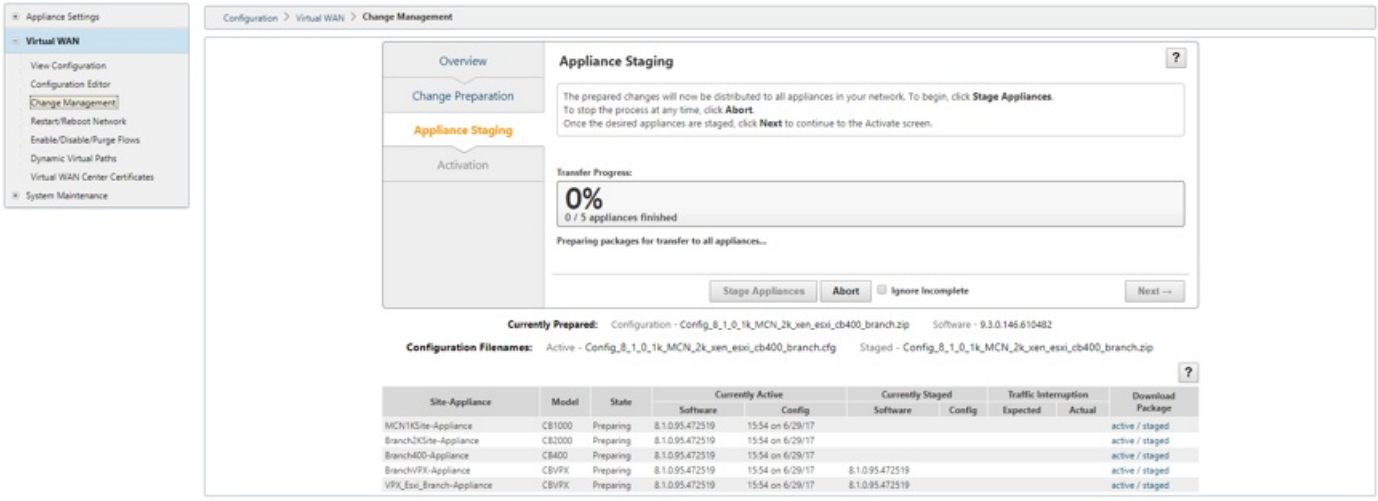

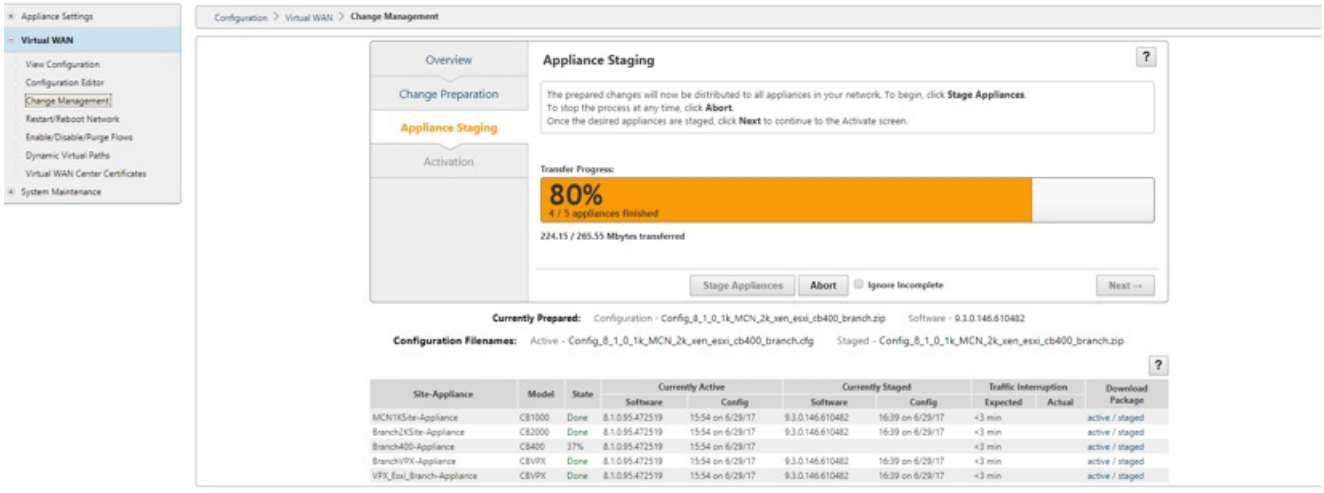

7. Click **Next** when Transfer Progress shows 100%, and button is enabled to proceed.

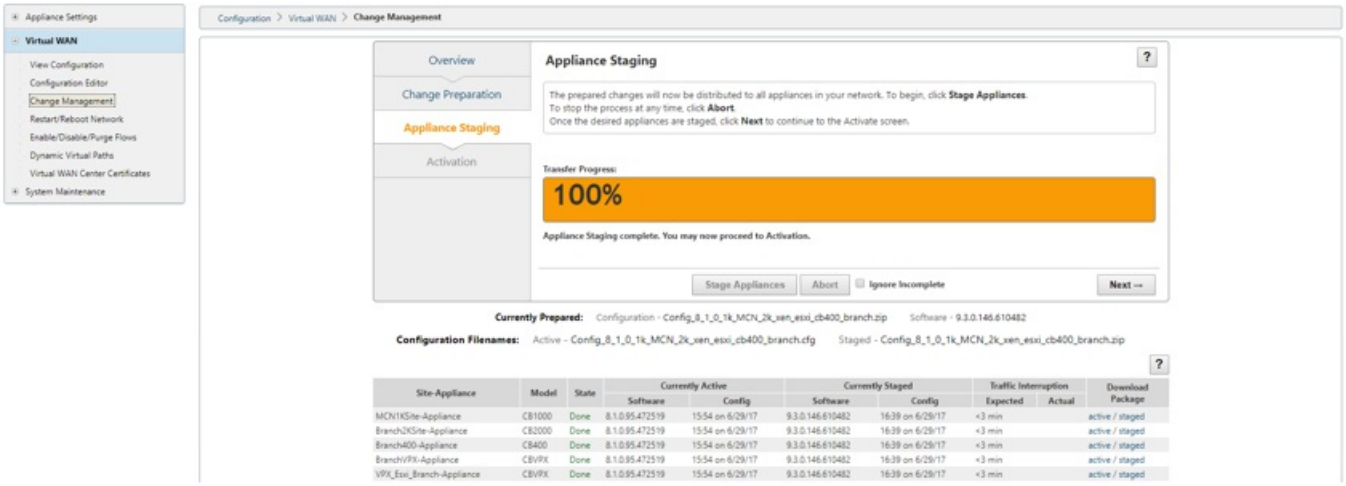

8. You are navigated to **Activation** page where you can activate staged software by clicking **Activate Staged** and confirm to start activation by clicking **OK** in pop message.

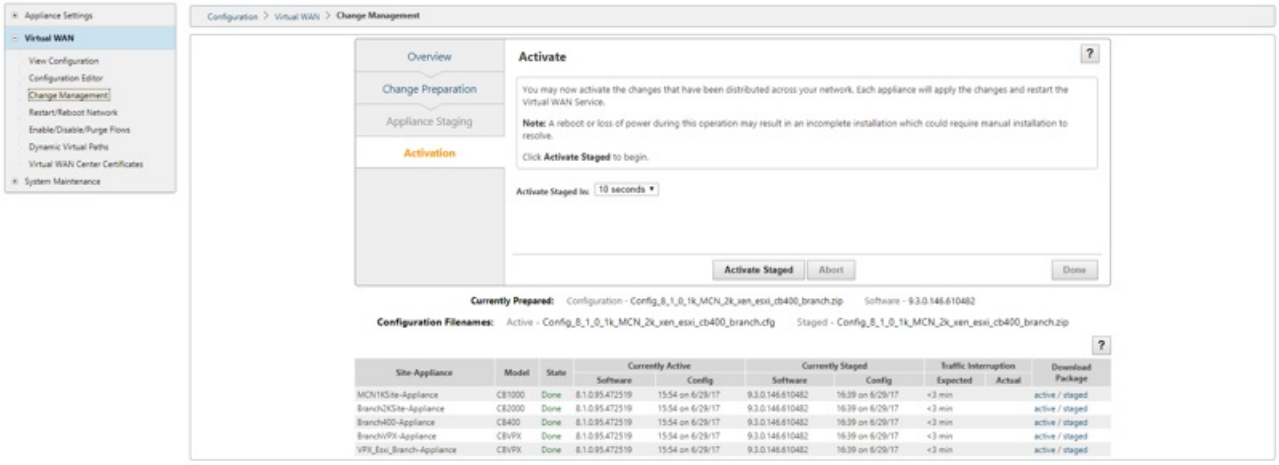

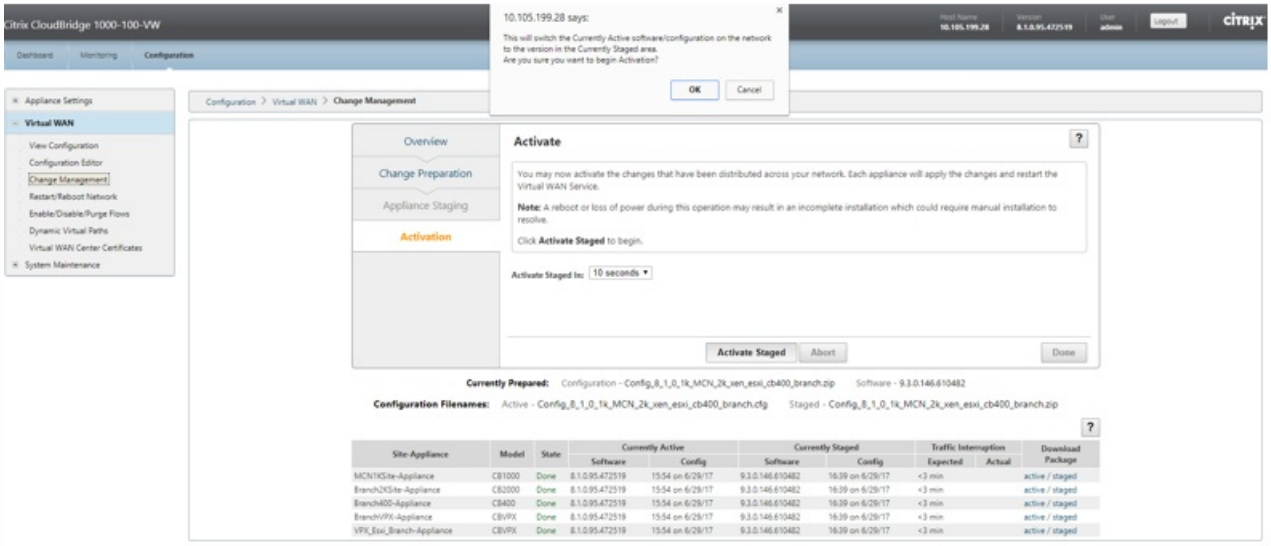

9. After completion of activation countdown of 180s click **Done** that gets enabled.

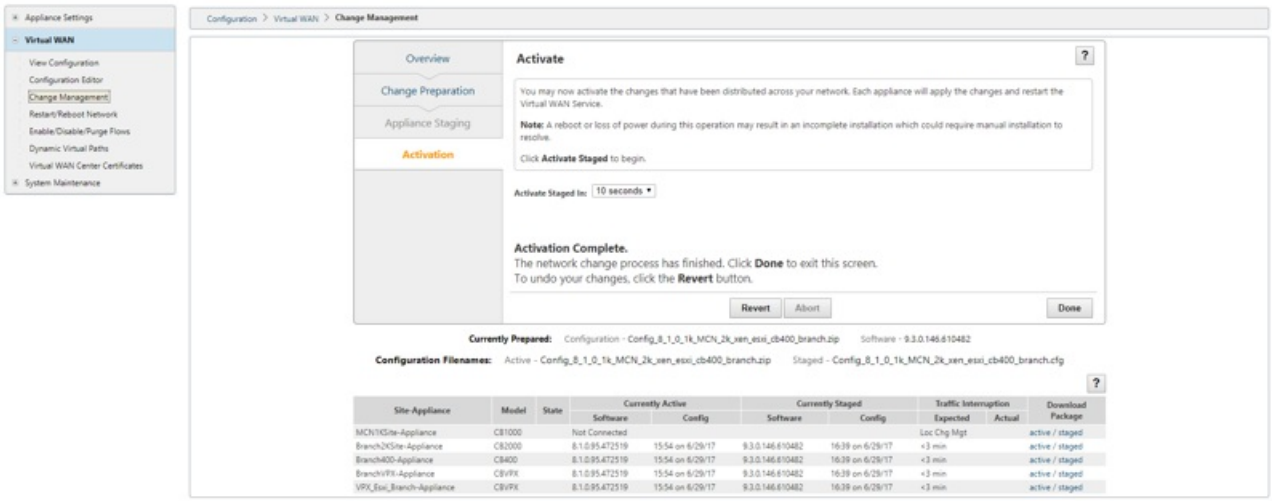

10. After the appliances are upgraded to 9.3.0 you need to perform Change Management once again this time by uploading single step upgrade package, ns-sdw-sw-9.3.0.x.zip after downloading the package from the download server.

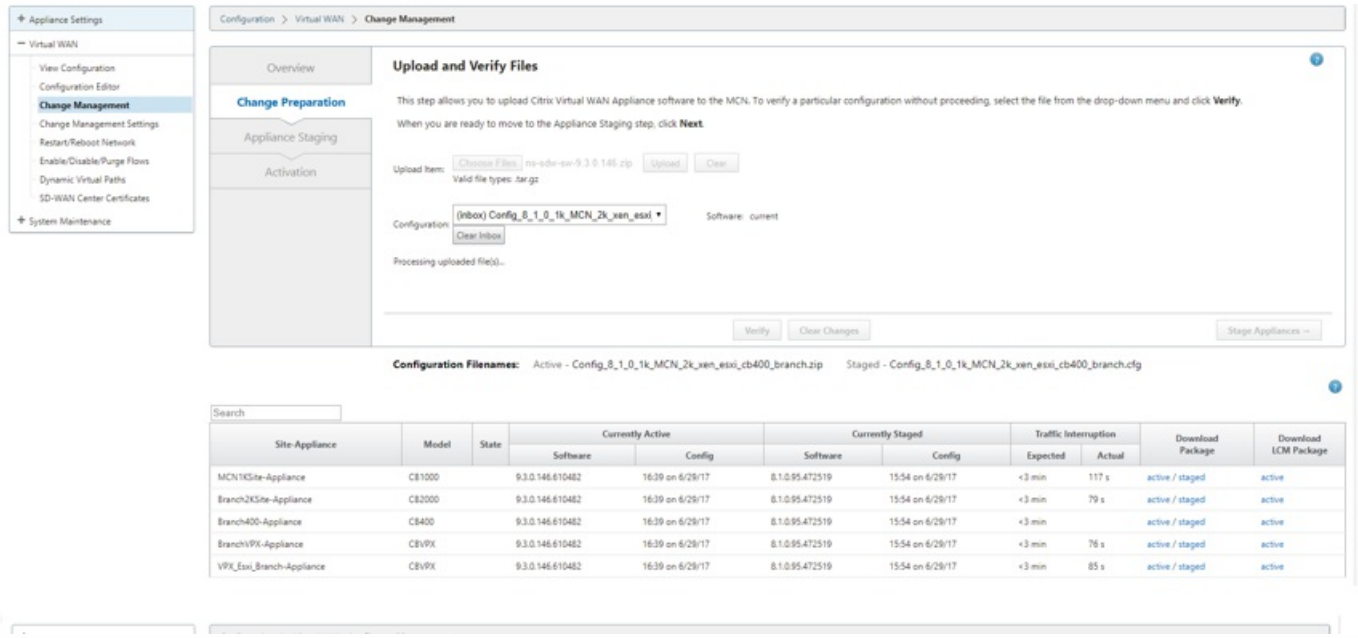

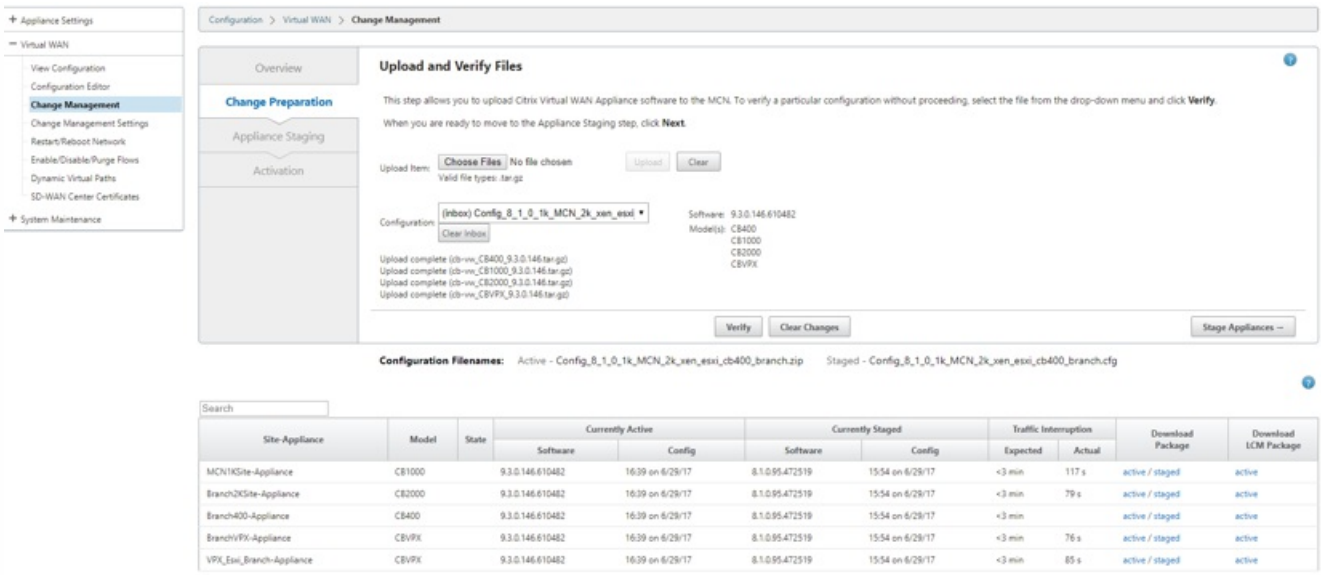

11. Click **Stage Appliances** once upload process is successful and relevant models are displayed that would be upgraded based on the configuration file that has information about each branch platform models. License agreement page pop-up for user to take action and proceed is displayed.

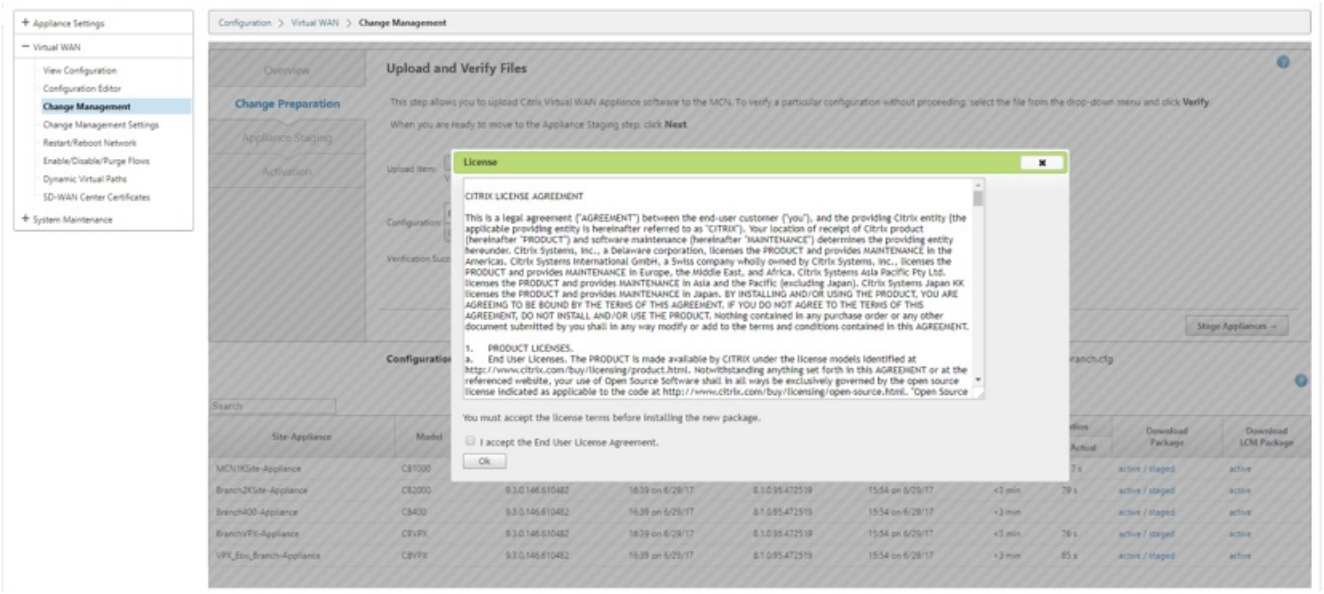

12. After accepting license agreement, you are navigated to **Appliance Staging** page which shows the status of package preparation and staging followed by transfer status for each branches.

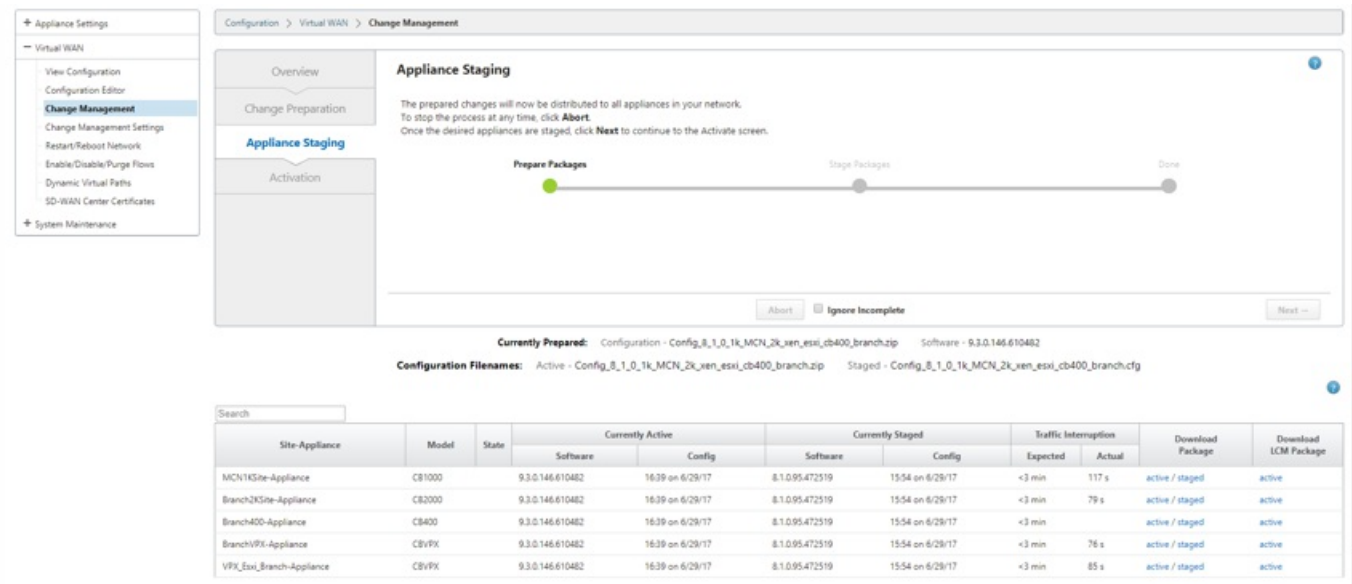

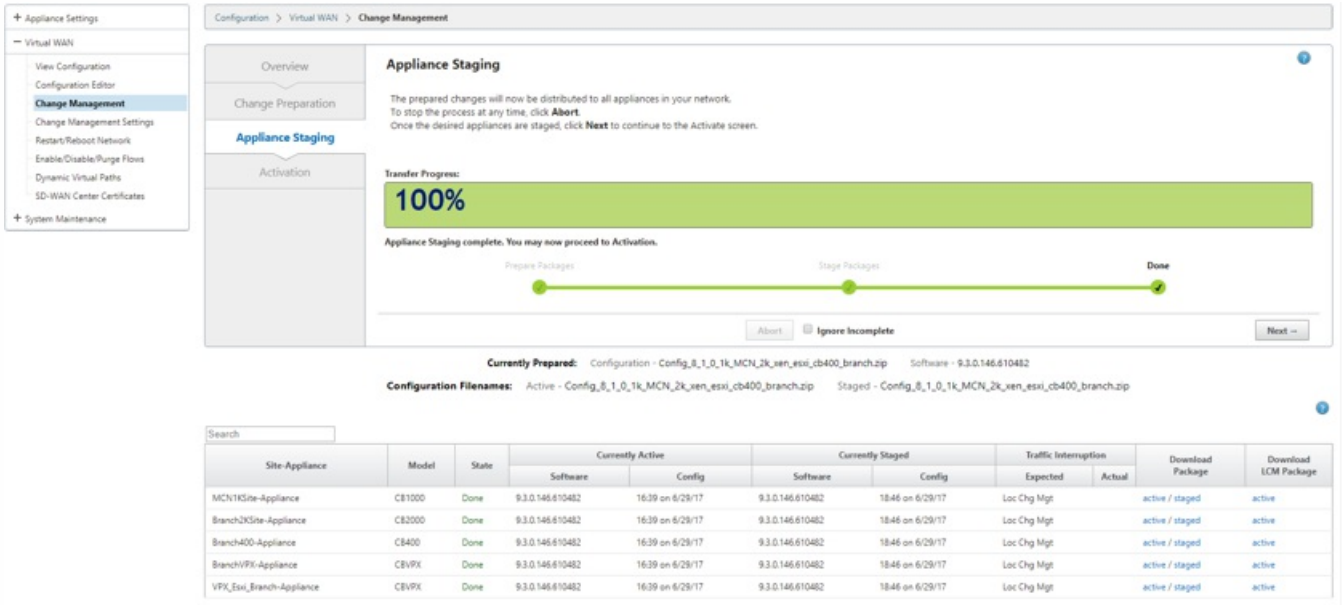

13. After completion of transfer, you are navigated to Activation page where you can click on **Activate Staged** button to active the staged software.

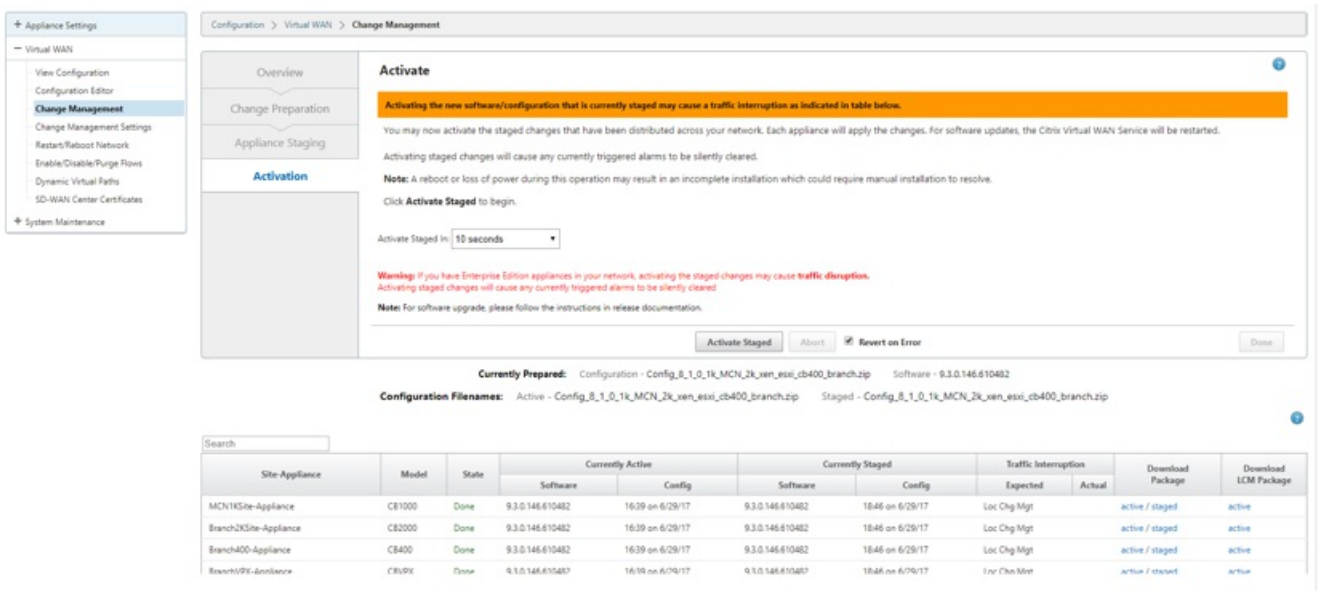

14. Click **done** once the countdown is completed and the button is enabled.

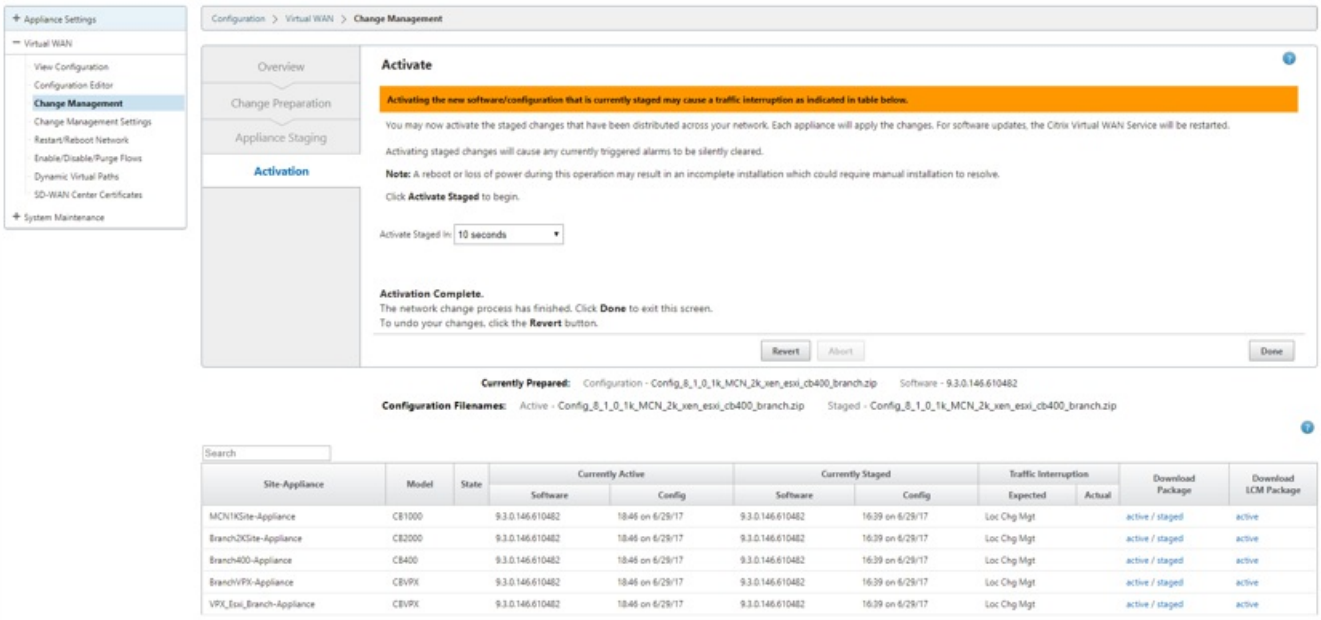

Note

Bare Metal appliances such as, SD-WAN – 410, 2100, 4100, and 5100 SD-WAN run the SD-WAN software only, and do not need any OS components packages.

15. Navigate to **Change Management** page and you can check the transfer status of WANOP, SVM , XenServer Hotfixes for applicable branches only.

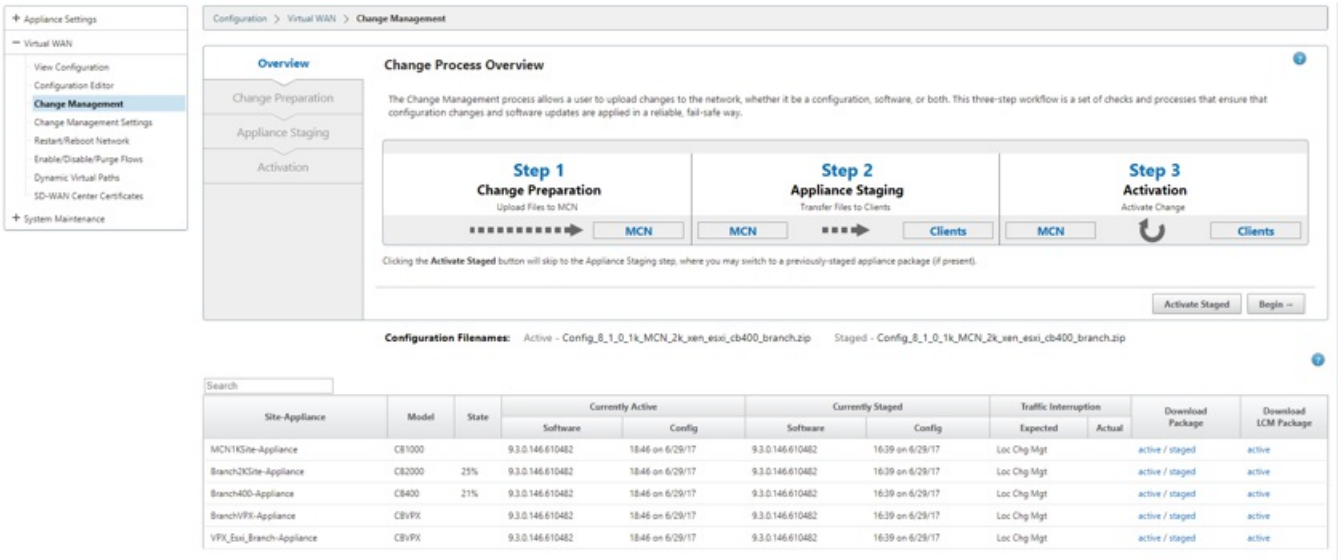

16. Navigate to **Change Management Settings** page to schedule the installation of software other than non-SDWAN like WANOP, SVM, XenServer Hotfixes. By default the MCN assigns schedules installation to be attempted every day at 21:20:00 based on software availability on the branches.

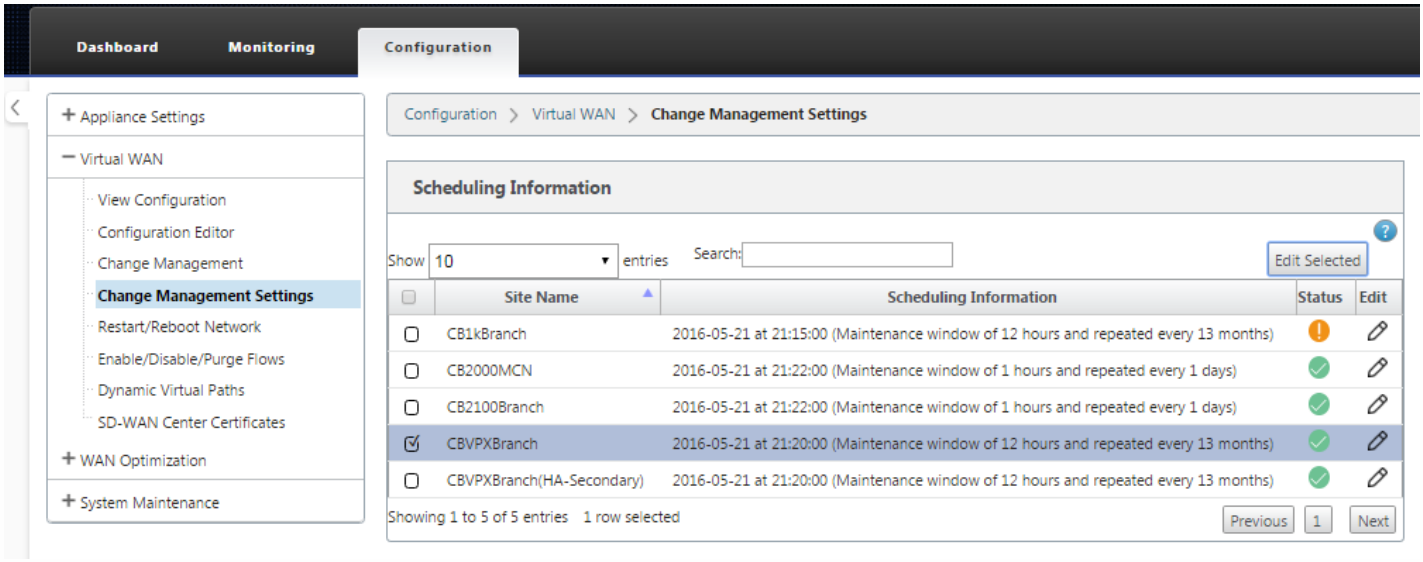

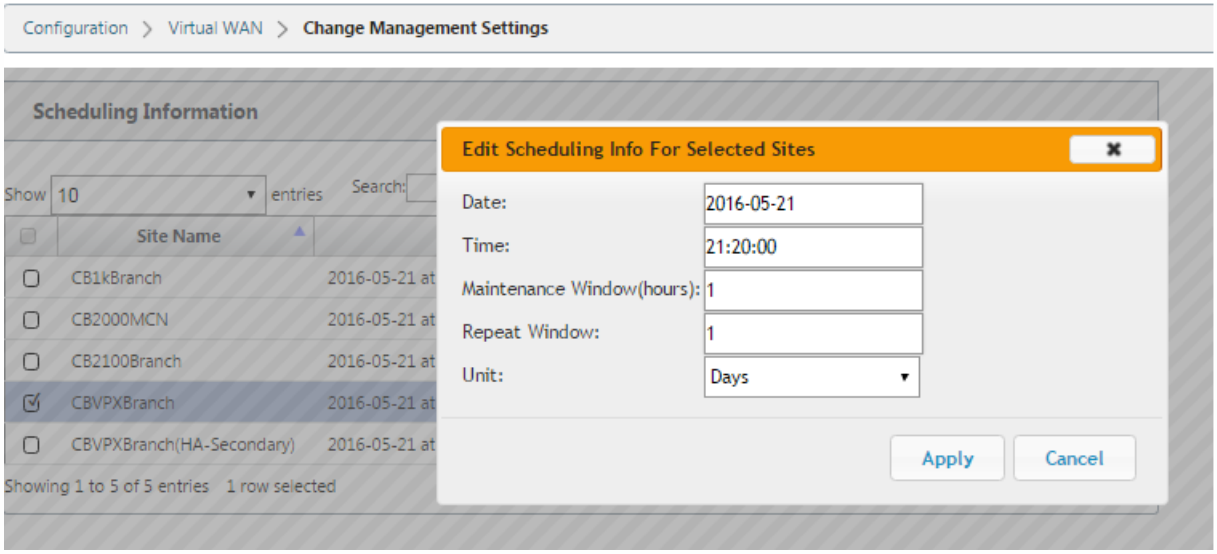

17. For detailed information or help on the scheduling information, you can click on help icon and get the information.

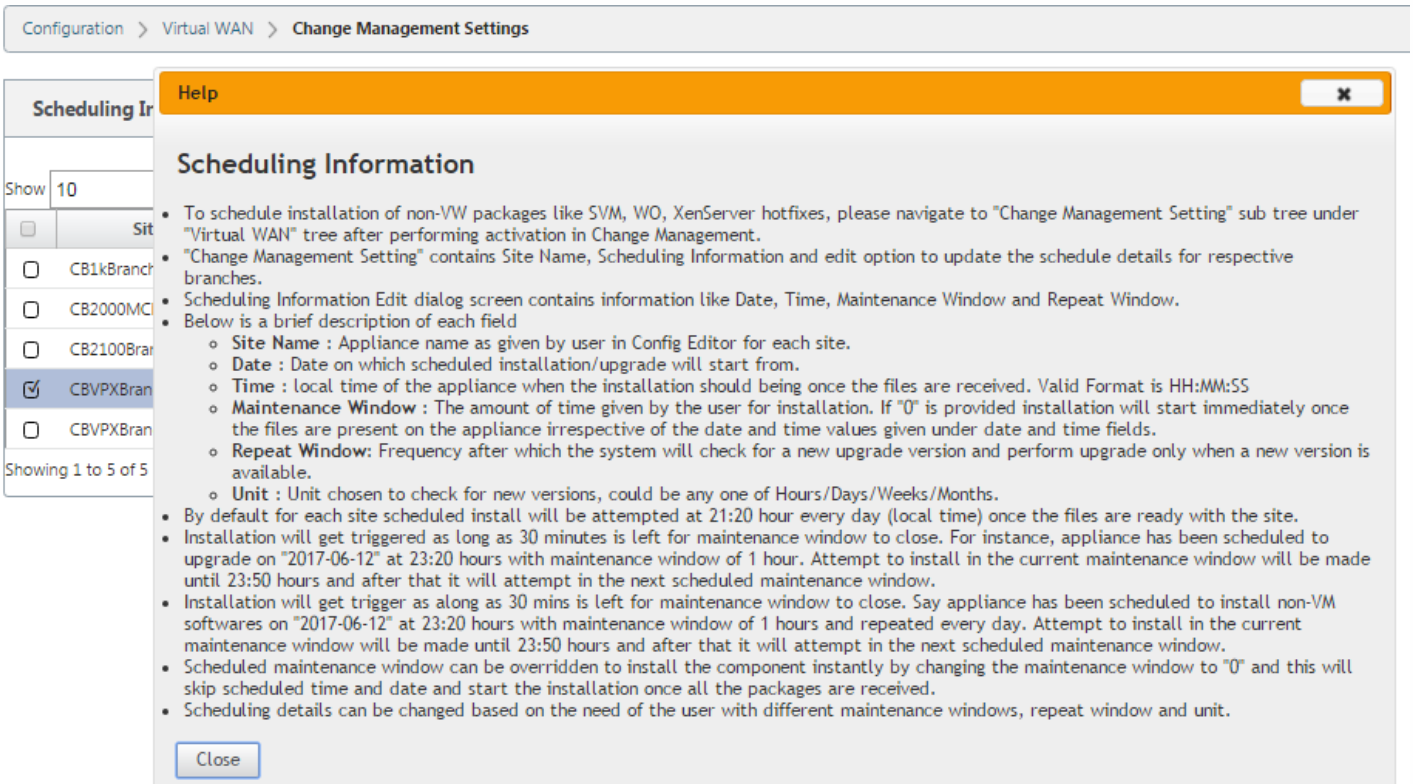

Upgrade to 9.3 Without Virtual WAN Configuration

Nov 16, 2017

> **Note**
>
> Upgrading to 9.3 release is a multi-step process. Virtual WAN software is upgraded centrally from the MCN appliance using tar.gz files. Refer to the upgrade instructions in the following sections.

> **Important**
>
> This upgrade procedure to software release 9.3 assumes that virtual paths are not established between the MCN and Branches.

How to Perform Single step Upgrade Using SD-WAN GUI

On the MCN:

- You can upload single step upgrade package (.zip) file.
- Proceed with Change Management workflow in the GUI.
- Schedule install components other than Virtual WAN software.

1. Prepare Configuration from **Configuration Editor** and **Save** configuration with valid name as shown below.

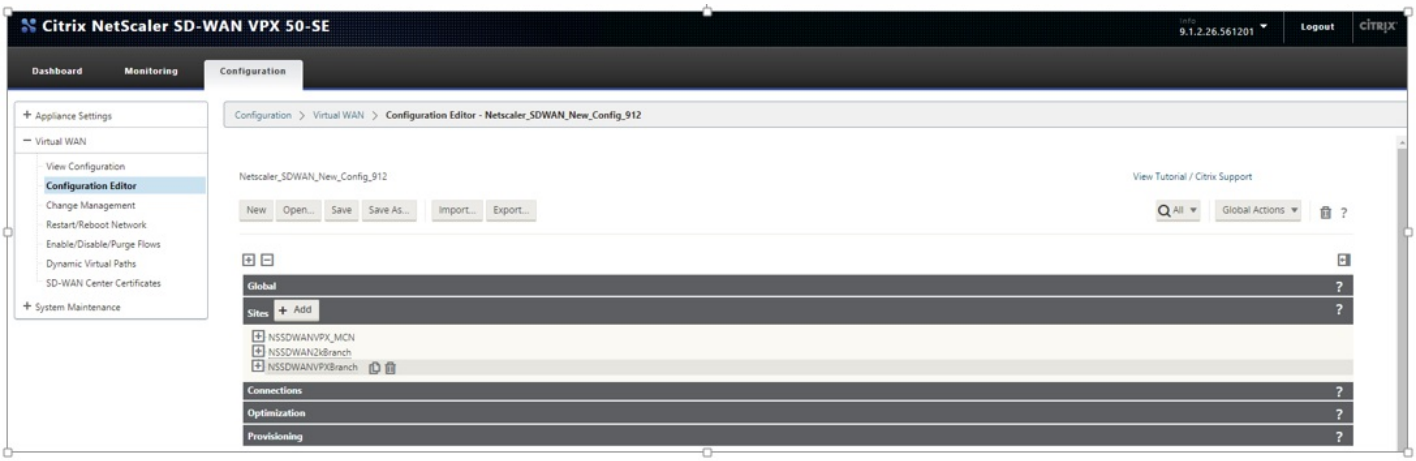

2. Export **Save Config** file to Change Management by selecting **Export** button and choose destination as **Change Management Inbox**.

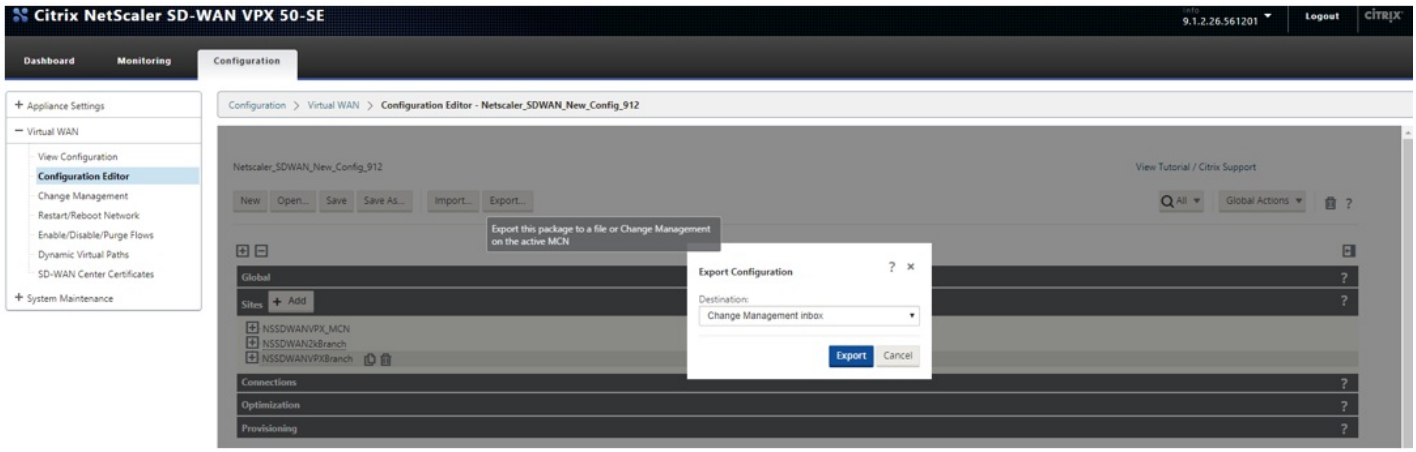

3. After clicking **Export**, click the **Change Management** hyperlink that you see to navigate to the **Change Management** page.

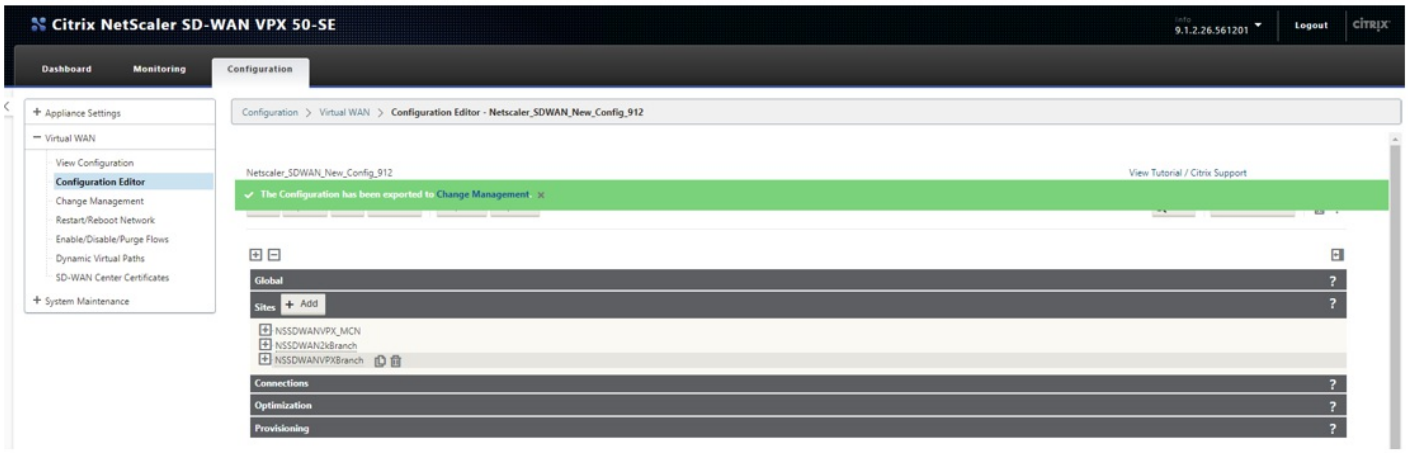

4. In the **Change Management** page, upload tar.gz software package after downloading from download server and provide the location from where to upload by selecting '**Choose Files**'.

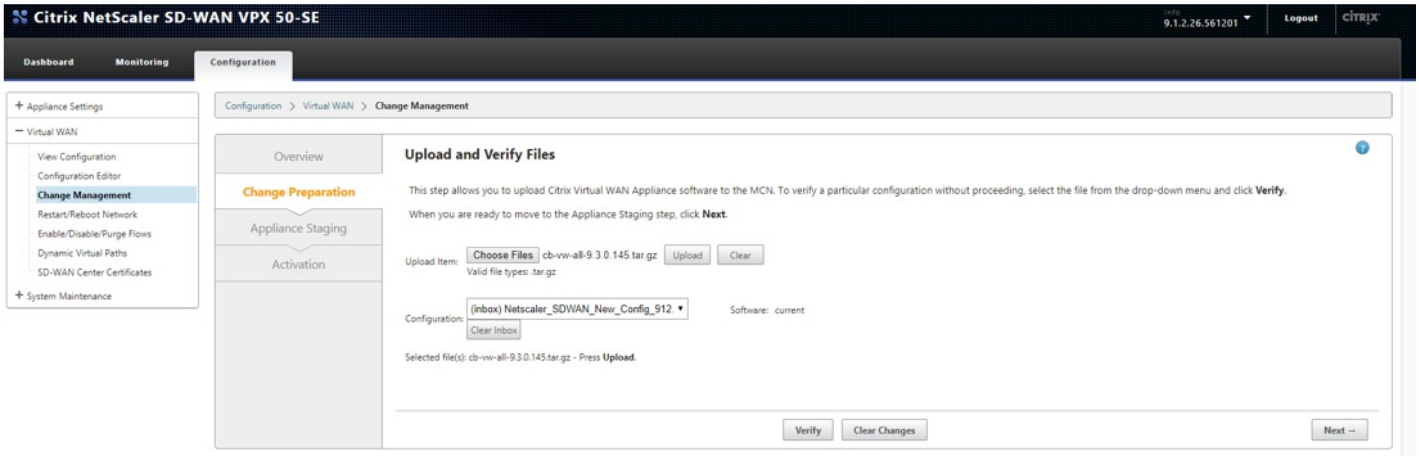

5. After selecting the file(s) to upload, click **upload** and a progress bar appears showing the current upload progress.

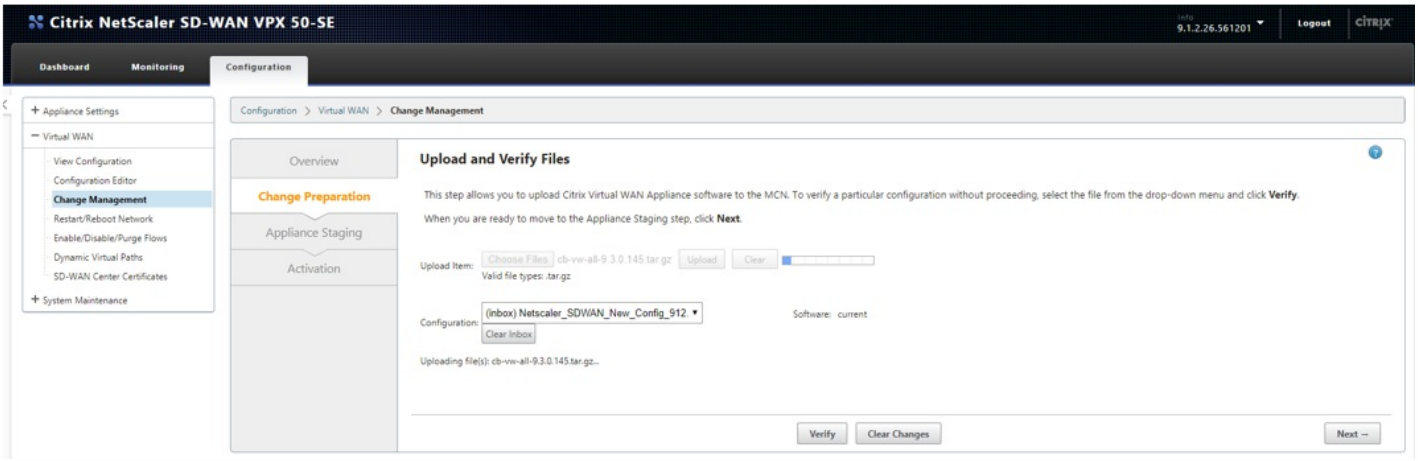

6. After upload is successful, the uploaded files are processed for any invalid errors.

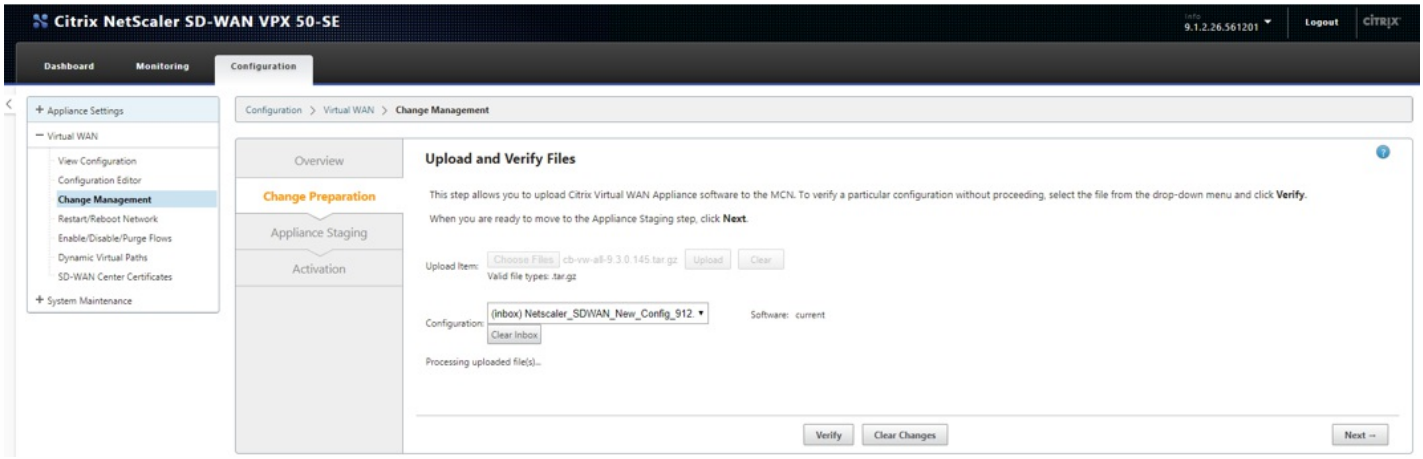

7. After upload processing is successful , information about for which platform models the software has been uploaded and processed is displayed.

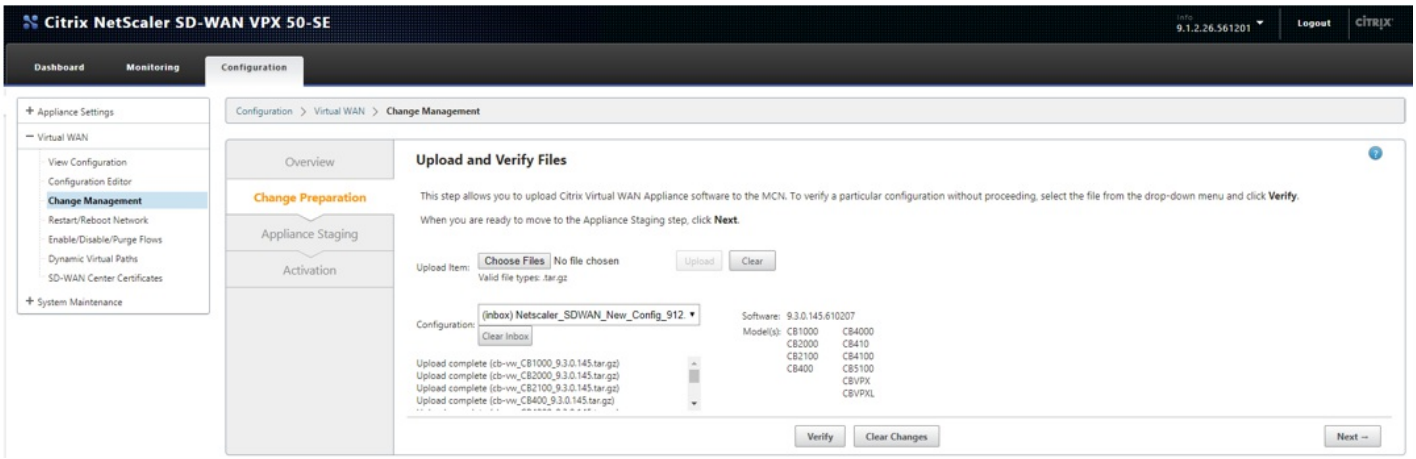

8. Click **Next** to proceed with Validation of Configuration file and appropriate message is shown based on the validation results and accept and proceed as displayed.

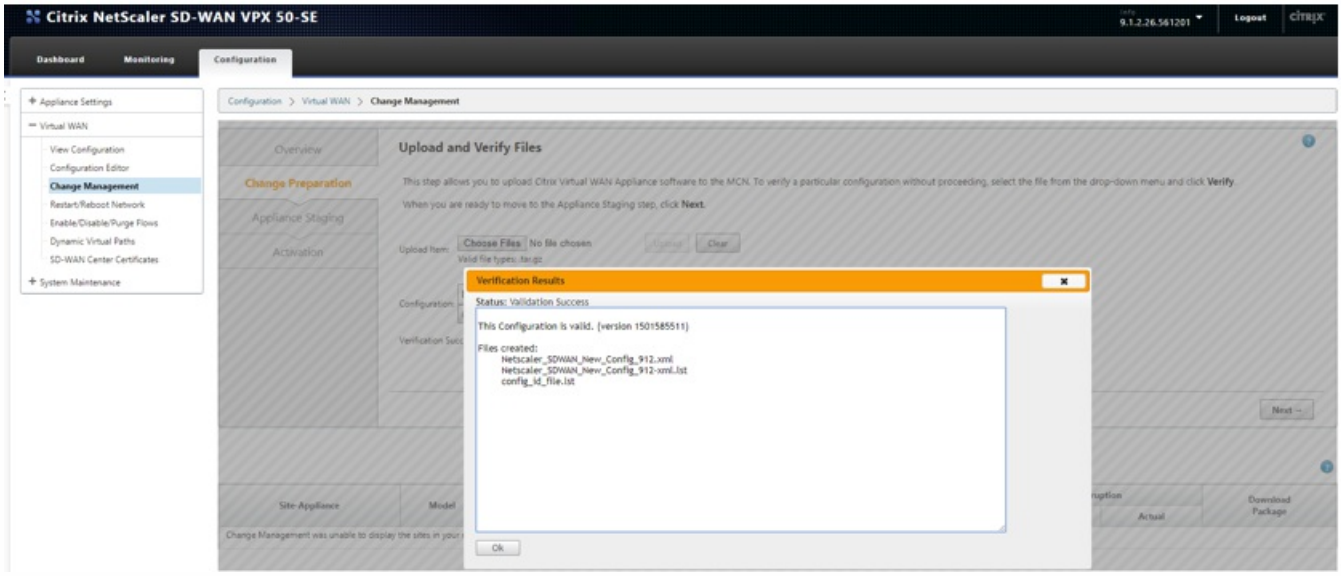

9. After accepting validation result, the License agreement page for user acceptance is displayed. Click on accept and proceed.

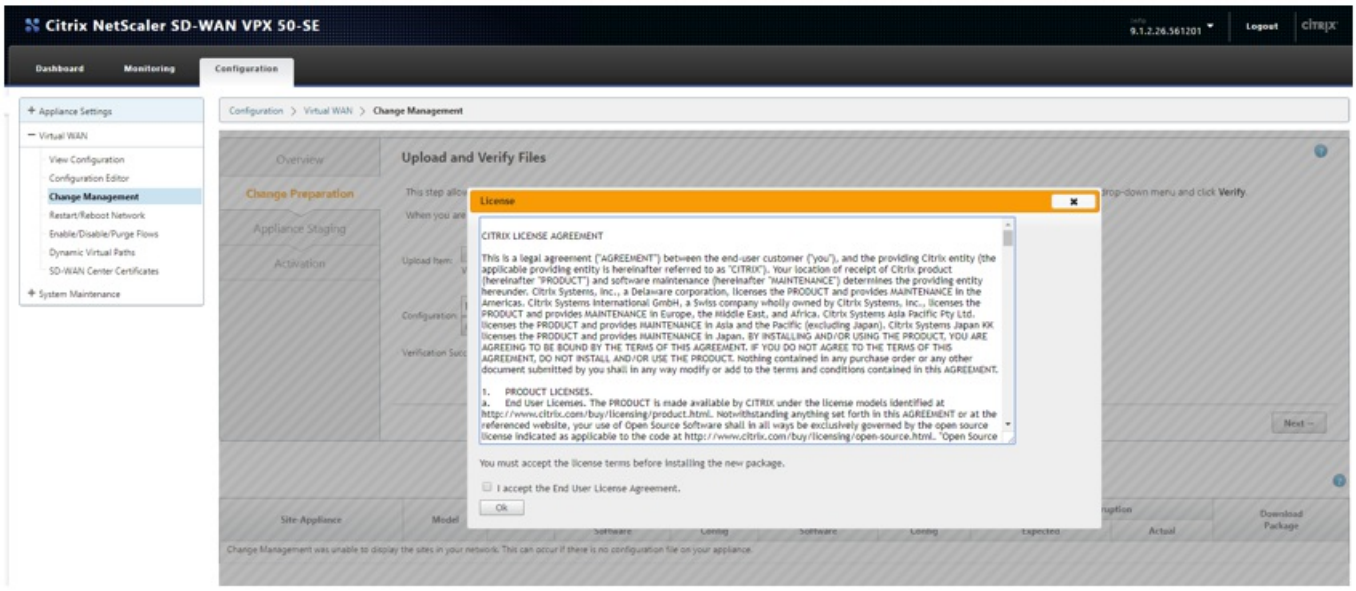

10. After accepting the license successfully you are navigated to Appliance Staging, click on **Stage Appliances** to proceed.

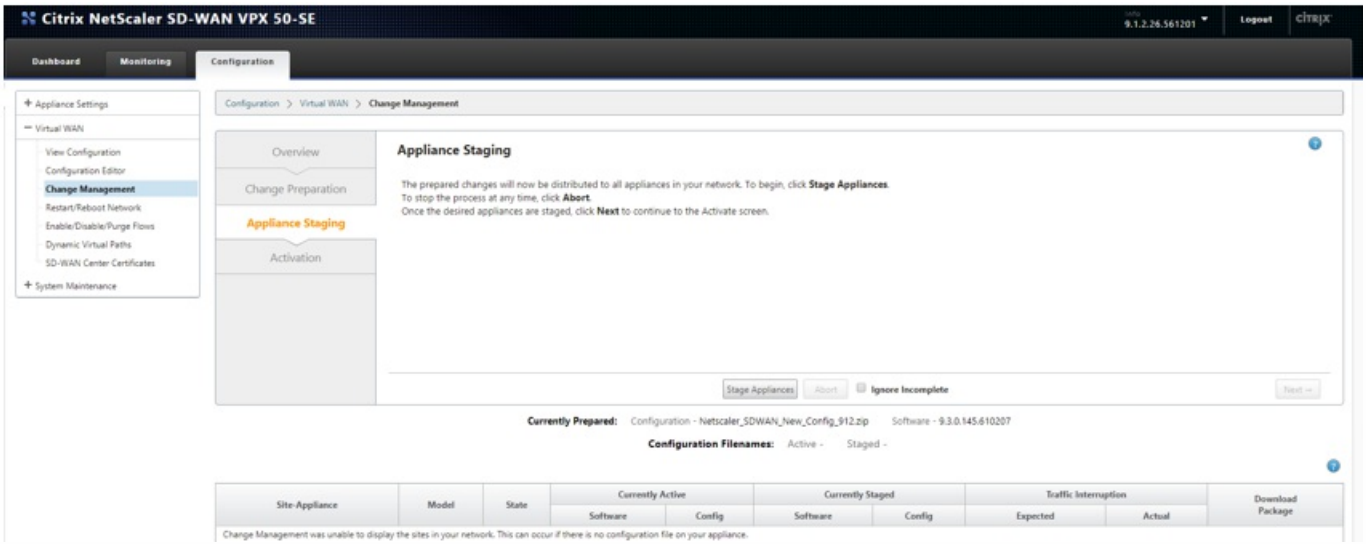

11. The transfer progress bar appears showing the current preparation and transfer status for each site.

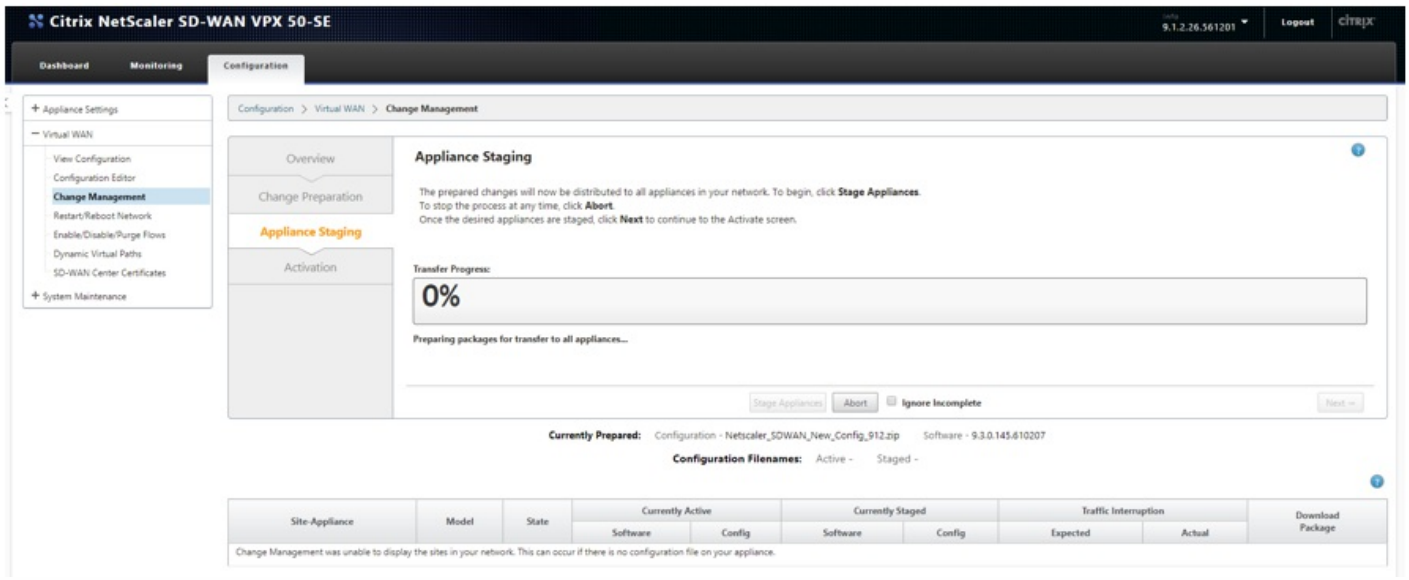

12. After preparing of package is completed, the Transfer Progress bar shows as 100% , click **Next** and proceed.

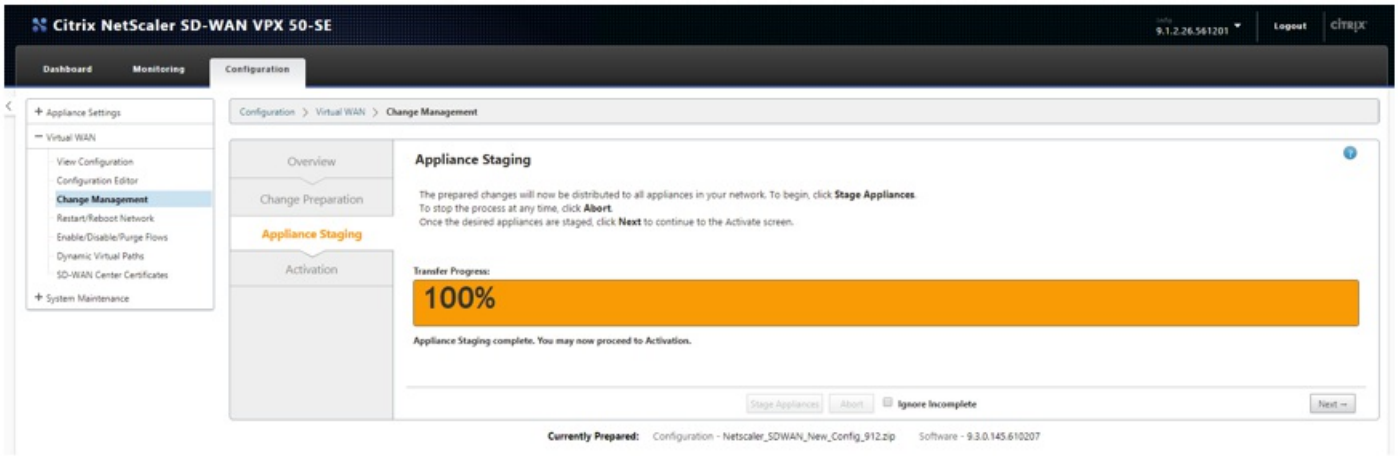

13. In the **Activation** page, staged software package can be activated.

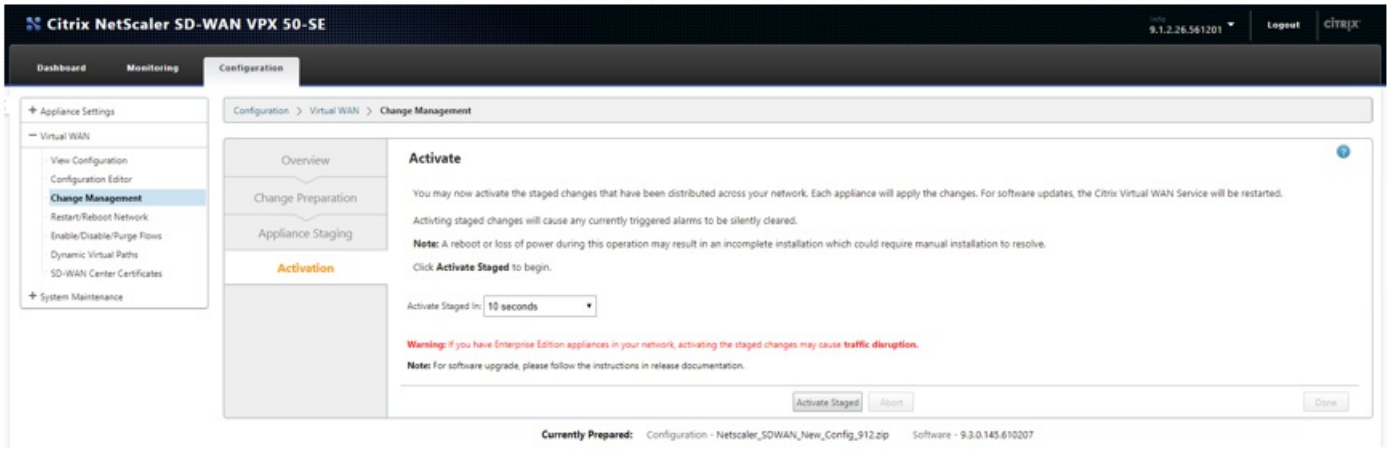

14. On clicking **Activate Staged**, user acceptance pop-up message is displayed. Accept and proceed as shown below.

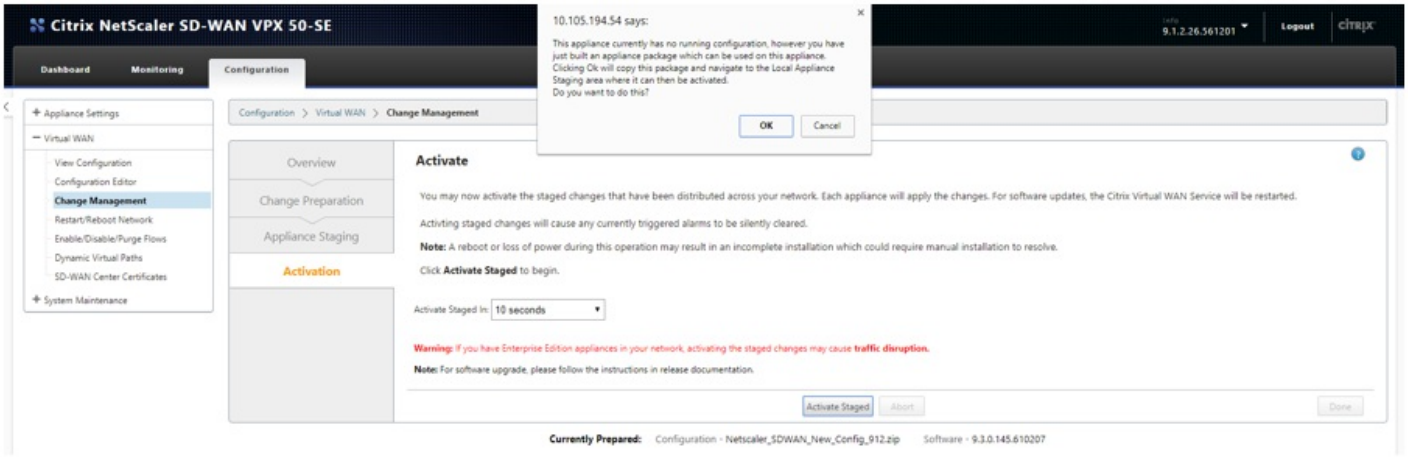

15. After accepting the message, click OK. The package gets copied to **Appliance Staging** state.

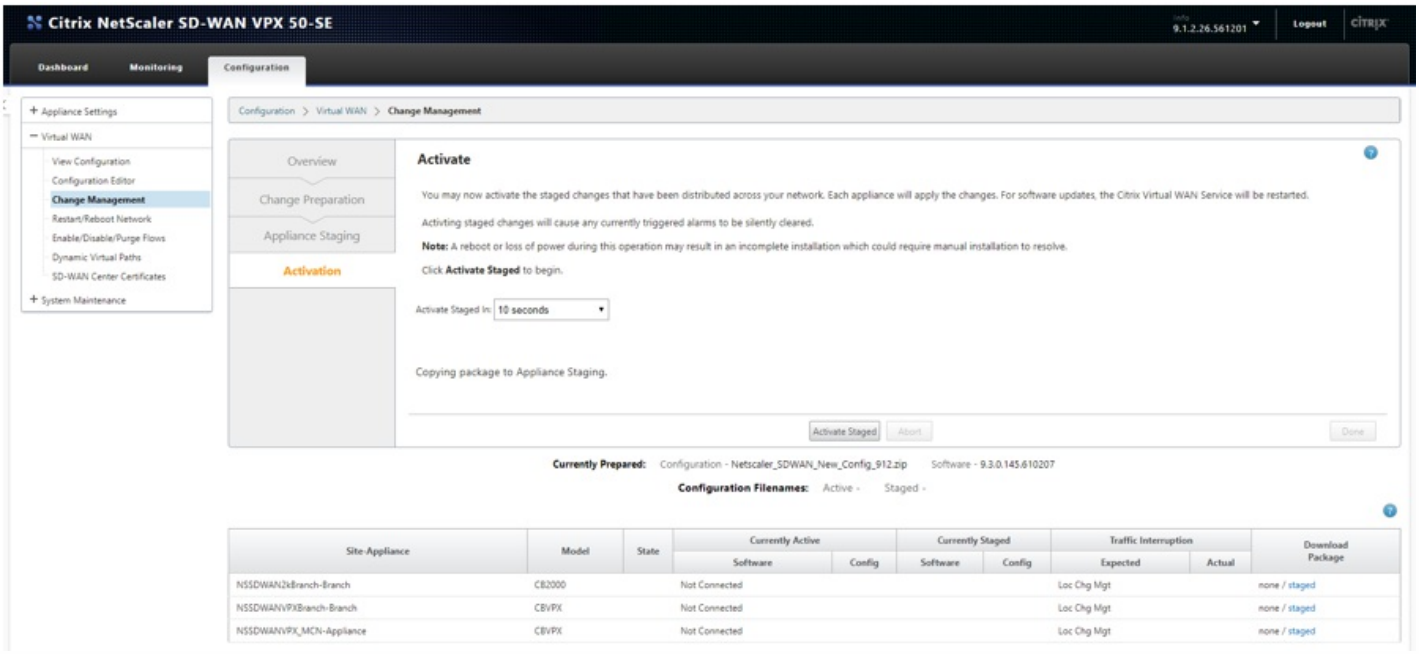

16. As this is the first time Appliance is being staged, you will be redirected to the **Local Change Management page** for activating the local appliance.

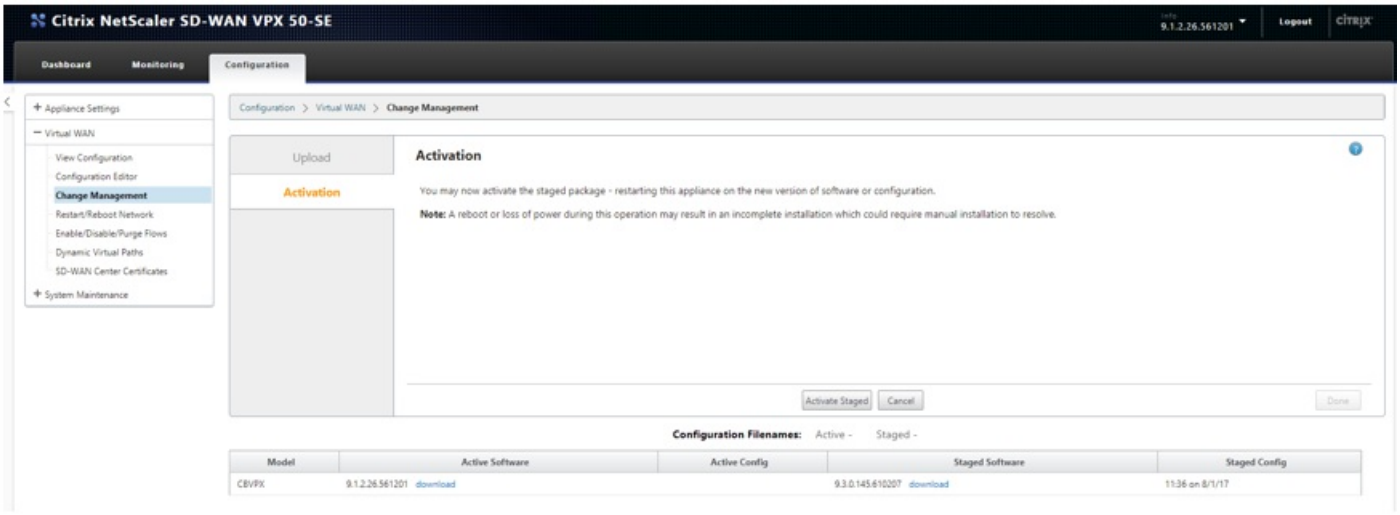

17. After clicking on Activate Staged in Local Change Management once again, user confirmation is requested to proceed further.

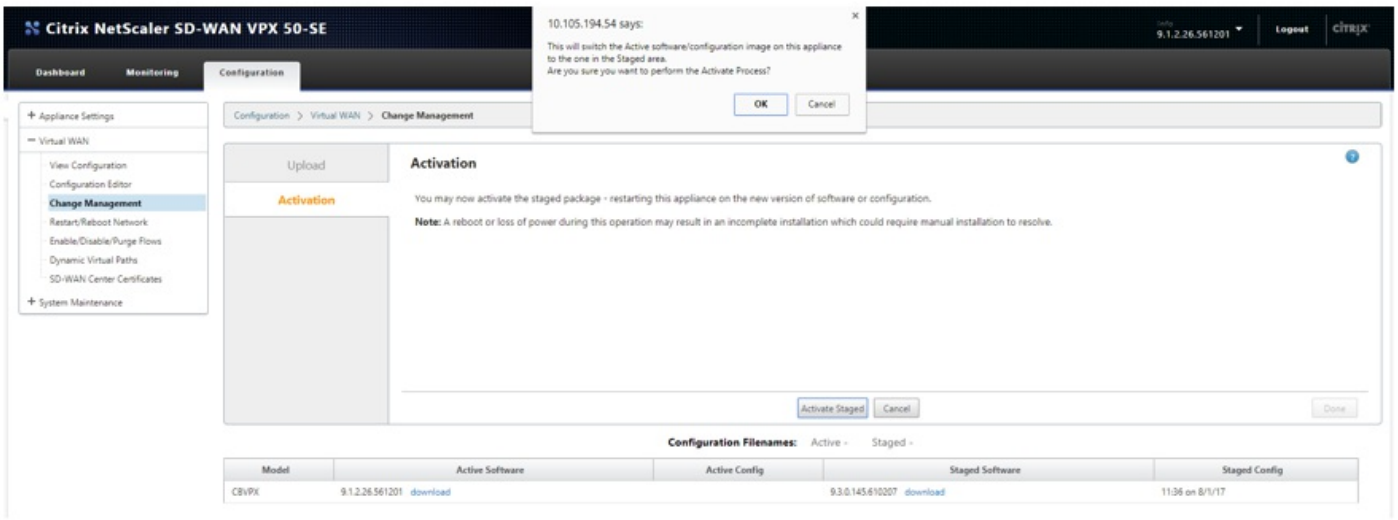

18. After accepting, activation starts with a countdown timer of 180s.

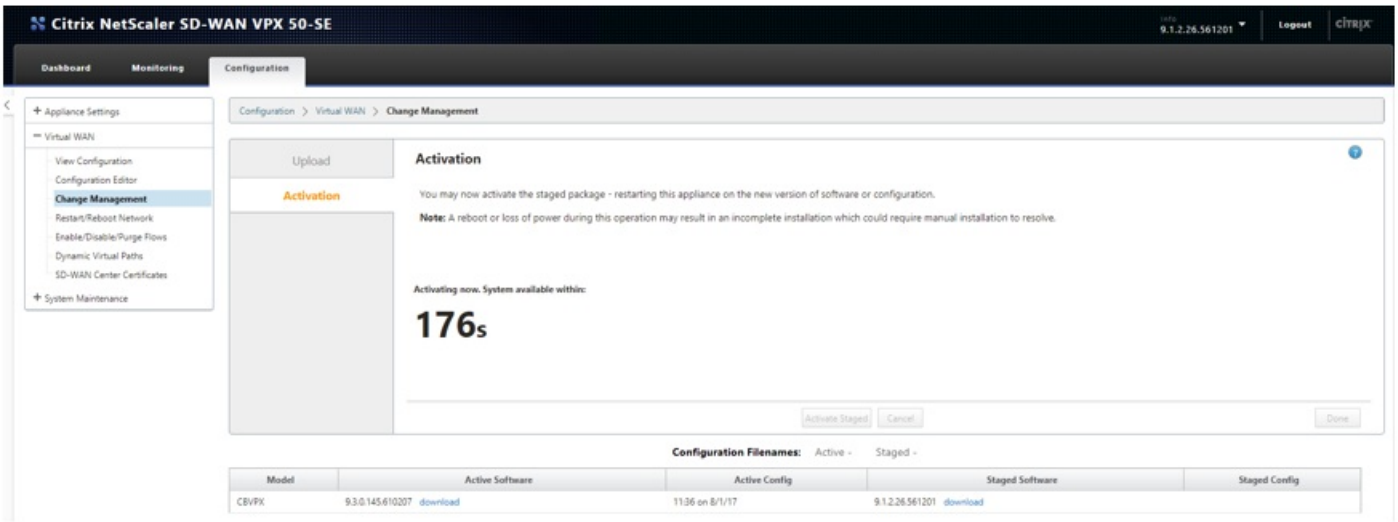

19. After the countdown timer expiration. a message indicating activation has completed is displayed.

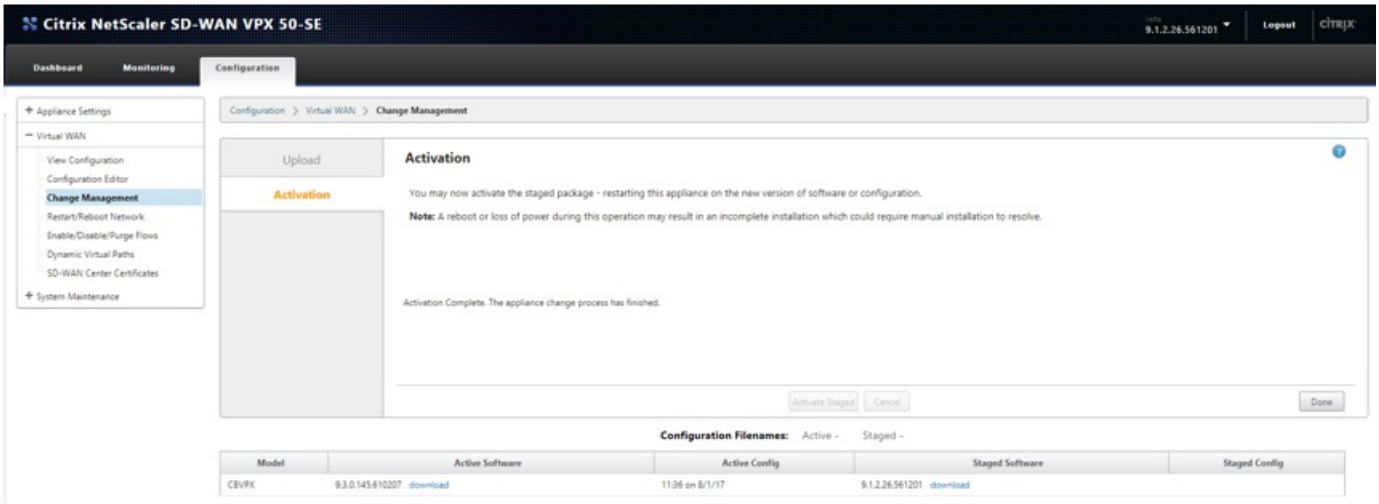

20. Navigate to **Change Management** page to download the local change management for respective branches that we need to bootstrap to the network with only VW software upgrade only, or else move to the next step as shown below.

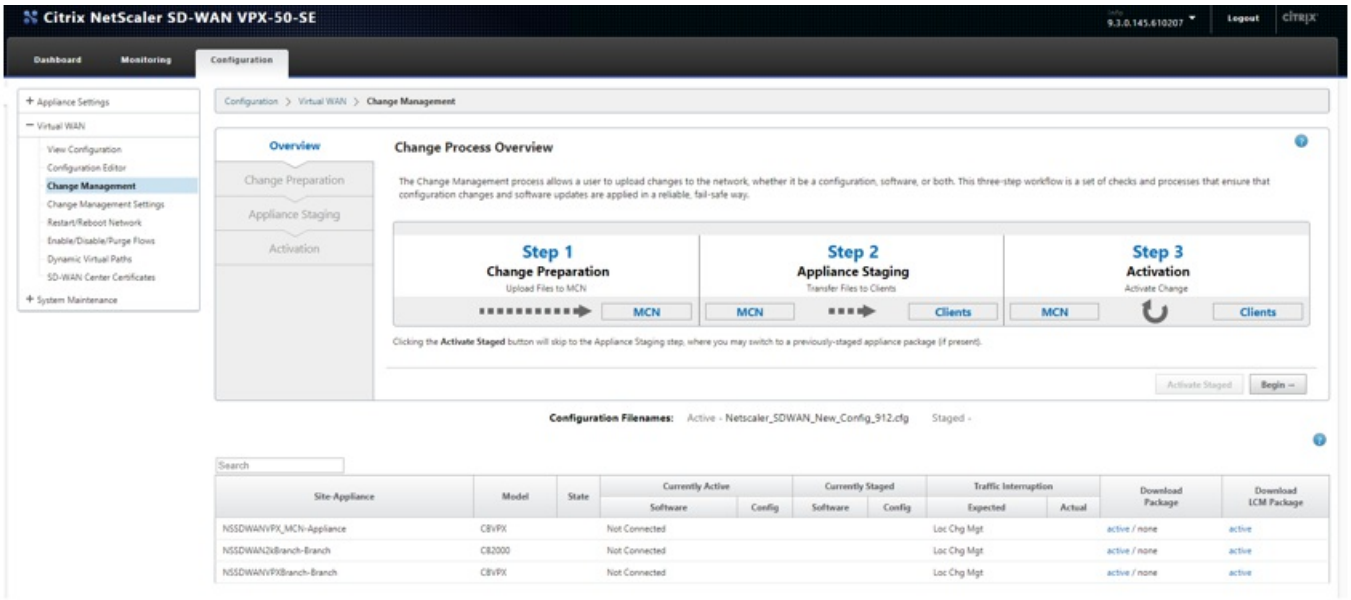

21. Perform **Change Management** once again, but this time the upload file would be of single step upgrade package with *ns-sdw-sw-9.3.x.zip* as shown below.

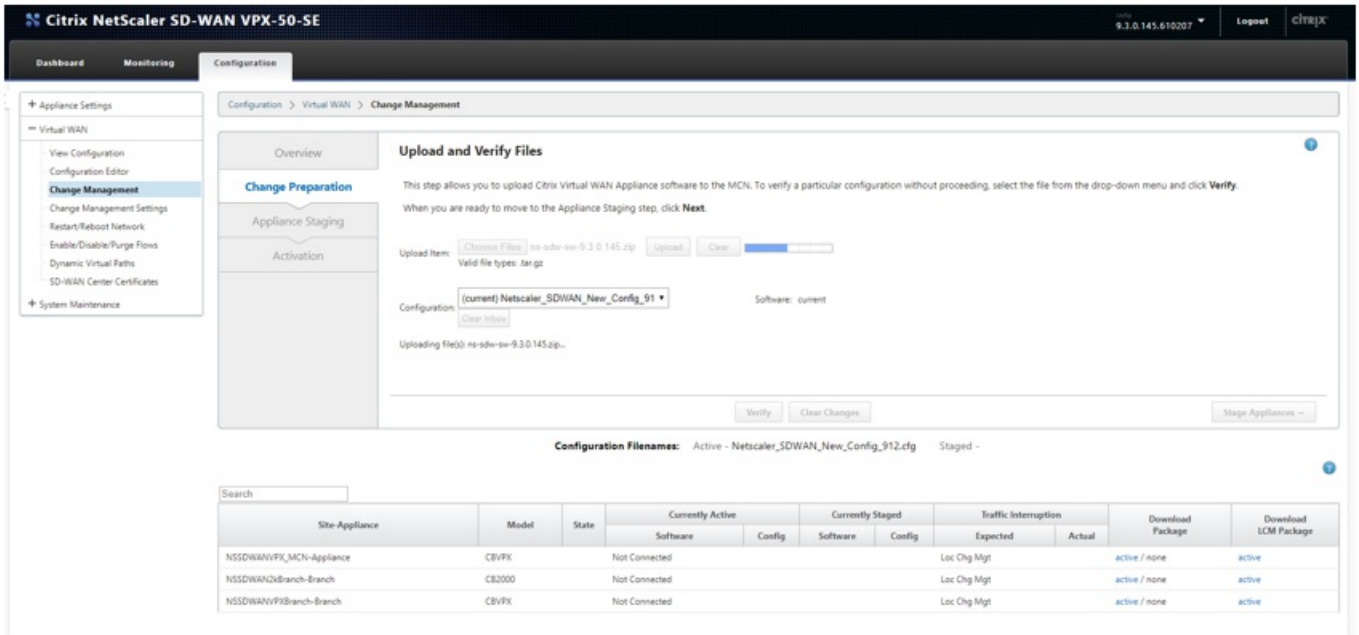

22. After upload is completed, the file is processed for any issues.

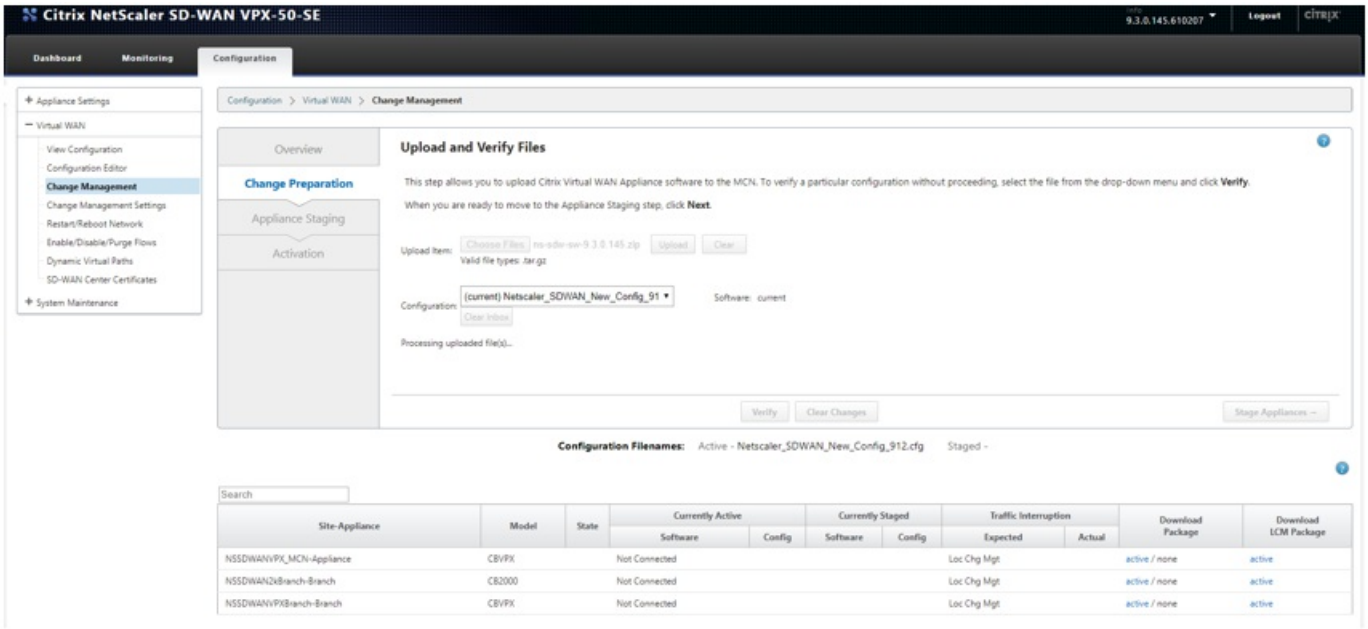

23. After the uploaded file is processed based on the configuration software for applicable models are uploaded to inbox as shown below.

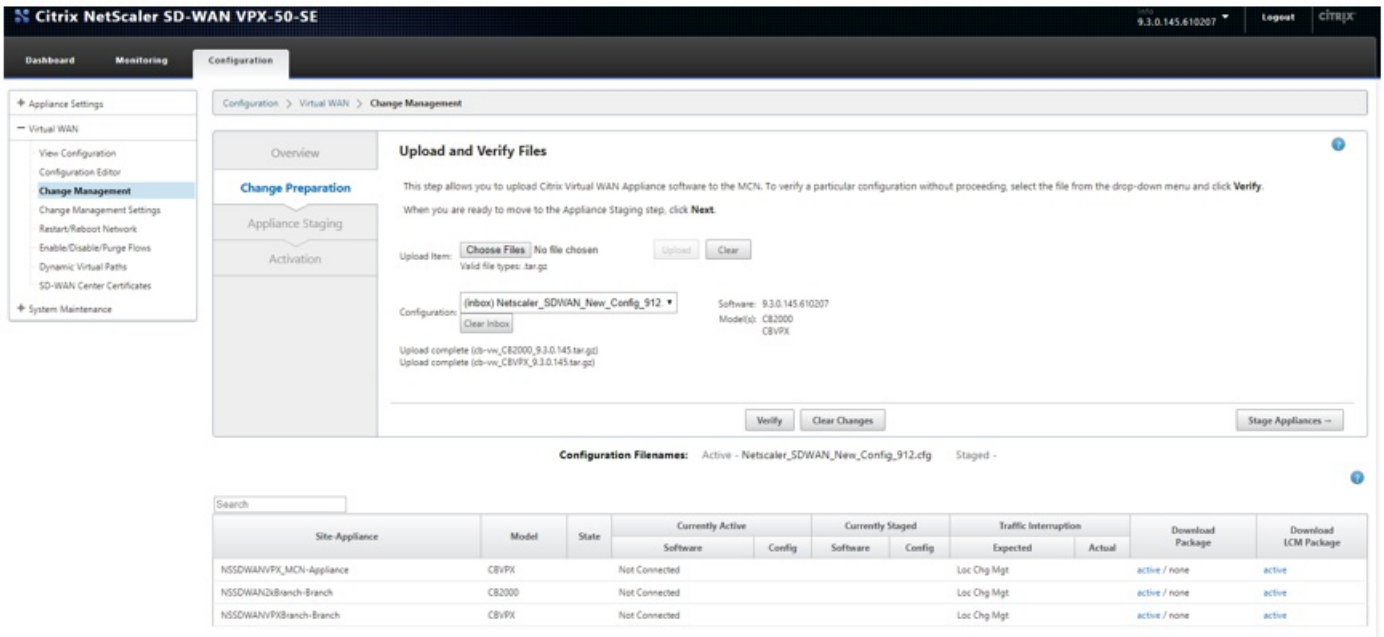

24. Click **Stage Appliances** to proceed with further steps as shown below.

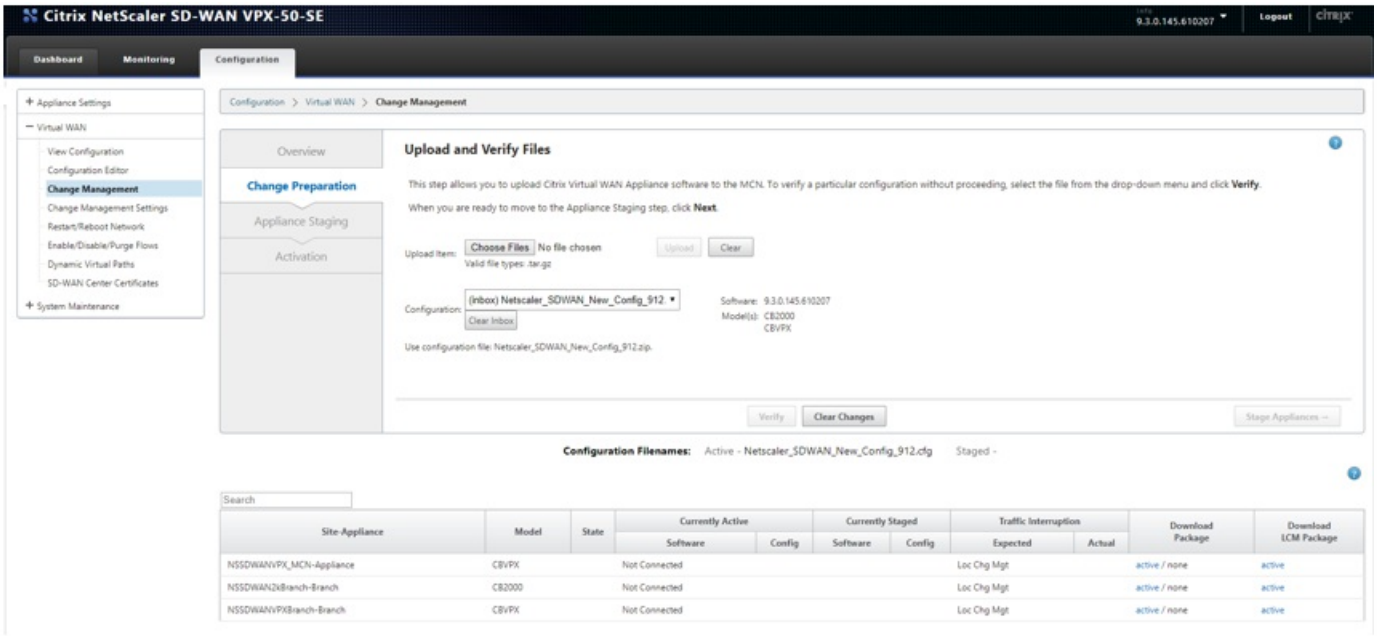

25. License agreement page is prompted when the user acceptance page is displayed.

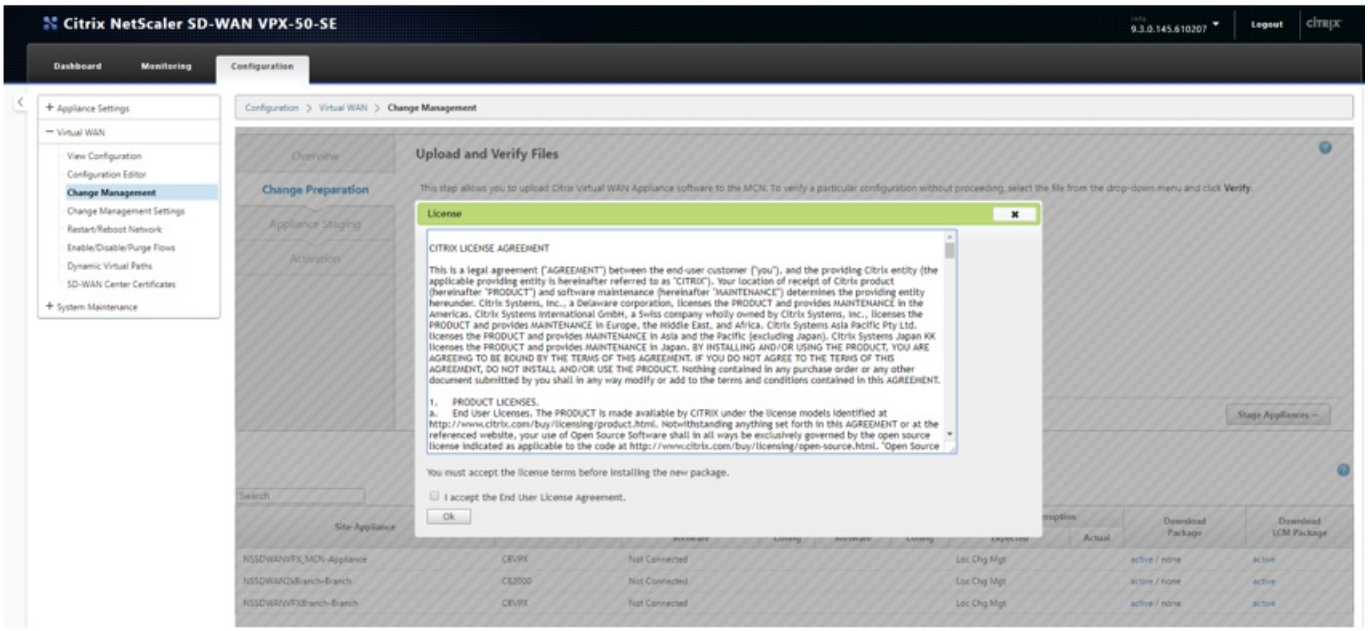

26. After accepting and clicking **OK,** user is redirected to **Appliance staging** page which shows the status for **Preparing and Staging the packages** for the concerned branches.

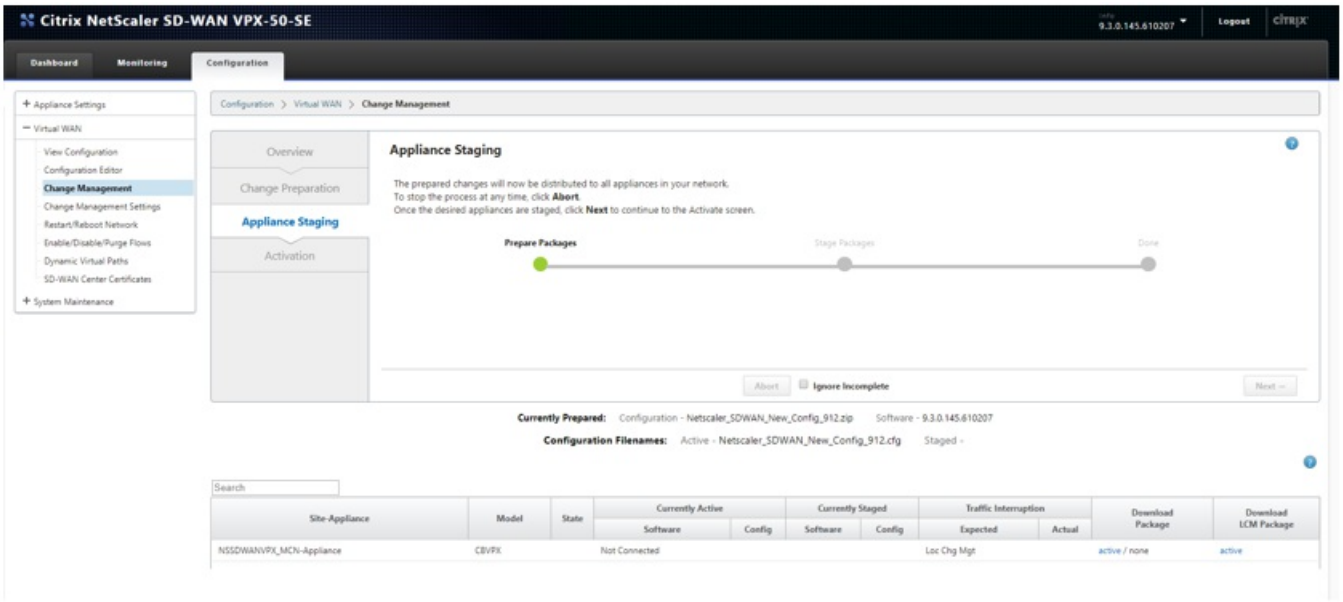

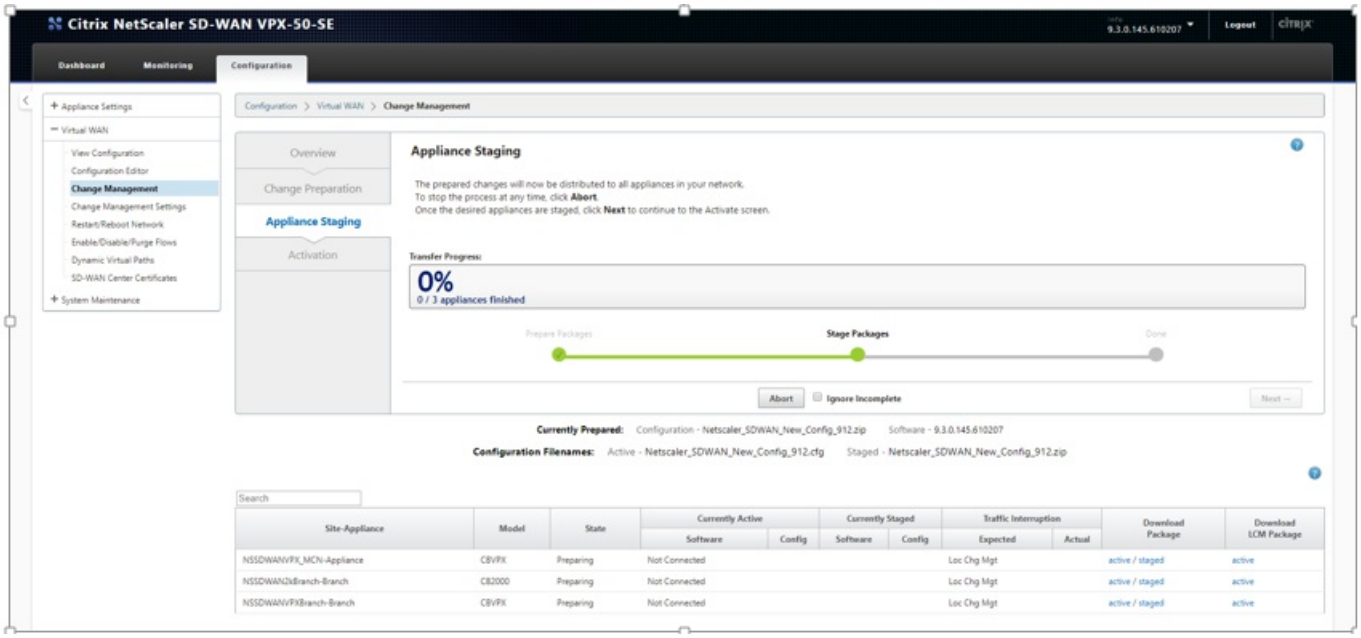

27. After transfer is completed, user clicks **Next** and is redirected to **Activation** page for further actions.

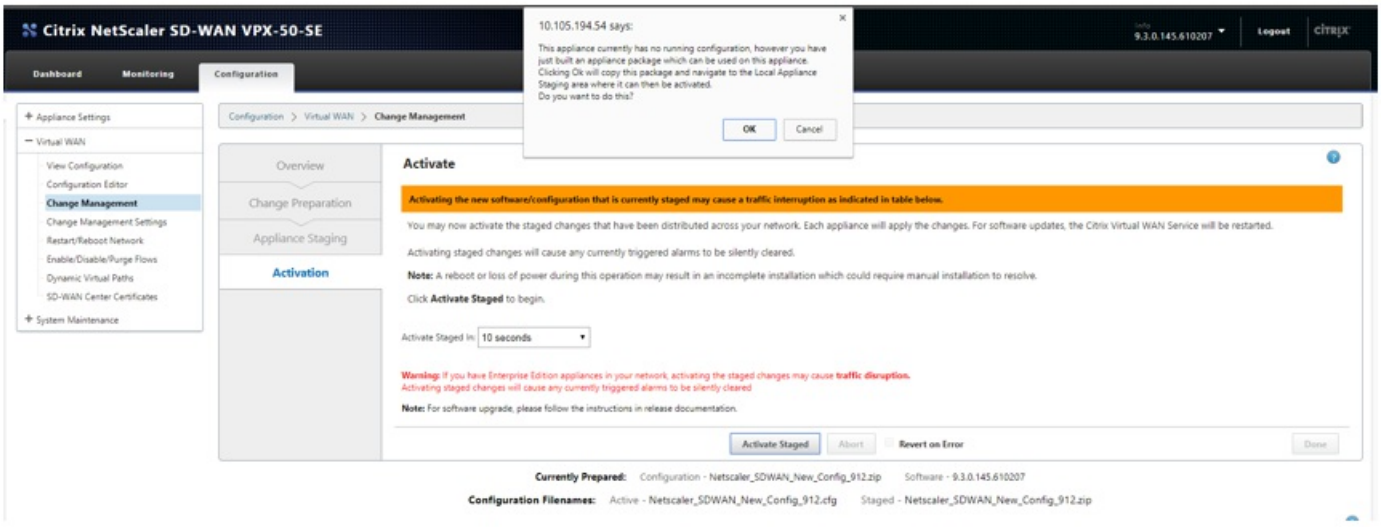

28. After completion of activation action in **Change Management**, the user is redirected to **Local Change Management** page as the upgrade is being done for the very first time on the MCN and no branches are yet bootstrapped to the network.

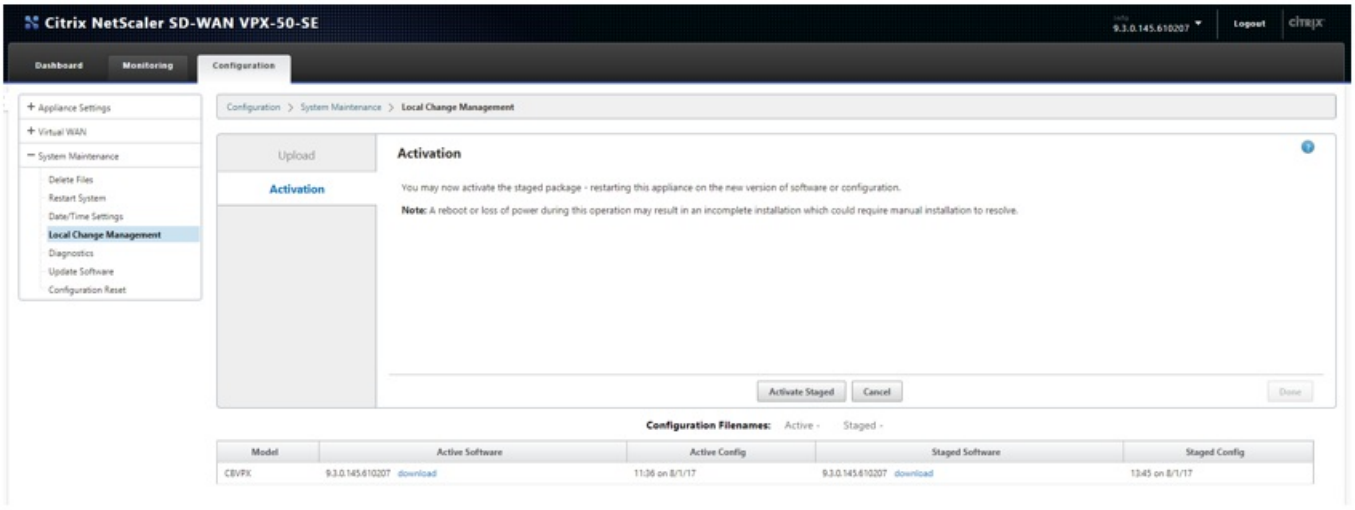

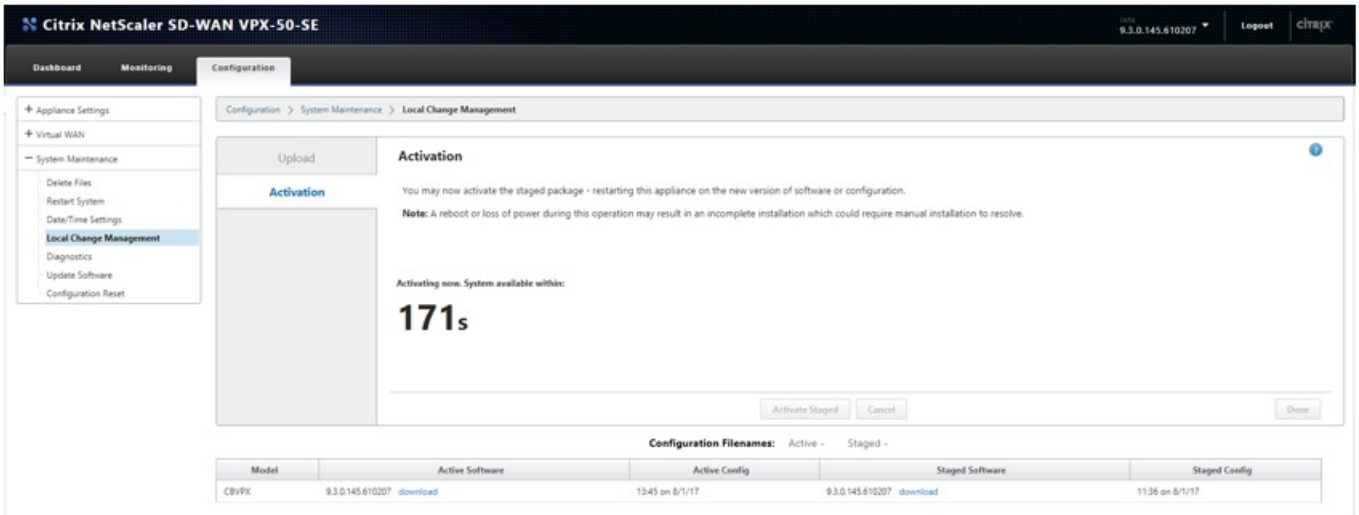

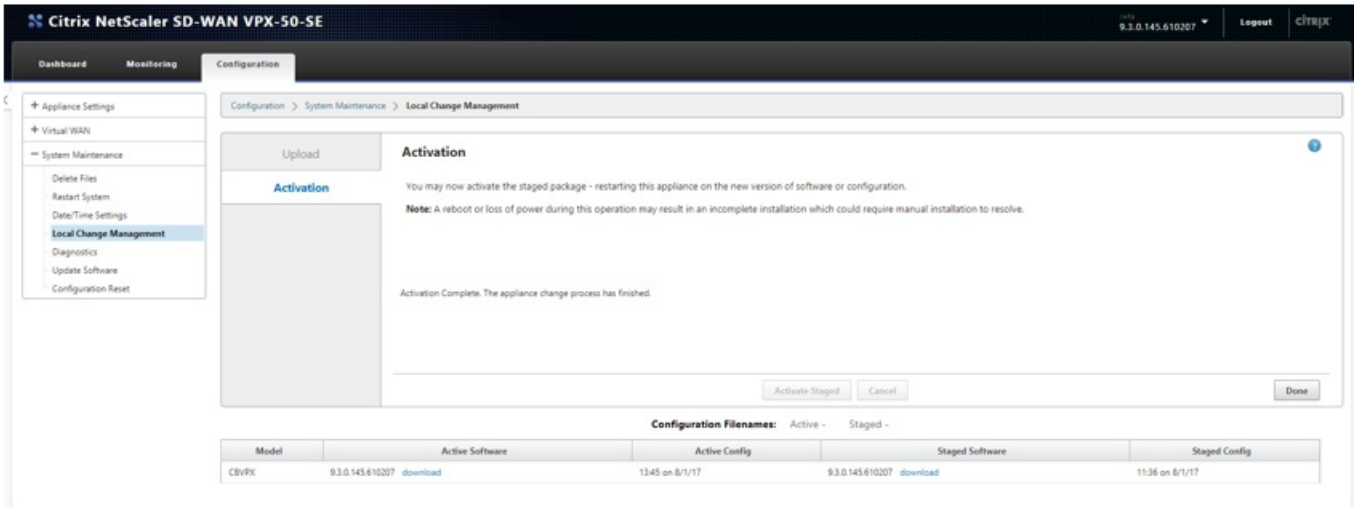

29. On completion of Local Change Management activation, user can go to Change Management page and download local change management package for corresponding branches from active hyperlink under '**Download Package**' column.

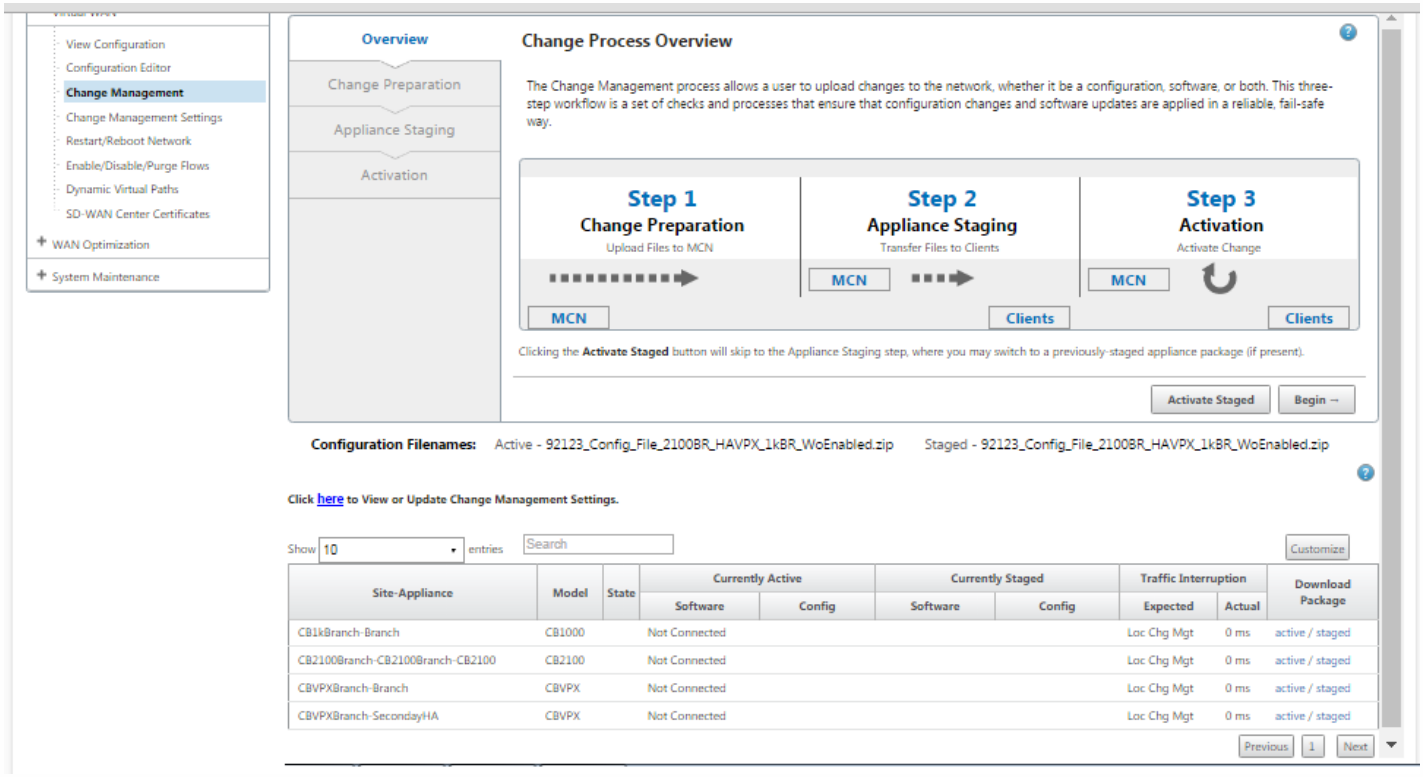

30. Login to the corresponding branch which needs to be bootstrapped into the network and navigate to **Local Change Management** page and upload the local change management package that was downloaded from MCN for this branch.

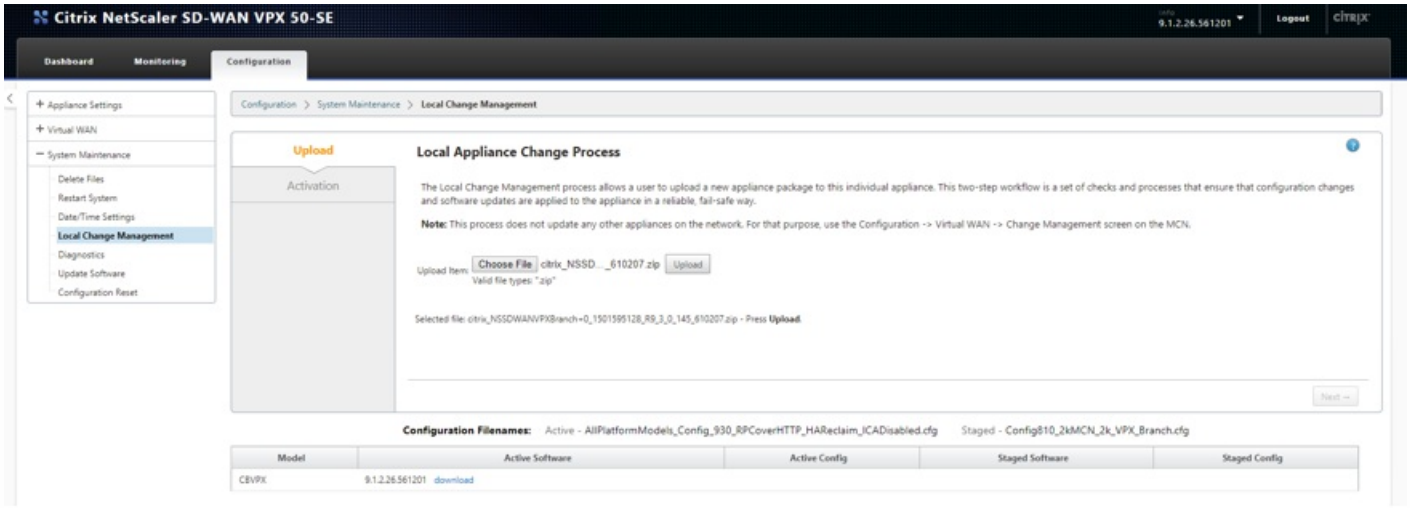

31. After file to path is chosen to be uploaded click **upload** and a upload status progress bar appears showing the status of upload as seen below.

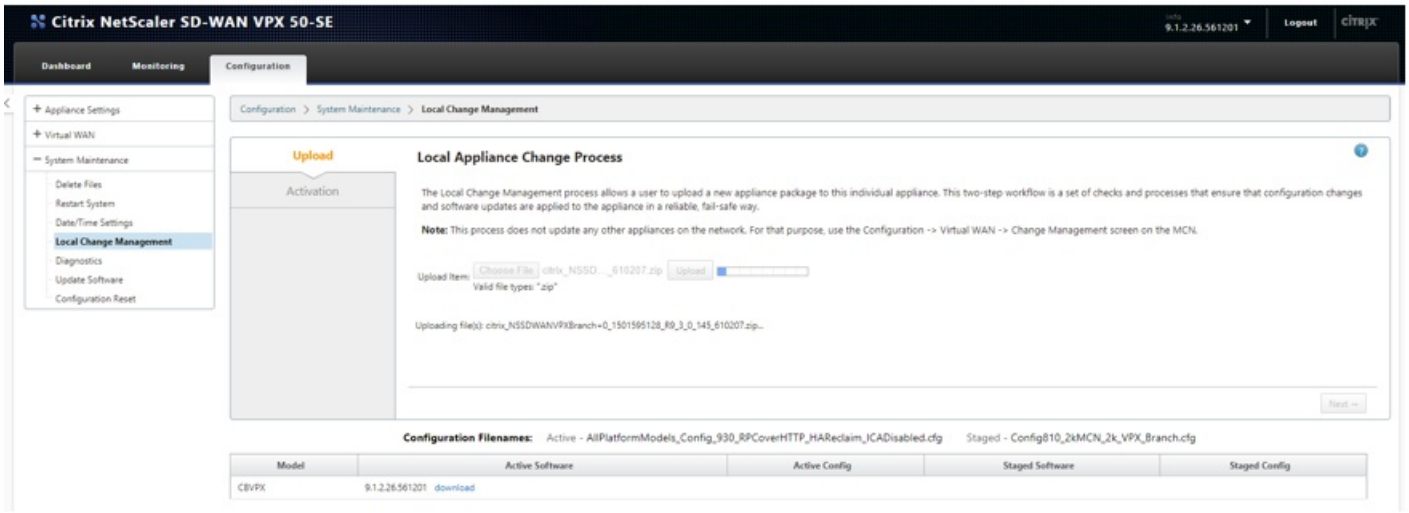

32. After upload is completed, you will receive a message saying 'Upload Complete'. Now click **Next**.

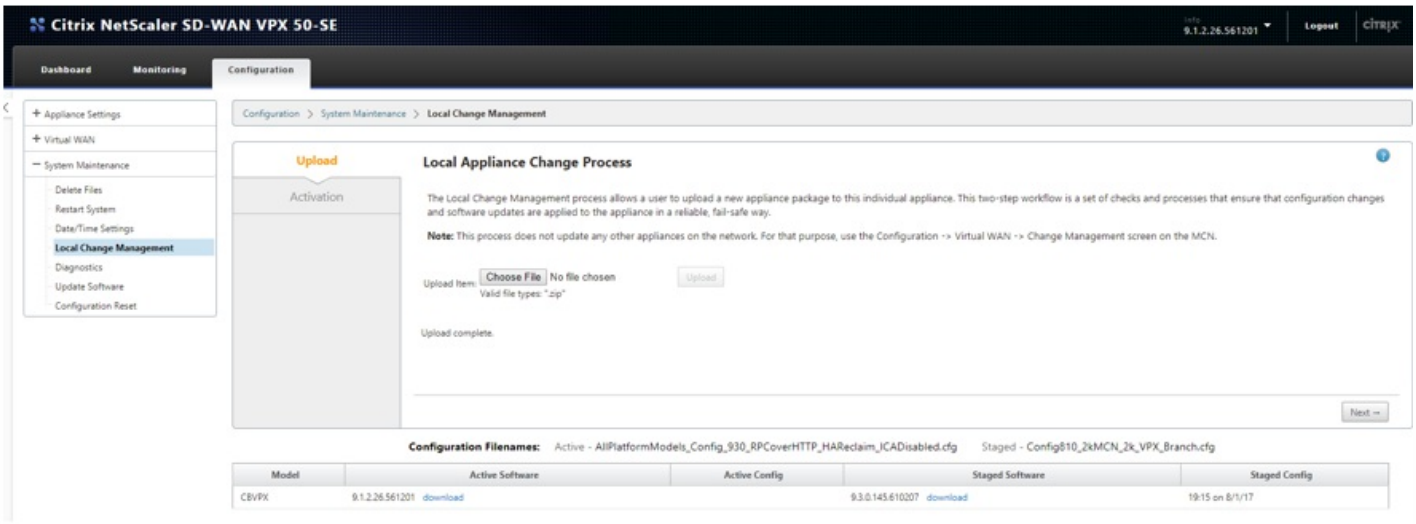

33. You are redirected to the activation page where you can choose to activate staged software or cancel as shown below.

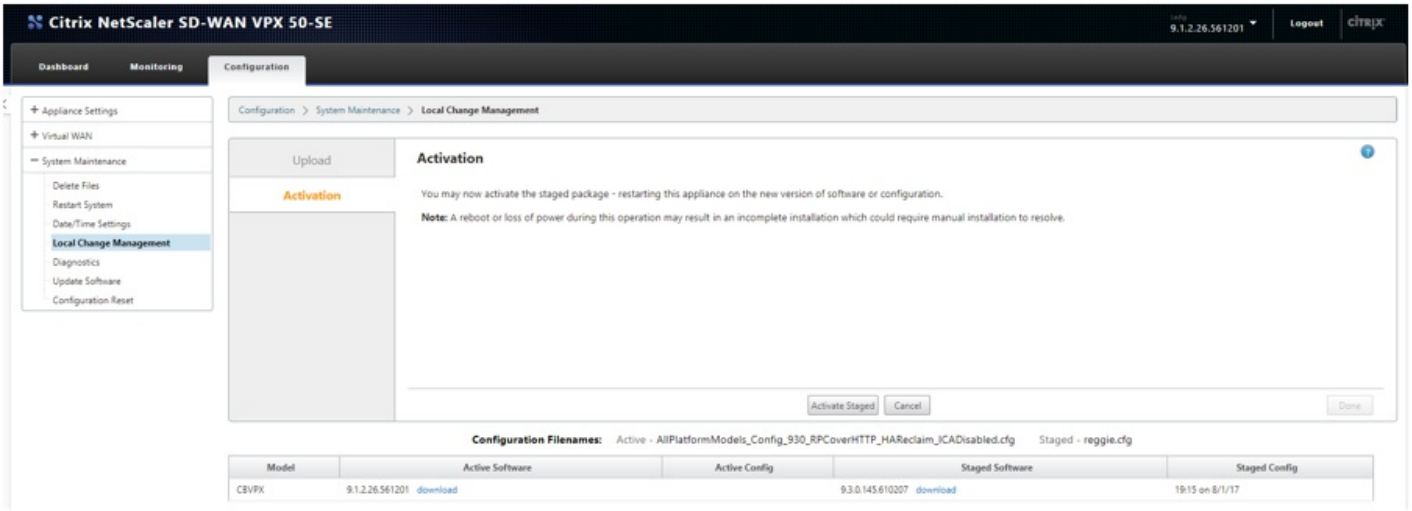

34. Clicking '**Activate Staged**', the user confirmation is asked to proceed.

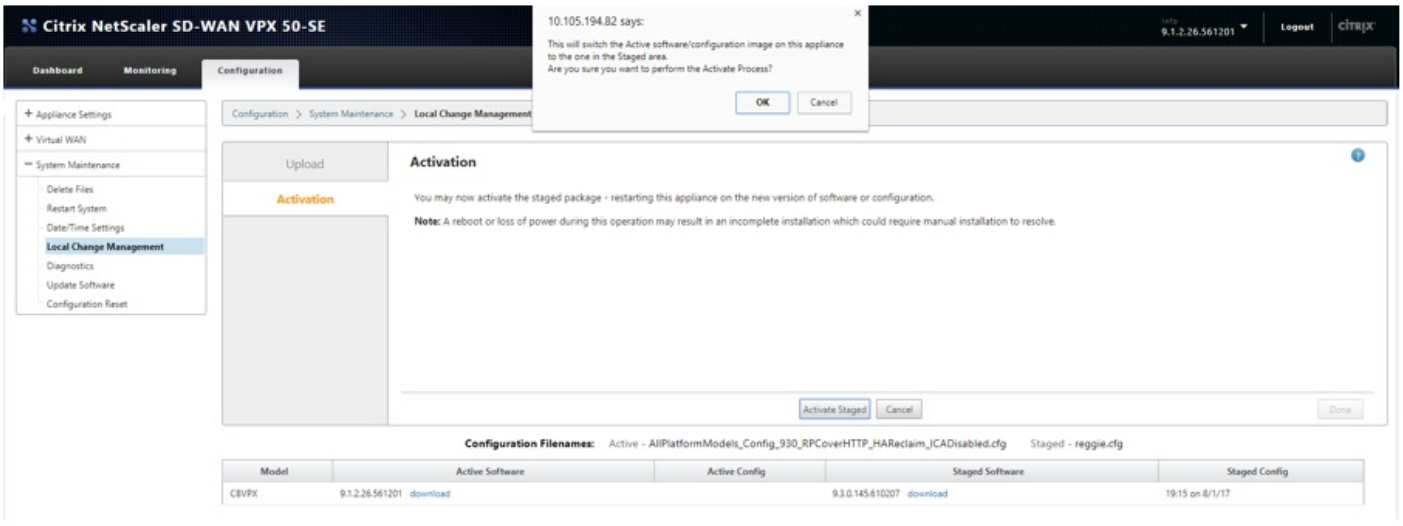

35. After accepting the message, an activation countdown is started with 180s as shown below.

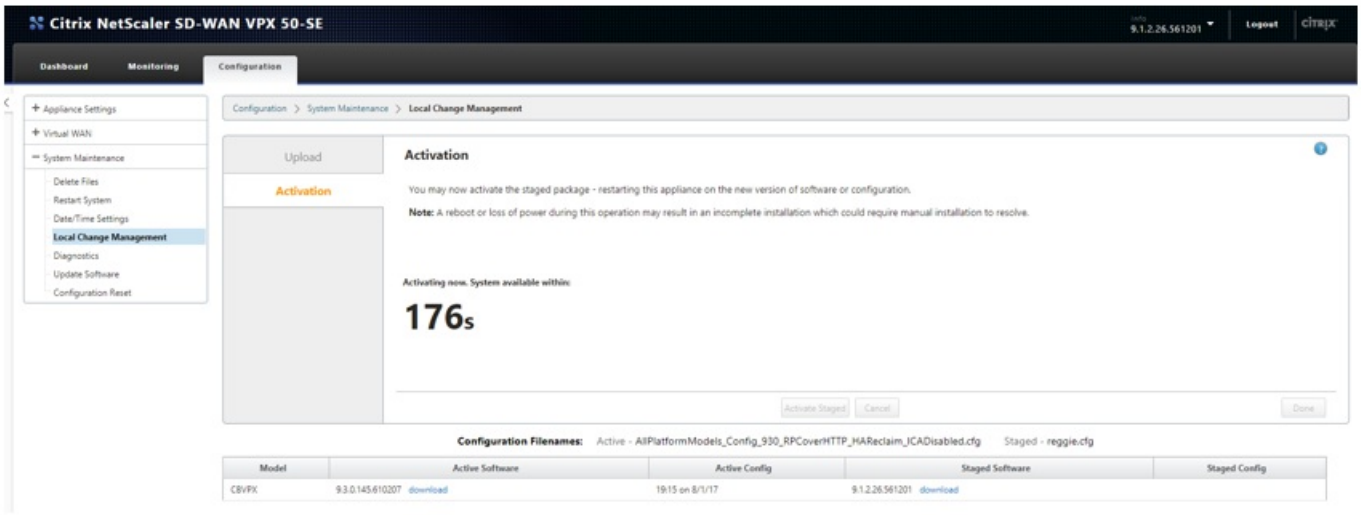

36. On completion of countdown a message is displayed, as stating **Activation Complete**. The appliance change process has completed, and you can click Done.

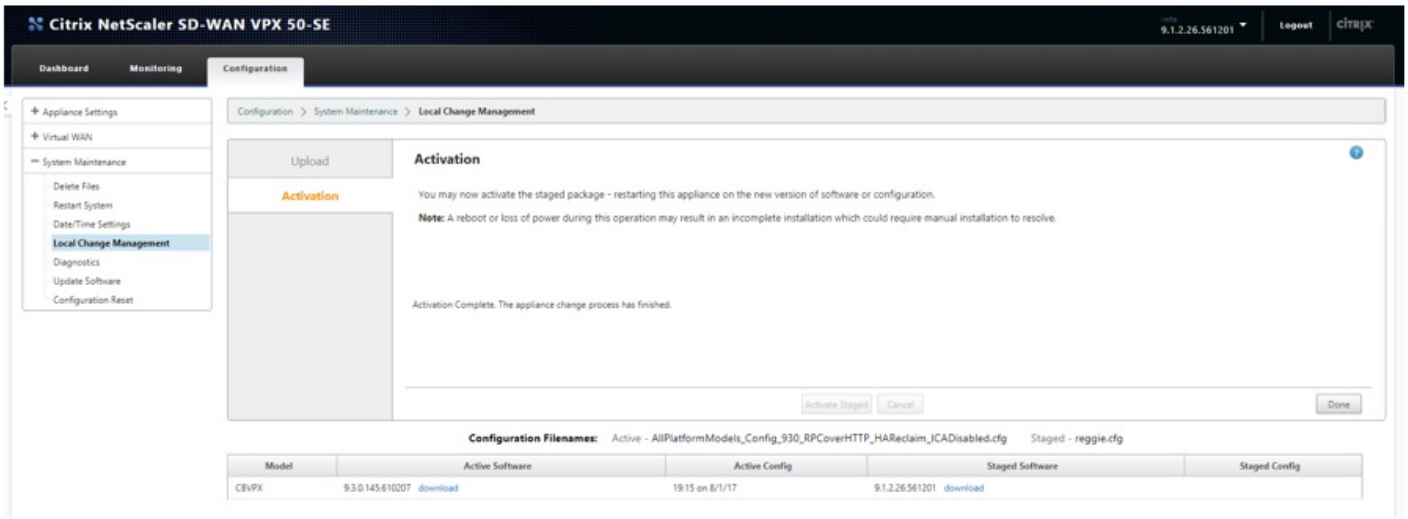

37. You can enable SD-WAN service and see that virtual paths are up between the branch and MCN.

Convert SD-WAN 1000 / 2000 WANOP Appliances to Enterprise Edition With USB

Dec 01, 2017

Only the SD-WAN 1000 and 2000 WANOP appliances can be converted to SD-WAN Enterprise Edition appliances.

- Ensure that you are converting the 1000 WANOP appliance only, and not the 1000 WANOP WS. The 1000 WANOP WS appliance does not support conversion to the SD-WAN Enterprise Edition appliance.
- Ensure that you have the default credentials to log into the existing *Dom-0 - root/nsroot*.

The conversion procedure is a two-step process involving the following steps:

- Insert enclosed USB stick into the Citrix SD-WAN appliance.
- Verify that the serial console is connected and proceed with the conversion process.

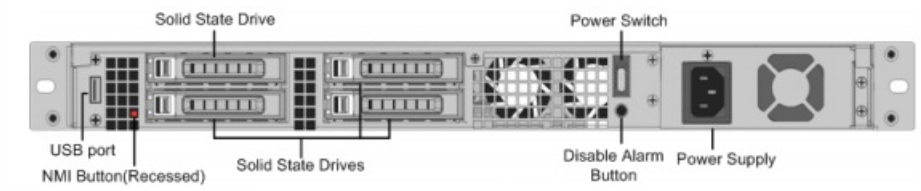

How To Convert With USB Stick

To upgrade the appliance with USB stick:

1. Insert the enclosed USB stick into the Citrix SD-WAN appliance.
2. Connect to the serial console of the appliance.
3. Reboot the appliance.
4. During the boot process, when you see the cursor moving across the screen, do the following:
    a. Press and hold the **ESC** key.
    b. Press and hold the **SHIFT** key.
    c. Press the number **1** key (SHIFT +1 = !) and release all keys.
    d. Repeat steps a, b, and c until the cursor stops moving.

Note

The above steps should be executed during the applaince reboot process. The key strokes should happen during BIOS post stage as described in step 4.

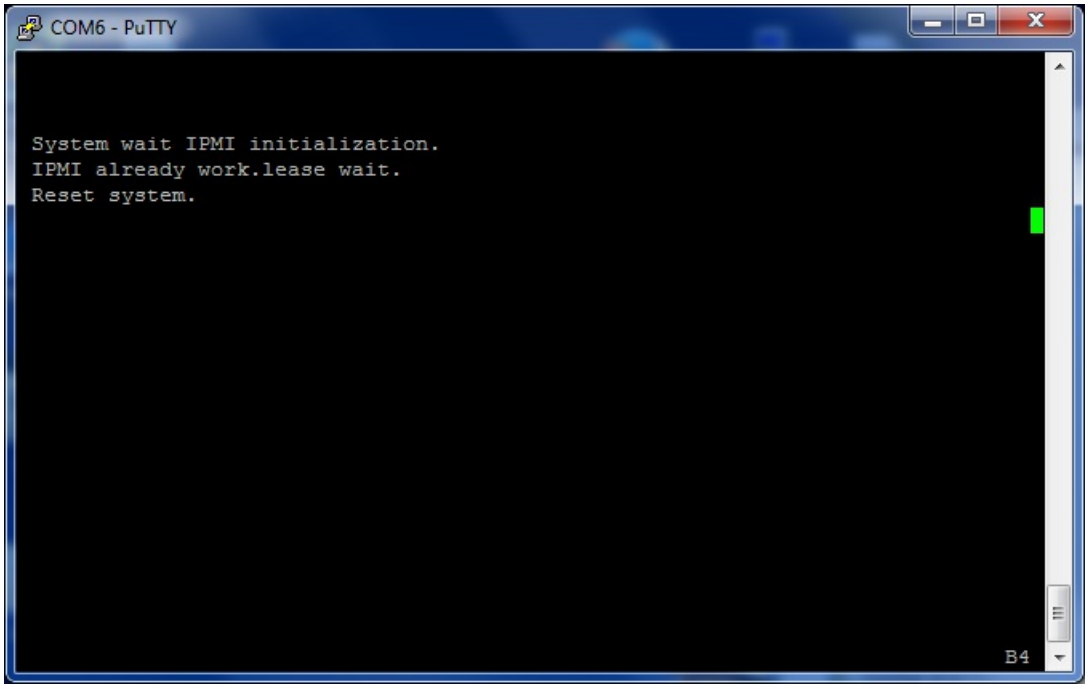

5. When BIOS loads, choose the external USB drive, for example; PNY USB 2.0 FD 1100 to boot the appliance. The external USB drive is shipped by Citrix if you have ordered for it.

6. You need to choose the platform edition which you want to use, if the platform supports more than one edition, such as SD-WAN 1000 and 2000. editions Choose the **Enterprise Edition** software upgrade option when prompted.

7. Click **<Yes>** to confirm.

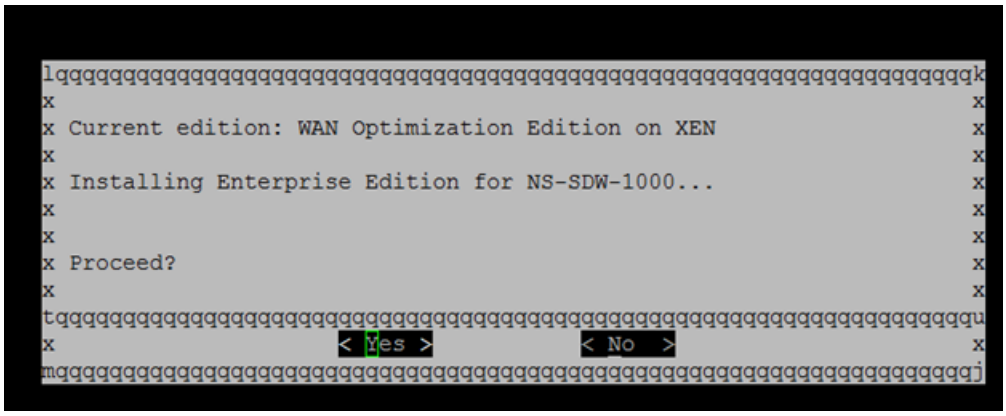

8. Once confirmed, the software package is copied from the USB.

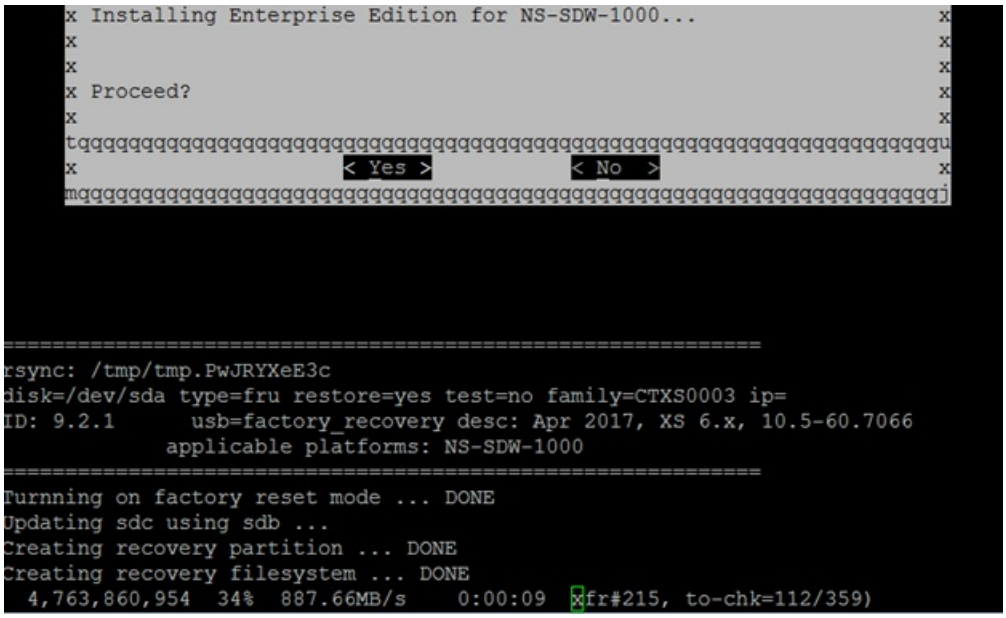

Upgrade process will take approximately 20-30 minutes to complete. The system reboots after 1-2 minutes and the login prompt is displayed. For the SD-WAN 1000 WANOP platform edition, upgrade process is approximately an hour since updating the internal USB drive itself takes around half an hour.

9. Unplug USB stick after the procedure is complete.

- For licensing about the NetScaler and NetScaler SD-WAN products, see the support link at: http://support.citrix.com/article/ctx131110
- For Documentation and Release Notes information about NetScaler SD-WAN, see; http://support.citrix.com/proddocs.

Converting Existing Appliance to Enterprise Edition Appliance

Aug 09, 2017

You can convert your existing appliance to Standard Edition if the Enterprise Edition license installation fails on the Enterprise Edition 1000 and 2000 appliances:

To perform platfrom conversion from Standard Edition to Enterprise Edition:

1. Export the Configuration locally.
2. Download the **Active Package** from the **Change Management** page.
3. Upgrade the appliance using the downloaded package from **System Maintenance** -> **Update Software** -> **Re-image Virtual WAN Appliance software**.
4. Click **Choose File** to provide the *cb-vw_CB1000_9.1.0.x.tar.gz* file.
5. Click **Upload**. Select **Accept** and click on **Install** to proceed.
6. Install the Enterprise Edition License.
7. Perform **Local Change Management** on the appliance using the downloaded active package in step 2 above.

Single-Step Upgrade for SD-WAN Appliances

Sep 27, 2017

> **Note**
>
> To use single step upgrade feature, the MCN must run a version which supports single step upgrade. For example; SD-WAN 9.3.x.

In release 9.3, a single step upgrade package using the SD-WAN GUI change management option to upgrade non-SD-WAN components in the network for all applicable platform editions has been introduced. The MCN distributes all necessary software components to the sites (Branch) in the network.

After the branch receives the upgrade component files, these can be installed at scheduled time intervals as specified by the user. If the scheduled time is not specified, it uses the default time which is set by MCN for all branches.

The MCN also generates packages for sites on demand. Download the active package from the active hyperlink under the **Download LCM package**. You can bring or boot strap a new site into the network using this package. All software components will be installed using the Single Bundle Upgrade (.tar.gz) package. The MCN keeps a copy of the software.

Pre-requisites to Perform Single Step Upgrade

- Have the SD-WAN 9.3 Virtual WAN software installed. Use the existing upgrade procedure to upgrade from any previous release version to the current 9.3 Virtual WAN software version. Do not use SBU (single bundle upgrade) procedure to upgrade from an older Virtual WAN software version to 9.3 version.
- After the SD-WAN 9.3 Virtual WAN software is installed, follow the single step upgrade (.zip) procedure to upgrade WANOP components and XenServer supplemental software packages. You can use the SD-WAN GUI to schedule upgrade of these additional components through the **Change Management Settings** page. The above pre-requisite does not apply for upgrading from release version 9.3 to the 9.3.x latest build version.

> **Note**
>
> Change management upload error occurs, if you attempt to perform single step upgrade (.zip file) from previous versions to 9.3.
>
> It is recommended that you use build version 9.3.0.x and above to use the single step upgrade procedure.

Review the following upgrade procedures to upgrade to software release version 9.3.

Upgrade to 9.3 with working Virtual WAN configuration

Upgrade to 9.3 without working Virtual WAN configuration

Following are the supported upgrade and downgrade scenarios for SD-WAN SE and EE appliances. It is assumed that we are upgrading Virtual WAN software first and then upgrading the other components after the required software is staged

at the branch sites or on the MCN.

You can download local change management package from MCN and apply it on the factory shipped appliance. After the local change management package is applied to the branch site boxes, all relevant components are upgraded immediately without waiting for maintenance window, if applicable.

1. The appliances currently in an active network with virtual paths up and running.

- In this case, the appliances receive packages from MCN. The components are installed all the files from MCN are received and it is in the scheduled time window.

2. The appliances are currently out of the network with virtual paths down.

- In this case, the process is similar to the appliances which are factory shipped. You need to download local change management from MCN and upload the package to the branch site appliances.

The appliance stages multiple files applicable at the branch site based on the appliance model and platform edition. The version information is reported by the branch site and/or configuration options, if applicable. The branch site appliances perform the upgrade utilizing the staged files. The non-Virtual WAN software components can be installed based on the preferences, manual and/or schedule.

Downgrading to a previous version of Virtual WAN software is supported. With single step upgrade process, you can install WANOP software packaged with a given Virtual WAN software version. You can only upgrade hotfix and/or SVM versions if the software versions in the packages are higher.

You can re-install the legacy software with the required configurations (using *tar.gz* files).

Downgrading to previous software version

If you upgraded an existing software version to release version 9.3 using the *tar.gz* upgrade process, you can downgrade the software version to a previous software version.

If you used the .zip (single step upgrade) procedure to upgrade to version 9.3, you cannot downgrade the software version to a previous software version.

Single Step upgrade in High Availability Deployment Mode

During single step upgrade if HA flip happens then, you need to switch back to the old primary appliance manually, or upload the single bundle package to the new primary appliance.

Before You Begin

Aug 09, 2017

This section outlines the hardware and software requirements for deploying Citrix NetScaler SD-WAN Standard and Enterprise Editions, and defines the platform dependencies. Also provided is a summary and overview of the SD-WAN appliance installation and deployment procedures.

For more information, refer to the following topics:

- System Requirements
- Acquiring the Netscaler SD-WAN Software Packages
- NetScaler SD-WAN Software Packages and Appliance Models
- NetScaler SD-WAN Appliance Packages
- Preparing for Your Deployment
- Installation and Configuration Information checklist

System Requirements

Aug 09, 2017

Hardware Requirements

Instructions for installing your NetScaler SD-WAN Appliances are provided in Setting up the SD-WAN Appliances.

Firmware Requirements

All SD-WAN appliance models in a Virtual WAN environment are required to be running the same NetScaler SD-WAN firmware release. For additional information, please contact NetScaler Customer Support.

The SD-WAN release requires all appliances on the SD-WAN network to install the same software release. Appliances running earlier CloudBridge software versions will not be able to establish a Virtual Path connection to the appliance running SD-WAN release 9.2.

Software Requirements

For details regarding license requirements, see Licensing.

Browser Requirements

Browsers must have cookies enabled, and JavaScript installed and enabled.

The NetScaler SD-WAN Management Web Interface is supported on the following browsers:

- Mozilla Firefox 35.0+ (Recommended version 43.x)
- Google Chrome 40.0+ (Recommended version 49.x)

Supported browsers must have cookies enabled, and JavaScript installed and enabled.

Acquiring the NetScaler SD-WAN Software Packages

Aug 09, 2017

This section provides information about downloading the NetScaler SD-WAN software packages.

> **Note**
>
> Before you download the software, you must obtain and register a NetScaler SD-WAN software license. For information, please see Licensing.

Downloading the Software Packages

There is a different NetScaler SD-WAN software package for each appliance model. You will need to download the appropriate software package for each appliance model you want to include in your network.

To download the NetScaler SD-WAN software packages, go to the following URL:

http://www.citrix.com/downloads.html

Instructions for downloading the software are provided on this site.

NetScaler SD-WAN Software Packages and Appliance Models

Aug 09, 2017

NetScaler SD-WAN Software Packages

There is a different Citrix NetScaler SD-WAN software package for each supported SD-WAN appliance model. You will need to acquire the appropriate package for each appliance model you plan to incorporate into your network.

Supported SD-WAN Appliance Models

There are two main categories of SD-WAN Appliances:

- NetScaler SD-WAN Appliance hardware models
  - Standard, WANOP, Standard Edition, and Enterprise Edition
- NetScaler SD-WAN VPX Virtual Appliances (NetScaler SD-WAN VPX)
  - Standard Edition and WANOP Edition

> **Note**
>
> All SD-WAN Appliance models in a Virtual WAN environment are required to be running the same SD-WAN firmware release. For additional information, please contact NetScaler SD-WAN Customer Support.

For a complete description of NetScaler SD-WAN Appliances, please refer to the following NetScaler SD-WAN datasheet:

https://www.citrix.com/content/dam/citrix/en_us/documents/data-sheet/netscaler-sd-wan-datasheet.pdf

NetScaler SD-WAN Standad Edition Hardware Appliances

Citrix NetScaler SD-WAN 9.2 supports the following SD-WAN standard edition hardware appliance models:

SD-WAN SE MODEL	ROLE
SD-WAN 410-SE	Small branch node appliance
SD-WAN 1000-SE	Large branch node appliance
SD-WAN 2000-SE	Large branch node appliance
SD-WAN 2100-SE	Large branch node appliance
SD-WAN 4100-SE	Data Center Master Control Node (MCN) appliance
SD-WAN 5100-SE	Data Center Master Control Node (MCN) appliance

The below figure illustrates the 410-SE appliance model.

The below figure illustrates NetScaler SD-WAN 5100-SE appliance model.

NetScaler SD-WAN WAN Optimization Hardware Appliances (SD-WAN WANOP)

Citrix NetScaler SD-WAN 9.2 supports the following SD-WAN WAN Optimization (WANOP) appliance models:

SD-WAN WANOP MODELS	ROLE
SD-WAN WANOP 800	Small branch node appliance
SD-WAN WANOP 1000 Windows Server	Large branch node appliance
SD-WAN WANOP 1000	Large branch node appliance
SD-WAN WANOP 2000	Large branch node appliance
SD-WAN WANOP 3000	Large branch node appliance
SD-WAN WANOP 4000	Data Center appliance
SD-WAN WANOP 5000	Data Center appliance

NetScaler SD-WAN VPX Virtual Appliances (SD-WAN VPX-SE)

Citrix NetScaler SD-WAN 9.2 supports the following SD-WAN VPX Virtual Appliance (VPX-SE) models:

SD-WAN VPX-SE MODELS	ROLE
SD-WAN VPX 20-SE	MCN or client node, small branch
SD-WAN VPX 50-SE	MCN or client node, small branch
SD-WAN VPX 100-SE	MCN or client node, small branch
SD-WAN VPX 200-SE	MCN or client node, small branch
SD-WAN VPX 500-SE	MCN or client node, small branch
SD-WAN VPX 1000-SE	MCN or client node, small branch

For more information, see About SD-WAN Virtual VPX.

NetScaler SD-WAN WANOP Virtual Appliances (SD-WAN VPX-WANOP)

Citrix NetScaler SD-WAN 9.2 supports the following SD-WAN WANOP Virtual Appliance (VPX-WANOP) models:

SD-WAN VPX WANOP MODELS	ROLE
SD-WAN WANOP VPX-2	Small branch node
SD-WAN WANOP VPX-6	Small branch node
SD-WAN WANOP VPX-10	Small branch node
SD-WAN WANOP VPX-20	Small branch node
SD-WAN WANOP VPX-50	Large branch node
SD-WAN WANOP VPX-100	Large branch node
SD-WAN WANOP VPX-200	Large branch node

NetScaler SD-WAN Enterprise Edition Hardware Appliances (SD-WAN EE)

Citrix NetScaler SD-WAN 9.2 supports the following SD-WAN Enterprise Edition appliance (SD-WAN EE) models:

SD-WAN EE MODELS	ROLE
SD-WAN 1000-EE	Large branch, Data Center appliance
SD-WAN 2000-EE	Large branch, Data Center appliance

NetScaler SD-WAN Appliance Packages

Aug 09, 2017

A SD-WAN appliance package contains the SD-WAN software package for a particular appliance model, bundled with a specific SD-WAN configuration package. The two packages are bundled together and distributed to the clients by means of the **Change Management** wizard in the Management Web Interface running on the Master Control Node (MCN).

If this is an initial installation, you must manually upload, stage, and activate the appropriate appliance package on each of the client appliances that will reside in your SD-WAN network. If you are updating the configuration for an existing SD-WAN deployment, the MCN automatically distributes and activates the appropriate appliance package on each of the existing clients, as soon as the virtual paths to the clients become operational.

Preparing for Your Deployment

Aug 09, 2017

It is strongly recommended that before beginning the installation, you first read through the Citrix CloudBridge Virtual WAN Deployment Planning Guide. This article discusses the essential Virtual WAN concepts and features, and provides guidelines for planning your deployment. You can find this document in the CloudBridge documentation section on the Citrix Documentation Portal (http://docs.citrix.com/).

Summary of Installation and Deployment Procedures

The following list outlines the steps and procedures involved in deploying the NetScaler SD-WAN Standard and Enterprise Editions.

1. Gather your NetScaler SD-WAN deployment information.
2. Set up the NetScaler SD-WAN Appliances.

   For each hardware appliance you want to add to your SD-WAN deployment, you must complete the following tasks:

   a. Set up the appliance hardware.

   b. Set the Management IP Address for the appliance and verify the connection.

   c. Set the date and time on the appliance.

   d. (Optional) Set the console session **Timeout** interval to a high or the maximum value.

   e. Upload and install the software license file on the appliance.

3. Set up the Master Control Node (MCN) site.

   a. Switch the Management Web Interface to MCN Console mode.

   b. Add and configure the MCN site.

   c. Configure the Virtual Interface Groups for the MCN site.

   d. Configure the Virtual IP Addresses for the MCN site.

   e. (Optional) Configure the GRE Tunnels for the MCN site.

   f. Configure the WAN Links for the MCN site.

   g. Configure the Routes for the MCN site.

   h. (Optional) Configure High Availability (HA) for the MCN site.

   i. (Optional) Configure Virtual WAN security and encryption.

   j. Name and save the MCN site configuration.

4. Set up the branch sites.

   a. Add the branch site.

b. Configure the Virtual Interface Groups for the branch site.

c. Configure the Virtual IP Addresses for the branch site.

d. (Optional) Configure the GRE Tunnels for the branch site.

e. Configure the WAN Links for the branch site.

f. Configure the Routes for the branch site.

g. (Optional) Configure High Availability (HA) for the branch site.

h. (Optional) Clone the new branch site to create and configure additional sites.

> **Note**
>
> Cloning the branch site is optional. The SD-WAN appliance models must be the same for both the original and the cloned sites. You cannot change the specified appliance model for a clone. If the appliance model is different for a site, you must manually add the site, by repeating steps (a) through (f).

i. Resolve any configuration Audit Alerts.

j. Save the new configuration.

5. Configure the Virtual Paths and Virtual Path Service between the MCN and the client sites.
6. (Optional/provisional) If your license includes WAN Optimization, enable and configure WAN Optimization.
   a. Enable WAN Optimization and configure the default Features settings.
   b. Configure the default Tuning Settings.
   c. Configure the default Application Classifiers.
   d. Configure the default Service Classes.
   e. Configure WAN Optimization for the branch sites.
7. Prepare the SD-WAN Appliance Packages on the MCN.
   a. Export the new configuration package to **Change Management** on the MCN.
   b. Generate and stage the Appliance Packages on the MCN.
8. Connect the client appliances to your network.
9. Install the SD-WAN Appliance Packages on the clients.
10. Enable the SD-WAN Service on each of the SD-WAN appliances in your network.
11. Use the **Monitoring** pages to verify the activation and check for any existing or potential configuration issues.

Installation and Configuration Information Checklist

Aug 09, 2017

Gather the following information for each NetScaler SD-WAN site you want to deploy:

- The licensing information for your product
- Required Network IP Addresses for each appliance to be deployed:
  * Management IP Address
  * Virtual IP Addresses
- Site Name
- Appliance Name (one per site)
- SD-WAN Appliance Model (for each appliance to be deployed)
- Deployment Mode (MCN or Client)
- Topology
- Gateway MPLS
- GRE Tunnel information
- Routes
- VLANs
- Bandwidth at each site for each circuit

Getting Started by Using NetScaler SD-WAN

Aug 09, 2017

Refer to the following sections to help you get familiarized with using NetScaler SD-WAN web interface, installing required appliance packages, connecting appliances, and setting up the SD-WAN network.

- NetScaler SD-WAN Management Web Interface
- Installing the Virtual WAN Appliance Packages
- Preparing the Virtual WAN Appliance Packages
- Connecting the Client Appliances to Your Network
- Setting up the SD-WAN Appliances

NetScaler SD-WAN Management Web Interface

Jan 20, 2018

This section provides basic navigation instructions, and a navigation roadmap of the Management Web Interface page hierarchy. Also provided are specific navigation instructions for the **Configuration Editor** and **Change Management wizard**.

Basic Navigation

The below figure outlines the basic navigation elements of the Management Web Interface, and the terminology used in this guide to identify them.

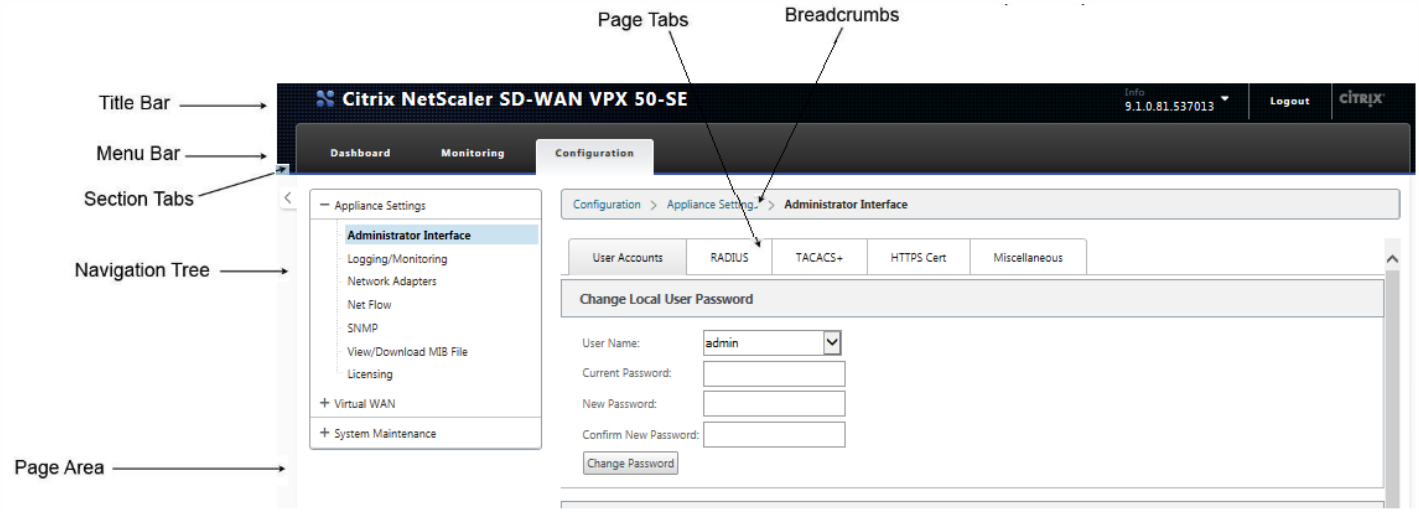

The basic navigation elements are as follows:

- **Title bar** – This is the dark grey bar at the top of all Management Web Interface screens. This displays the appliance model number, Host IP Address for the appliance, the version of the software package currently running on the appliance, and the user name for the current login session. The title bar also contains the **Logout** button for terminating the session.
- **Main menu bar** – This is the light blue bar displayed below the title bar on every Management Web Interface screen. This contains the section tabs for displaying the navigation tree and pages for a selected section.
- **Section tabs** – The section tabs are located in the blue main menu bar at the top of the page. These are the top-level categories for the Management Web Interface pages and forms. Each section has its own navigation tree for navigating the page hierarchy in that section. Click a section tab to display the navigation tree for that section.
- **Navigation tree** – The navigation tree is located in the left blue and grey pane, below the main menu bar. This displays the navigation tree for a section. Click a section tab to display the navigation tree for that section. The navigation tree offers the following display and navigation options:
  - Click a section tab to display the navigation tree and page hierarchy for that section.
  - Click + (plus sign) next to a branch in the tree to reveal the available pages for that branch topic.
  - Click a page name to display that page in the page area.

- Click – (minus sign) next to a branch item to close the branch.

- **Breadcrumbs** – This displays the navigation path to the current page. The breadcrumbs are located at the top of the page area, just below the main menu bar. Active navigation links display in blue font. The name of the current page is displayed in black bold font.

- **Page area** – This is the page display and work area for the selected page. Select an item in the navigation tree to display the default page for that item.

- **Page tabs** – Some pages contain tabs for displaying additional child pages for that topic or configuration form. These are usually located at the top of the page area, just below the breadcrumbs display. In some cases (as for the **Change Management** wizard), tabs are located in the left pane of the page area, between the navigation tree and the work area of the page.

- **Page area resizing –** For some pages, you can grow or shrink the width of the page area (or sections of it) to reveal additional fields in a table or form. Where this is the case, there will be a grey, vertical resize bar on the right border of a page area pane, form, or table. Roll your cursor over the resize bar until the cursor changes to a bi-directional arrow. Then click and drag the bar to the right or left to grow or shrink the area width.

  If the resize bar is not available for a page, you can click and drag the right edge of your browser to display the full page.

Management Web Interface Navigation Tree Hierarchy

Top Level Section Tab	Tree Level 1	Tree Level 2
Dashboard		
Monitoring		
	Statistics	
	Flows	
	Performance Reports	
	QoS Reports	
	Usage Reports	
	Availability Reports	
	Appliance Reports	
	+ WAN Optimization	
		• Connections • Compression • Usage Graph • AppFlow • Filesystem (CIFS/SMB) • Citrix (ICA/CGP) • ICA Advanced • Outlook (MAPI) • Partners
Configuration		
	+ Appliance Settings	
		• Administrator Interface • Logging/Monitoring • Network Adaptors • NetFlow • SNMP • Licensing
	+ Virtual WAN	
		• View Configuration • Configuration Editor (*MCN only*) • Change Management (*MCN only*) • Restart/Reboot Network • Enable/Disable/Purge Flows • Dynamic Virtual Paths • Virtual WAN Certificates
	System Maintenance	
		• Delete Files • Restart System • Date/Time Settings • Local Change Management • Diagnostics • Update Software • Configuration Reset • Factory Reset

The Management Web Interface Dashboard

Click the **Dashboard** section tab to display basic information for the local appliance.

The **Dashboard** page displays the following basic information for the appliance:

- System status
- Virtual Path service status
- Local appliance software package version information

The below figure_sh**o**ws a sample Master Control Node (MCN) appliance **Dashboard** display.

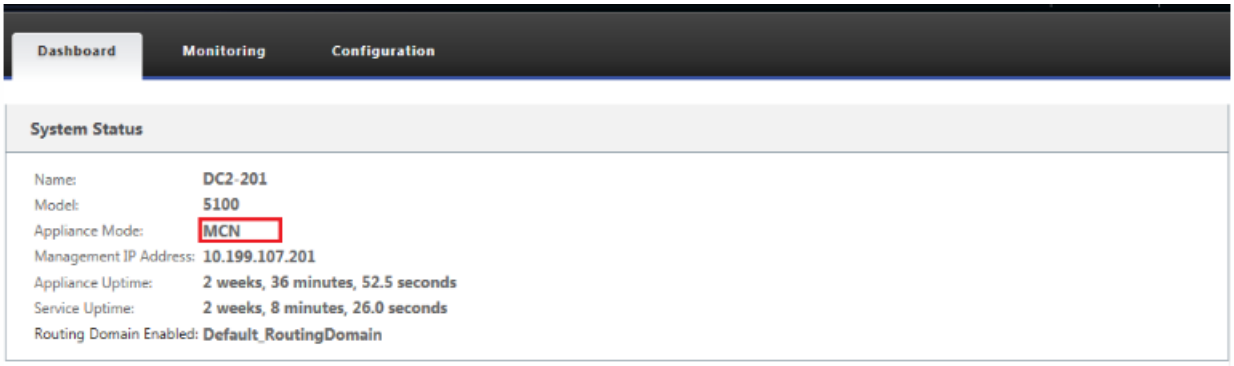

The below figure shows a sample client appliance Dashboard display.

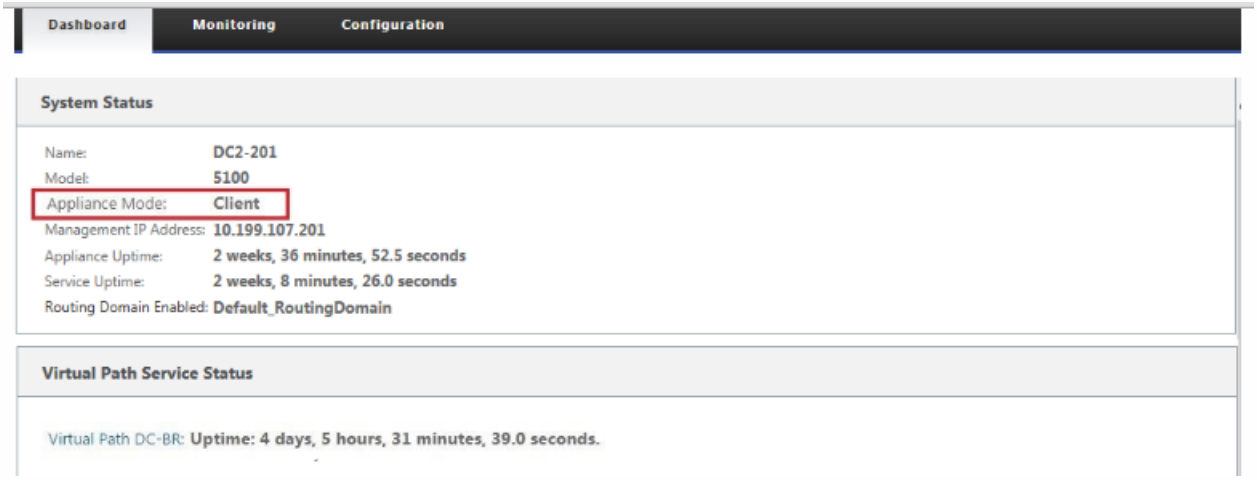

The Configuration Editor

The **Configuration Editor** enables you to add and configure CloudBridge Virtual WAN Appliance sites, connections, optimization, and provisioning, and to create and define the Virtual WAN Configuration.

The **Configuration Editor** is available when the Management Web Interface is in MCN Console mode, only. By default, the Management Web Interface on a new appliance is set to Client mode. You must change the mode setting to MCN Console before you can access the Configuration Editor. For instructions, see the section Switching the Management Web Interface to MCN Console Mode.

To navigate to the **Configuration Editor**, do the following:

1. Log into the Management Web Interface on the MCN appliance.
2. Select the **Configuration** tab.
3. In the navigation tree, click **+** next to the **Virtual WAN** branch in the tree. This displays the available pages for the **Virtual WAN** category.
4. In the Virtual WAN branch of the tree, select **Configuration Editor**.

The below figure outlines the basic navigation and page elements of the **Configuration Editor**, and the terminology used to identify them.

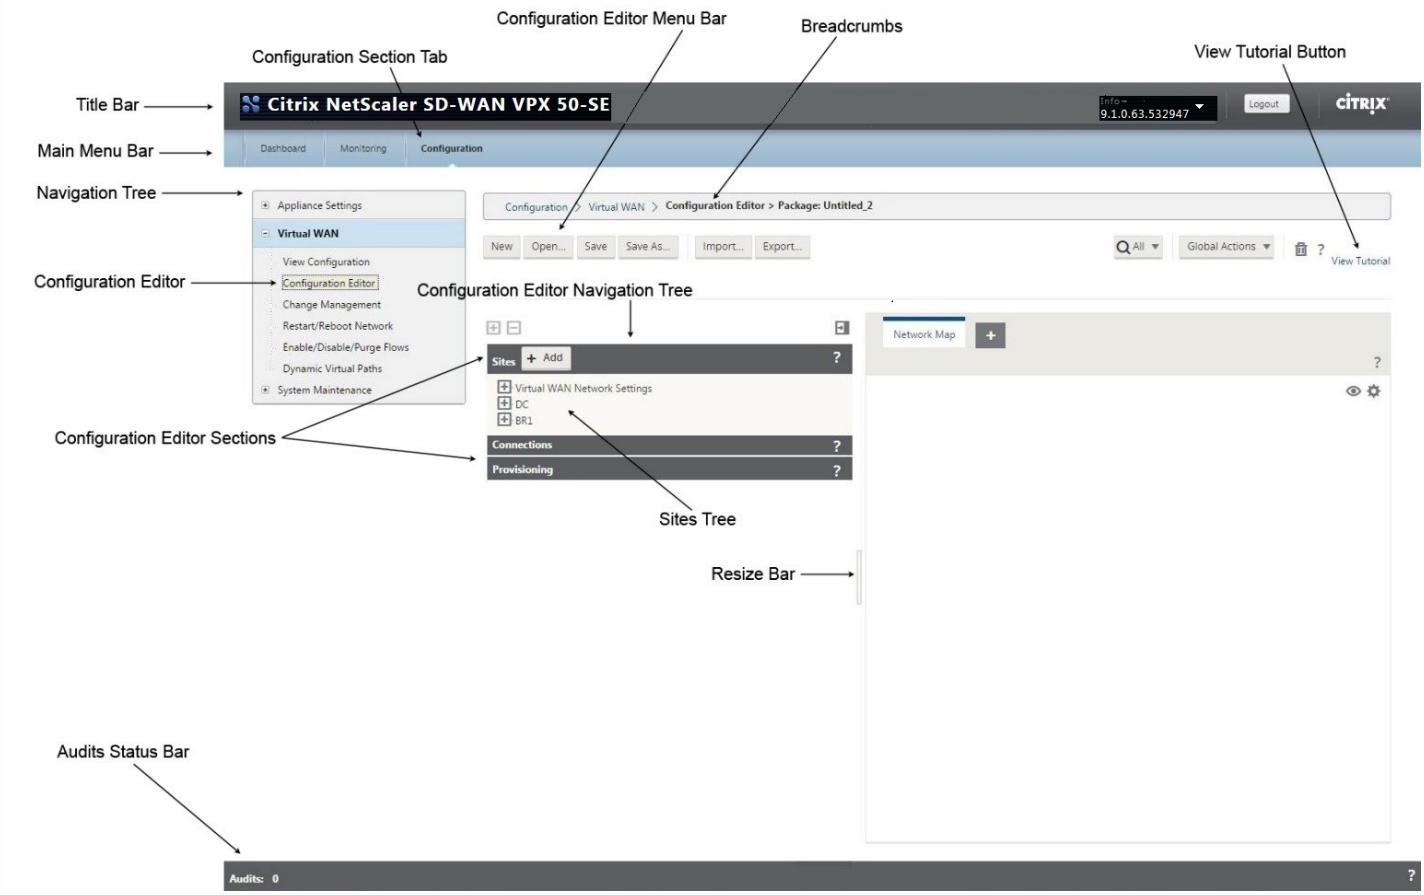

The following describes the primary **Configuration Editor** navigation elements referenced in this guide:

- **Configuration Editor menu bar** – This is located at the top of the page area, just below the breadcrumbs links. The menu bar contains the primary activity buttons for **Configuration Editor** operations. In addition, at the far right edge of the menu bar is the **View Tutorial** link button for initiating the **Configuration Editor** tutorial. The tutorial steps you through a series of bubble descriptions for each element of the **Configuration Editor** display.
- **Configuration Editor sections tree** – This is the stack of dark grey bars located in the left pane of the **Configuration Editor** page area. Each grey bar represents a top-level section. Click **+** at the left of a section name to reveal the sub-branches for that section.
- **Sections tree branches** – Click **+** (plus sign) at the left of a section name in the sections tree to open a section branch. Click – (minus sign) to close a branch. Each section branch contains one or more sub-branches of configuration categories and forms, which in turn may contain additional child branches and forms.
- **Sites tree** – This lists the site nodes that have been added to the configuration currently opened in the **Configuration Editor**. In the section tree, click **+** at the left of **Sites** to open the **Sites** tree. Click **+** to the left of a site name to open the branch for that site. Click – (minus sign) to close a branch. For detailed instructions on navigating and using the **Sites** tree and configuration forms, see the following sections:
  - Setting up the Master Control Node (MCN) Site.
  - Adding and Configuring the Branch Sites.
- **Audits status bar** – This is the dark grey bar at the bottom of the **Configuration Editor** page, and spanning the entire width of the Management Web Interface screen. The **Audits** status bar is available only when the **Configuration**

Editor is open. An Audit Alert icon (red dot or goldenrod delta) at the far left of the status bar indicates one or more errors present in the currently-opened configuration. Click the status bar to display a complete list of all unresolved Audit Alerts for that configuration.

- **Resize bar** – The resize bar is the thin, grey, vertical bar located on the right border of the main page area pane, and is available in most of the **Configuration Editor** pages. You can use the resize bar to grow or shrink the width of the page area to reveal or truncate content in a table, tree, or form. Roll your cursor over the resize bar until the cursor changes to a bi-directional arrow. Then click and drag the bar to the right or left to grow or shrink the area width.

  If the resize bar is not available for a page area, you can click and drag the right edge of your browser to display the full page.

Change Management Wizards

The **Change Management** wizards guide you through the process of uploading, downloading, staging, and activating the CloudBridge Virtual WAN software and configuration on the Master Control Node (MCN) appliance and client appliances. There are two versions of the **Change Management** wizard, one for Virtual WAN system-wide ("global") change management, and one for local change management, as follows:

- **MCN (Global) Change Management wizard – The MCN Global Change Management** wizard is the primary (main) version, and is available in the MCN Management Web Interface, only. Use this to generate the Virtual WAN Appliance Packages to be deployed for each type of Virtual WAN Appliance in your network. You can also use the wizard to automatically propagate configuration changes to Virtual Appliances already deployed in your Virtual WAN. Basic navigation instructions are provided in the section, "Using the MCN Global Change Management Wizard" below. Instructions for using the MCN global **Change Management** wizard to create the Appliance Packages are provided in the section Preparing the Virtual WAN Appliance Packages on the MCN.

- **Local Change Management wizard – The Local Change Management** wizard is available in the Management Web Interface running on both the MCN and on all client node appliances. Use this to upload, stage, and activate the appropriate Virtual WAN Appliance Package on a local appliance to be added to your Virtual WAN. You can also use this wizard to upload an updated Appliance Package specifically to the local MCN, or to an individual, local Virtual WAN Appliance already deployed in your network. Instructions for navigating and using the **Local Change Management** wizard are provided in the section Installing the Virtual WAN Appliance Packages on the Clients.

Using the MCN Global Change Management Wizard

To open the MCN Global **Change Management** Wizard, do the following:

1. Log into the Management Web Interface on the MCN appliance.
2. Select the **Configuration** tab.
3. In the navigation tree, click **+** next to the **Virtual WAN** branch in the tree.
4. In the **Virtual WAN** branch. select **Change Management**.

This displays the first page of the **Change Management** wizard, the **Change Process Overview** page, as shown in the below figure.

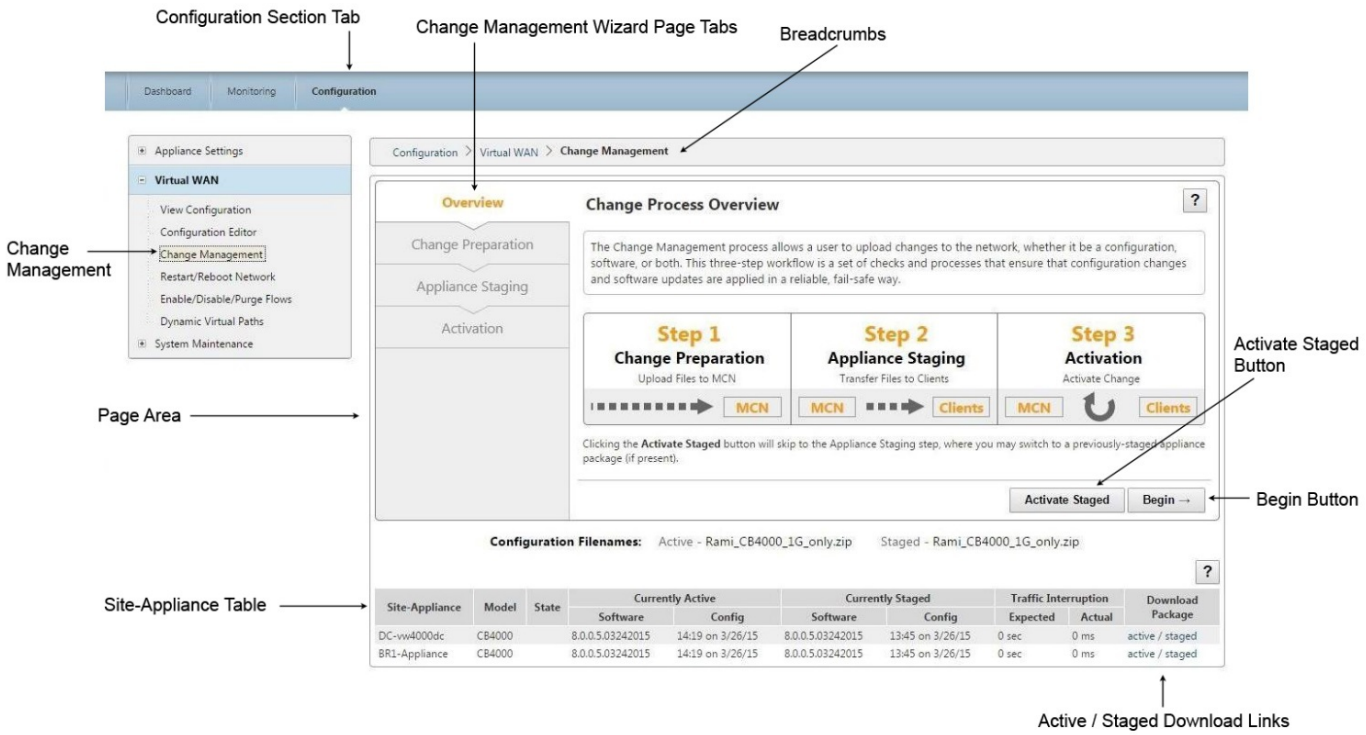

5. To start the wizard, click Begin.

For complete instructions on using the wizard to upload, stage, and activate the SD-WAN software and configuration on the appliances, see the following sections:

- Preparing the Virtual WAN Appliance Packages on the MCN
- Installing the Virtual WAN Appliance Packages on the Clients

The **Change Management** wizard contains the following navigation elements:

- **Page area** – This displays the forms, tables, and activity buttons for each page of the **Change Management** wizard.
- **Change Management wizard page tabs** – The page tabs are located in the left pane of the page area on each page of the wizard. Tabs are listed in the order that the corresponding steps occur in the wizard process. When a tab is active, you can click it to return to a previous page in the wizard. If a tab is active, the name displays in blue font. Grey font indicates an inactive tab. Tabs are inactive until all dependencies (previous steps) have been fulfilled without error.
- **Appliance-Site table** – This is located at the bottom of the wizard page area, on most wizard pages. The table contains information about each configured appliance site, and links for downloading the active or staged Appliance Packages for that appliance model and site. A package in this context is a Zip file bundle containing the appropriate NetScaler SD-WAN software package for that appliance model, and the specified configuration package. The **Configuration Filenames** section above the table shows the package name for the current active and staged packages on the local appliance.
- **Active/Staged download links** – These are located in the **Download Package** field (far right column) of each entry in the **Appliance-Site** table. Click a link in an entry to download the active or staged package for that appliance site.

- **Begin button** – Click **Begin** to initiate the **Change Management** wizard process and proceed to the **Change Preparation** tab page.
- **Activate Staged button** – If this is not an initial deployment, and you want to activate the currently staged configuration, you have the option of proceeding directly to the **Activation** step. Click **Activate Staged** to proceed directly to the Activation page and initiate activation of the currently staged configuration.

One Touch Start

Aug 09, 2017

Once touch start feature is introduced in SD-WAN 9.3 release, this allows you to easily and quickly configure your SD-WAN appliance as a Client on first time start up.

The one touch start option is displayed when your appliance boots up for the first time.

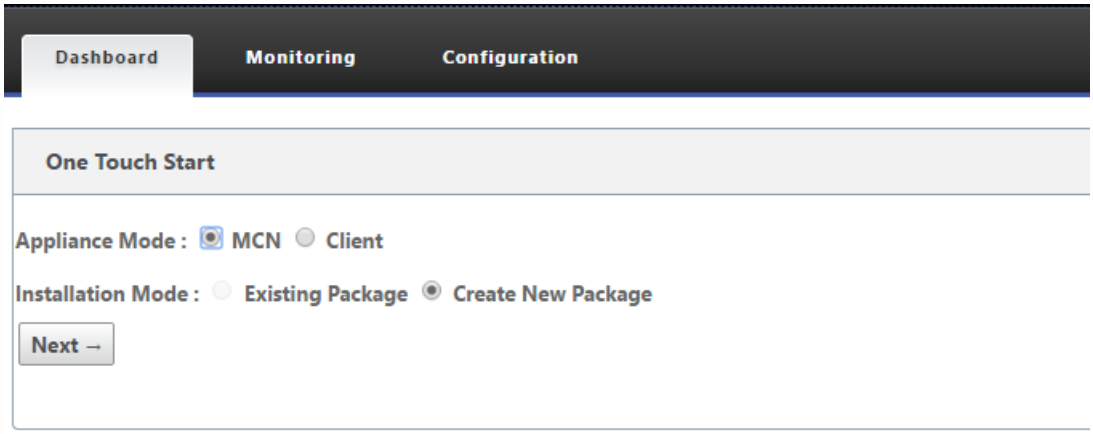

> ## Note
>
> For configuring the SD-WAN appliance as an MCN, create a new configuration or import an existing configuration using the **Configuration Editor**. For more information see, Preparing the SD-WAN Appliance Packages on the MCN.

To configure your SD-WAN appliance as a client using an existing configuration file:

1. Select **Client** as the appliance mode.
2. Select **Existing Package** installation mode. Administrator must periodically save the configuration of the MCN in order to make use of an existing package of the MCN.
3. Click **Choose File** to select the configuration package from your local computer.
4. Click **Upload and Install**.

To configure your SD-WAN appliance as a client using Local Change Management:

1. Select **Client** as the appliance mode.
2. Select **Create New Package** to upload the configuration package for this appliance using Local change management. The package can be downloaded from the MCN appliance from the change Management screen.
3. Click **Next**.

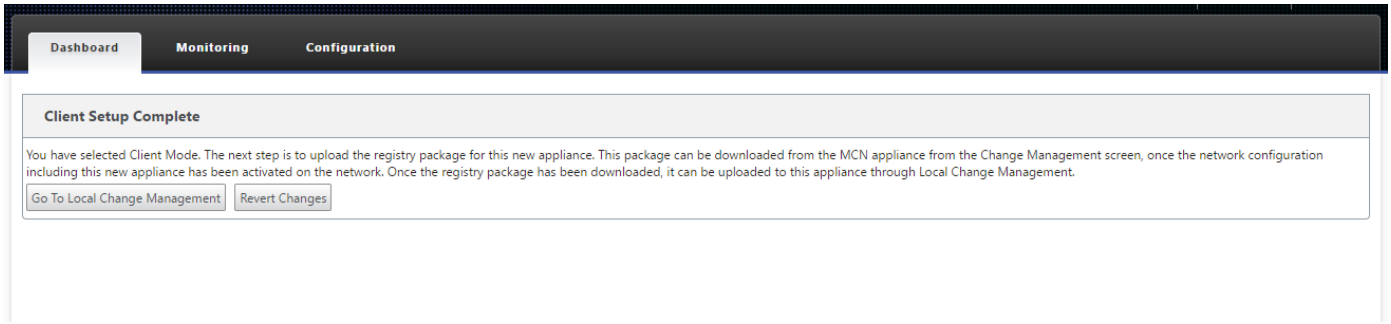

4. Click **Go To Local Change Management**.
5. Follow the procedure in the topic Installing the SD-WAN Appliance Packages on the Clients.

Installing the SD-WAN Appliance Packages on the Clients

Aug 09, 2017

After you have prepared the Appliance Packages and connected the MCN, and the branch Site Administrators have connected their respective client appliances to the LAN and WAN, the next step is to upload and activate the appropriate SD-WAN Appliance Package on each client. The Change Management wizard guides you through this process.

To install and activate the software and configuration on a client appliance, do the following

1. On a connected PC, open a browser and log onto the MCN appliance Management Web Interface.

   Enter the Management IP Address for the MCN in the browser address field. This displays the Management Web Interface **Dashboard** page for the MCN appliance.

2. Select the **Configuration** tab.

3. In the navigation pane on the left, select **Virtual WAN** and then select **Change Management**.

   This displays the **Change Process Overview** page (the first page of the **Change Management** wizard).

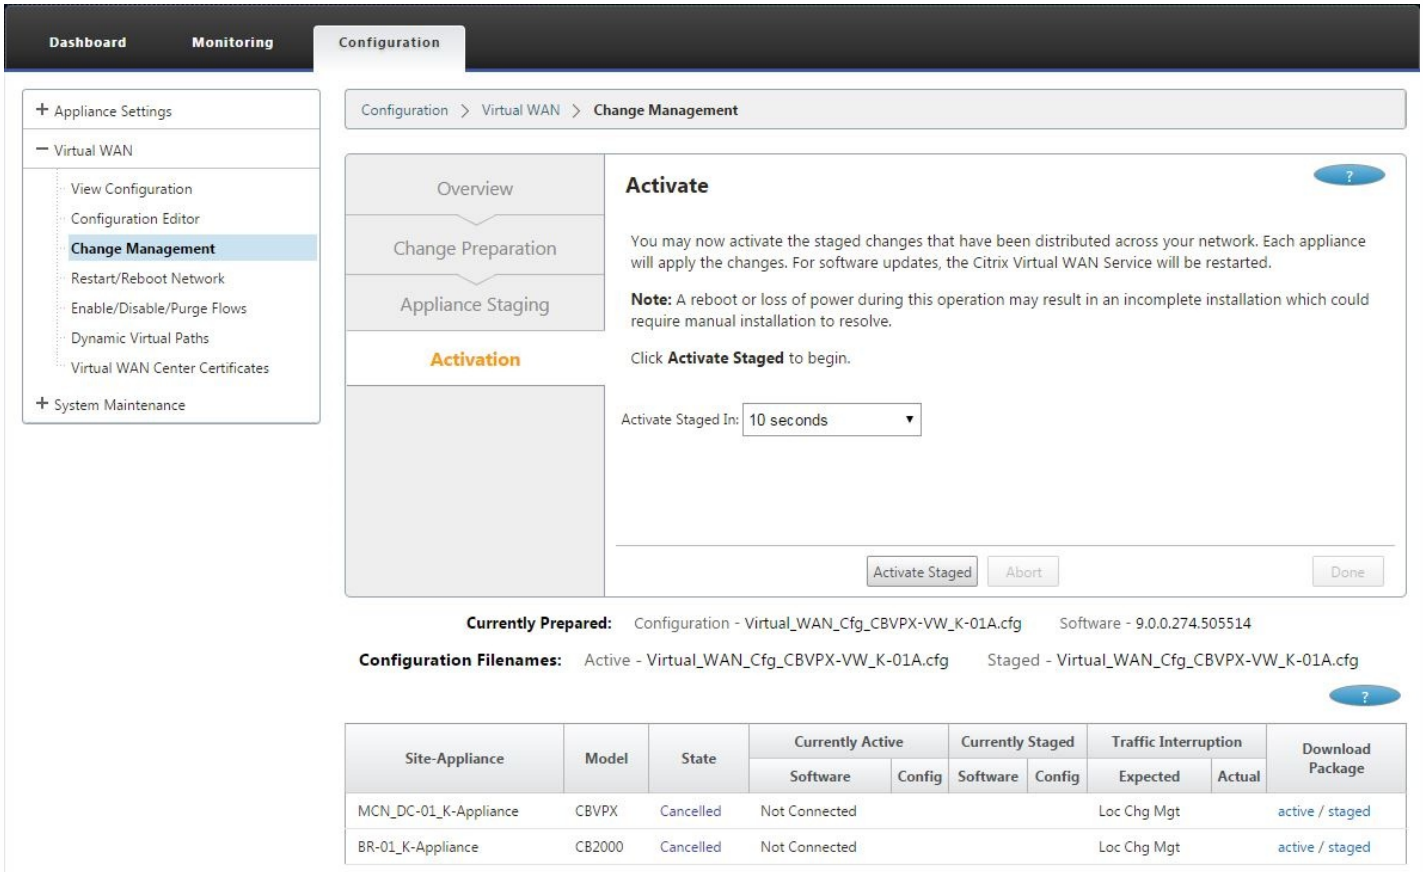

At the bottom of this page, you will see a table listing the individual sites and appliances. At the far right of the table in the **Download Package** column, are links for the **Active** (if available) and **Staged** Appliance Packages.

> **Note**
>
> If this is an initial installation, the **Active** links are not yet available, and are replaced by a plain text marker **none**.

4. Click the **Staged** link for the package you want to download.

In the **Site-Appliance** table, locate the entry for your site appliance, and click the **Staged** link in the **Download Package** column of that entry. A file browser for selecting the download location (on the local PC) displays.

5. Select the download location and click **OK.**
6. (Optional.) After the download completes, log out of the MCN Management Web Interface.
7. Open a browser, and enter the IP Address for the client to which you want to upload the Appliance Package Zip file.

> **Note**
>
> Please ignore any browser certificate warnings for the CloudBridge Management Web Interface.

This opens the NetScaler SD-WAN Management Web Interface Login screen on the client appliance.

8. Enter the Administrator user name and password and click **Login**.

The default Administrator user name is *admin*; the default password is *password*.

This displays the Management Web Interface **Dashboard** page for the client appliance.

Note

If this is an initial installation, or if you have temporarily disabled the Virtual WAN Service on this appliance, you will see a goldenrod Audit Alert icon with a status message indicating that the Virtual WAN Service is currently inactive or disabled. You can ignore this alert for now. The alert will remain on the **Dashboard** page until you manually start the service, after completing the installation.

9. Select the **Configuration** tab.
10. Open the System Maintenance branch in the navigation tree (left pane), and select Local Change Management.

    This displays the **Local Appliance Change Process Upload** page for uploading an Appliance Package.

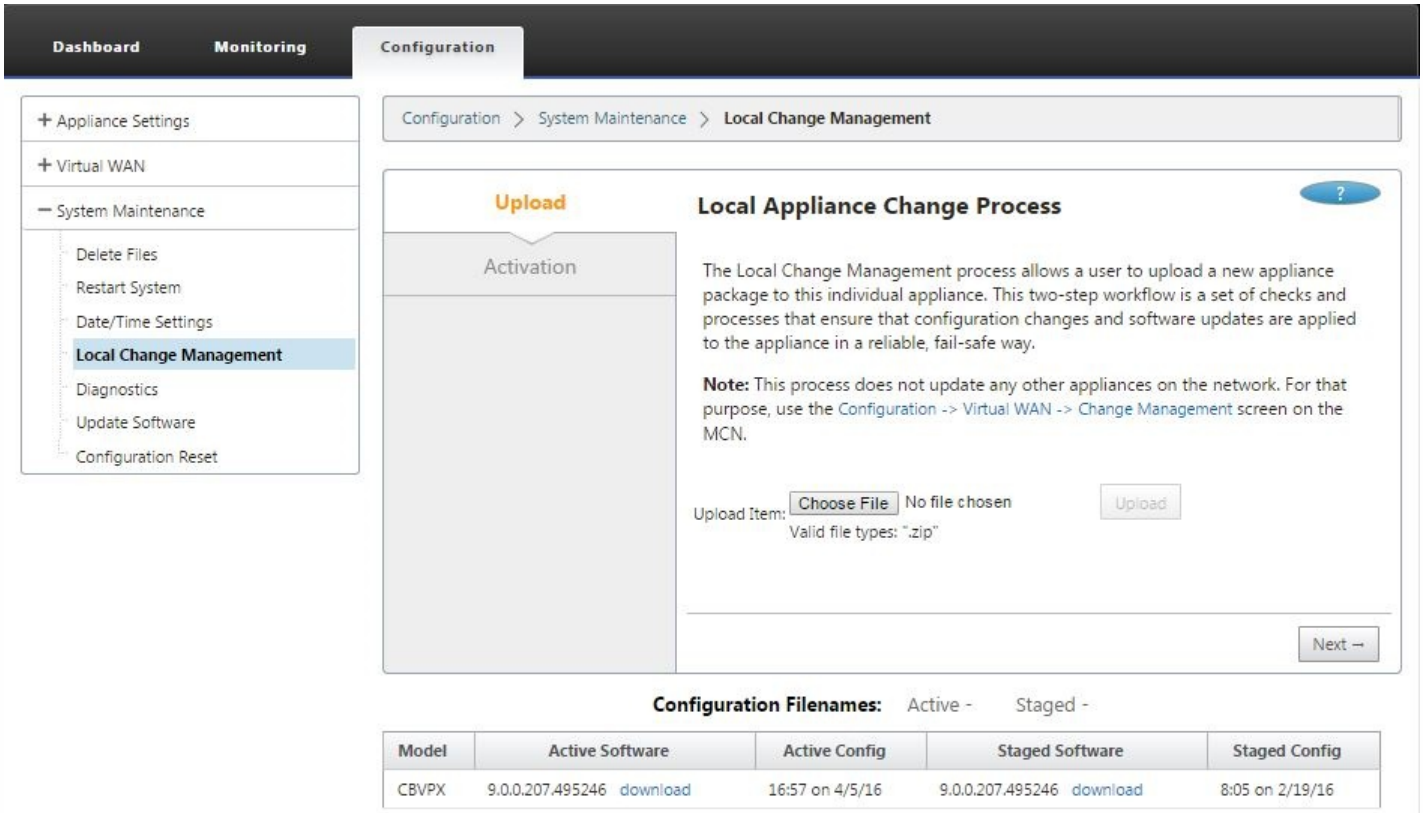

11. Click **Choose File** next to the **Upload Item** label.

    This opens a file browser for selecting the Appliance Package you want to upload to the client.

12. Navigate to the SD-WAN Appliance Package Zip file you just downloaded from the MCN, select it, and click **OK**.

13. Click **Upload**.

The upload process takes a few seconds to complete. When completed, a status message displays (left middle of page), stating **Upload complete.**

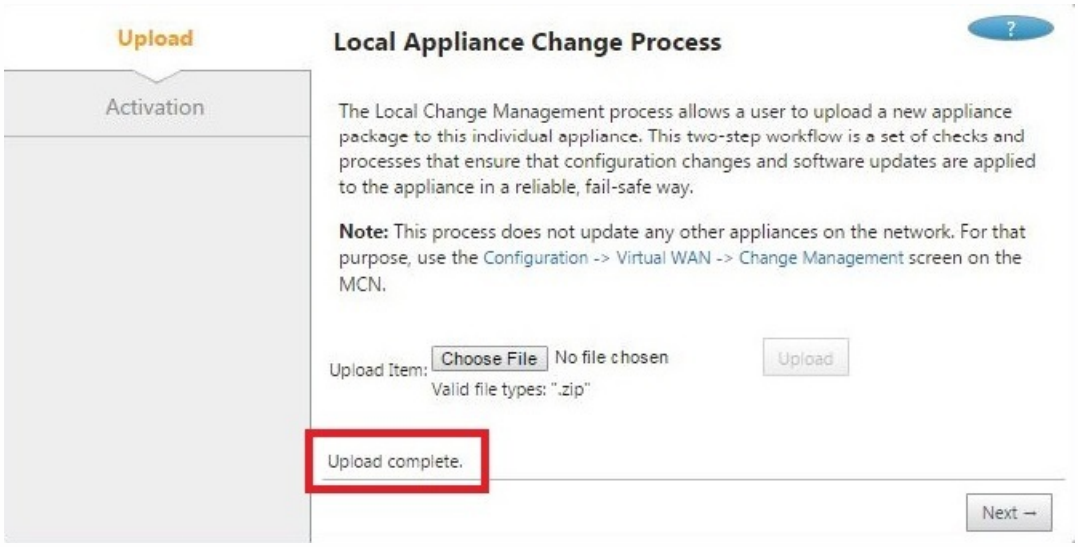

14. Click **Next**.

    This uploads the specified software package, and displays the Local Change Management **Activation** page.

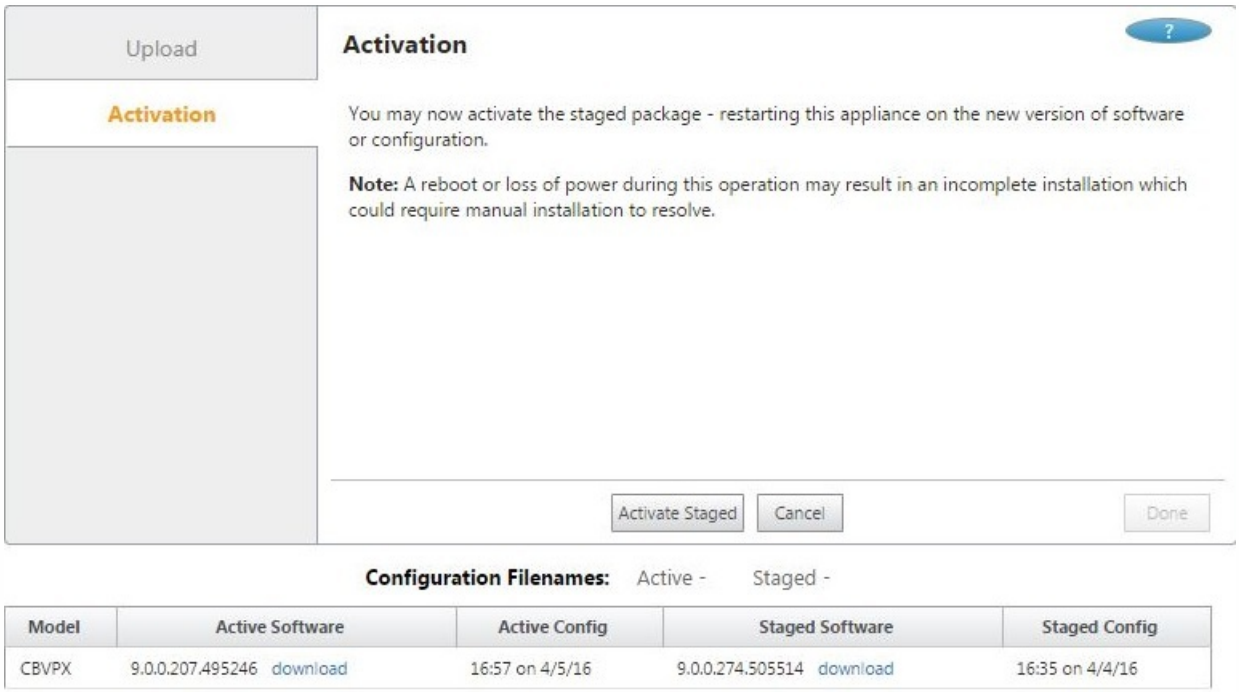

15. Click **Activate Staged**.

This displays a dialog box prompting you to confirm the activation operation.

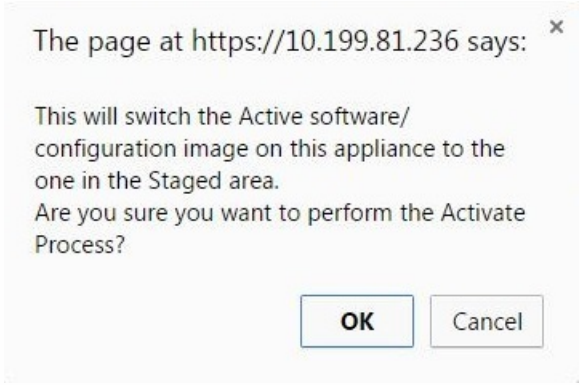

16. Click **OK**.

This activates the newly-installed package and, if this is not an initial deployment, starts the Virtual WAN Service on the client appliance. This process takes several seconds, during which a progress status message displays.

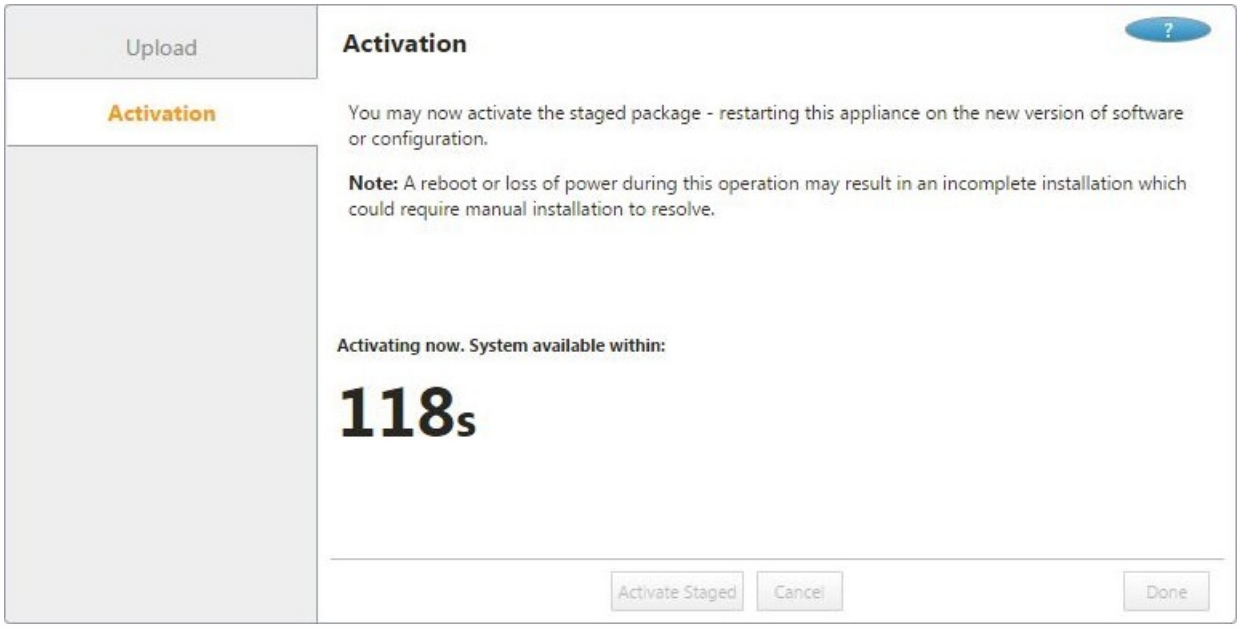

When the activation completes, a status message displays stating **Activation complete**, and the **Done** button becomes available.

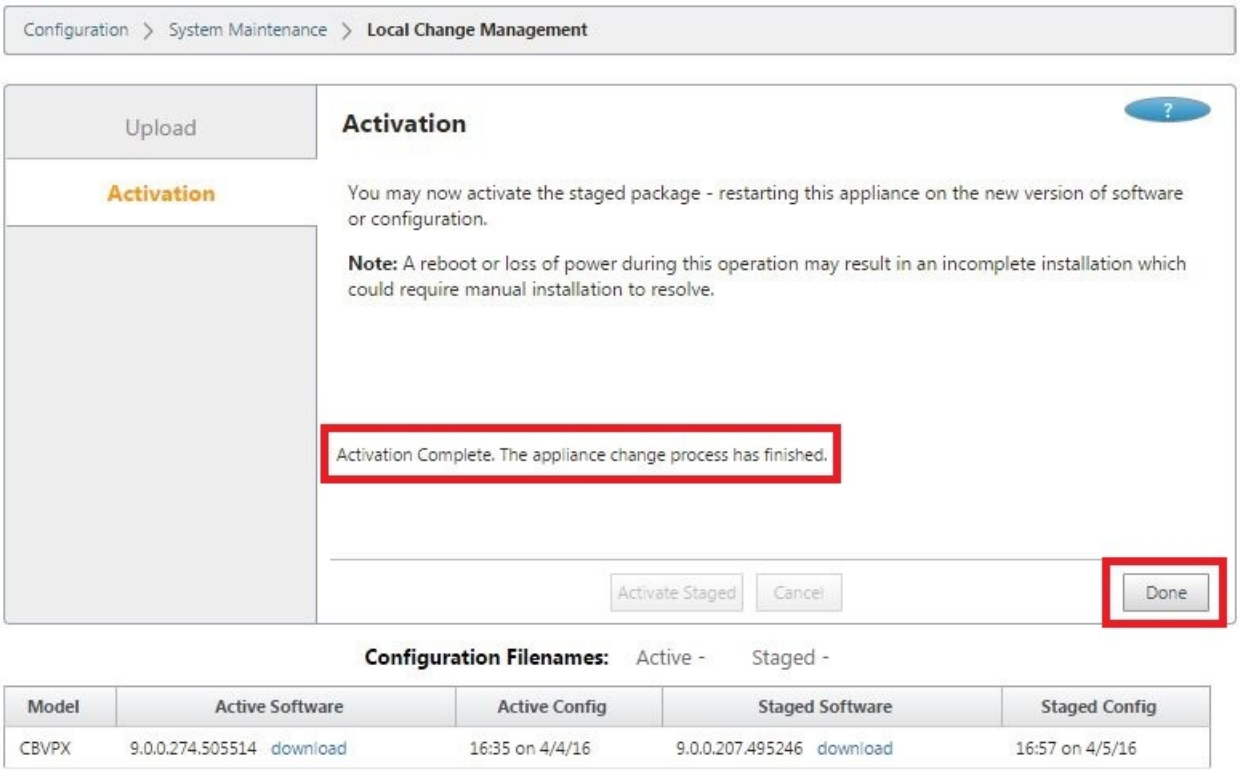

17. Click **Done** to exit the wizard and view the activation results.

After the activation completes, click **Done** on the **Activation** page to return to the Management Web Interface **Dashboard** page.

If this is not an initial deployment, this page should now display updated information for the currently active version of the software package, the OS partition, and the status of the CloudBridge Virtual Path. If this is an initial installation, there will be a goldenrod Audit Alert icon, along with a status message indicating that the Virtual WAN Service is currently inactive or disabled. In this case, you must manually enable the service, as described in Enabling the Virtual WAN Service.

The below figure shows a sample client **Dashboard** page displaying the alert icon and status message.

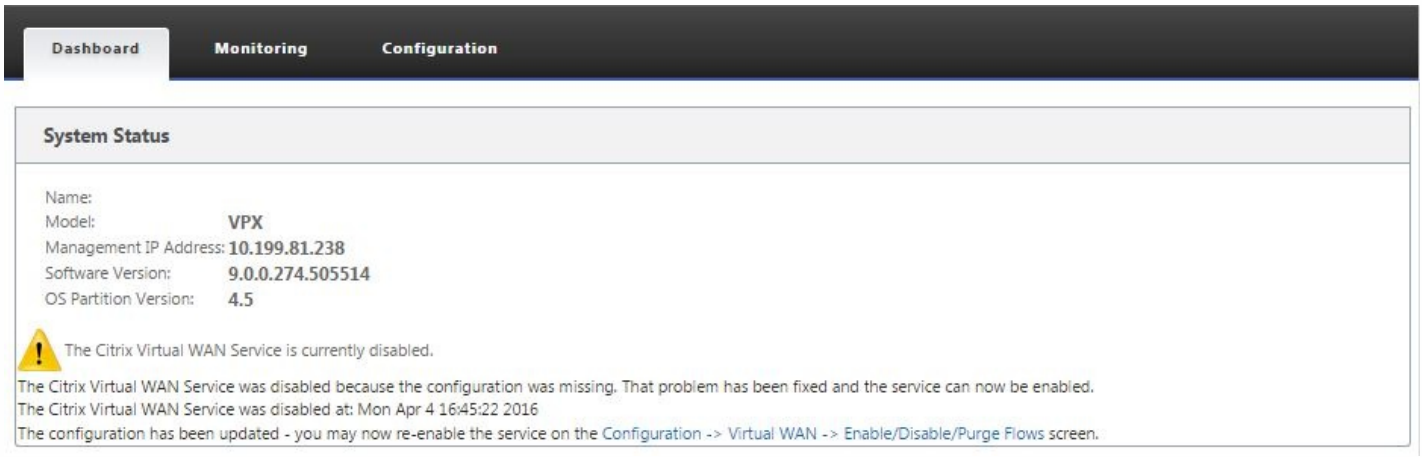

The final step to complete an initial SD-WAN deployment, is to enable the Virtual WAN Service. Instructions are provided in the section Enabling the Virtual WAN Service.

Preparing the SD-WAN Appliance Packages on the MCN

Aug 09, 2017

The next step is to prepare the SD-WAN Appliance Packages for distribution to the client nodes. This involves the following two procedures:

1. Export the Configuration Package to Change Management.

   Before you can generate the Appliance Packages, you must first export the completed configuration package from the **Configuration Editor** to the global **Change Management** staging inbox on the MCN. Instructions are provided in the section Exporting the Configuration Package to Change Management.

2. Generate and stage the Appliance Packages.

   After you have added the new configuration package to the **Change Management** inbox, you can generate and stage the Appliance Packages. To do this, you will use the **Change Management** wizard in the Management Web Interface on the MCN. Instructions are provided in the section Generating and Staging the SD-WAN Appliance Packages.

Exporting the Configuration Package to Change Management

Aug 09, 2017

Before you can generate the Appliance Packages, you must first export the completed configuration package to the Management Web Interface **Change Management** system.

To export the configuration package to **Change Management**, do the following:

1. In the **Configuration Editor** page, click **Export** (at the top of the page).

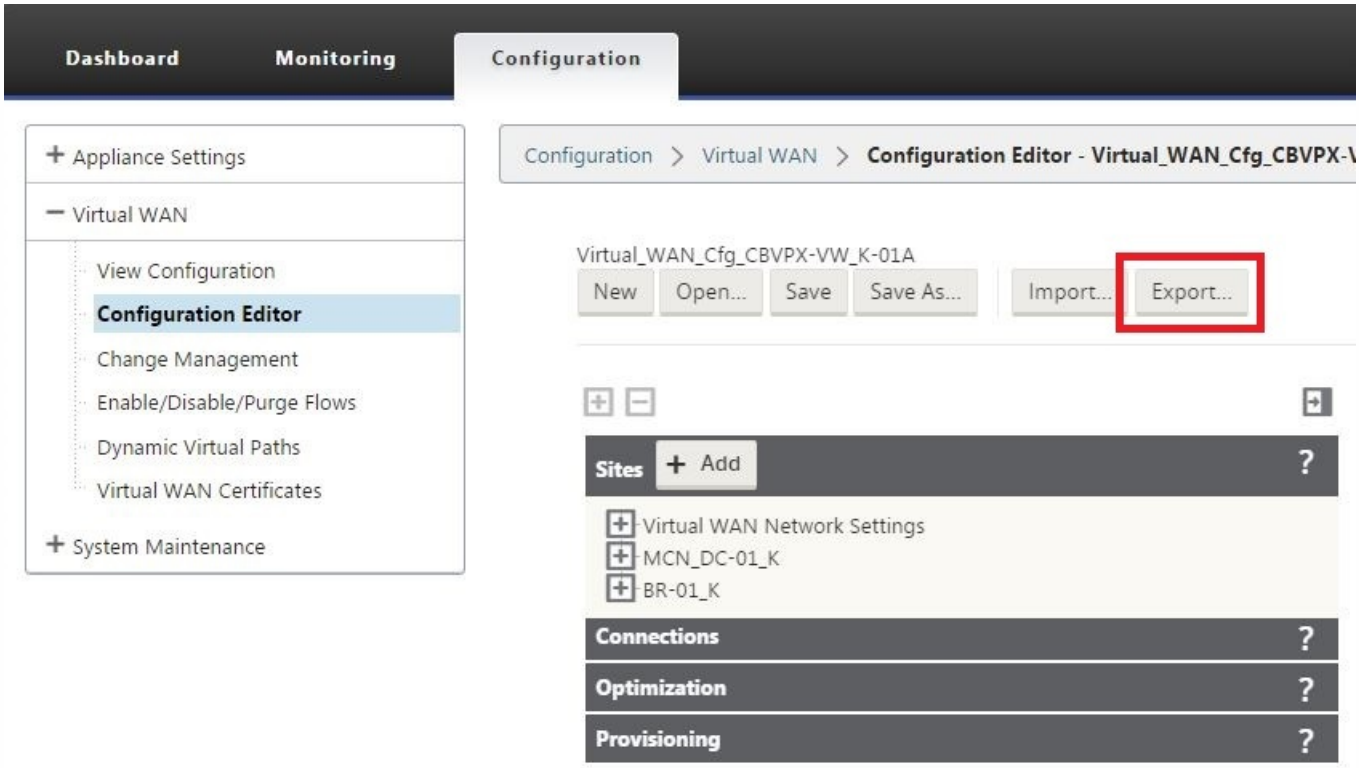

This opens the **Export Configuration** dialog box.

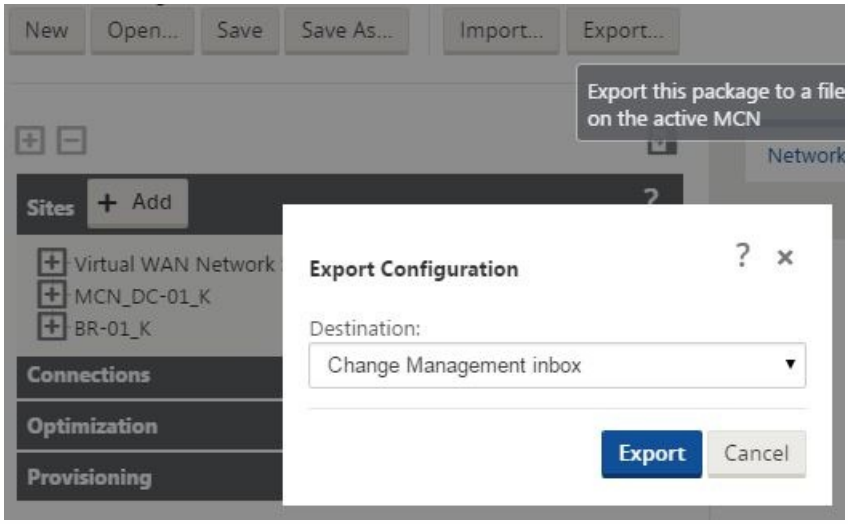

2. Select **Change Management** Inbox as the export destination.

   Use the drop-down menu in the **Destination: field to make your selection.**

3. Click **Export.**

   When the export operation completes, a green success status message displays at the top of the page.

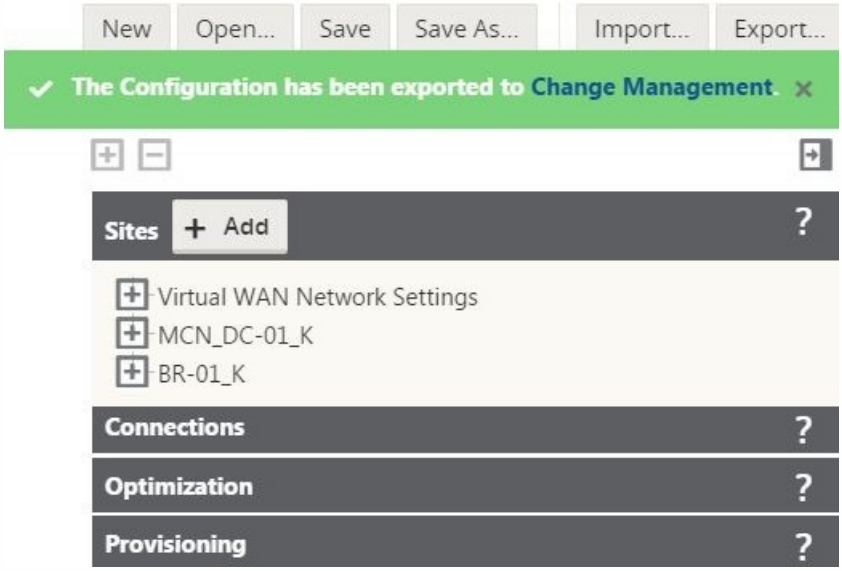

Tip

You can click the blue **Change Management** link in the success message to go directly to the **Change Preparation** -- **Upload and Verify Files** page (second page) of the **Change Management** wizard. You will need to navigate to this page to perform the next step in the configuration process. However, the success message displays for only a few seconds, after which you must use the navigation tree to open the wizard and then step through to this page. Instructions are provided in the next section.

You are now ready to upload the SD-WAN software packages to the MCN Appliance, and prepare the Appliance Packages for distribution to the client nodes.

Generating and Staging the SD-WAN Appliance Packages

Aug 09, 2017

After you have prepared the configuration using the configuration editor and exported the configuration package to the change management inbox, the next step is to prepare the SD-WAN Appliance Packages for distribution to the client nodes. To do this, use the **Change Management** wizard in the Management Web Interface on the MCN.

There is a different SD-WAN software package for each SD-WAN Appliance model. An Appliance Package consists of the software package for a specific model, bundled with the configuration package you want to deploy. Consequently, a different Appliance Package must be prepared and generated for each appliance model in your network.

> ### Note
>
> If you have not already downloaded the required SD-WAN software packages to a PC connected to your network, you will need to do so now. For information on acquiring and downloading the software, see the section Acquiring the SD-WAN Software Packages.

To upload and install the package and configuration to the MCN, do the following:

1. Log into the Management Web Interface on the MCN appliance.

> ### Note
>
> You will be uploading the software packages you previously downloaded to the connected PC. For convenience, you might want to use this same PC to connect to the MCN again.

2. Select the **Configuration** tab.
3. In the left pane, open the **Virtual WAN** section, and select **Change Management**.

   This displays the first page of the **Change Management** wizard, the **Change Process Overview** page.

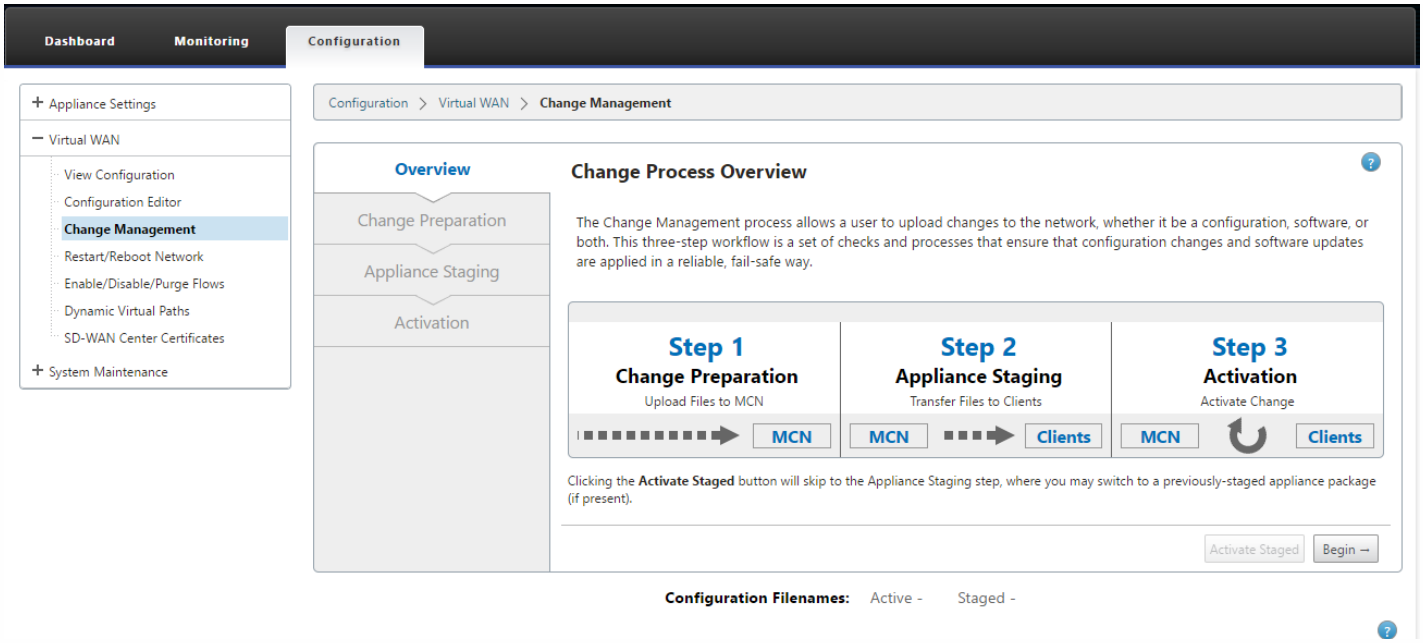

4. Click **Begin**.

This displays the **Change Preparation** page for uploading and verifying the specified configuration and software package(s).

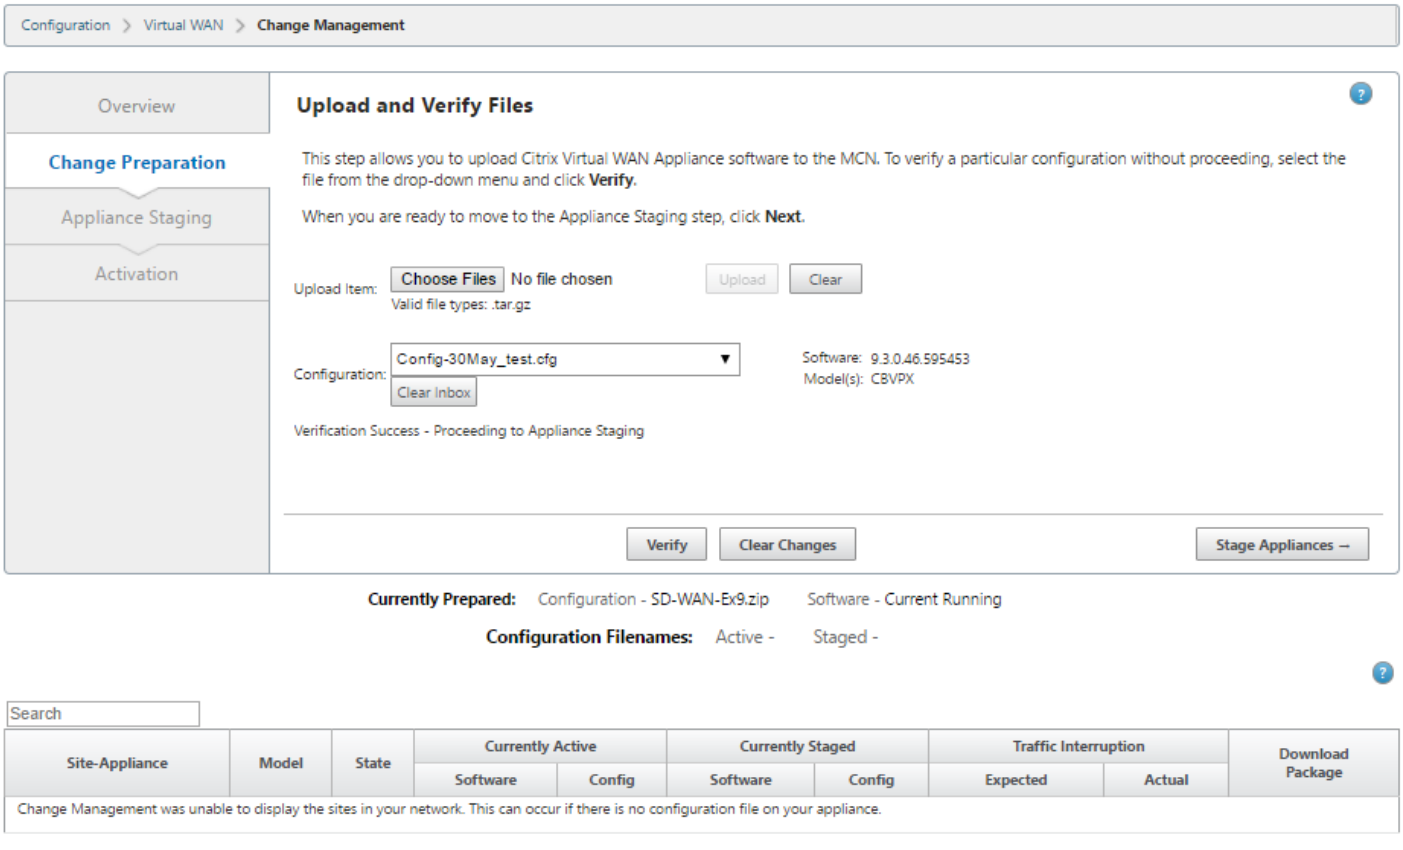

5. Upload each of the SD-WAN software packages required for your network.

Note

There is a different software package for each SD-WAN Appliance model. Before proceeding with this step, make sure you have downloaded a copy of the appropriate Virtual WAN software package for each of the different appliance models in your network. For information on downloading the software packages, see the section Acquiring the SD-WAN Software Packages.

For each SD-WAN software package you want to deploy, do the following:

a. Click **Choose File** next to the **Upload Item** field.

This opens a file browser for selecting a SD-WAN software package to upload.

b. Select a SD-WAN software package, and click **OK**.

Navigate to the SD-WAN software packages you downloaded earlier to the local PC, and select the package to upload.

c. Click **Upload**.

d. Repeat steps (a) through (c) for each of the SD-WAN software packages required for your network.

6. In the **Configuration** field drop-down menu, select the new configuration package that you just exported to **Change Management**.

7. Click **Stage Appliance**.

This initiates the following actions:

- Transfers the selected software package and configuration to the MCN.
- Generates an Appliance Package for each appliance model identified in the selected configuration.
- Adds the new Appliance Packages to the list of available packages in the Site-Appliance table.
- Stages the new configuration and appropriate software package on the MCN.

Note

If this is an initial deployment, only the MCN is updated and staged at this time. If you are updating an existing deployment and the Virtual Paths are already functioning between the deployed sites, this also distributes the appropriate Appliance Packages to the deployed client nodes, and initiates staging on those nodes. However, if you are adding new client nodes to an existing Virtual WAN deployment, you still must manually upload, stage, and activate the appropriate Appliance Package on each new client, as outlined in the remaining steps in this procedure.

The **License** page appears.

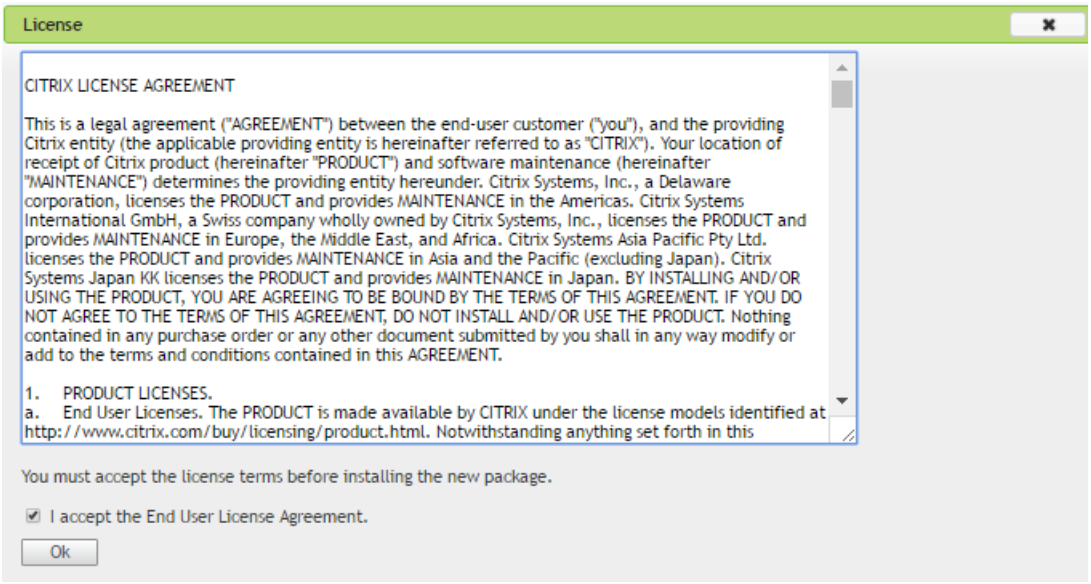

8. Select **I accept the End User License Agreement** and click **OK.**

   This dismisses the **License** page and proceeds to the **Appliance Staging** page.

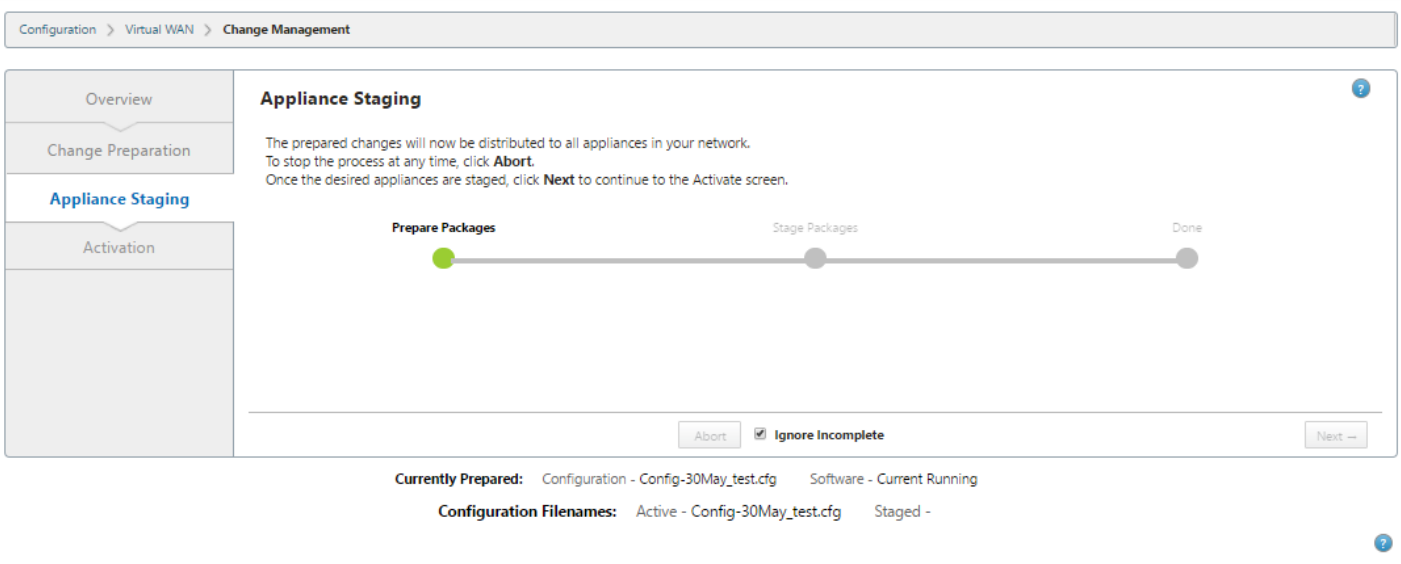

Note

Select Ignore incomplete, when adding additional sites to the network or if the site is in **not connected** state. This indicates that the client sites should be ignored for this staging operation, and only the MCN should be updated and staged.

When the staging operation completes, the **Site-Appliance** table is populated with the newly staged Appliance Packages information.

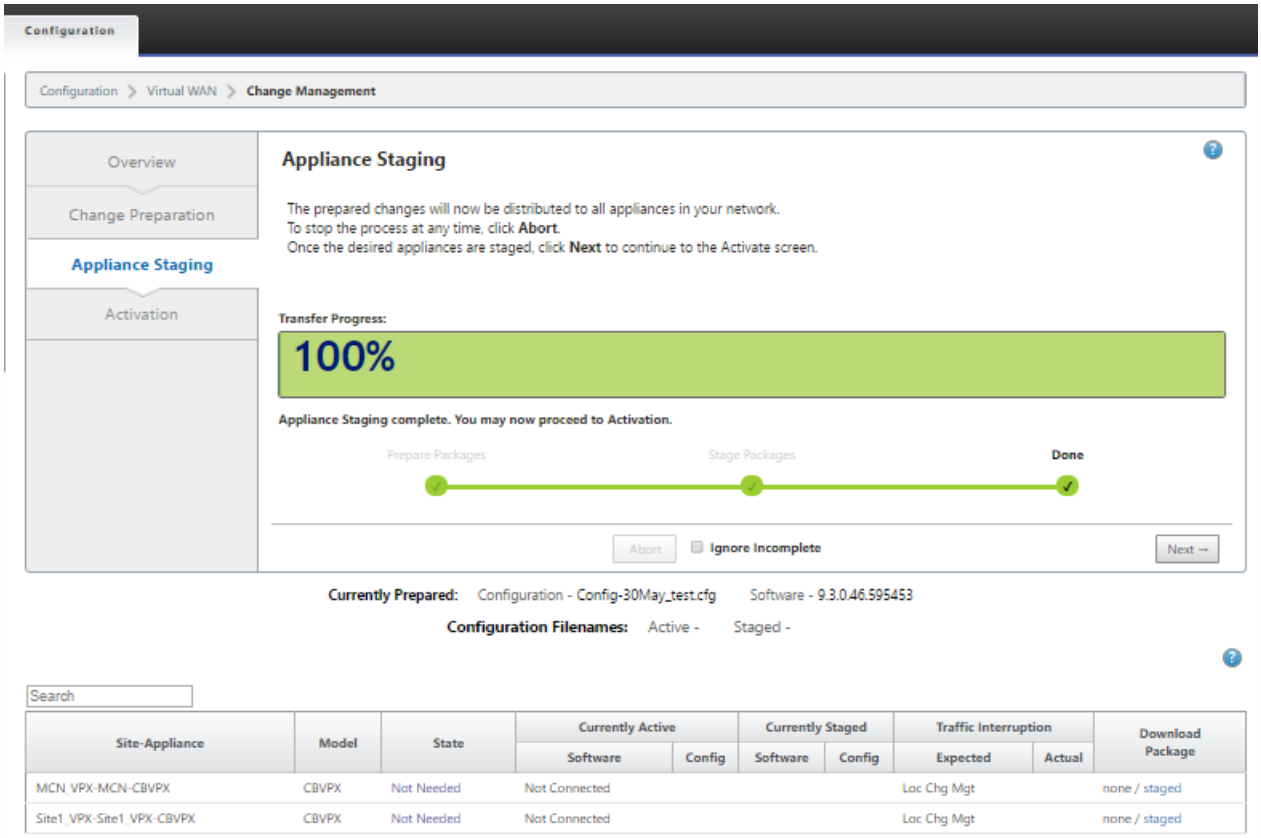

9. Click **Next**.

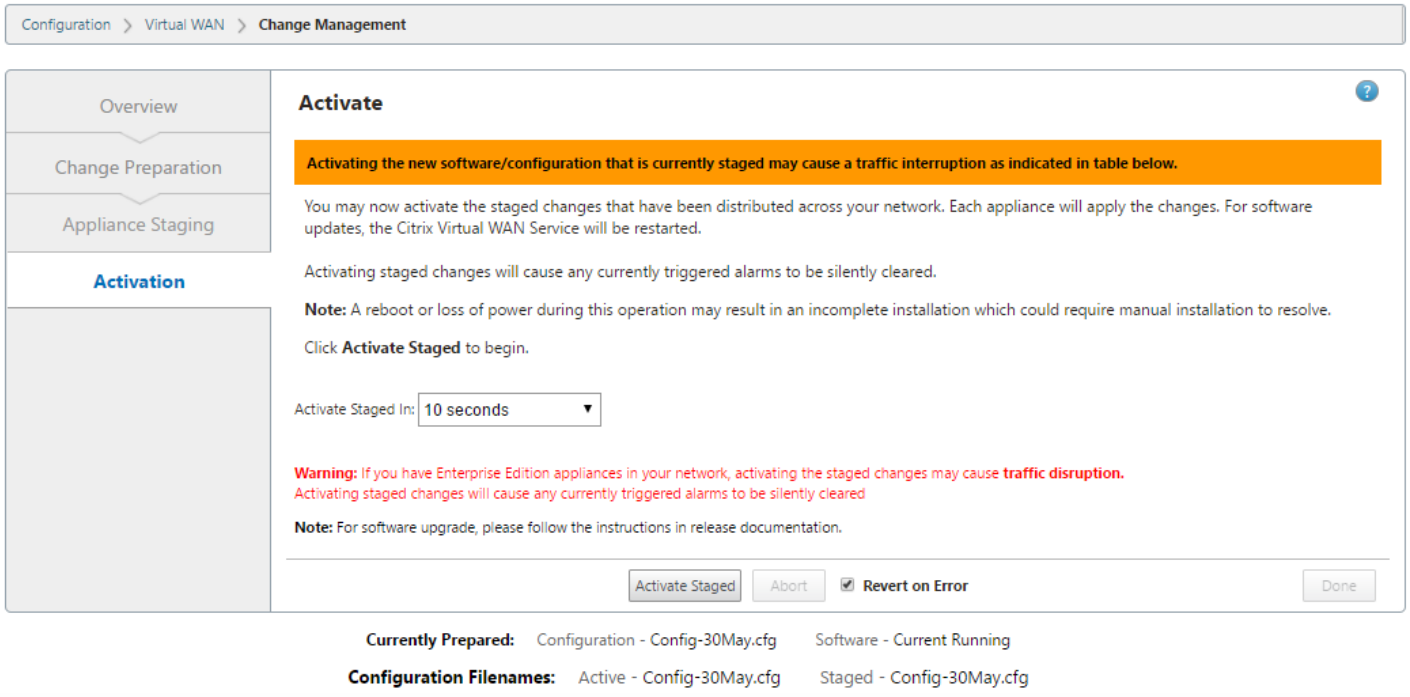

10. Select **Revert on Error** to revert to previous application package on encountering some error. For more information, see Configuration Rollback.

11. Click **Activate Staged**.

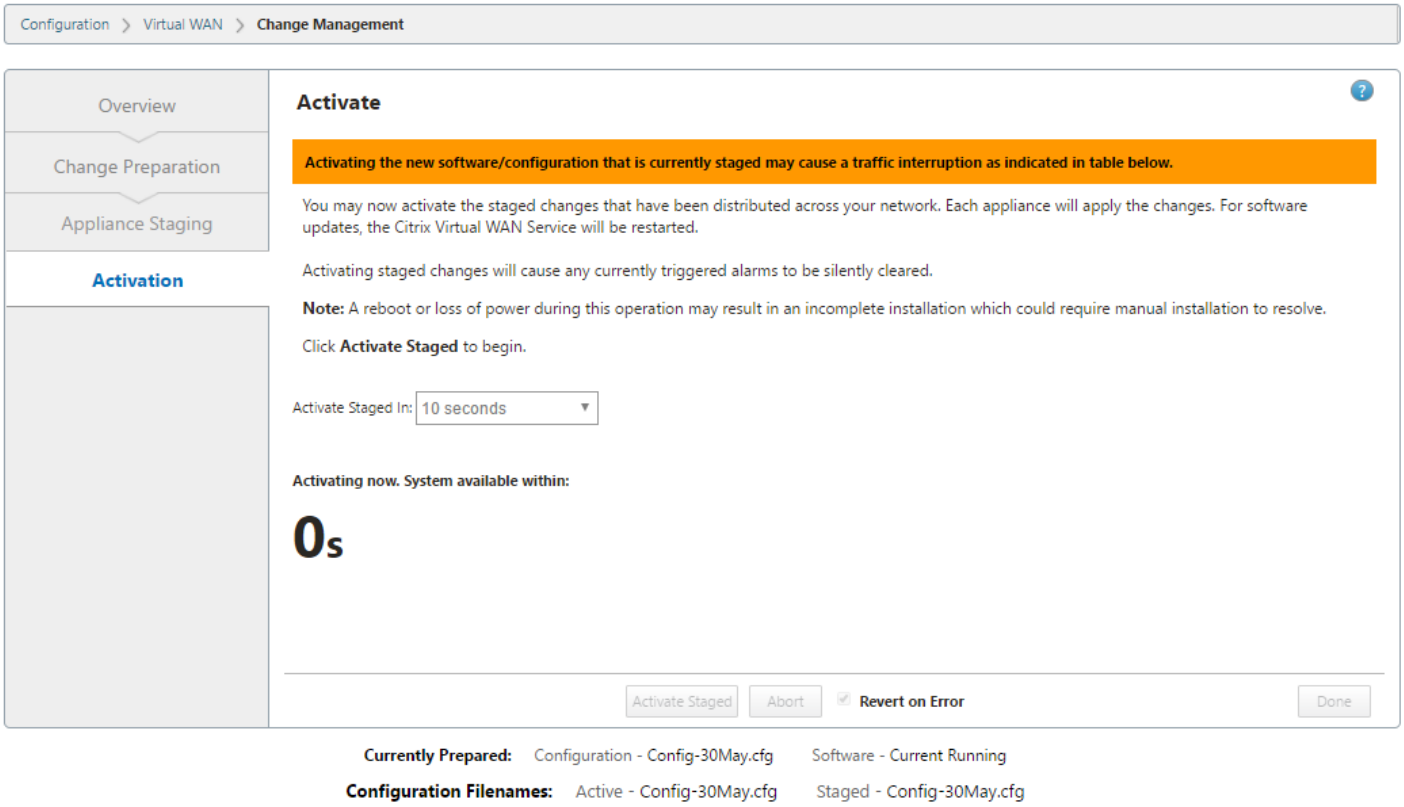

The results and next steps will differ at this point, depending on whether this is an initial configuration or you are updating or replacing an existing configuration, as follows:

- **If you are updating or changing the configuration on an existing deployment**:

If this is not an initial configuration, this activates the new configuration and the appropriate Appliance Package on the MCN appliance. The appropriate Appliance Package is then distributed to and automatically activated on each client in your SD-WAN. (This may take several seconds to complete.)

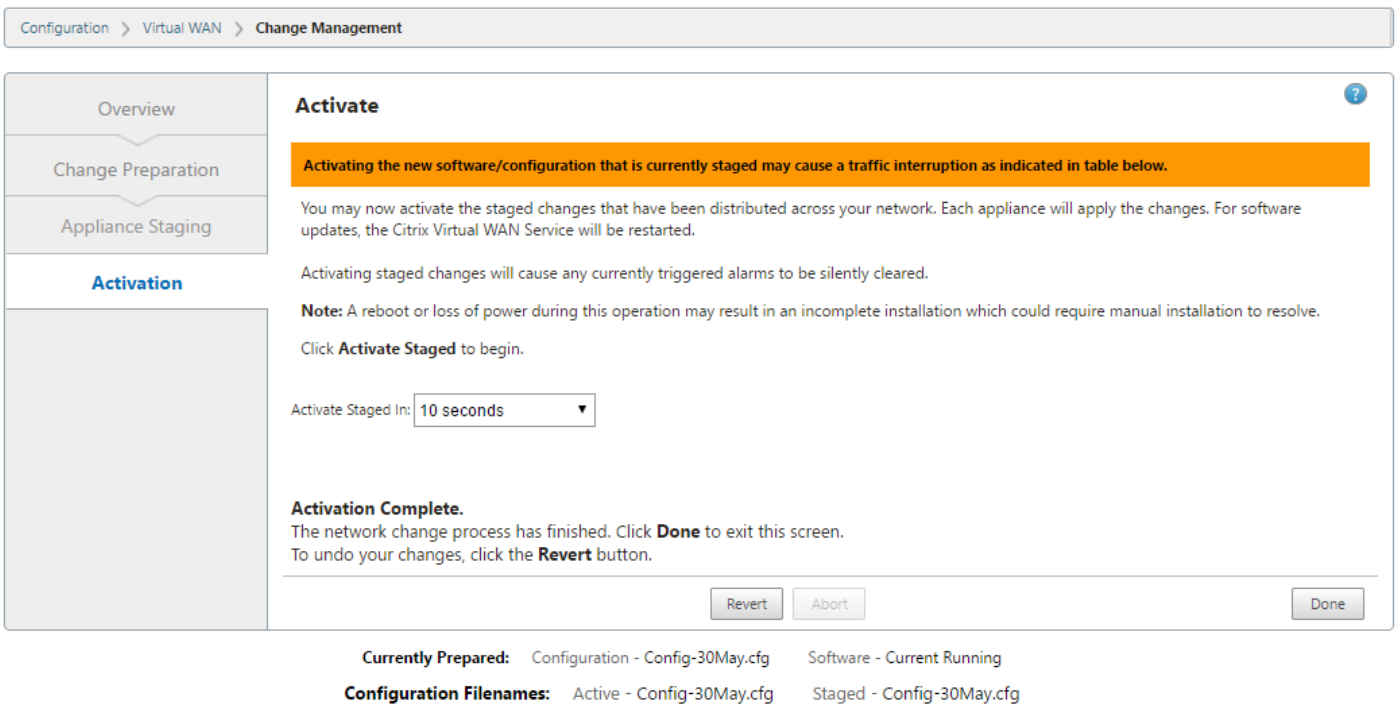

When the activation completes, an **Activation complete** status message appears, and the **Done** button is enabled. In addition, the **Configuration Filenames** status line (above the table) now displays the name of newly-activated package in the **Active** field

Click **Done** and proceed to one of the following:

* If you are not adding any new nodes to your SD-WAN, this completes the preparation, distribution, and activation of the new Appliance Packages in your SD-WAN. You can proceed directly to Enabling the Virtual WAN Service.

* If you want to add new client nodes to your SD-WAN, please proceed to Connecting the Client Appliances to Your Network.

- **If you are activating an initial configuration:**

If this is an initial configuration, the new configuration package will not be activated at this point, and there are some extra steps you must perform. The next step is to copy the configuration package to the Local Appliance Staging area, in preparation for staging and activating the configuration package on the MCN.

Do the following:

a. Once you click **Activate Staged**, the following message appears .

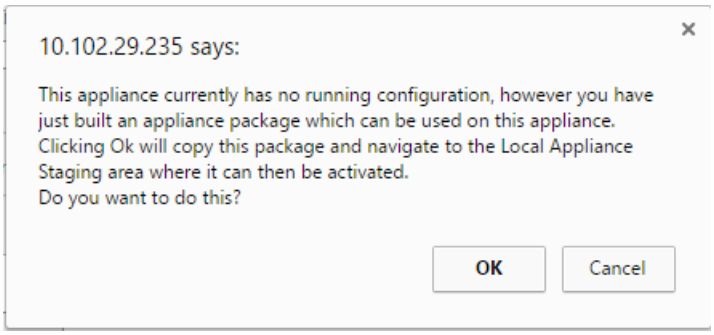

b. Click **OK**.

c. Click **Activate staged**.

This displays a dialog box asking you to confirm the activation operation.

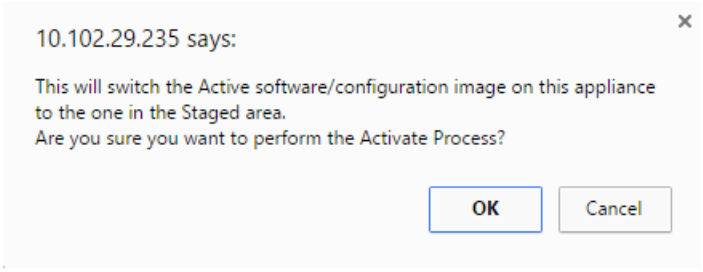

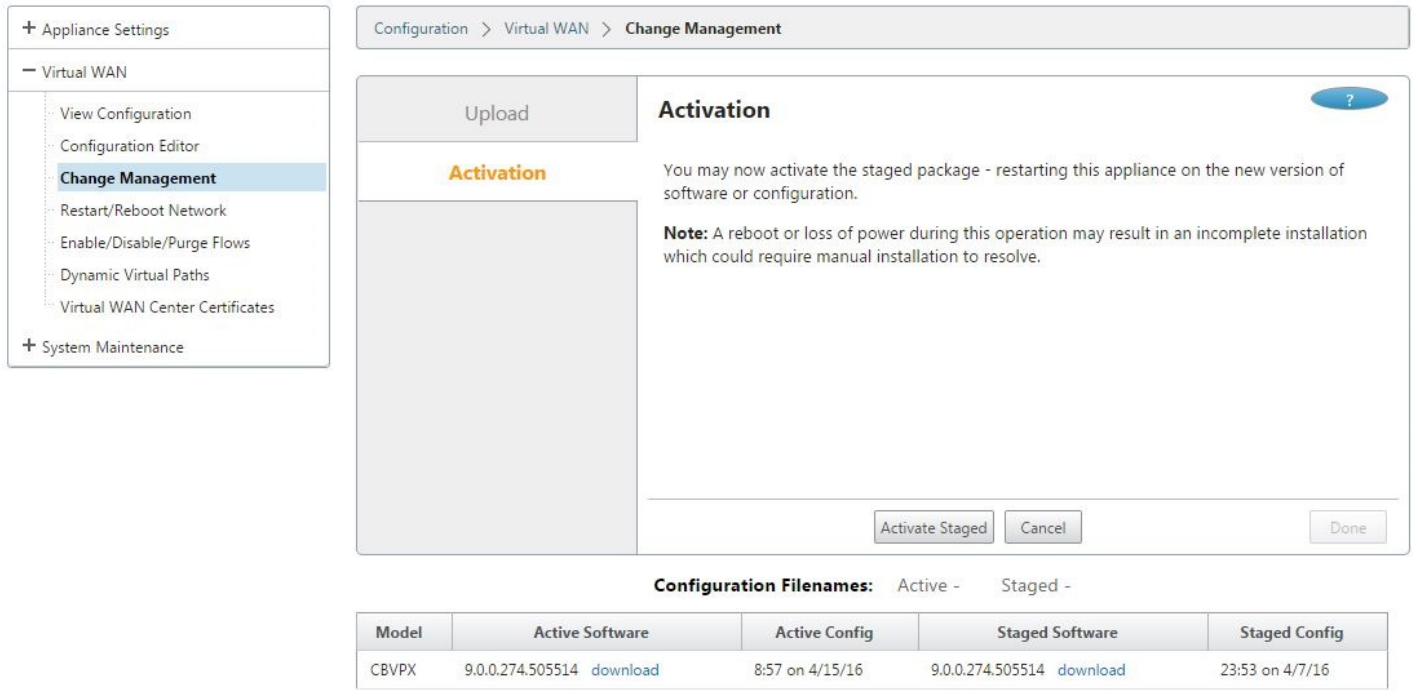

d. Click **OK**.

This initiates activation of the staged configuration package. This process takes several seconds, during which a progress status message displays.

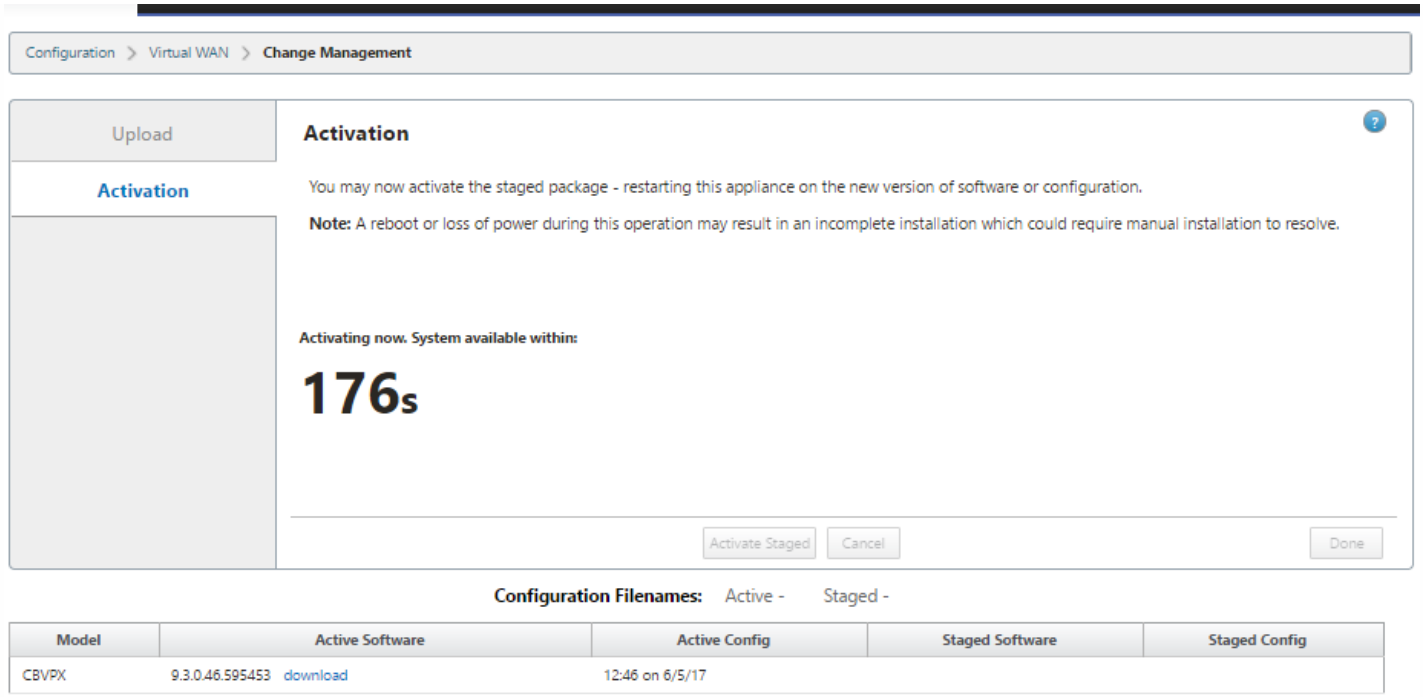

When the activation completes, a status message displays stating activation complete, and the **Done** button is enabled.

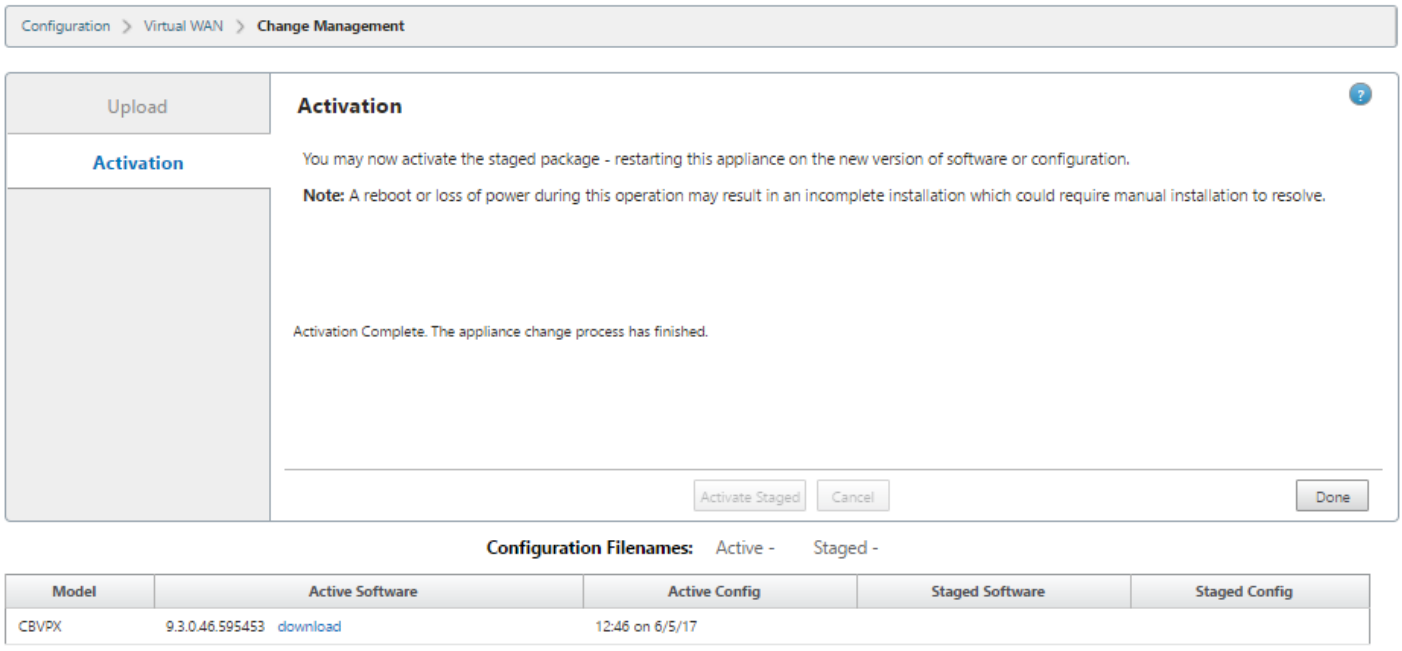

e. Click **Done**.

This proceeds to the Management Web Interface **Dashboard** page, where you can view the activation results.

You have now completed the preparation of the SD-WAN Appliance Packages on the MCN. Proceed to Connecting the Client Appliances to Your Network.

Tip

The **Change Management wizard** has an option to search the site- appliance table. This allows you to easily look up sites on a large network with multiple sites and download the required staged configuration. You can also search for error states, for example: 'Fail' or 'Not connected'. This will give you a list of all the sites in that state.

Connecting the Client Appliances to Your Network

Aug 09, 2017

If this is an initial deployment, or you are adding client nodes to an existing SD-WAN, the next step is for the branch Site Administrators to connect the client appliances to the network at their respective branch sites. This is in preparation for uploading and activating the appropriate SD-WAN Appliance Packages to the clients. You will need to contact each branch Site Administrator to initiate and coordinate these procedures.

To connect the site appliances to the SD-WAN, Site Administrators should do the following:

1. If you have not already done so, set up the client appliances.

   For each appliance you want to add to your SD-WAN, you will need to do the following:

> **Note**
>
> Instructions for each of these tasks are provided in Setting up the SD-WAN Appliances.

   a. Set up the SD-WAN Appliance hardware and any SD-WAN VPX Virtual Appliances (SD-WAN VPX-SE) you will be deploying.

   b. Set the Management IP Address for the appliance and verify the connection.

   c. Set the date and time on the appliance.

   d. Upload and install the software license file on the appliance.

2. Connect the appliance to the branch site LAN.

   Connect one end of an Ethernet cable to a port configured for LAN on the SD-WAN Appliance, and the other end of the cable to the LAN switch.

3. Connect the appliance to the WAN.

   Connect one end of an Ethernet cable to a port configured for WAN on the SD-WAN Appliance, and the other end of the cable to the WAN router.

   The next step is for the branch Site Administrators to install and activate the appropriate SD-WAN Appliance Package on their respective clients.

Setting up the SD-WAN Appliances

Aug 09, 2017

These procedures must be completed for each appliance you want to add to your SD-WAN. Consequently, this process will require some coordination with your Site Administrators across your network, to ensure the appliances are prepared and ready to deploy at the proper time. However, once the Master Control Node (MCN) is configured and deployed, you can add client appliances (client nodes) to your SD-WAN at any time.

For each appliance you want to add to your Virtual WAN, you will need to do the following.

1. Set up the SD-WAN Appliance hardware and any SD-WAN VPX Virtual Appliances (SD-WAN VPX-VW) you will be deploying.
2. Set the Management IP Address for the appliance and verify the connection.
3. Set the date and time on the appliance.
4. Set the console session **Timeout** threshold to a high or the maximum value.

> ### Warning
>
> If your console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you set the console session **Timeout** interval to a high value when creating or modifying a configuration package, or performing other complex tasks.

5. Upload and install the software license file on the appliance.

For instructions on installing a SD-WAN Virtual Appliance (SD-WAN VPX), see the following sections:

- About SD-WAN VPX.
- Installing and Deploying a SD-WAN VPX-VW on ESXi.
- Differences Between a SD-WAN VPX-SE and SD-WAN WANOP VPX Installation.

Setting up the Appliance Hardware

Aug 09, 2017

To set up your NetScaler SD-WAN Appliance hardware, do the following:

1. Set up the chassis.

   NetScaler SD-WAN Appliances can be installed in a standard rack. For desktop installation, place the chassis on a flat surface. Make sure that there is a minimum of two inches of clearance at the sides and back of the appliance, for proper ventilation.

2. Connect the Power.

   a. Make sure the power switch is set to Off.
   b. Plug the power cord into the appliance and an AC outlet.
   c. Press the power button located on the front of the appliance.

3. Connect the appliance Management Port to a personal computer.

   You will need to connect the appliance to a PC in preparation for completing the next procedure, setting the Management IP Address for the appliance.

> ## Note
>
> Before you connect the appliance, make sure the Ethernet port is enabled on the PC. Use an Ethernet cable to connect the SD-WAN Appliance Management Port to the default Ethernet port on a personal computer.

NetScaler SD-WAN 400-SE Management Port

The NetScaler SD-WAN 400-SE Management Port is the bottom far right port labeled **MGMT**, on the back of the chassis. The default IP Address for the Management Port is 192.168.100.1.

The below figure shows the location of the NetScaler SD-WAN 400-SE Management Port.

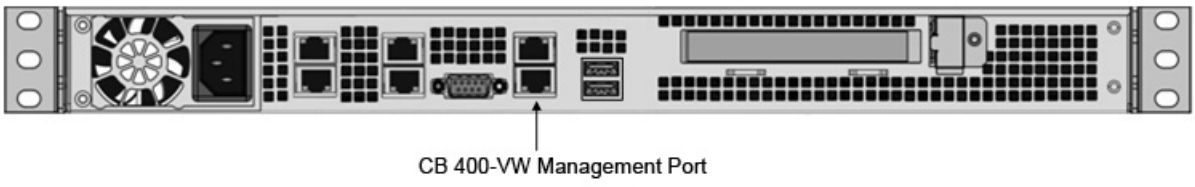

NetScaler SD-WAN 1000-SE Management Port

The NetScaler SD-WAN 1000-SE Management Port is the bottom far right port labeled MGMT, on the back of the chassis. The default IP Address for the Management Port is 192.168.100.1.

The below figure shows the location of the NetScaler SD-WAN 1000-SE Management Port.

NetScaler SD-WAN 2000-SE Management Port

The NetScaler SD-WAN 2000-SE Management Port is the bottom-left port labeled 0/1, on the front of the chassis. The default IP Address for the Management Port is 192.168.100.1.

The following figure shows the location of the NetScaler SD-WAN 2000-SE Management Port.

NetScaler SD-WAN 4000-SE Management Port

The NetScaler SD-WAN 4000-SE Management Port is the bottom-left port labeled 0/1, on the front of the chassis. The default IP Address for the Management Port is 192.168.100.1.

The below figure shows the location of the NetScaler SD-WAN 4000-SE Management Port.

NetScaler SD-WAN 5100-SE Management Port

The NetScaler SD-WAN 5100-SE Management Port is the bottom-left port labeled 0/1, on the front of the chassis. The default IP Address for the Management Port is 192.168.100.1.

The below figure shows the location of the NetScaler SD-WAN 5100-SE Management Port.

NetScaler SD-WAN VPX-SE Management Port

The NetScaler SD-WAN VPX-SE Virtual Appliance is a Virtual Machine, so there is no physical Management Port. However, if you did not configure the Management IP Address for the SD-WAN VPX-SE when you created the VPX Virtual Machine, you will need to do so now, as outlined in the section, Configuring the Management IP Address for the SD-WAN VPX-SE.

Also see the section Setting the Management IP Addresses for the Appliances.

Setting the Management IP Addresses for the Appliances

Aug 09, 2017

To enable remote access to a NetScaler SD-WAN appliance, you must specify a unique Management IP Address for the appliance. To do so, you must first connect the appliance to a personal computer. You can then open a browser on the PC and connect directly to the Management Web Interface on the appliance, where you can set the Management IP Address for that appliance. The Management IP Address must be unique for each appliance.

The procedures are different for setting the Management IP Address for a hardware SD-WAN Appliance and a VPX Virtual Appliance (NetScaler SD-WAN VPX-SE). For instructions for configuring the address for each type of appliance, see the following:

- **SD-WAN VPX Virtual Appliance –** See the sections, Configuring the Management IP Address for the SD-WAN VPX-SE and Differences Between a SD-WAN VPX-SE and SD-WAN WANOP VPX Installation.
- **SD-WAN hardware appliance –** See the section Setting the Management IP Address for a Hardware SD-WAN Appliance.

Setting the Management IP Address for a SD-WAN Appliance

Aug 09, 2017

To configure the Management IP Address for a hardware SD-WAN Appliance, do the following:

> **Note**
>
> You must repeat the following process for each hardware appliance you want to add to your network.

1. If you are configuring a hardware SD-WAN Appliance, physically connect the appliance to a PC.

   If you have not already done so, connect one end of an Ethernet cable to the Management Port on the appliance, and the other end to the default Ethernet port on the PC.

> **Note**
>
> Make sure the Ethernet port is enabled on the PC you are using to connect to the appliance.

2. Record the current Ethernet port settings for the PC you will be using to set the appliance Management IP Address.

   You will need to change the Ethernet port settings on the PC before you can set the appliance Management IP Address. Be sure to record the original settings so you can restore them after configuring the Management IP Address.

3. Change the IP Address for the PC.

   On the PC, open your network interface settings and change the IP Address for your PC to the following:

   192.168.100.50

4. Change the Subnet Mask setting on your PC to the following:

   255.255.0.0

5. On the PC, open a browser and enter the default IP Address for the appliance.

> **Note**
>
> It is recommended that you use Google Chrome browser when connecting to a SD-WAN Appliance.

Enter the following IP Address in the address line of the browser:

192.168.100.1

Note

Please ignore any browser certificate warnings for the CloudBridge Management Web Interface.

This opens the NetScaler SD-WAN Management Web Interface Login screen on the connected appliance, as shown in the below figure.

6. Enter the Administrator user name and password, and click **Login**.

   - Default Administrator user name: *admin*
   - Default Administrator password: *password*

Note

It is strongly recommended that you change the default password as soon as possible. Be sure to record the password in a secure location, as password recovery might require a configuration reset.

After you have logged into the Management Web Interface, the **Dashboard** page displays, as shown below.

The first time you log into the Management Web Interface on an appliance, the **Dashboard** displays an Alert icon (goldenrod delta) and alert message indicating that the Virtual WAN Service is disabled, and the license has not been installed. For now, you can ignore this alert. The alert will be resolved after you have installed the license, and completed the configuration and deployment process for the appliance.

Below figure shows a sample **Dashboard** after the Virtual WAN has been fully configured and deployed.

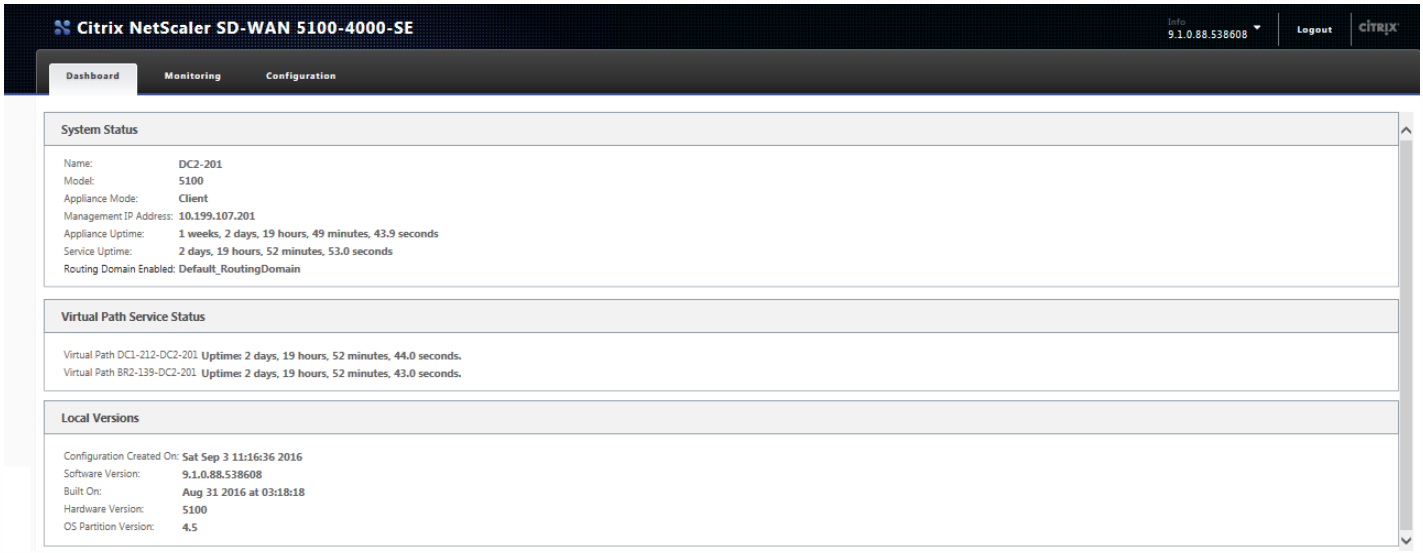

7. In the main menu bar, select the **Configuration** section tab.

This displays the **Configuration** navigation tree in the left pane of the screen. The **Configuration** navigation tree contains the following three primary branches:

- **Appliance Settings**
- **Virtual WAN**
- **System Maintenance**

When you select the **Configuration** tab, the **Appliance Settings** branch automatically opens, with the **Administrator Interface** page preselected by default, as shown in the below figure.

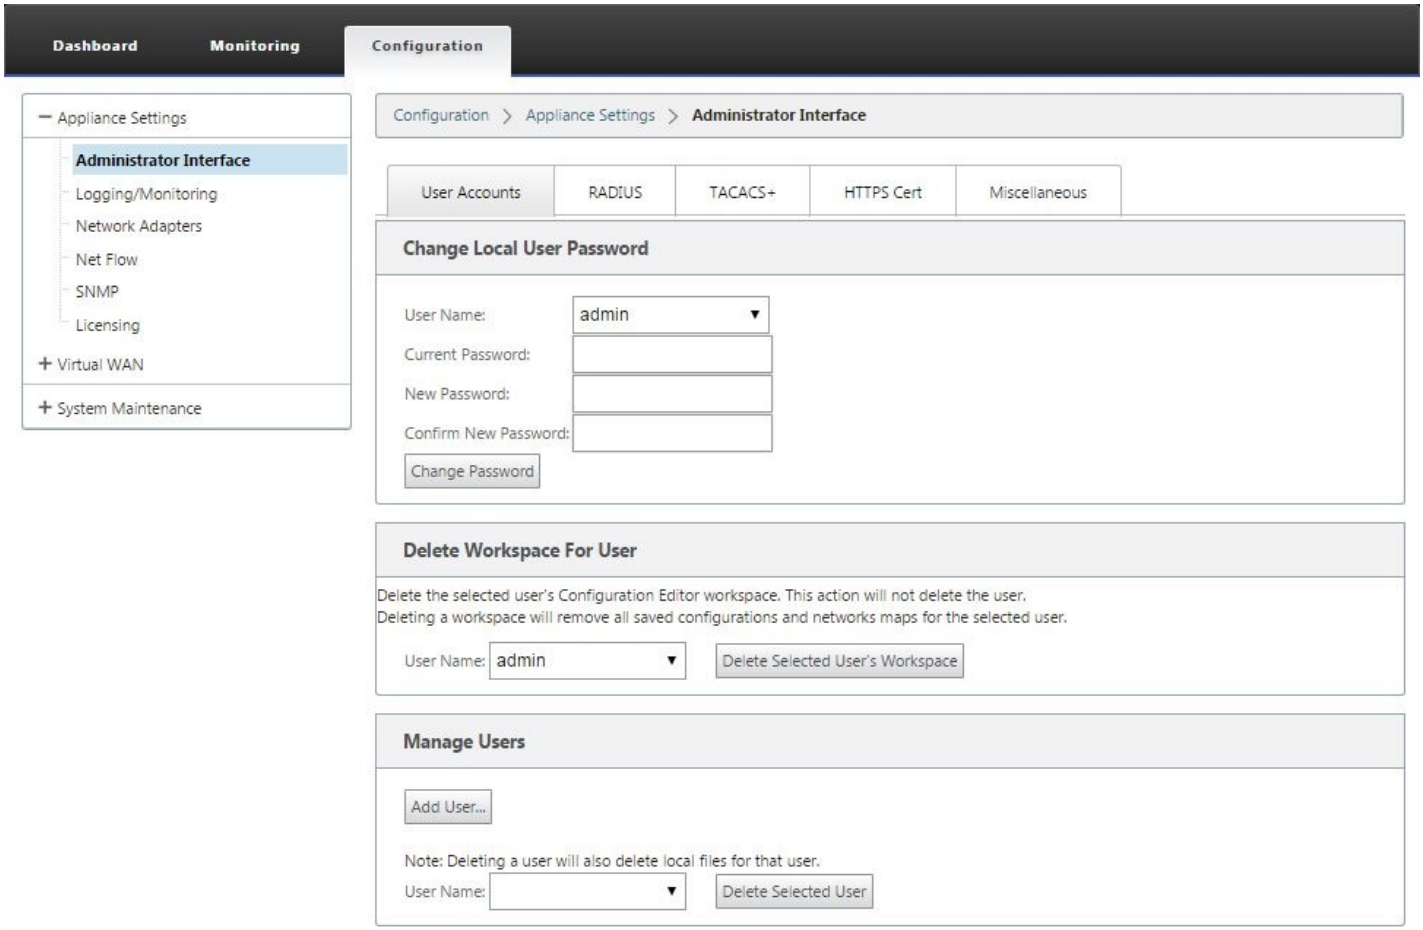

8. In the **Appliance Settings** branch of the navigation tree, select **Network Adaptors**.

This displays the **Network Adaptors** settings page with the **IP Address** tab preselected by default, as shown in the below figure.

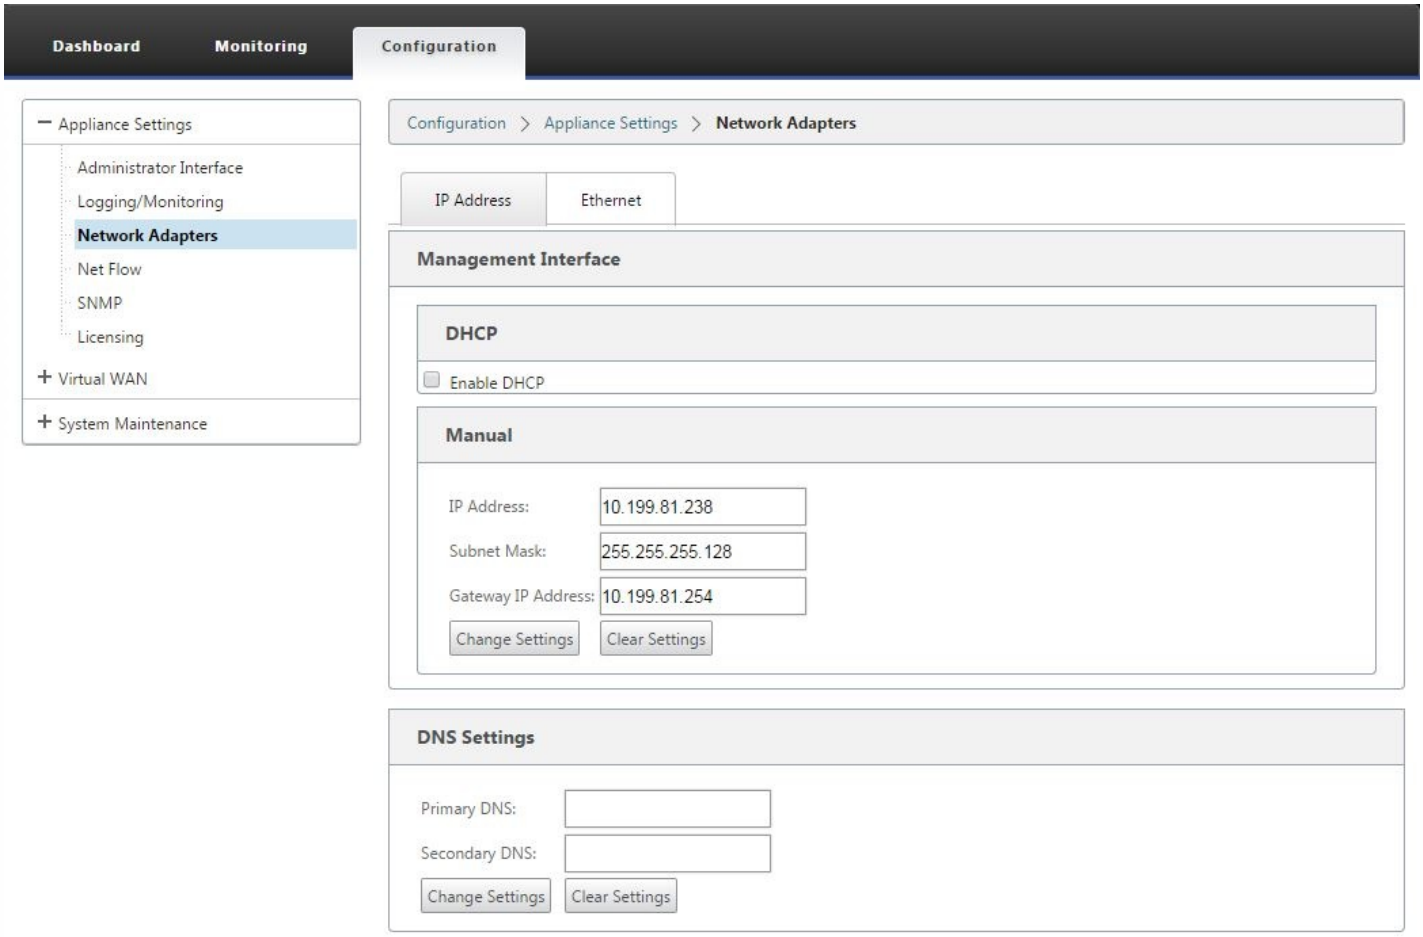

9. In the **IP Address** tab page, enter the following information for the CloudBridge Virtual WAN Appliance you want to configure.

- IP Address
- Subnet Mask
- Gateway IP Address

Note

The Management IP Address must be unique for each appliance.

10. Click **Change Settings**.
    A confirmation dialog box displays, prompting you to verify that you want to change these settings.
11. Click **OK**.
12. Change the network interface settings on your PC back to the original settings.

Note

Changing the IP Address for your PC automatically closes the connection to the appliance, and terminates your login session on the Management Web Interface.

13. Disconnect the appliance from the PC and connect the appliance to your network router or switch.

    Disconnect the Ethernet cable from the PC, but do not disconnect it from your appliance. Connect the free end of the cable to your network router or switch.

    The SD-WAN Appliance is now connected to and available on your network.

14. Test the connection.

    On a PC connected to your network, open a browser and enter the Management IP Address you just configured for the appliance.

    If the connection is successful, this displays the **Login** screen for the NetScaler SD-WAN Management Web Interface on the appliance you just configured.

> ## Tip
>
> After verifying the connection, do not log out of the Management Web Interface. You will be using it to complete the remaining tasks outlined in the subsequent sections.

You have now set the Management IP Address of your SD-WAN Appliance, and can connect to the appliance from any location in your network.

Setting the Date and Time on an SD-WAN Appliance

Aug 09, 2017

Before installing the SD-WAN software license on an appliance, you must set the date and time on the appliance.

> **Note**
>
> You must repeat this process for each appliance you want to add to your network.

To set the date and time, do the following:

1. Log into the Management Web Interface on the appliance you are configuring.
2. In the main menu bar, select the **Configuration tab.**
   This displays the **Configuration** navigation tree in the left pane of the screen.
3. Open the **System Maintenance branch in the navigation tree.**
4. Under the **System Maintenance branch, select Date/Time Settings**.
   This displays the Date/Time Settings page, as shown below.

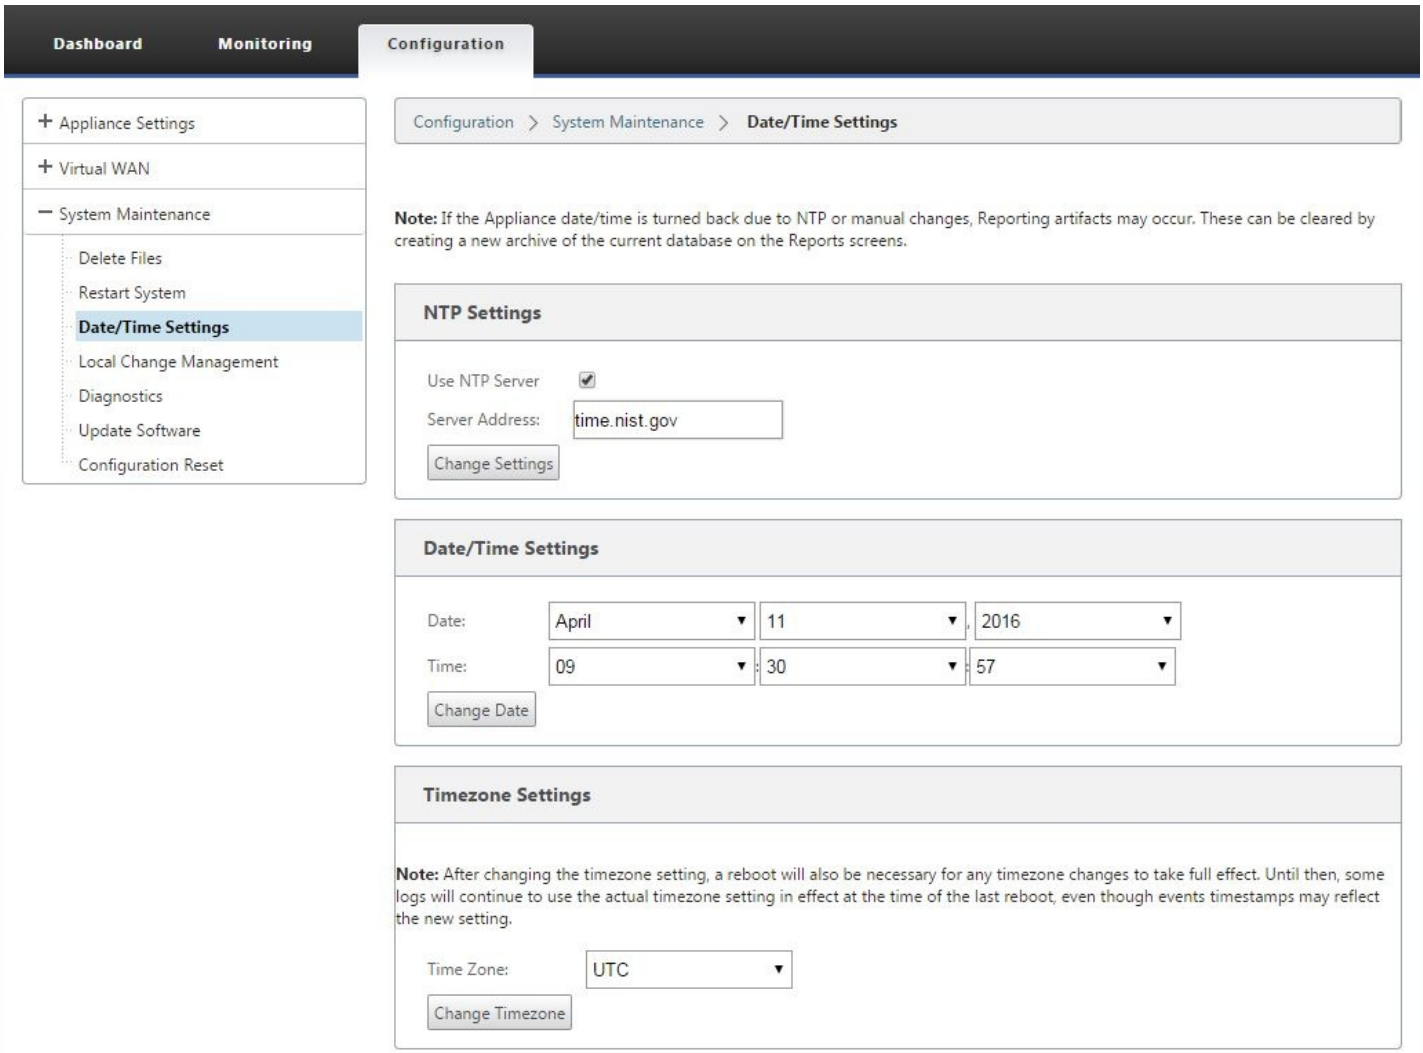

5. Select the time zone from the **Time Zone** field drop-down menu at the bottom of the page.

> **Note**
>
> If you need to change the time zone setting, you must do this before setting the date and time, or your settings will not persist as entered.

6. Click **Change Timezone**.

   This updates the time zone and recalculates the current date and time setting, accordingly. If you set the correct date and time before this step, then your settings will no longer be correct.

When the time zone update completes, a success Alert icon (green check mark) and status message displays in the top section of the page.

7. (Optional) Enable NTP Server service.
   a) Select **Use NTP Server**.
   b) Enter the server address in the **Server Address** field.
   c) Click **Change Settings**.
      A success Alert icon (green checkmark) and status message displays when the update completes.
8. Select the month, day, and year from the **Date** field drop-down menus.
9. Select the hour, minutes, and seconds from the **Time** field drop-down menus.
10. Click **Change Date**.

> **Note**
>
> This updates the date and time setting, but does not display a success Alert icon or status message.

The next step is to set the console session **Timeout** threshold to the maximum value. This step is optional, but strongly recommended. This prevents the session from terminating prematurely while you are working on the configuration, which could result in a loss of work. Instructions for setting the console session **Timeout** value are provided in the following section. If you do not want to reset the timeout threshold, you can proceed directly to the section, Uploading and Installing the SD-WAN Software License File.

> **Warning**
>
> If your console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning.

Setting the Console Session Timeout Interval (Optional)

Aug 09, 2017

If your console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you set the console session **Timeout** interval to a high value when creating or modifying a configuration package, or performing other complex tasks. The default is 60 minutes; the maximum is 9999 minutes. For security reasons, you should then reset it to a lower threshold after completing those tasks.

To reset the console session **Timeout** interval, do the following:

1. Select the **Configuration** tab, and then select the **Appliance Settings** branch in the navigation tree.

   This displays the **Appliance Settings** page, with the **User Accounts** tab preselected by default.

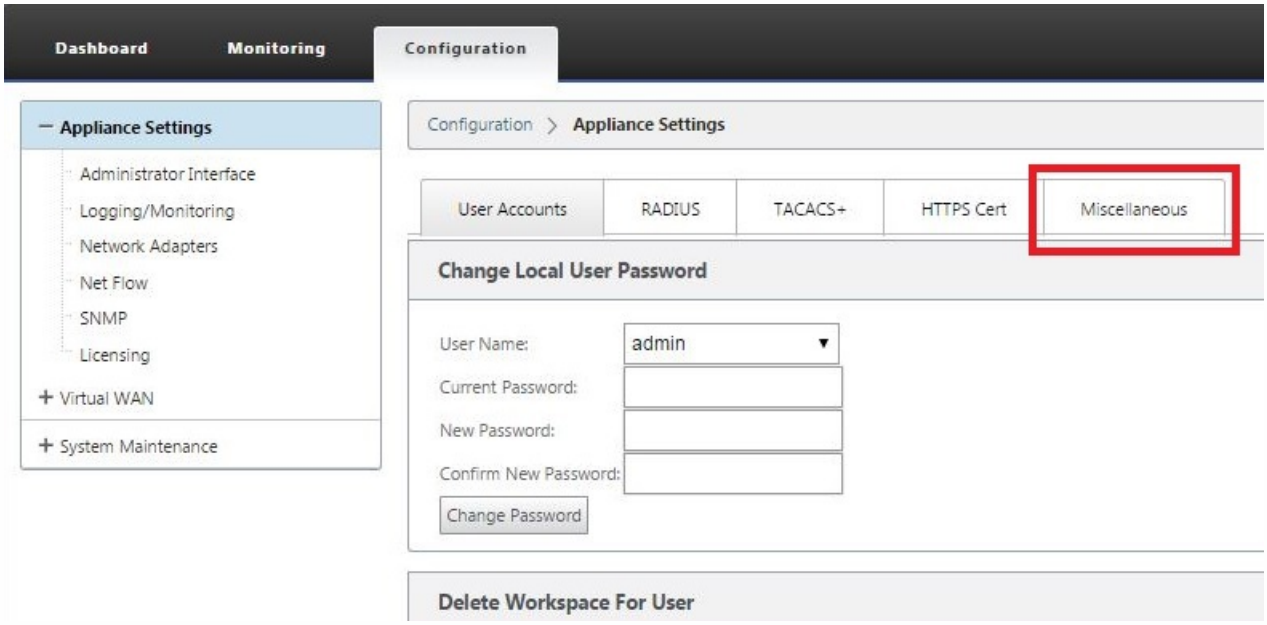

2. Select the **Miscellaneous** tab (far right corner).

   This displays the **Miscellaneous** tab page.

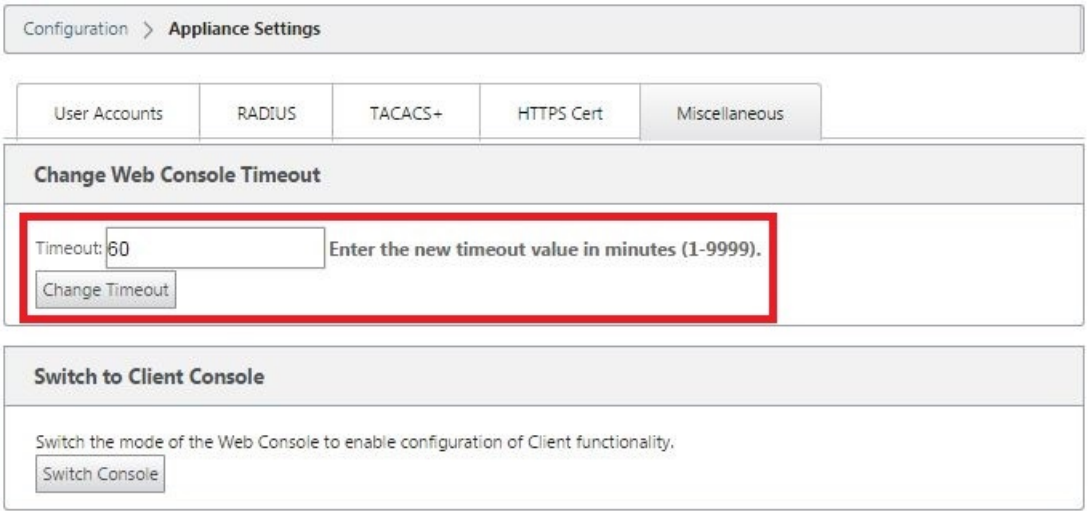

3. Enter the console **Timeout** value.

In the **Timeout** field of the **Change Web Console Timeout** section, enter a higher value (in minutes) up to the maximum value of 9999. The default is 60, which is usually much too brief for an initial configuration session.

> ## Note
>
> For security reasons, be sure to reset this value to a lower interval after completing the configuration and deployment.

4. Click **Change Timeout**.

This resets the session **Timeout** interval, and displays a success message when the operation completes.

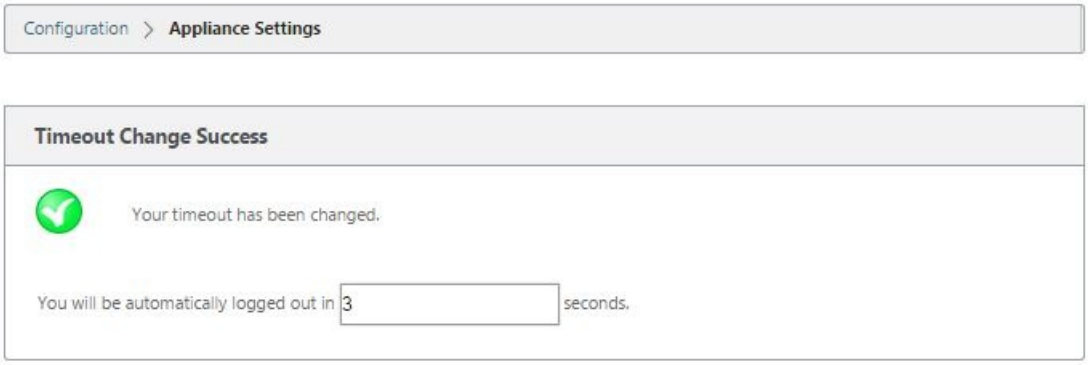

After a brief interval (a few seconds), the session is terminated and you are automatically logged out of the Management Web Interface. The Login page page appears.

5. Enter the Administrator user name (*admin*) and password (*password*), and click **Login**.

   The next step is to upload and install the SD-WAN software license file on the appliance.

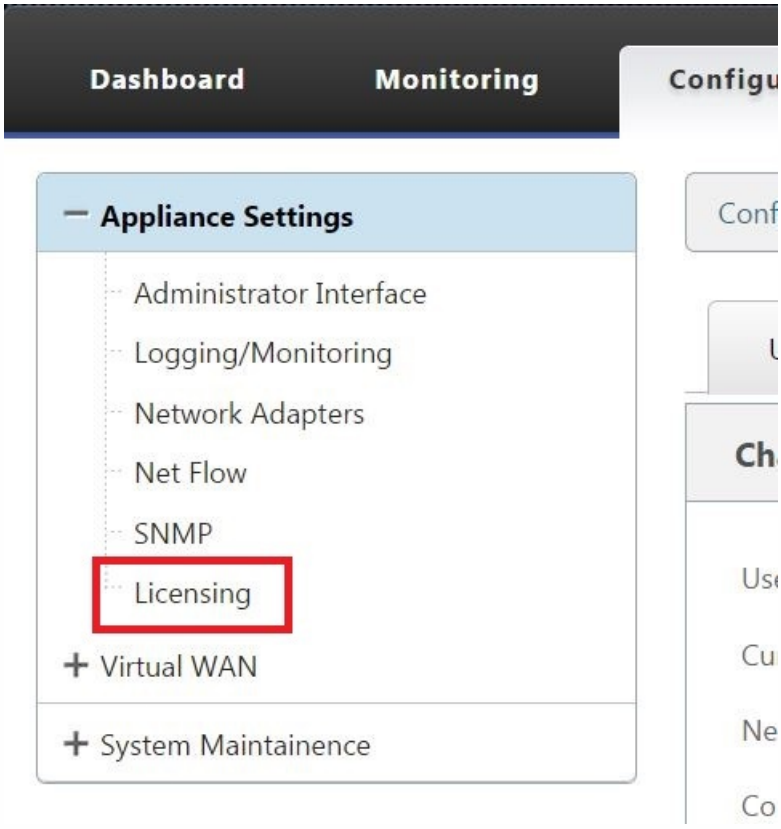

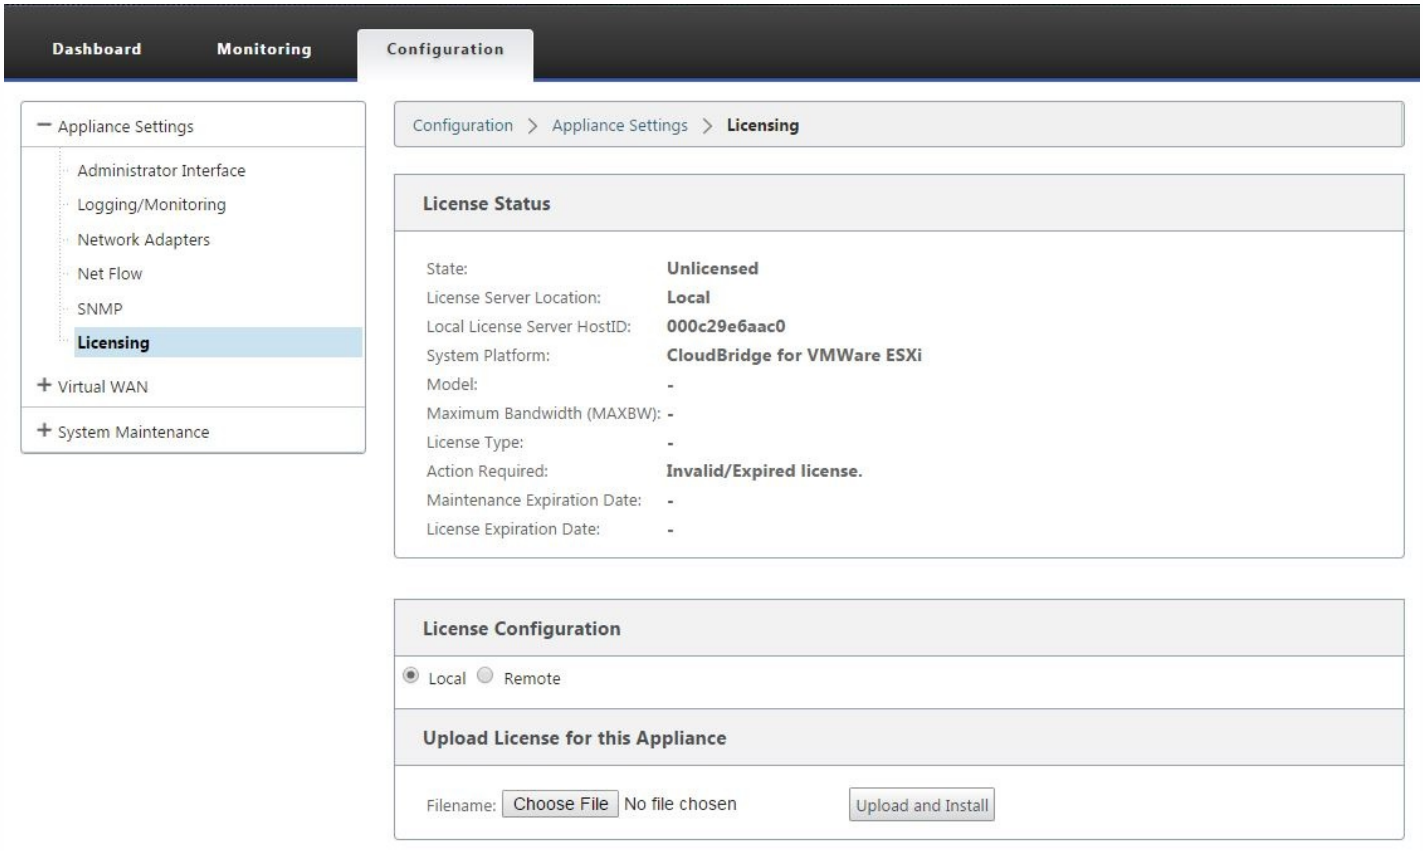

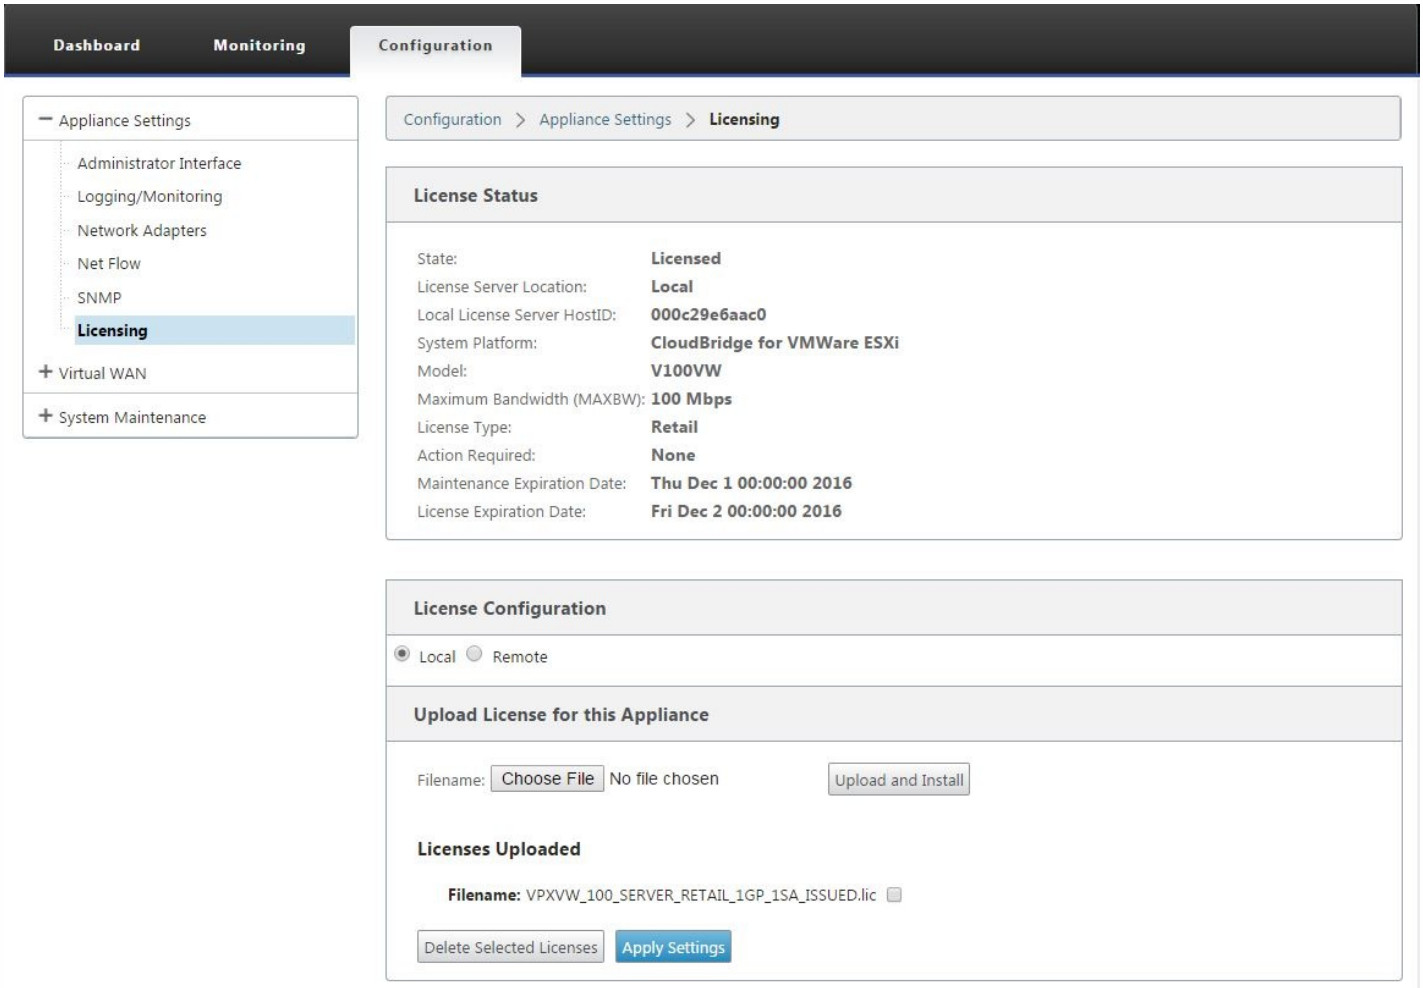

-

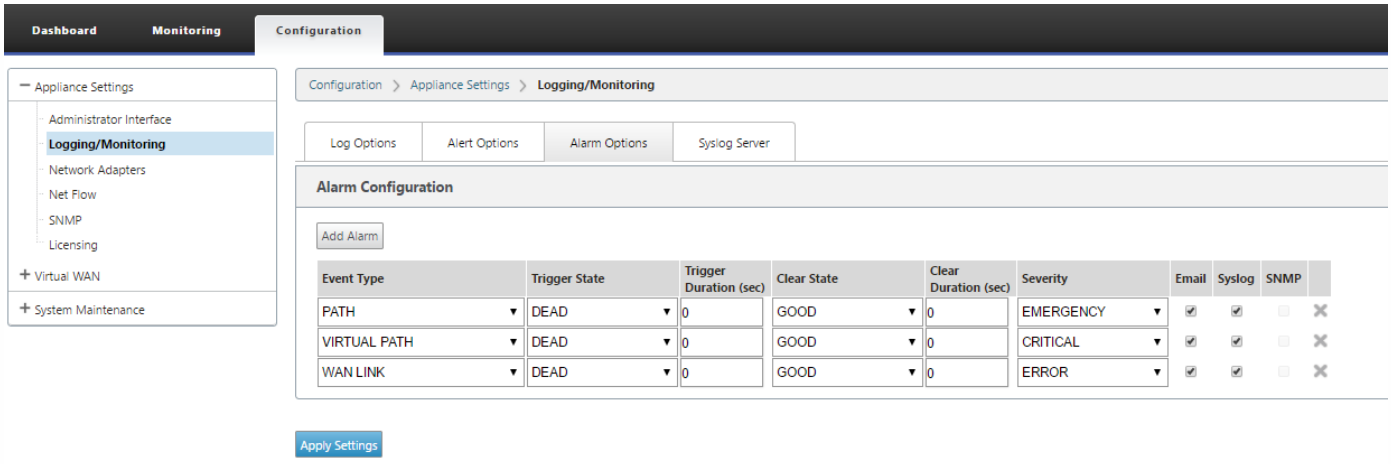

-
-
-
-
-
-
-

-
-

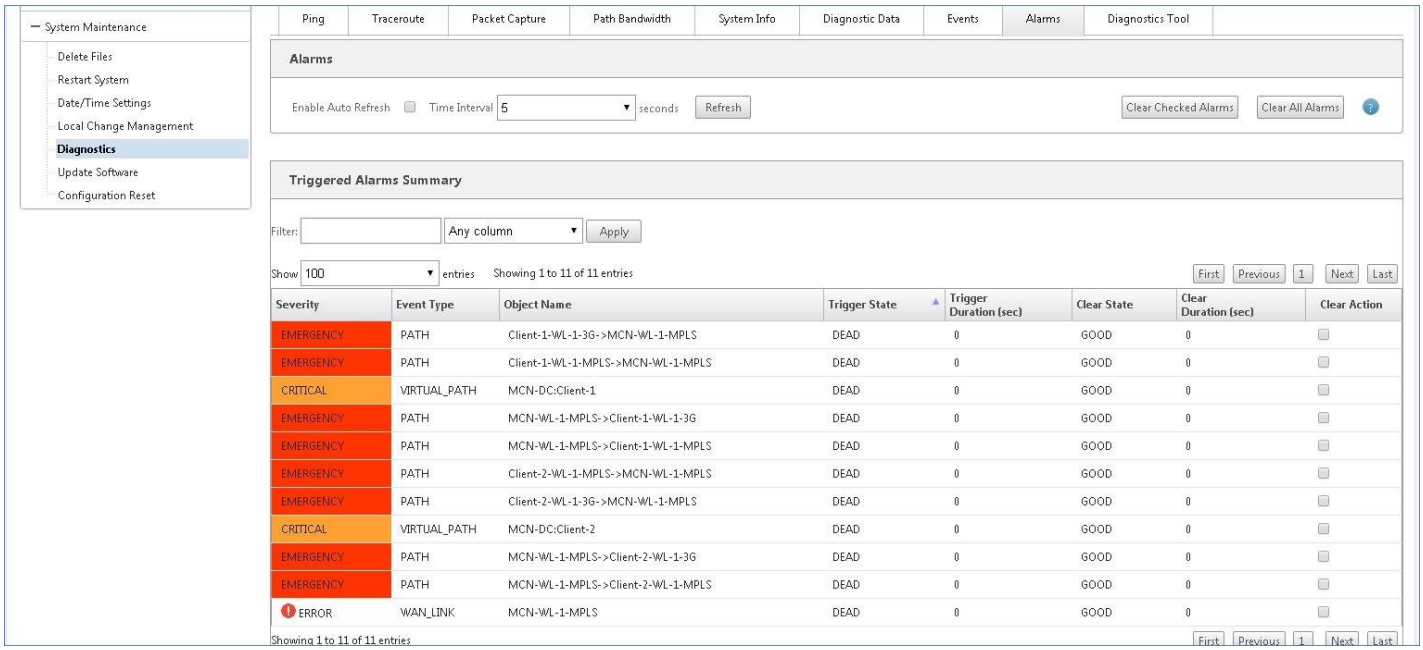

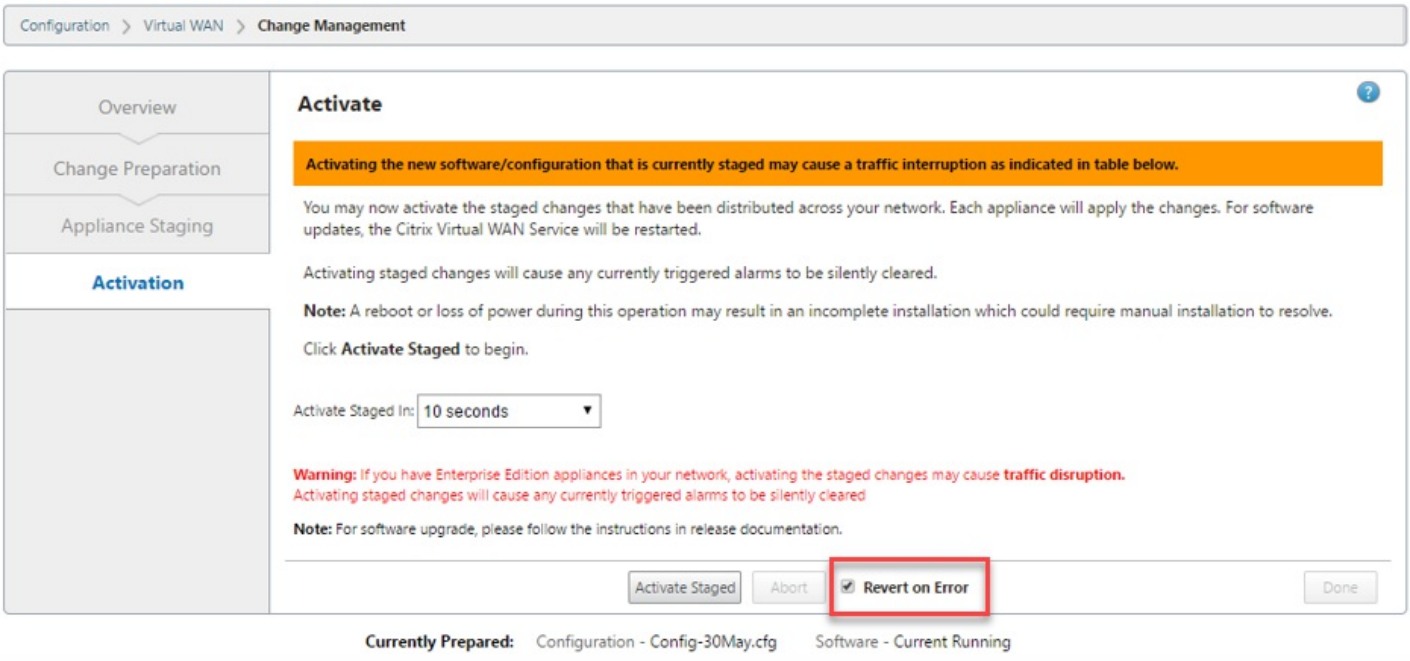

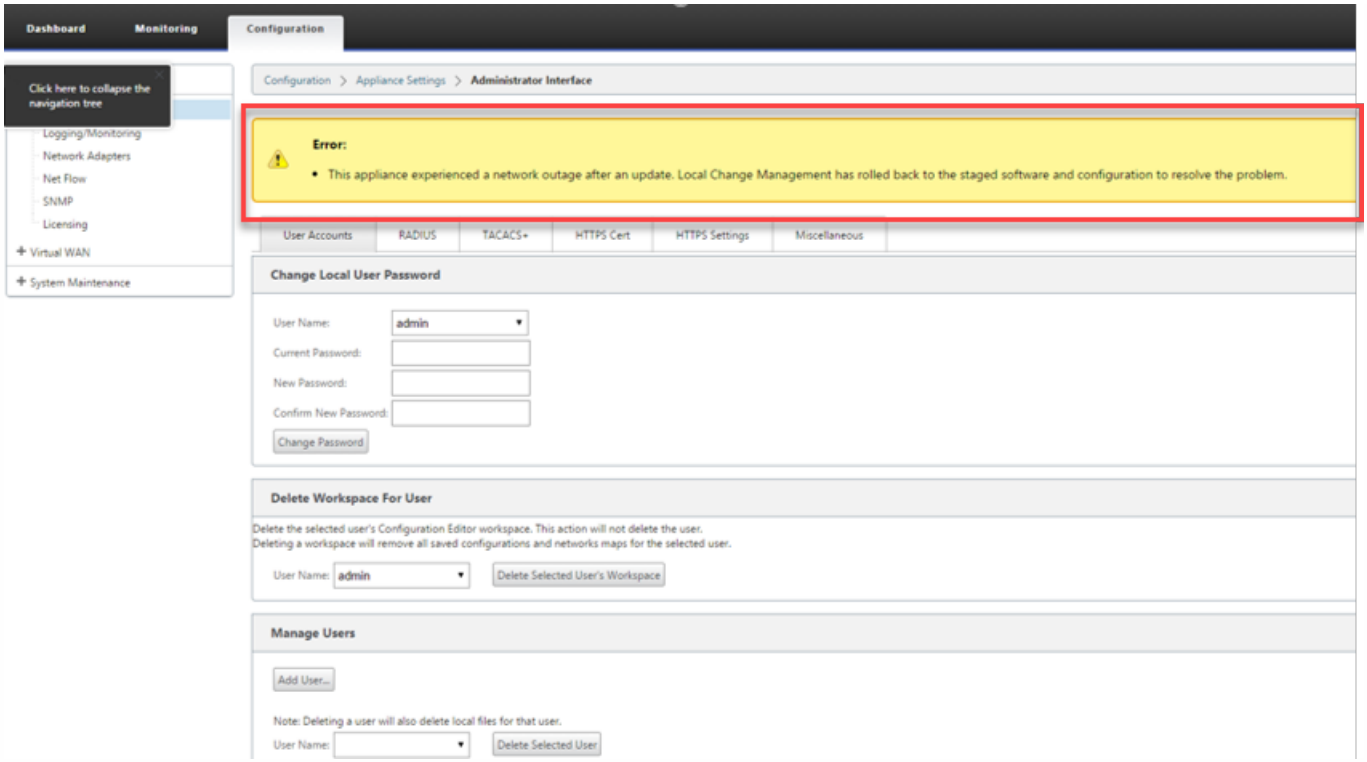

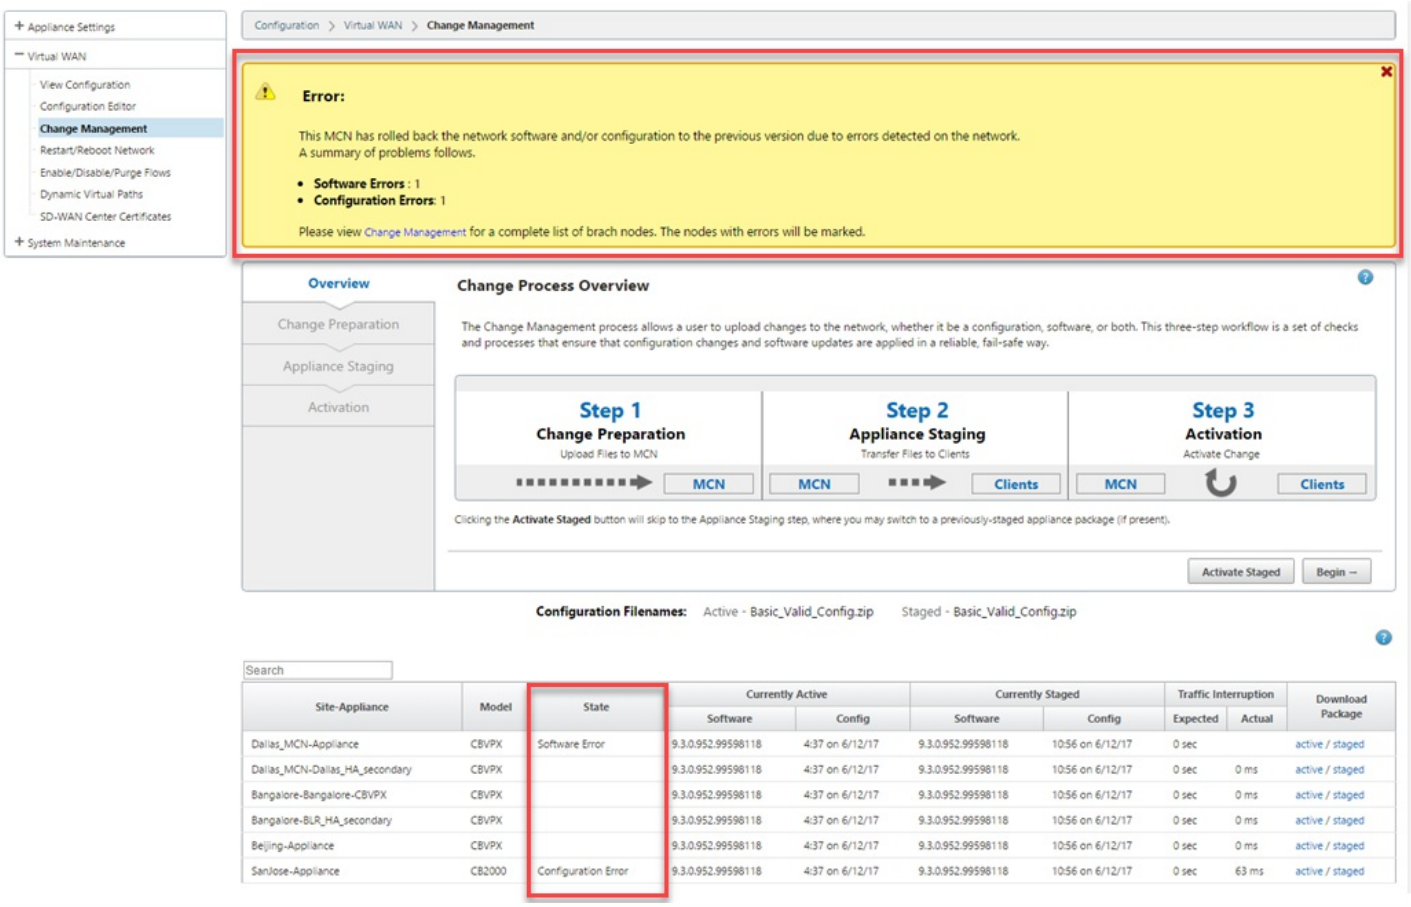

-
-

-
-

-
-
-
-
-

-
-
-

-

-

-

-

-
-

-

-

-
-
-
-
-
-
-

-
-

1. Open a browser and navigate to the ESXi server that will host your vSphere Client and VPX-SE Virtual Machine (VM) instance at: https://my.vmware.com/group/vmware/evalcenter?p=free-esxi6

The **VMware ESXi downloads** page displays,

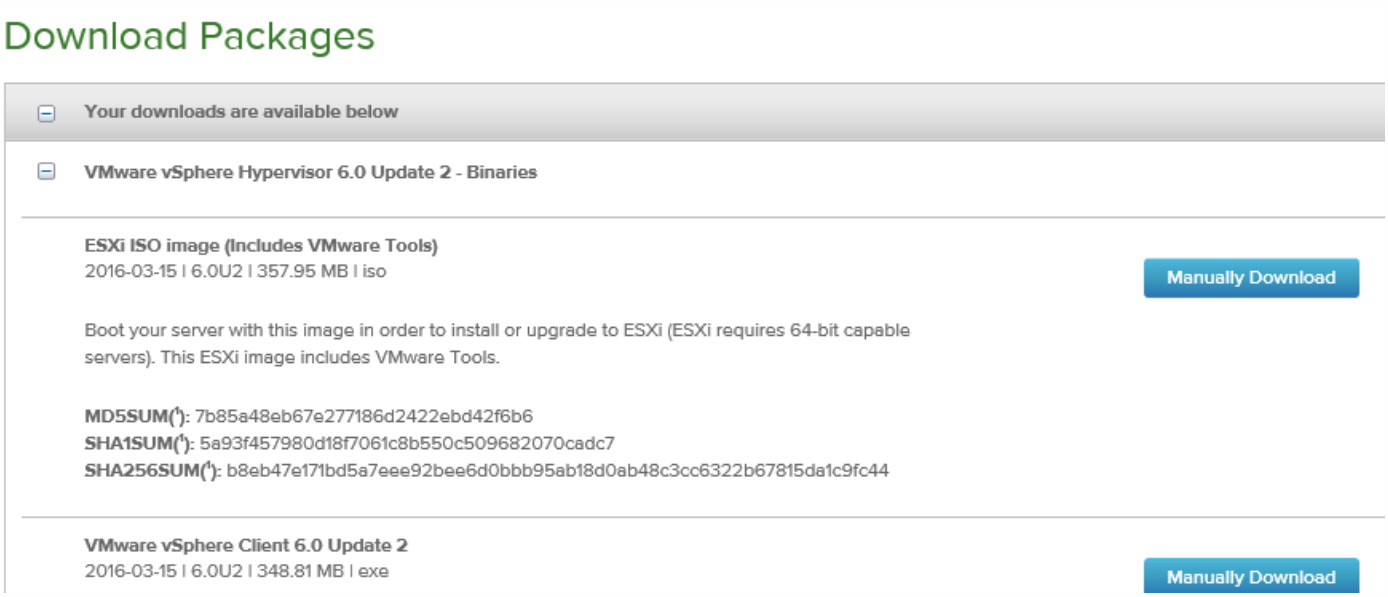

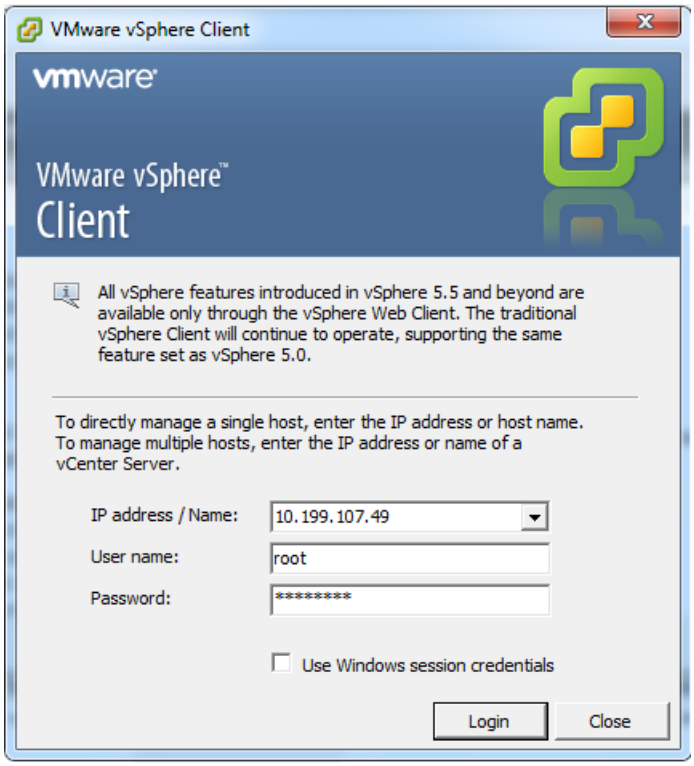

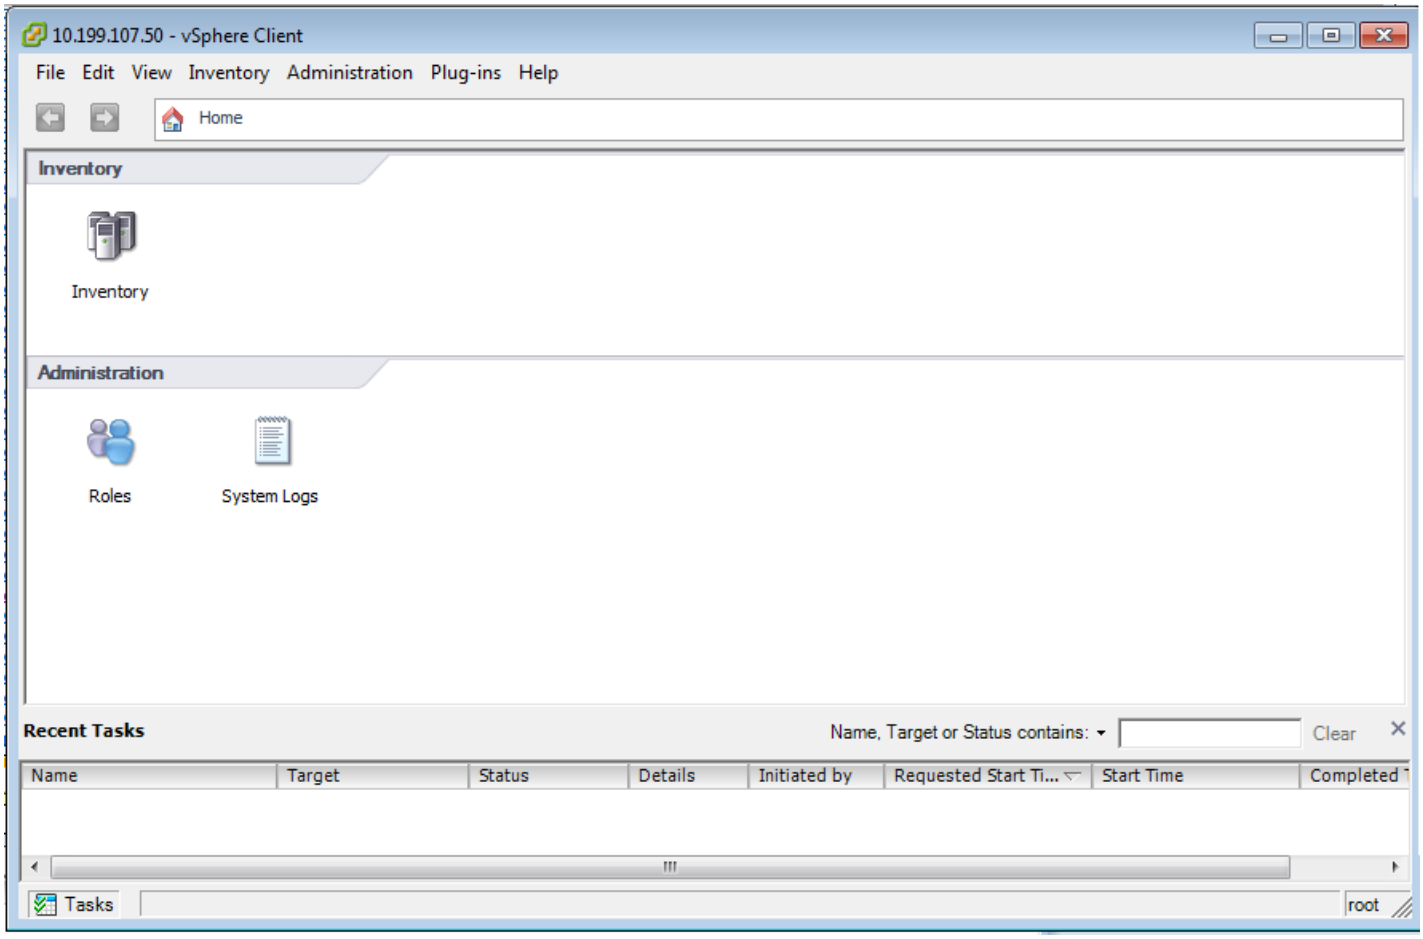

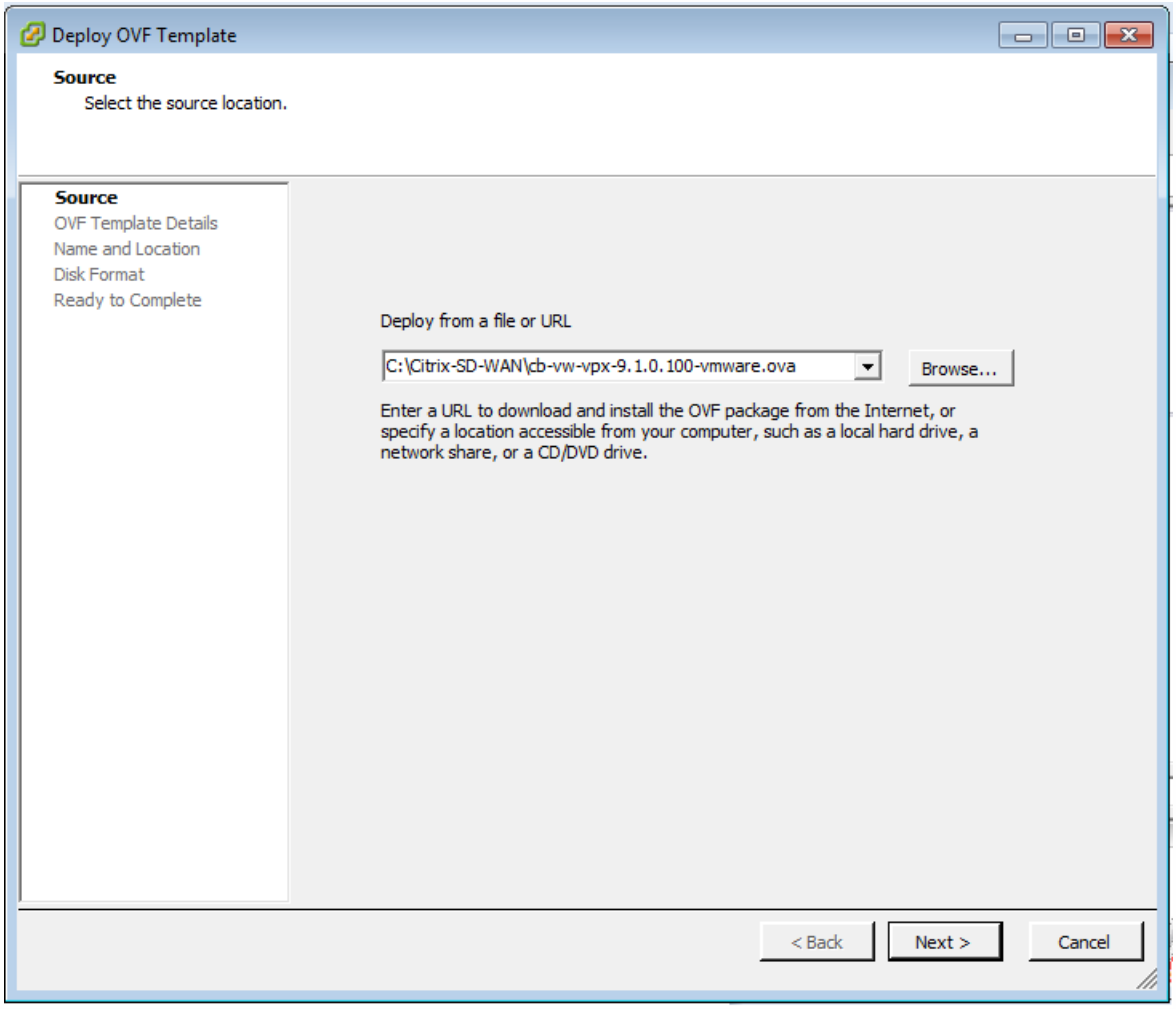

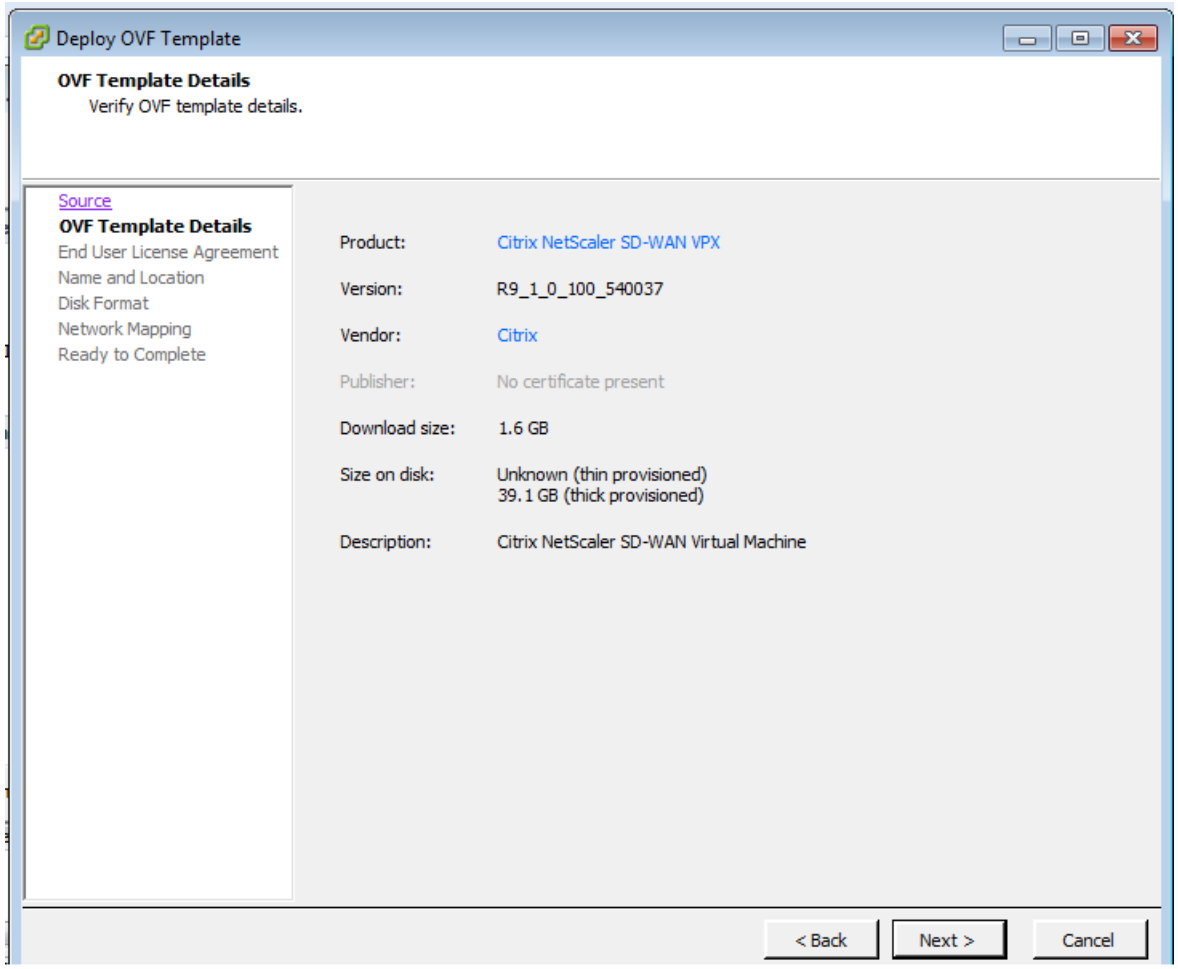

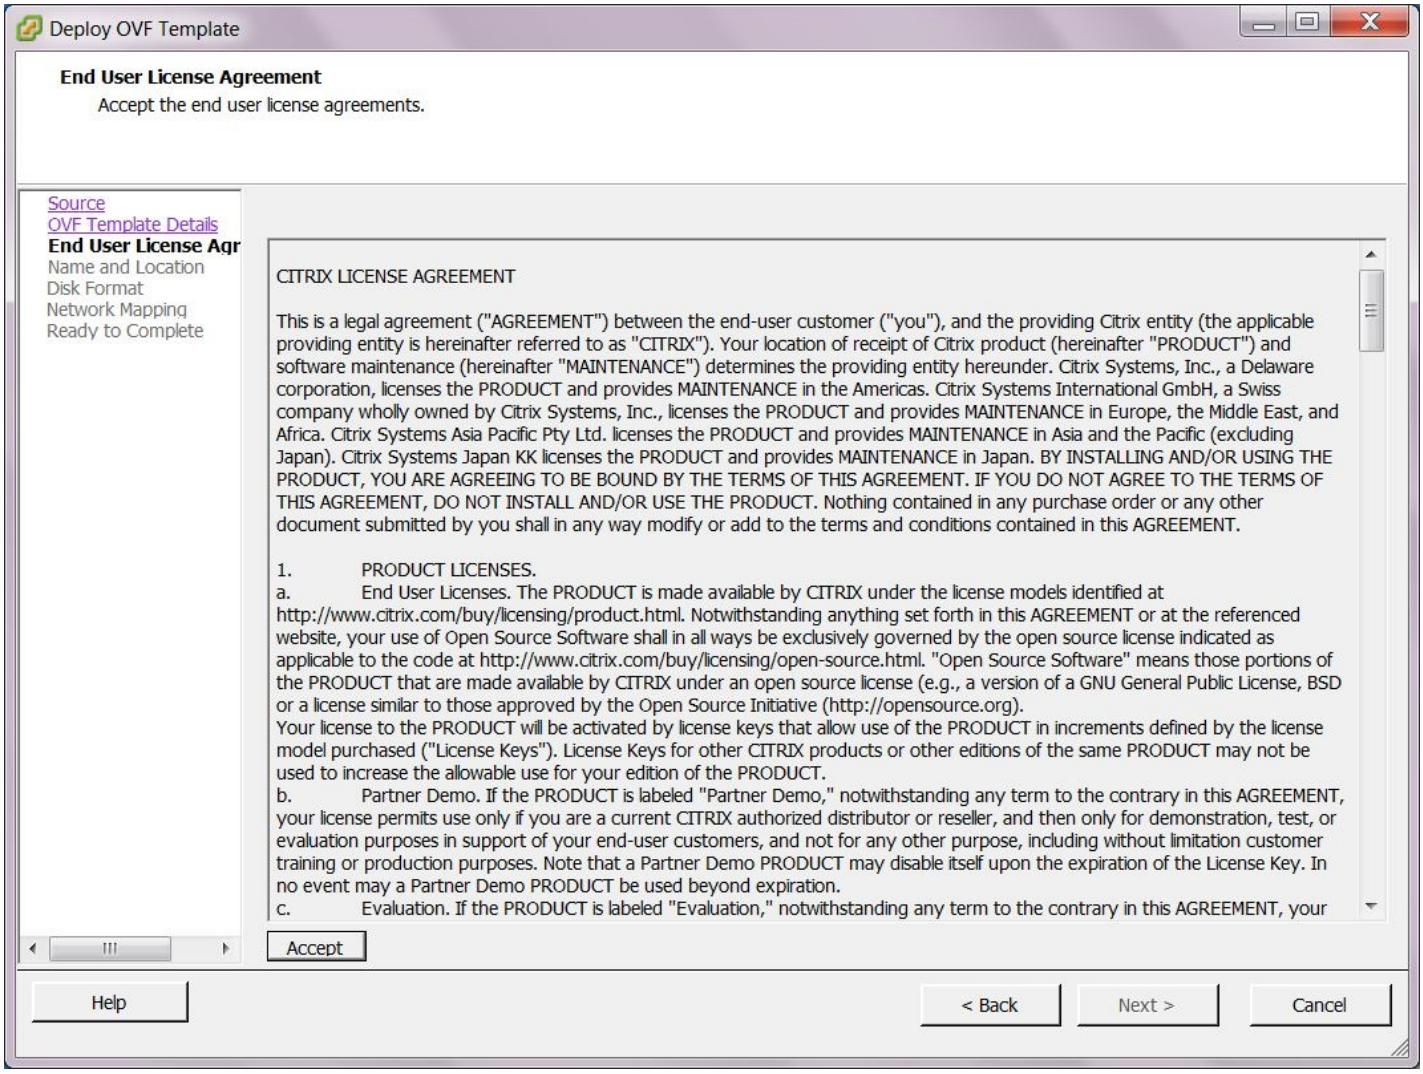

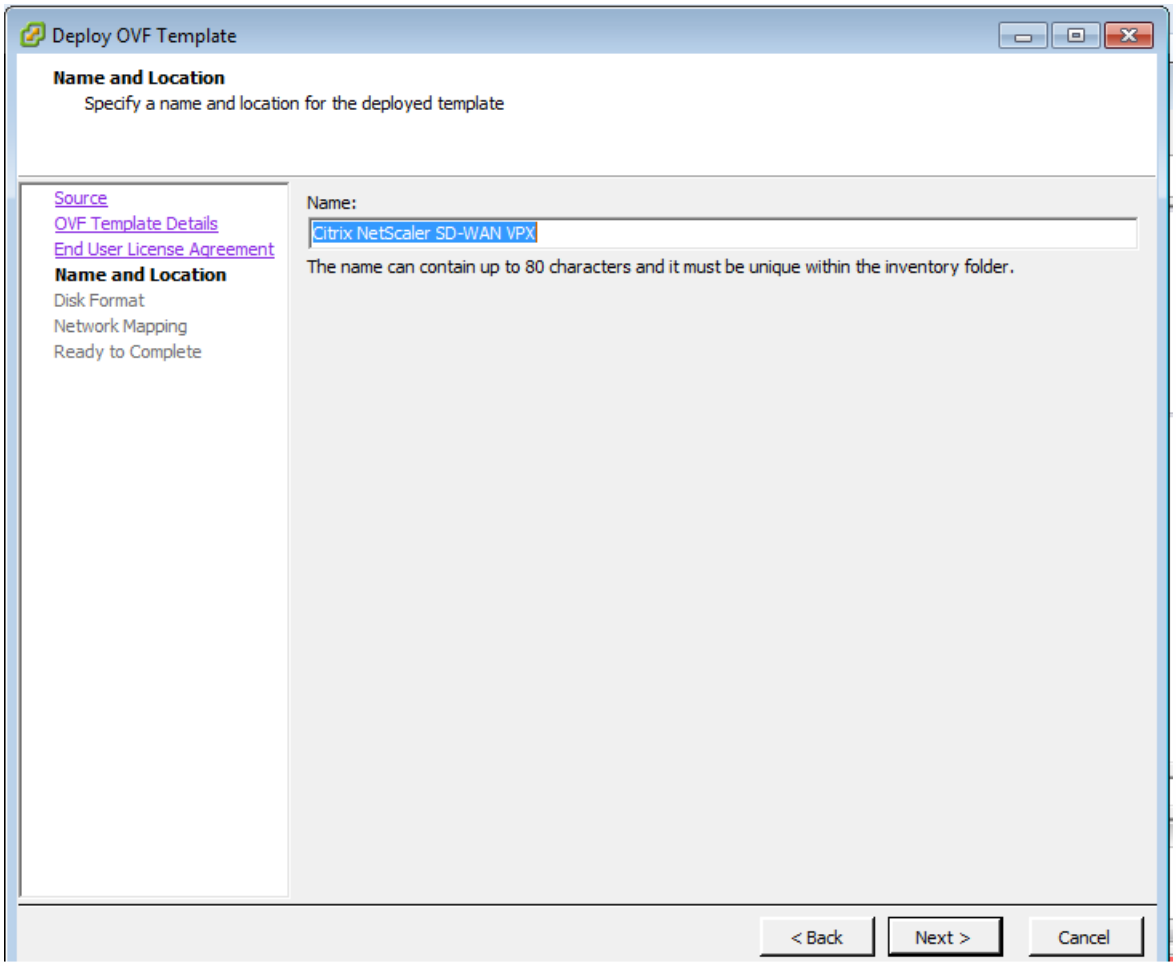

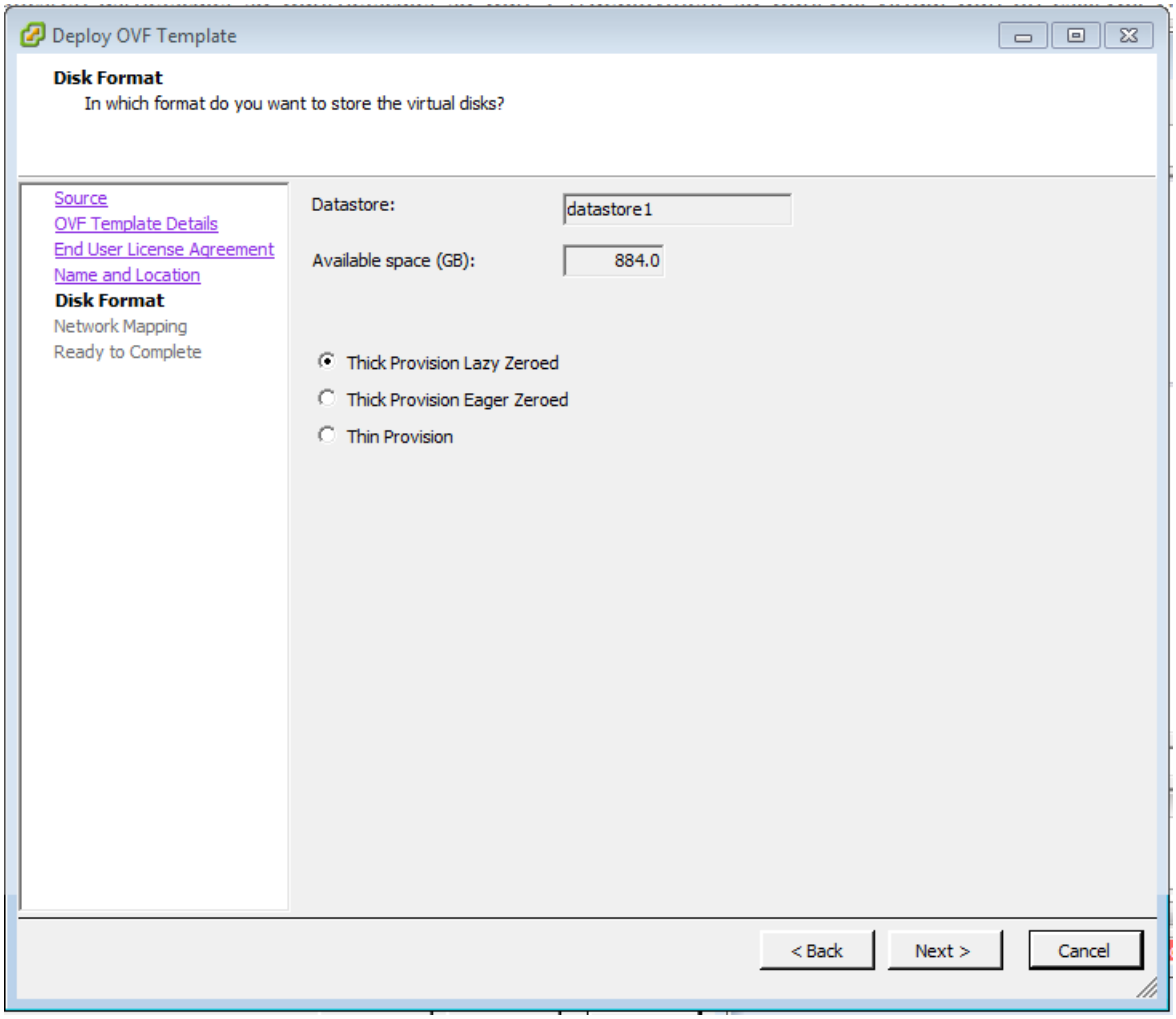

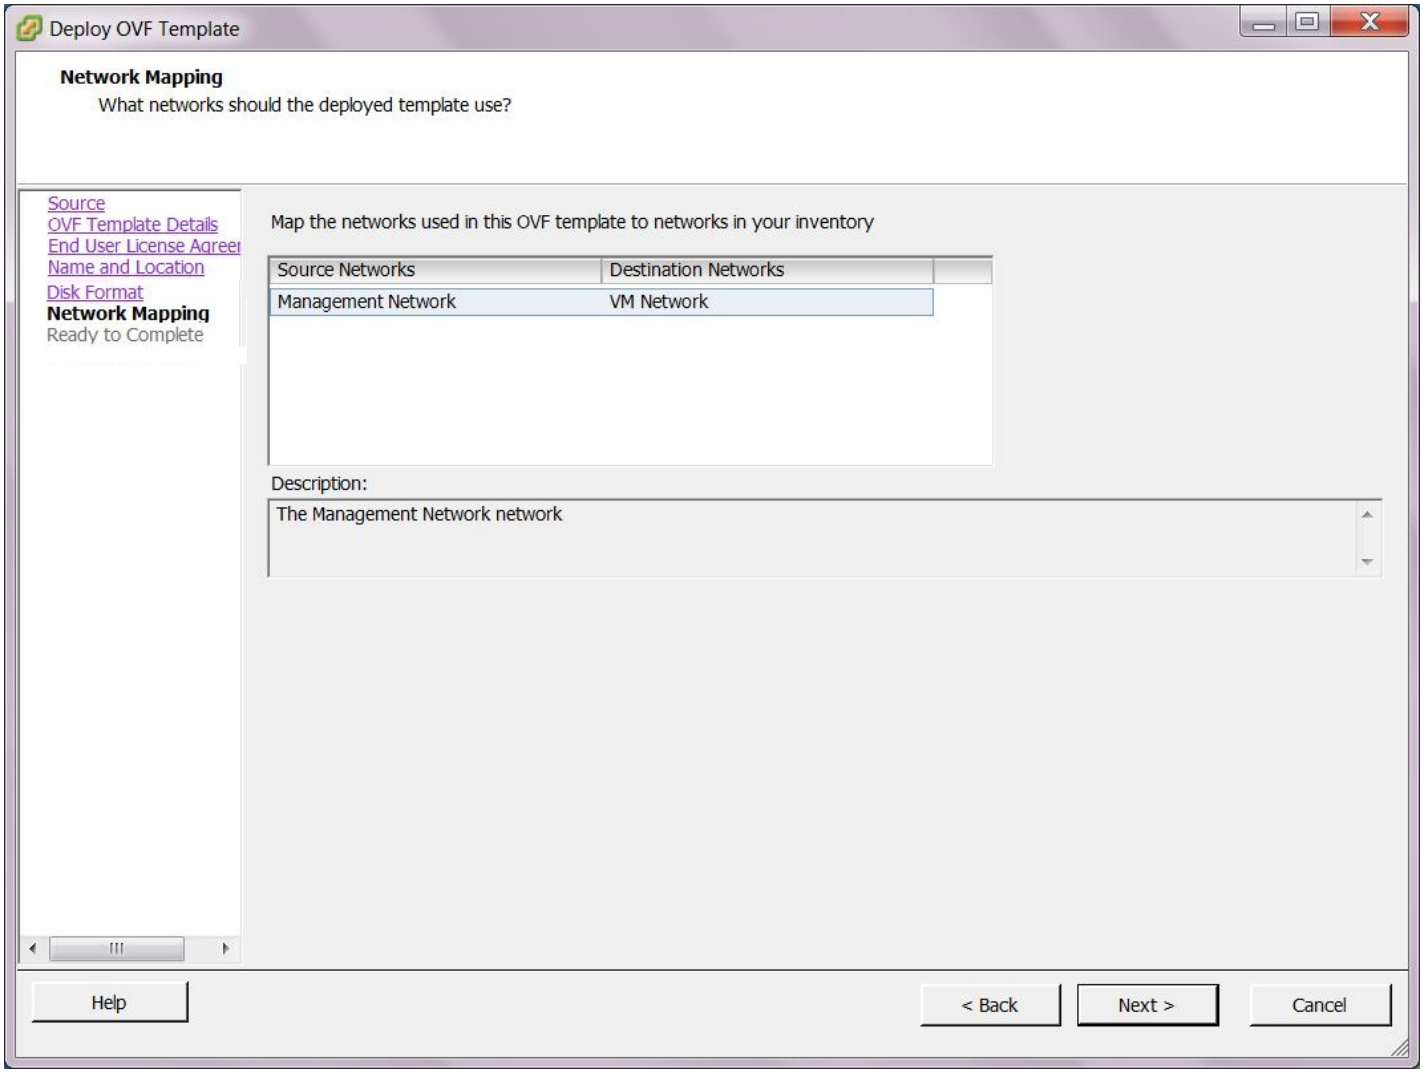

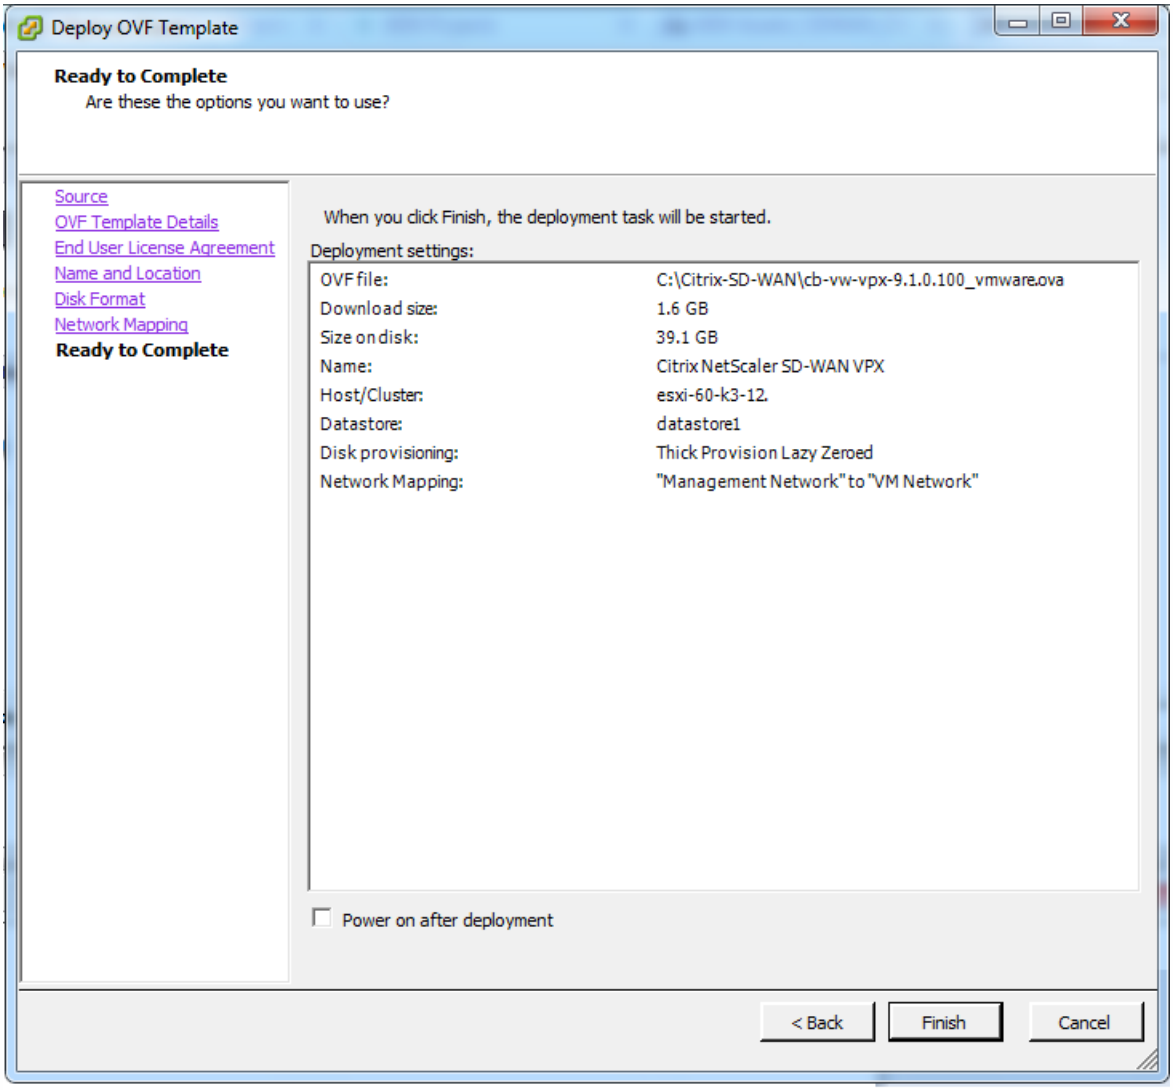

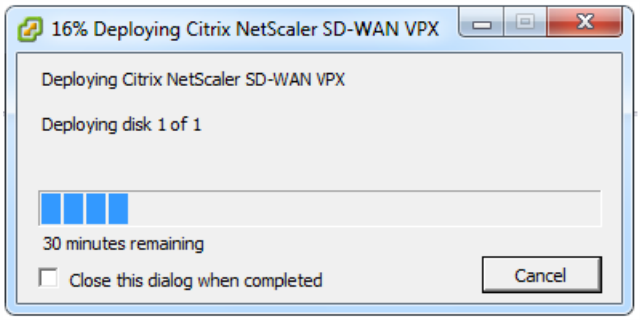

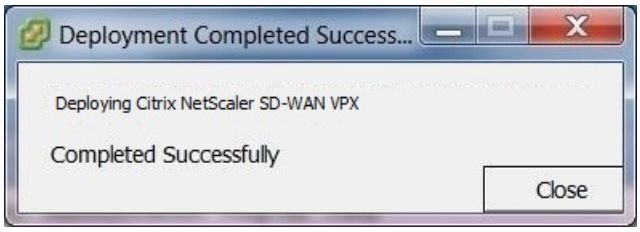

-
-

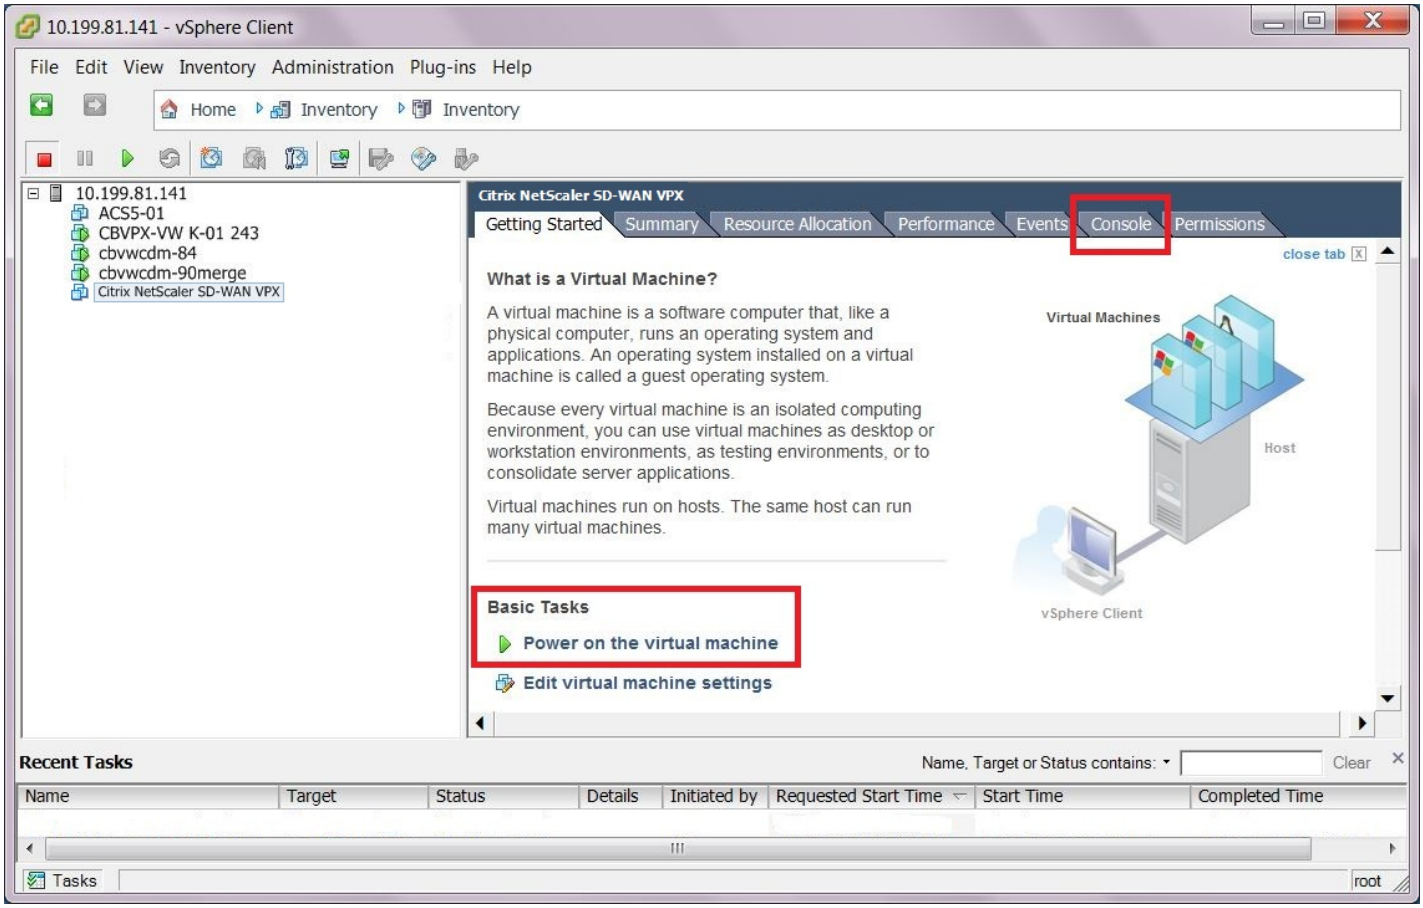

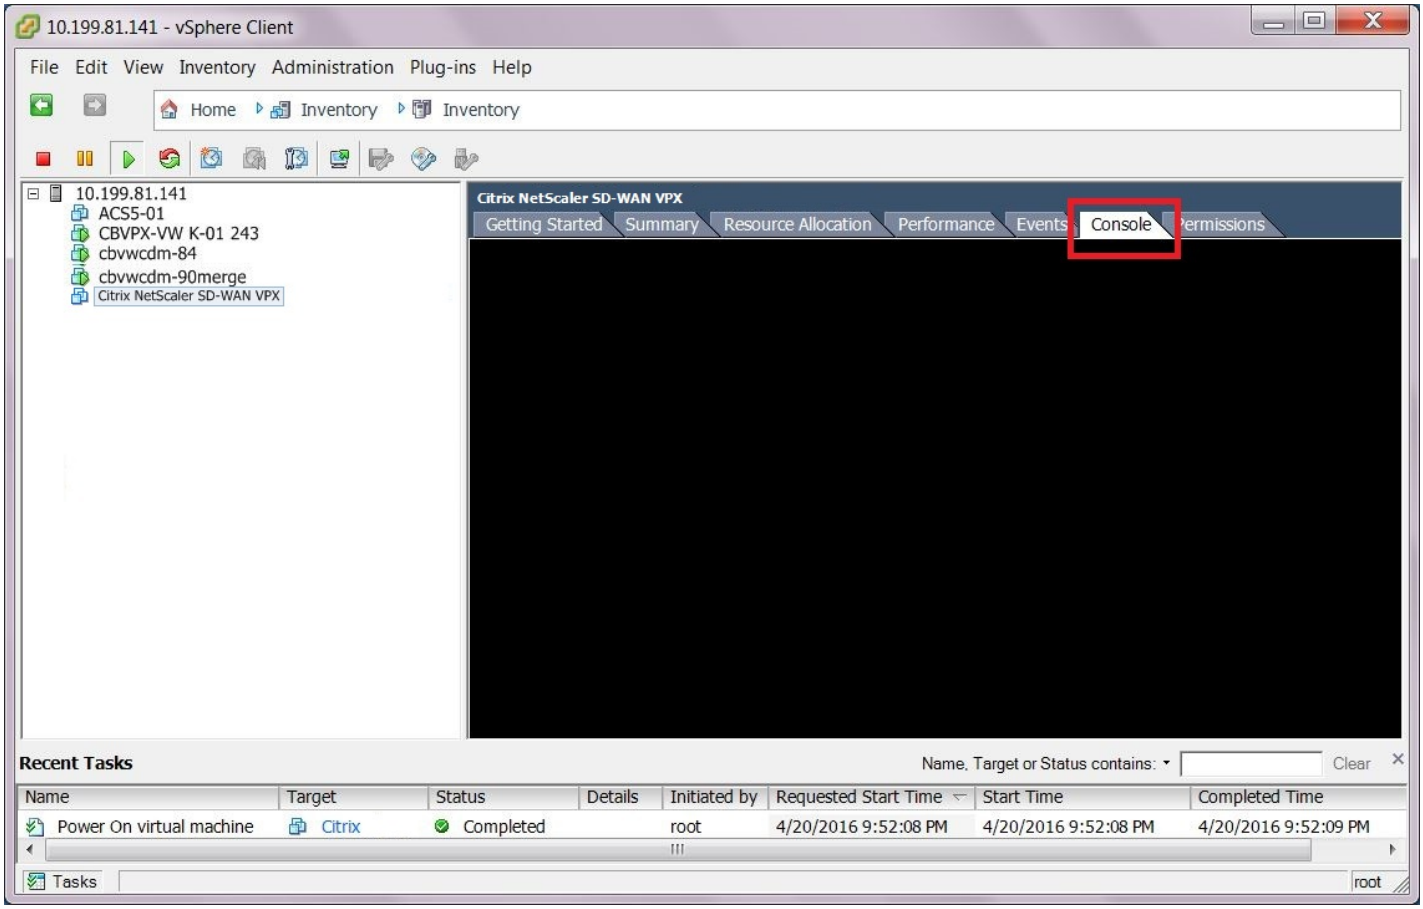

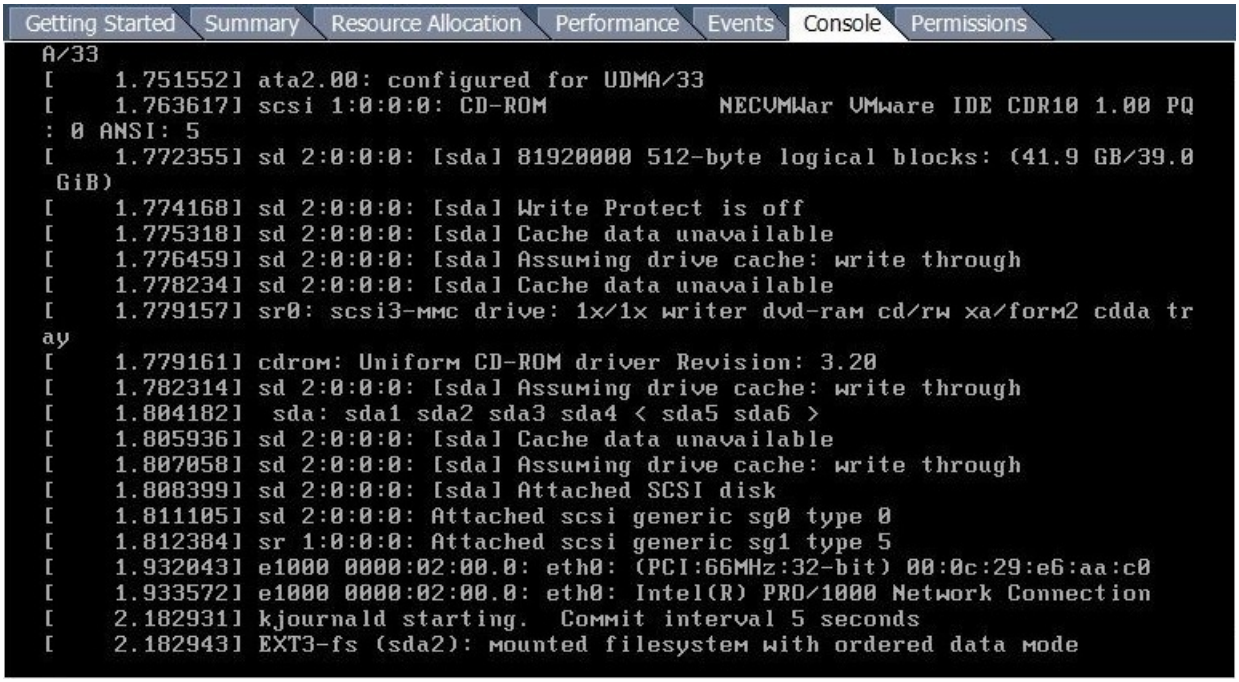

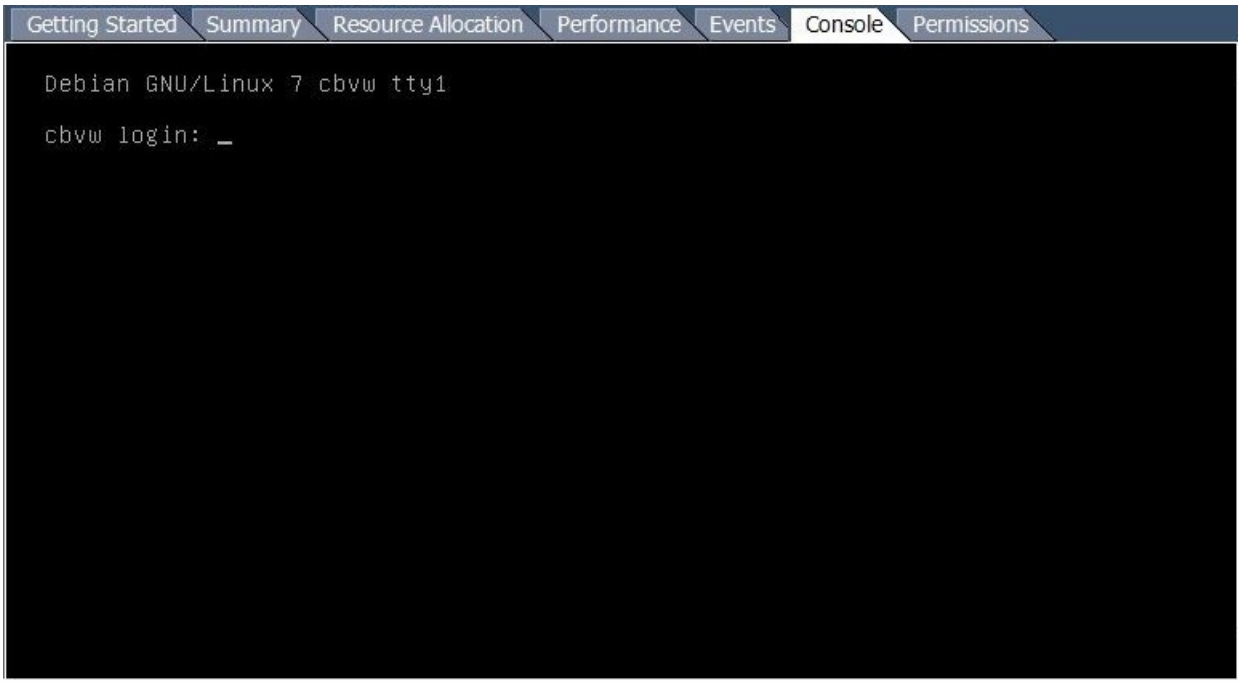

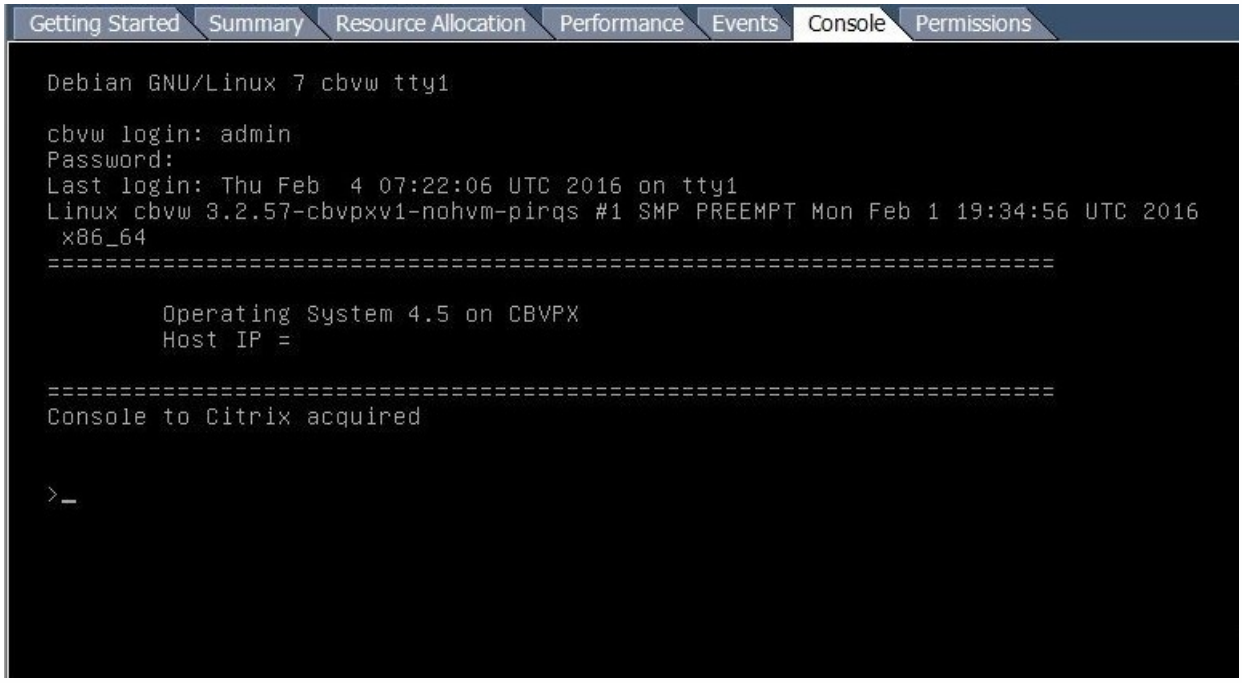

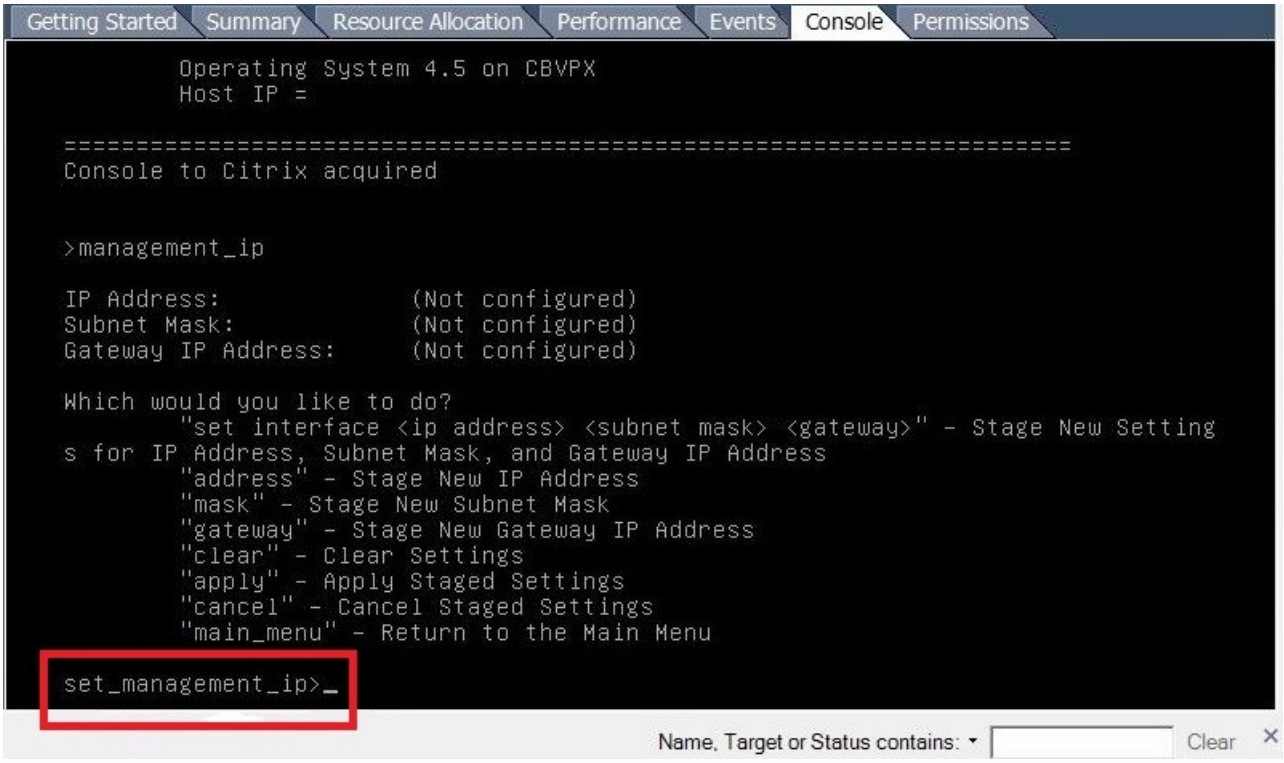

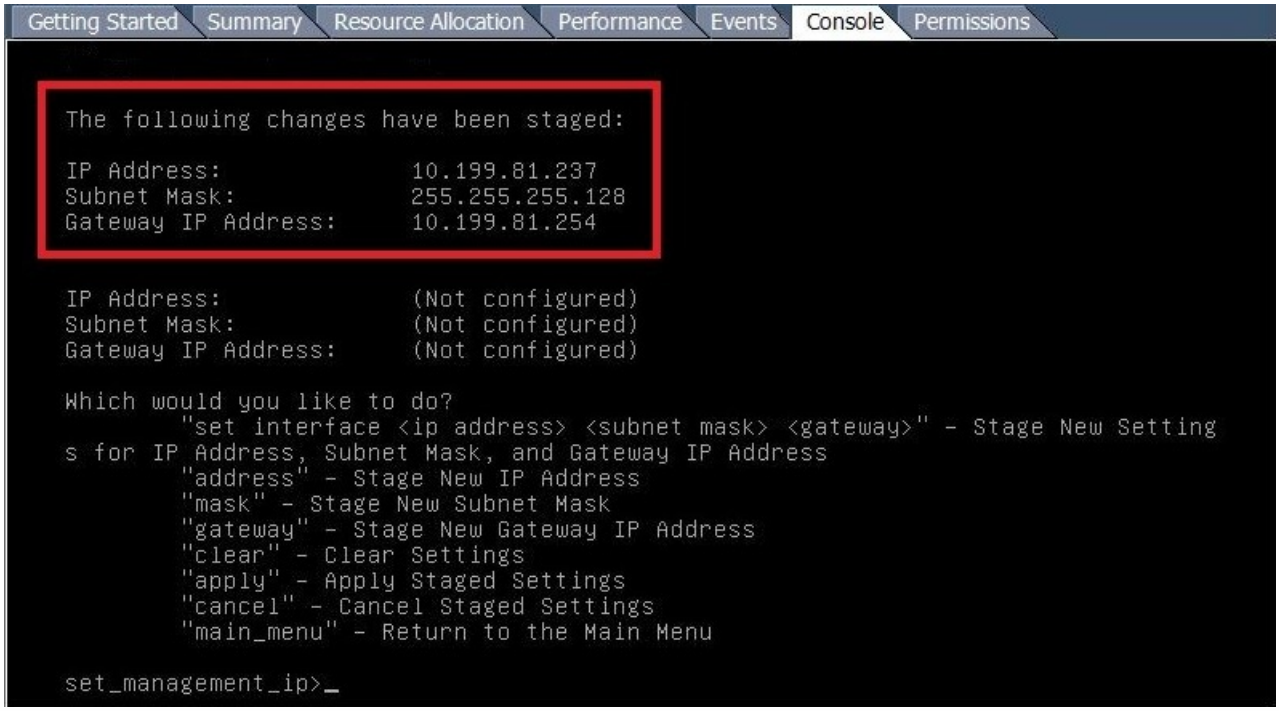

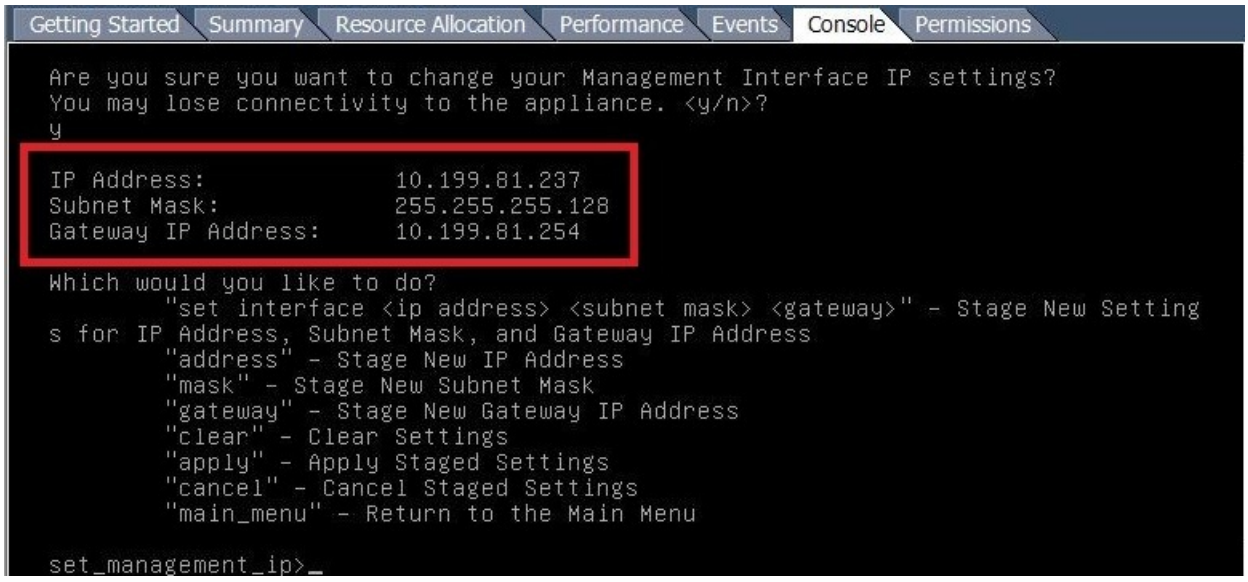

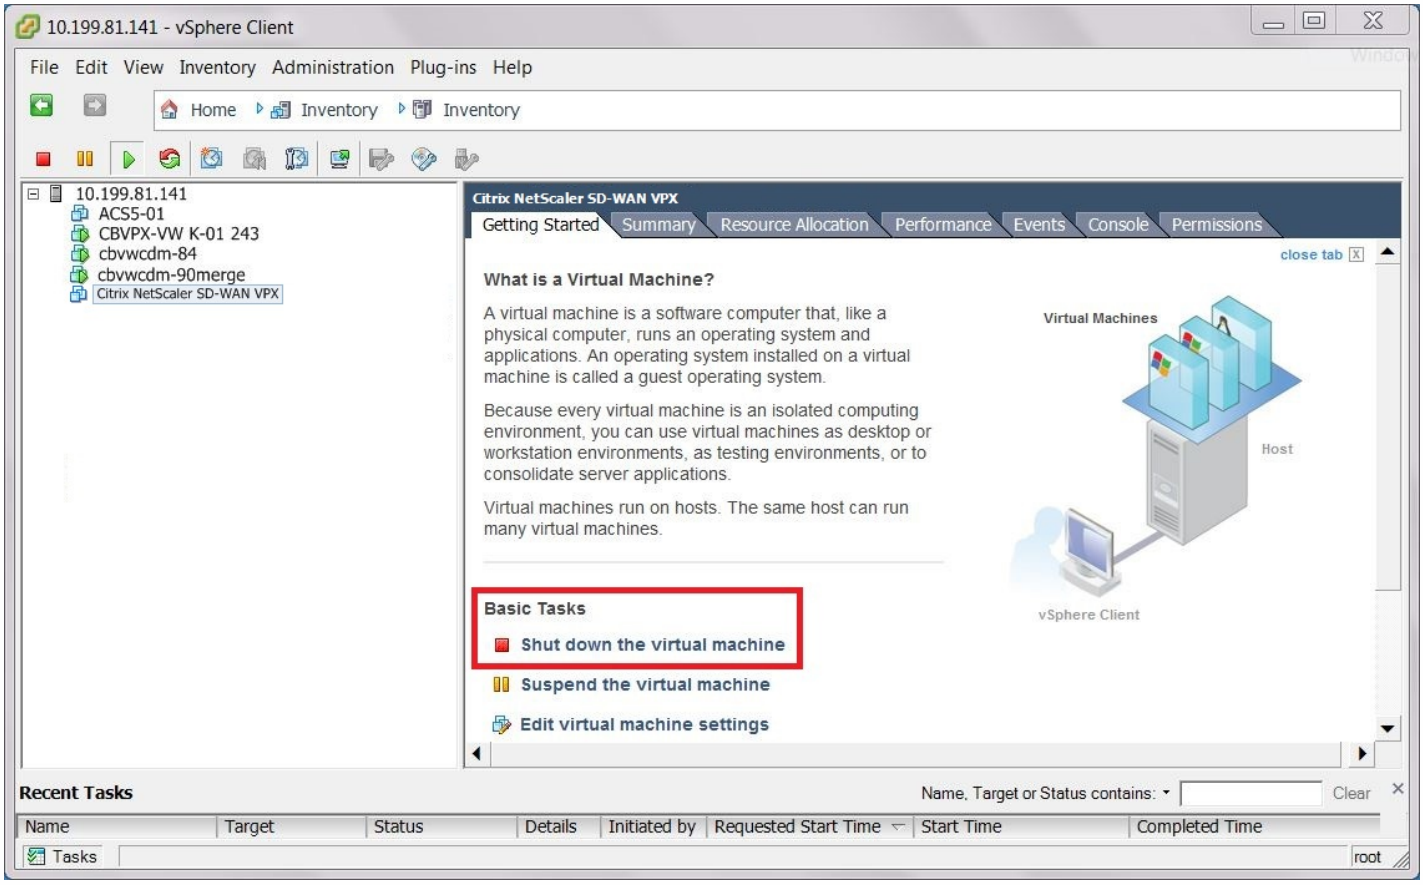

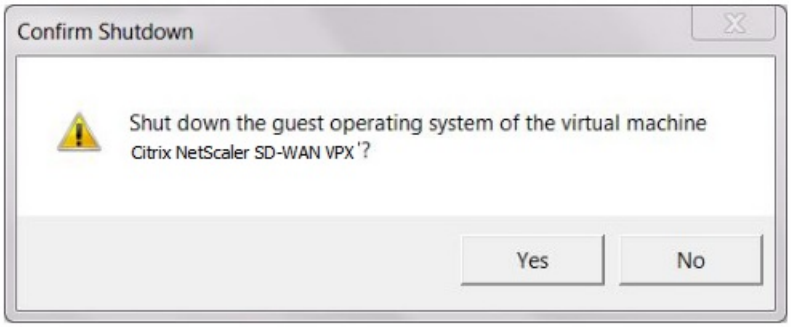

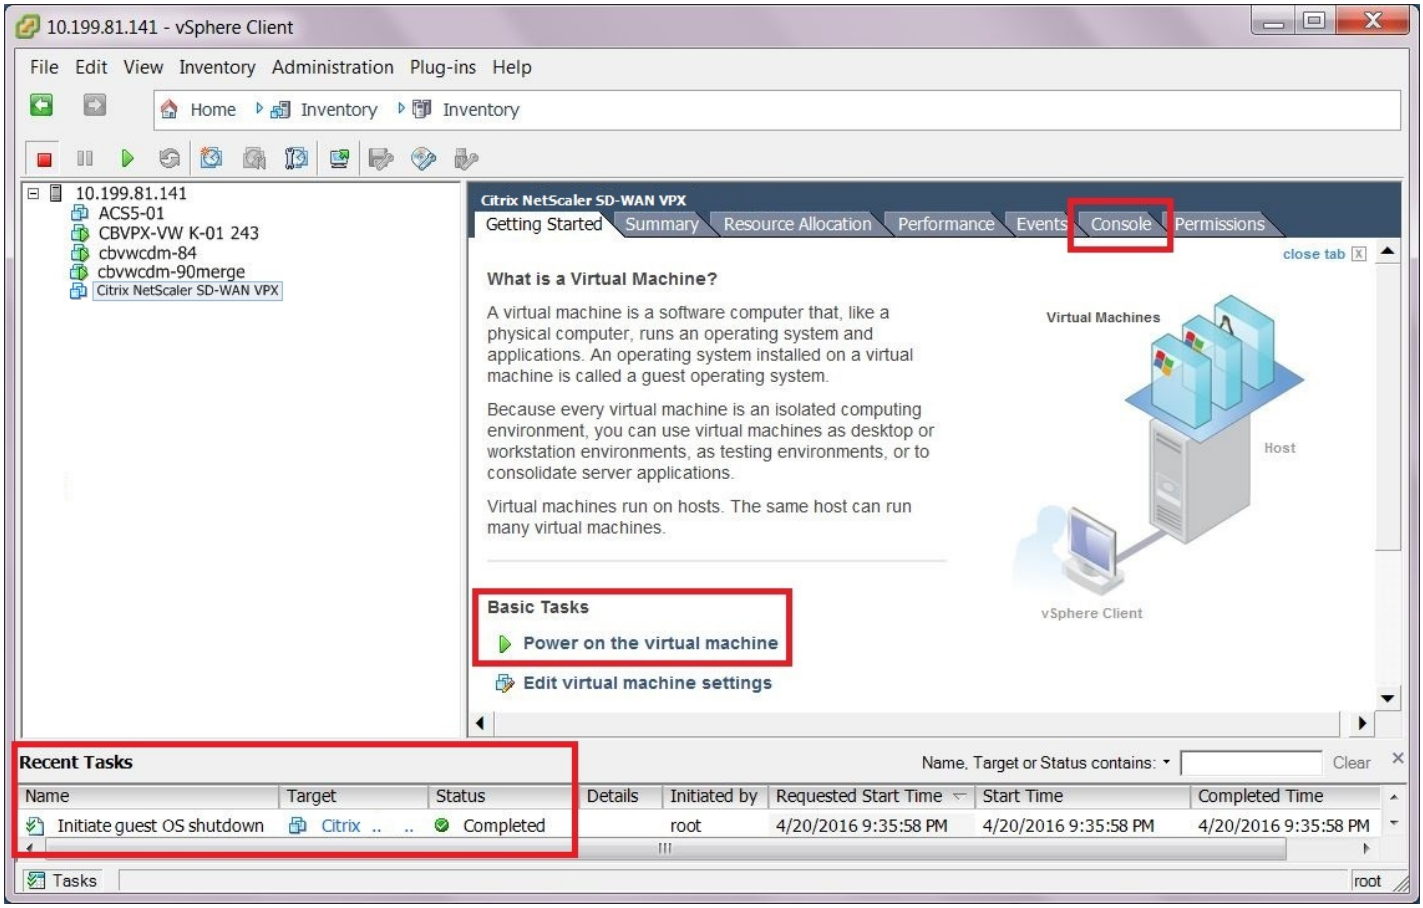

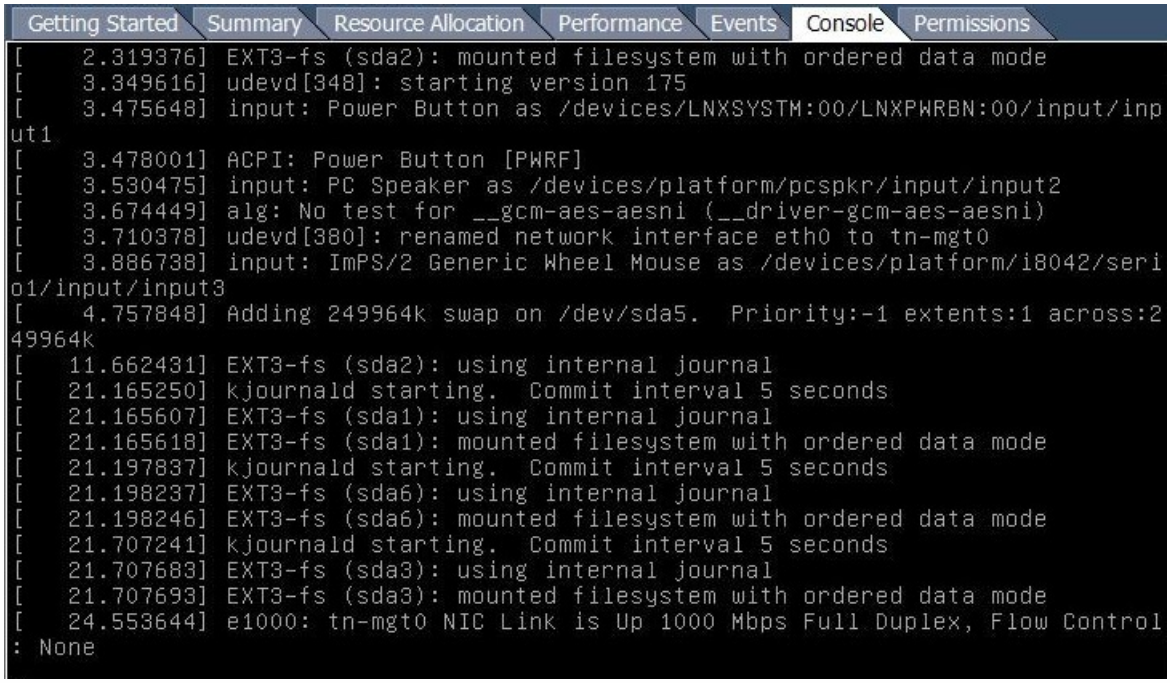

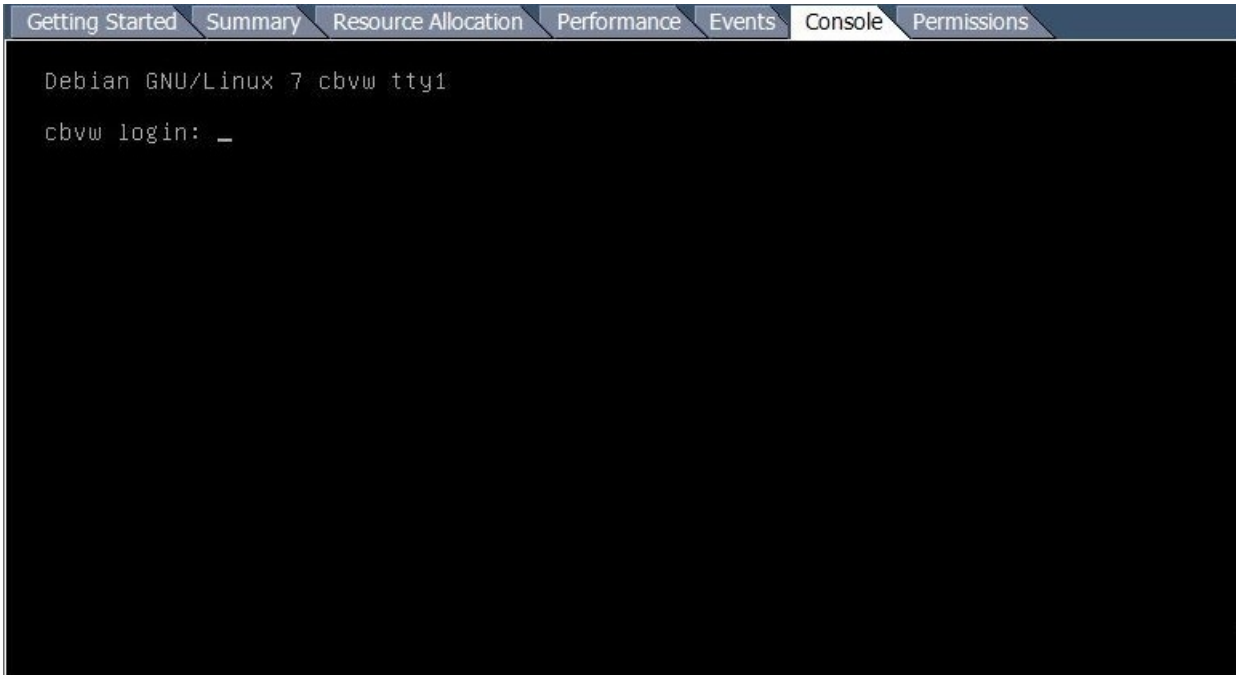

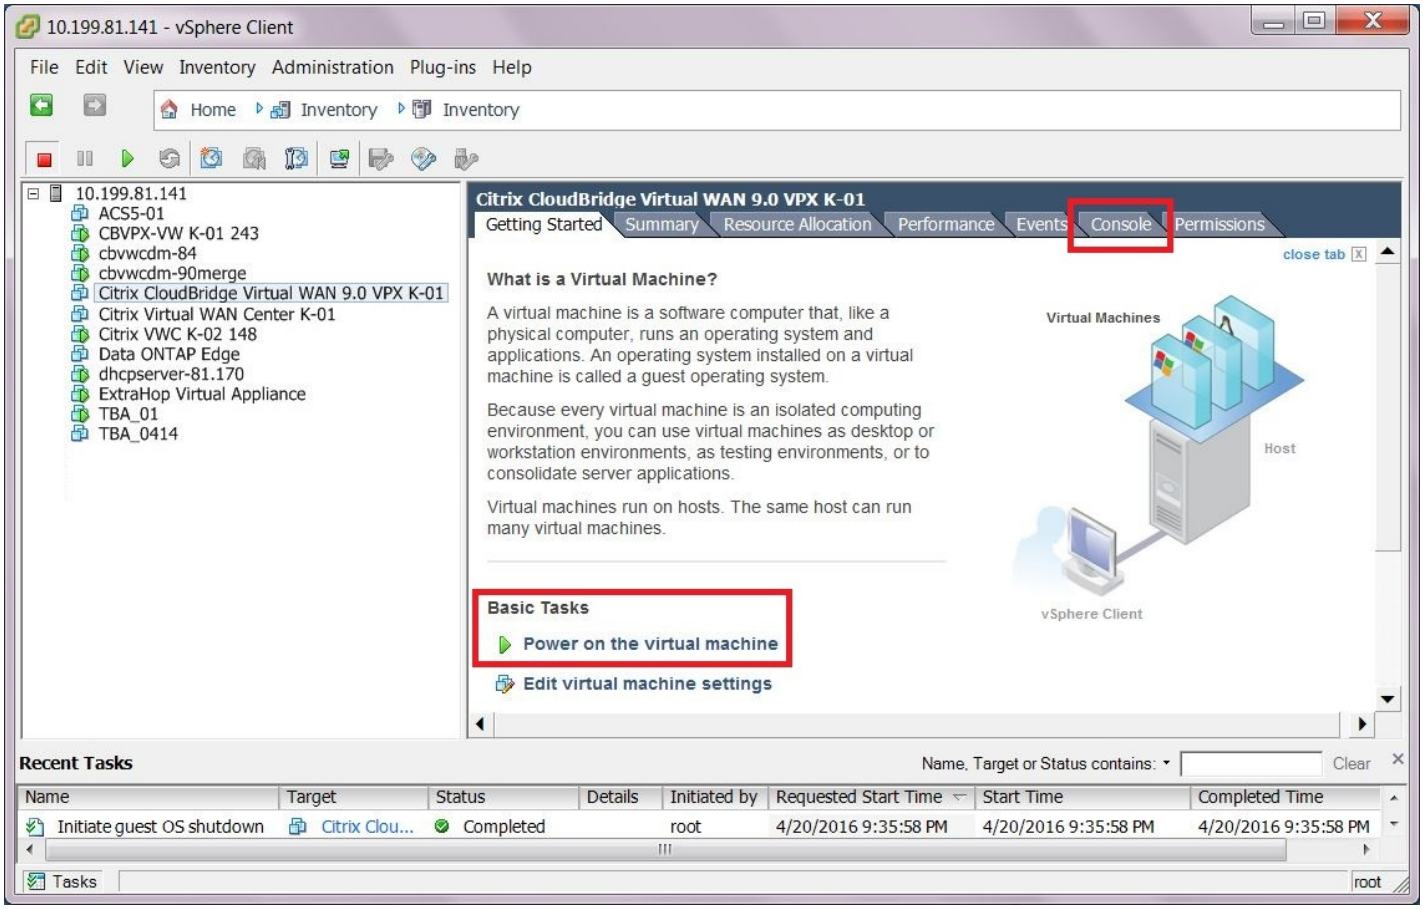

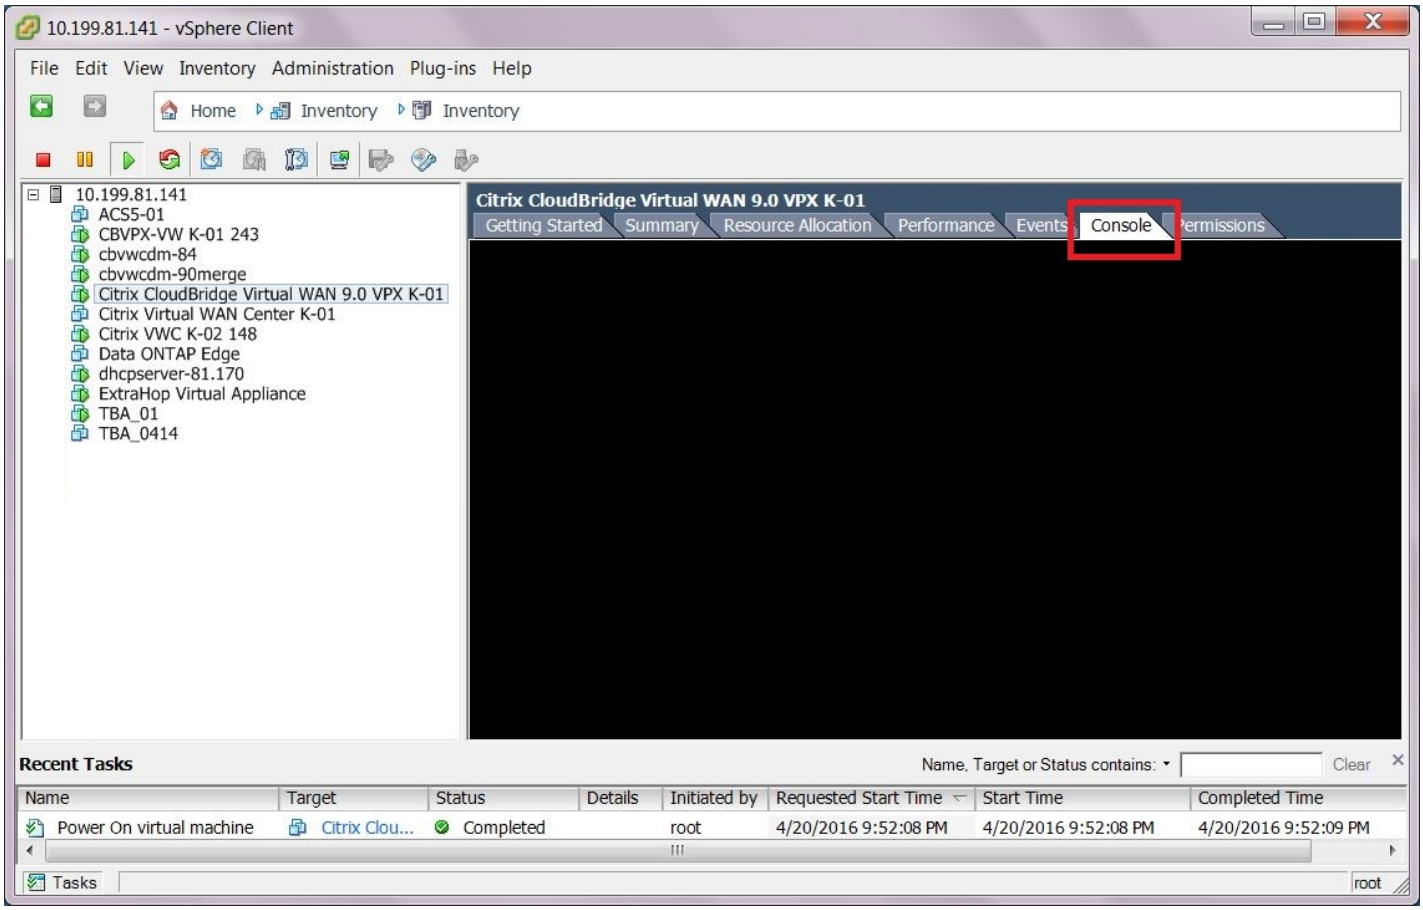

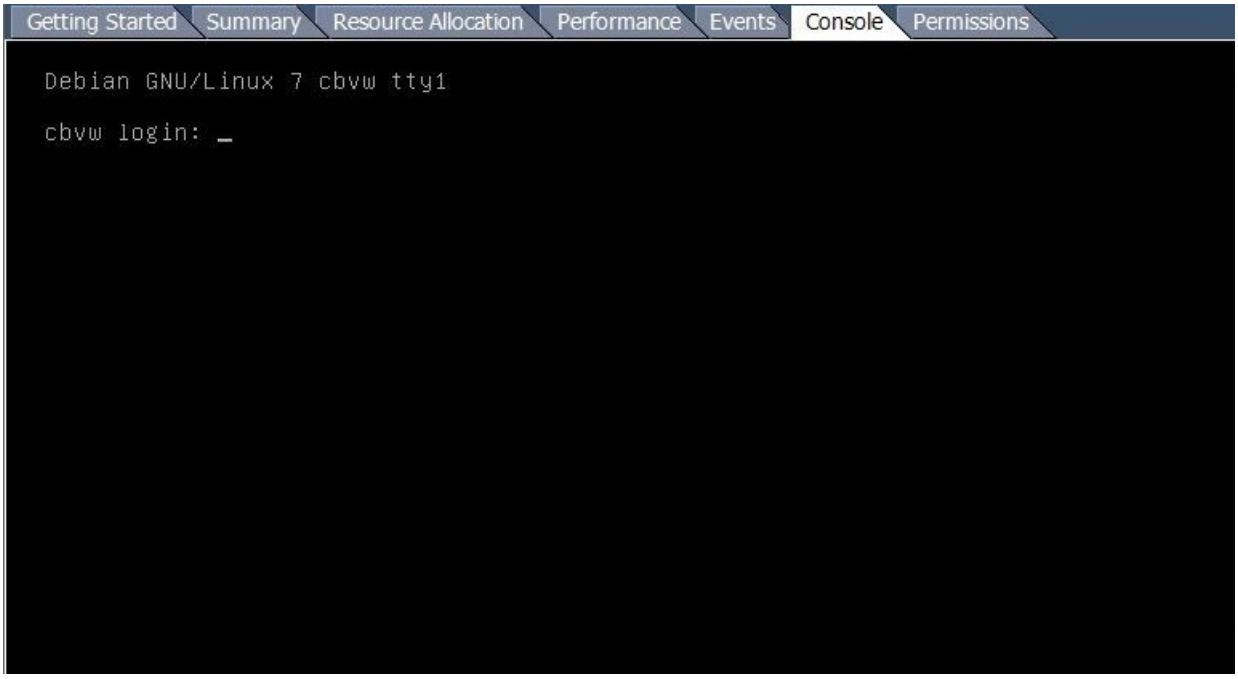

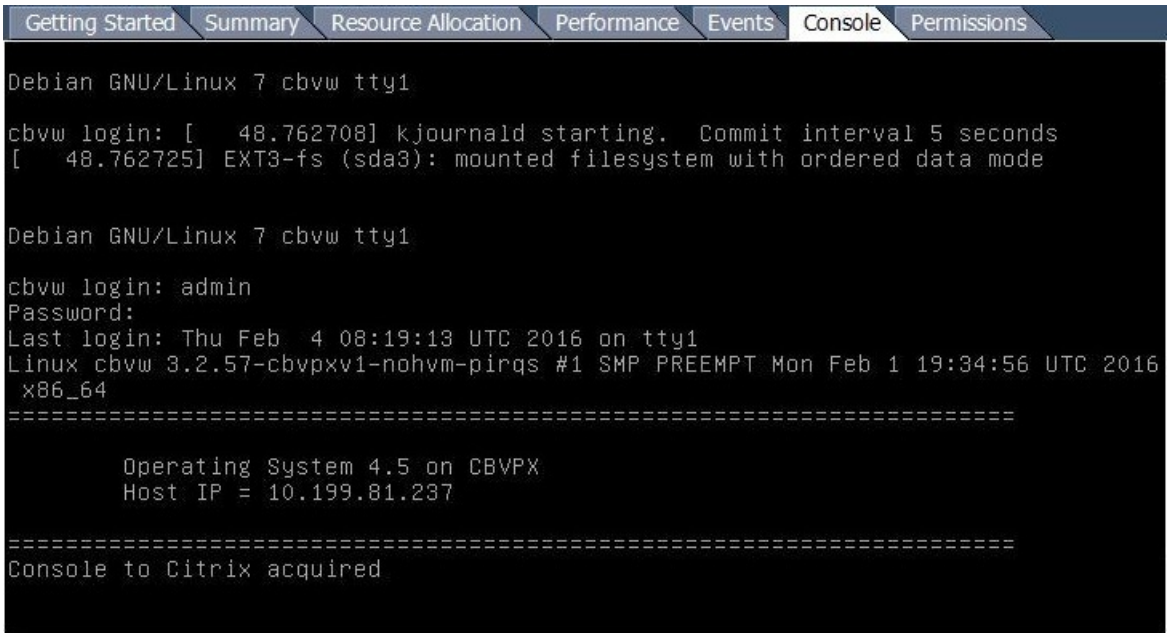

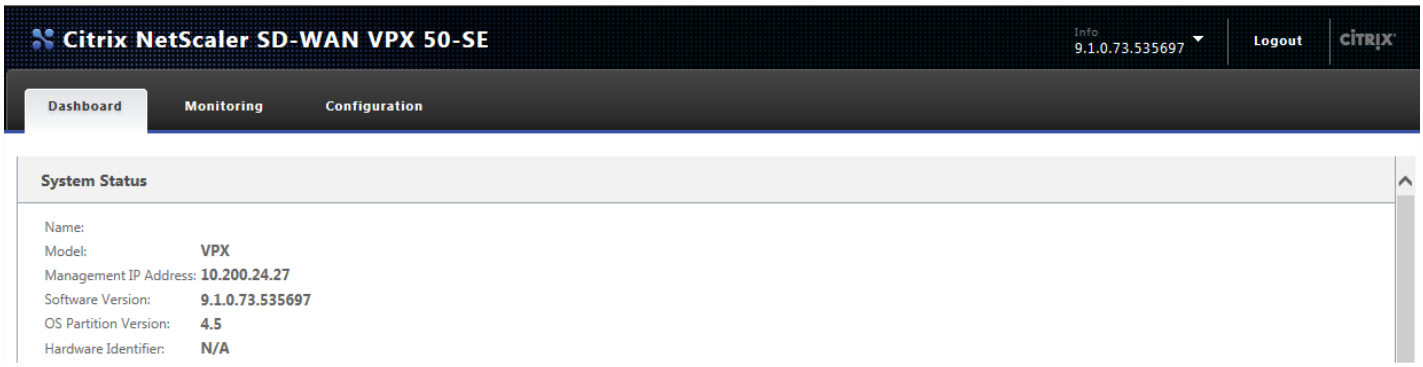

-
-

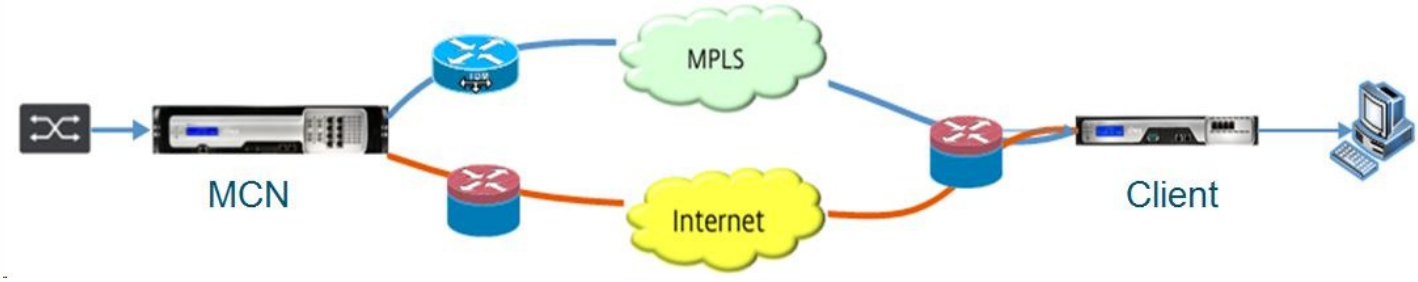

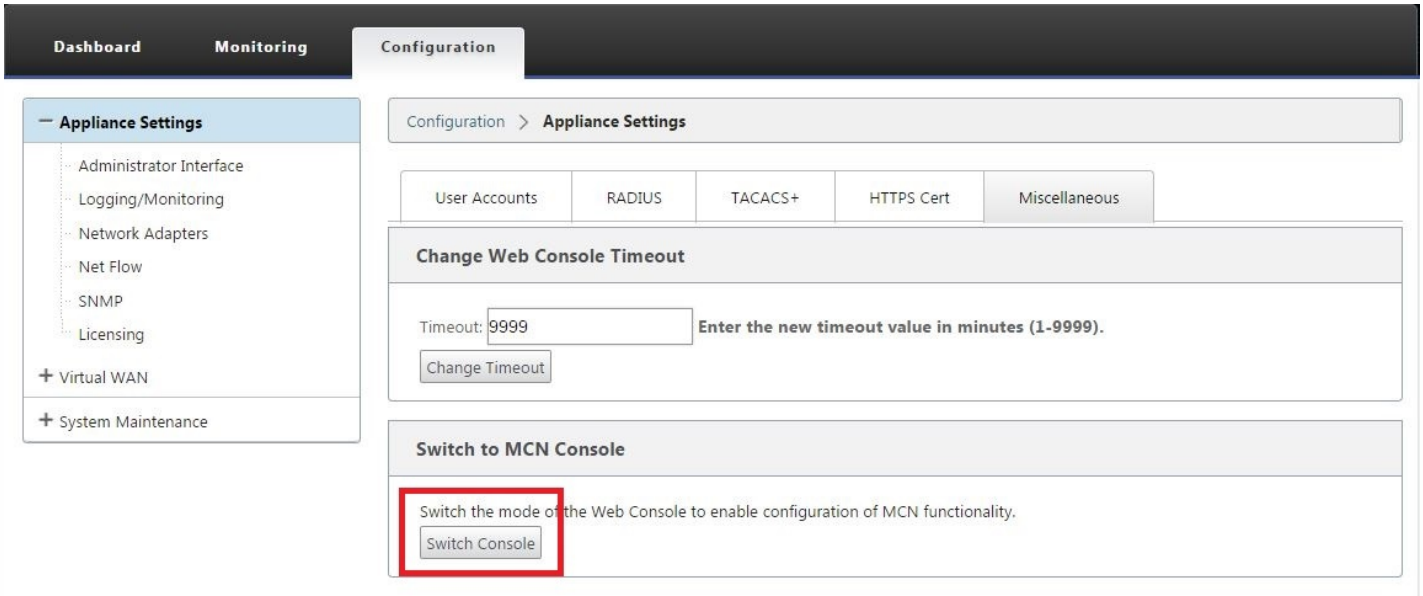

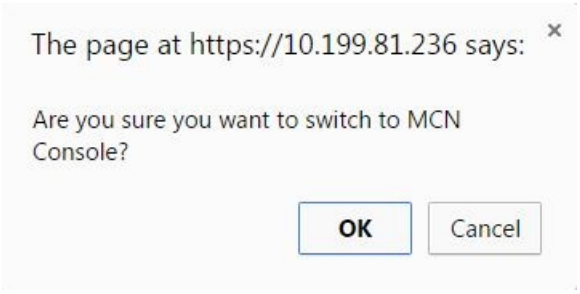

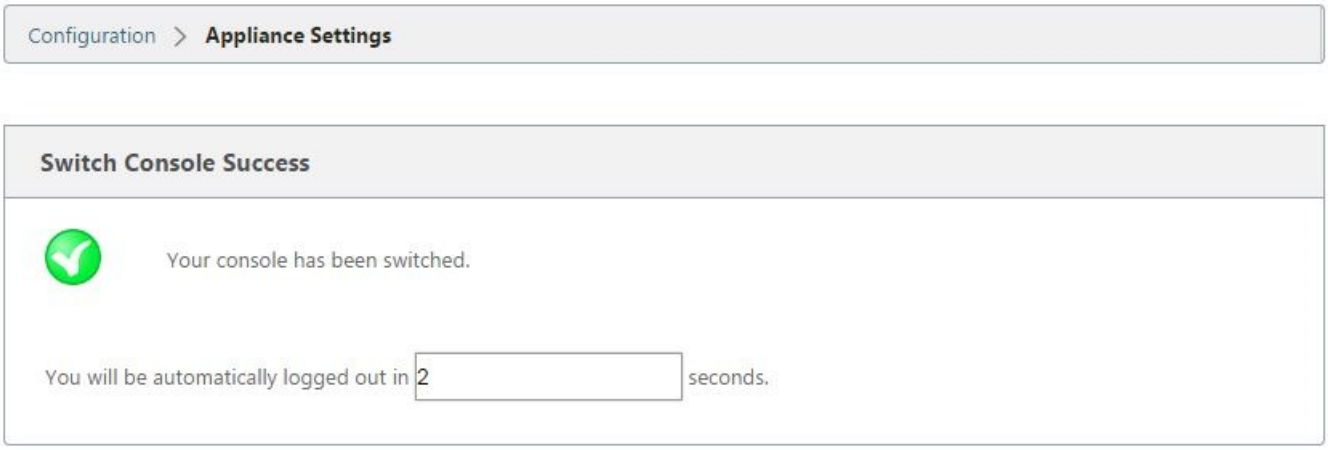

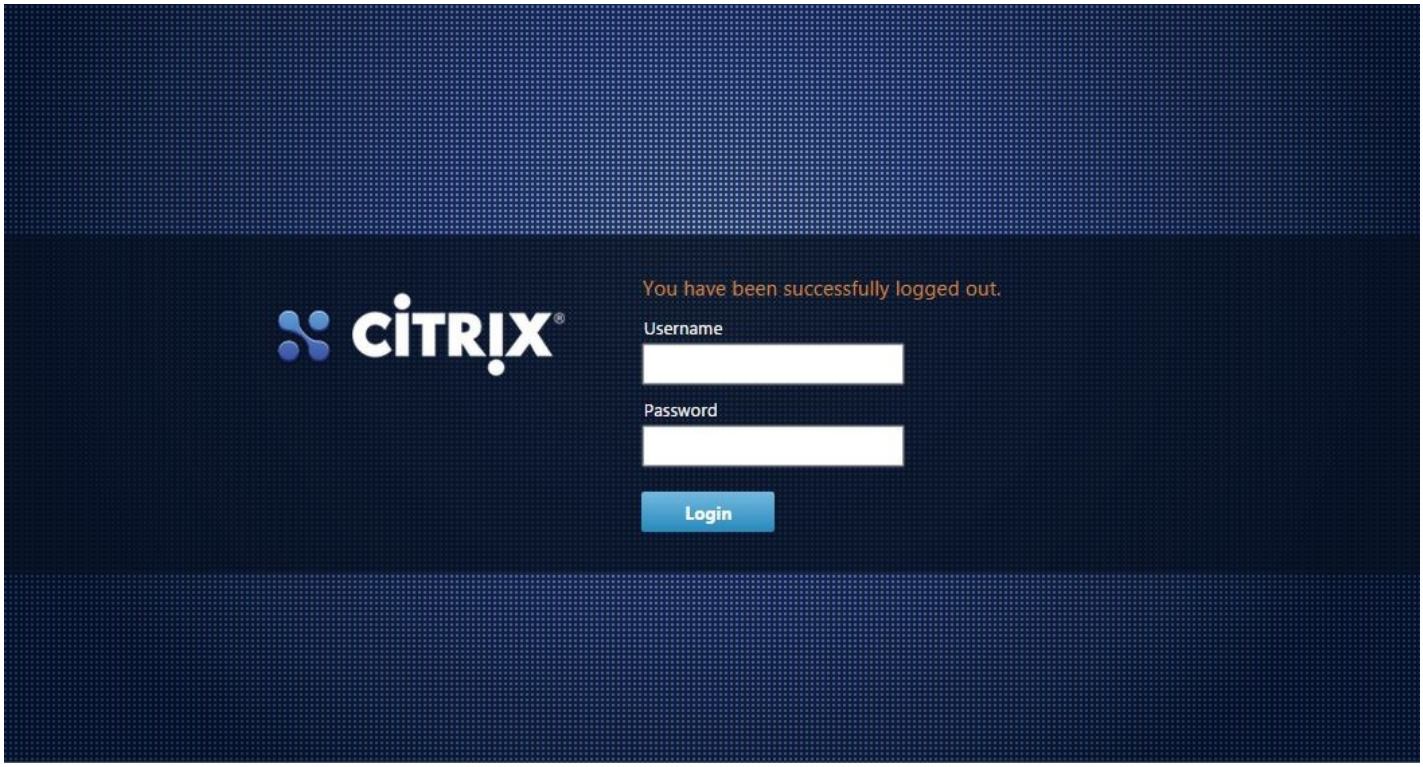

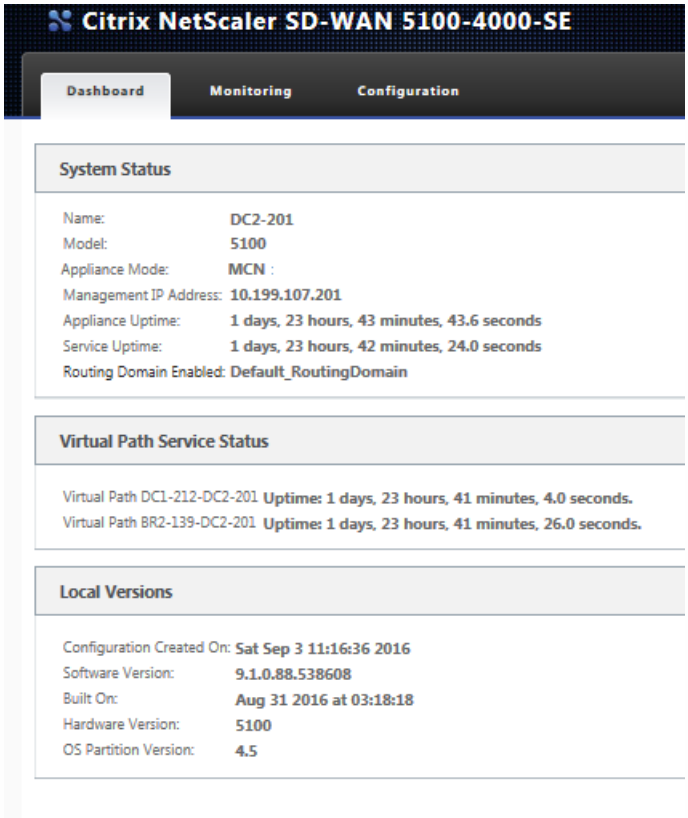

How to Add the MCN Site

Aug 09, 2017

The first step is to open a new configuration package, and add the MCN site to the new configuration.

> ## Note
>
> It is strongly recommended that you save the configuration package often, or at key points in the configuration. Instructions are provided in the section Naming, Saving, and Backing Up the MCN Site Configuration.

> ## Warning
>
> If the console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you set the console session Timeout interval to a high value when creating or modifying a configuration package, or performing other complex tasks. The default is 60 minutes; the maximum is 9999 minutes. For security reasons, you should then reset it to a lower threshold after completing those tasks. For instructions, see the section Setting the Console Session Timeout Interval (Optional).

To add and begin configuring the MCN appliance site, do the following:

1. In the navigation tree, open the **Virtual WAN** branch and select **Configuration Editor**.

> ## Note
>
> The **Configuration Editor** is available in **MCN Console** mode, only. If the **Configuration Editor** option is not available in the Virtual WAN branch of the navigation tree, please see section, Switching the Management Web Interface to MCN Console Mode, for instructions on changing the console mode.

This displays the **Configuration Editor** main page (middle pane).

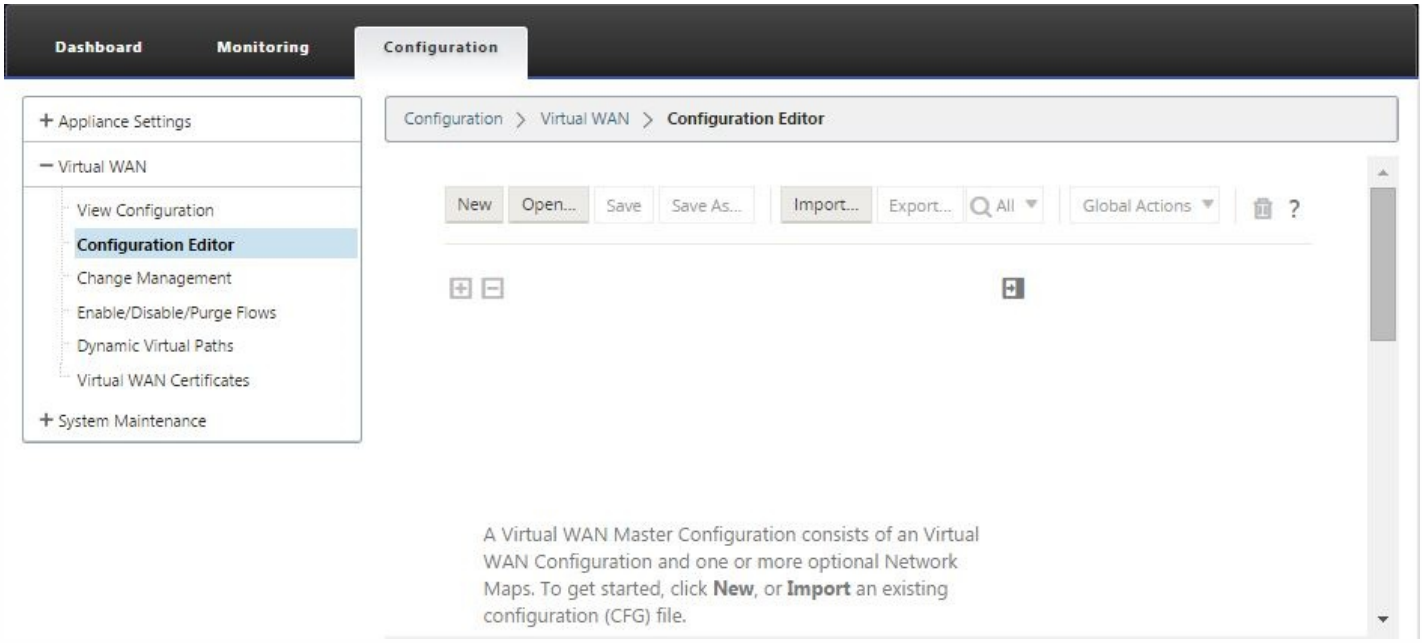

2. Click **New** to start defining a new configuration.

   This displays the **New** configuration settings page.

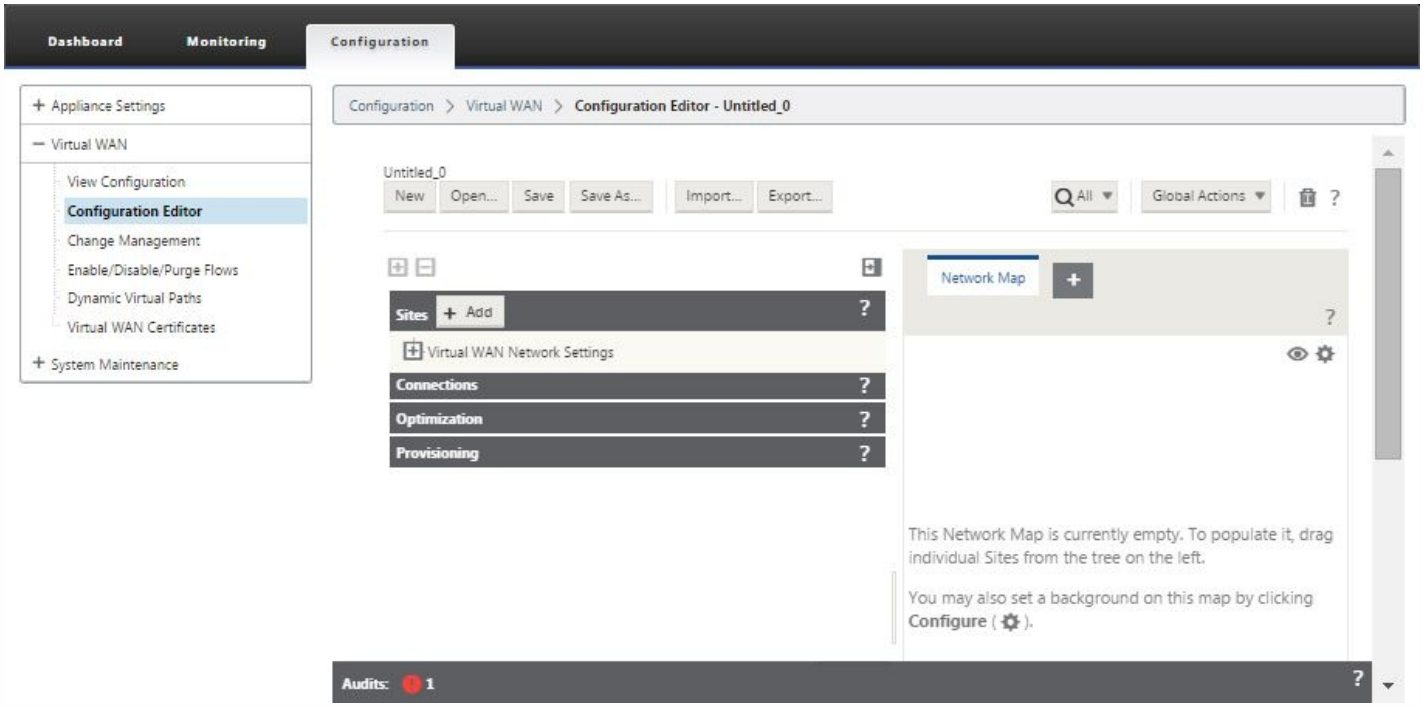

3. Click **Add** in the **Sites** bar to begin adding and configuring the MCN site.

   This displays the **Add Site** dialog box.

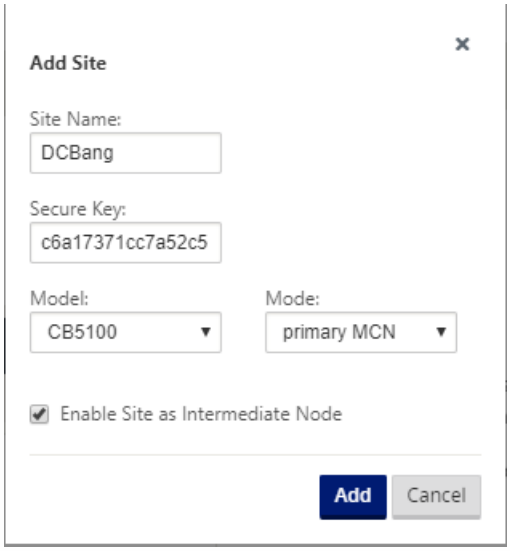

4. Enter the site information.

   Do the following:

   a. Enter the **Site Name**, **Appliance Name**, and **Secure Key**.

   b. Select the appliance **Model**.

> ## Note
>
> The Model options menu lists the generic model names for the supported appliance models. The generic names do not include the –Standard Edition model suffix, but do correspond to the equivalent SD-WAN Appliance models. Select the corresponding model number for this SDWAN Appliance model. (For example, select NetScaler 4000 if this is a NetScaler SD-WAN 4000-SE appliance.)

   c. Select **primary MCN** as the mode.

> ## Note
>
> You can add another site as the secondary MCN to support MCN redundancy. The secondary MCN continuosly monitors the health of the primary MCN, if the primary MCN fails, the secondary MCN assumes the role of the MCN. To confgure a site as a secondary MCN, in the **Mode** field secelect **Secondary MCN**.

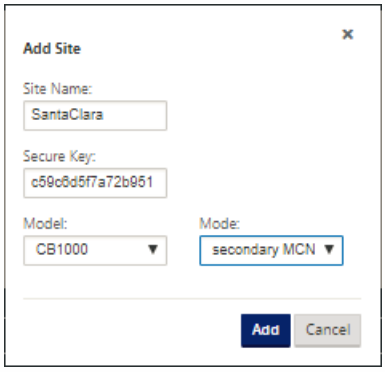

d. Select **Enable Site as Intermediate Node,** If this option is enabled, the site serves as a mediator for the creation and deletion of Dynamic virtual paths between two or more sites connected to this site.

> ## Note
>
> Entries cannot contain spaces and must be in Linux format.

5. Click **Add** to add the site.

   This adds the new site to the **Sites** tree, and displays the **Basic Settings** configuration form for the new site.

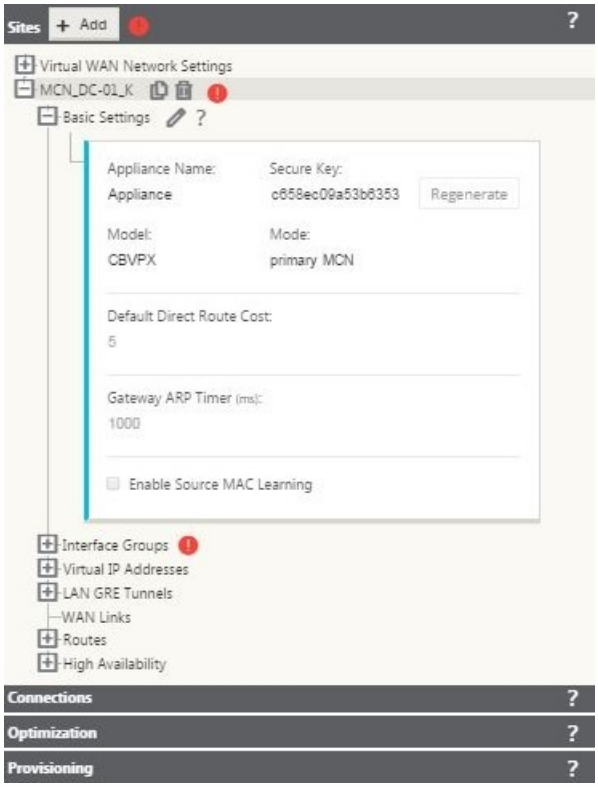

> ## Note
>
> After you click **Add**, audit warnings appear indicating that further action is required. A red dot or goldenrod delta icon indicates an error in the section where it appears. You can use these warnings to identify errors or missing configuration information. Roll your cursor over an audit warning icon to display a short description of the error(s) in that section. You can also click the dark grey **Audits** status bar (bottom of page) to display a complete list of all unresolved audit warnings.

6. Enter the basic settings for the new site, or accept the defaults.

7. (Optional, strongly recommended) Save the configuration-in-progress.

   If you cannot complete the configuration in one session, you can save it at any time, so you can return to complete it later. The configuration is saved to your workspace on the local appliance. To resume working in a saved configuration, click **Open** in the **Configuration Editor** menu bar (top of page area). This displays a dialog box for selecting the configuration

you want to modify.

> **Note**
>
> As an extra precaution, it is recommended that you use **Save As**, rather than **Save**, to avoid overwriting the wrong configuration package.

To save the current configuration package, do the following:

1. Click **Save As** (at the top of the **Configuration Editor** middle pane).

   This opens the **Save As** dialog box.

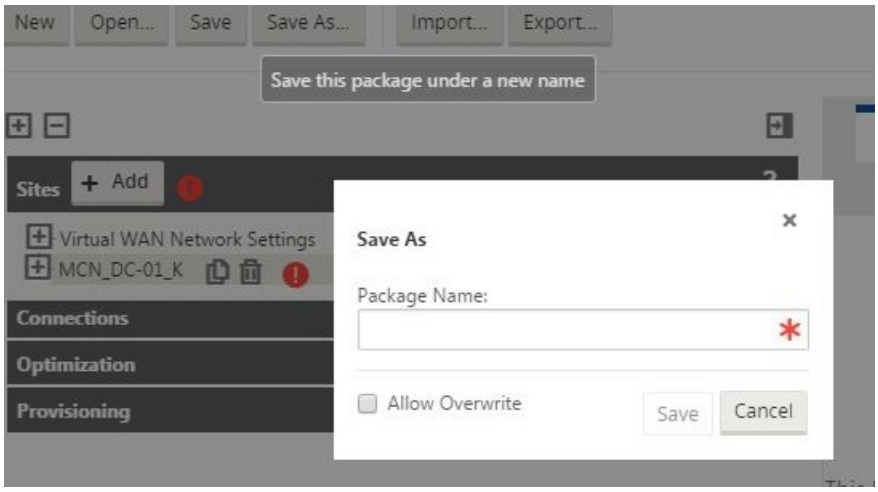

2. Enter the configuration package name.

> **Note**
>
> If you are saving the configuration to an existing package, be sure to select **Allow Overwrite** before saving.

3. Click **Save**.

How to Configure Virtual Interface Groups for the MCN Site

Aug 09, 2017

After adding the new MCN site, the next step is to create and configure the Virtual Interface Groups for the site.

The following are some guidelines for configuring Virtual Interface groups:

- Use logical names that will best describe the group.
- Trusted networks are networks that are protected behind a Firewall.
- Virtual Interfaces associate interfaces to Fail to Wire (FTW) pairs.
- Single WAN interfaces cannot be in an FTW pair.

> **Note**
>
> For additional guidelines and information on configuring Virtual Interface Groups, see the Virtual Routing and Forwarding section.

To add a Virtual Interface Group to the new MCN site, do the following:

1. Continuing in the **Sites** tree of the **Configuration Editor**, click **+** next to the name of the site you just added.

   This opens the configuration branches for the new site.

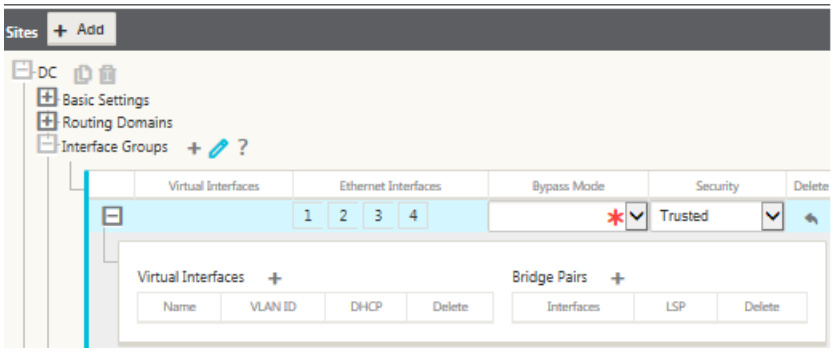

2. Click **+** to the left of the **Interface Groups** branch.

   This displays the **Interface Groups** table for the site.

3. Click **+** to the right of **Interface Groups**.

   This adds a new blank group entry to the table and opens it for editing.

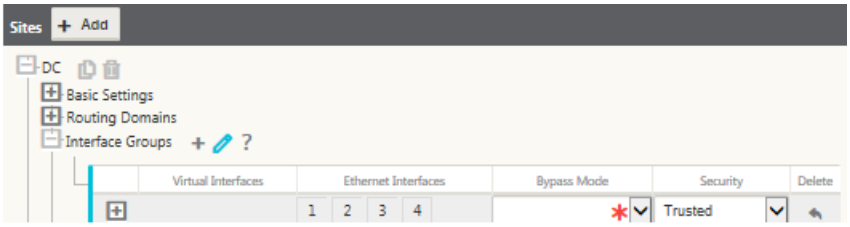

4. Select the Ethernet Interfaces to include in the group.

Under **Ethernet Interfaces**, click a box to include/exclude that interface. You can select any number of interfaces to include in the group. A goldenrod highlight indicates an included interface.

5. Select the **Bypass Mode** from the drop-down menu (no default).

The **Bypass Mode** specifies the behavior of bridge-paired interfaces in the Virtual Interface Group, in the event of an appliance or service failure or restart. The options are: **Fail-to-Wire** or **Fail-to-Block**.

6. Select the Security Level from the drop-down menu.

This specifies the security level for the network segment of the Virtual Interface Group. The options are: **Trusted** or **Untrusted**. Trusted segments are generally protected by a firewall (default is Trusted).

7. Click **+** at the left edge of the new blank entry.

This displays the **Virtual Interfaces** and **Bridge Pairs** tables.

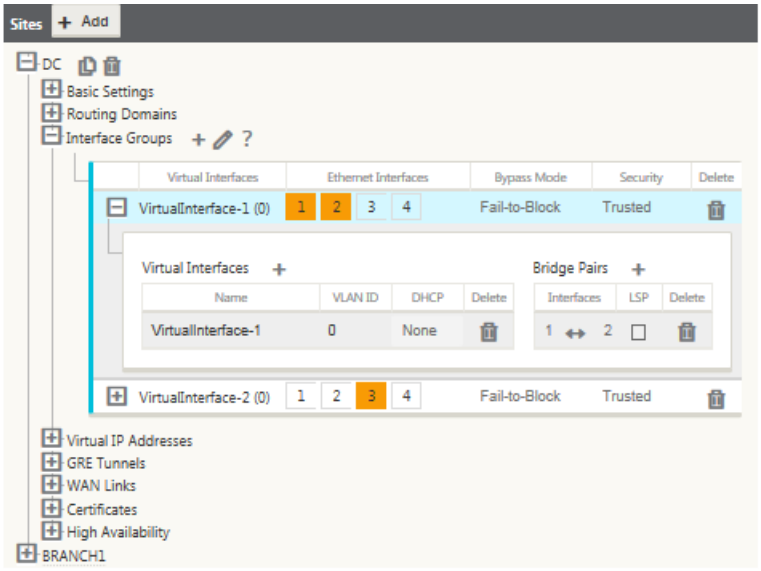

8. Click **+** to the right of **Virtual Interfaces**.

This reveals the **Name** and the **VLAN ID** lds.

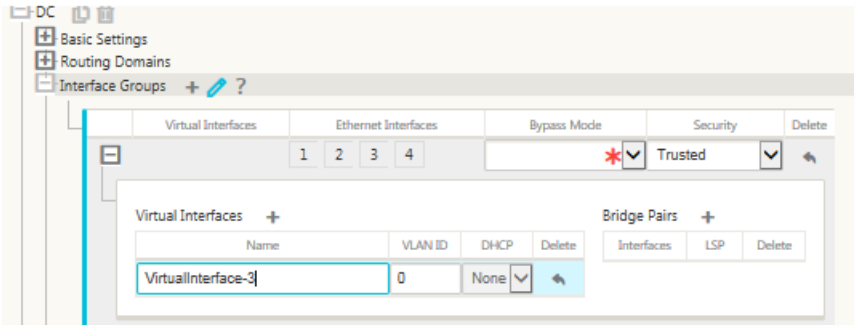

9. Enter the **Name** and **VLAN ID** for this Virtual Interface Group.

   Name – This is the name by which this Virtual Interface will be referenced.

   VLAN ID – This is the ID for identifying and marking traffic to and from the Virtual Interface. Use an ID of 0 (zero) for native/untagged traffic.

10. Click **+** to the right of **Bridge Pairs**.

    This adds a new **Bridge Pairs** entry and opens it for editing.

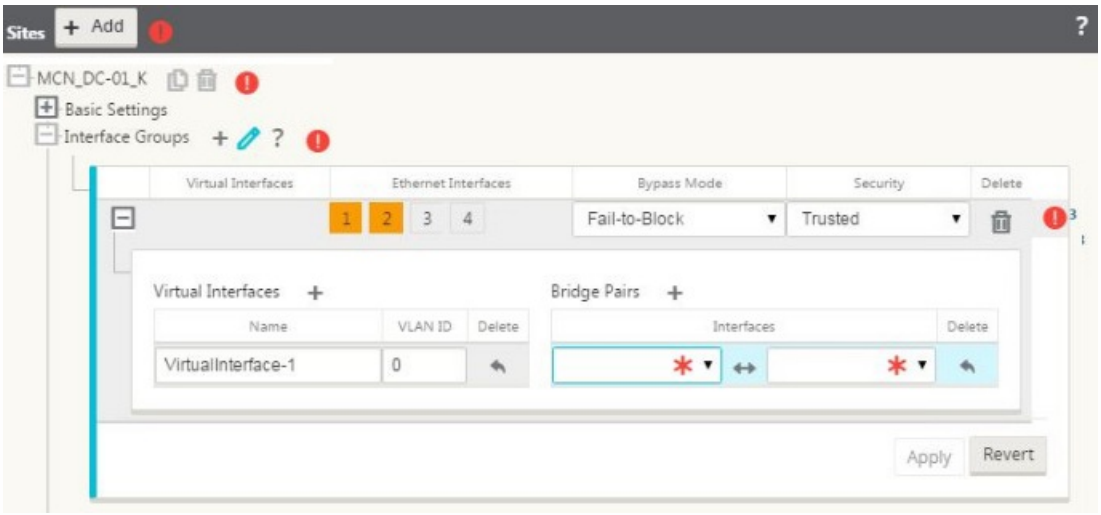

11. Select the Ethernet interfaces to be paired from the drop-down menus.

    To add more pairs, click **+** next to **Bridge Pairs** again.

12. Click **Apply**.

    This applies your settings and adds the new Virtual Interface Group to the table.

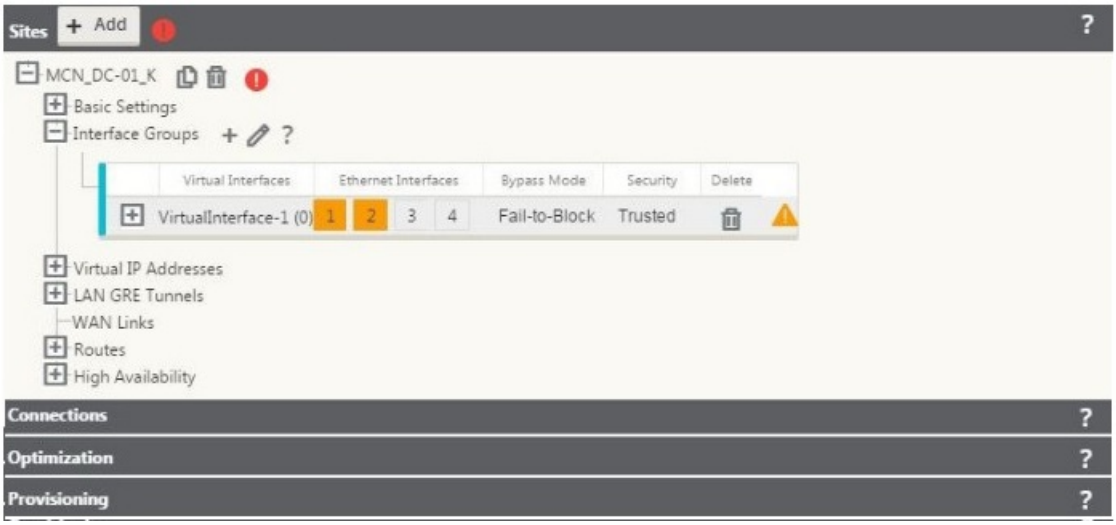

Note

At this stage, you will see a yellow delta Audit Alert icon, to the right of the new Virtual Interface Group entry. This is because you have not yet configured any Virtual IP Addresses (VIPs) for the site. For now, you can ignore this alert, as it will be resolved automatically when you have properly configured the VIPs for the site.

13. To add more Virtual Interface Groups, click **+** to the right of the **Interface Groups** branch, and proceed as above.

How to Configure Virtual IP Addresses for the MCN Site

Aug 09, 2017

The next step is to configure the Virtual IP Addresses for the site, and assign them to the appropriate group.

1. Continuing in the **Sites** tree for the new MCN site, click **+** to the left of the **Virtual IP Addresses** branch.

   This displays the **Virtual IP Addresses** table for the new site.

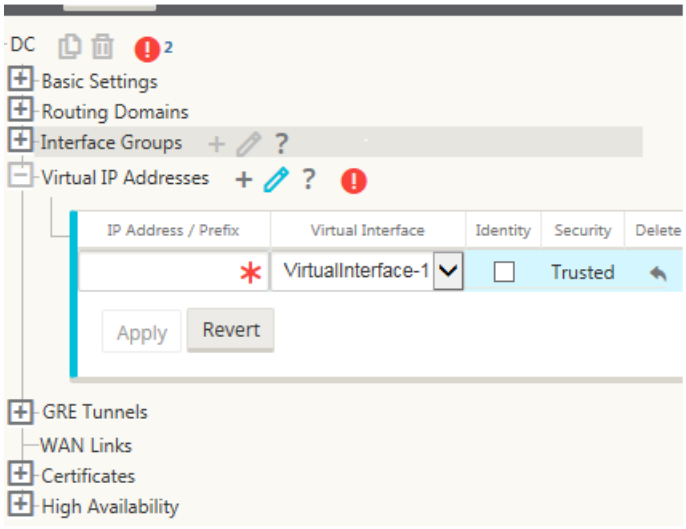

2. Click **+** to the right of **Virtual IP Addresses** to add an address.

   This opens the form for adding and configuring a new Virtual IP Address.

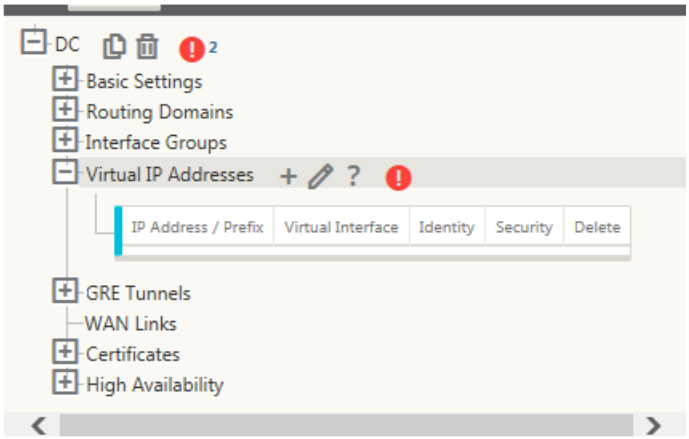

3. Enter the **Virtual IP Address** / **Prefix** information, and select the **Virtual Interface** with which the address is associated.

   The Virtual IP Address must include the full host address and netmask.

> ## Note
>
> You can click **+** again to add more Virtual IP Address entries before applying your settings.

4. Click **Apply**.

   This adds the address information to the site and includes it in the site **Virtual IP Addresses** table.

5. To add more Virtual IP Addresses, click **+** to the right of the **Virtual IP Addresses** branch, and proceed as above.

How to Configure GRE Tunnels for the MCN Site (Optional)

Aug 09, 2017

The SD-WAN GRE Tunnels settings enable you to configure SD-WAN Appliances to terminate GRE tunnels on the LAN. If you do not want to configure this site as a GRE Tunnel termination node, you can skip this step, and proceed to the section, Configuring the WAN Links for the MCN Site.

To configure a GRE Tunnel, do the following:

1. Continuing in the site tree for the new MCN site, click **+** to the left of the **GRE Tunnels** branch label.

   This opens the **GRE Tunnels** table for the new site.

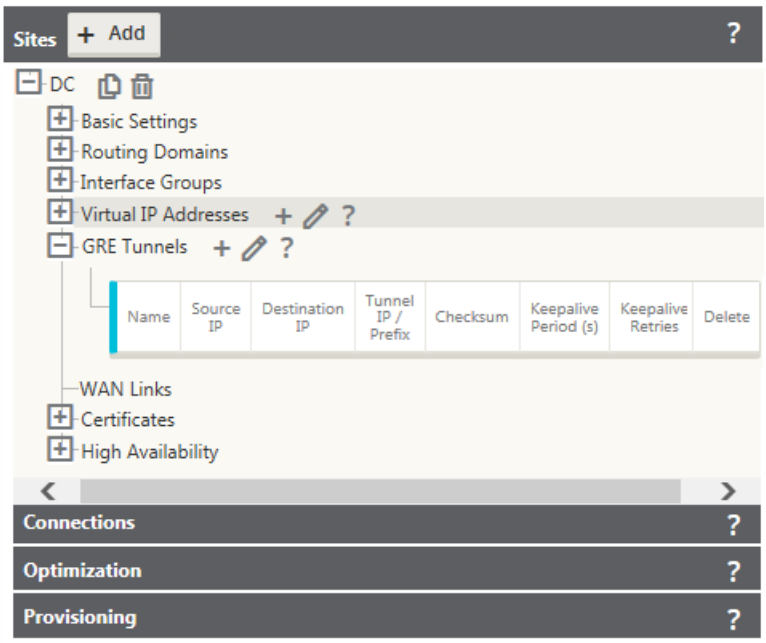

2. Click **+** to the right of the **GRE** Tunnels.

   This adds a new blank GRE Tunnel entry to the table and opens it for editing.

3. Configure the GRE Tunnel settings.

   Enter the following:

   - **Name** – Enter a name for the new GRE tunnel, or accept the default. The default uses the following naming format:

     Appliance-Tunnel-<number>

Where <*number*> is the number of GRE Tunnels configured for this site, incremented by one.

- **Source IP –** Select a source IP Address for the tunnel from the drop-down menu for this field. The menu options will be the list of Virtual Interfaces configured for this site. You must configure at least one Virtual Interface before you can configure a GRE Tunnel. For instructions, see Configuring the Virtual Interface Groups for the MCN Site and Configuring the Virtual IP Addresses for the MCN Site.

- **Destination IP** – Enter the destination IP Address for the tunnel.

- **Tunnel IP / Prefix –** Enter the tunnel IP Address and prefix.

- **Checksum** – Select this to enable Checksum for the tunnel GRE header.

- **Keepalive Period(s)** – Enter the wait time interval (in seconds) between keepalive messages. If configured to 0, no keepalive packets will be sent, but the tunnel will remain up. The default is 10.

- **Keepalive Retries –** Enter the number of keepalive retries the Virtual WAN Appliance should attempt before it brings down the tunnel. The default is 3.

4. Click **Apply**.

   This submits your settings and adds the new GRE Tunnel to the table.

5. To configure additional GRE Tunnels, click **+** to the right of the **GRE Tunnels** branch label, and proceed as above.

   The next step is to configure the WAN links for the MCN site.

How to Configure WAN Links for the MCN Site

Mar 14, 2018

The next step is to configure the WAN links for the site.

1. Continuing in the site tree for the new MCN site, click the WAN Links branch label.

> ## Note
>
> At this point in a new configuration, there are no WAN links to form a table, and therefore no Open (**+**) icon to the left of the WAN Links branch. However, if links exist, the **+** active icon is available. If so, click **+** to the left of the **WAN Links** branch to display the table. This also reveals the Add (**+**), Edit (pencil), Delete (trashcan), and Help (?) active icons to the right of the **WAN Links** branch.

This reveals the Add (**+**) and Help (**?**) active icons to the right of the **WAN Links** label.

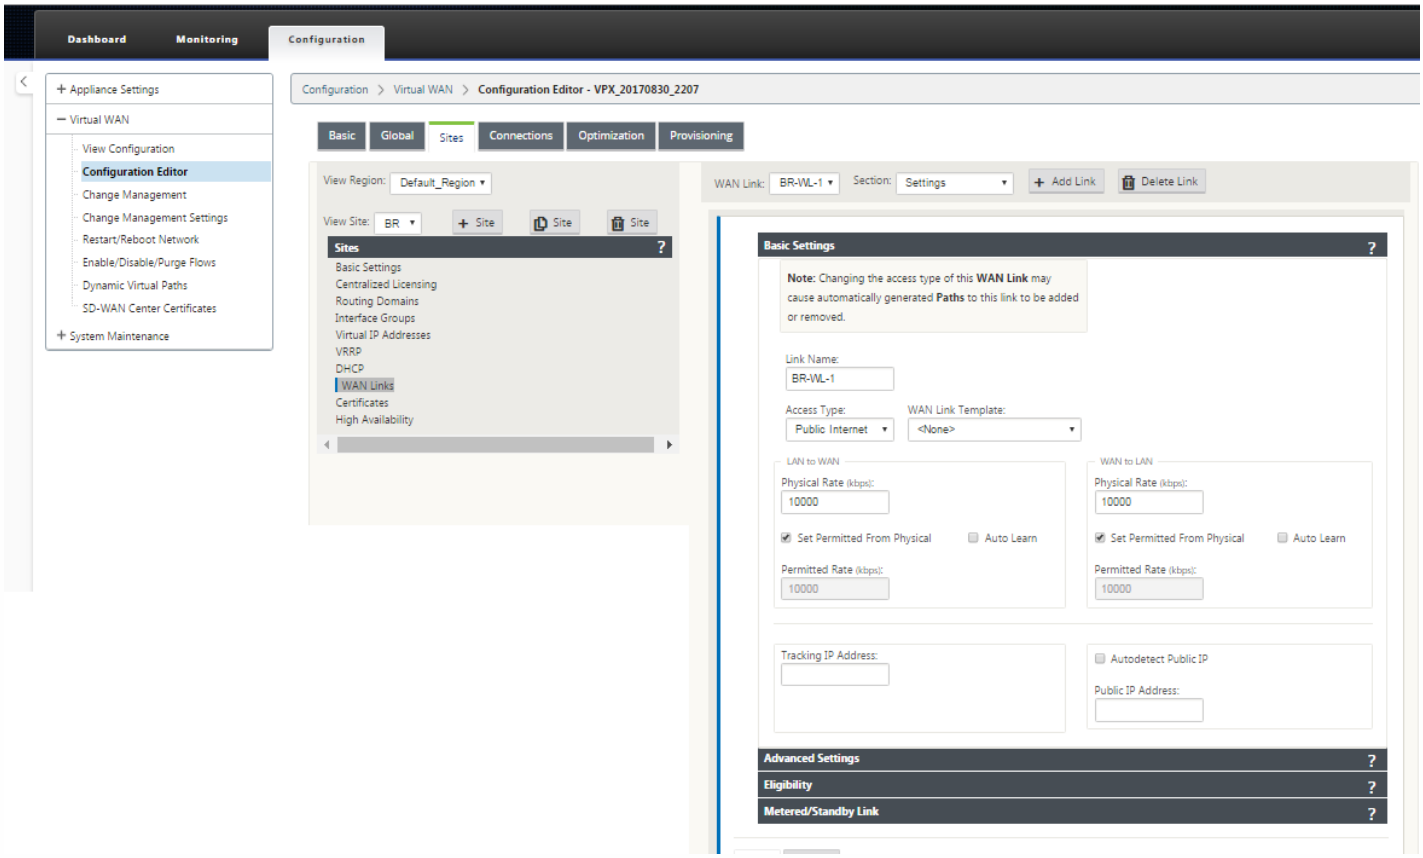

2. Click **+** to the right of the **WAN Links** branch to add a new WAN link.

   This opens the **Add WAN Link** dialog box.

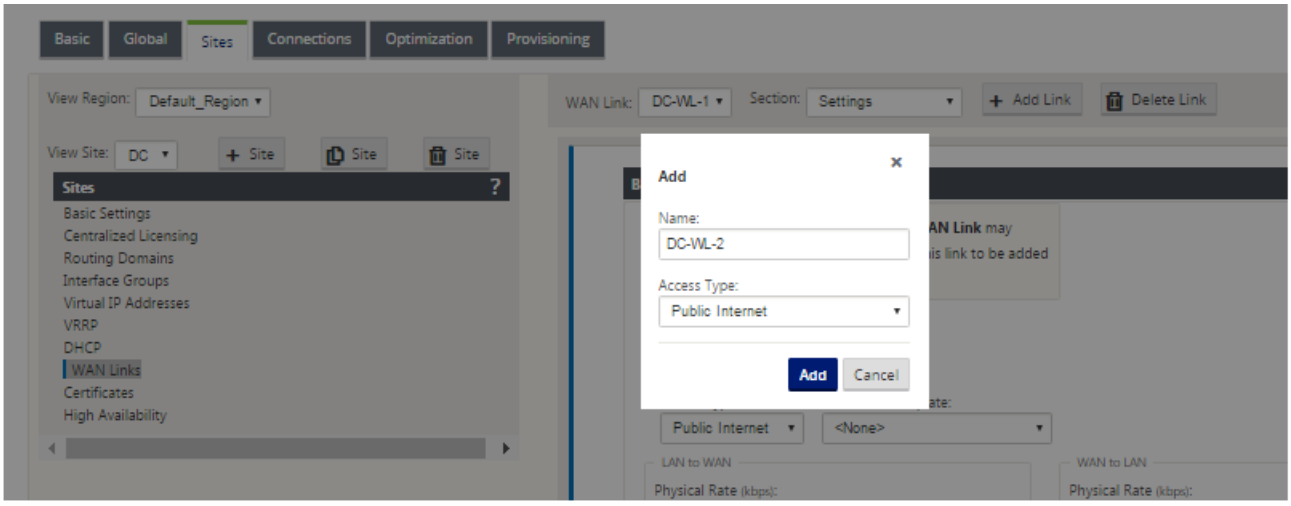

3. (Optional) Enter a name for the WAN Link if you do not want to use the default.

   The default is the site name, appended with the following suffix:

   -WL-<number>

   Where <number> is the number of WAN Links for this site, incremented by one.

4. Select the **Access Type** from the drop-down menu.

   The options are **Public Internet**, **Private Intranet**, or **Private MPLS.**

5. Click **Add**.

This displays the **WAN Links** table, adds the new unconfigured link to the table, and opens the **Basic Settings** configuration form for the link.

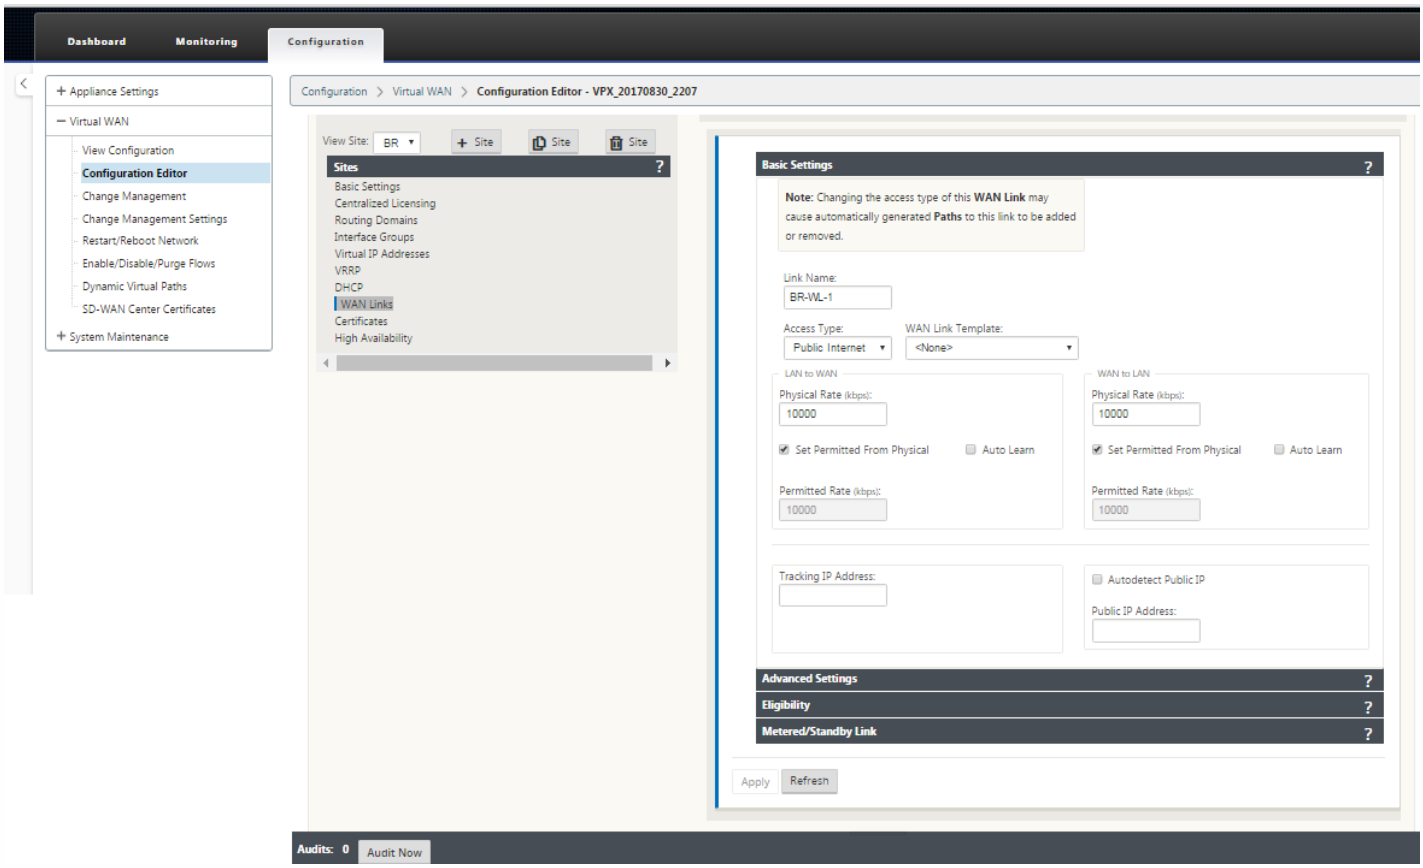

Auto Learn of bandwidth consumption

Auto learn runs on system startup and repeats every 5 minutes until a successful result is observed. Auto learn also runs after any WAN link configuration changes are made from the config editor.

You can execute tests manually or schedule tests in the SD-WAN GUI. Results from these tests should also apply to the permitted rate when the test is successful and auto learn is enabled.

When using auto learn on large networks, if config change restarts then all sites run tests simultaneously on the MCN, causing high bandwidth usage leading to inaccurate results. It is recommended that you sechdule bandwidth tests once or twice a day, typically when traffic volume is low.

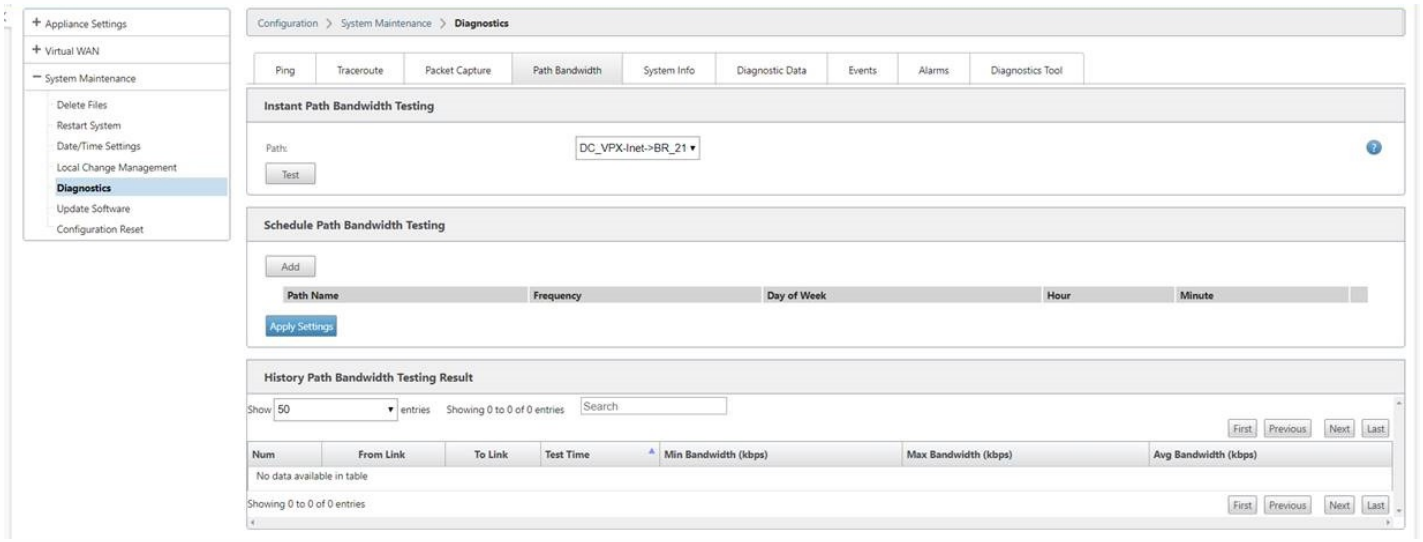

6. Enter the link details for the new WAN link.

Some guidelines are as follows:

- Some Internet links might be asymmetrical.

- Misconfiguring the permitted speed can adversely affect performance for that link.

- Avoid using burst speeds that surpass the Committed Rate.

- For Internet WAN links, be sure to add the Public IP Address.

7. Click the grey **Advanced Settings** section bar. This opens the **Advanced Settings** form for the link.

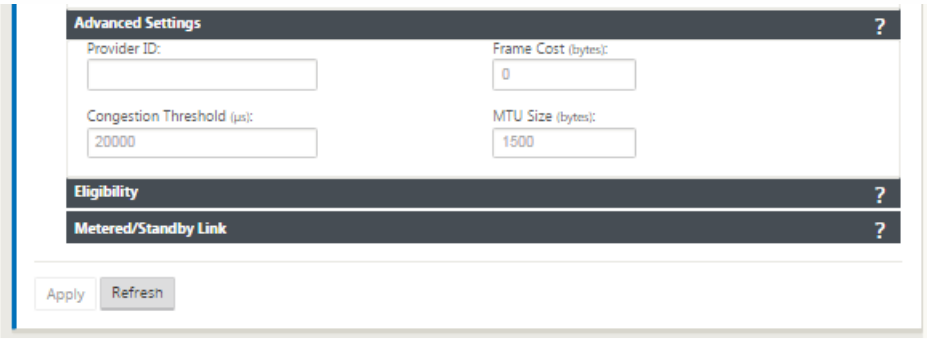

8. Enter the **Advanced Settings** for the link.

Enter the following:

- **Provider ID** – (Optional) Enter a unique ID number from 1-100 to designate WAN Links connected to the same service provider. Virtual WAN uses the Provider ID to differentiate paths when sending duplicate packets.

- **Frame Cost (bytes)** – Enter the size (in bytes) of the header/trailer added to each packet; for example, the size in bytes of added Ethernet IPG or AAL5 trailers.

- **Congestion Threshold** – Enter the congestion threshold (in microseconds) after which the WAN link will throttle packet transmission to avoid further congestion.

- **MTU Size (bytes)** – Enter the largest raw packet size (in bytes), not including the Frame Cost.

9. Click the grey **Eligibility** section bar. This opens the **Eligibility** settings form for the link.

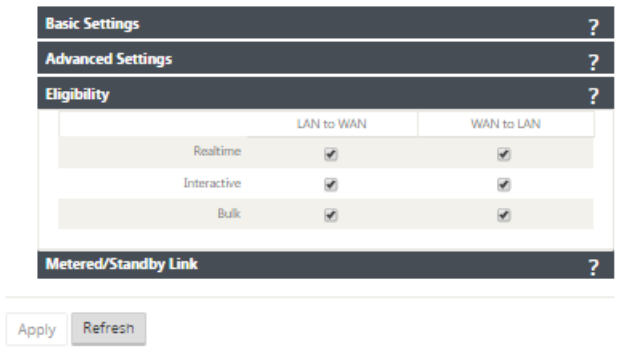

10. Select the **Eligibility** settings for the link.

11. Click the grey **Metered Link** section bar. This opens the **Metered Link** settings form for the link.

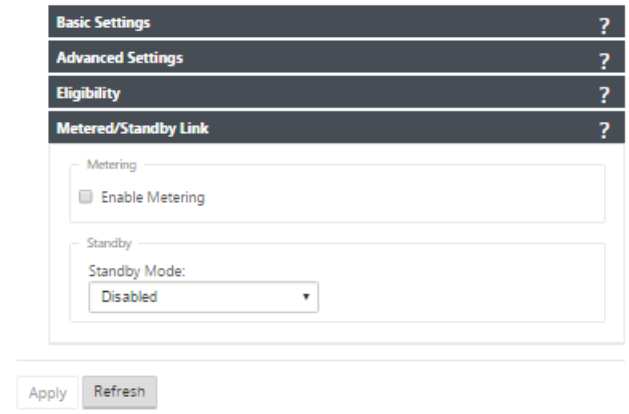

12. (Optional) Select **Enable Metering** to enable metering for this link. This displays the **Enable Metering** settings fields.

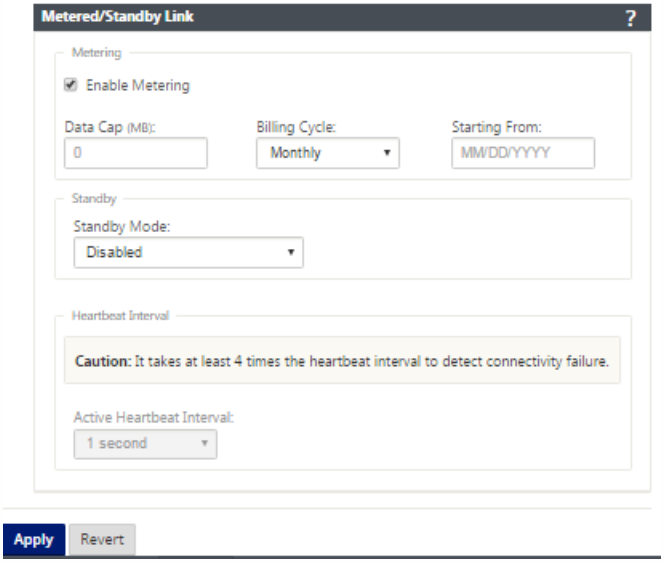

13. Configure the metering settings for the link.

    Enter the following:

    - **Data Cap (MB)** – Enter the data cap allocation for the link, in megabytes.

    - **Billing Cycle** – Select either Monthly or Weekly from the drop-down menu.

    - **Starting From** – Enter the start date of the billing cycle.

    - **Set Last Resort** – Select this to enable this link as a link of last resort in the event of a failure of all other available links. Under normal WAN conditions, Virtual WAN sends only minimal traffic over metered links, for the purpose of checking link status. However, in the event of a failure, SD-WAN can use active metered links as a last resort for forwarding production traffic.

14. Click **Apply**. This applies your specified settings to the new WAN link.

The next step is to configure the Access Interfaces for the new WAN link. An Access Interface consists of a Virtual Interface, WAN endpoint IP Address, Gateway IP Address, and Virtual Path Mode defined collectively as an interface for a specific WAN link. Each WAN link must have at least one Access Interface.

15. Click **Access Interfaces**. This opens the **Access Interfaces** table for the site.

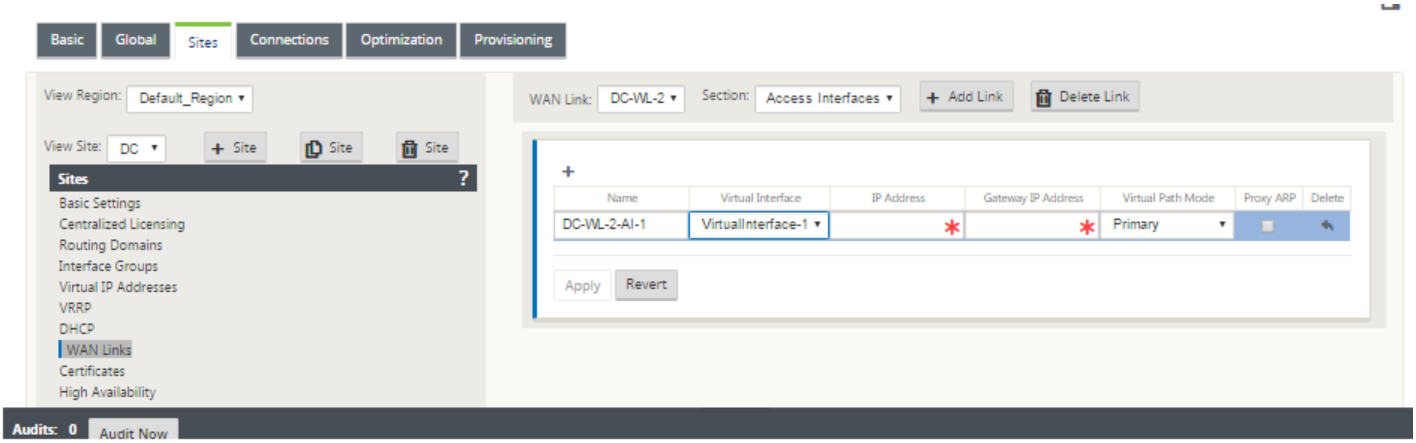

16. Enter the **Access Interfaces** settings for the link.

Note

Each WAN link must have at least one Access Interface.

Enter the following:

- **Name** – This is the name by which this Access Interface will be referenced. Enter a name for the new Access Interface, or accept the default. The default uses the following naming convention:

WAN_link_name-AI-number

Where *WAN_link_name* is the name of the WAN link you are associating with this interface, and number is the number of Access Interfaces currently configured for this link, incremented by 1.

Note

If the name appears truncated, you can place your cursor in the field, then click and hold and roll your mouse right or left to see the truncated portion.

- **Virtual Interface** – This is the Virtual Interface this Access Interface will use. Select an entry from the drop-down menu of Virtual Interfaces configured for this branch site.

- **IP Address** – This is the IP Address for the Access Interface endpoint from the appliance to the WAN.

- **Gateway IP Address –** This is the IP Address for the gateway router.

- **Virtual Path Mode** – This specifies the priority for Virtual Path traffic on this WAN link. The options are: **Primary**, **Secondary**, or **Exclude**. If set to **Exclude**, this Access Interface will be used for Internet and Intranet traffic, only.

- **Proxy ARP** – Select the checkbox to enable. If enabled, the Virtual WAN Appliance replies to ARP requests for the Gateway IP Address, when the gateway is unreachable.

17. Click **Apply**. This applies your settings and adds the new Access Interface entry to the **Access Interfaces** table.

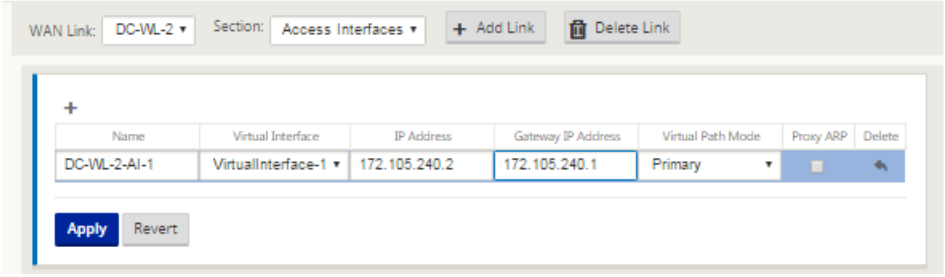

You have now finished configuring the new WAN link. Repeat these steps to add and configure additional WAN links for the site.

The next step is to add and configure the routes for the site.

How to Configure Routes for the MCN Site

Aug 09, 2017

To add and configure the routes for the site, do the following:

1. Continuing in the **Connections** tree for the new MCN site, click **+** to the left of **Routes**.

   This displays the **Routes** table for the site.

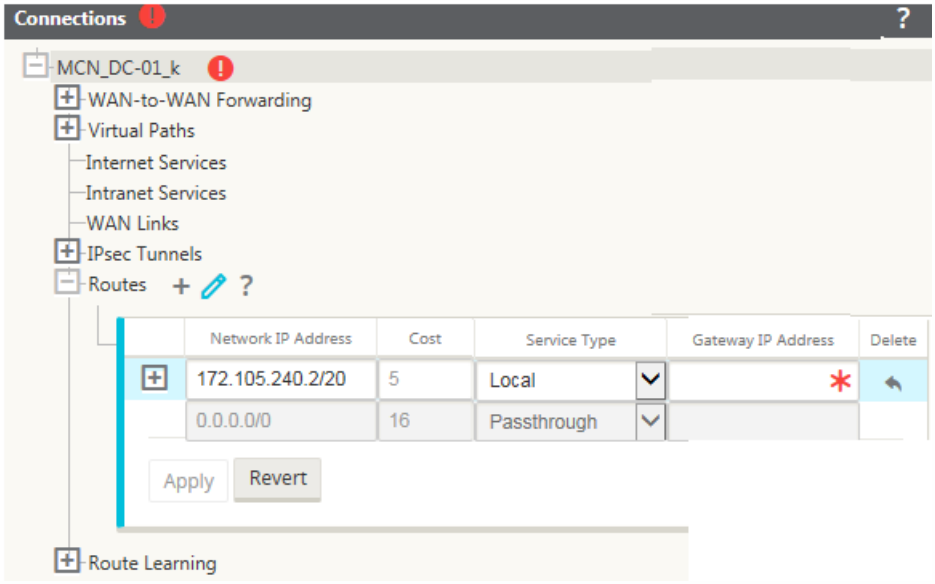

2. Click **+** to the right of the **Routes** branch to add a route.

   This opens the **Routes** table for editing and adds a blank route entry to the table (top entry).

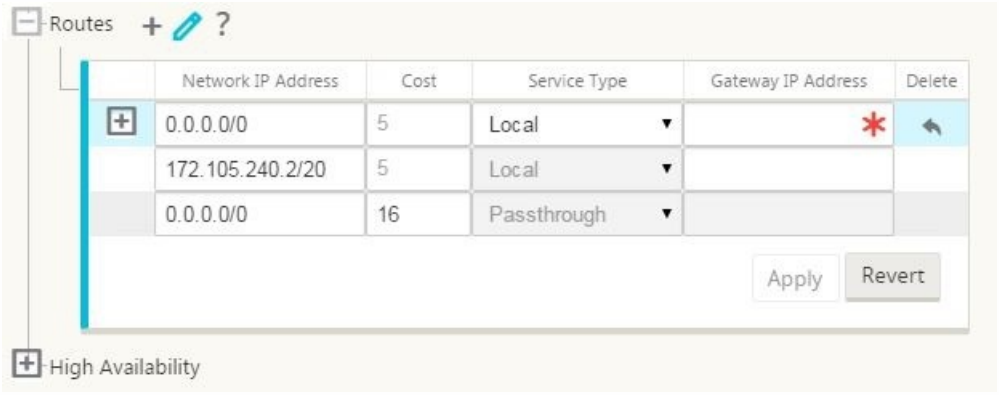

3. Enter the route configuration information for the new route.

   Enter the following:

   - **Network IP Address** – Enter the Network IP Address.

   - **Cost** – Enter a weight from 1 to 15 for determining the route priority for this route. Lower-cost routes take precedence over higher-cost routes. The default value is 5.

   - **Service Type** – Select the service type for the route from the drop-down menu for this field. The options are as follows:

* **Virtual Path** – This service manages traffic across the Virtual Paths. A Virtual Path is a logical link between two WAN links. It comprises a collection of WAN Paths combined to provide high service-level communication between two SD-WAN nodes. This is accomplished by constantly measuring and adapting to changing application demand and WAN conditions. SD-WAN Appliances measure the network on a per-path basis. A Virtual Path can be static (always exists) or dynamic (exists only when traffic between two SD-WAN Appliances reaches a configured threshold).

* **Internet** – This service manages traffic between an Enterprise site and sites on the public Internet. Traffic of this type is not encapsulated. During times of congestion, the SD-WAN actively manages bandwidth by rate-limiting Internet traffic relative to the Virtual Path, and Intranet traffic according to the SD-WAN configuration established by the Administrator.

* **Intranet** – This service manages Enterprise Intranet traffic that has not been defined for transmission across a Virtual Path. As with Internet traffic, it remains unencapsulated, and the SD-WAN manages bandwidth by rate-limiting this traffic relative to other service types during times of congestion. Note that under certain conditions, and if configured for Intranet Fallback on the Virtual Path, traffic that ordinarily travels by means of a Virtual Path may instead be treated as Intranet traffic, in order to maintain network reliability.

* **Passthrough** – This service manages traffic that is to be passed through the Virtual WAN. Traffic directed to the Passthrough Service includes broadcasts, ARPs and other non-IPv4 traffic, as well as traffic on the Virtual WAN Appliance local subnet, specifically-configured subnets, or Rules applied by the Network Administrator. This traffic is not delayed, shaped or modified by the SD-WAN. Consequently, you must ensure that Passthrough traffic does not consume substantial resources on the WAN links that the SD-WAN Appliance is configured to use for other services.

* **Local** – This service manages IP traffic local to the site that matches no other service. SD-WAN ignores traffic sourced and destined to a local route.

* **GRE Tunnel** – This service manages IP traffic destined for a GRE tunnel, and matches the LAN GRE tunnel configured at the site. The GRE Tunnel feature enables you to configure SD-WAN Appliances to terminate GRE tunnels on the LAN. For a route with service type GRE Tunnel, the gateway must reside in one of the tunnel subnets of the local GRE tunnel.

- **Gateway IP Address** – Enter the Gateway IP Address for this route.

4. Click **+** to the left of the new entry.

   This opens the Eligibility Settings form for the route.

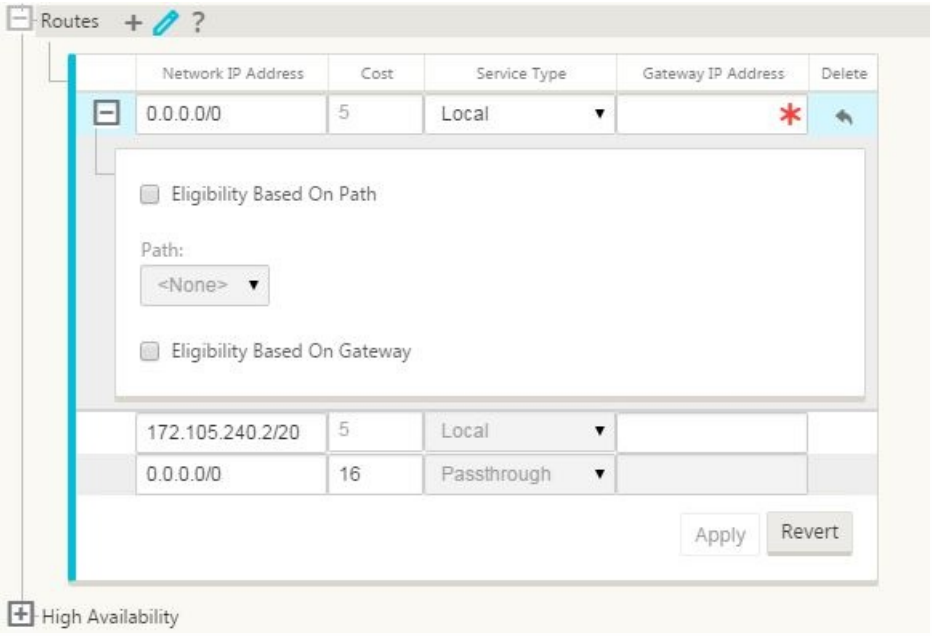

5. Enter the route **Eligibility** settings.

   The settings are as follows:

   - **Next Hop Site** – This indicates the remote site to which Virtual Path packets will be directed.

   - **Eligibility Based on Path** (checkbox) – (Optional) If enabled, the route will not receive traffic when the selected path is down.

   - **Path** – This specifies the path to be used for determining route eligibility.

9. Click **Apply**.

> ## Note
>
> After you click **Apply**, audit warnings might appear indicating that further action is required. A red dot or goldenrod delta icon indicates an error in the section where it appears. You can use these warnings to identify errors or missing configuration information. Roll your cursor over an audit warning icon to display a short description of the error(s) in that section. You can also click the dark grey **Audits** status bar (bottom of page) to display a complete list of all audit warnings.

10. To add more routes for the site, click **+** to the right of the **Routes** branch, and proceed as above.

You have now finished entering the primary configuration information for the new MCN site. The following two sections provide instructions for additional optional steps:

- Configuring High Availability (HA) for the MCN Site (Optional).
- Enabling and Configuring Virtual WAN Security and Encryption (Optional).

If you do not want to configure these features at this time, you can proceed directly to the section Naming, Saving, and Backing Up the MCN Site Configuration.

How to Enable and Configure Virtual WAN Security and Encryption (Optional)

Aug 09, 2017

To enable and configure Virtual WAN security and encryption, do the following:

> **Note**
>
> Enabling Virtual WAN security and encryption is optional.

1. At the top of the **Global** tree of the **Configuration Editor**, click **+** to the left of the **Virtual WAN Network Settings** branch.

   This opens the branch and displays the **Global Security Settings** configuration form.

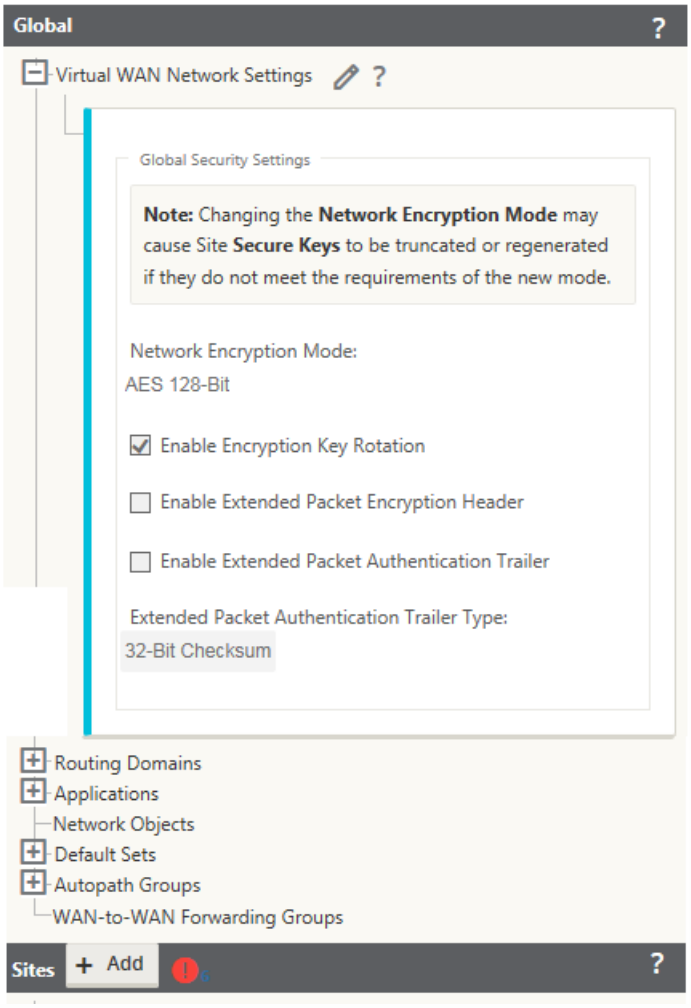

2. Click Edit (pencil icon) to enable editing for the form.

3. Enter your global security settings.

The options are as follows:

- **Network Encryption Mode** – This is the encryption algorithm used for encrypted paths. Select one of the following from the drop-down menu: **AES 128-Bit** or **AES 256-Bit**.

- **Enable Encryption Key Rotation** – When enabled, encryption keys are rotated at intervals of 10 to 15 minutes.

- **Enable Extended Packet Encryption Header** – When enabled, a 16 byte encrypted counter is prepended to encrypted traffic to serve as an initialization vector, and randomize packet encryption.

- **Enable Extended Packet Authentication Trailer** – When enabled, an authentication code is appended to the contents of the encrypted traffic to verify that the message is delivered unaltered.

- **Extended Packet Authentication Trailer Type** – This is the type of trailer used to validate packet contents. Select one of the following from the drop-down menu: **32-Bit Checksum** or **SHA-256**.

4. Click **Apply** to apply your settings to the configuration.

This completes the configuration of the MCN site. The next step is to name and save the new MCN site configuration (optional, but strongly recommended), as described in the following section.

> ## Warning
>
> If your console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you save the configuration package often, or at key points in the configuration.

Naming, Saving, and Backing Up the MCN Site Configuration

Aug 09, 2017

The next step is to name and save the new configuration, referred to as a configuration package. This step is optional at this point in the configuration, but strongly recommended. The configuration package will be saved to your workspace on the local appliance. You then have the option to log out of the Management Web Interface and continue the configuration process at a later time. However, if you log out, you will need to reopen the saved configuration when you resume. Instructions for opening a saved configuration are provided in the section Loading a Saved Configuration Package into the Configuration Editor.

Warning

If the Console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you save the configuration package often, or at key points in the configuration.

Tip

As an extra precaution, it is recommended that you use Save As, rather than Save, to avoid overwriting the wrong configuration package.

1. Click **Save As** (at the top of the **Configuration Editor** middle pane).

   This opens the **Save As** dialog box.

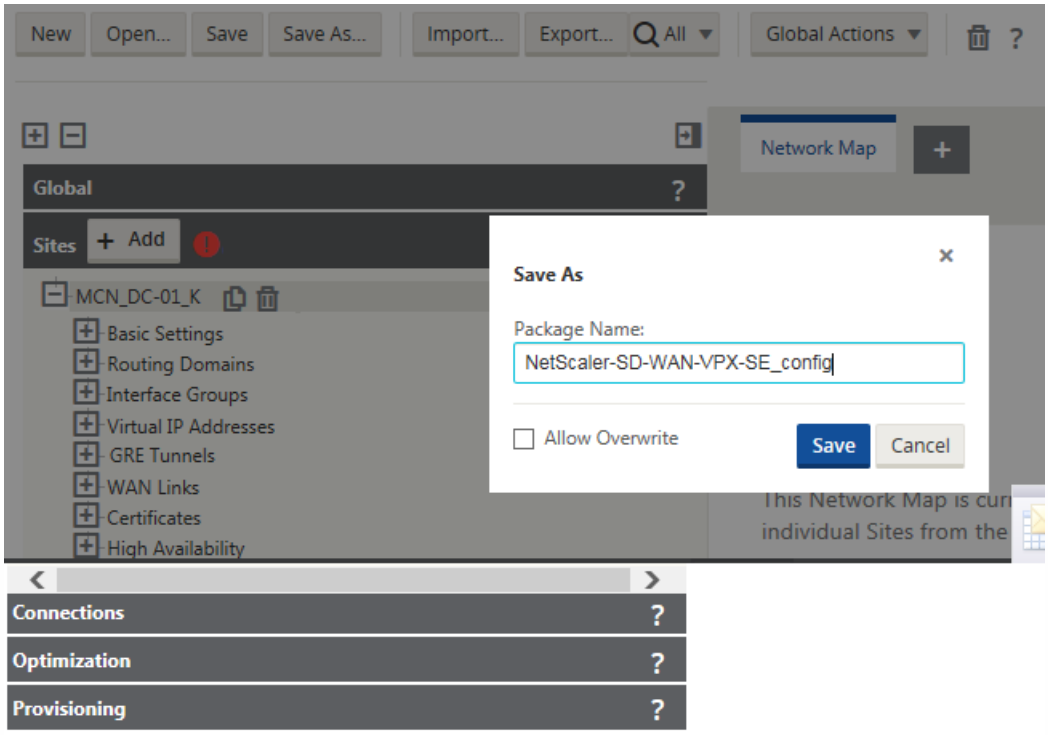

2. Enter the configuration package name.

> ## Note
>
> If you are saving the configuration to an existing configuration package, be sure to select **Allow Overwrite** before saving.

3. Click **Save**.

> ## Note
>
> After saving the configuration file, you have the option to log out of the Management Web Interface and continue the configuration process at a later time. However, if you log out, you will need to reopen the saved configuration when you resume. Instructions are provided in the section, Loading a Saved Configuration Package into the Configuration Editor.

You have now completed the MCN site configuration, and created a new SD-WAN configuration package. You are now ready to add and configure the branch sites. Instructions are provided in Adding and Configuring the Branch Sites.

How to Export Backup Copy of the Configuration Package (Optional)

Aug 09, 2017

In addition to saving the configuration-in-progress to your appliance workspace, is recommended that you also periodically back up the configuration to your local PC.

To export the current configuration package to your PC, do the following:

1. Click **Export**.

   This displays the **Export Configuration** dialog box.

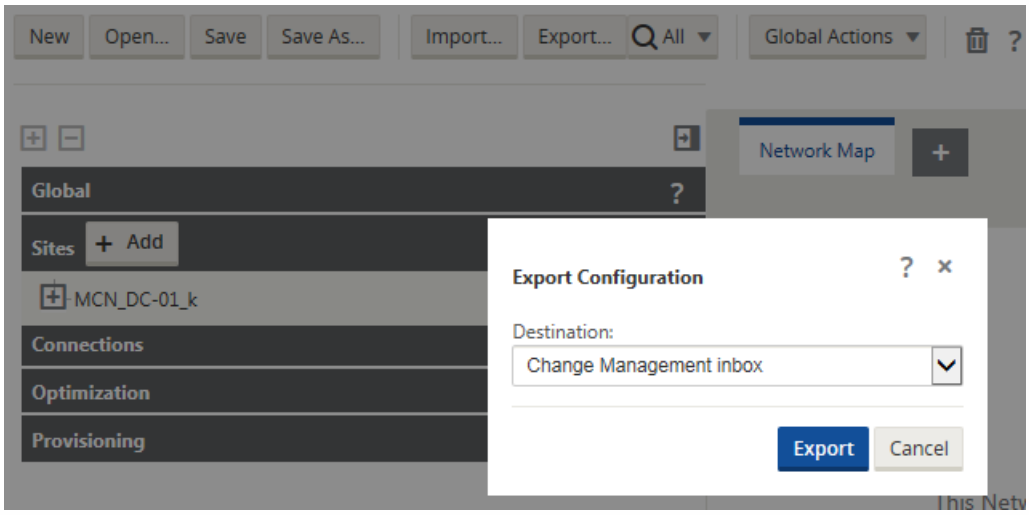

2. Select **File Download** from the **Destination:** drop-down menu.

   This reveals the **Include Network Map** option, which is selected by default.

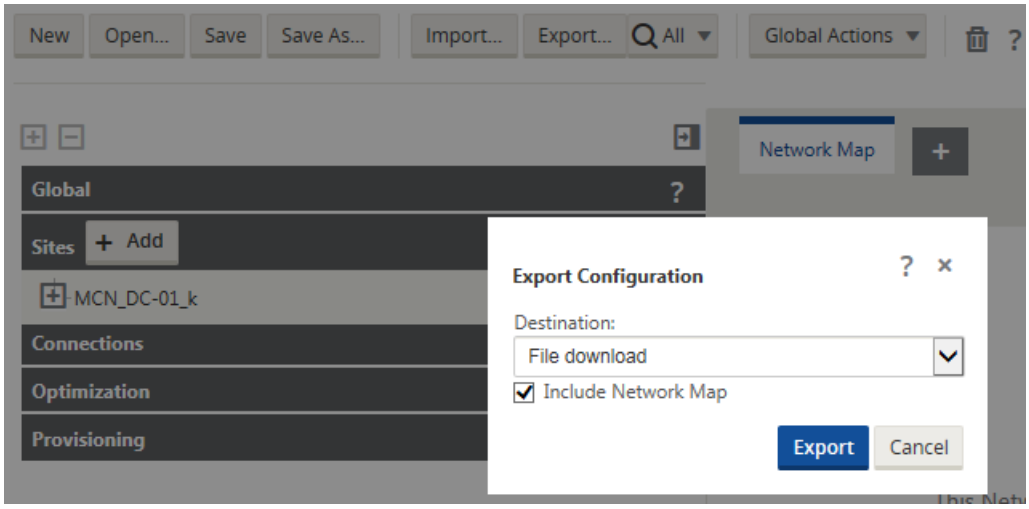

3. Accept the default, and click **Export**.

This includes the **Network Map** information in the configuration package, and opens a file browser for specifying the name and location for saving the configuration.

4. Navigate to the save location on your PC and click **Save**.

   This saves the configuration package to your PC.

> ## Note
>
> To recover a backed-up configuration package, you can use an **Import** operation to import the package from your PC and load it into the **Configuration Editor**. You can then save the imported package to your Management Web Interface workspace for future use. Instructions are provided in the section Importing a Backed up Configuration Package into the Configuration Editor.

How to Load Saved Configuration Package into the Configuration Editor

Aug 09, 2017

To resume work on a saved configuration package, you must first open the package and load it into the **Configuration Editor.**

To load a saved configuration package, do the following:

1. Log back into the Management Web Interface, and navigate to the **Configuration Editor**.

   To open the **Configuration Editor**, do the following**:**

   a. Select the **Configuration** tab at the top of the page to open the **Configuration** navigation tree (left pane).

   b. In the navigation tree, click **+** to the left of the **Virtual WAN** branch to open that branch.

   c. In the **Virtual WAN** branch, select Configuration Editor.

      This opens the **Configuration Editor** main page for a new session.

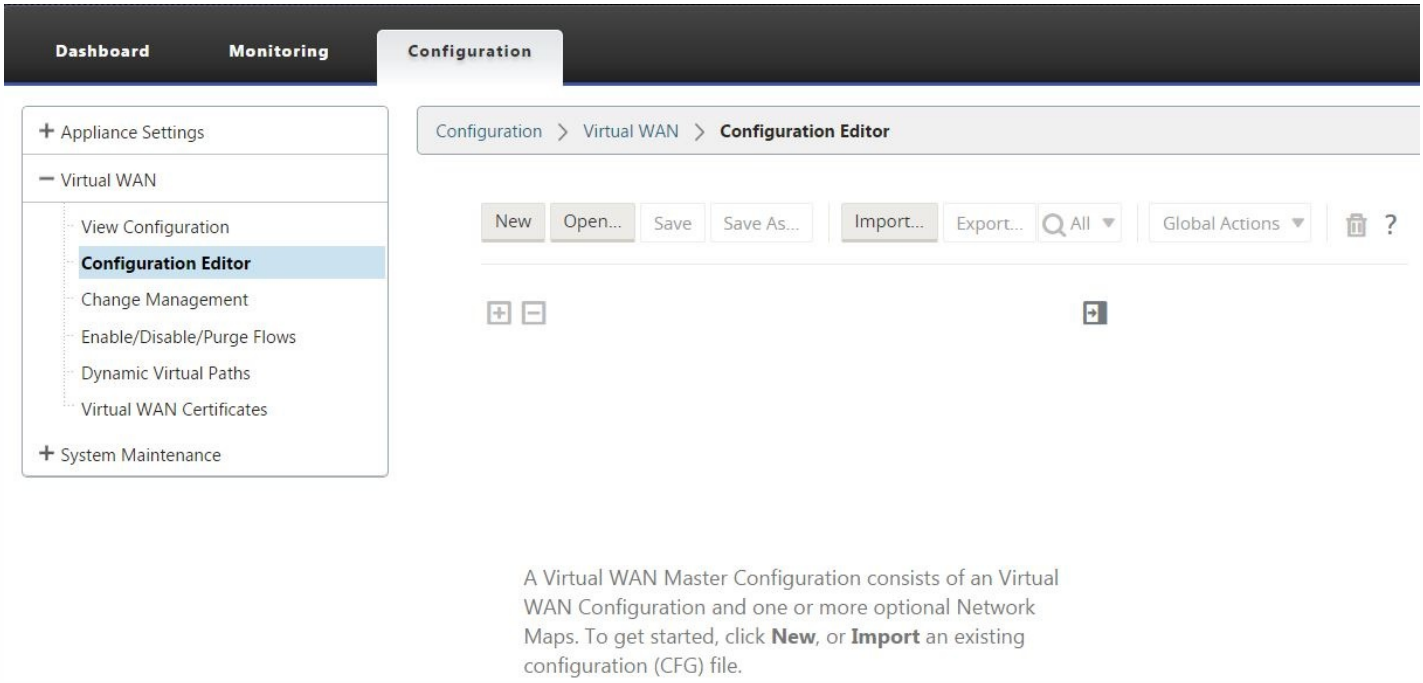

If you have just logged back into the Management Web Interface, the **Configuration Editor** initially opens for a new session, with no configuration package loaded. You have the option of starting a new configuration (**New**), opening an existing saved configuration (**Open**), or importing (**Import)** and then opening **(Open**) a configuration previously backed up to your local PC.

2. Click **Open**.

   This displays the **Open Configuration Package** dialog box.

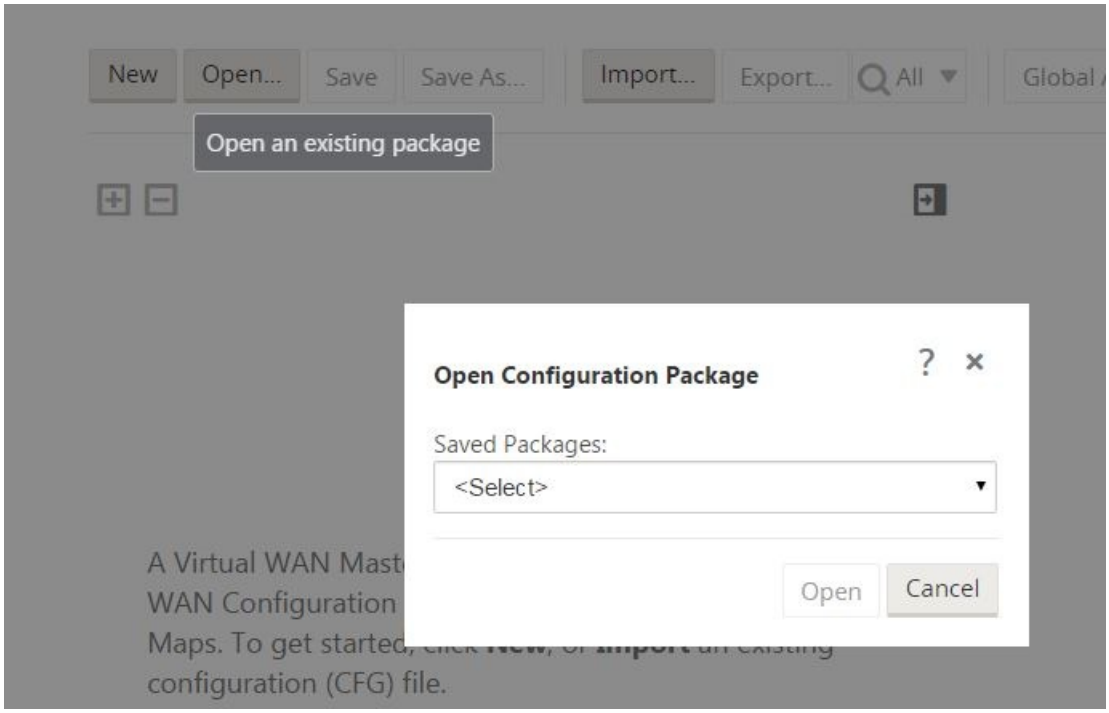

3. Select the package to open from the **Saved Packages** drop-down menu.

> ## Note
>
> If you have just opened the **Configuration Editor**, it might take a few seconds or a minute or two for the **Saved Packages** menu to be populated, depending on the number of configurations you have saved to your workspace. If so, in the interim, the **Saved Packages** menu field might display the message **No saved packages**. If this occurs, click **Cancel** to close the dialog box, wait a few moments, and click **Open** again to reopen the dialog box.

4. Click **Open**.

> ## Note
>
> This opens the specified Configuration Package and loads it into the **Configuration Editor** for editing, only. This does not stage or activate the selected configuration to the local appliance.

How to Import Backed up Configuration Package into the Configuration Editor

Aug 09, 2017

In some cases, you might want to revert to an earlier version of a Configuration Package. If you have saved a copy of the earlier version to your local PC, you can import it back into the Configuration Editor, and then open it for editing. If this is not an initial deployment, you can also import an existing Configuration Package from the global Change Management inbox on the current MCN. Instructions for both of these procedures are provided below.

To import a Configuration Package, do the following:

1. Open the **Configuration Editor**.

   To open the editor, do the following:

   a. Select the **Configuration** tab at the top of the page to open the **Configuration** navigation tree (left pane).

   b. In the navigation tree, click **+** to the left of the **Virtual WAN** branch to open that branch.

   c. In the **Virtual WAN** branch, select **Configuration Editor.**

2. In the **Configuration Editor** menu bar, click **Import.**

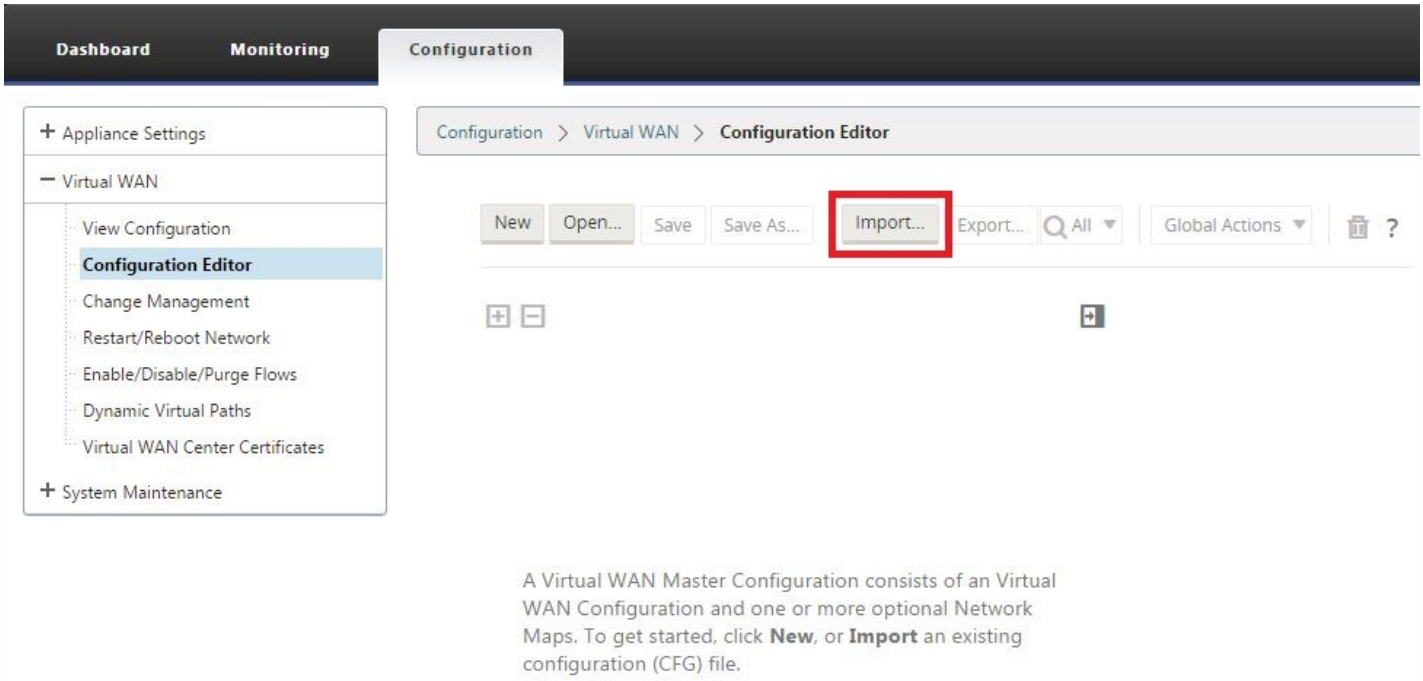

   The Import **Virtual WAN Configuration** dialog box appears.

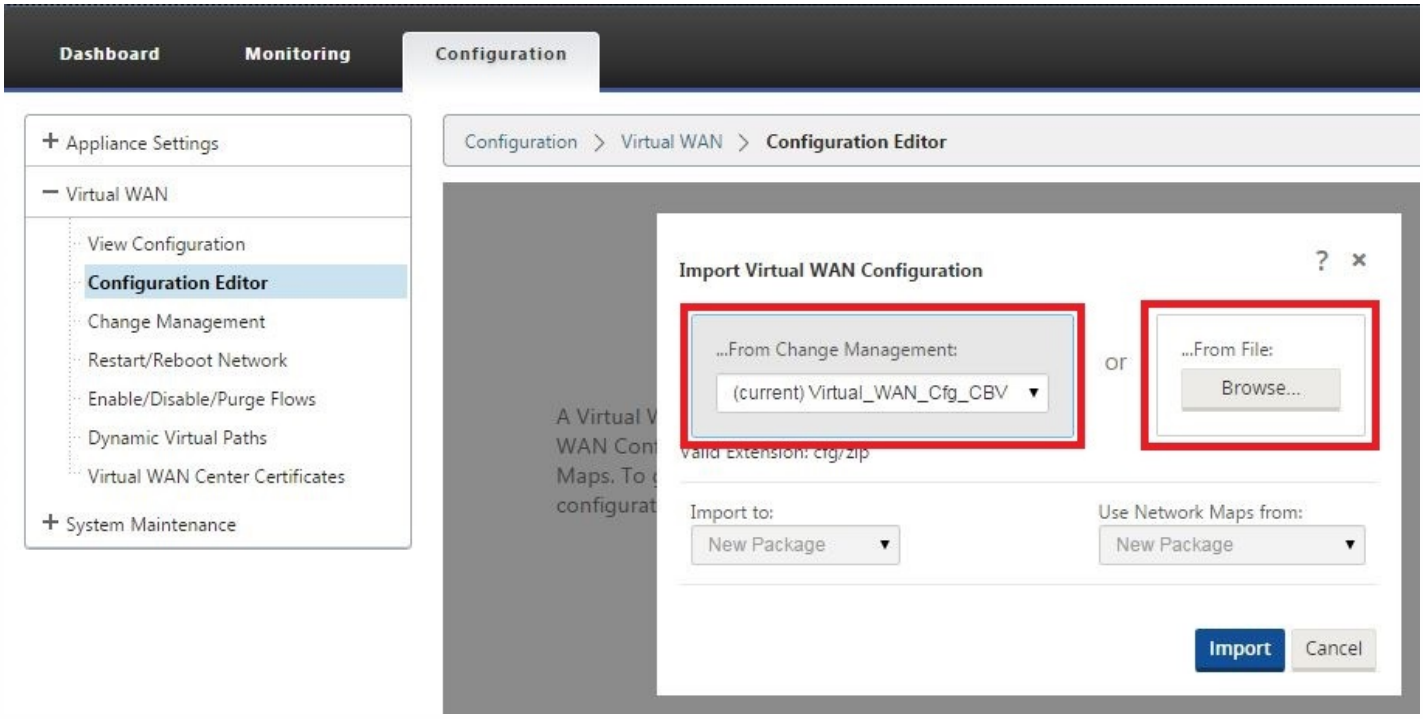

3. Select the location from which to import the package.

- **To import a Configuration Package from Change Management:** Select the package from the **From Change Management** drop down menu (top left corner).

- **To import a Configuration Package from your local PC:** Click **Browse** to open a file browser on your local PC. Select the file and click **OK**.

4. Select the import destination (if applicable).

If a Configuration Package is already open in the **Configuration Editor,** then the **Import to:** drop down menu will be available.

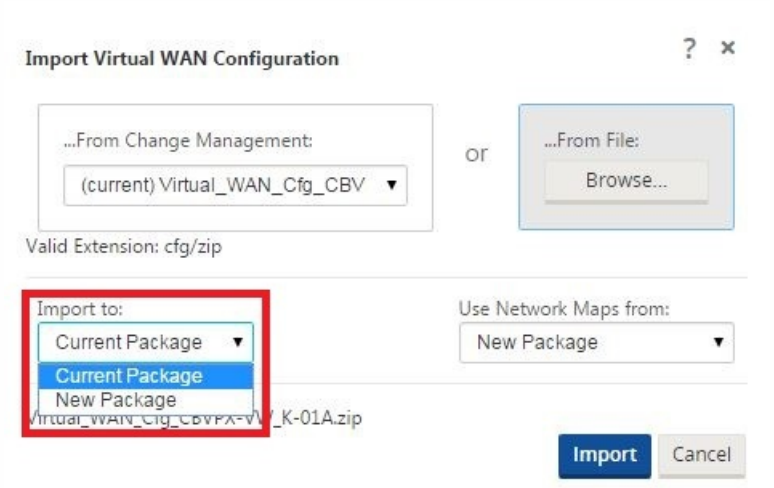

Select one of the following options:

- **Current Package –** Select this to replace the contents of the currently opened Configuration Package with the contents of the imported package, and retain the name of the opened package. However, the contents of the saved version of the current package will not be overwritten until you explicitly save the modified package. If you use **Save As** to

save the package, select **Allow Overwrite** to enable overwriting of the previous version.

- **New Package** – Select this to open a new, blank Configuration Package, and populate it with the contents of the imported package. The new package automatically takes the same name as the imported package.

5. Specify which network maps to include (if applicable).

If a Configuration Package is already open in the **Configuration Editor**, then the **Use Network Maps:** drop-down menu will be available.

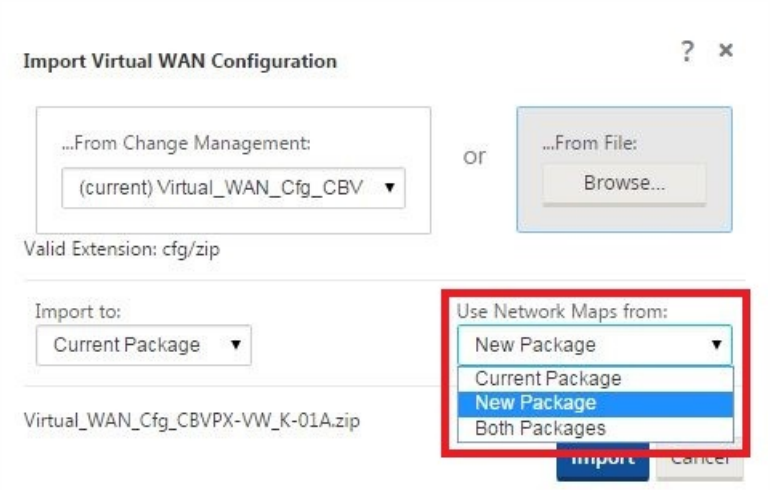

Select one of the following options:

- **Current Package** – This retains the network maps currently configured in the package currently open in the Configuration Editor, and discards any network maps from the imported package.

- **New Package** – This replaces the network maps currently configured in the currently open package with the network maps (if any) from the imported package.

- **Both Packages** – This includes all network maps from both the current and the imported package.

6. Click **Import**.

This loads the imported file into the **Configuration Editor**, according to your specifications.

> ## Note
>
> If a package of the same name already exists in your workspace, then the **Name Conflict** dialog box displays.

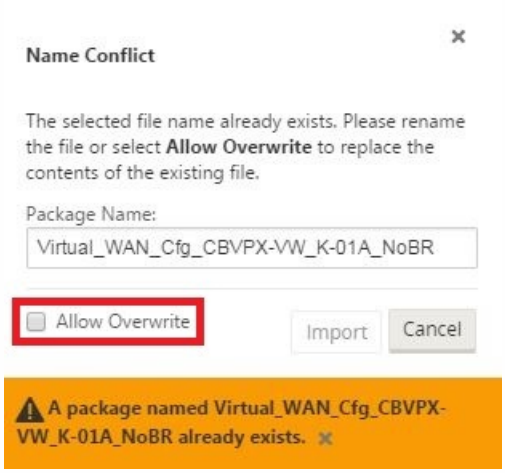

To specify the name to use for the imported package, do one of the following:

- Enter a different name in the **Package Name** field to rename the new package and enable the **Import** button. The imported package is loaded into the **Configuration Editor** with the specified name. The package name is saved to your workspace at this time, but the package contents will not be saved to your workspace until you explicitly save the package.

- Select **Allow Overwrite** to confirm that you want to retain the existing name and enable overwriting of the contents of the saved package. However, the contents of the saved version of the current package will not be overwritten until you explicitly save the modified package.

This also enables the **Import** button in the **Name Conflict** dialog box. Click **Import** to complete the import operation.

Adding and Configuring the Branch Sites

Aug 09, 2017

This chapter provides instructions for adding and configuring the branch sites. The procedure for adding a branch site is very similar to creating and configuring the MCN site. However, some of the configuration steps and settings do vary slightly for a branch site. In addition, once you have added an initial branch site, for sites that have the same appliance model you can use the **Clone** (duplicate) feature to streamline the process of adding and configuring those sites.

As with creating the MCN site, to set up a branch site you must use the **Configuration Editor** in the Management Web Interface on the MCN appliance. The **Configuration Editor** is available only when the interface is set to **MCN Console** mode.

Supplemental Branch Site Deployment Information

In addition to this guide, the following CloudBridge Knowledge Base support articles are also recommended:

- CloudBridge Virtual WAN PBR Mode Deployment Steps (CTX201577)

  http://support.citrix.com/article/CTX201577

- CloudBridge Virtual WAN Gateway Mode Deployment Steps (CTX201576)

  http://support.citrix.com/article/CTX201576

Overview of Branch Site Configuration Procedures

The steps to complete this process are as follows:

1. Add the branch site.
2. Configure the Virtual Interface Groups for the branch site.
3. Configure the Virtual IP Addresses for the branch site.
4. (Optional) Configure the LAN GRE Tunnels for the branch site.
5. Configure the WAN Links for the branch site.
6. Configure the Routes for the branch site.
7. (Optional) Configure High Availability for the branch site.
8. (Optional) Clone the new branch site to create and configure additional sites.

> ### Note
>
> Cloning the site is optional. The Virtual WAN appliance models must be the same for both the original and the cloned sites. You cannot change the specified appliance model for a clone. If the appliance model is different for a site, you must manually add the site.

9. Resolve any configuration Audit Alerts.

10. Save the completed configuration.

How to Add the Branch Site

Aug 09, 2017

To add a new branch site to the **Sites** table and begin configuring the site, do the following:

> ## Note
>
> If you logged out of the MCN after creating and saving the new configuration package, you will need to log back in and reopen the configuration before you can continue. To do so, click **Open** in the **Configuration Editor** menu bar (top of page area). This displays a dialog box for selecting the configuration you want to modify.

1. Continuing in the **Configuration Editor**, click **Add** in the **Sites** bar to begin adding and configuring the new branch site.

   This displays the **Add Site** dialog box.

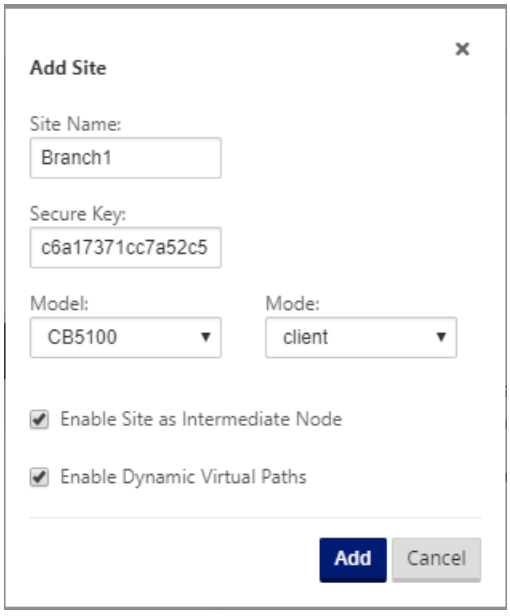

2. Enter the following site information.

> ## Note
>
> Entries cannot contain spaces and must be in Linux format.

- **Site Name** – Enter a name for the site.

- **Appliance Name** – Enter the name you want to assign to the appliance.

- **Secure Key** – This is a hexadecimal key of 8 to 32 digits used for encryption and membership verification in the SD-WAN Appliance. By default, this field is prefilled with an automatically generated security key. Accept the default or enter a custom key in hexadecimal format.

- **Model –** Select the appliance model from the drop-down menu.

- **Mode** – Select client as the mode.

- **Enable Site as Intermediate Node** – If this option is enabled, the site serves as a mediator for the creation and deletion of Dynamic virtual paths between two or more sites connected to this site.

- **Enable Dynamic Virtual Paths -** If this option is enabled, Dynamic virtual paths will be allowed between this site and other sites connected through an existing intermediate node.

3. Click **Add** to add the site.

   This adds the new site to the **Sites** tree, and opens the **Basic Settings** configuration form for the site.

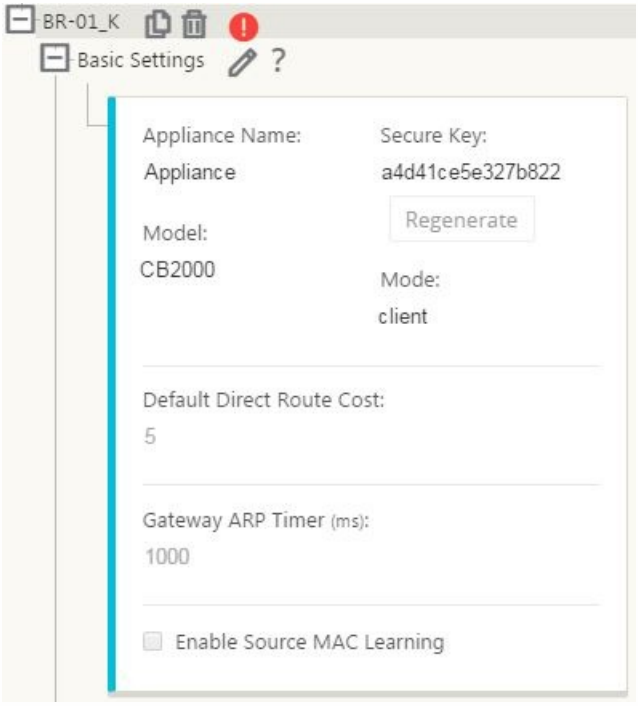

4. Click the Edit (pencil) icon to enable editing for the form.

5. Enter the basic settings for the site, and click **Apply**.

   The next step is to add and configure the Virtual Interface Groups for the new site.

How to Configure Virtual Interface Groups for the Branch Site

Aug 09, 2017

After adding the new site, the next step is to create and configure the Virtual Interface Groups for the site.

To add a Virtual Interface Group to the new site, do the following:

1. In the **Sites** navigation tree, click **+** next to the name of the site you just added.

   This opens the configuration branches for the new site.

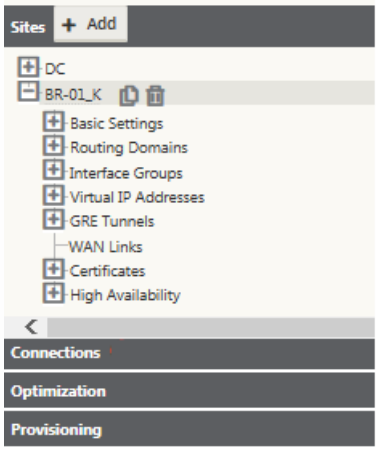

2. Click **+** to the left of the **Interface Groups** branch.

   This displays the **Interface Groups** table for the site.

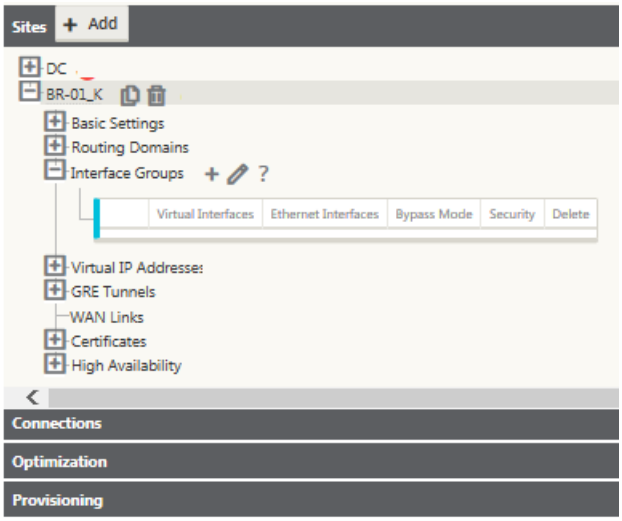

3. Click **+** to the right of **Interface Groups**.

   This adds a new blank group entry to the table and opens it for editing.

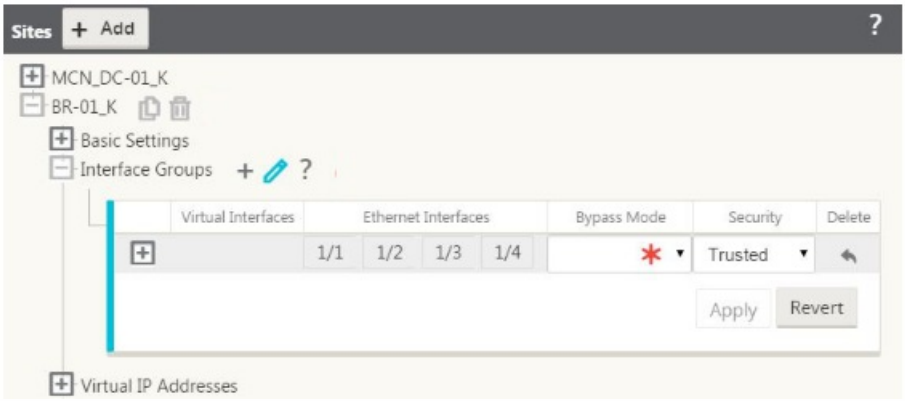

4. Select the **Ethernet Interfaces** to include.

Under **Ethernet Interfaces**, click a box to include/exclude that interface. You can select any number of interfaces to include in the group. A goldenrod highlight indicates an included interface.

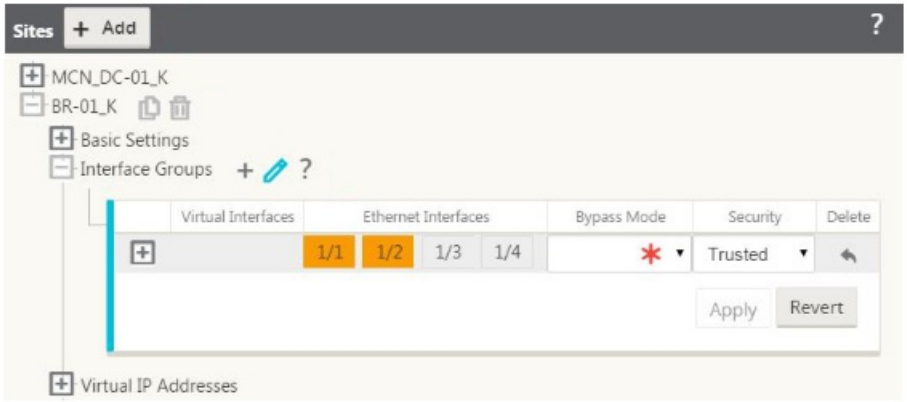

5. Select the **Bypass Mode** and **Security** level from the drop-down menus.

The **Bypass Mode** specifies the behavior of bridge-paired interfaces in the Virtual Interface Group, in the event of an appliance or service failure or restart. The options are: **Fail-to-Wire** or **Fail-to-Block**.

6. Select the **Security Level** from the drop-down menu**.**

This specifies the security level for the network segment of the Virtual Interface Group. The options are: **Trusted** or **Untrusted**. Trusted segments are generally protected by a firewall (default is Trusted).

7. Click **+** at the left edge of the new blank entry.

This reveals the **Virtual Interfaces** and **Bridge Pairs** fields.

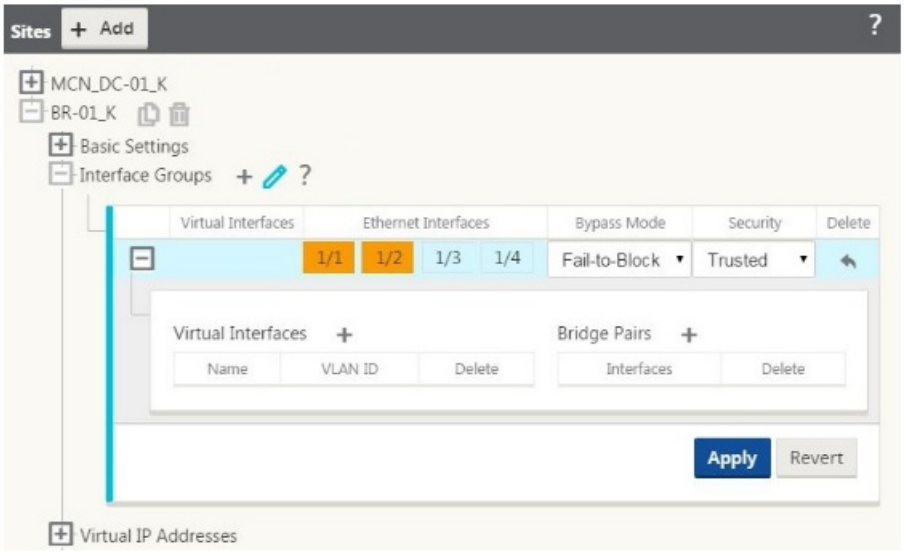

Tip

You can resize the tree pane to reveal any truncated contents. To do so, roll your cursor over the resize bar at the right edge of the tree area. When the cursor changes to a bi-directional arrow, click and drag the bar to the right or left to grow or shrink the pane width.

8. Click **+** to the right of **Virtual Interfaces**.

   This opens the **Name** and **VLAN ID** fields for editing.

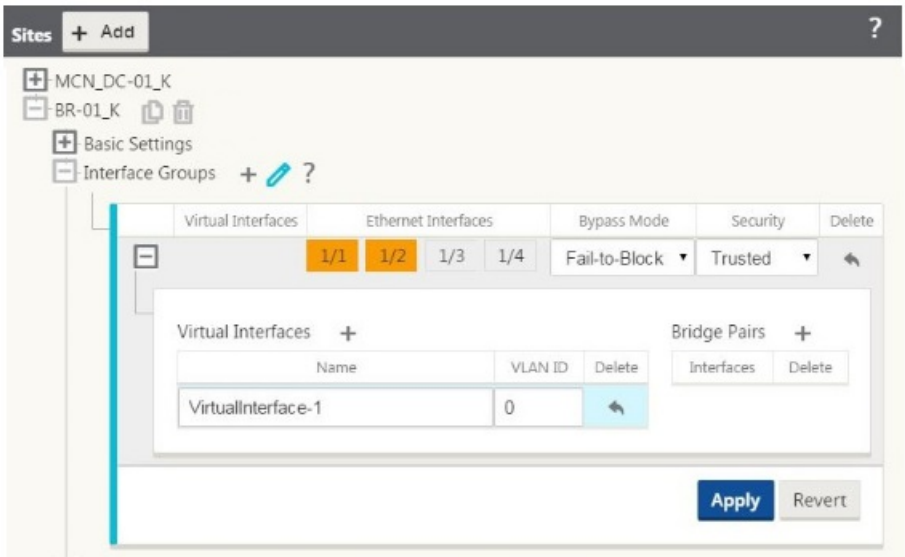

9. Enter the **Name** and **VLAN ID** for this Virtual Interface Group.
10. Click **+** to the right of **Bridge Pairs**.

    This adds a new **Bridge Pairs** entry and opens it for editing.

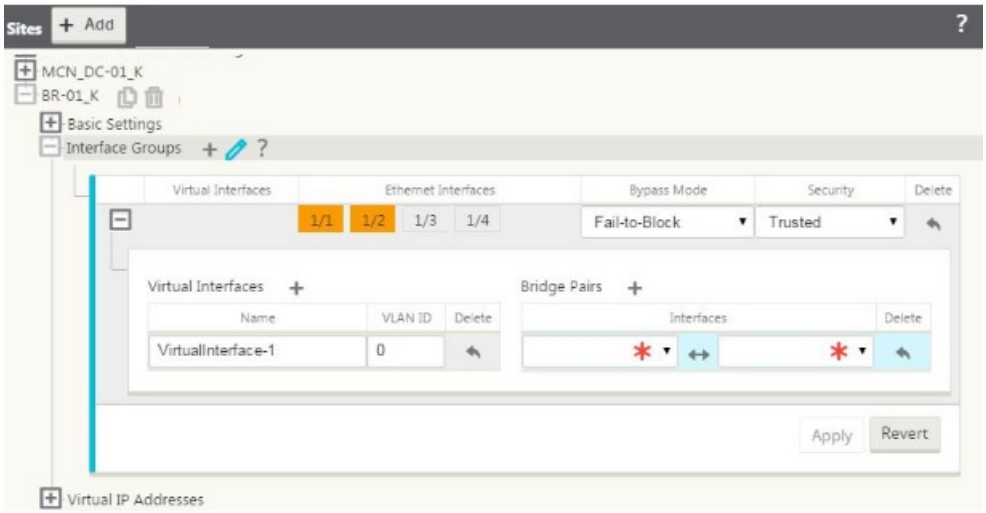

11. Select the interfaces to be paired from the drop-down menus.

    To add more pairs, click **+** next to the **Bridge Pairs** field again.

12. Click **Apply**.

    This applies your settings and adds the new Virtual Interface Group to the table.

Note

At this stage, you will see a yellow delta Audit Alert icon, to the right of the new Virtual Interface Group entry. This is because you have not yet configured any Virtual IP Addresses (VIPs) for the site. For now, you can ignore this alert, as it will be resolved automatically when you have properly configured the VIPs for the site.

13. To add more Virtual Interface groups, click **+** to the right of the **Interface Groups** branch, and proceed as above.

How to Configure Virtual IP Addresses for the Branch Site

Aug 09, 2017

The next step is to configure the Virtual IP Addresses for the site, and assign them to the appropriate group.

1. Continuing in the site tree for the new site, click **+** to the left of the **Virtual IP Addresses** branch.

   This displays the **Virtual IP Addresses** table for the new site.

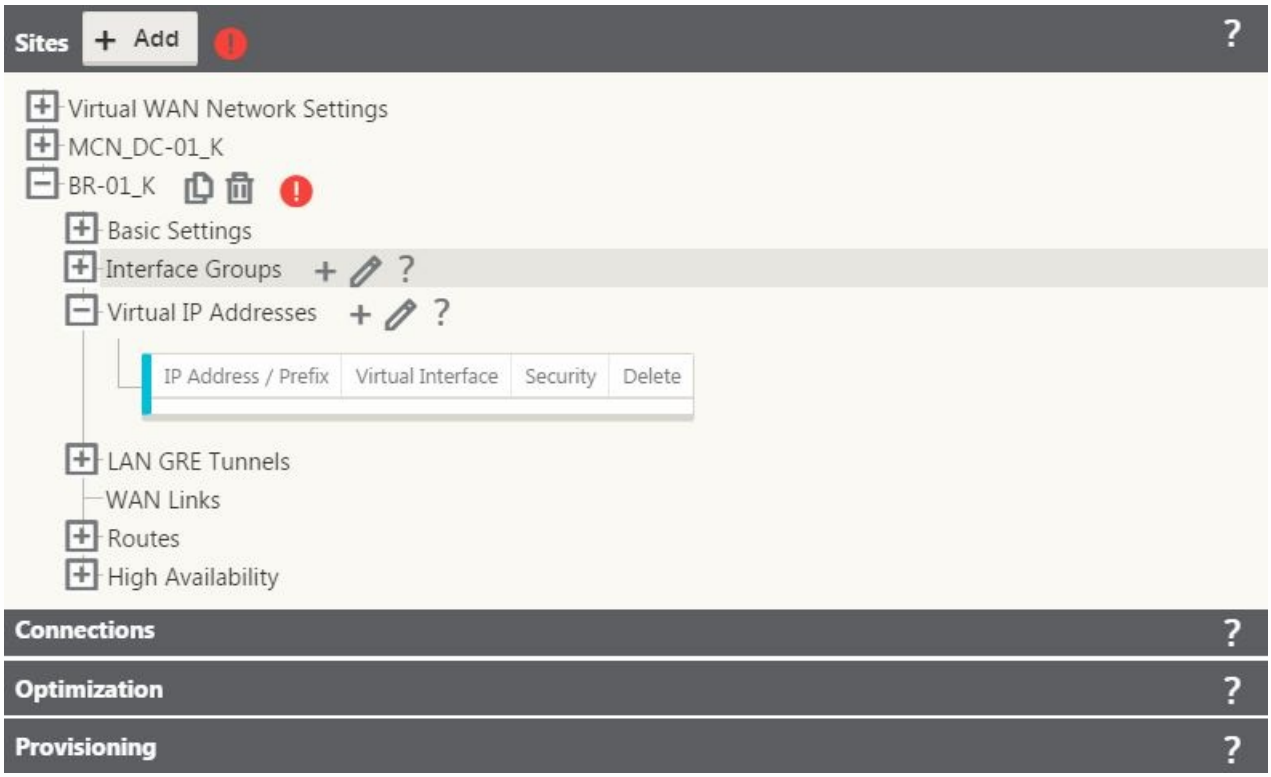

2. Click **+** to the right of the **Virtual IP Addresses** branch to add an address.

   This opens the form for adding and configuring a new Virtual IP Address.

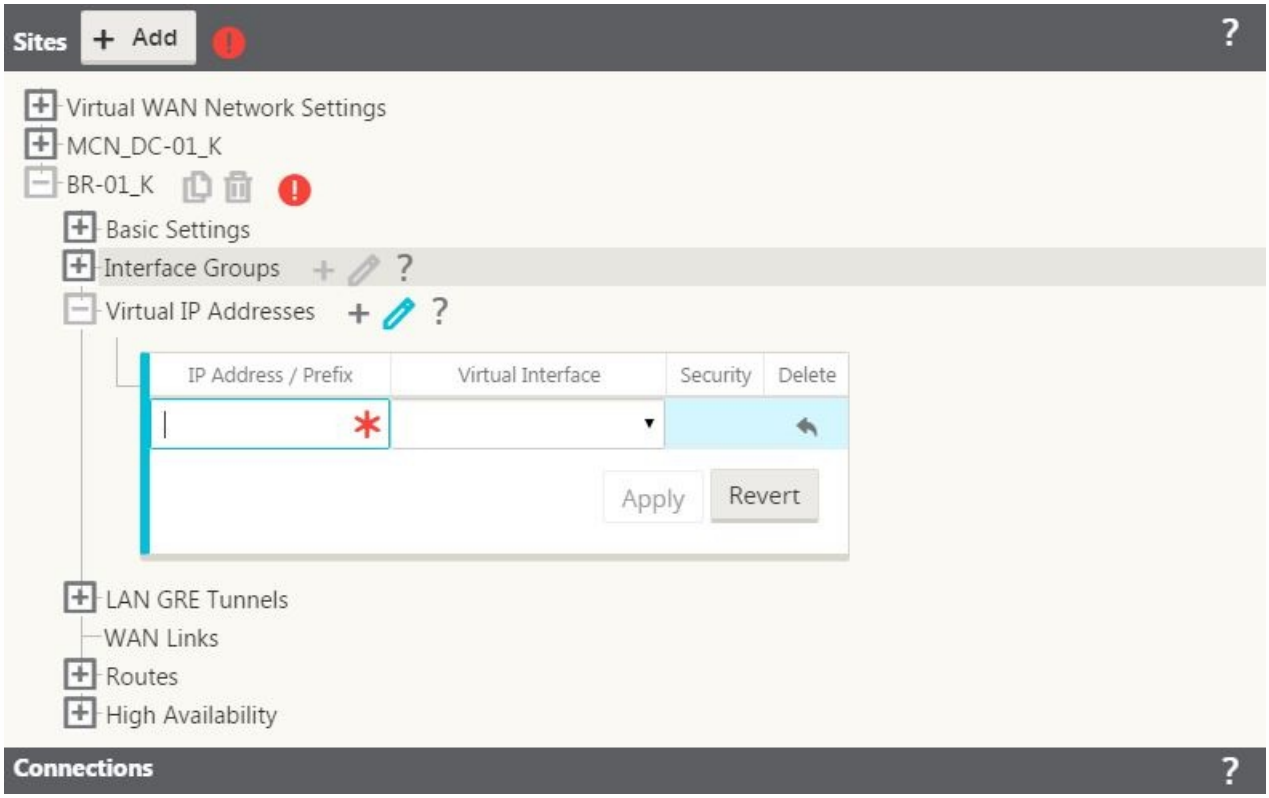

3. Enter the **Virtual IP Address / Prefix** information, and select the **Virtual Interface** (Virtual Interface Group) with which the address is associated.

   The Virtual IP Address must include the full host address and netmask.

> ## Note
>
> You can click **+** again to add more Virtual IP Address entries before applying your settings.

4. Click **Apply**.

   This adds the address information to the site configuration and includes it in the site **Virtual IP Addresses** table.

5. To add more Virtual IP Addresses, click **+** to the right of the **Virtual IP Addresses** branch, and proceed as above.

How to Configure GRE Tunnels for the Branch Site

Aug 09, 2017

The Virtual WAN LAN GRE Tunnels settings enable you to configure Virtual WAN Appliances to terminate GRE tunnels on the LAN. If you do not want to configure this branch site as a LAN GRE Tunnel termination node, you can skip this step, and proceed to the section, Configuring the WAN Links for the Branch Site.

To configure a LAN GRE Tunnel for the branch site, do the following:

1. Continuing in the site tree for the new branch site, click **+** to the left of the **LAN GRE Tunnels** branch label.

   This opens the **LAN GRE Tunnels** table for the new site.

2. Click **+** to the right of the **LAN GRE Tunnels**.

   This adds a new blank LAN GRE Tunnel entry to the table and opens it for editing.

3. Configure the LAN GRE Tunnel settings.

Enter the following:

- **Name** – Enter a name for the new LAN GRE tunnel, or accept the default. The default uses the following naming format:

Appliance-Tunnel-<number>

Where <*number*> is the number of LAN GRE Tunnels configured for this site, incremented by one.

- **Source IP** – Select a Source IP Address for the tunnel from the drop-down menu for this field. The menu options will be the list of Virtual IP Addresses that you configured for this site. You must configure at least one Virtual Interface and one Virtual IP Address before you can configure a LAN GRE Tunnel. For instructions, see the sections, Configuring the Virtual Interface Groups for the Branch Site and Configuring the Virtual IP Addresses for the Branch Site.

- **Destination IP** – Enter the destination IP Address for the tunnel.

- **Tunnel IP / Prefix** – Enter the tunnel IP Address and prefix.

- **Checksum** – Select this to enable Checksum for the tunnel GRE header.

- **Keepalive Period(s)** – Enter the wait time interval (in seconds) between keepalive messages. If configured to 0, no keepalive packets will be sent, but the tunnel will remain up. The default is 10.

- **Keepalive Retries** – Enter the number of keepalive retries the Virtual WAN Appliance should attempt before it brings down the tunnel. The default is 3.

4. Click **Apply**.

This submits your settings and adds the new LAN GRE Tunnel entry to the table.

5. To configure additional LAN GRE Tunnels, click **+** to the right of the **LAN GRE Tunnels** branch label, and proceed as above.

The next step is to configure the WAN links for the branch site.

How to Configure WAN Links for the Branch Site

Aug 09, 2017

1. Continuing in the site tree for the new site, click the **WAN Links** branch label.

   This reveals the Add (**+**) and Help (**?**) active icons to the right of the **WAN Links** branch.

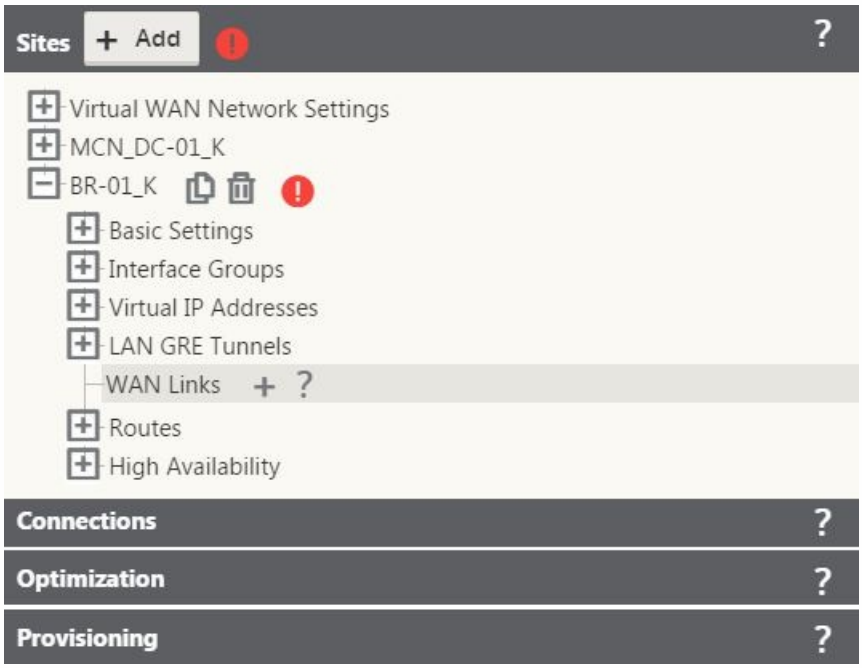

> ## Note
>
> At this point in a new configuration, there are no WAN links to form a table, and therefore no Open (**+**) icon to the left of the **WAN Links** branch label. However, if links exist, the **+** icon is available. If so, you can click **+** to the left of the **WAN Links** branch to display the table. This also reveals the Add (**+**), Edit (pencil), Delete (trashcan), and Help (**?**) active icons to the right of the **WAN Links** branch.

2. Click **+** to the right of the **WAN Links** branch to add a new WAN link.

   This opens the **Add WAN Link** dialog box.

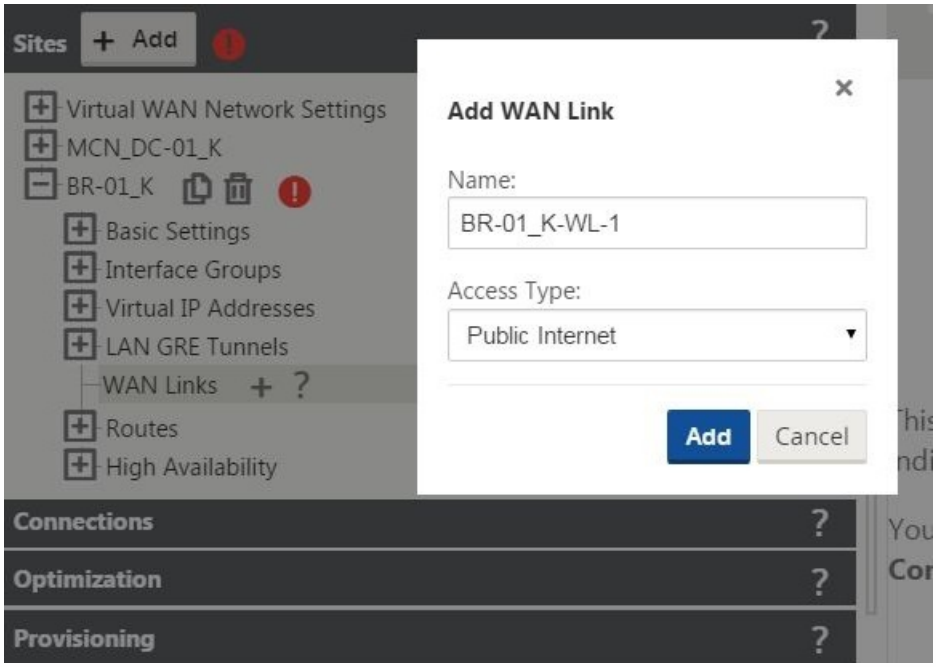

3. (Optional) Enter a name for the WAN Link if you do not want to use the default.

   The default is the site name, appended with the following suffix:

   -WL-<number>

   Where <number> is the number of WAN Links for this site, incremented by one.

4. Select the **Access Type** from the drop-down menu.

   The options are **Public Internet** or **Private Intranet.**

5. Click **Add**.

This displays the **WAN Links** table, adds the new un-configured link to the table, and opens the **Basic Settings** configuration form for the link.

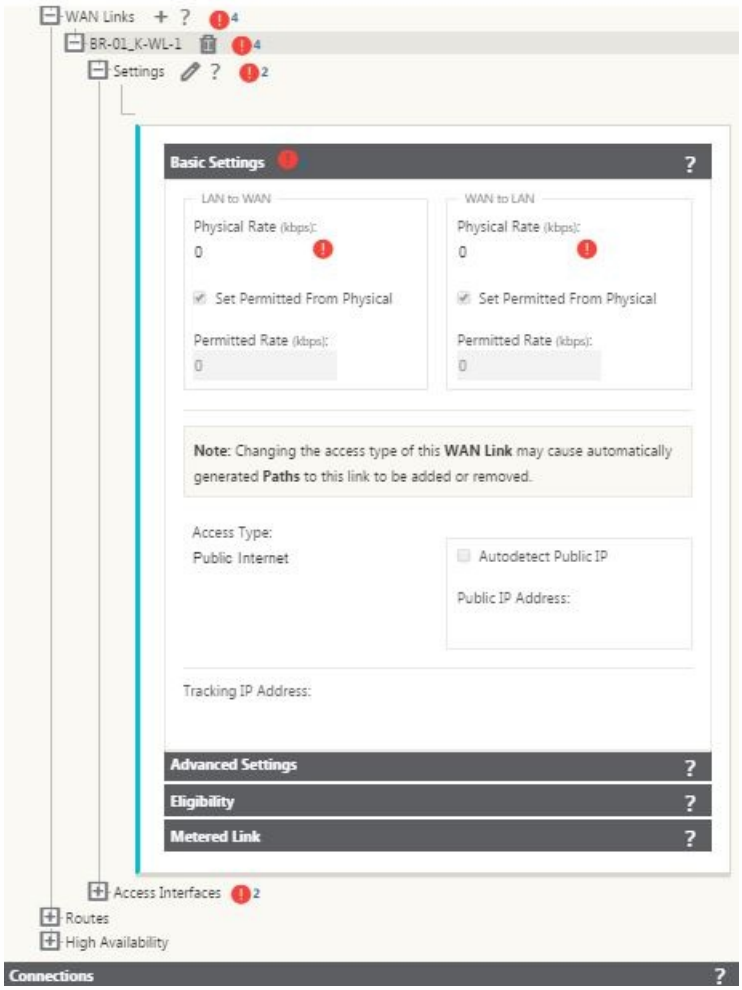

Note

At this point, you will see some additional Audit Alerts (red dot icons) display in various sections of the **WAN Links** configuration form. This is because you have not yet configured the settings for the new WAN link. You can ignore these for now, as these will be resolved automatically as you complete the configuration of the new WAN link.

6. Click the Edit (pencil) icon to the right of the **Settings** branch to enable editing of the form**.**

   This enables editing for the form, and reveals the **Apply** and **Close buttons.**

7. Enter the path information for the new WAN link.

   Some guidelines are as follows:

   - Some Internet links might be asymmetrical.
   - Misconfiguring the permitted speed can adversely affect performance for that path.
   - Avoid using burst speeds that surpass the Committed Rate.
   - For Internet WAN link paths, be sure to add the Public IP Address.

8. Click the grey **Advanced Settings** section bar.

This opens the **Advanced Settings** form for the link.

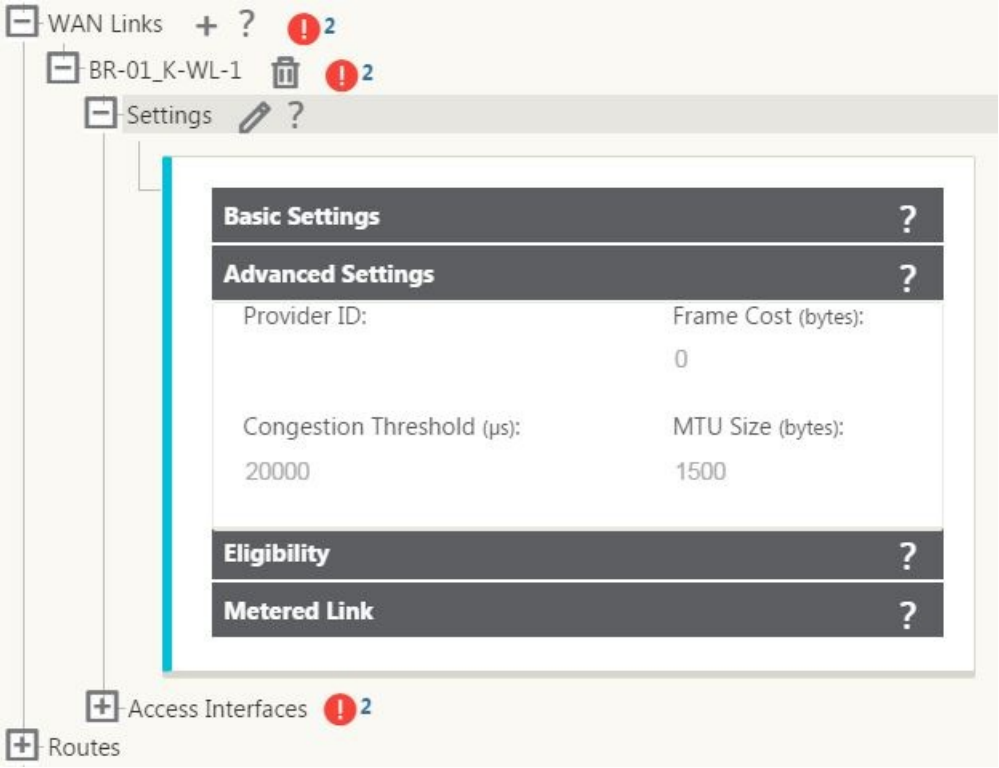

9. Enter the **Advanced Settings** for the link.
10. Click the grey **Eligibility** section bar.

    This opens the **Eligibility** settings form for the link.

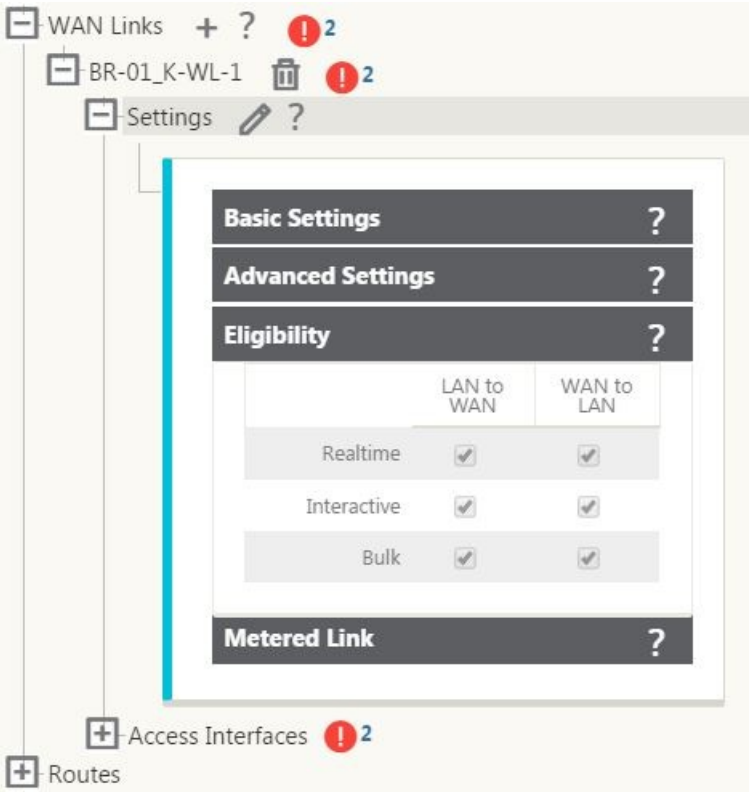

11. Select the **Eligibility** settings for the link.

12. Click the grey **Metered Link** section bar.

    This opens the **Metered Link** settings form for the link.

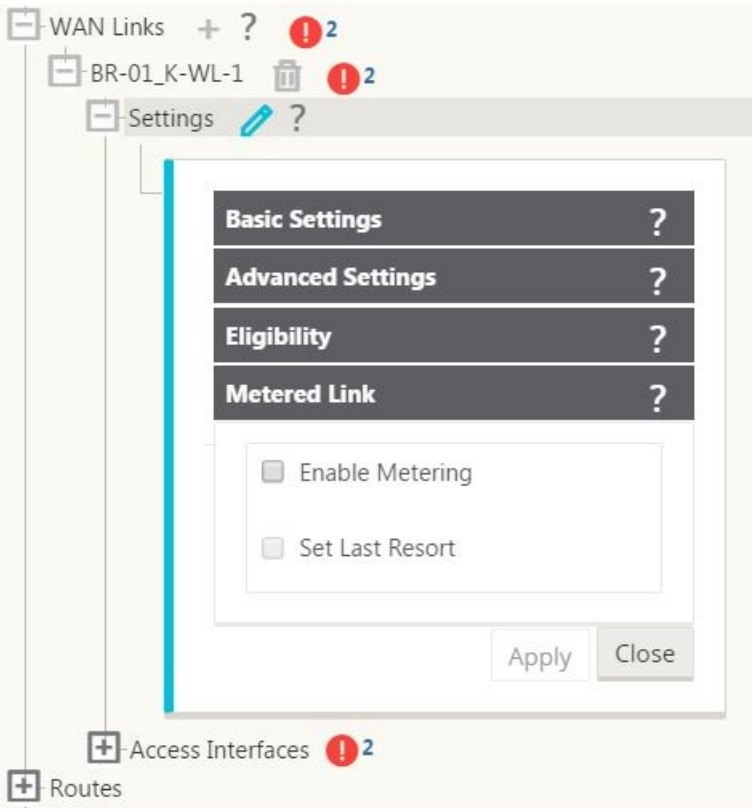

13. (Optional) Select **Enable Metering** to enable metering for this link.

    This displays the **Enable Metering** settings fields.

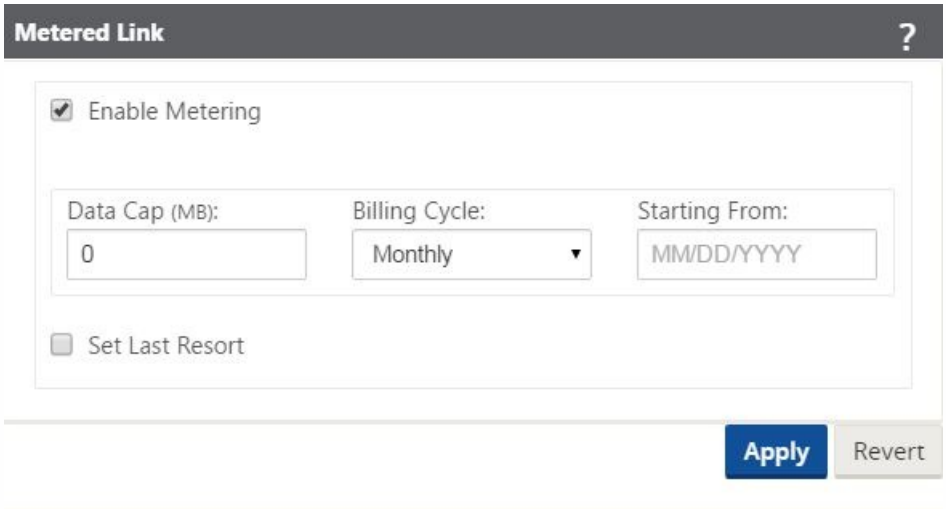

14. Configure the metering settings for the link.

    Enter the following:

    - **Data Cap (MB)** – Enter the data cap allocation for the link, in megabytes.

- **Billing Cycle** – Select either **Monthly** or **Weekly** from the drop-down menu.

- **Starting From** – Enter the start date of the billing cycle.

- **Set Last Resort** – Select this to enable this link as a link of last resort in the event of a failure of all other available links. Under normal WAN conditions, Virtual WAN sends only minimal traffic over metered links, for the purpose of checking link status. However, in the event of a failure, Virtual WAN can use active metered links as a last resort for forwarding production traffic.

15. Click **Apply**.

This applies your specified settings to the new WAN link.

The next step is to configure the Access Interfaces for the new WAN link. An Access Interface consists of a Virtual Interface, WAN endpoint IP Address, Gateway IP Address, and Virtual Path Mode defined collectively as an interface for a specific WAN link. Each WAN link must have at least one Access Interface.

16. Click **+** next to the **Access Interfaces** branch in the configuration tree for the link**.**

This opens the **Access Interfaces** table for the site.

17. Click **+** to the right of the **Access Interfaces** branch to add an interface.

This adds a blank entry to the table and opens it for editing.

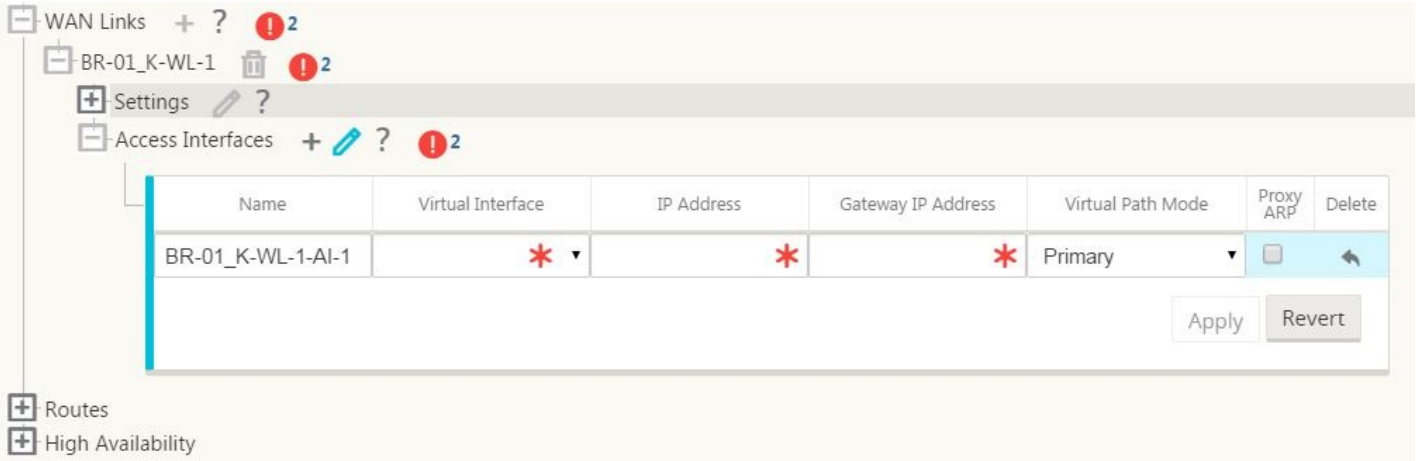

18. Enter the **Access Interfaces** settings for the link.

> **Note**
>
> Each WAN link must have at least one Access Interface.

Enter the following:

- **Name** – This is the name by which this Access Interface will be referenced. Enter a name for the new Access Interface, or accept the default. The default uses the following naming convention:

WAN_link_name-AI-number

Where *WAN_link_name* is the name of the WAN link you are associating with this interface, and number is the number of Access Interfaces currently configured for this link, incremented by 1.

> **Note**
>
> If the name appears truncated, you can place your cursor in the field, then click and hold and roll your mouse right or left to see the truncated portion.

- **Virtual Interface** – This is the Virtual Interface this Access Interface will use. Select an entry from the drop-down menu of Virtual Interfaces configured for this branch site.

- **IP Address** – This is the IP Address for the Access Interface endpoint from the appliance to the WAN.

- **Gateway IP Address –** This is the IP Address for the gateway router.

- **Virtual Path Mode** – This specifies the priority for Virtual Path traffic on this WAN link. The options are: **Primary, Secondary**, or **Exclude**. If set to **Exclude**, this Access Interface will be used for Internet and Intranet traffic, only.

- **Proxy ARP** – Select the checkbox to enable. If enabled, the Virtual WAN Appliance replies to ARP requests for the Gateway IP Address, when the gateway is unreachable.

19. Click **Apply.**

This applies your settings and adds the new Access Interface entry to the **Access Interfaces** table.

You have now finished configuring the new WAN link. Repeat these steps to add and configure additional WAN links for the

site.

The next step is to add and configure the routes for the site.

How to Configure Routes for the Branch Site

Aug 09, 2017

To add and configure the routes for the site, do the following:

1. Continuing in the site tree for the new site, click **+** to the left of the **Routes** branch.

   This displays the Routes table for the site.

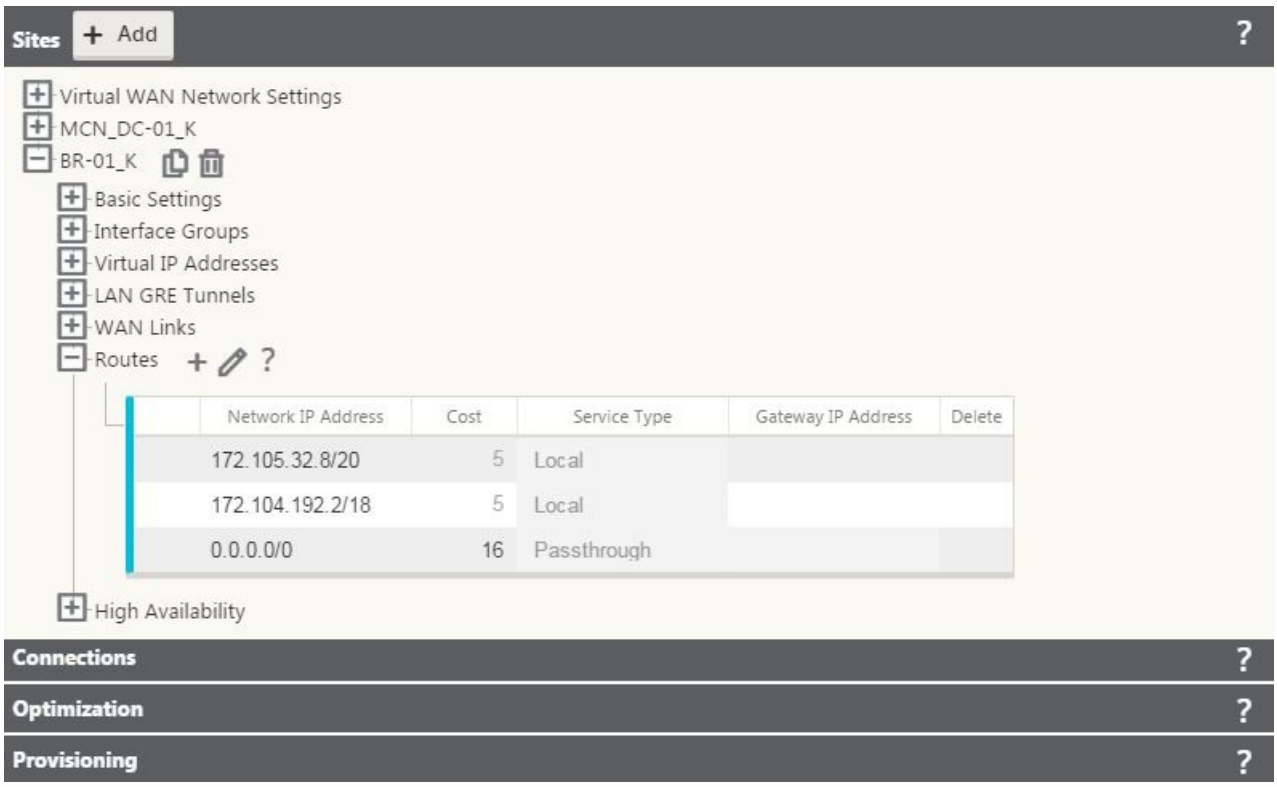

2. Click **+** to the right of the **Routes** branch to add a route.

   This opens the **Routes** table for editing and adds a blank route entry to the table.

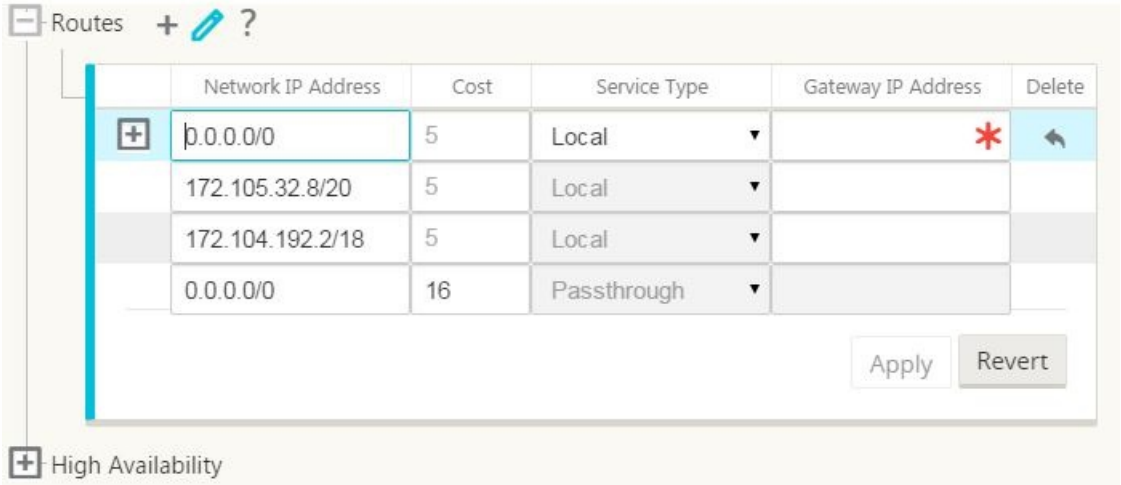

3. Enter the route configuration information and click **Apply**.

> **Note**
>
> After you click **Apply**, audit warnings appear indicating that further action is required. A red dot or goldenrod delta icon indicates an error in the section where it appears. You can use these warnings to identify errors or missing configuration information. Roll your cursor over an audit warning icon to display a short description of the error(s) in that section. You can also click the dark grey **Audits** status bar (bottom of page) to display a complete list of all audit warnings.

4. To add more routes for the site, click + to the right of the **Routes** branch, and proceed as above.
5. (Recommended.) Save your changes to the configuration.

> **Note**
>
> As an extra precaution, it is recommended that you use **Save As**, rather than **Save**, to avoid overwriting the wrong configuration package. Be sure to select **Allow Overwrite** before saving to an existing configuration, or your changes will not be saved.

You have now completed the required steps for configuring a client site. There are also some additional, optional steps you can choose to complete, before proceeding with the next phase of the deployment. A list of these steps and links to instructions are provided below. If you do not want to configure these features at this time, you can proceed directly to Preparing the Virtual WAN Appliance Packages on the MCN.

The optional steps are as follows:

- **Configure High Availability** – High Availability refers to a configuration in which two Virtual WAN Appliances at a site serve in an Active/Standby partnership capacity for redundancy purposes. If you are not implementing High Availability for this site, you can skip this step. For instructions, see Configuring High Availability (HA) for the Branch Site (Optional).
- **Clone the new branch site** – You have the option of cloning the branch site you just configured, and using that as a template for adding another site. The appliance models for the original site and the clone must be the same. For instructions, see Cloning the Branch Site (Optional).
- **Configure WAN Optimization –** If your CloudBridge Virtual WAN license includes WAN Optimization features, you have the option of enabling and adding these features to your configuration. To do so, you must complete the **Optimization** section in the **Configuration Editor**, and save the modified configuration. For instructions, proceed to Enabling and Configuring WAN Optimization.

How to Clone the Branch Site (Optional)

Aug 09, 2017

This section provides instructions for cloning the new branch site for use as a partial template for adding more branch sites.

> ## Note
>
> Cloning the site is optional. The Virtual WAN appliance models must be the same for both the original and the cloned sites. You cannot change the specified appliance model for a clone. If the appliance model is different for a site, you must manually add the site, as instructed in the previous sections.

Cloning a site streamlines the process of adding and configuring additional branch nodes. When a site is cloned, the entire set of configuration settings for the site are copied and displayed in a single form page. You can then modify the settings according to the requirements of the new site. Some of the original settings can be retained, where applicable. However, most of the settings must be unique for each site.

To clone a site, do the following:

1. In the **Sites** tree (middle pane) of the **Configuration Editor**, click **+** to the left of the branch site you want to duplicate.

   This opens that site branch in the **Sites** tree, and reveals the Clone button (double page icon) and Delete button (trashcan icon).

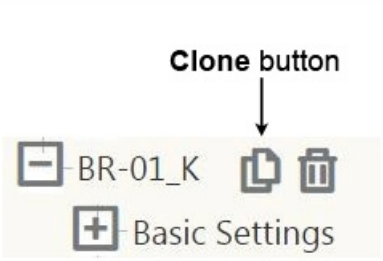

2. Click the **Clone** icon to the right of the branch site name in the tree.

   This opens the **Clone Site** configuration page.

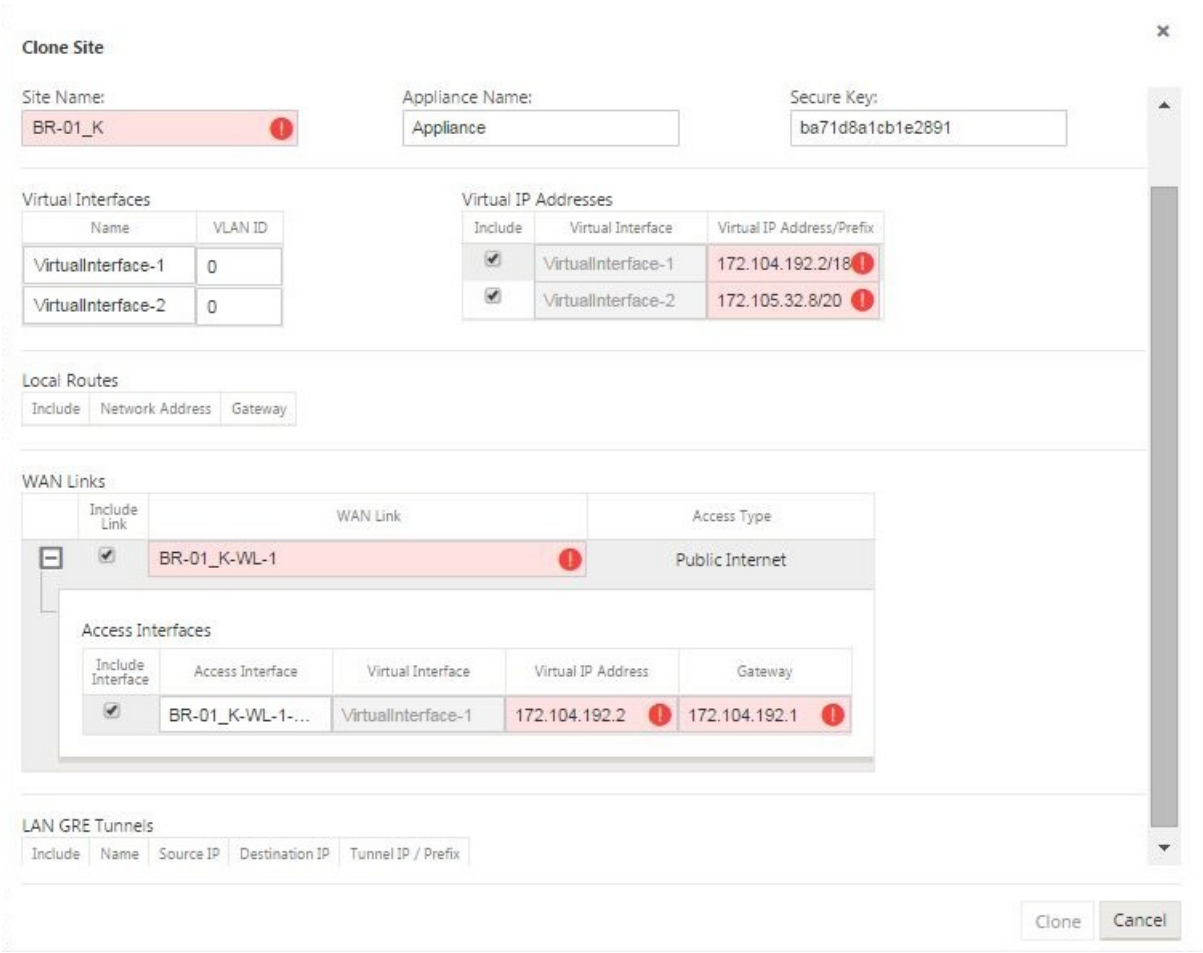

3. Enter the configuration parameter settings for the new site.

A pink field with an Audit Alert icon (red dot) indicates a required parameter setting that must have a value different than the setting for the original cloned site. In most cases, this value must be unique.

> ## Tip
>
> To further streamline the cloning process, use a consistent, pre-defined naming convention when naming the clones.

4. Resolve any Audit Alerts.

To diagnose an error, roll your cursor over the Audit Alert icon (red dot or goldenrod delta) to reveal bubble help for that specific alert.

5. Click **Clone** (far right corner) to create the new site and add it to the **Sites** table.

> ## Note
>
> The **Clone** button remains unavailable until you have entered all of the required values, and the new site configuration is error-free.

6. (Optional.) Save your changes to the configuration.

> **Note**
>
> As an extra precaution, it is recommended that you use **Save As**, rather than **Save**, to avoid overwriting the wrong configuration package. Be sure to select **Allow Overwrite** before saving to an existing configuration, or you changes will not be saved.

Repeat the steps up to this point for each branch site you want to add.

After you have finished adding all of the sites, the next step is to check the configuration for Audit Alerts, and make corrections or additions as needed.

How to Resolve Configuration Audit Alerts

Aug 09, 2017

An Audit Alert icon (a red dot or goldenrod delta) next to an item indicates a configuration error or missing parameter information for that item. A number next to the icon indicates the number of associated errors for that alert. To see bubble help for a particular alert, roll your cursor over the alert icon. This displays a brief description of the specific errors flagged by that alert. You must resolve all Audit Alerts in the configuration, or you will not be able to verify, stage, and activate the configuration package, later in the deployment process.

Resolving all of the Audit Alerts (if any), completes the **Sites** phase of the configuration. The next step is to save the completed **Sites** configuration.

How to Save the Completed Sites Configuration

Aug 09, 2017

The next step is to save the completed Sites configuration. The configuration will be saved to your workspace on the local appliance.

> ## Warning
>
> If the console session times out or you log out of the Management Web Interface before saving your configuration, any unsaved configuration changes will be lost. You must then log back into the system, and repeat the configuration procedure from the beginning. For that reason, it is strongly recommended that you save the configuration package often, or at key points in the configuration.

> ## Note
>
> As an extra precaution, it is recommended that you use **Save As**, rather than **Save**, to avoid overwriting the wrong configuration package.

To save the current configuration package, do the following:

1. Click **Save As** (at the top of the **Configuration Editor** middle pane).

   This opens the **Save As** dialog box,.

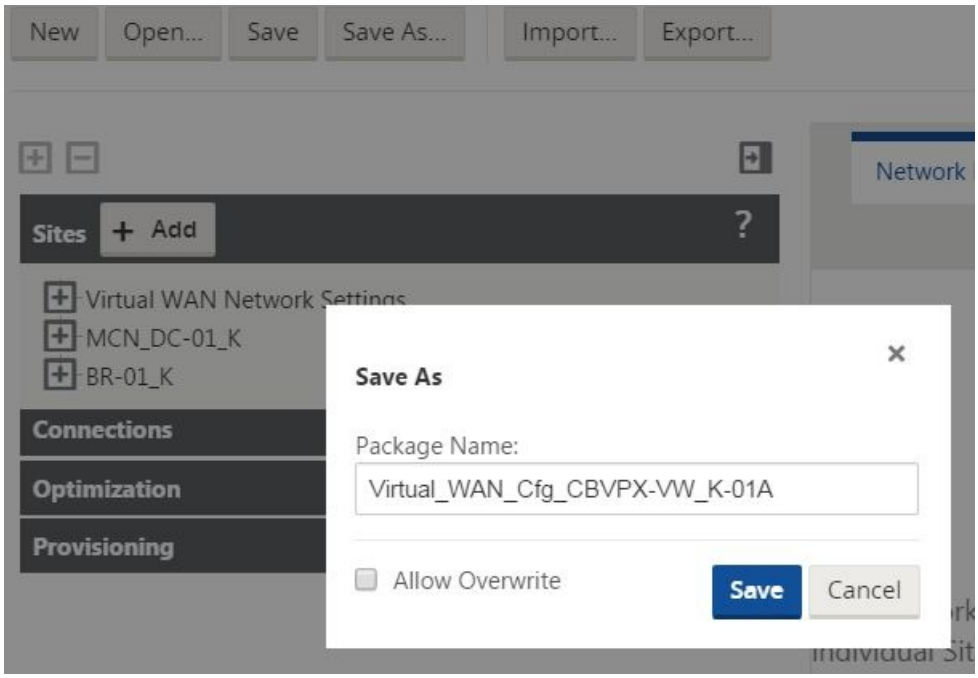

2. Enter the configuration package name.

> **Note**
>
> If you are saving the configuration to an existing configuration package, be sure to select **Allow Overwrite** before saving.

3. Click **Save**.

> **Note**
>
> After saving the configuration file, you have the option to log out of the Management Web Interface and continue the configuration process at a later time. However, if you log out, you will need to reopen the saved configuration when you resume. Instructions are provided in the section Loading a Saved Configuration Package into the Configuration Editor.

The next step is to configure the Virtual Paths and Virtual Path Service between the MCN and the client sites. Instructions are provided in the Configuring the Virtual Path Service Between the MCN and Client Sites.

Deployment use Cases

Aug 09, 2017

Following are some of the use case scenarios implemented by using NetScaler SD-WAN appliances:

- Deploying SD-WAN in Gateway Mode
- Deploying SD-WAN in PBR mode (Virtual Inline Mode)
- Dynamic Paths for Branch to Branch Communication
- Static WAN Paths
- Building an SD-WAN Network
- Routing for LAN Segementation
- Utilizing Enterprise Edition Appliance to Provide WAN Optimization Services Only

Deploying SD-WAN in Gateway Mode

Aug 09, 2017

To deploy SD-WAN in a Gateway Mode:

This article provides step-by-step procedure to configure a SD-WAN appliance in Gateway mode in a sample network setup. Inline deployment is also described for the branch side to complete the configuration.

Gateway mode places the SD-WAN appliance physically in the path (two-arm deployment) and requires changes in the existing network infrastructure to make the SD-WAN appliance the default gateway for the entire LAN network for that site.

> ## Note
>
> An SD-WAN deployed in Gateway mode acts as a Layer 3 device and cannot perform fail-to-wire. All interfaces involved will be configured for "Fail-to-block". In the event of appliance failure, the default gateway for the site will also fail, causing an outage until the appliance and default gateway are restored.

Topology

DataCenter in Gateway Deployment

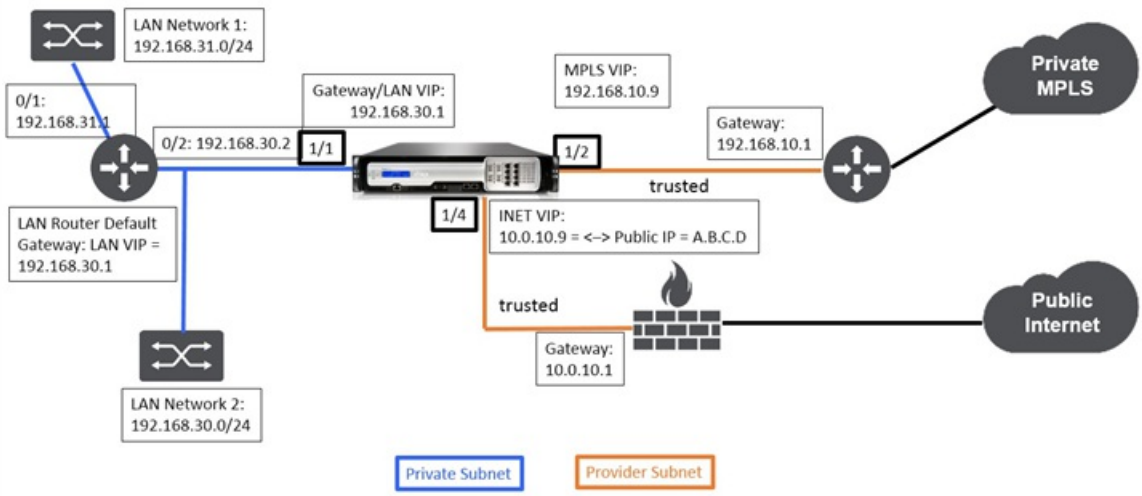

Branch in Inline Deployment

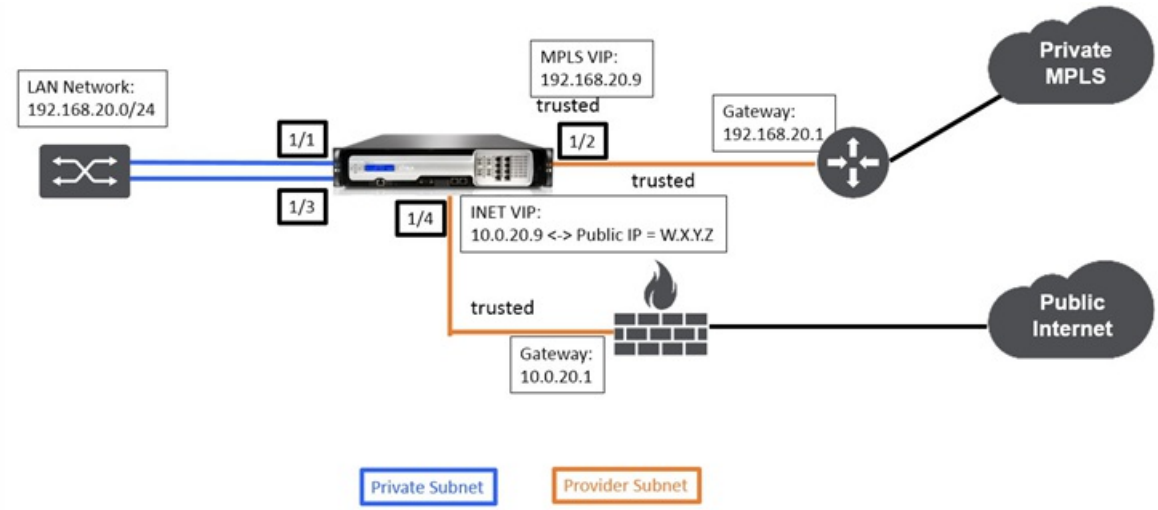

Deployment Requirements

Deployment requirements and related information is described below to assist you in building the configuration.

Site Name	DataCenter Site	Branch Site
Appliance Name	A_DC1	A_BR1
Management IP	172.30.2.10/24	172.30.2.20/24
Security Key	If any	If any
Model/Edition	4000	2000
Mode	Gateway	Inline
Topology	2 x WAN Path	2 x WAN Path
VIP Address	192.168.10.9/24 – MPLS 10.0.10.9/24 – Internet (Public IP – A.B.C.D) 192.168.30.1/24 - LAN	192.168.20.9/24 - MPLS 10.0.20.9/24 – Internet (Public IP – W.X.Y.Z)
Gateway MPLS	192.168.10.1	192.168.20.1
Gateway Internet	10.0.10.1	10.0.20.1
Link Speed	MPLS – 100 Mbps Internet – 20 Mbps	MPLS – 10 Mbps Internet – 2 Mbps
Route	Network IP Address - 192.168.31.0/24 Service Type - Local Gateway IP Address - 192.168.30.2	If any
VLANs	If any	If any

Configuration Pre-requisites

- Enable SD-WAN appliance as a Master Control Node.
- Configuration is done only on the Master Control Node (MCN) of the SD-WAN appliance.

To enable an appliance as a Master Control Node:

1. In the NetScaler SD-WAN web management interface, navigate to **Configuration** > **Appliance Settings** > **Administrator Interface** > **Miscellaneous tab** > **Switch Console**.

> **Note**
>
> If "Switch to Client Console" is displayed, then the appliance is already in MCN mode. There should only be one active MCN in a SD-WAN network.

2. Start Configuration by navigating to **Configuration** > **Virtual WAN** > **Configuration Editor**. Click the **New** to begin configuration.

Datacenter Site Gateway Mode Configuration

Following are the high-level configuration steps to configure Datacenter site Gateway deployment:

1. Create a new DC site.
2. Populate Interface Groups based on connected Ethernet interfaces.
3. Create Virtual IP address for each virtual interface.
4. Populate WAN links based on physical rate and not burst speeds using Internet and MPLS Links.
5. Populate Routes if there are additional subnets in the LAN infrastructure.

To create a new DC site

1. Navigate to **Configuration Editor** - > **Sites**, and click the "**+" Add** button.
2. Populate the fields as shown below.
3. Keep default settings unless instructed to change.

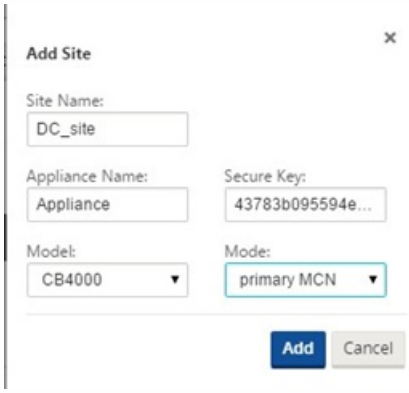

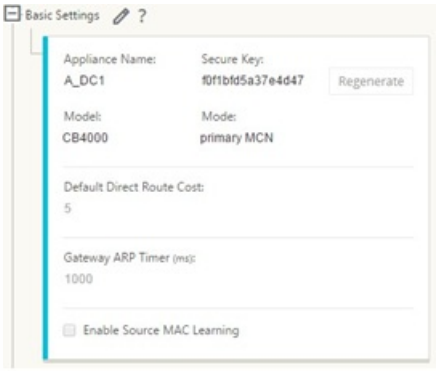

To configure interface groups based on connected Ethernet interfaces

1. In the **Configuration Editor**, navigate to **Sites → [Site Name] → Interface Groups**. Click **"+"**to add interfaces intended to be used. For Gateway Mode, each Interface Group is assigned a single Ethernet interface.
2. Bypass mode is set to **fail-to-block** since only one Ethernet/physical interface is used per virtual interface. There are also no Bridge Pairs.
3. In this example three Interfaces Groups are created, one facing the LAN and two others facing each respective WAN Link. Refer to the sample "DC Gateway Mode" topology above and populate the Interface Groups fields as shown below.

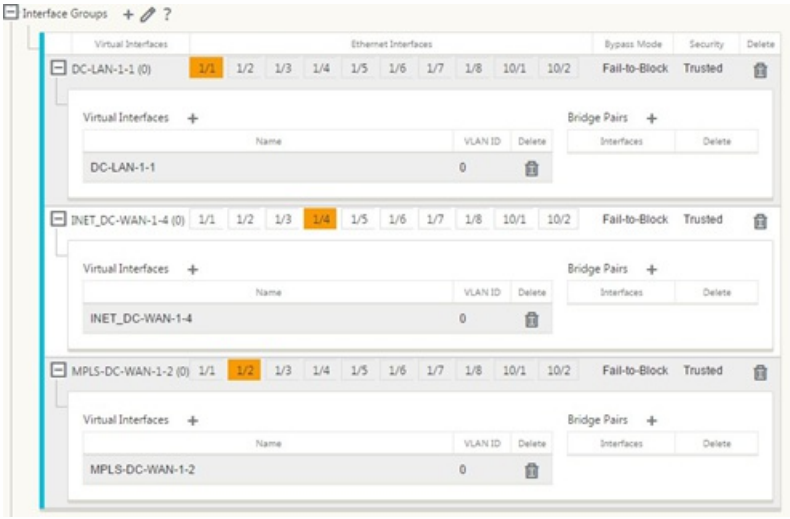

To create Virtual IP (VIP) address for each virtual interface

1. Create a VIP on the appropriate subnet for each WAN Link. VIPs are used for communication between two SD-WAN appliances in the Virtual WAN environment.
2. Create a Virtual IP Address to be used as the Gateway address for the LAN network

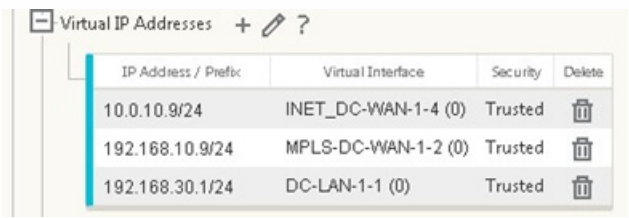

To populate WAN links based on physical rate and not on burst speeds using Internet link

1. Navigate to **WAN Links**, click the **"+"** button to add a WAN Link for the Internet link.

2. Populate Internet link details, including the supplied Public IP address as shown below. Note that **AutoDetect Public IP** cannot be selected for SD-WAN appliance configured as MCN.
3. Navigate to **Access Interfaces**, click the **"+"** button to add interface details specific for the Internet link.
4. Populate Access Interface for IP and gateway addresses as shown below.

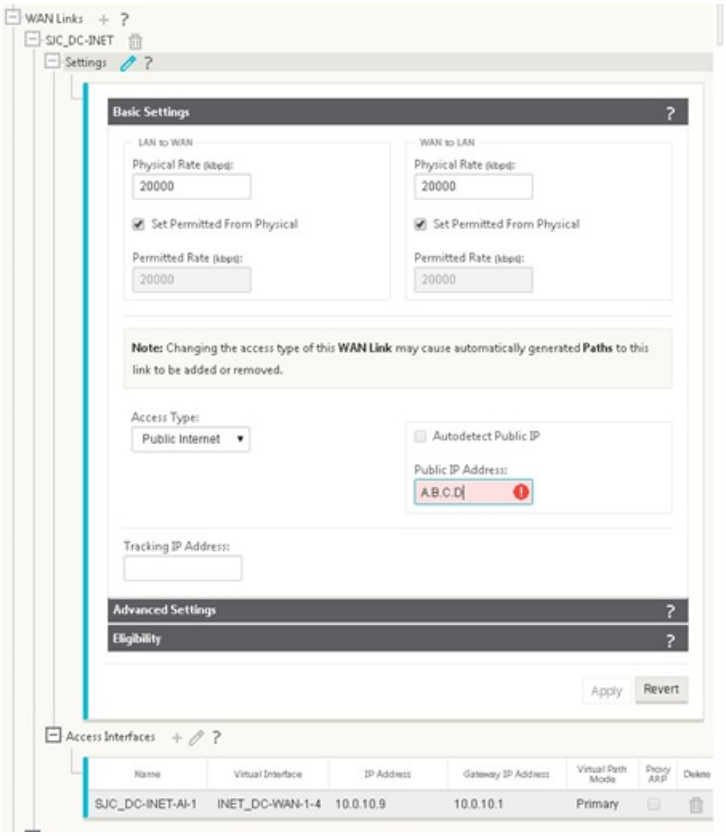

To create MPLS Link

1. Navigate to **WAN Links**, click the **"+"** button to add a WAN Link for the MPLS link.
2. Populate MPLS link details as shown below.
3. Navigate to **Access Interfaces**, click the **"+"** button to add interface detail specific for the MPLS link.
4. Populate Access Interface for IP and gateway addresses as shown below.

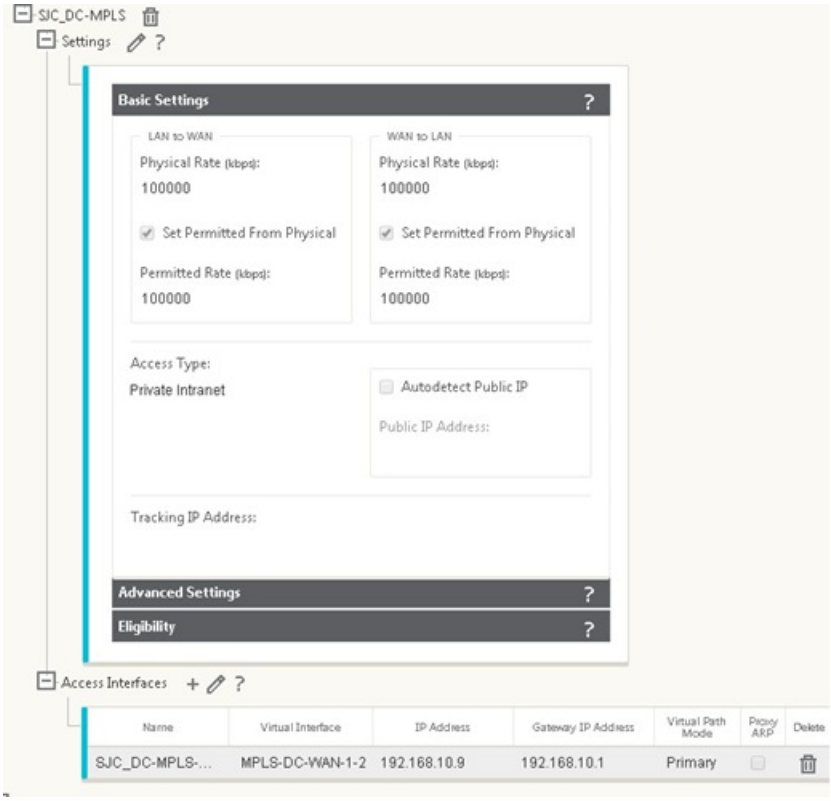

To populate Routes

Routes are auto-created based on the above configuration. The DC LAN sample topology shown above has an additional LAN subnet which is **192.168.31.0/24**. A route needs to be created for this subnet. Gateway IP address must be in the same subnet as the DC LAN VIP as shown below.

Branch Site Inline Deployment Configuration

Following are the high-level configuration steps to configure Branch site for Inline deployment:

1. Create a new Branch site.
2. Populate Interface Groups based on connected Ethernet interfaces.
3. Create Virtual IP address for each virtual interface.
4. Populate WAN links based on physical rate and not burst speeds using Internet and MPLS Links.
5. Populate Routes if there are additional subnets in the LAN infrastructure.

To create a new Branch site

1. Navigate to **Configuration Editor** - > **Sites**, and click the "**+" Add** button.
2. Populate the fields as shown below.
3. Keep default settings unless instructed to change.

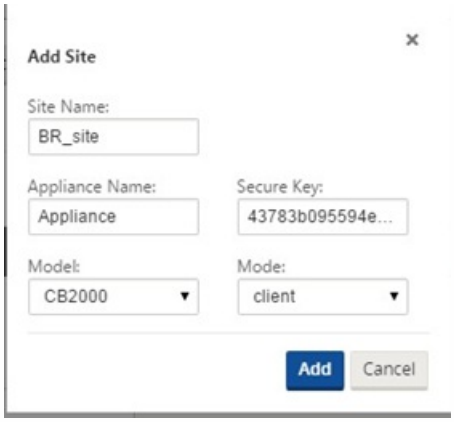

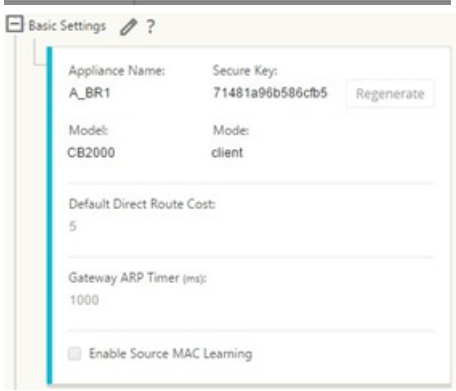

To populate interface groups based on connected Ethernet interfaces

1. In the **Configuration Editor**, navigate to **Sites → [Client Site Name] → Interface Groups**. Click **"+"**to add interfaces intended to be used. For Inline Mode, each Interface Group is assigned two Ethernet interfaces.
2. Bypass mode is set to **fail-to-wire** and Bridge Pair is created using the two Ethernet interfaces.
3. Refer to the sample "Remote Site Inline Mode" topology above and populate the Interface Groups fields as shown below.

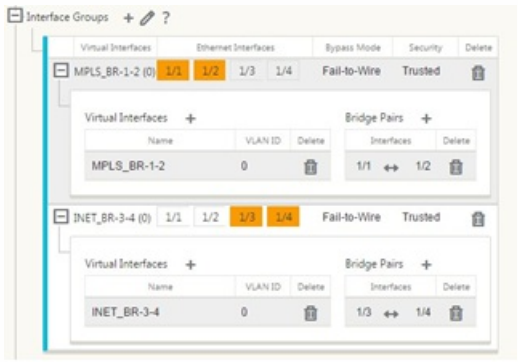

To create Virtual IP (VIP) address for each virtual interface

1. Create a Virtual IP address on the appropriate subnet for each WAN Link. VIPs are used for communication between two SD-WAN appliances in the Virtual WAN environment.

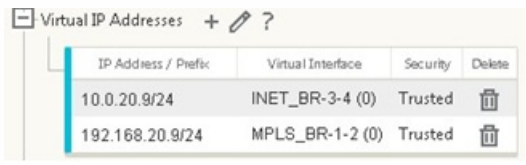

To populate WAN links based on physical rate and not on burst speeds using Internet link

1. Navigate to **WAN Links**, click the **"+"** button to add a WAN Link for the Internet link.
2. Populate Internet link details, including the AutoDetect Public IP address as shown below.
3. Navigate to **Access Interfaces**, click the **"+"** button to add interface details specific for the Internet link.
4. Populate Access Interface for IP address and gateway as shown below.

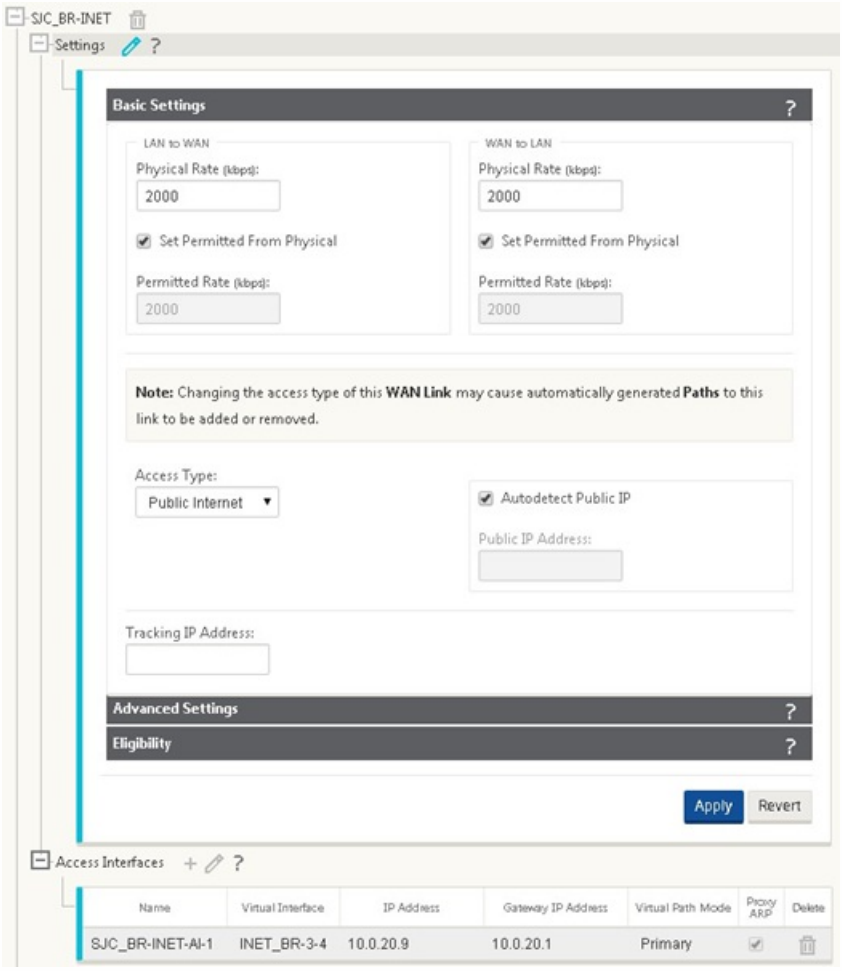

To create MPLS Link

1. Navigate to WAN Links, click the **"+"** button to add a WAN Link for the MPLS link.
2. Populate MPLS link details as shown below.
3. Navigate to Access Interfaces, click the **"+"** button to add interface details specific for the MPLS link.
4. Populate Access Interface for IP address and gateway as shown below.

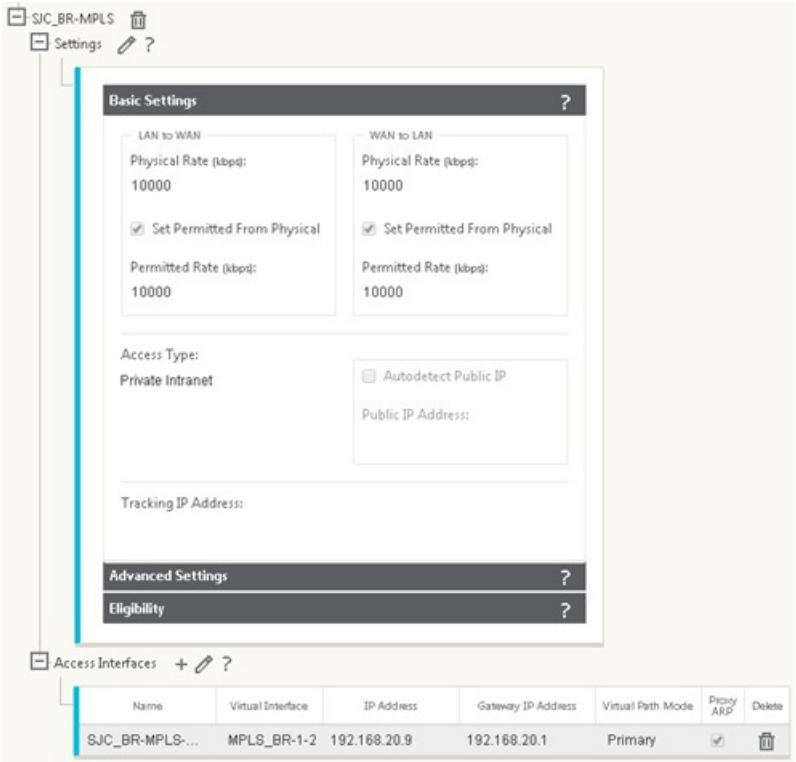

To populate Routes

Routes are auto-created based on above configuration. In case there are additional subnets specific to this remote branch office, then specific routes need to be added identifying which gateway to direct traffic to in order to reach those backend subnets.

Resolving Audit Errors

After completing configuration for DC and Branch sites, you will be alerted to resolve audit error on both DC and BR sites.

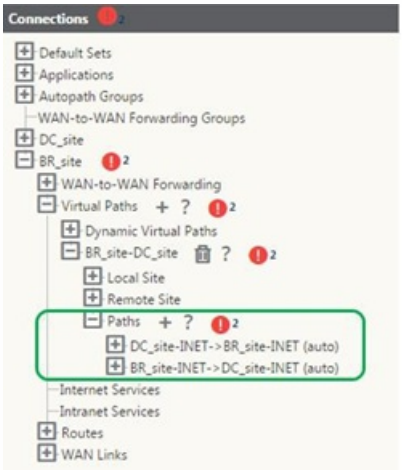

By default, the system will generate paths for WAN Links defined as access type Public Internet. You would be required to use the auto-path group function or enable paths manually for WAN Links with an access type of Private Internet. Paths

for MPLS links can be enabled by clicking on the Add operator (in the green rectangle).

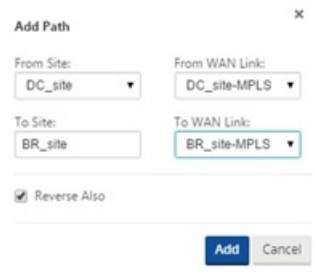

After completing all the above steps, proceed to Preparing the SD-WAN Appliance Packages on the MCN topic.

Deploying SD-WAN in PBR mode (Virtual Inline Mode)

Aug 17, 2017

In virtual inline mode, the router uses policy based routing rules to redirect incoming and outgoing WAN traffic to the appliance, and the appliance forwards the processed packets back to the router.

The following article describes the step-by-step procedure to configure two SD-WAN (SD-WAN SE) appliances:

- Data Center Appliance in PBR mode (Virtual Inline Mode)
- Branch Appliance in Inline mode
- PBR needs to be configured either at the core switch or further upstream at the router. The router must monitor the health of the SD-WAN appliance so that the appliance can be bypassed if it fails.
- Virtual Inline Mode places the SD-WAN appliance physically out of path (one-arm deployment) i.e. only a single Ethernet interface to be used (Example: Interface 1/1) with bypass mode set to fail-to-block (FTB).

NetScaler SD-WAN appliance needs to be configured to pass traffic to the proper gateway. Traffic intended for the Virtual Path is directed towards the SD-WAN appliance and then encapsulated and directed to the appropriate WAN link.

Gathering Information for Configuration

- Accurate network diagram (example diagram show below) of your local and remote site(s) including:

  - Local and Remote WAN links and their bandwidths in both directions, their subnets, Virtual IP Addresses and Gateways from each link, Routes, and VLANs.

- Deployment Table (example diagram shown below)

Data Center Topology – PBR mode (Virtual Inline Mode)

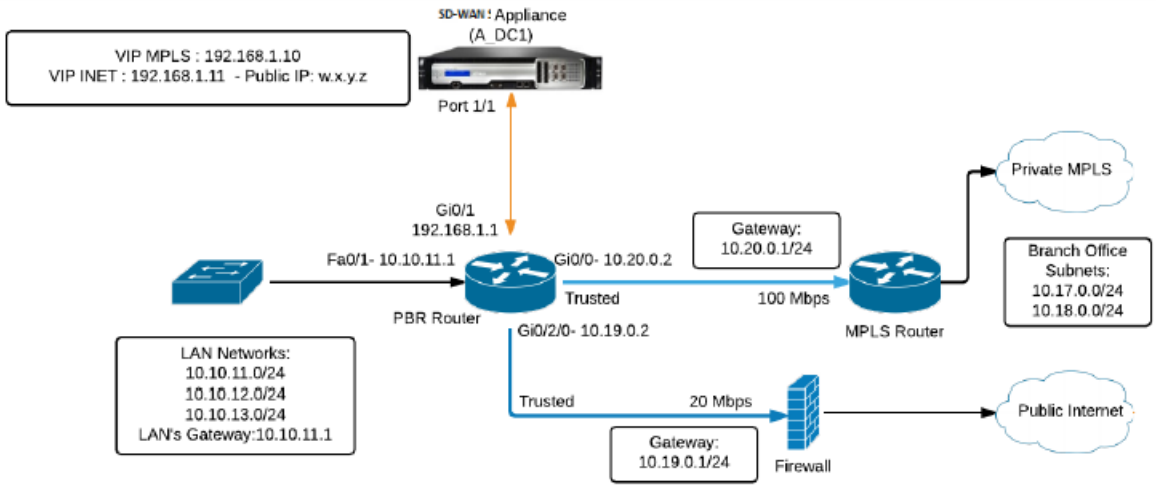

Branch Topology – Inline Mode

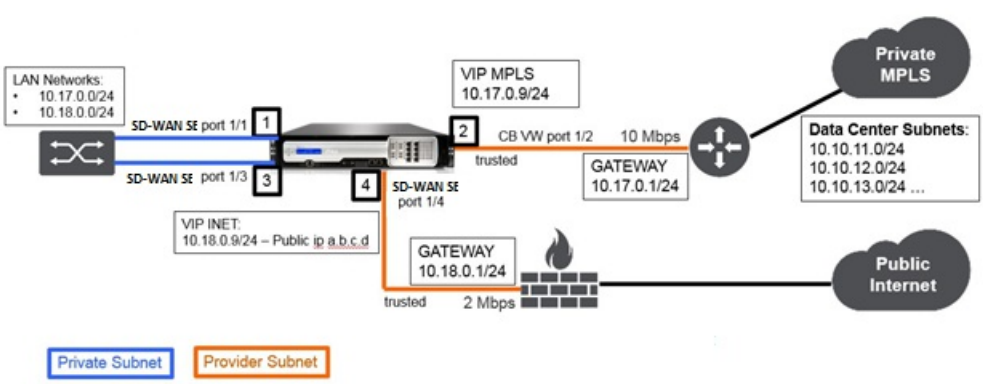

Site Name	DataCenter Site	Branch Site
Appliance Name	A_DC1	A_BR1
Management IP	172.30.2.10/24	172.30.2.20/24
Security Key	If any	If any
Model/Edition	4000	2000
Mode	PBR mode (Virtual Inline Mode)	Inline
Topology	2 x WAN Path	2 x WAN Path
VIP Address	192.168.1.10/24 – MPLS 192.168.1.11/24 – Internet*Public IP w.x.y.z	10.17.0.9/24 - MPLS 10.18.0.9/24 – Internet *Public IP a.b.c.d
Gateway MPLS	10.20.0.1	10.17.0.1
Gateway Internet	10.19.0.1	10.18.0.1
Link Speed	MPLS – 100 Mbps Internet – 20 Mbps	MPLS – 10 Mbps Internet – 2 Mbps
Route	Need to add a route on the SD-WAN SE Appliance on how to reach the LAN Subnets (10.10.11.0/24, 10.10.12.0/24, 10.10.13.0/24, etc) through any of the physical interfaces: Gi0/1 - 192.168.1.1 Configuration > Virtual WAN > Configuration Editor > SJC_DC > Routes. In this example interface 192.168.1.1 was used: • n/w address: 10.10.13.0/24, 10.10.12.0/24, 10.10.11.0/24 • service type: local • gateway IP address: 192.168.1.1	No additional routes were added

Site Name	**DataCenter Site**	**Branch Site**

Steps to configure a site in Virtual Inline Mode:

- Enable the MCN functionality.
- Create a New site.
- Create an Interface Group and Virtual Interfaces.
- Assign Virtual IP Address to Virtual Interfaces.
- Create WAN Links and assign IP address.
- Add Routes.
- Troubleshooting.
- Policy Based Routing configuration on the PBR Router.

Configuration Pre-requisites

- Enable SD-WAN appliance as a Master Control Node.
- Configuration is done only on the Master Control Node (MCN) of the SD-WAN appliance.

To enable an appliance as a Master Control Node:

1. In the NetScaler SD-WAN web management interface, navigate to **Configuration** > **Appliance Settings** > **Administrator Interface** > **Miscellaneous tab** > **Switch Console**.

> **Note**
>
> If "Switch to Client Console" is displayed, then the appliance is already in MCN mode. There should only be one active MCN in a SD-WAN network.

2. Enable Virtual WAN Service.

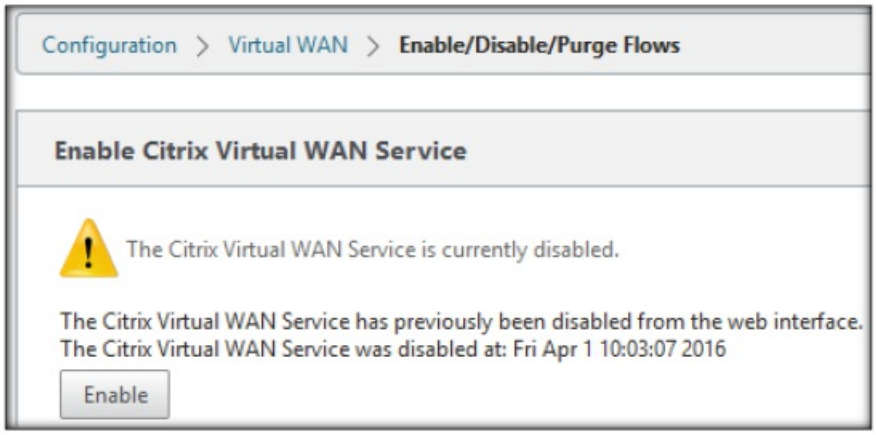

3. Start Configuration by navigating to **Configuration** > **Virtual WAN** > **Configuration Editor**. Click **New** to begin configuration.

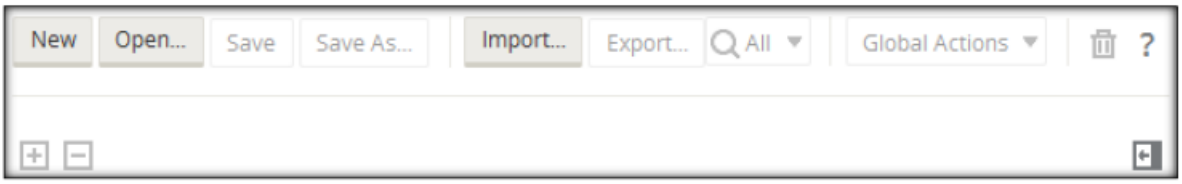

This operation will create an Untitled_1 initial configuration file which can be renamed [optional] later using the **Save As** button.

Following are the high-level configuration steps to configure Datacenter site in PBR deployment mode:

1. Create a new DC site.
2. Configure Interface Groups based on connected Ethernet interfaces.
3. Configure Virtual IP address for each virtual interface.
4. Populate WAN links based on physical rate and not burst speeds using Internet and MPLS Links.
5. Populate Routes if there are additional subnets in the LAN infrastructure.

Datacenter Site PBR Mode Configuration

To create a new DC site

1. Navigate to Configuration Editor - > Sites, and click the "+" Add button.
2. Populate the fields as shown below.
3. Keep default settings unless instructed to change.

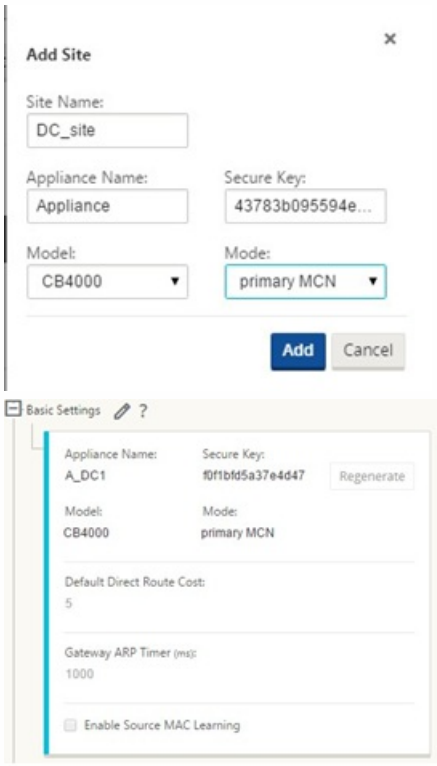

To configure interface groups based on connected Ethernet interfaces

1. In the Configuration Editor, navigate to Sites → [Site Name] → Interface Groups. Click “+”to add interfaces intended to be used. In PBR mode, configuration on only a single Ethernet interface is used i.e. interface connecting the upstream router providing PBR policy implications (Example- Interface 1/1).
2. Bypass mode is set to fail-to-block since only one Ethernet/physical interface is used per virtual interface. There are also no Bridge Pairs.
3. In this example, expand Virtual Interfaces + option and configure the Virtual Interfaces.

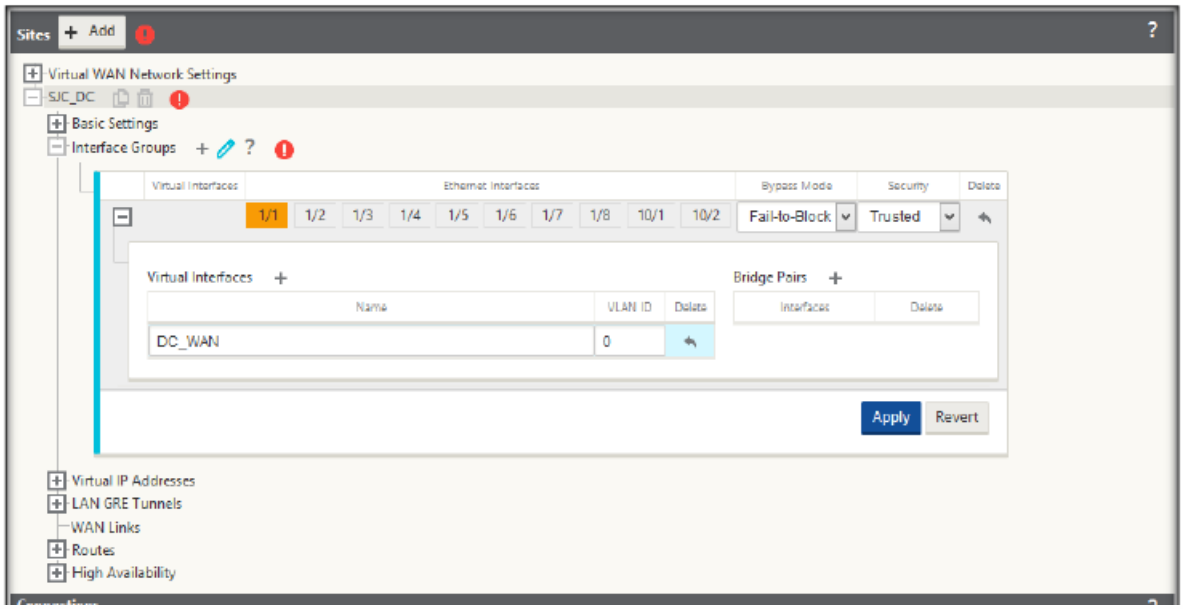

To create Virtual IP (VIP) address for each virtual interface

1. Create a **Virtual IP Address** on the appropriate subnet for each WAN Link. VIPs are used for communication between two SD-WAN appliances in the Virtual WAN environment.

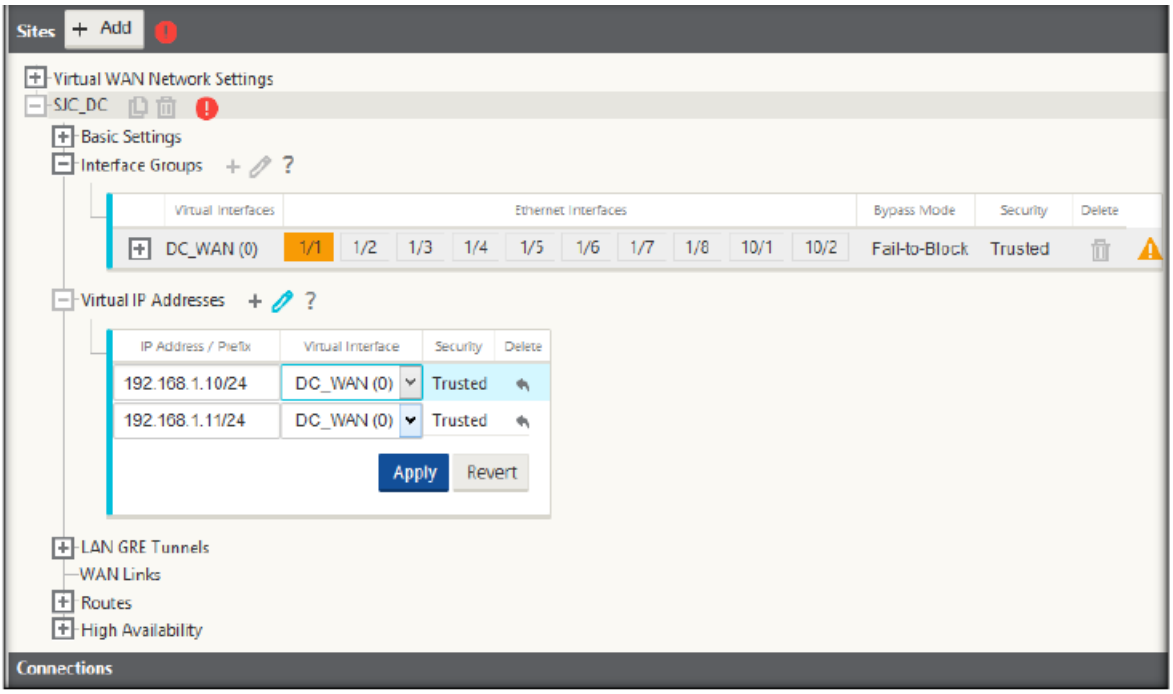

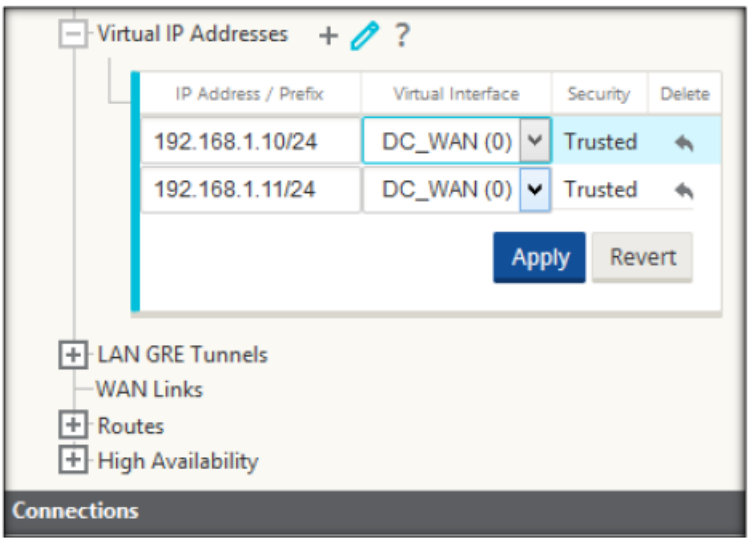

To populate WAN links based on physical rate and not on burst speeds using Internet and MPLS link:

1. Navigate to **WAN Links**, click the "**+**" button to add a WAN Link for the Internet link.
2. Populate Internet link details, including the supplied Public IP address as shown below. Note that **AutoDetect Public IP** cannot be selected for SD-WAN appliance configured as MCN.
3. Navigate to **Access Interfaces**, click the "+" button to add interface details specific for the Internet link.
4. Populate Access Interface for IP and gateway addresses as shown below. The **Proxy ARP** is not checked for less than two Ethernet interfaces.

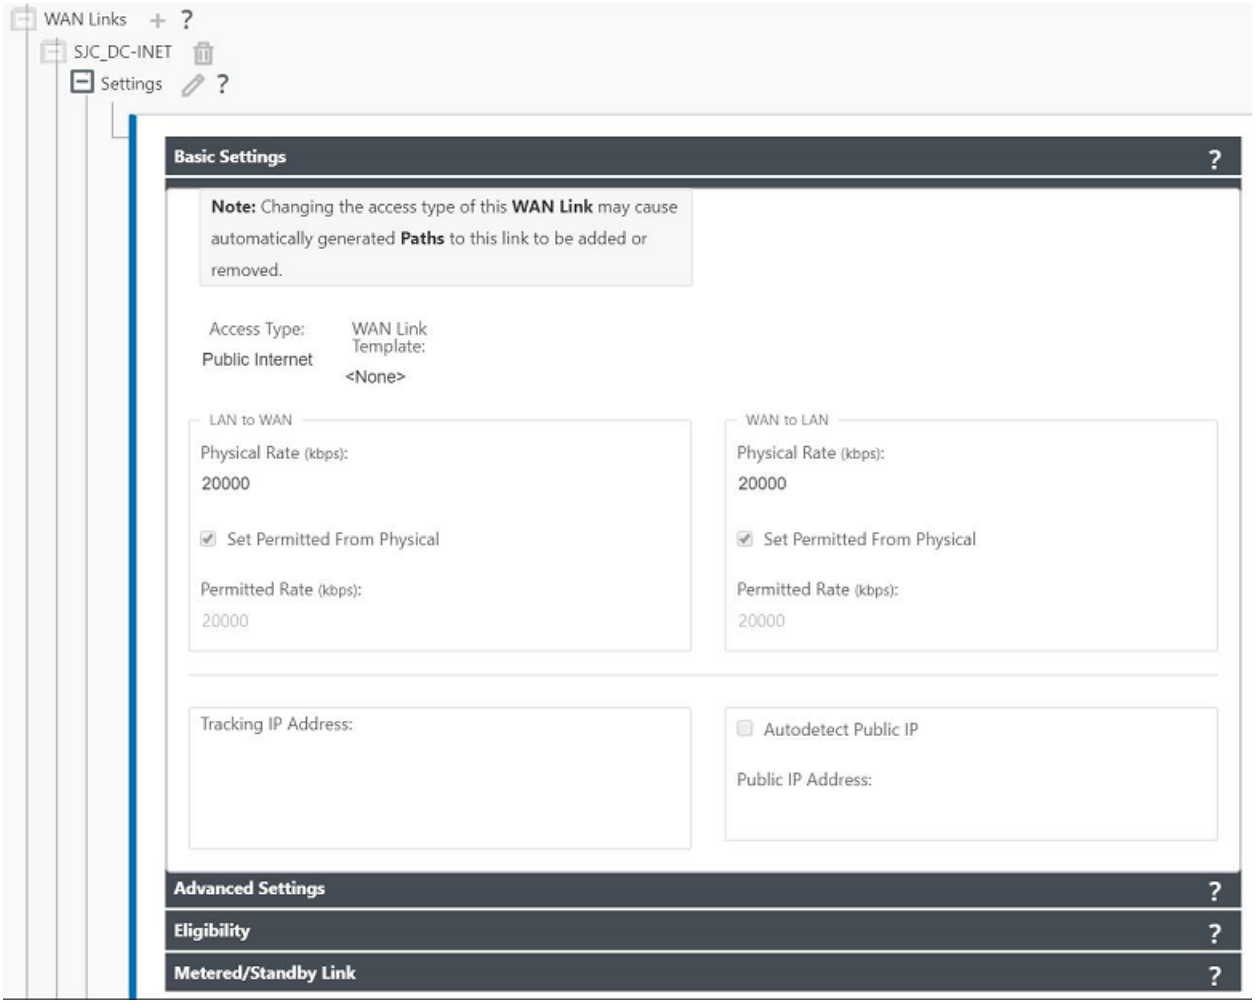

To create MPLS Link

1. Navigate to **WAN Links**, click the "**+**" button to add a WAN Link for the MPLS link.
2. Populate MPLS link details as shown below.
3. Navigate to **Access Interfaces**, click the "+" button to add interface detail specific for the MPLS link.
4. Populate Access Interface for MPLS Virtual IP and gateway addresses as shown below.

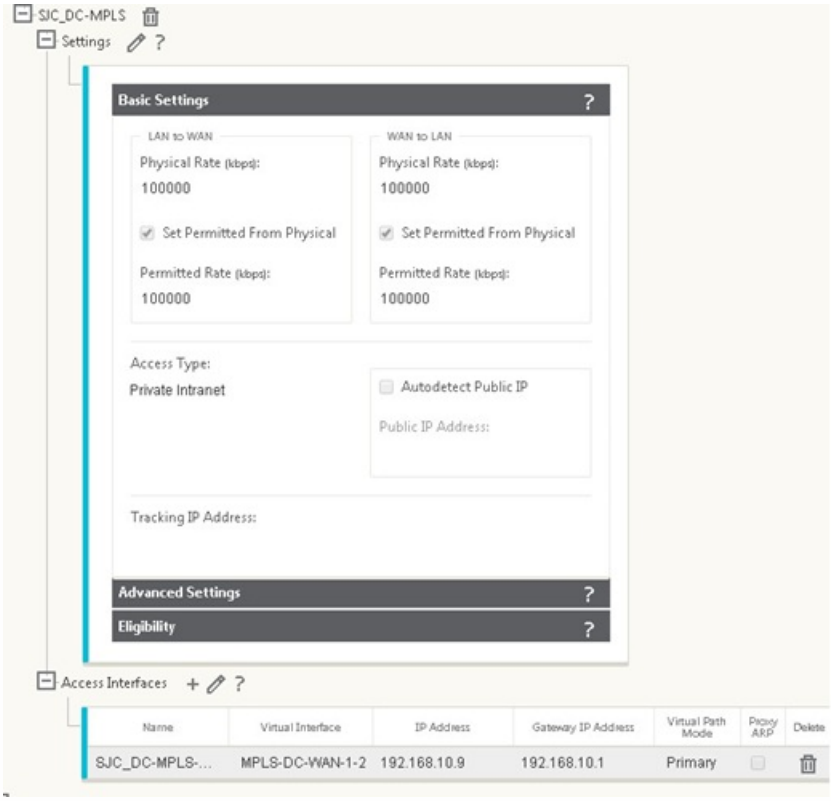

Note

The Proxy ARP is not checked for less than two Ethernet interfaces.

To populate Routes

On the Data center site, add a route on the SD-WAN SEE appliance to reach the LAN Subnets (10.10.11.0/24, 10.10.12.0/24, 10.10.13.0/24, etc) through any of the physical interfaces:

0/1/0.1 – 192.168.1.1 on VLAN 10

0/1/0.2 – 192.168.2.1 on VLAN 20

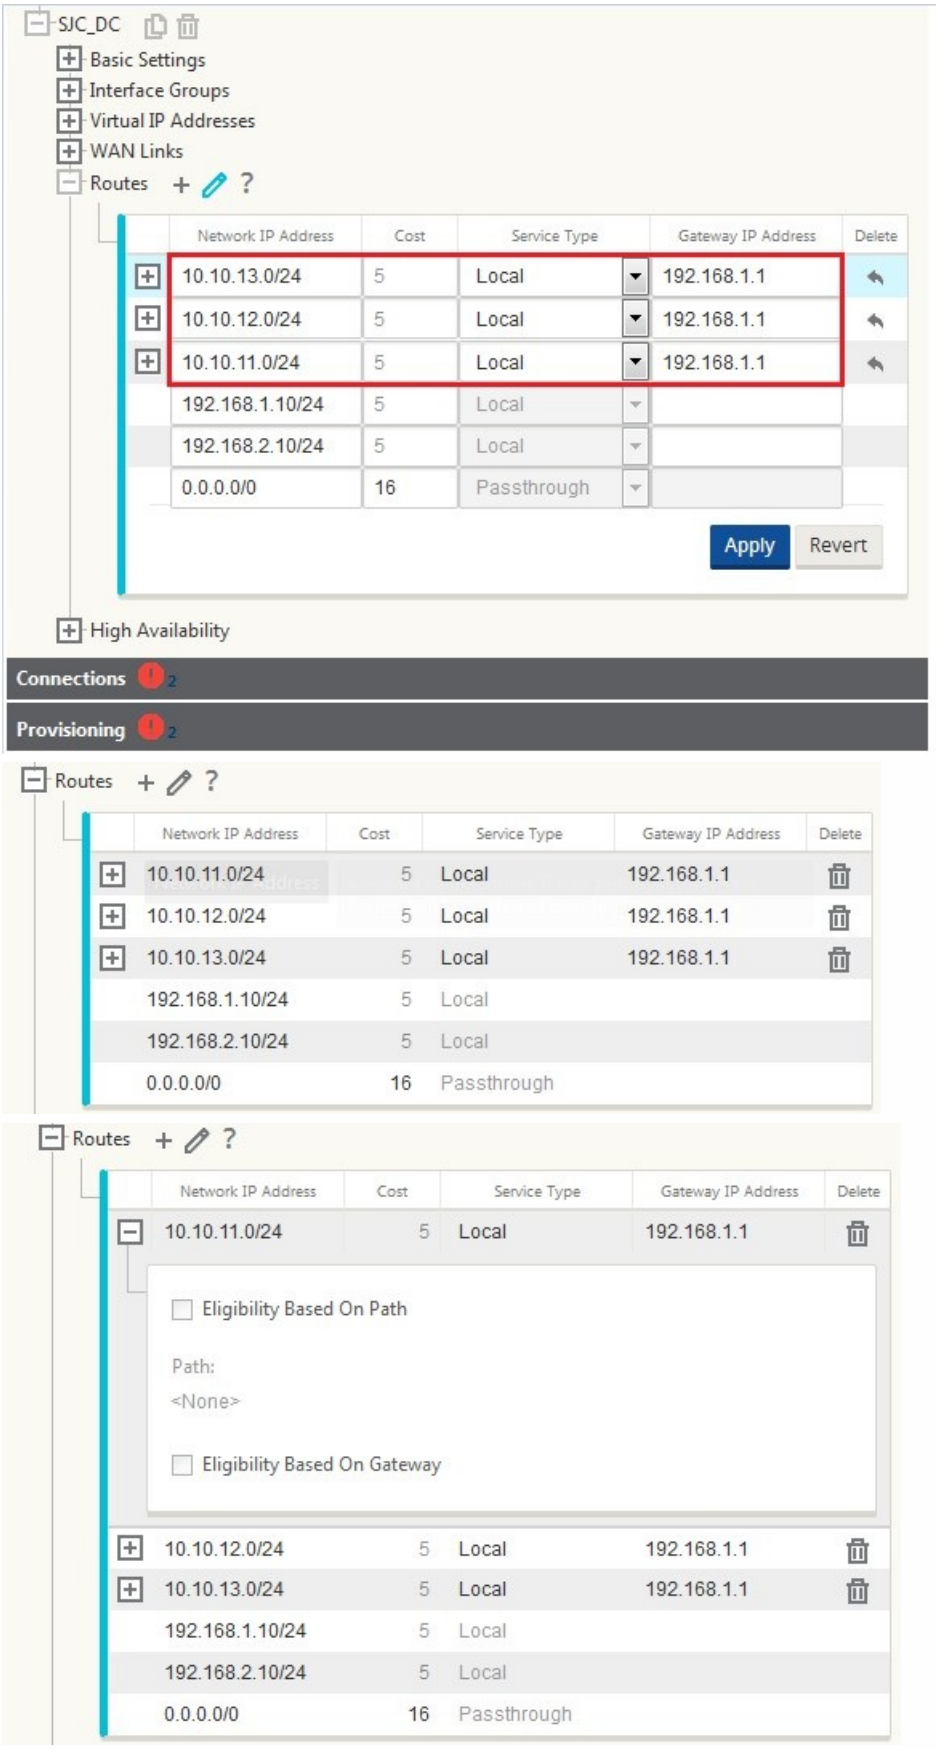

Branch Site Inline Deployment Configuration

Following are the high-level configuration steps to configure Branch site for Inline deployment:

1. Create a new Branch site.

2. Populate Interface Groups based on connected Ethernet interfaces.

3. Create Virtual IP address for each virtual interface.

4. Populate WAN links based on physical rate and not burst speeds using Internet and MPLS Links.

- Virtual Interface “INTERNET” configured on Bridge pair 1/3 and 1/4
- Virtual Interface “MPLS” configured con Bridge Pair 1/1 and 1/2

5. Populate Routes if there are additional subnets in the LAN infrastructure.

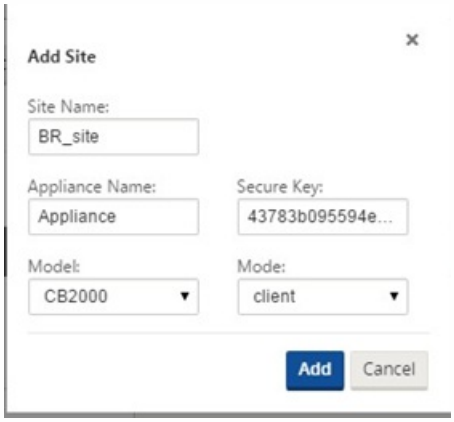

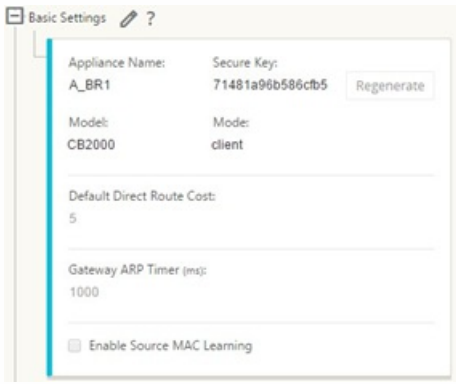

To populate interface groups based on connected Ethernet interfaces

1. In the **Configuration Editor**, navigate to **Sites → [Client Site Name] → Interface Groups**. Click **“+**”to add interfaces intended to be used. For Inline mode configuration, four Ethernet interface are used; interface pair 1/3, 1/4 and interface pair 1/1 and 1/2.
2. Bypass mode is set to fail-to-wire since two Ethernet/physical interfaces are used per virtual interface. There are two bridge Pairs.
3. Populate WAN links based on physical rate and not burst speeds using Internet and MPLS Links.

- Virtual Interface “INTERNET” configured on Bridge pair 1/3 and 1/4
- Virtual Interface “MPLS” configured con Bridge Pair 1/1 and 1/2.

4. Refer to the sample “Remote Site Inline Mode” topology above and populate the Interface Groups fields as shown below.

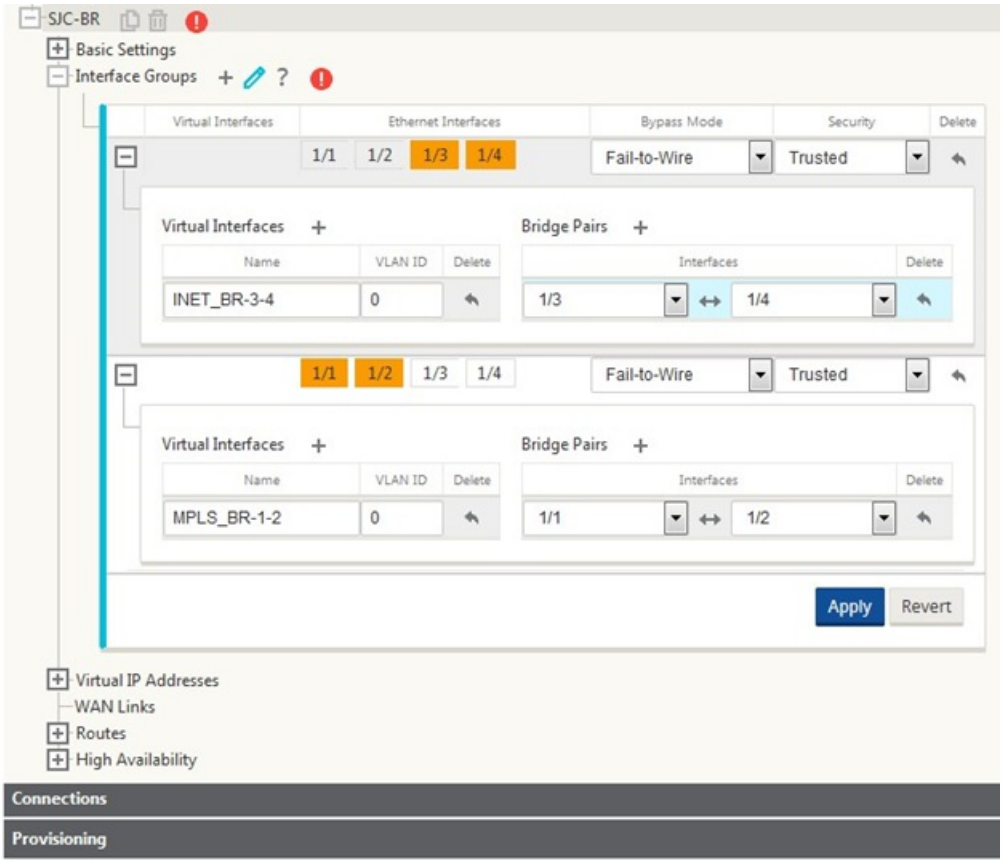

To create Virtual IP (VIP) address for each virtual interface

1. Create a Virtual IP address on the appropriate subnet for each WAN Link. VIPs are used for communication between two SD-WAN appliances in the Virtual WAN environment.

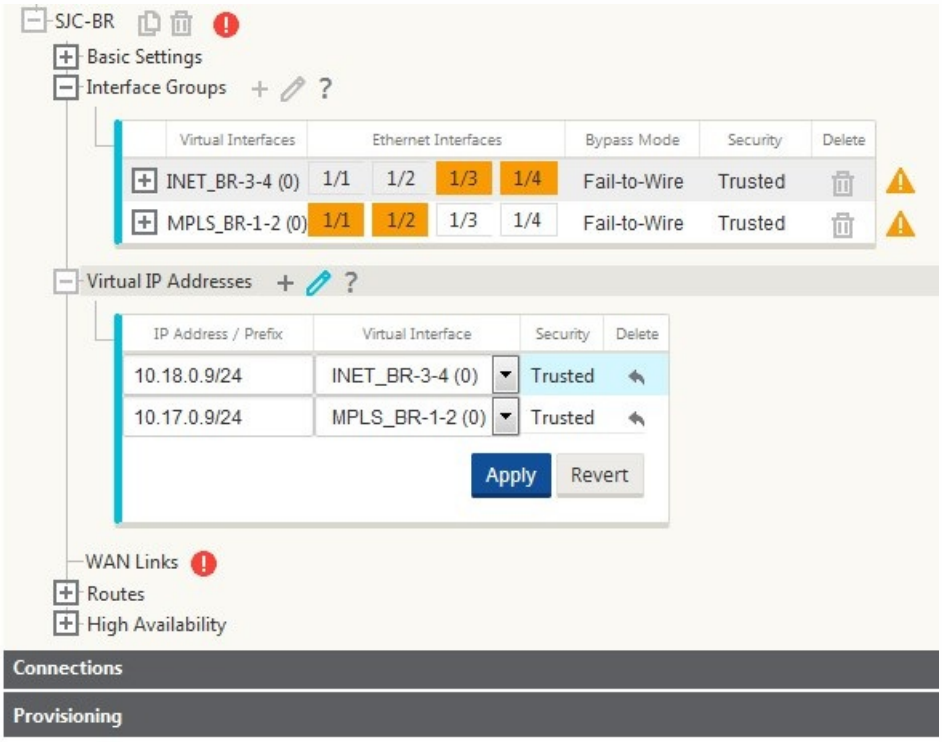

To populate WAN links based on physical rate and not on burst speeds using Internet link

1. Navigate to **WAN Links**, click the "**+**" button to add a WAN Link for the Internet link.
2. Populate Internet link details, including the **AutoDetect Public IP address** as shown below.
3. Navigate to **Access Interfaces**, click the "**+**" button to add interface details specific for the Internet link.
4. Populate Access Interface for Virtual IP address and gateway as shown below.

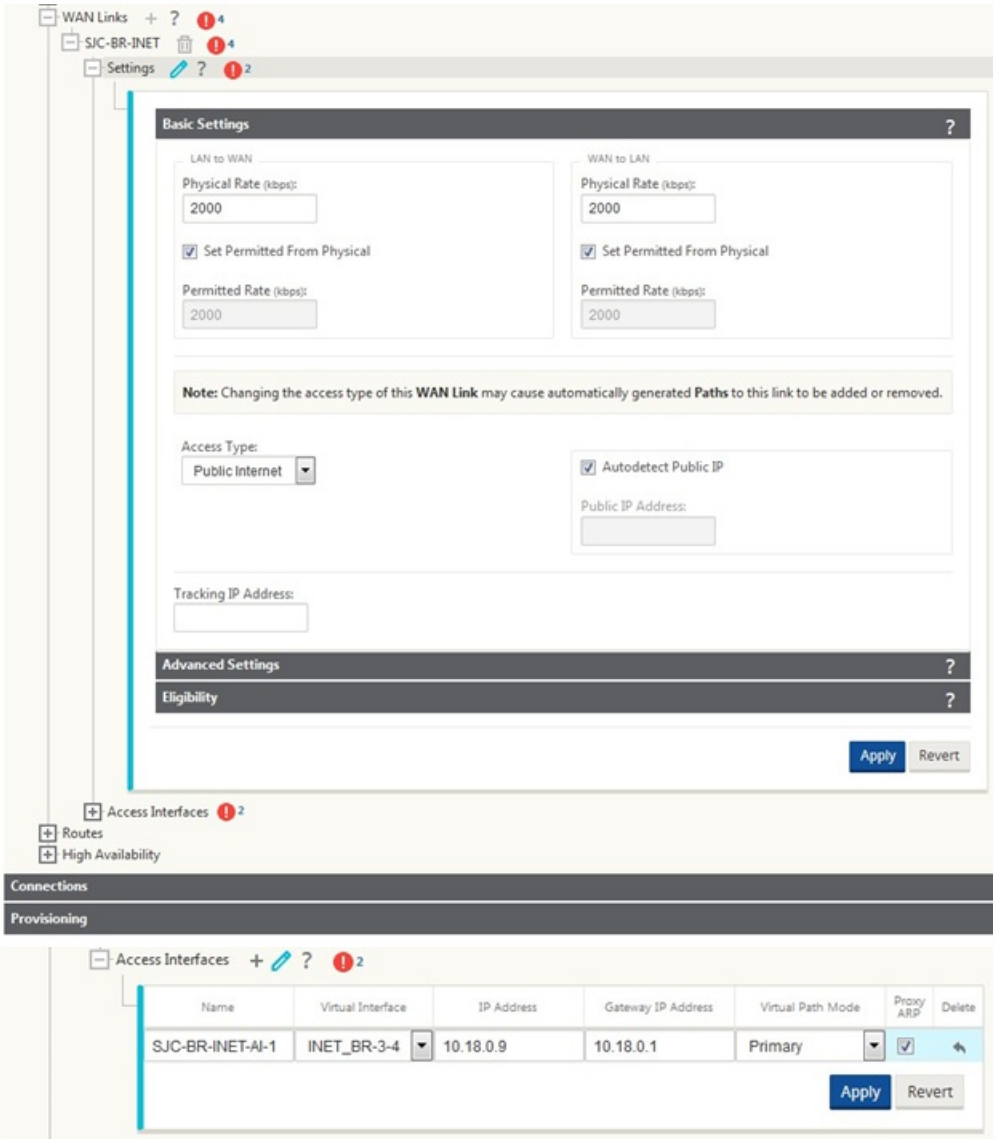

To create MPLS Link

1. Navigate to **WAN Links,** click the "**+**" button to add a WAN Link for the MPLS link.
2. Populate MPLS link details as shown below.
3. Navigate to **Access Interfaces**, click the "**+**" button to add interface details specific for the MPLS link.
4. Populate Access Interface for Virtual IP address and gateway as shown below.

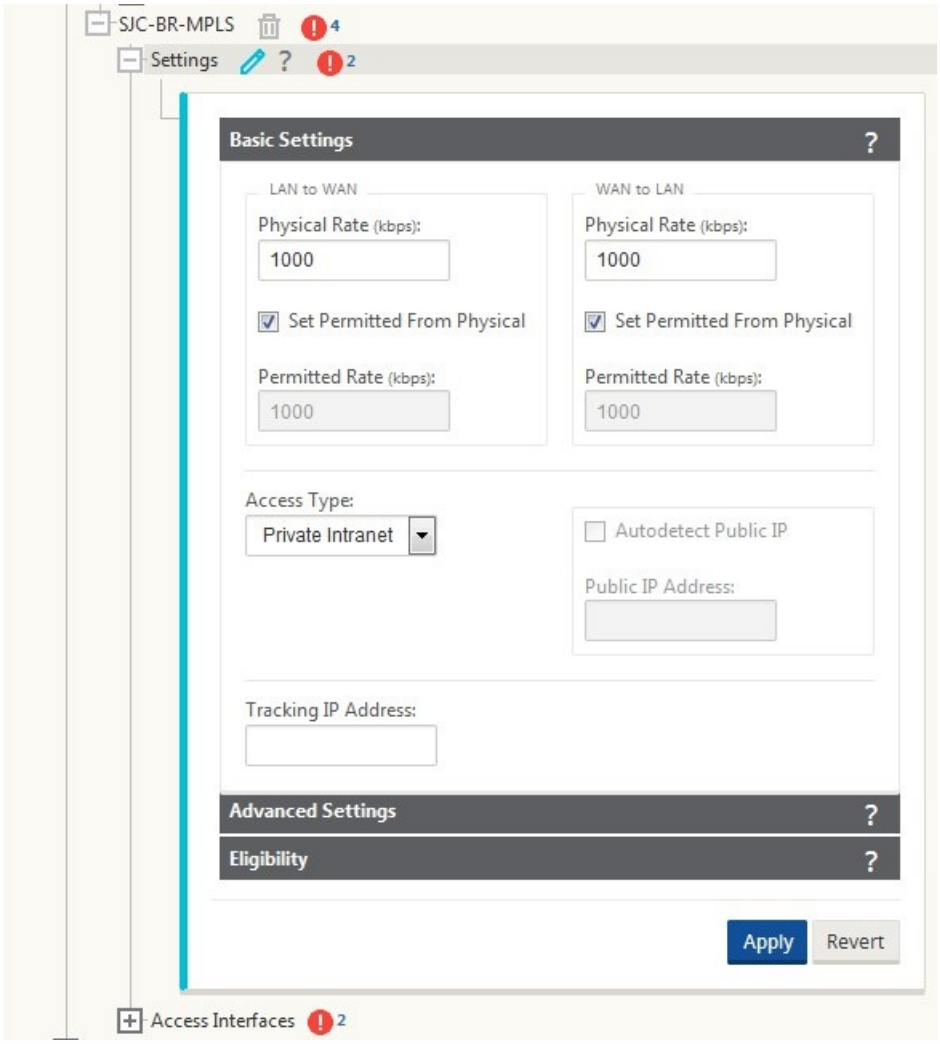

To populate Routes

Routes are auto-created based on above configuration. In case there are additional subnets specific to this remote branch office, then specific routes need to be added identifying which gateway to direct traffic to in order to reach those backend subnets.

Resolving Audit Errors

After completing configuration for DC and Branch sites, you will be alerted to resolve audit error on both DC and BR sites.

In this example, we will resolve the Audit Error related to Private Intranet WAN Link [SJC_DC-MPLS].

Note: By default the system will generate paths for WAN Links defined as access type Public Internet (highlighted).

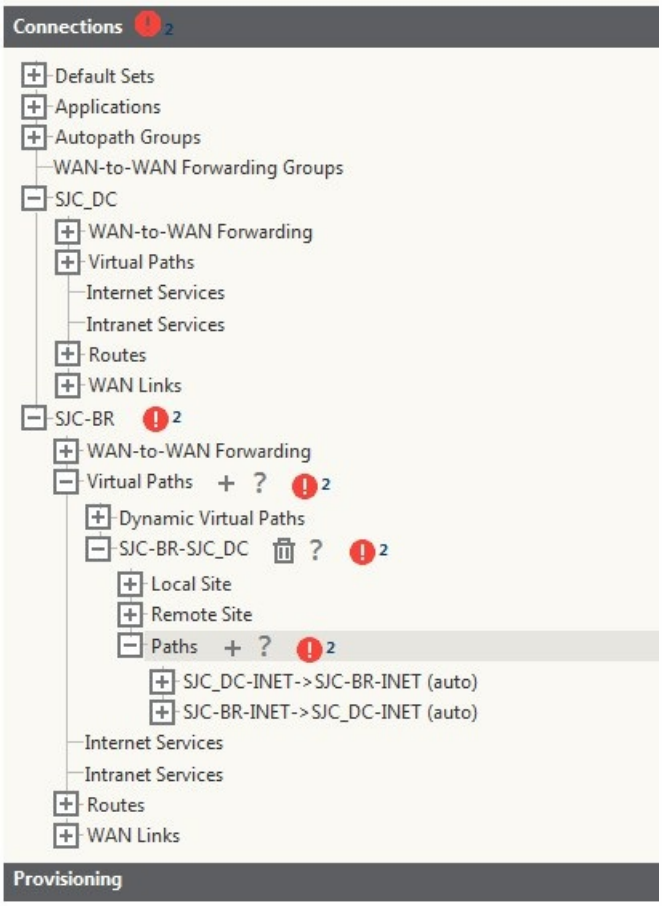

By default, the system will generate paths for WAN Links defined as access type Public Internet. You would be required to use the auto-path group function or enable paths manually for WAN Links with an access type of Private Internet. Paths for MPLS links can be enabled by clicking on the Add operator (in the green rectangle).

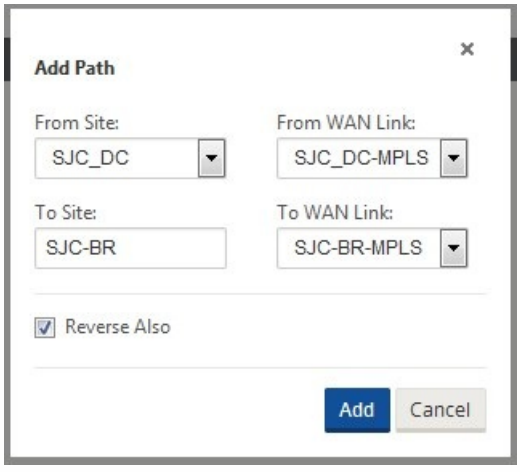

Create an Autopath Group

1. Click on the [+] sign next to Autopath Groups.
2. Configure the Autopath Group created as per requirement and click Apply.

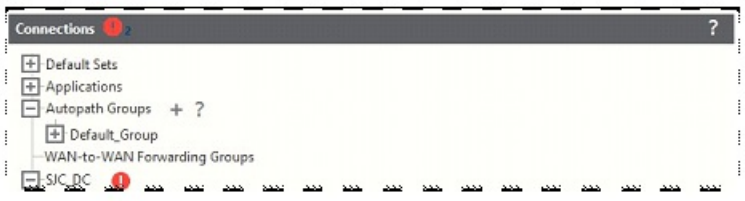

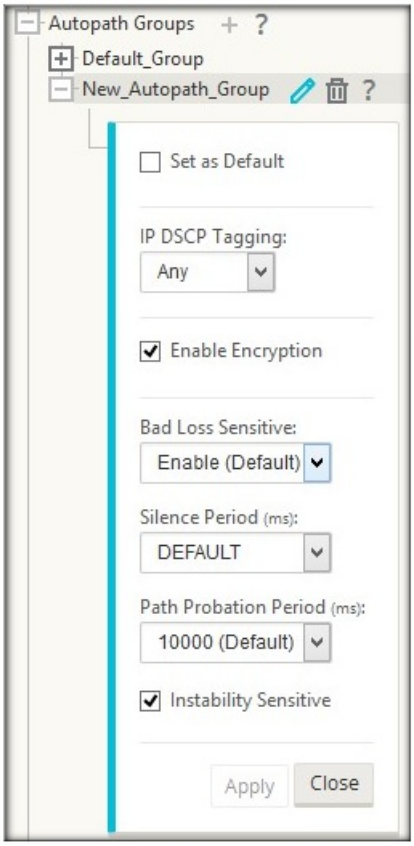

3. Rename the Autopath Group [Optional].

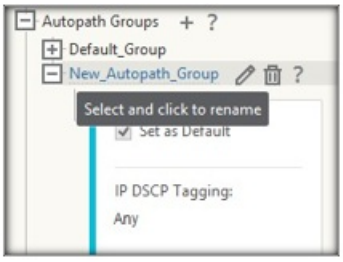

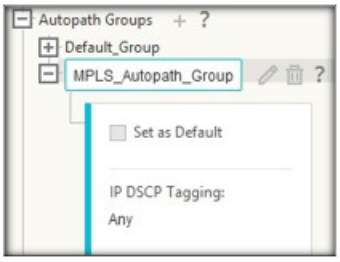

4. Map the Autopath Group to the Virtual Paths of Intranet WAN links at respective sites.

No two Autopath Groups can be marked as default. If marked would lead to an Audit Error.

After mapping the Autopath Group to the Virtual Paths of Intranet WAN, the paths should be automatically

populated (highlighted).

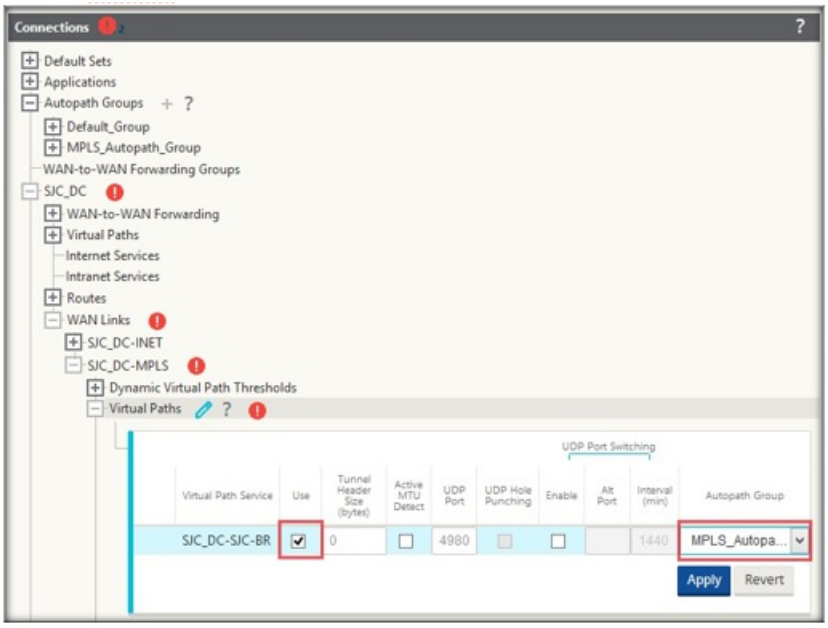

Manually add WAN links with access type Private Intranet

1. Select the Virtual Paths under WAN Links for respective sites and no Autopath Group would be mapped.
2. Click the [+] sign next to Paths to add Virtual Paths manually.

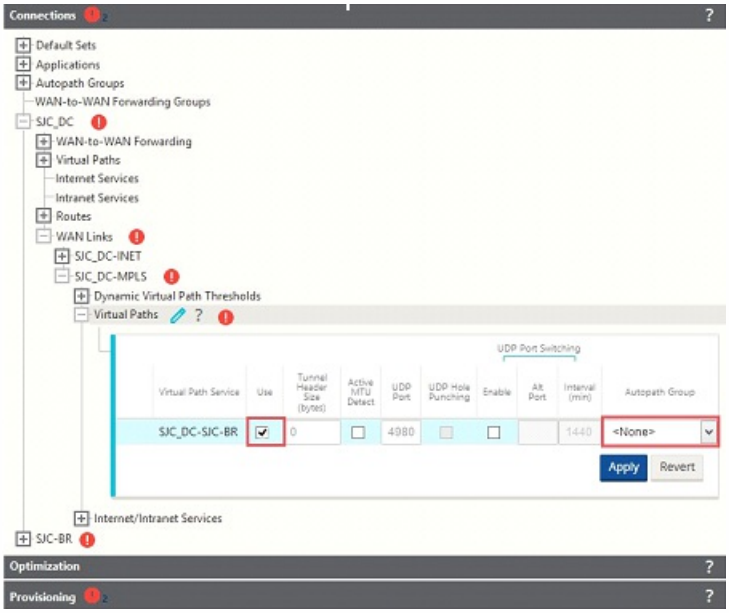

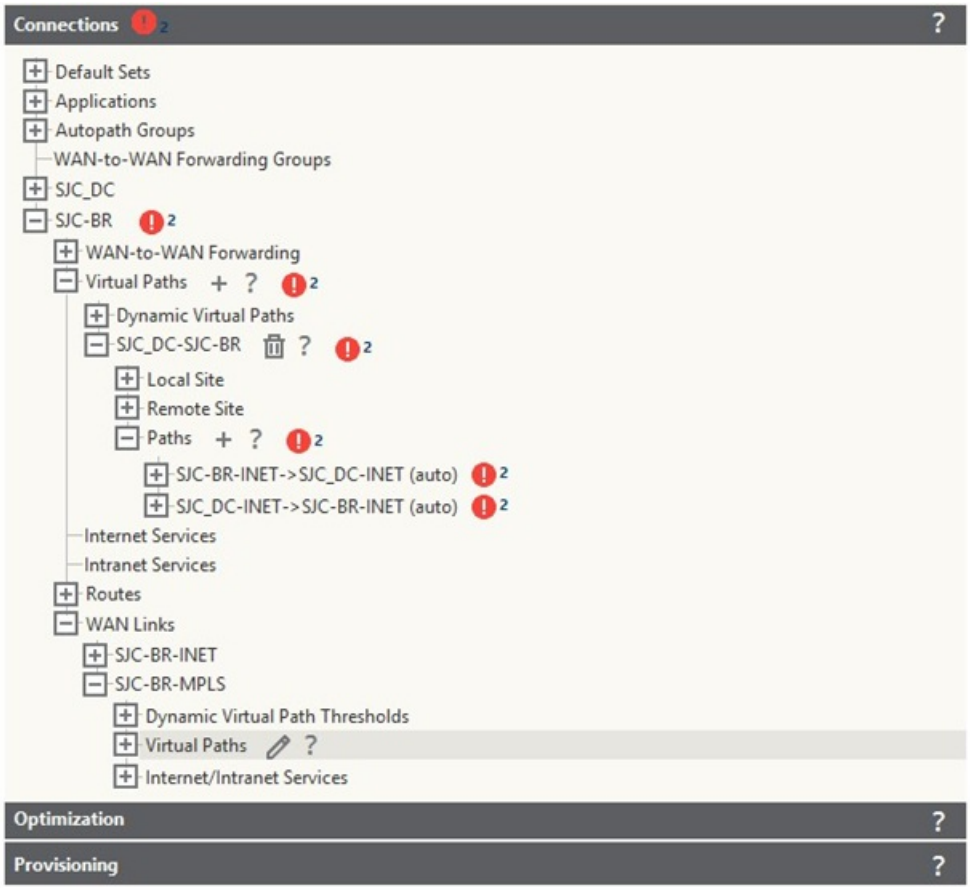

3. Select the Virtual Paths WAN Links for each site.

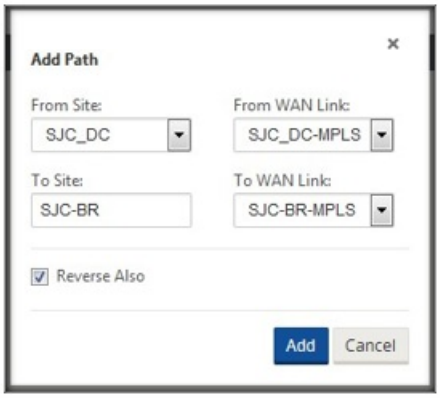

After manually adding the virtual paths for WAN links with access type Private Intranet, it gets populated under Paths (highlighted).

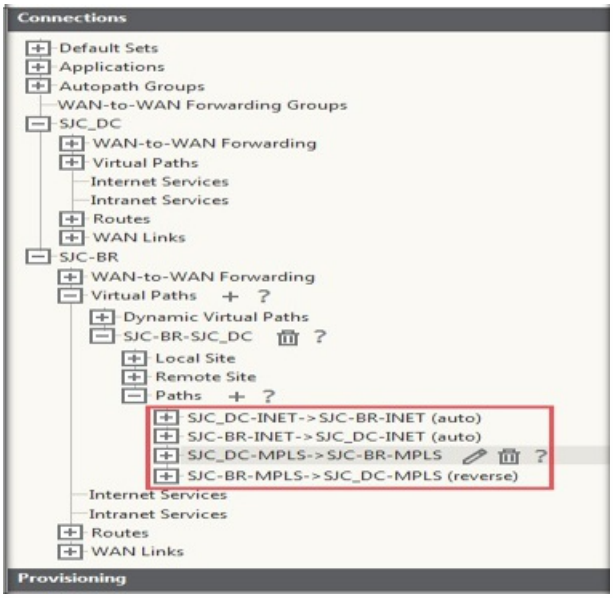

After completing all the above steps, proceed to Preparing the SD-WAN Appliance Packages on the MCN topic.

Policy Based Routing configuration on the PBR Router

Interface connected to the LAN

- Router# configure terminal
- Router(config)# interface FastEthernet0/1
- Router(config-if)# description ToLAN
- Router(config-if)# ip address 10.10.11.1 255.255.255.0
- Router(config-if)# duplex auto
- Router(config-if)# speed auto

Interface connect to the MPLS WAN Link

- Router# configure terminal
- Router(config)# interface GigabitEthernet0/0
- Router(config-if)# description To-MPLS-WAN
- Router(config-if)# ip address 10.20.0.2 255.255.255.0
- Router(config-if)# duplex auto
- Router(config-if)# speed auto

Interface connected to the INET WAN Link

- Router# configure terminal
- Router(config)# interface GigabitEthernet0/2/0
- Router(config-if)# description To-INET-WAN
- Router(config-if)# ip address 10.19.0.2 255.255.255.0
- Router(config-if)# duplex auto
- Router(config-if)# speed auto

Note: Interface GigabitEthernet0/1 on the PBR router is connected to the SD-WAN port 1/1, it is in 1-arm mode and this one port will serve traffic for MPLS and INET links.

- Router# configure terminal
- Router(config)# interface GigabitEthernet0/1
- Router(config-if)# description To-SDWAN-link
- Router(config-if)# ip address 192.168.1.1 255.255.255.0

Static Route Configuration (Route to the client/remote subnets):

- MPLS 10.17.0.0/24 via next hop WAN router MPLS 10.20.0.1
- INET 10.18.0.0/24 via next hop WAN router/FW INET 10.19.0.1

- Router# configure terminal
- Router(config)# ip route 10.17.0.0 255.255.255.0 10.20.0.1
- Router(config)# ip route 10.18.0.0 255.255.255.0 10.19.0.1

Route Map Definition

Access Control List Configuration:

Configure ACL's to define the traffic to be sent to and from the SD-WAN appliance.

1- From LAN to SD-WAN Appliance

As per topology, the LAN subnets are 10.10.11.0/24, 10.10.12.0/24, 10.10.13.0/24, etc. To send traffic from LAN to the SD-WAN, configure a unidirectional ACL (from LAN to any).

- Router# configure terminal
- Router(config)# ip access-list extended server_side
- Router(config)# permit ip 10.10.0.0 0.0.255.255 any

2- From SD-WAN Appliance to physical WAN Links

- Router# configure terminal
- Router(config)# ip access-list extended MPLS_Link
- Router(config)# permit ip 192.168.1.10 0.0.0.0 any
- Router# configure terminal
- Router(config)# ip access-list extended INET_Link
- Router(config)# permit ip 192.168.1.11 0.0.0.0 any

Route Map Configuration:

Define the route-map matching the ACL's.

Route map for LAN traffic:

Next hop will be any of SD-WAN Virtual IP's (VIP).

MPLS VIP 192.168.1.10

INET VIP 192.168.1.11

In this case, we chose MPLS VIP 192.168.1.10 as next hop and also added a health check to make sure if the SD-WAN fails, traffic is not routed to it.

- Router# configure terminal
- Router(config)# route-map server_side_VW_PBR permit 10
- Router(config-route-map)# match ip address server_side
- Router(config-route-map)# set ip next-hop verify-availability 192.168.1.10 10 track 123

Note: The above command configures the route map to verify the reachability of the tracked object. The tracking process provides the ability to track individual objects, such as ICMP ping reachability, routing adjacency, an application running on a remote device, a route in the Routing Information Base (RIB) or to track the state of an interface line protocol.

Route map for WAN traffic:

Next hop will be MPLS Router and Firewall for respective WAN links.

- Router# configure terminal
- Router(config)# route-map WAN_VW_PBR permit 20
- Router(config-route-map)# match ip address MPLS_Link
- Router(config-route-map)# set ip next-hop verify-availability 10.20.0.1 20 track 124
- Router# configure terminal
- Router(config)# route-map WAN_VW_PBR permit 30
- Router(config-route-map)# match ip address INET_Link
- Router(config-route-map)# set ip next-hop verify-availability 10.19.0.1 30 track 125

Apply the Route Map to the interface:

Router# configure terminal

- Router(config)# interface FastEthernet0/1
- Router(config-if)# ip policy route-map server_side_VW_PBR
- Router(config-if)# duplex auto
- Router(config-if)# speed auto
- Router# configure terminal
- Router(config)# interface GigabitEthernet0/1
- Router(config-if)# ip policy route-map WAN_VW_PBR
- Router(config-if)# duplex auto
- Router(config-if)# speed auto

MPLS Router Configuration (Gateway 10.20.0.1)

- Add route on MPLS router to reach MPLS VWAN VIP on the Data Center.
- MPLS VIP subnet 192.168.1.0/24 via next hop PBR router MPLS link 10.20.0.2
- Router# configure terminal
- Router(config)# ip route 192.168.1.0 255.255.255.0 10.20.0.2

Firewall Configuration (Gateway 10.19.0.1)

Add route on Firewall to reach INET VWAN VIP on the Data Center.

INET VIP subnet 192.168.1.0/24 via next hop PBR router INET link 10.19.0.2

- Router# configure terminal
- Router(config)# ip route 192.168.1.0 255.255.255.0 10.19.0.2

Building a SD-WAN Network

Aug 09, 2017

To build an SD-WAN overlay network without the need to statically build SD-WAN overlay route tables:

1. Create a WAN Path tunnel across each WAN link between two SD-WAN appliances.

2. Configure Virtual IP to represent the endpoint for each WAN link. You can establish encrypted WAN paths through the current L3 Network.

3. Aggregate 2, 3, and 4 WAN paths (physical links) into a single Virtual Path allowing packets to traverse the WAN utilizing the SD-WAN overlay network instead of the existing underlay which is least intelligent and cost inefficient.

SD-WAN Routing Components and Network Topology

- Local – subnet resides at this site (advertised to SD-WAN environment)
- Virtual Path – sent through Virtualized Path to the selected site appliance
- Intranet – sites with no SD-WAN appliance
- Internet – internet bound traffic
- Pass-through – untouched traffic, in one bridge interface out the other
- Default route (0.0.0.0/0) defined. Used for pass-through traffic not captured by the SD-WAN overlay route table, or utilized at the MCN to instruct clients sites to forward all traffic back to MCN node for back-haul of internet traffic.

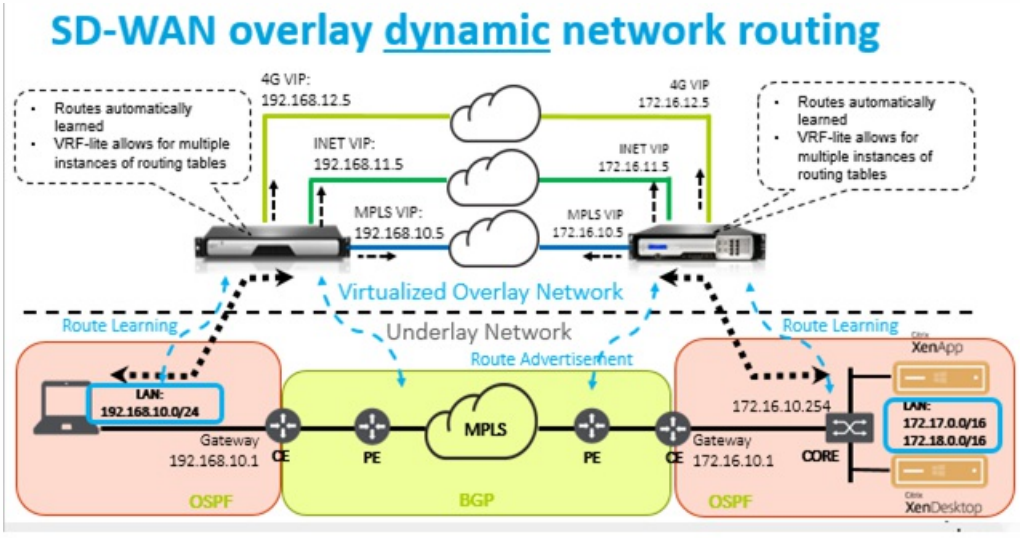

Dynamic Paths for Branch to Branch Communication

Aug 09, 2017

With demand for VoIP and video conferencing, the traffic is increasingly moving between offices. It is inefficient to set up full mesh connections through datacenters which can be time consuming.

With NetScaler SD-WAN, you do not need to configure paths between every office. You can enable the Dynamic Path feature and the SD-WAN solution automatically creates paths between offices on demand. The session initially uses an existing fixed path. And as bandwidth and time threshold is met, a path is created dynamically if that new path has better performance characteristics than the fixed path. Session traffic is transmitted through the new path. This results in efficient usage of resources. Paths exist only when they are needed and reduce the amount of traffic getting transmitted to and from the datacenter.

Additional benefits of SD-WAN network include:

- Bandwidth and PPS thresholds to allow branch to branch connections
- Reduce bandwidth requirements in and out of data center while minimizing latency
- Paths created on demand depend on set thresholds
- Dynamically release network resources when not required
- Reduce load on the Master Control Node and latency

Branch to Branch Communication Using Dynamic Paths

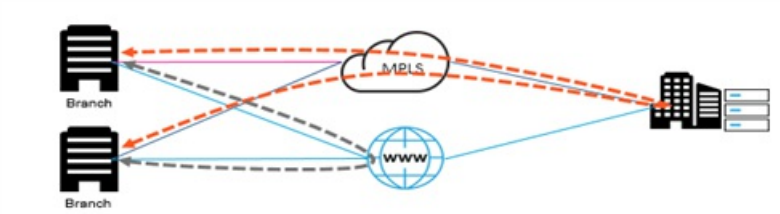

SD-WAN Network with Dynamic Path

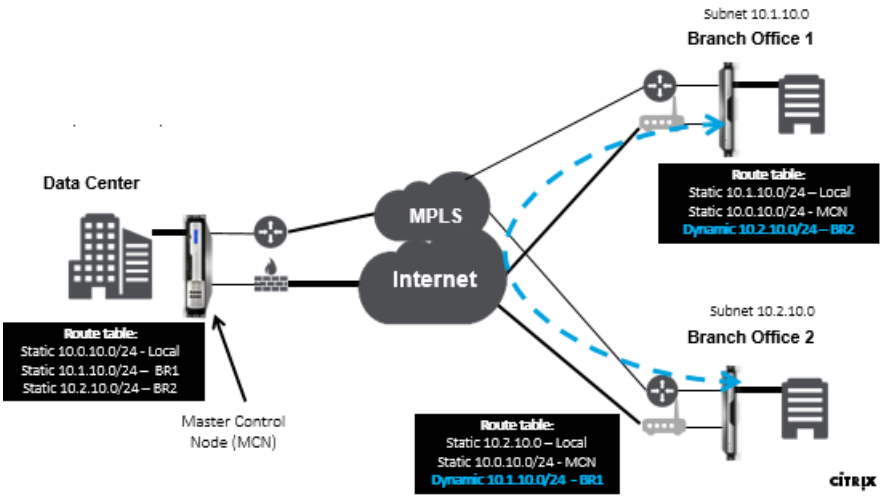

- Dynamic virtual paths are used for large scale deployments, such as Enterprises

- Smaller deployments use Static virtual paths and any-to-any virtual paths
- Always use Static virtual paths between two Data Centers (DC to DC)
- Not all WAN paths need to be configured for using Dynamic virtual path
- Each SD-WAN appliance has limited number of Dynamic virtual paths (8 dynamic lowest limit, 8 static lowest limit = total 16) that can be configured.

How to Enable Dynamic Virtual Path in the SD-WAN GUI

To enable dynamic virtual paths:

1. In the NetScaler SD-WAN GUI, under the **Connections** pane, create a WAN to WAN Forwarding Group.
2. Navigate to **Connections → [Client Site Name] → WAN to WAN Forwarding**.

    a) Enable **WAN to WAN Forwarding** to enable the site to serve as a proxy for multi-hop site to site.

    b) Enable **Site as Intermediate Node**

3. Navigate to **Connections → Remote Site→ WAN to WAN Forwarding**.

    a) Enable WAN to WAN Forwarding to enable the site to serve as a proxy for multi-hop site to site.

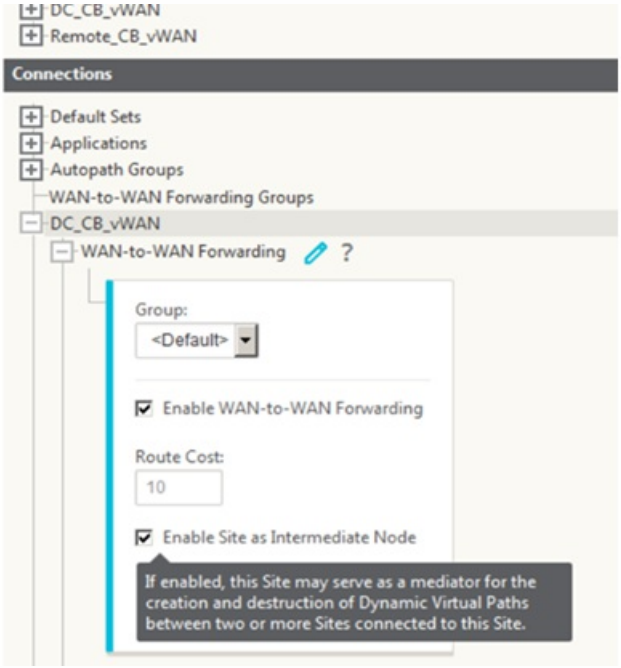

4. Navigate to **Connections → Remote Site → Virtual Path → Dynamic Virtual Path**.

    a) Enable **Dynamic Virtual Paths**.

    b) Set the maximum number of dynamic paths.

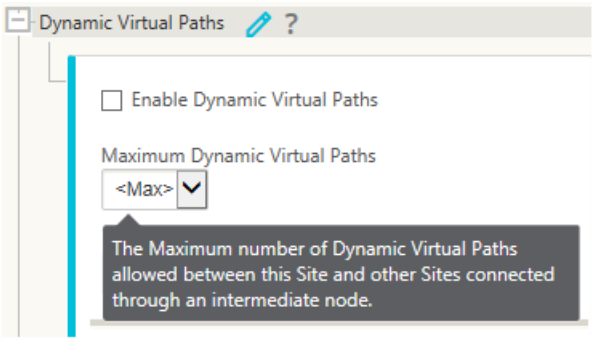

How to Create a Dynamic Virtual Path

- Configuration determines when a Dynamic Virtual Path is active or down.
- Configure sample packet count (pps) or bandwidth (kbps) within a timeframe.
- Can be set Globally or with WAN Link configured at the Intermediate Node.

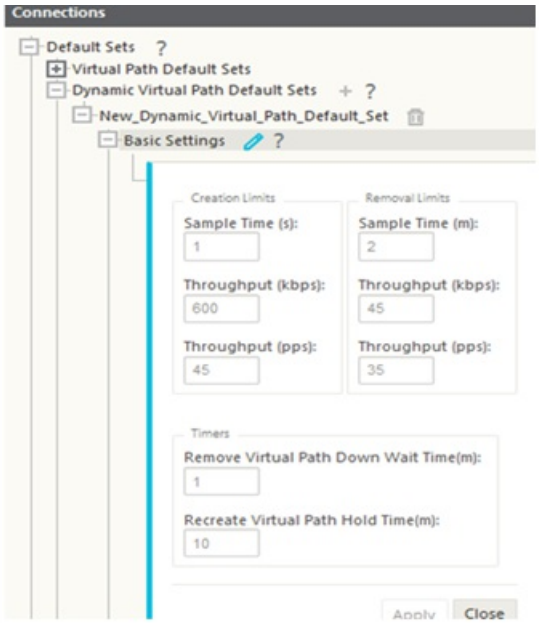

Configuring Static WAN Paths

Aug 09, 2017

With WAN to WAN enabled on the MCN, remote site routes are advertised by MCN.

- Clients are aware of MCN local routes as well as other client site routes
- From client perspective, all routes are considered as MCN routes

When WAN-to-WAN forwarding is not enabled on the MCN, Branch to Branch communication issues are encountered in the customer network.

SD-WAN appliances running in client mode are not aware of other branches subnets until WAN-to-WAN forwarding is enabled on MCN. Once this option is enabled, branch SD-WAN nodes become aware of other branch subnets and all the traffic destined to other branches is forwarded to MCN. MCN routes it to the correct destination.

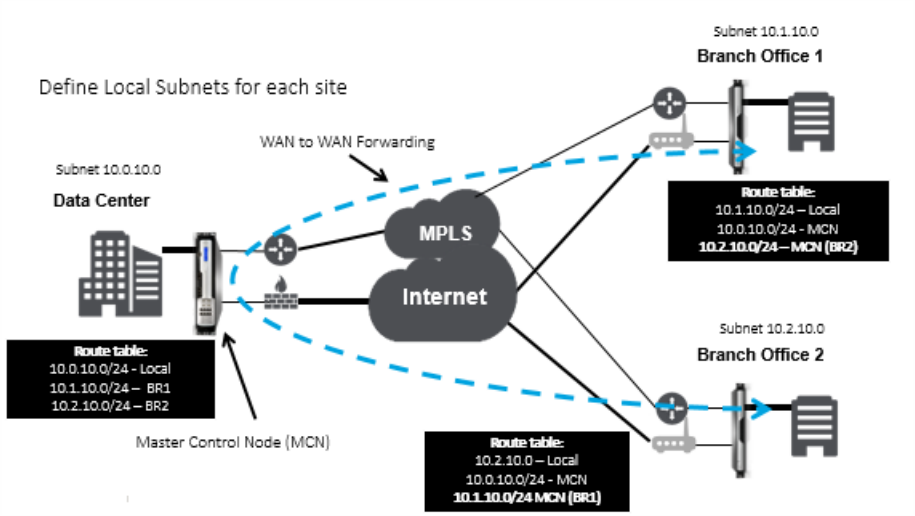

Routing Support for LAN Segmentation

Aug 09, 2017

The SD-WAN Standard and Enterprise Edition appliances implement LAN segmentation across distinct sites where either appliance is deployed. The appliances recognize and maintain a record of the LAN side VLANs available, and configure rules around what other LAN segments (VLANs) can connect to at a remote location with another SD-WAN Standard or Enterprise Edition appliance.

The above capability is implemented through the use of a Virtual Routing and Forwarding (VRF) table that is maintained in the SD-WAN Standard or Enterprise Edition appliance, which keeps track of the remote IP address ranges accessible to a local LAN segment. This VLAN-to-VLAN traffic would still traverse the WAN through the same pre-established Virtual Path between the two appliances (no new paths need to be created).

An example use case for this functionality is that a WAN administrator may be able to segment local branch networking environment through a VLAN, and provide some of those segments (VLANs) access to DC-side LAN segments that have access to the internet, while others may not obtain such access. The configuration of the VLAN-to-VLAN associations is achieved through the MCN's Configuration Editor in the SD-WAN management web interface.

Utilizing Enterprise Edition Appliance to Provide WAN Optimization Services Only

Aug 09, 2017

The SD-WAN Enterprise Edition appliances contain fully featured WAN Optimization functionality in addition to WAN Virtualization. Some customers prefer to implement WAN Optimization functionality before migrating to SD-WAN services. This deployment use case provides the steps to utilize Enterprise Edition appliances to utilize WAN optimization services.

NetScaler SD-WAN Product Platform Editions include the following appliances:

- SD-WAN: SD-WAN Standard Edition appliance
- Enterprise: SD-WAN Enterprise Edition appliance
- WANOP: SD-WAN WANOP Edition appliance

To integrate Enterprise Edition appliances into an existing distributed WANOP network, you need to configure SD-WAN (Physical or Virtual) appliance at the DC site as the MCN. The SD-WAN appliance manages all configuration of the network. A Virtual Path is established between the Branch site and MCN at the DC site. This Virtual Path is only used for sending control traffic between the appliances. At the branch appliance, the data traffic is processed as an intranet service. The intranet traffic is not encapsulated and traverses over existing WAN link to reach the DC site. A WANOP appliance at the DC site needs to be in the traffic path to provide end-to-end traffic optimization.

For customer sites that do not have SD-WAN hardware appliance at the head-end, VPX appliances in a HA pair (two Virtual WAN VPXs) can be used as MCN in one-arm mode. For the one-arm mode, PBR rules on the third-party router are required to redirect traffic to the SD-WAN appliance.

This document assumes that the DC site appliances are deployed in HA mode for redundancy. However, note that HA mode is not mandatory for this deployment

Prerequisites

- A pair of WANOP appliances and a pair of SD-WAN appliances deployed in HA mode at the DC site.
- An Enterprise Edition appliance at the Branch site.

Network Topology

SD-WAN Standard Edition and WANOP Appliances in PBR Deployment

In the below illustration, both the SD-WAN SE and WAN OP appliances at the DC site are deployed in one-arm mode. The SD-WAN appliance supports PBR deployment while the WANOP appliance supports both PBR and WCCP. The control traffic (Virtual Path traffic) received from WAN at the DC site will be redirected to the SD-WAN appliance by the PBR Router. The data traffic will be redirected to WAN Optimization appliance by the PBR Router.

Traffic flow for WAN to DC LAN:

- CE (Customer Edge) Router -> PBR Router -> SD-WAN -> PBR Router -> LAN
- CE (Customer Edge) Router -> PBR Router - > WAN OPT - > PBR Router- > LAN

The same traffic flow will be followed in the reverse direction.

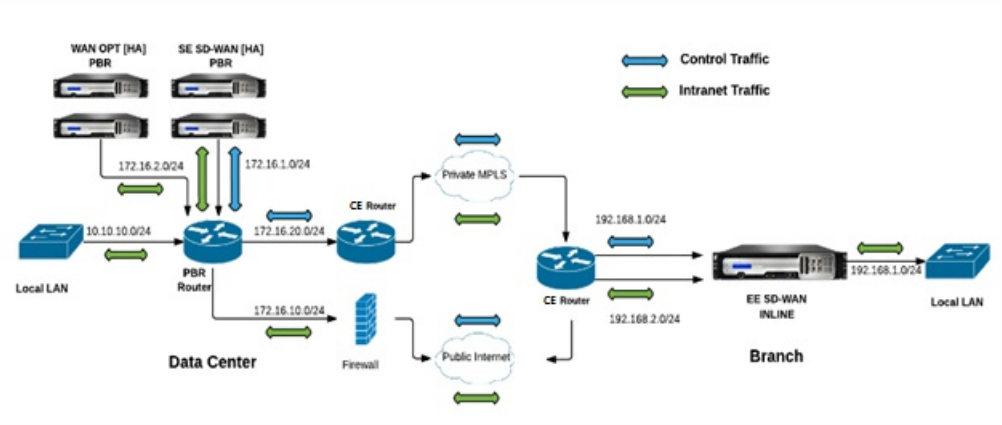

SD-WAN Standard Edition in PBR mode and WANOP in Inline Deployment

In the below illustration, the SD-WAN appliance at the DC site is deployed in one-arm mode while the WANOP appliance is deployed in inline mode.

The control traffic (Virtual Path traffic) received from WAN at the DC site will be redirected to the SD-WAN appliance by the PBR Router. The data traffic will be forwarded to WAN Optimization appliance (inline) by the PBR Router.

Traffic flow for WAN to DC LAN:

- CE (Customer Edge) Router - > PBR Router - > SD-WAN -> PBR Router - > LAN
- CE (Customer Edge) Router - > PBR Router -> WAN OPT - > LAN

The same traffic flow will be followed in the reverse direction.

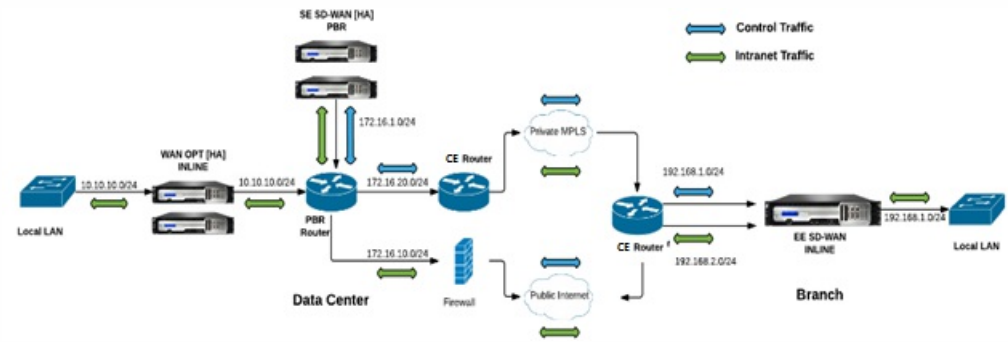

Configuration Steps

1. Configure the SD-WAN Appliance at DC [MCN] to establish Virtual Paths between DC and Branch sites.

   See, http://docs.citrix.com/en-us/netscaler-sd-wan/9-1/configuration-topics/configuring-virtual-path-service-between-mcn-client-sites.html

2. Configure Intranet Service at the DC site.

   a) On the MCN (DC site), go to **Configuration** > **Virtual WAN** > **Configuration Editor** > **Connections** > **Site (DC)** > **Intranet Services**. Click the **[+]** sign to add an Intranet Service.

   b) Select the **WAN Link(s)** for Intranet Service, and then click **Apply**.

c) Navigate to Routes under the same **Site (DC)**, click **[+]** sign to add the remote network with cost lower than 5, and select click **Add**.

For example - Enter **192.168.1.0/24** in the **Network IP address** field with cost 4 and select **Service Type** as **Intranet**.

> **Note**
>
> Cost at each site should be less than 5 for the intranet route to take precedence.

3. Configure Intranet Service at the Branch site.

    a) Repeat sub-steps a to c from **step 2** above on the Branch site.

For example - Enter **172.16.1.0/24** in the Network IP address field with cost 4 and select **Service Type** as **Intranet**.

4. Perform **Change Management** to upload and distribute configuration to the Branch site.

    See, Exporting configuration package and change management

    By default, the traffic is sent from Branch to DC through the Virtual Path.

> **Note**
>
> The PBR router needs to be configured to redirect traffic as per the deployment steps provided.

For more information about configuring WAN Optimization, refer to the CloudBridge 9.1 documentation at: Enabling-configuring-wan-optimization

SD-WAN SE/EE Appliance in Hairpin Deployment Mode

Aug 09, 2017

With a hairpin deployment, you can implement use of a Remote Hub site for internet access through backhaul or hairpin when local internet services are unavailable or are experiencing a slow traffic. You can leverage high bandwidth routing between client sites by allowing backhauling from specific sites.

The purpose of a hairpin deployment from a non-WAN to a WAN forwarding site is to provide customers with a more efficient deployment process and more streamlined technical implementation. Customers will have the ability to use a remote hub site for internet access when needs arise, and can route flows through the virtual path to the SD-WAN network.

For example, consider an administrator with multiple SD-WAN Sites, A and B. Site A has poor internet service. Site B has usable internet service, with which you want to backhaul traffic from site A to site B only. You can try to accomplish this without the complexity of strategically weighted route costs and propagation to sites that should not receive the traffic.

Also, the route table is not shared across all sites in a Hairpin deployment. For example, if traffic is hairpin'ned between Site A and Site B through Site C, then only Site C would be aware of site A's and B's routes. Site A and Site B will not share each other's route table unlike in WAN-to-WAN forwarding.

When traffic is Hairpin'ned between Site A and Site B through Site C, the static routes are required to be added in Site A and Site B indicating that the next hop for both the sites is the intermediate Site C.

WAN-to-WAN Forwarding and Hairpin deployment have certain differences, namely:

a. Dynamic Virtual Paths are not configured. At all times, the intermediate site will see all the traffic between the two sites.

b. Does not participate in WAN-to-WAN Forwarding groups.

WAN-to-WAN Forwarding and Hairpin deployment are mutually exclusive. Only one of them can be configured at any given point in time.

NetScaler SD-WAN SE/EE and VPX (virtual) appliances support hairpin deployment. You can now configure a 0.0.0.0/0 route to hairpin traffic between two locations without affecting any additional locations. If hairpinning used for intranet traffic, specific Intranet routes are added to the client site to forward intranet traffic through the virtual path to the hairpin site. Enabling WAN-to-WAN forwarding to accomplish hairpin functionality is no longer required.

You can configure hairpin deployment through the SD-WAN web management interface from the configuration editor.

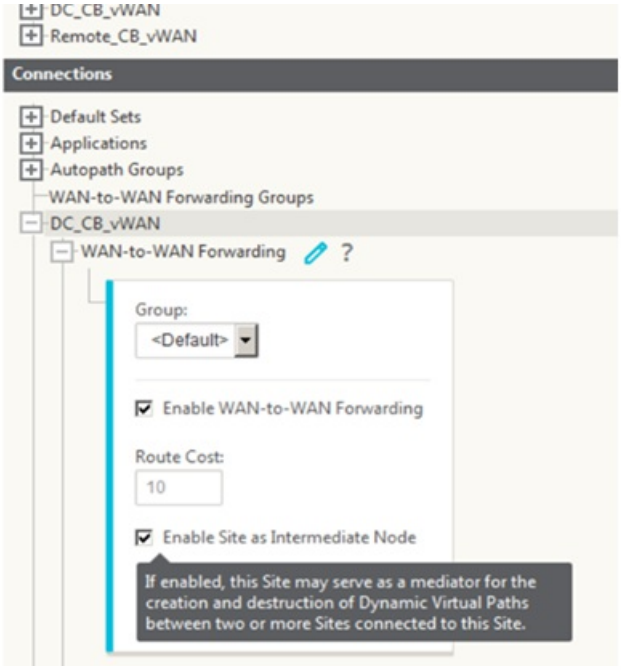

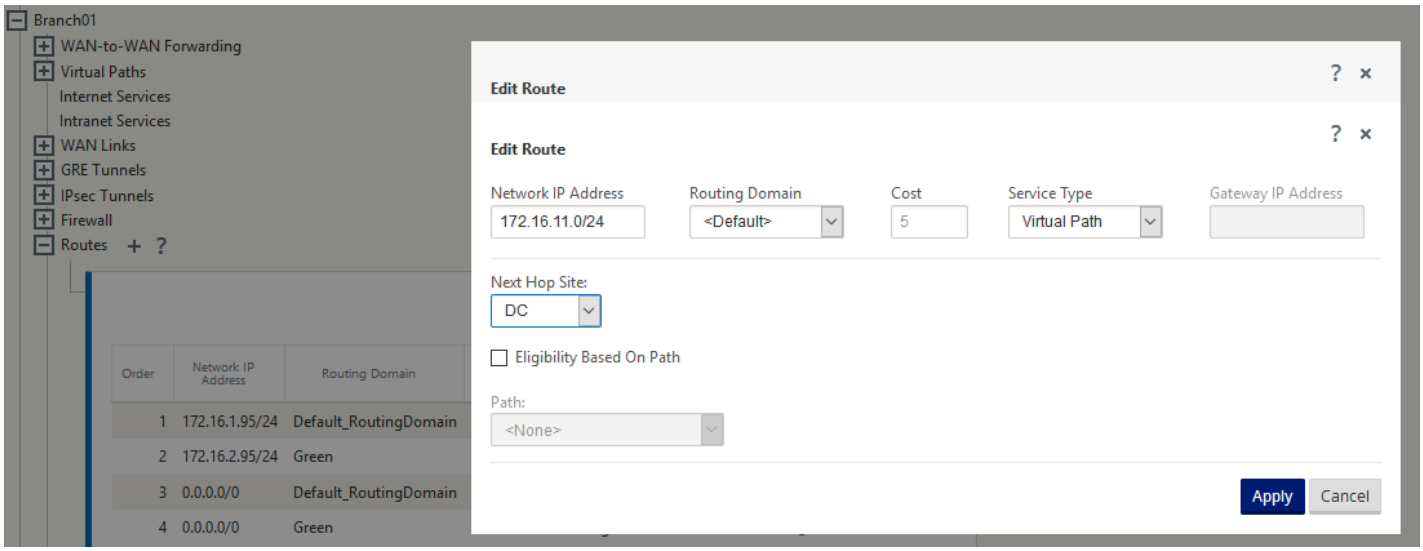

Two Box Mode

Jan 18, 2018

Two box mode is a WCCP one-arm based deployment where the SD-WAN SE appliance acts as a WCCP router and the SD-WAN WANOP (4000/5000/4100/5100) appliances act as WCCP clients and help establish WCCP convergence. This way all the virtual path/Intranet service oriented TCP packets reaching the SD-WAN SE appliance get redirected to the SDWAN-WANOP appliance for optimization benefits thereby providing both SD-WAN SE and WANOP benefits for the customer traffic.

Two Box mode is supported only on the following appliance models:

- SD-WAN SE appliances – 4000, 4100, and 5100
- SDWAN-WANOP appliances – 4000, 5000, 4100, and 5100

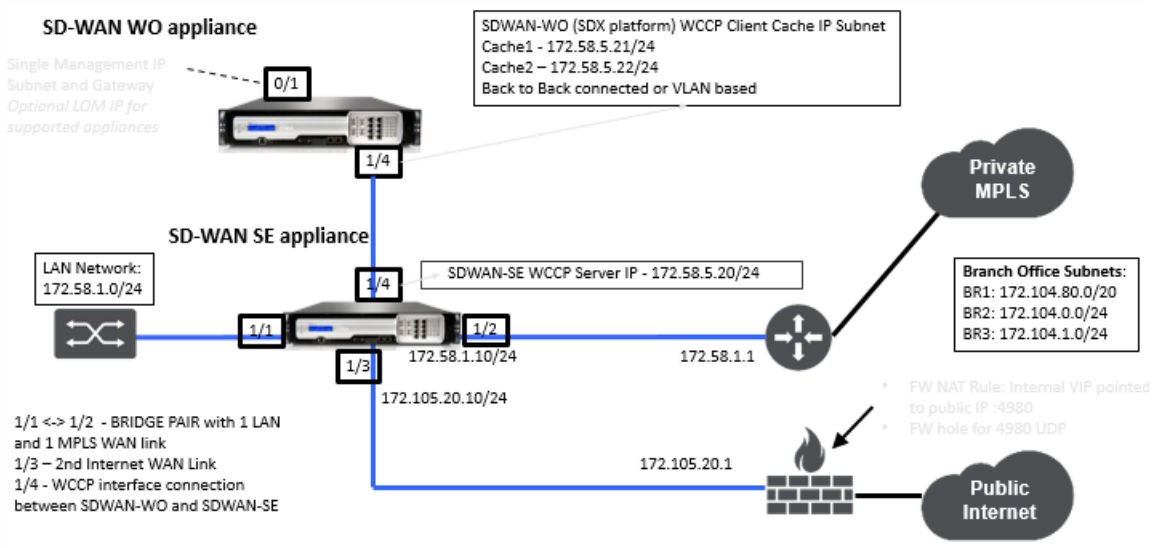

Note

High Availability and WCCP deployment modes are not accessible when Two Box mode is enabled. However, these deployment modes will be available for the user to administer.

Important

- Although the legacy WCCP deployment is disabled when Two Box Mode is enabled, the Service Group convergence can only be verified from the WCCP monitoring page. There is no separate GUI page under the monitoring section for the Two Box Mode.
- If WCCP process running on the Standard Edition appliance reboots multiple times within a short interval of time, for example; 3 times in a minute then Service Group shuts down automatically. In such scenario, to get the WCCP convergence on the WANOP appliance, re-enable the WCCP feature in the WANOP appliance web GUI.
- When there is a change in the WCCP configuration or WAN optimization related to configuration on the Standard Edition appliance, the external WANOP appliance reboots. For example, enabling/disabling the WCCP checkbox in the Interface Group of config editor followed by Change Management process, restarts the WANOP appliance as well.

Note

Also, note the following points to consider when implementing the two box mode:

- When a routing domain is selected to be redirected to the WANOP appliance from the Configuration Editor, it should be added in the Interface Group for which WCCP is enabled.
- The same routing domain's traffic should be selected on the partner site as well. For example; **MCN** > **Branch01** to observe WAN optimization benefits.
- If a routing domain is selected in the interface group on which WCCP is enabled, another interface group which contains the bridged interfaces should have the same routing domain configured. Only if WCCP enabled interface group has the routing domain configured it is not enough to transmit the end-to-end traffic flowing with WAN optimization benefits.

NetScaler SD-WAN Standard Edition

To configure two-box mode solution in the Standard Edition appliance at the DC or Branch site:

1. In the NetScaler SD-WAN SE web management interface, go to **Configuration** > **Virtual WAN** > **Configuration Editor**. Open an existing configuration package or create a new package.

2. In the chosen configuration package, go to the **Advanced** tab to view the configuration details.

3. Open **Global** settings and expand **Routing Domains** to view that the **Redirect to WANOP** checkbox is enabled.

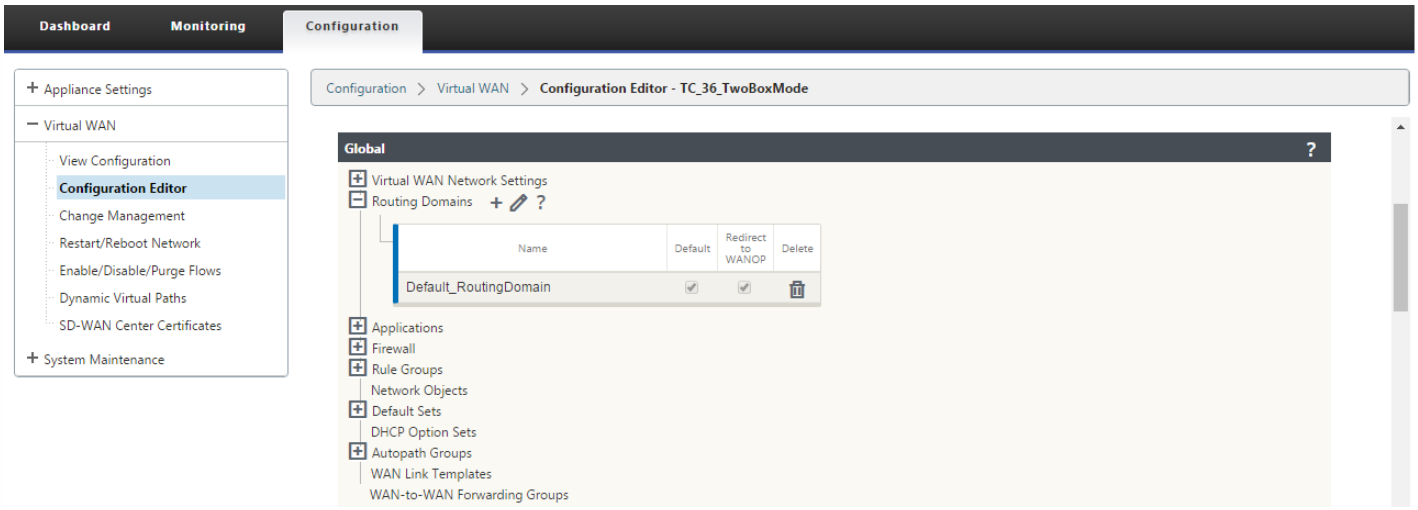

4. Expand DC to enable **WCCP** for the **Virtual Interface** under **Interface Group** settings that will signify which virtual network interface the appliance will be enabled for.

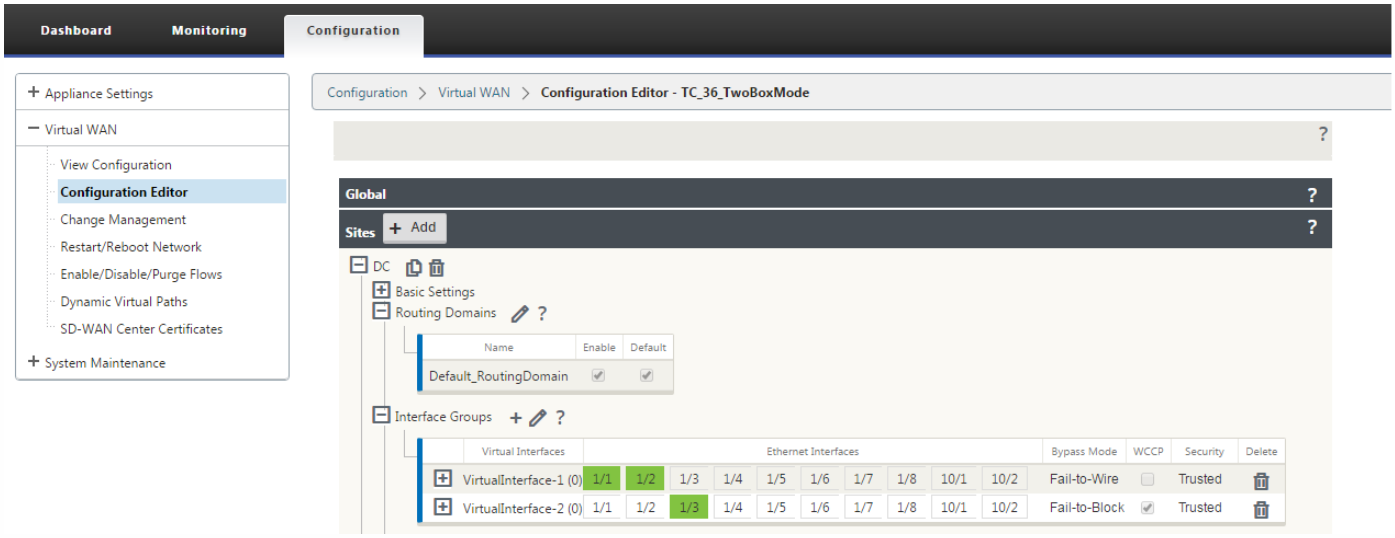

5. Expand **Sites+ Add** to view the Branch routing domain and interface group settings. Under the Branch site, the **Redirect to WANOP** checkbox is enabled for Routing Domains.

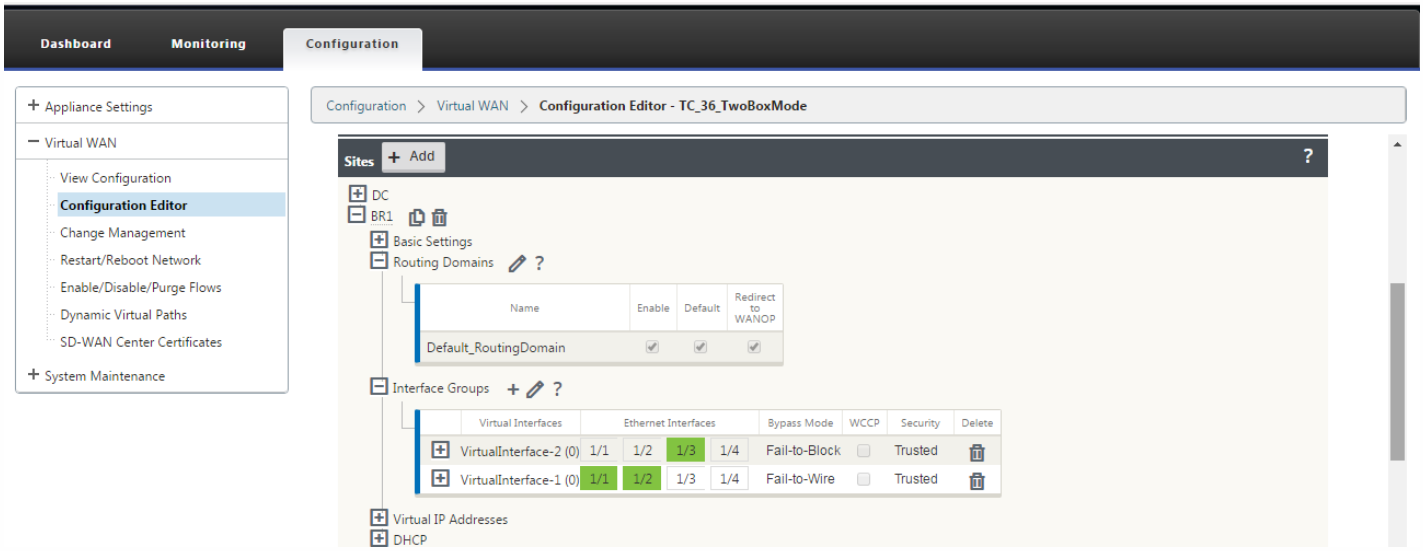

Note

The WCCP listener should be enabled only for those virtual network interfaces which have only ONE Ethernet Interface configured. This indicates that WCCP Listener should not be enabled on a BRIDGED Pair. It is intended to be enabled on the ONE-ARM interface between the SD-WAN SE and SD-WAN WANOP appliances.

NetScaler SD-WAN WANOP Configuration

To configure two-box deployment mode in the SD-WAN WANOP appliance web GUI:

1. In the NetScaler SD-WAN WANOP web management interface, go to **Configuration** > **Appliance Settings** > **Advanced Deployments** > **Two Box Solution**.

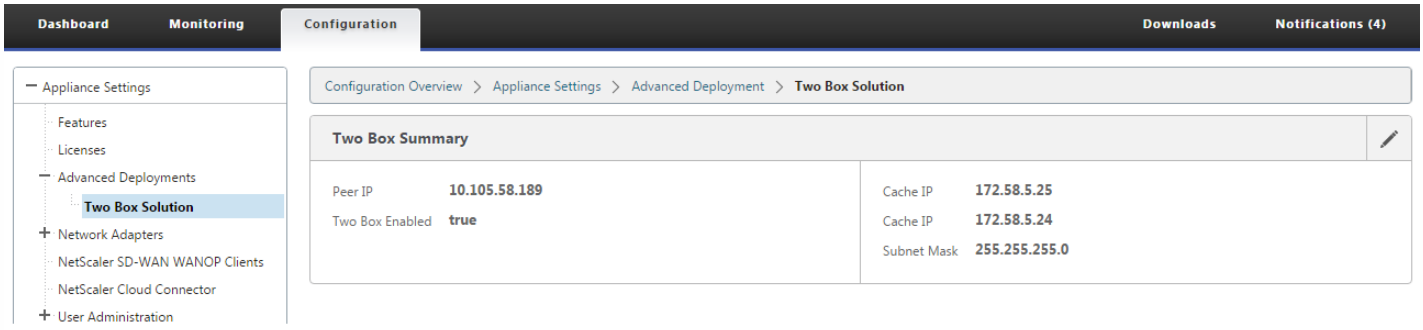

2. Click the Edit icon to edit the two box mode settings. Information dialog about **Cache IPs** is displayed. Click **O**K.

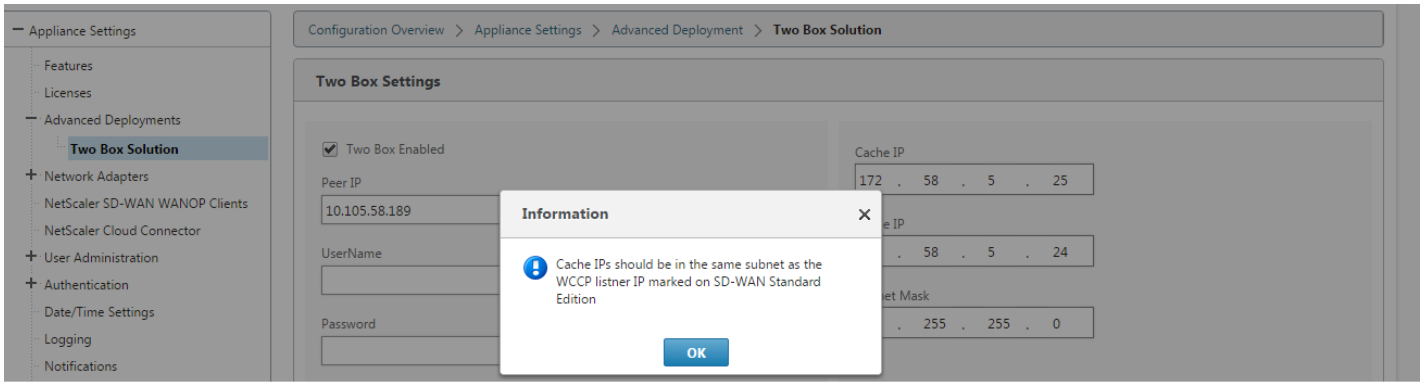

3. Enable the **Two Box Enabled** checkbox.

4. Enter the **Peer IP**. Peer IP is the Netscaler SD-WAN Standard Edition appliance IP address.

5. Enter the user credentials and click **Apply**.

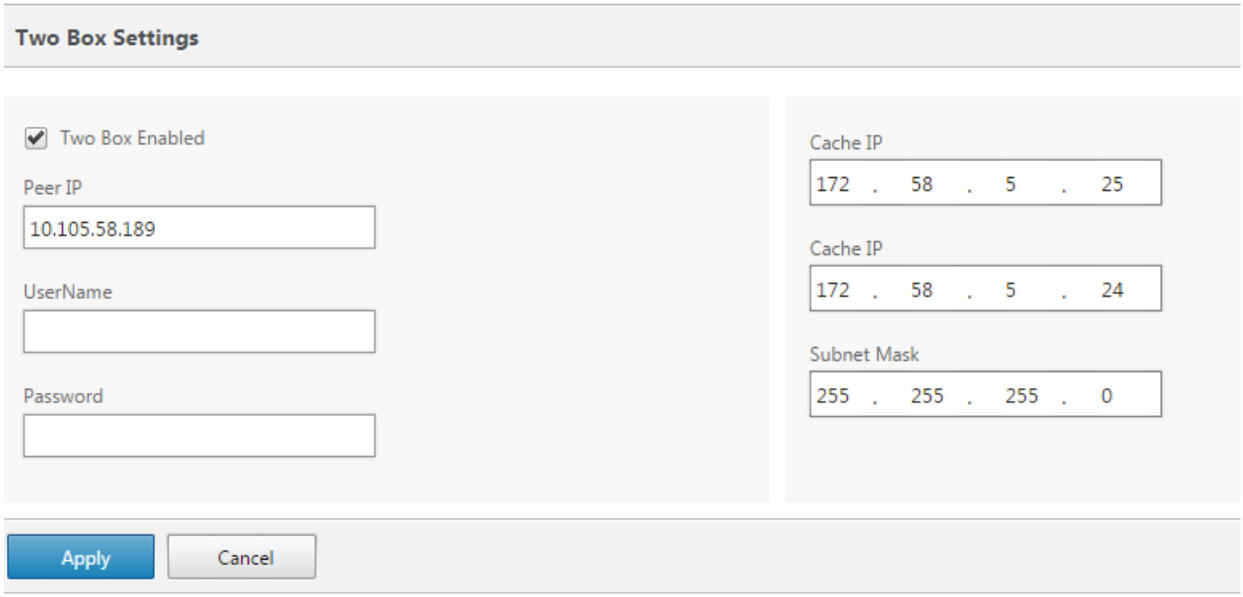

Two Box Mode Configuration and Manageability

Following are some of the two box mode configuration and manageability points to consider for deployment:

* SD-WAN WANOP configurations mentioned below can be configured from SD-WAN SE configuration editor as a unified pane

- SERVICE CLASS
- APPLICATION CLASSIFIER
- FEATURES
- SYSTEM TUNING

Monitoring

You can monitor SD-WAN WANOP traffic directly using the Monitoring page of the SD-WAN SE appliance's web UI. This allows for a single pane monitoring of both the SDWAN-SE and SDWAN-WO appliances while processing data traffic. You can view the connection details, secure partner details and so on under the WAN Optimization node in the SDWAN-SE UI.

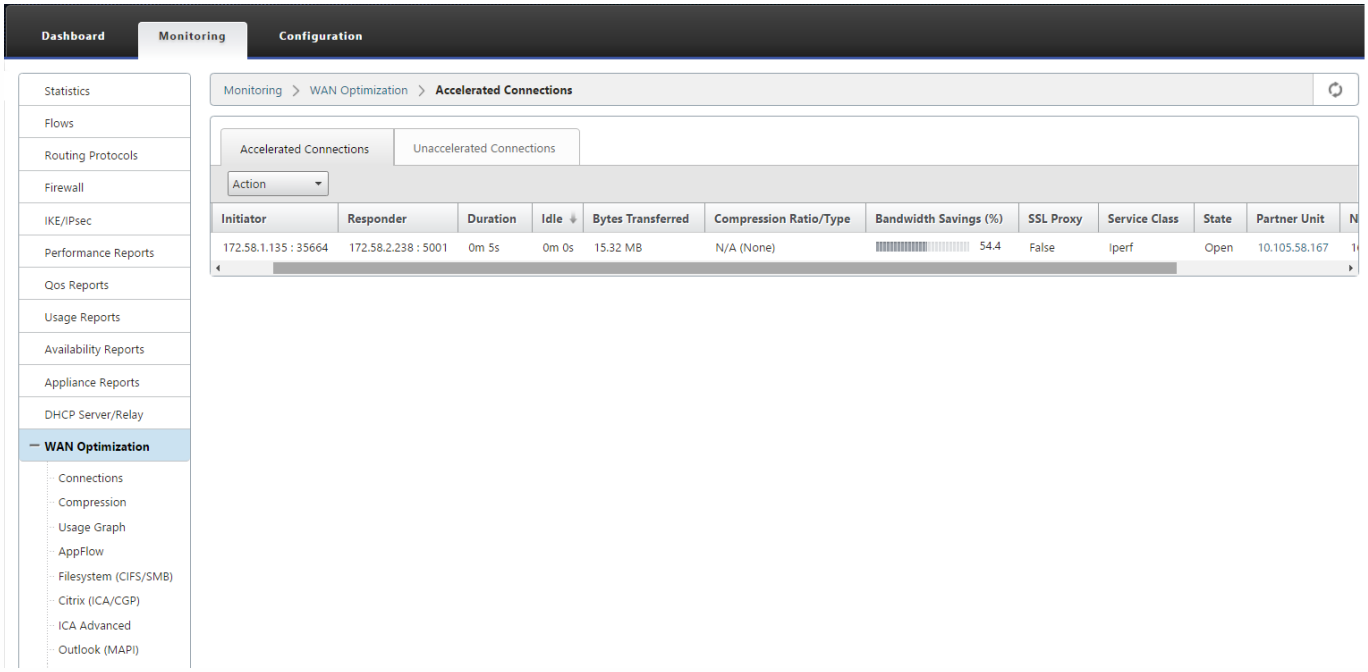

Configuration

You can configure APPFLOW directly from the SD-WAN-SE **Configuration** page under **APPFLOW** node. This enables SD-WAN-SE to act as a single pane for configuration of APPFLOW and other data processing configuration attributes such as Service Class, Application Classifiers and so on. The configuration done on the SD-WAN-SE reflects on the SD-WAN-WO configuration, maintaining seamless APPFLOW functionality support.

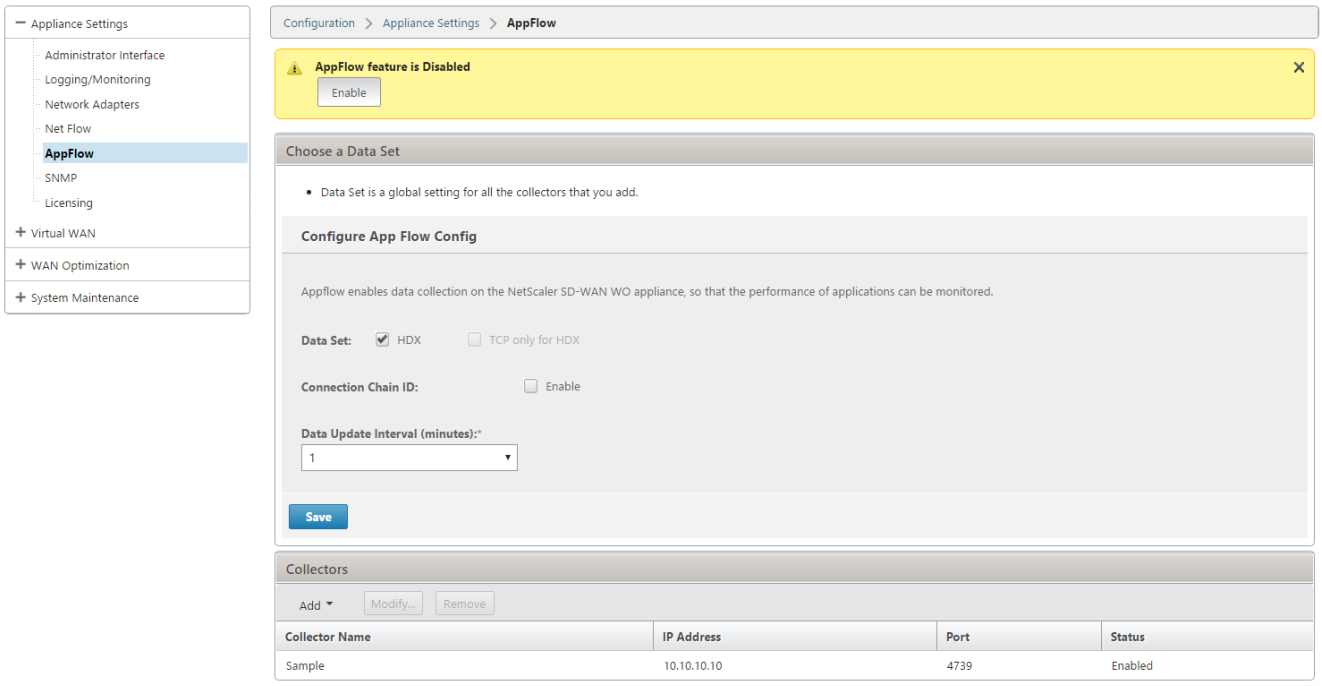

SD-WAN WANOP already discovered by an insight center, if used in Two Box Mode should be isolated and not configured using MAS/Insight until this mode is turned off. This is because the configuration of WANOP for traffic processing is governed by the SD-WAN SE appliance in the Two Box Mode.

Advanced Optimizations or Secure Acceleration should be directly configured on the SD-WAN-SE appliance like we would configure on the SD-WAN-WO appliance. This helps maintain a single pane of configuration of configurations like Domain Join or Secure Acceleration/SSL Profile creation for Advanced optimizations or SSL Proxy.

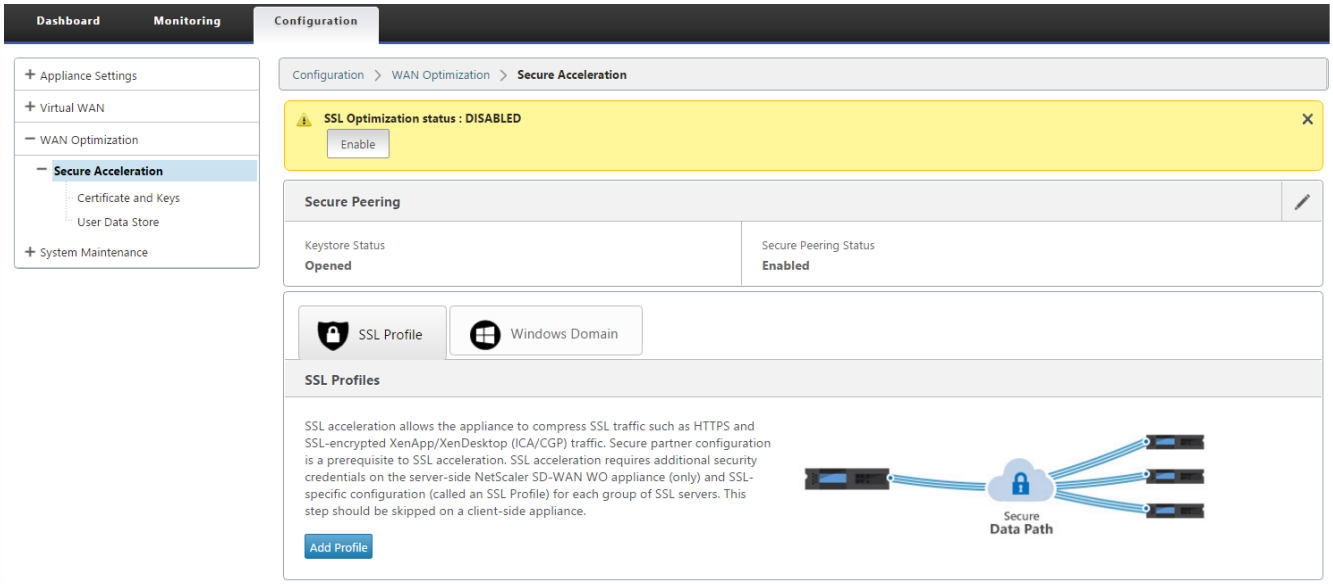

* Licensing should be separately managed for each of SD-WAN SE and SD-WAN WANOP appliances.

* Software Upgrade should be separately managed for each of of SD-WAN SE and SD-WAN WANOP appliances with the respective software packages. For example, tar.gz for SD-WAN SE and upgrade upg for SD-WAN WANOP.

* Data path integration should be configured between SD-WAN SE and External WANOP appliances through the WCCP deployment mode.

- At data path level both WCCP and Virtual WAN features are offered through data path integration between WANOP and SE externally in one-arm mode to obtain optimization benefits.

Unified Configuration and Monitoring

When you enable the two box mode with SD-WAN SE and SD-WAN-WANOP appliances, you can view the configuration in the SD-WAN SE appliance similar to how you can view two box configuration with the SD-WAN-EE appliance.

a. Go to **Configuration** > **Virtual WAN** > **WAN Optimization**

b. Appflow node under **Configuration** > **Appliance Settings**

c. WAN Optimzation node under Configuration.

This information is redirected from the SD-WAN WANOP appliance which is in Two box mode with the SD-WAN SE appliance.

Configuration related to WANOP, such as SSL Acceleration and AppFlow can now be performed from SD-WAN SE web GUI.

Traffic related statistics, such as Connections, Compression, CIFS/SMB , ICA Advanced, MAPI and partners can now be monitored from SD-WAN SE web GUI under **Monitoring** > **WAN Optimization** similar to the SD-WAN Enterprise edition appliance.

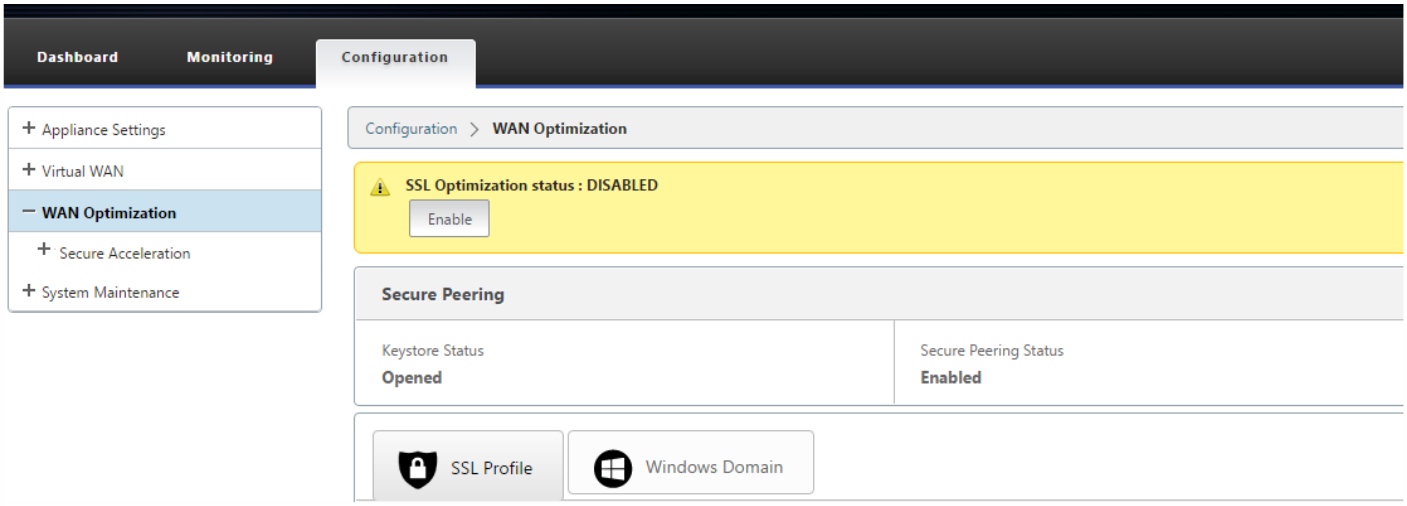

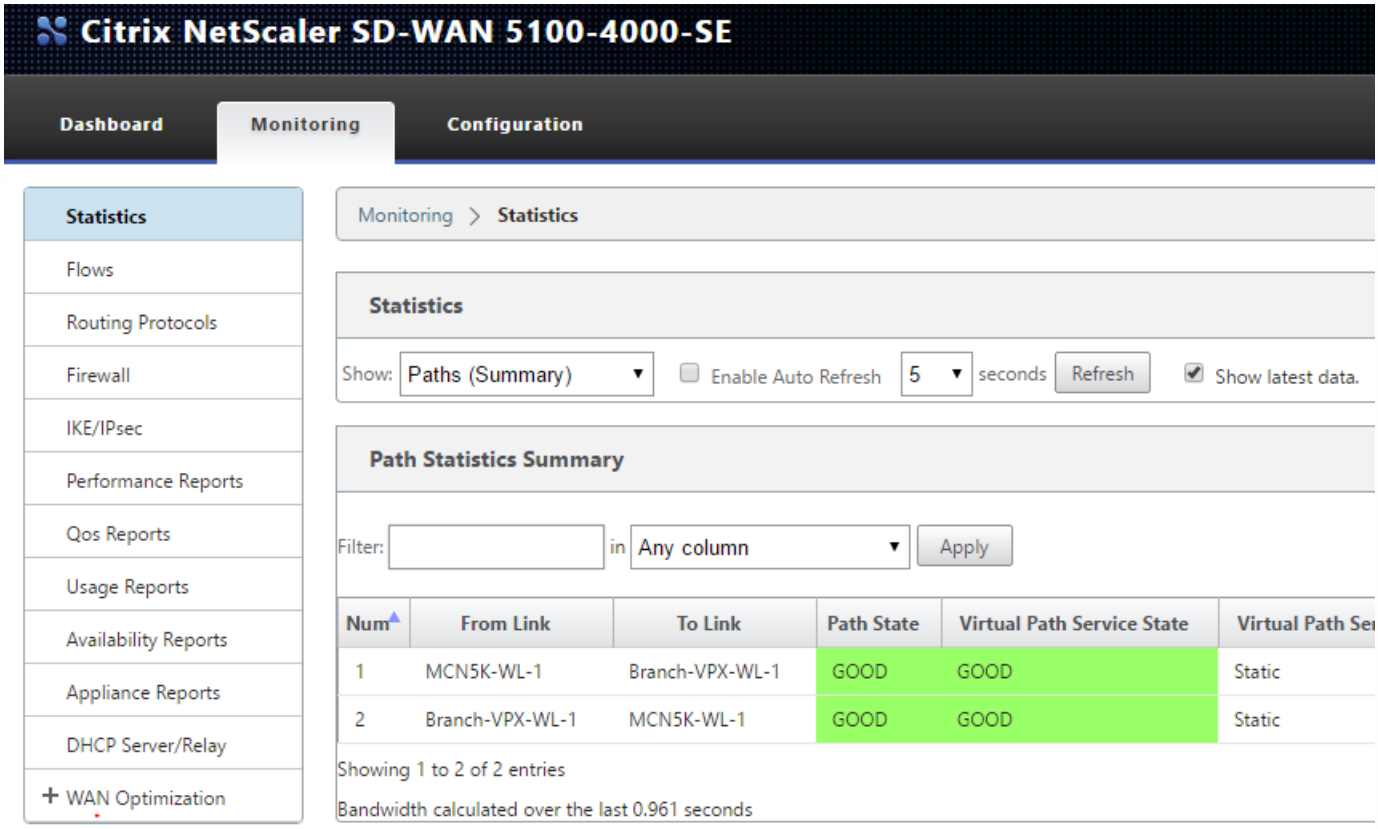

Management IP Address Change for SD-WAN WANOP Appliance in Two Box Mode

To change the management IP address of SD-WAN-WANOP appliance in Two box mode:

1. Execute command *clear_wo_sync* on the SD-WAN SE appliance. This ensures that the SD-WAN WANOP IP address information is cleared for GUI redirection.
2. Disable and enable Two box mode config on the SD-WAN WANOP appliance. The new IP address (changed IP) of SD-WAN WANOP appliance is sent to SD-WAN SE. The new changed IP address is displayed in the URL redirection pages.

> **Note**
>
> The management IP address is used for peer IP configuration.

Disable Two Box mode on SD-WAN WANOP Appliance

To disable or decouple the SD-WAN WANOP and SD-WAN SE appliances from the Two Box mode:

a. Disable the Two Box mode from SD-WAN WANOP appliance.

b. It is expected to still see the SD-WAN WANOP appliance two box mode pages in the SD-WAN SE web GUI. To clear these pages, execute the command: *clear_wo_sync*.

SD-WAN Overlay Routing

Oct 12, 2017

Introduction

NetScaler SD-WAN provides resilient and robust connectivity between remote sites, data centers, and cloud networks. The SD-WAN solution can accomplish this by establishing tunnels between SD-WAN appliances in the network enabling connectivity between sites by leveraging route tables that overlay the existing underlay network. SD-WAN route tables can fully replace or coexist with the existing routing infrastructure. This article below provides detailed routing configuration within the NetScaler SD-WAN network.

NetScaler SD-WAN Route Table

The SD-WAN configuration allows static route entries for specific sites, and route entries learned from the underlay network through supported routing protocols; such as OSPF, eBGP, and iBGP. Routes are not only defined by their next hop but by their service type. This determines how the route is forwarded. Below are the main service types in use:

- **Local Service:** This service denotes any route or subnet local to the SD-WAN appliance. This includes the Virtual Interface subnets (automatically creates local routes), and any local route defined in the route table (with a local next hop). The route is advertised to other SD-WAN appliances that have a Virtual Path to this local site where this route is configured when trusted as a partner.

> **Note**
>
> Be cautious when adding default routes, and summary routes as local routes as these can result in virtual path routes at other sites. Always check the route tables to make sure the correct routing is in effect.

- **Virtual Path** – This denotes any local route learned from a remote SD-WAN site; that is what is reachable down the virtual paths. These routes are normally automatic, however a virtual path route can be added manually at a site. Any traffic for this route is forwarded to the defined Virtual Path for this destination route (subnet).
- **Intranet** – This service denotes routes that are reachable through a private WAN link (MPLS, P2P, VPN etc.). For example, a remote branch that is on the MPLS network but does not have an SD-WAN appliance. It is assumed that these routes need to be forwarded to a certain WAN router. Intranet Service is not enabled by default. Any traffic matching this route (subnet) is classified as intranet for this appliance for delivery to a site that does not have an SD-WAN solution.

> **Note**
>
> Notice that when adding an Intranet route there is no next hop, but rather a forward to an Intranet Service. The Service is associated with a given WAN link.

- **Internet** – This is similar to Intranet but is used to define traffic flowing to public Internet WAN links rather than private WAN links. One unique difference is that the Internet service can be associated with multiple WAN links and set to load

balance (per flow) or be active/backup. A default Internet routes gets created when internet service is enabled (it is off by default). Any traffic matching this route (subnet) is classified as Internet for this appliance for delivery to public internet resources.

> **Note**
>
> Internet Service routes can be advertised to the other SD-WAN appliances or prevented from being exported depending on whether you are backhauling Internet access over the Virtual Paths.

- **Passthrough** – This service acts as a last resort or override service when an appliance is in-line mode. If a destination IP address fails to match with any other route, then the SD-WAN appliance simply forwards it onto the WAN link next hop. A default route: 0.0.0.0/0 cost of 16 pass-through route is created automatically. Passthrough does not work when the SD-WAN appliance is deployed out of path or in Edge/Gateway mode. Any traffic matching this route (subnet) is classified as passthrough for this appliance. It is recommended that passthrough traffic be limited as much as possible.

> **Note**
>
> Passthrough can be useful when conducting a POCs to avoid having to configure a lot of routing, however be very careful in production because SD-WAN does not account for WAN link utilization for traffic sent to passthrough. It is also helpful when troubleshooting issues and you want to take a certain IP flow out of delivery over the Virtual Path.

- **Discard** - This is not a service but a last resort route that drops the packets if it matches. Normally this does not occur expect when the SD-WAN appliance is deployed out of path. You must have an Intranet service or local route as a catch all route, otherwise the traffic will be discarded as there is no passthrough service (even though a passthrough default route will be present).

  The SD-WAN Configuration Editor enables route table customization for each available site:

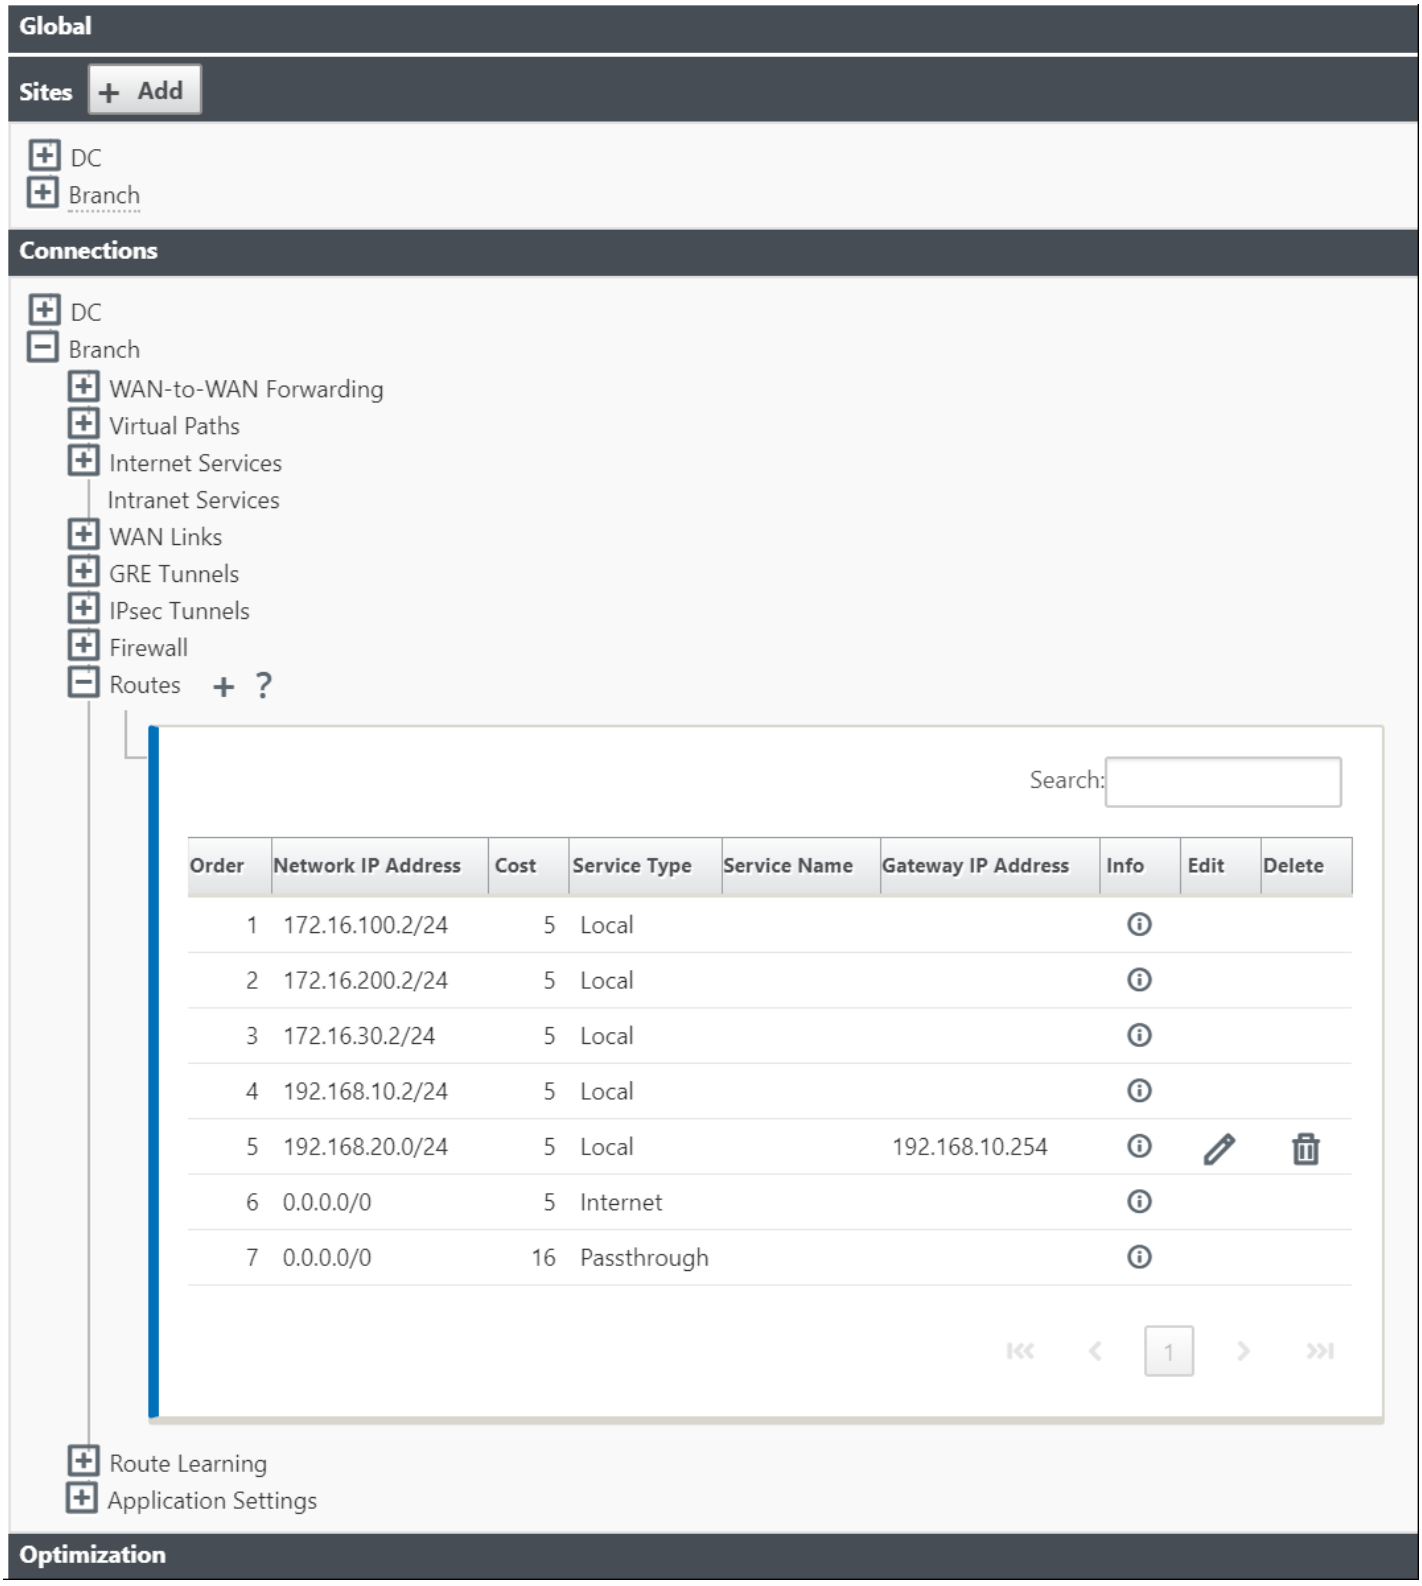

Route table entries are populated from different inputs:

- Configured Virtual IP Address (VIP) auto-populate as Service Type Local route. The Configuration Editor will prevent the same VIP assignment to different site nodes.
- Internet Services enabled at a local site will auto-populate a default route (0.0.0.0/0) locally for direct internet breakout.
- Admin defined static routes on a per site basis, which will also be defined as a Service Type Local route.

- A default (0.0.0.0/0) catch all route with cost 16 defined as Passthrough

Administrators can configure one of the above routes, but also include a service type, next hop, or gateway depending on the service type, in addition to route cost. A default route cost will automatically be added to each route type (reference the table below for default route costs). Additionally, only trusted routes are advertised to other SD-WAN appliances. Untrusted routes are only used by the local appliance.

Client node routes are only advertised to the MCN node and no other client nodes by default. In order for client node routes to be visible to another client nodes WAN to WAN Forwarding needs to be enabled at the MCN node.

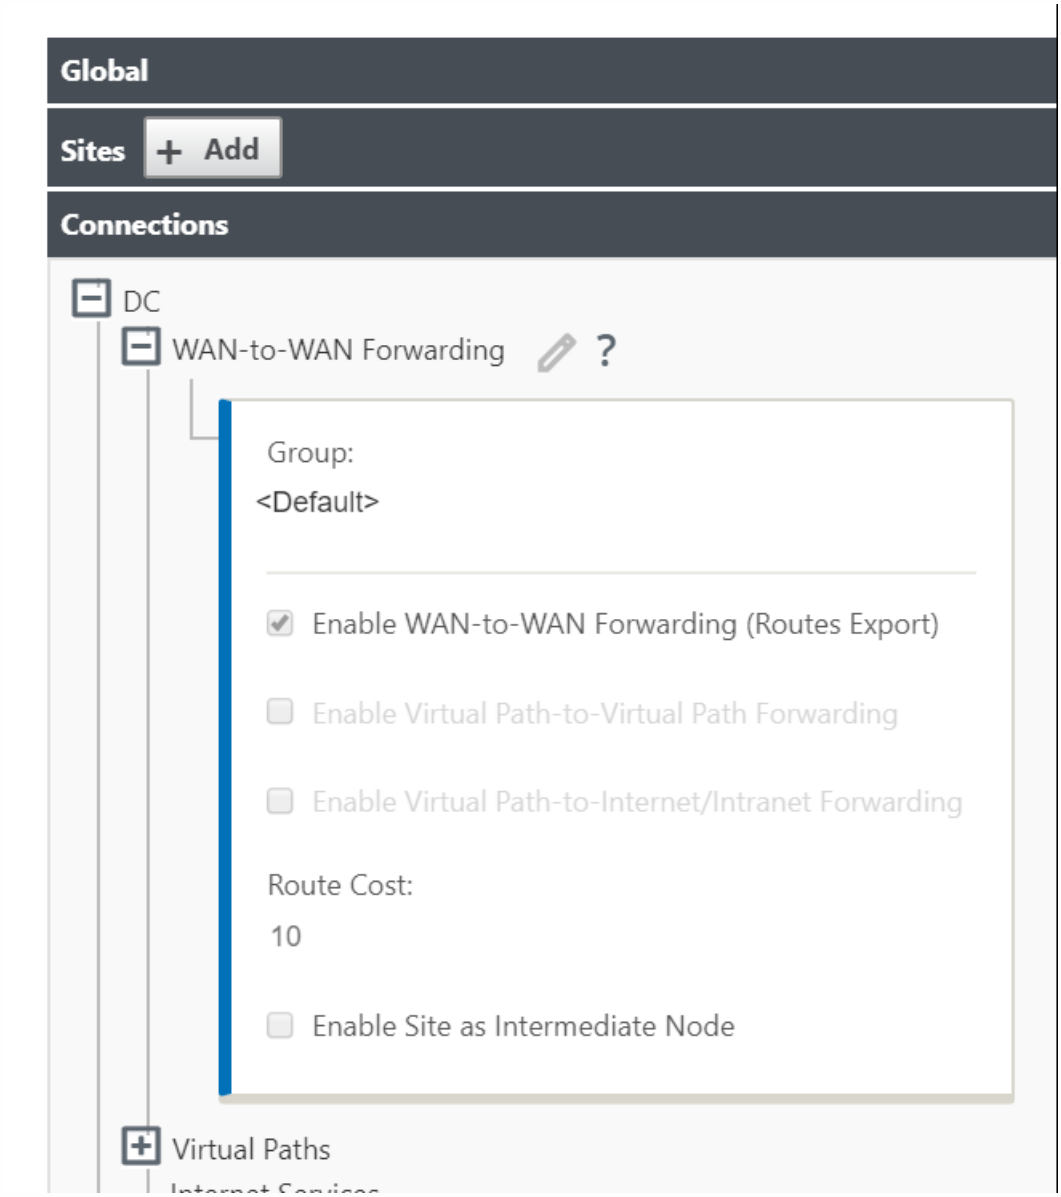

With WAN-to-WAN Forwarding (Routes Export) enabled, the MCN site will share the advertised routes to all clients participating in the SD-WAN overlay. Turning on this feature enables IP connectivity between hosts at different client node sites with the communication traveling through the MCN. The route table for the local client node can be monitored on the **Monitoring** > **Statistics** page with Routes selected for the **Show** drop-down.

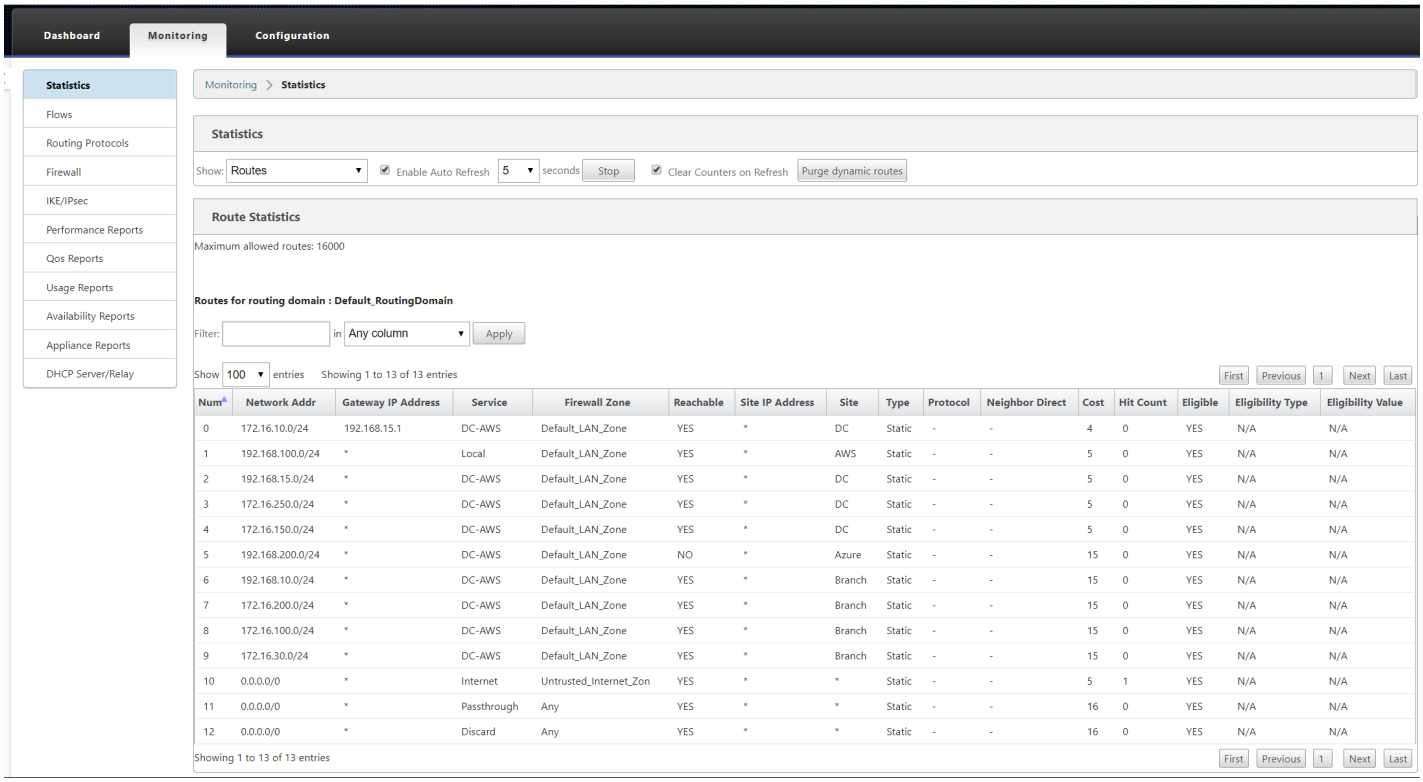

Each route for remote branch office subnets is advertised as a Service through the Virtual Path connecting through the MCN, with the Site column populated with the client node where the destination resides as a local subnet.

In the below example, with "**WAN-to-WAN Forwarding** (Routes Export)" enabled, Branch A has a route table entry for the Branch B subnet (10.2.2.0/24) through the MCN as a next hop.

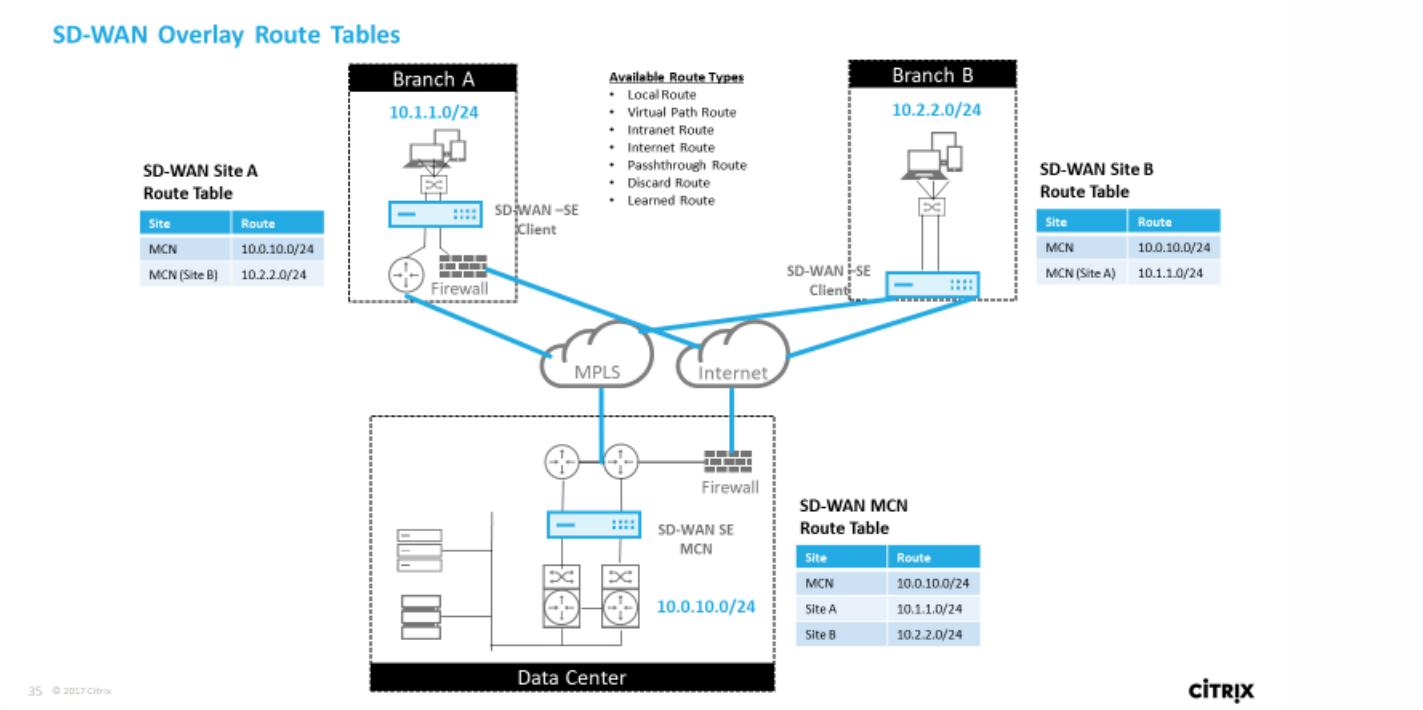

How NetScaler SD-WAN Traffic Matches on Defined Routes

The match process for defined routes on NetScaler SD-WAN is based on longest prefix match for destination subnet (similar to a router operation). The more specific the route, the higher the change on it being matched. Beyond that,

sorting is done in the following order:

1. Longest prefix match
2. Cost
3. Service

Therefore a /32 route always precedes a /31 route. For two /32 routes, a cost 4 route always precedes a cost 5 route. For two /32 cost 5 routes, routes are chosen based on ordered IP host. Service order is as follows: Local, Virtual Path, Intranet, Internet, Passthrough, Discard.

As an example, consider the following two routes below:

- 192.168.1.0/24 Cost 5
- 192.168.1.64/26 Cost 10

A packet destined for 192.168.1.65 host would use the latter route even though the cost is higher. Based on this, it is common for configuration to be in place for only the routes intended to be delivered over the Virtual Path overlay with other traffic falling into catch all routes such as a default route to the passthrough service.

Routes can be configured in a site node route table that have the same prefix. The tie break then goes to the route cost, the service type (Virtual Path, Intranet, Internet, etc.) and the next hop IP.

NetScaler SD-WAN Routing Packet Flow

- LAN to WAN (Virtual Path) Traffic Route Matching:
  1. Incoming traffic is received by the LAN interface and is processed.
  2. The received frame is compared to the route table for the longest prefix match.
  3. If a match is found, the frame is then processed by the rule engine and a flow is created in the flow database.
- WAN to LAN (Virtual Path) Traffic Route Matching:
  1. Virtual Path traffic is received by SD-WAN from the tunnel and is processed.
  2. The appliance compares the source IP address to see if the source is local.
     - If yes – then WAN eligible and match IP destination to routing table/Virtual Path.
     - If no – then WAN to WAN forwarding enabled check.
  3. (WAN to WAN Forwarding disabled) Forward to LAN based on local routes.
  4. (WAN to WAN Forwarding enabled) Forward to Virtual Path based on route table.
- Non Virtual Path Traffic:
  1. Incoming traffic is received on LAN interface and is processed.
  2. The received frame is compared to the route table for the longest prefix match.

3. If a match is found, the frame is then processed by the rule engine and a flow is created in the flow database.

NetScaler SD-WAN Routing Protocol Support

NetScaler SD-WAN release 9.1 introduced OSPF and BGP routing protocols into the configuration. Introducing routing protocols to SD-WAN enabled easier integration of SD-WAN in more complex underlay networks where routing protocols are actively in use. With the same routing protocols enabled on SD-WAN, configuration of subnets denoted to make use of the SD-WAN overlay was made easier. In addition, the routing protocols enable communication between SD-WAN and non-SD-WAN sites with direct communication to existing customer edge routers using the common routing protocol. NetScaler SD-WAN participating in routing protocols operating in the underlay network can be done regardless of the deployment mode of SD-WAN (Inline mode, Virtual Inline mode, or Edge/Gateway mode). Additionally, SD-WAN can be deployed in "learn only" mode where SD-WAN can receive routes but not advertise routes back to the underlay. This can be useful when introducing the SD-WAN solution into a network where the routing infrastructure is complex or uncertain.

Important

It is very easy to accidently leak unwanted route, if you are not careful.

The SD-WAN Virtual Path route table works as an External Gateway Protocol (EGP), very similar to BGP (think site-to-site). For example, when SD-WAN advertises routes from the SD-WAN appliance to OSPF they are typically considered external to site and protocol.

Note

Be aware of environments that have IGPs across the entire infrastructure (across the WAN) as it does complicate how SD-WAN advertised routes are used. EIGRP is extensively used in the market and SD-WAN does not interop with that protocol.

One challenge in introducing Routing Protocols to an SD-WAN deployment is that the route table is not available until the SD-WAN service is enabled and operation in the network, therefore it is not recommended to enable advertise routes from the SD-WAN appliance initially. Use the import and export filters for a gradual introduction of routing protocols on SD-WAN.

Let us take a closer look by reviewing the following example:

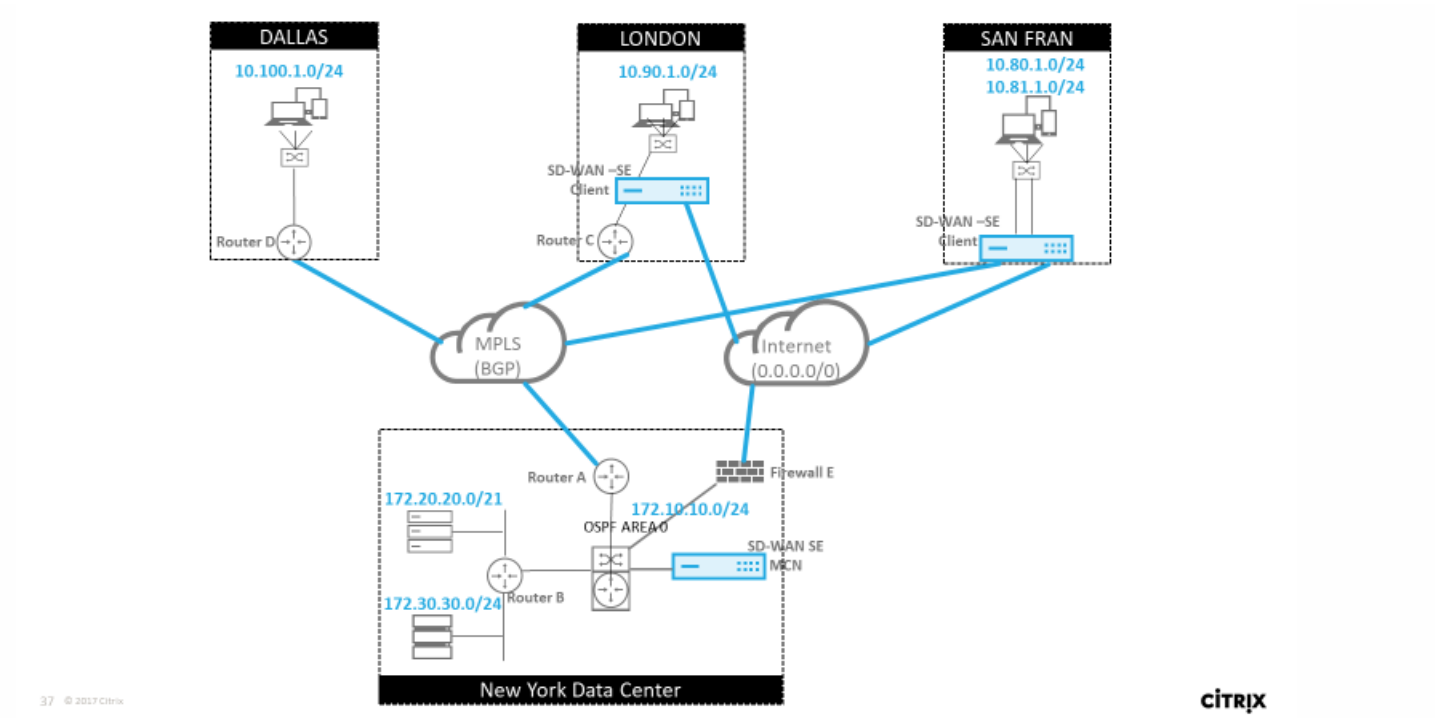

In this example, we will examine a routing protocol use case. The above network has four locations; New York, Dallas, London, and San Francisco. We will deploy SD-WAN appliances at three of these locations, and utilize SD-WAN to create a hybrid WAN network where MPLS and Internet WAN Links will be used to provide a Virtualized WAN. Since Dallas will not have an SD-WAN device, we need to consider how to best integrate with existing route protocols to that site to ensure full connectivity between underlay and SD-WAN overlay networks.

In the example network, eBGP is used between all four locations across the MPLS network. Each location has its own Autonomous System Number (ASN).

In the New York Data Center, OSPF is running to advertise the core Data Center subnets to the remote sites and also announce a default route from the New York Firewall (E). In this example, all internet traffic is backhauled to the datacenter, even though London and San Francisco Branches have a path to the internet.

The San Francisco site also should be noted to not have a router. SD-WAN will be deployed in Edge/Gateway mode with that appliance being the default gateway for the San Francisco subnet and also participating in eBGP to the MPLS.

- With the New York Data Center, take note that the SD-WAN is deployed in Virtual Inline mode. The intent is to participate in the existing OSPF routing protocol to get traffic forwarded to the appliance as the preferred gateway.
- The London site is deployed in traditional inline mode. The upstream WAN Router (C) will still be the default gateway for the London subnet.
- The San Francisco site will be a newly introduced site to this network and the SD-WAN is planned to be deployed in Edge/Gateway mode and act as the default gateway for the new San Francisco subnet.

First. we will take a look at some of the existing underlay route tables before implementing SD-WAN.

New York Core Router B:

```
vyos@VYOS-ROUTER-B-CORE:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O>* 0.0.0.0/0 [110/10] via 172.10.10.3, eth1, 00:08:56
O>* 10.90.1.0/24 [110/20] via 172.10.10.1, eth1, 00:21:02
O>* 10.100.1.0/24 [110/20] via 172.10.10.1, eth1, 00:21:02
C>* 127.0.0.0/8 is directly connected, lo
O   172.10.10.0/24 [110/10] is directly connected, eth1, 1d20h00m
C>* 172.10.10.0/24 is directly connected, eth1
C>* 172.20.20.0/24 is directly connected, eth2
C>* 172.30.30.0/24 is directly connected, eth3
C>* 192.168.65.0/24 is directly connected, eth0
```

The local New York subnets (172.x.x.x) are available on router B as directly connected, and from the route table we identify that the default route is 172.10.10.3 (Firewall E). Also, we can see that Dallas (10.90.1.0/24) and London (10.100.1.0/24) subnets are available via 172.10.10.1 (MPLS Router A). Note the route costs indicate they were learned from eBGP.

> **Note**
>
> In the example provided, San Francisco is not listed as a route, because we have not yet deployed the site with SD-WAN in Edge/Gateway mode for that network.

```
vyos@VYATTA-ROUTER-A:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O>* 0.0.0.0/0 [110/10] via 172.10.10.3, eth1, 00:09:52
B>* 10.90.1.0/24 [20/1] via 192.168.10.2, eth2, 1d23h09m
B>* 10.100.1.0/24 [20/1] via 192.168.10.3, eth2, 1d23h10m
C>* 127.0.0.0/8 is directly connected, lo
O   172.10.10.0/24 [110/10] is directly connected, eth1, 1d20h01m
C>* 172.10.10.0/24 is directly connected, eth1
O>* 172.20.20.0/24 [110/20] via 172.10.10.2, eth1, 00:21:58
O>* 172.30.30.0/24 [110/20] via 172.10.10.2, eth1, 00:21:58
C>* 192.168.10.0/24 is directly connected, eth2
O   192.168.65.0/24 [110/20] via 172.10.10.2, 1d19h57m
C>* 192.168.65.0/24 is directly connected, eth0
```

For the New York WAN Router (A), OSPF learned routes and routes learned across the MPLS through eBGP are listed routes. Note the route costs. BGP is lower administrative domain and cost by default 20/1 compared to OSPF 110/10.

Dallas Router D:

For the Dallas WAN Router (D) all routes are learned across the MPLS.

```
vyos@VYATTA-ROUTER-D:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B>* 0.0.0.0/0 [20/10] via 192.168.10.1, eth2, 00:10:17
B>* 10.90.1.0/24 [20/1] via 192.168.10.2, eth2, 1d23h10m
C>* 10.100.1.0/24 is directly connected, eth1
C>* 127.0.0.0/8 is directly connected, lo
B>* 172.10.10.0/24 [20/1] via 192.168.10.1, eth2, 1d23h10m
B>* 172.20.20.0/24 [20/20] via 192.168.10.1, eth2, 00:22:17
B>* 172.30.30.0/24 [20/20] via 192.168.10.1, eth2, 00:22:17
C>* 192.168.10.0/24 is directly connected, eth2
C>* 192.168.65.0/24 is directly connected, eth0
```

> ## Note
>
> In this example, you can ignore the 192.168.65.0/24 subnet. This is a management network and not pertinent to the example. All the Routers are connected to the management subnet but the is not advertised in any routing protocol.

In NetScaler SD-WAN, we can add the SD-WAN overlay by enabling OSPF on the SD-WAN located in the New York site under **Connections** > **Site** > **Route Learning** > **OSPF** > **Basic Settings**:

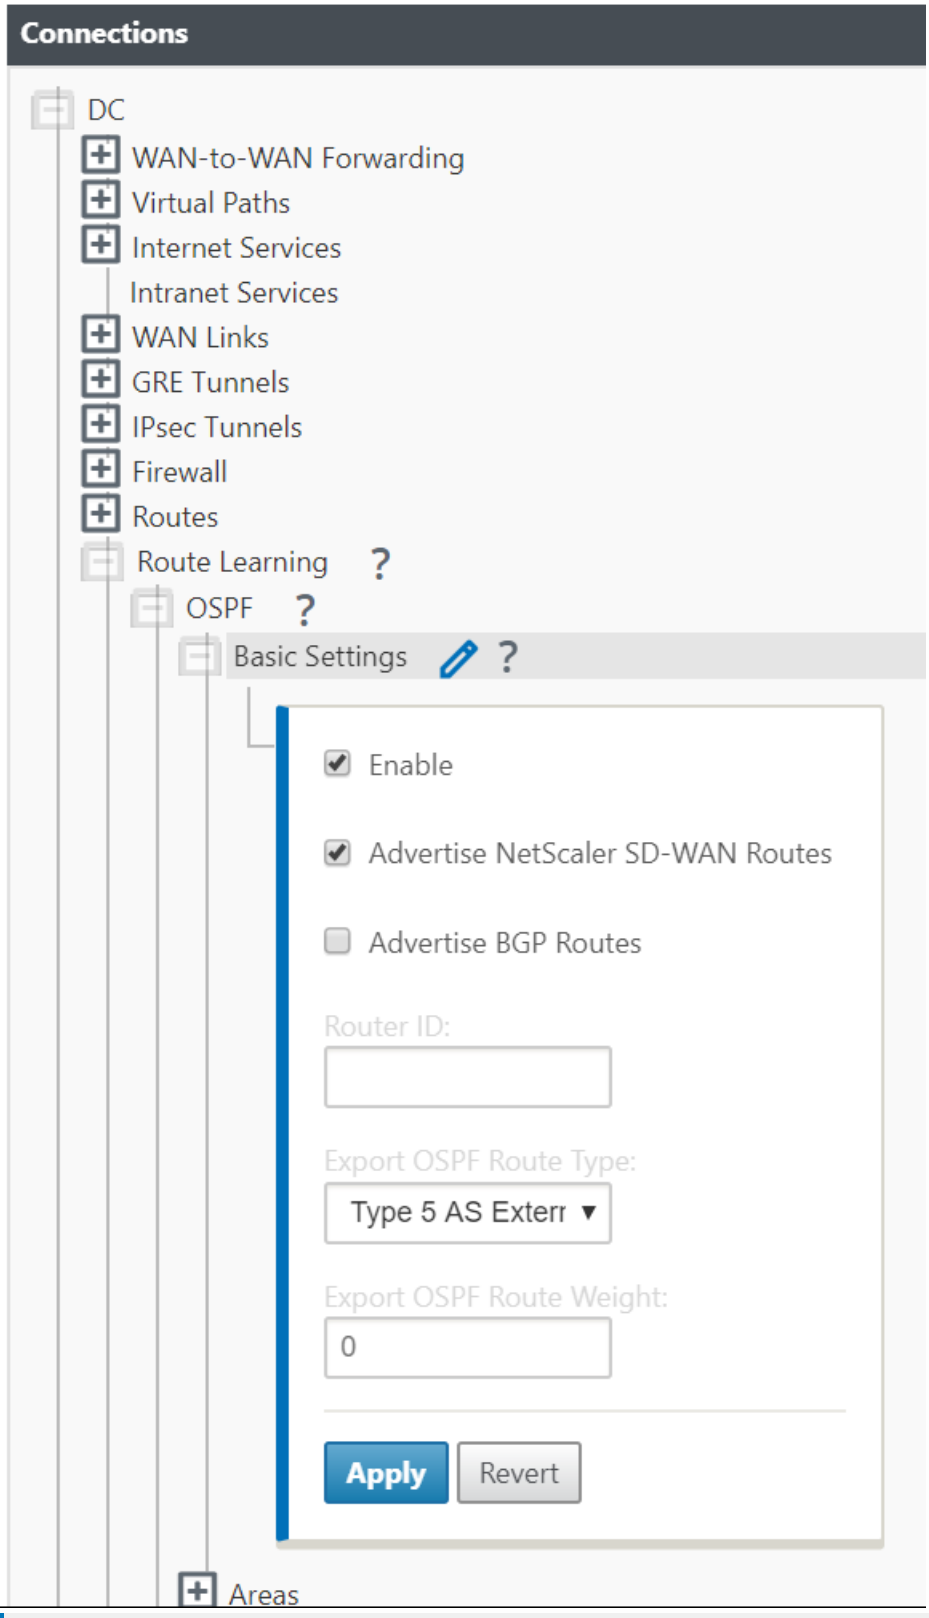

Note

The **Export OSPF Route Type** is Type 5 External by default. This is because SD-WAN routing table is considered external to the OSPF protocol and so OSPF will prefer a route learned internal (intra-area), therefore routes advertised by SD-WAN may not take precedence.

When OSPF is used across the WAN (i.e. MPLS networks), then this can be changed to Type 1 intra-area. OSPF areas can be configured directly below the **Basic Settings** node.

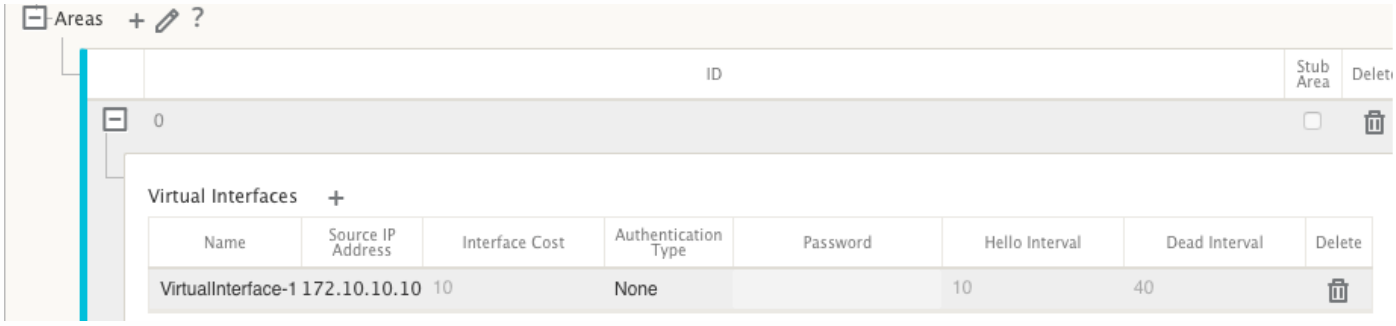

Area 0 added with the local network derived from the Virtual Interface (172.10.10.0), all other settings were left default.

For the new San Francisco site, we will need to enable eBGP since it will be directly connected to the MPLS network and operating as the customer edge route for the site. BGP can be enabled under **Connections** > **Site** > **Route Learning** > **BGP** > **Basic Settings.**

Note the Autonomous System number of 13.

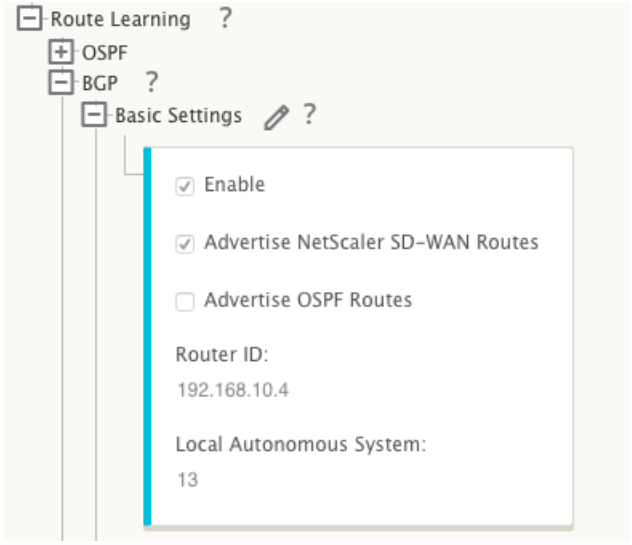

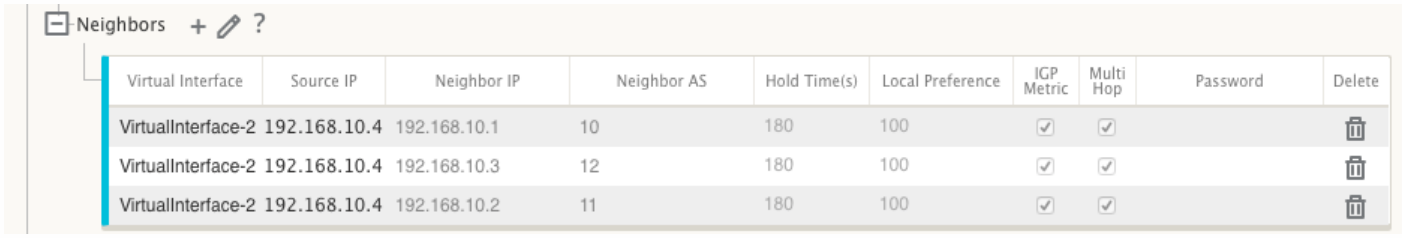

The eBGP peers with each other location. Note that each ASN is different.

It is important to understand how routes are passed between the Virtual Path routing table and the dynamic route protocols in use. It is easy to create routing loops or advertise routes in an adverse way. The filter mechanism gives us the ability to control what gets into and out of the routing table. We will consider each location in turn.

- The San Francisco location has two local subnets **10.80.1.0/24** and **10.81.1.0/24**. We want to advertise them through eBGP so that sites like Dallas can still reach the San Francisco site over the underlay network and also sites like London and San Francisco can still reach San Francisco over the Virtual Path overlay network. We also want to learn from eBGP reachability to all sites in case the SD-WAN Virtual Path overlay goes down and the environment needs to fall back to using

just the MPLS. We also do not want to re-advertise anything SD-WAN learns from eBGP to the SD-WAN routers. In order to accomplish this, the filters need to be configured as follows:

- Import all routes from eBGP. Do not re-advertise/export routes to SD-WAN appliances.

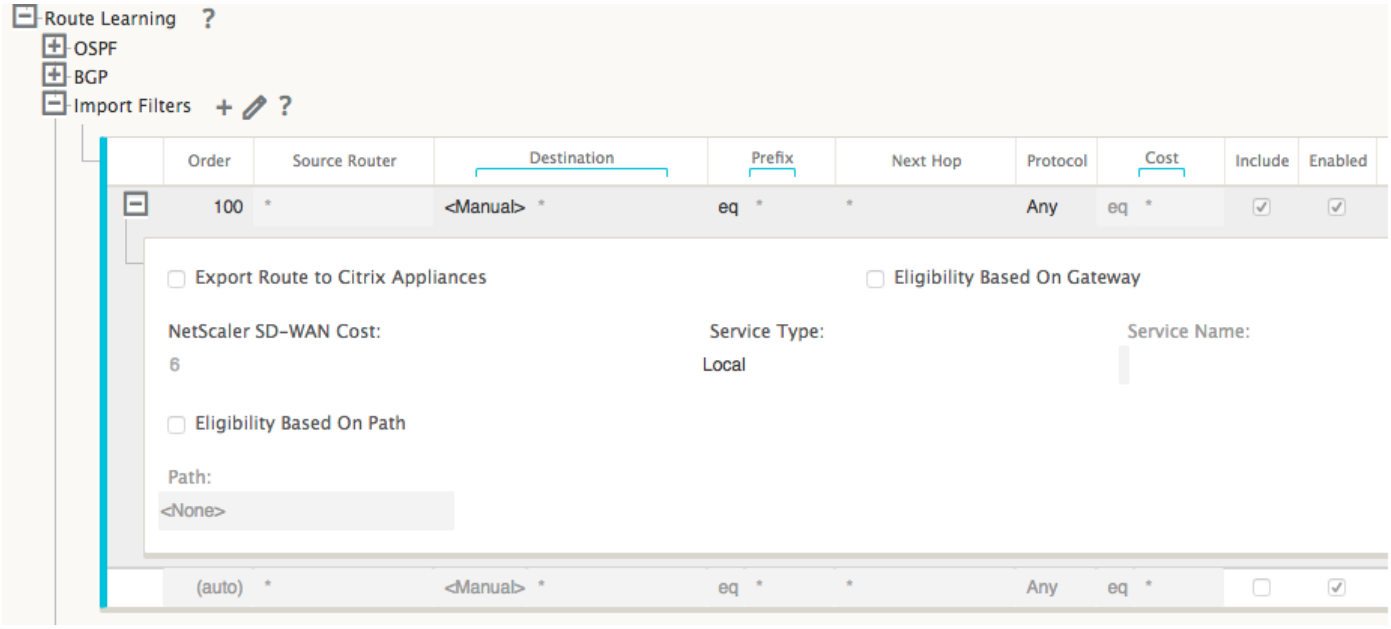

- Export local routes to eBGP

The default rule for export is to export everything. Rule 200 is used to override the fault rule in order to not re-advertise the routes. Any route matching any prefix SD-WAN has learned across the Virtual Paths.

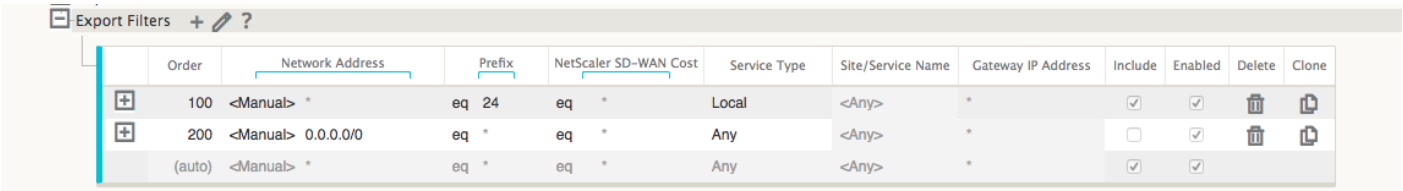

After the NetScaler SD-WAN appliances have been deployed, we can take a refreshed look at the route tables for the BGP router at the Dallas site. We see 10.80.1.0/24 and 10.81.1.0/24 subnets are being seen correctly through eBGP from the San Francisco SD-WAN.

Dallas Router D:

```
vyos@VYATTA-ROUTER-D:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B>* 0.0.0.0/0 [20/10] via 192.168.10.1, eth2, 00:00:01
B>* 10.80.1.0/24 [20/0] via 192.168.10.4, eth2, 3d19h07m
B>* 10.81.1.0/24 [20/0] via 192.168.10.4, eth2, 3d19h07m
B>* 10.90.1.0/24 [20/1] via 192.168.10.2, eth2, 4d23h38m
C>* 10.100.1.0/24 is directly connected, eth1
C>* 127.0.0.0/8 is directly connected, lo
B>* 172.10.10.0/24 [20/1] via 192.168.10.1, eth2, 4d23h38m
B>* 172.20.20.0/24 [20/20] via 192.168.10.1, eth2, 00:00:01
B>* 172.30.30.0/24 [20/20] via 192.168.10.1, eth2, 00:00:01
B   192.168.10.0/24 [20/0] via 192.168.10.4 inactive, 3d19h07m
C>* 192.168.10.0/24 is directly connected, eth2
C>* 192.168.65.0/24 is directly connected, eth0
```

Further, the NetScaler SD-WAN route table can be viewed on the **Monitoring** > **Statistics** > **Show Routes** page.

San Francisco NetScaler SD-WAN:

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: in Any column Apply

Show 100 entries Showing 1 to 16 of 16 entries First Previous 1 Next Last

Num	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
0	10.81.1.0/24	10.80.1.20	Local	YES	*	SFO	Static	-	-	5	0	YES	N/A	N/A
1	10.80.1.0/24	*	Local	YES	*	SFO	Static	-	-	5	0	YES	N/A	N/A
2	192.168.10.0/24	*	Local	YES	*	SFO	Static	-	-	5	122	YES	N/A	N/A
3	172.10.10.0/24	*	NYC-SFO	YES	*	NYC	Static	-	-	5	0	YES	N/A	N/A
4	172.30.30.0/24	192.168.10.1	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
5	172.20.20.0/24	192.168.10.1	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
6	172.10.10.0/24	192.168.10.1	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
7	10.100.1.0/24	192.168.10.3	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
8	10.90.1.0/24	192.168.10.2	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
9	172.20.20.0/24	*	NYC-SFO	YES	*	NYC	Dynamic	Virtual WAN	YES	6	0	YES	N/A	N/A
10	10.100.1.0/24	*	NYC-SFO	YES	*	NYC	Dynamic	Virtual WAN	YES	6	0	YES	N/A	N/A
11	172.30.30.0/24	*	NYC-SFO	YES	*	NYC	Dynamic	Virtual WAN	YES	6	0	YES	N/A	N/A
12	0.0.0.0/0	192.168.10.1	Local	YES	*	SFO	Dynamic	BGP	-	6	0	YES	N/A	N/A
13	0.0.0.0/0	*	NYC-SFO	YES	*	NYC	Dynamic	Virtual WAN	YES	6	0	YES	N/A	N/A
14	0.0.0.0/0	*	Passthrough	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
15	0.0.0.0/0	*	Discard	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

Showing 1 to 16 of 16 entries First Previous 1 Next Last

NetScaler SD-WAN shows all the routes learned, including routes available through the Virtual Path overlay.

Let us consider 172.10.10.0/24, which is located in New York Data Center. This route is being learned in two ways:

- As a Virtual Path route (Num 3), service = NYC-SFO with a cost of 5 and type static. This is a local subnet being advertised by SD-WAN appliance in New York. It is static in that it is either directly connected to the appliance or it is a

manual static route entered in the configuration. It is reachable because the Virtual Path between the sites is in a working/up state.

- As an advertised route through BGP (Num 6), with a cost of 6. This is now considered a fallback route.

Since the prefix is equal and cost is different, SD-WAN will use the Virtual Path route unless it becomes unavailable, in which case the fallback route is learned through BGP.

Now, let us consider the route 172.20.20.0/24.

- This is learned as a Virtual Path route (Num 9) but has a type of dynamic and a cost of 6. This means the remote SD-WAN appliance learned this route through a routing protocol, in this case OSPF. By default the route cost is higher.
- SD-WAN also learns this route through BGP with the same cost, so in this case this route may be preferred over the Virtual Path route.

In order to ensure correct routing, we must increase the BGP route cost to make sure if we have a Virtual Path route and it will be the preferred route. This can be done by adjusting the import filter route weight to be higher than the default of 6.

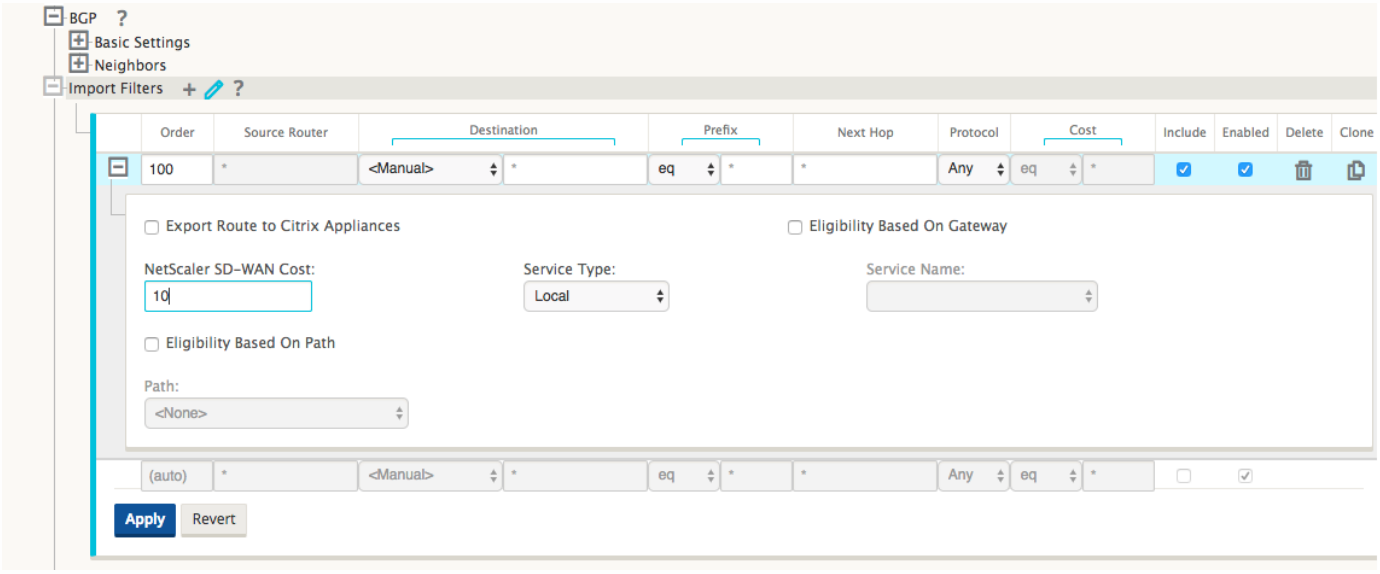

After making the adjustment, we can refresh the SD-WAN route table on the San Francisco appliance to see the adjusted route costs. Make use of the filter option to focus the displayed list.

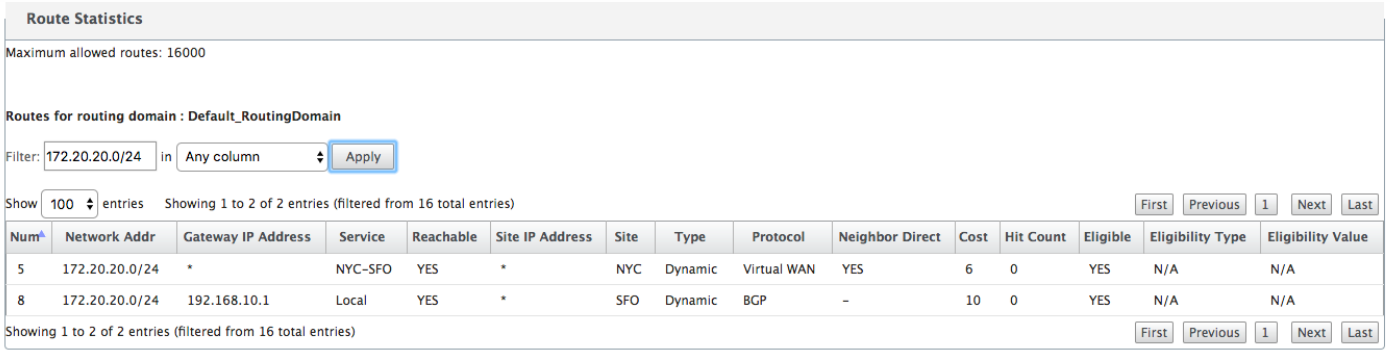

Finally, let us look at the learned default route on the San Francisco SD-WAN. We want to backhaul all internet traffic to New York. We can see that we will send it using the Virtual Path, if it is up, or through the MPLS network as a fallback.

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: 0.0.0.0/0 in Any column Apply

Show 100 entries Showing 1 to 4 of 4 entries (filtered from 16 total entries) First Previous 1 Next Last

Num	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
12	0.0.0.0/0	*	NYC-SFO	YES	*	NYC	Dynamic	Virtual WAN	YES	6	0	YES	N/A	N/A
13	0.0.0.0/0	192.168.10.1	Local	YES	*	SFO	Dynamic	BGP	-	10	0	YES	N/A	N/A
14	0.0.0.0/0	*	Passthrough	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
15	0.0.0.0/0	*	Discard	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

Showing 1 to 4 of 4 entries (filtered from 16 total entries) First Previous 1 Next Last

We also see a passthrough and discard route with cost 16. These are automatic routes that cannot be removed. If the device is inline, the passthrough route is used as a last resort so if a packet cannot be matched to a more specific route, SD-WAN will pass it along to the next hop of the interface group. If the SD-WAN is out of path or in edge/gateway mode, there is no passthrough service, in which case SD-WAN drops the packet using the default discard route. The Hit Count indicates the number of packets that are hitting each route, which can be valuable when troubleshooting.

Now focusing on the New York site, we want to get traffic destined for remote sites (London and San Francisco) to be directed to the SD-WAN appliance when the Virtual Path is active.

There are multiple subnets available in the New York site:

- 172.10.10.0/24 (directly connected)
- 172.20.20.0/24 (advertised via OSPF from the core router B)
- 172.30.30.0/24 (advertised via OSPF from the core router B)

We also are required to still provide traffic flow to Dallas (10.100.1.0/24) through MPLS.

Lastly, we want all internet bound traffic route to the Firewall E through 172.10.10.3 as a next hop. SD-WAN learns this default route through OSPF and will need to advertise this across the Virtual Path. The filters for the New York site are:

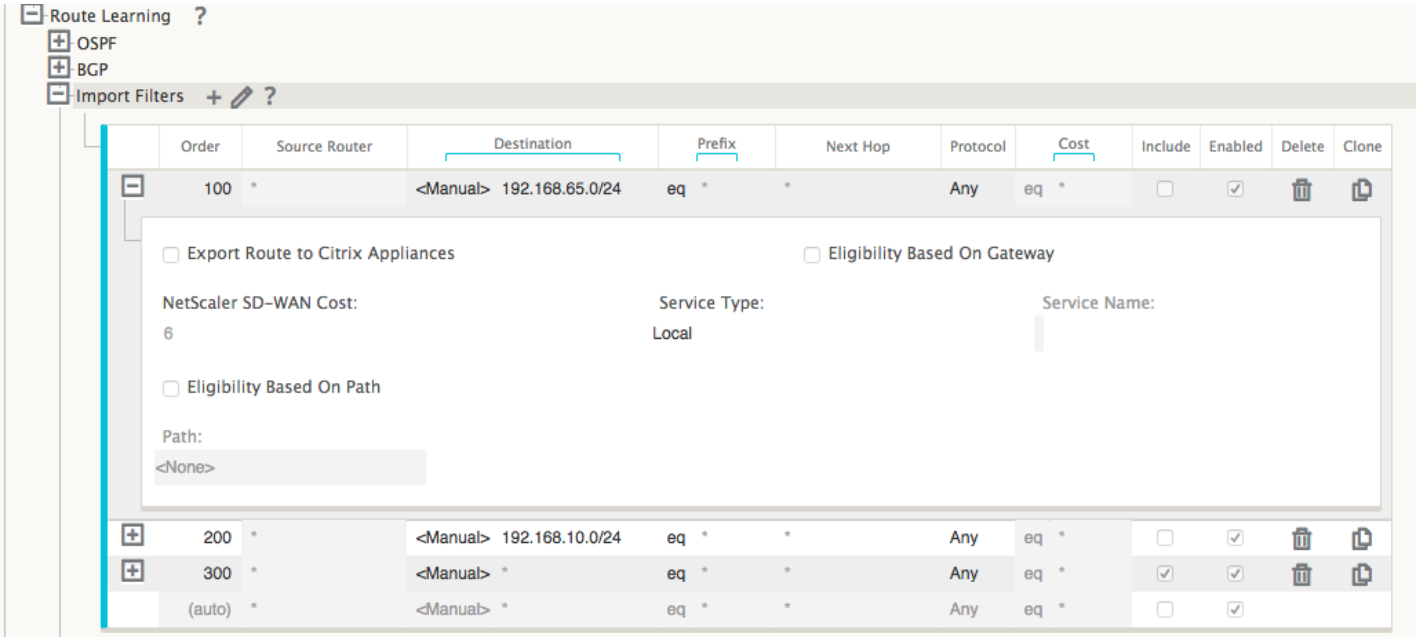

The New York SD-WAN site will import all routes for the management network. This can be ignored. We can focus on filter 200.

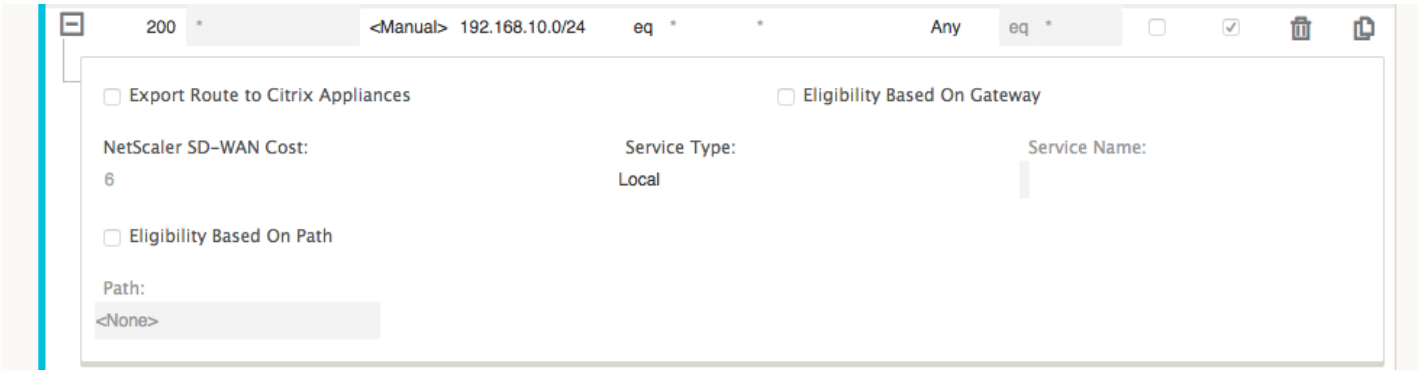

Filter 200 is used to import 192.168.10.0/24 (our MPLS core) for reachability but it is not advertised across the Virtual Path overlay. All other routes are then included.

For the export filters, we can exclude route for 192.168.10.0/24. This is because, as a directly connected subnet in San Francisco site, we cannot filter this route out at the source, so it is suppressed at this end.

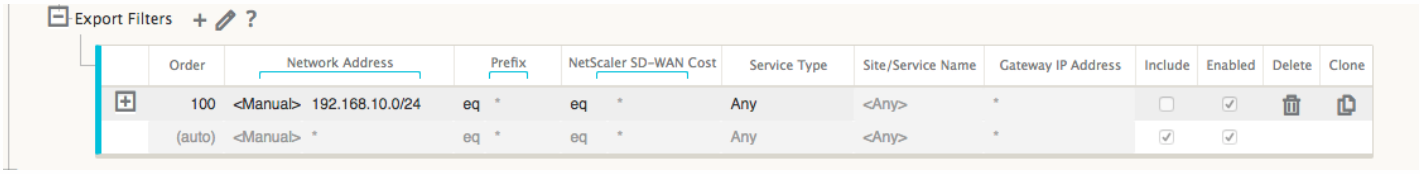

Now let us review the refreshed route table starting at the core route in New York site.

New York Router B:

```
vyos@VYOS-ROUTER-B-CORE:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O>* 0.0.0.0/0 [110/10] via 172.10.10.3, eth1, 4d22h22m
O>* 10.80.1.0/24 [110/15] via 172.10.10.10, eth1, 3d19h49m
O>* 10.81.1.0/24 [110/15] via 172.10.10.10, eth1, 3d19h49m
O>* 10.90.1.0/24 [110/15] via 172.10.10.10, eth1, 3d19h50m
O>* 10.100.1.0/24 [110/20] via 172.10.10.1, eth1, 4d22h22m
C>* 127.0.0.0/8 is directly connected, lo
O   172.10.10.0/24 [110/10] is directly connected, eth1, 4d22h22m
C>* 172.10.10.0/24 is directly connected, eth1
C>* 172.20.20.0/24 is directly connected, eth2
C>* 172.30.30.0/24 is directly connected, eth3
C>* 192.168.65.0/24 is directly connected, eth0
```

We can see the subnets for San Francisco (10.80.1.0 & 10.81.1.0) and London (10.90.1.0) now being advertised via the New York SD-WAN Appliance (172.10.10.10). The route 10.100.1.0/24 is still being advertised through the underlay MPLS Router A. Let us review the New York site SD-WAN route table.

New York site SD-WAN Route Table:

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: in Any column Apply

Show 100 entries Showing 1 to 11 of 11 entries First Previous 1 Next Last

Num	Network Addr	Gateway IP Address	Service	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
0	172.10.10.0/24	*	Local	YES	*	NYC	Static	-	-	5	0	YES	N/A	N/A
1	10.90.1.0/24	*	NYC-LON	YES	*	LON	Static	-	-	5	0	YES	N/A	N/A
2	10.81.1.0/24	10.80.1.20	NYC-SFO	YES	*	SFO	Static	-	-	5	0	YES	N/A	N/A
3	10.80.1.0/24	*	NYC-SFO	YES	*	SFO	Static	-	-	5	0	YES	N/A	N/A
4	192.168.10.0/24	*	NYC-SFO	YES	*	SFO	Static	-	-	5	0	YES	N/A	N/A
5	172.30.30.0/24	172.10.10.2	Local	YES	*	NYC	Dynamic	OSPF	-	6	0	YES	N/A	N/A
6	172.20.20.0/24	172.10.10.2	Local	YES	*	NYC	Dynamic	OSPF	-	6	0	YES	N/A	N/A
7	10.100.1.0/24	172.10.10.1	Local	YES	*	NYC	Dynamic	OSPF	-	6	0	YES	N/A	N/A
8	0.0.0.0/0	172.10.10.3	Local	YES	*	NYC	Dynamic	OSPF	-	6	0	YES	N/A	N/A
9	0.0.0.0/0	*	Passthrough	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
10	0.0.0.0/0	*	Discard	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

We can see the correct routes for both the local subnets learned via OSPF, a route to Dallas site learned from the MPLS Router A and the remote subnets for San Francisco and London sites. Lastly, let us look at the MPLS Router A. This router is participating in OSPF and BGP.

```
vyos@VYATTA-ROUTER-A:~$ show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O>* 0.0.0.0/0 [110/10] via 172.10.10.3, eth1, 00:04:12
O   10.80.1.0/24 [110/15] via 172.10.10.10, 00:04:13
B>* 10.80.1.0/24 [20/0] via 192.168.10.4, eth2, 00:05:09
O   10.81.1.0/24 [110/15] via 172.10.10.10, 00:04:13
B>* 10.81.1.0/24 [20/0] via 192.168.10.4, eth2, 00:05:09
O   10.90.1.0/24 [110/15] via 172.10.10.10, 00:04:13
B>* 10.90.1.0/24 [20/1] via 192.168.10.2, eth2, 00:05:11
S>* 10.90.1.10/32 [5/0] via 192.168.10.2, eth2
B>* 10.100.1.0/24 [20/1] via 192.168.10.3, eth2, 00:04:28
C>* 127.0.0.0/8 is directly connected, lo
O   172.10.10.0/24 [110/10] is directly connected, eth1, 00:05:24
C>* 172.10.10.0/24 is directly connected, eth1
O>* 172.20.20.0/24 [110/20] via 172.10.10.2, eth1, 00:04:12
O>* 172.30.30.0/24 [110/20] via 172.10.10.2, eth1, 00:04:12
B   192.168.10.0/24 [20/0] via 192.168.10.4 inactive, 00:05:09
C>* 192.168.10.0/24 is directly connected, eth2
O   192.168.65.0/24 [110/20] via 172.10.10.2, eth1, 00:04:12
C>* 192.168.65.0/24 is directly connected, eth0
```

From the route table this Router A is learning the remote subnets through BGP and OSPF with the Administrative distance and cost of the BGP route (20/5) being lower than OSPF (110/10) and hence preferred. In this example, network where there is only one core route, this may not cause concern. However, traffic arriving here would be delivered via the MPLS

network rather than being sent to the SD-WAN Appliance (172.10.10.10). If we want to maintain complete routing symmetry, we would need a route map to adjust the AD/Metric cost so that there is route preference from the route coming from 172.10.10.10 rather than the route learned via eBGP.

Alternatively, a "backdoor" route can be configured to force the router to prefer the OSPF route over the BGP route. astly, notice the static route for the SD-WAN Virtual IP address to the London site SD-WAN appliance.

```
S>* 10.90.1.10/32 [5/0] via 192.168.10.2, eth2
```

This is necessary to ensure the Virtual Path is re-routed back to the New York site SD-WAN appliance if the MPLS path goes down. Since there is a route for the 10.90.1.0/24 being advertised via 172.10.10.10 (New York SD-WAN). It is also recommended to create an override service rule to drop any UDP 4980 packets at the SD-WAN appliance to prevent the Virtual Path from coming back to itself.

Dynamic Virtual Paths

Dynamic Virtual Paths can be allowed between two client nodes to build on-demand virtual paths for direct communication between two sites. The advantage of a dynamic virtual path is that traffic can flow directly from one client node to second without having to traverse the MCN or two virtual paths, which could add latency to the traffic flow. Dynamic virtual paths are built and removed dynamically based on user-defined traffic thresholds. These thresholds are defined as either packets per second (pps) or bandwidth (kbps). This functionality enables a dynamic full mesh SD-WAN overlay topology.

Once the thresholds for dynamic virtual paths are meet, the client nodes dynamically create their virtualized path to one another leveraging all available WAN paths between the sites and make full use of it in the following manner:

- Send Bulk data if any exists and verify no loss, then
- Send Interactive data and verify no loss, then
- Send Real Time data after the Bulk and Interactive data are considered stable (no loss or acceptable levels)
- If there is no Bulk or interactive data send Real Time Data after the Dynamic Virtual Path has been stable for a period of time
- If the user data falls below the configured thresholds for a user defined period of time, the dynamic virtual path is torn down

  Dynamic Virtual Paths have the concept of an Intermediate site. The intermediate site could be an MCN site or any other site in the network that has Static Virtual Path configured and connected to two or more other client nodes. Another design consideration requirement is to have WAN-to-WAN Forwarding enabled, allowing all routes from all sites to be advertised to the client nodes where the dynamic virtual path is desired. "**Enable Site as Intermediate Node**" must be enabled in addition to **WAN-to-WAN Forwarding** in order for this intermediate site to monitor client node communication and to dictate when the dynamic path needs to be established and torn down.

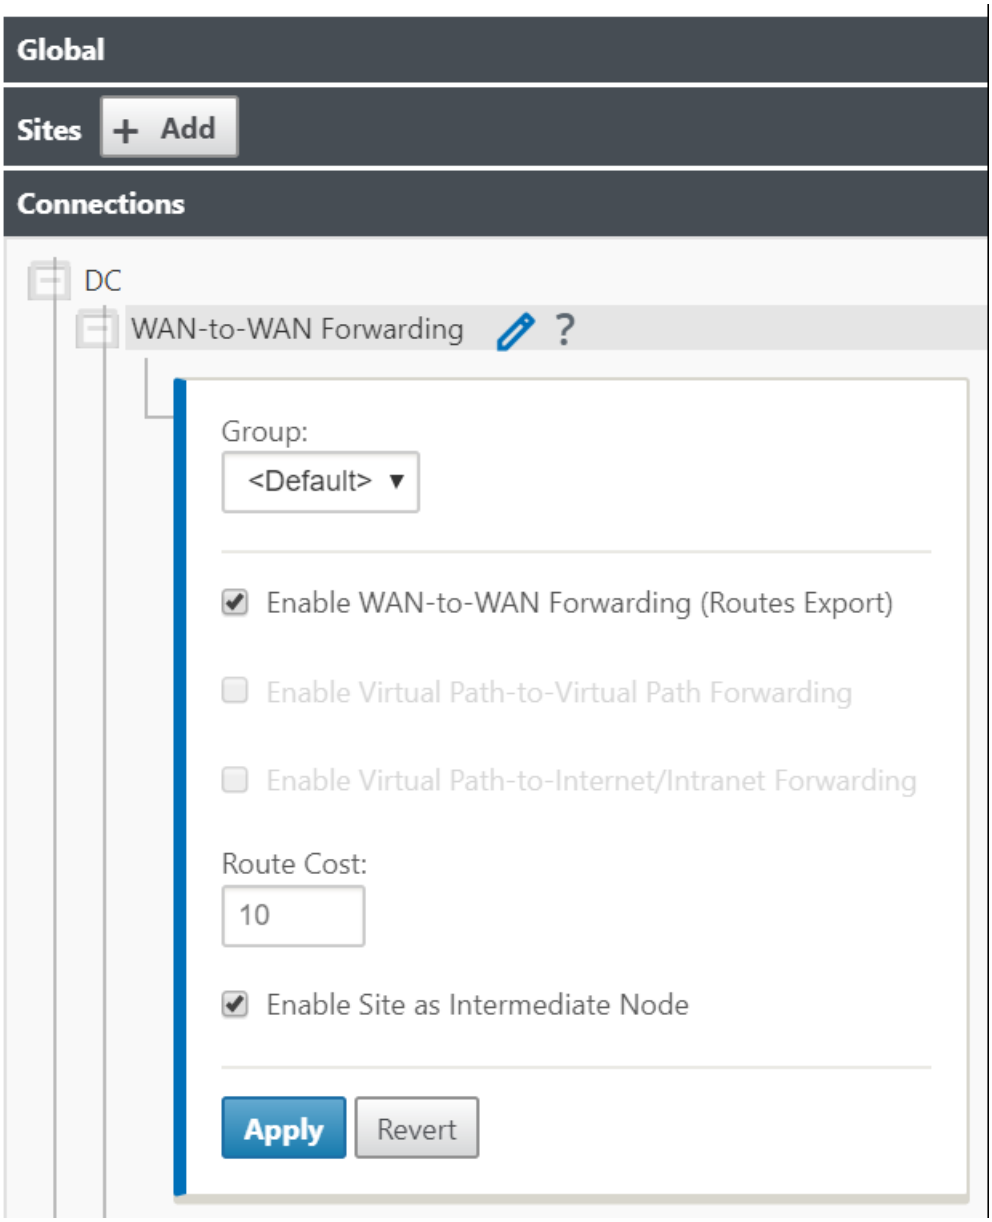

Multiple WAN-to-WAN Forwarding Groups can are allowed in the SD-WAN configuration, enabling full control to path establishment between certain client nodes and not others.

Multiple WAN to WAN Forwarding Groups

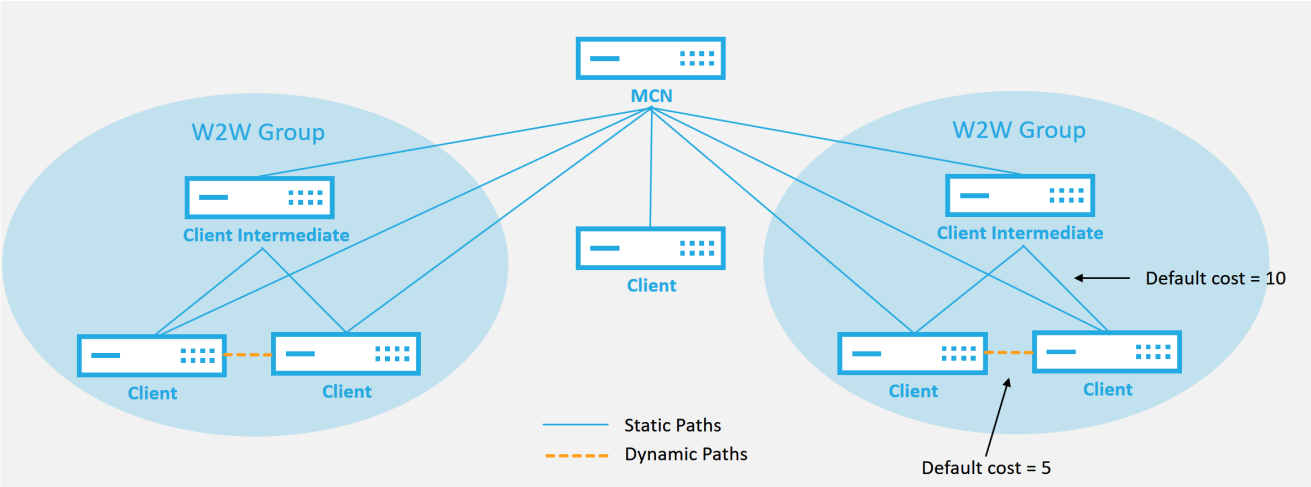

WAN to WAN Forwarding Group:

- A network can have multiple WAN to WAN Forwarding Groups
- Direct dynamic path will have a lower cost then through the intermediate node

CITRIX

For client nodes to operate as Intermediate sites, a static Virtual Path is required to be configured between it and the clients that are associated with that **WAN-to-WAN Forwarding Group**. In addition, client nodes will need **Enable Dynamic Virtual Path** option turned on for each client node.

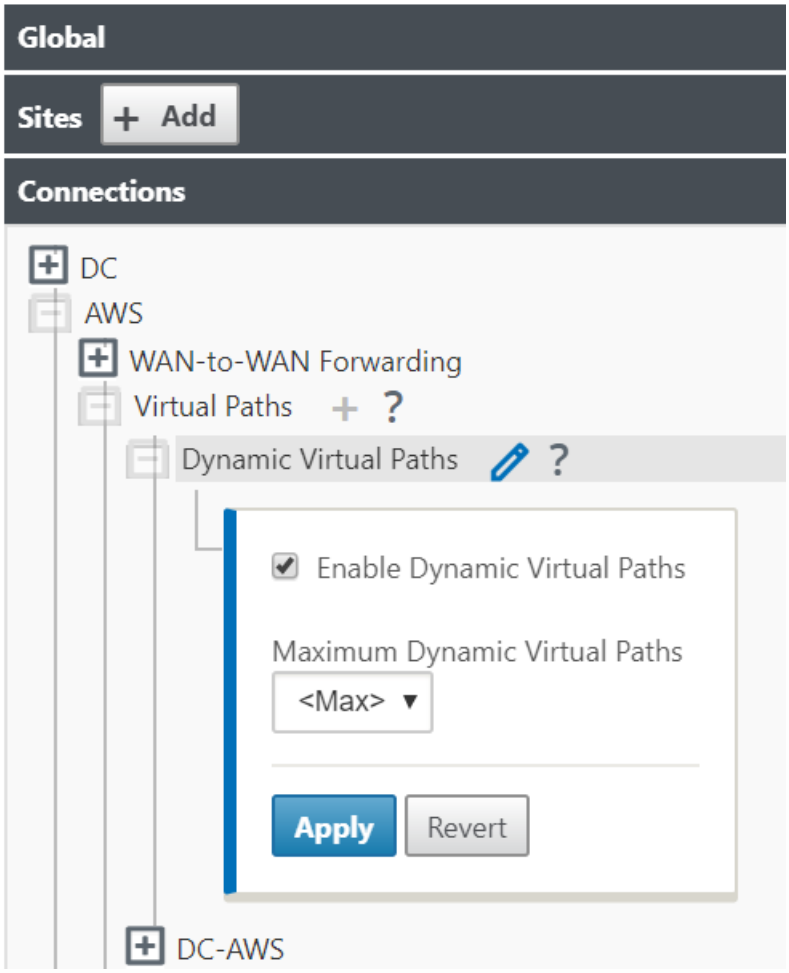

Each SD-WAN device will have its own unique route table with the following details defined for each route:

- Num – order of route of this appliance based on match process (lowest Num processed first)
- Network address – subnet or host address
- Gateway if required
- Service – what service is applied for this route
- Firewall Zone – the firewall zone classification of the route
- Reachable – Identifies if the Virtual Path state is active for this site
- Site – The name of the site where the route is expected to exist
- Type – Identification of route type (Static or Dynamic)
- Neighbor Direct
- Cost - cost of the specific route
- Hit Count – how many times the route has been used per packet. This would be used to verify that a route is being hit correctly.
- Eligible

- Eligibility Type
- Eligibility Value

Below is an example SD-WAN site route table:

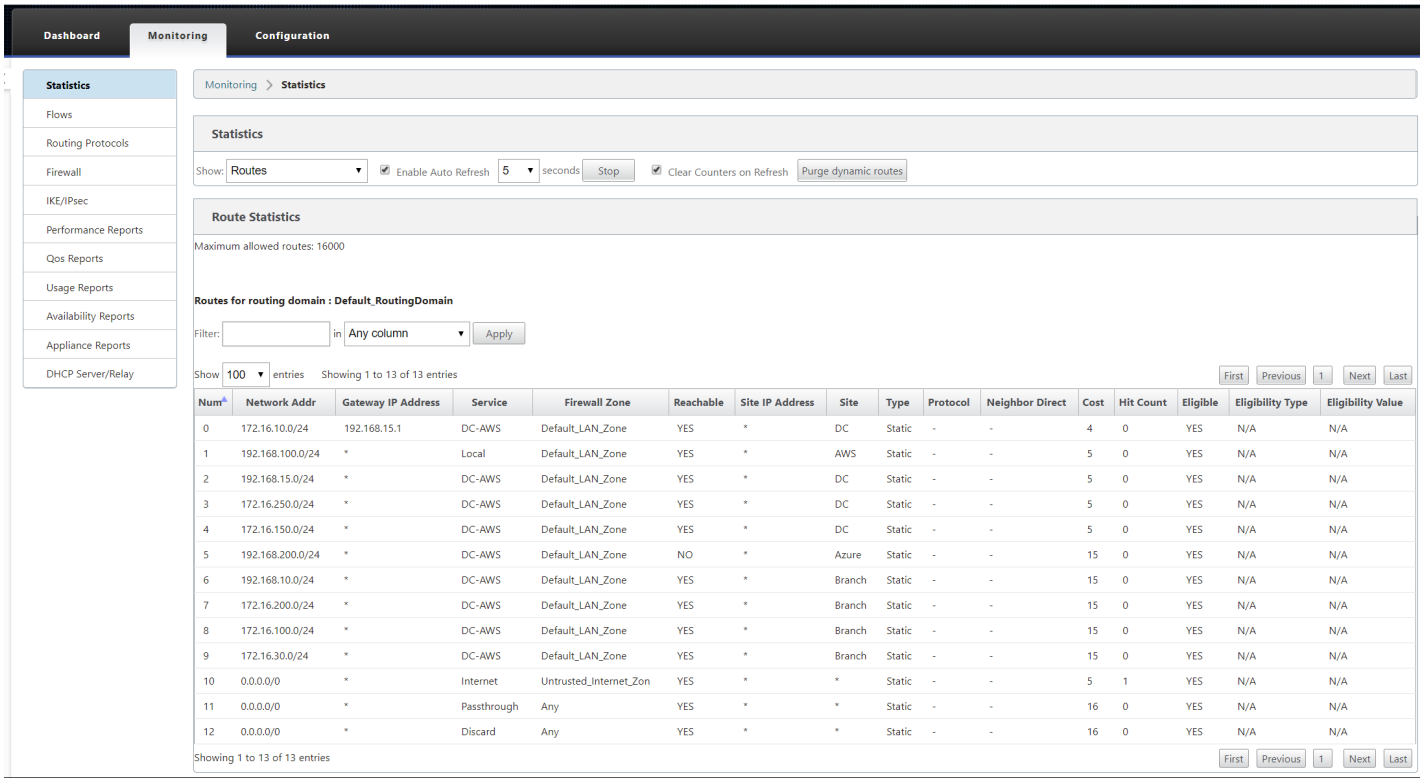

Notice from the above SD-WAN route table that there are additional elements not normally availability in traditional routers. Most notable is the “Reachable” column, which renders the route either active or inactive (yes/no) depending on the WAN path state. Routes listed here are suppressed based on various states of the service (the Virtual Path being down as an example). Other events that can force a route to be ineligible are path down state, next hop unreachable or WAN link down.

From the above table, we can see fourteen defined routes. A description of the routes or groups of routes is described below:

- Route 0 – On the MCN this is a Host subnet route that resides at the DC site. 172.16.10.0/24 resides in the DC LAN and 192.168.15.1 is the gateway on the LAN that is the next hop that will get to that subnet.
- Route 1 – This is a local route to this SD-WAN device that displaying the route table.
- Route 2-4 – These are the subnets that are part of the virtual interfaces configured for the DC site SD-WAN. These subnets are derived from the trusted virtual interfaces defined.
- Route 5 – This is a shared route to another client node that is shared by the MCN with a Reachability status of No due to the down Virtual Path between that site and the MCN.
- Route 6-9 – These routes exist at another client site. For this route, a Virtual Path route is created for matching WAN ingress traffic destined for the remote site on the Virtual Path.
- Route 10 – With the Internet Service defined, the system adds a catch all route for direct internet breakout for this local

site.

- Route 11 – Passthrough is default route the system always adds to allow packets to flow through in case there is no match on any existing routes. The Passthrough is not groomed, typically local broadcasts and ARP traffic will be mapped to this service.
- Route 12 – Discard is default route the system always adds to drop anything undefined.

Default Route Cost Values:

- WAN to WAN Forwarding – 10
- Default Direct Route Cost – 5
- Auto Generated Routes – 5
- Virtual Path – 5
- Local – 5
- Intranet – 5
- Internet – 5
- Passthrough – 5
- Optional – route is 0.0.0.0/0 defined as a service level

After defining these routes, it is important to understand how the traffic flows using the defined routes. These traffic flows will be broken into the following flows:

- LAN to WAN (Virtual Path) – Traffic going into the SD-WAN overlay tunnel
- WAN to LAN (Virtual Path) – Traffic existing the SD-WAN overlay tunnel
- Non-Virtual Path Traffic – Traffic routed to the underlay network

The default route cost can be altered on a per-site basis. The configuration can be found under **Sites** > **Site node** > **Basic Settings** node:

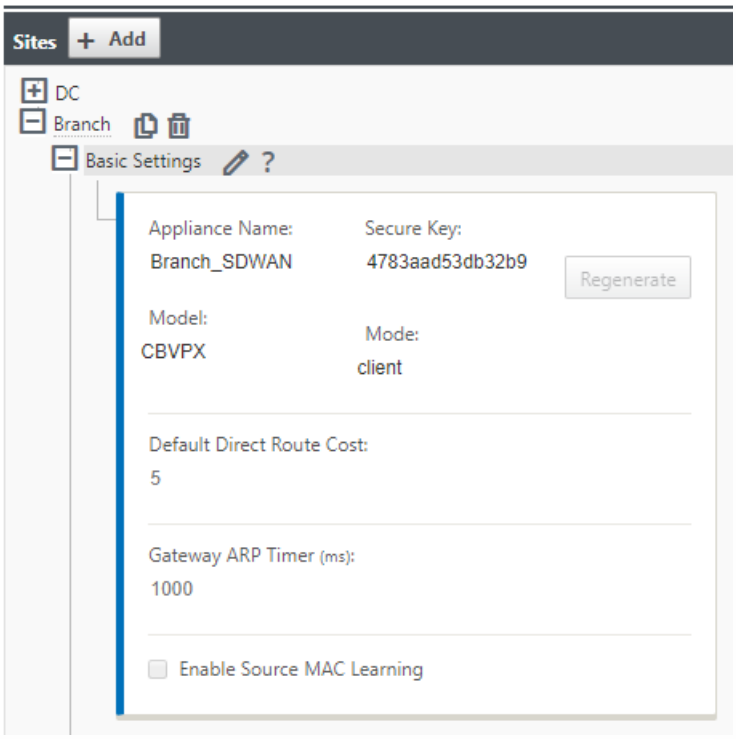

Static routes can be defined per site under the **Connections** > **Site** > **Route**s node:

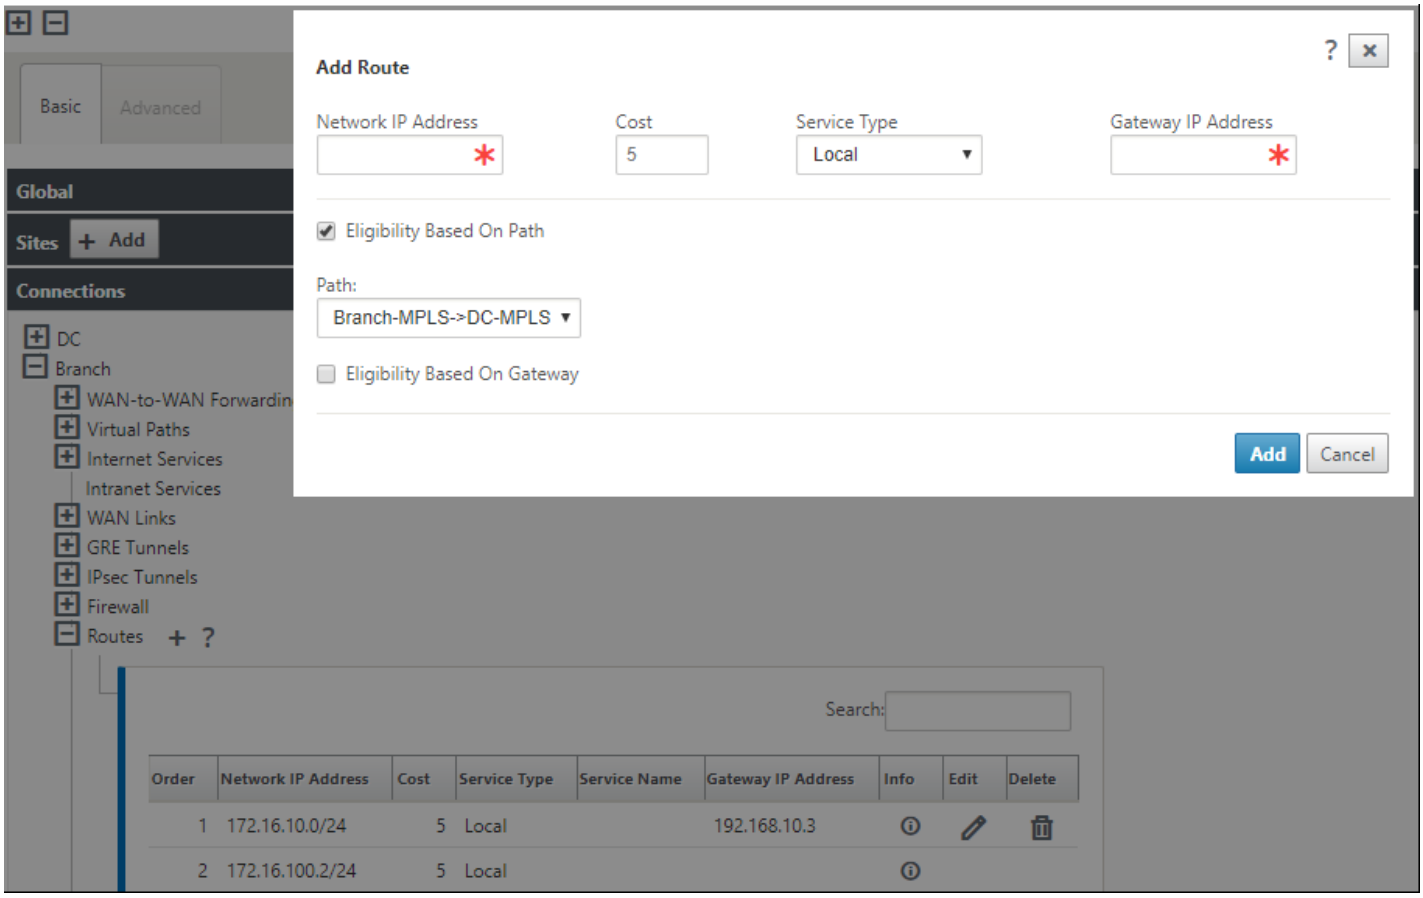

You will notice that routes can be tied to the Virtual Path or Gateway IP availability. Internet routes can be exported to the Virtual Path overlay or not depending on desired behavior. You can also create static Virtual Path routes to force traffic to a Virtual Path even though we are not getting the prefix advertised to SD-WAN (i.e. a higher cost route of last resort). SD-WAN can also suppress local subnets from being advertised by making the Virtual IP Address (VIP) private.

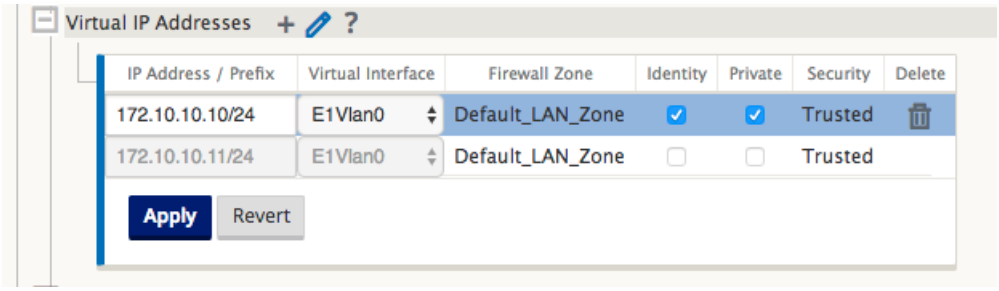

Note

The configuration does require at least one non-private VIP in each route domain.

Intranet and Internet Routes

For the Intranet and Internet service types, the user must have defined a SD-WAN WAN Link to support those types of services. This is a pre-requisite for any defined routes for either of these services. If the WAN link is not defined to support the Intranet Service it will be considered as a local route. The Intranet, Internet and Passthrough routes are only relevant to the site/appliance they are configured for.

When defining Intranet, Internet or Passthrough routes the following are design considerations:

- Must have service defined on the WAN link (Intranet/Internet – required)
- For Intranet/Internet must have gateway defined for the WAN link
- Relevant to local SD-WAN device
- Intranet routes can be learned via the Virtual Path but are done so at a higher cost
- With Internet Service, there is automatically a default route created (0.0.0.0/0) catch all route with a max cost
- Do no assume Passthrough will work, this should be tested/verified, also test with Virtual Path down/disabled to verify desired behavior
- Route tables are static unless route learning feature is enabled

As of NetScaler SD-WAN release 9.3, please find below the maximum supported limit for multiple routing parameters:

- Maximum Routing Domains: 255
- Maximum Access Interfaces per WAN Link: 64
- Maximum BGP neighbors per site: 255
- Maximum OSPF area per site: 255
- Maximum Virtual Interfaces per OSPF area: 255
- Maximum Route Learning import filters per site: 512

- Maximum Route Learning export filters per site: 512
- Maximum BGP routing policies: 255
- Maximum BGP community string objects: 255

High Availability Deployment

Apr 12, 2018

This topic covers the High Availability (high availability) deployments and configurations supported by SD-WAN appliances (Standard Edition and Enterprise Edition).

SD-WAN appliances can be deployed in high availability configuration as a pair of appliances in Active/Standby roles. There are three modes of high availability deployment:

- Parallel Inline high availability
- Fail-to-Wire high availability
- One-Arm high availability

These high availability deployment modes are similar to Virtual Router Redundancy Protocol (VRRP) and use a proprietary SD-WAN protocol. Both Client Nodes (Clients) and Master Control Nodes (MCNs) within an SD-WAN network can be deployed in a high availability configuration as long as the selected SD-WAN platform model supports high availability.

In high availability configuration, one SD-WAN appliance at the site is designated as the Active appliance and is continuously monitored by the Standby appliance. Configuration is mirrored across both appliances. When the Standby appliance loses connectivity with the Active appliance for a defined period, the Standby appliance assumes the identity of the Active appliance and takes over the traffic load. Depending on the deployment mode, this fast failover has minimal impact on the application traffic passing through the network.

High Availability Deployment Modes

One-Arm mode:

In One-Arm mode, the high availability appliance pair is outside of the data path. Application traffic is redirected to the appliance pair with Policy Based Routing (PBR). One-Arm mode is implemented when a single insertion point in the network is not feasible or to counter challenges of fail-to-wire. In the following illustration, the Standby appliance can be added to the same VLAN or subnet as the Active appliance and the router.

In One-Arm mode, it is recommended that the SD-WAN appliances do not reside in the data network subnets. The virtual path traffic does not have to traverse the PBR and avoids route loops. The SD-WAN appliance and router have to be directly connected, either through an Ethernet port or be in the same VLAN.

IP SLA Monitoring for Fall Back

The active traffic flows even if the virtual path is down, as long because one of the SD-WAN appliances is active. The SD-WAN appliance redirects traffic back to the router as Intranet traffic. However, if both active/standby SD-WAN appliances become inactive, the router tries to redirect traffic to the appliances. IP SLA monitoring can be configured at the router to disable PBR, if the next appliance is not reachable. This allows the router to fall back to perform a route lookup and forward packets appropriately.

Parallel Inline high availability mode:

In Parallel Inline high availability mode, the SD-WAN appliances are deployed alongside each other, inline by using the data path. Only one path through the Active appliance is used. It is important to note that
bypass interface groups are configured to be fail-to-block and not fail-to-wire so that you don't get bridging loops during a failover.

The high availability state can be monitored through the inline interface groups, or through a direct connection between the appliances. External Tracking can be used to monitor the reachability of the upstream or downstream network infrastructure. For example; switch port failure) to direct high availability state change, if needed.

If both active and standby SD-WAN appliances are disabled or fail, a tertiary path can be used directly between the switch and router. This path must have a higher spanning tree cost than the SD-WAN paths so that it is not used under normal conditions. Failover in parallel inline high availability mode is very quick and nearly hitless, because no physical state change occurs. Fallback to the tertiary path is not hitless and can block traffic for 5-30 seconds depending on the spanning tree configuration. If there are out of path connections to other WAN Links, both appliances must be connected to them.

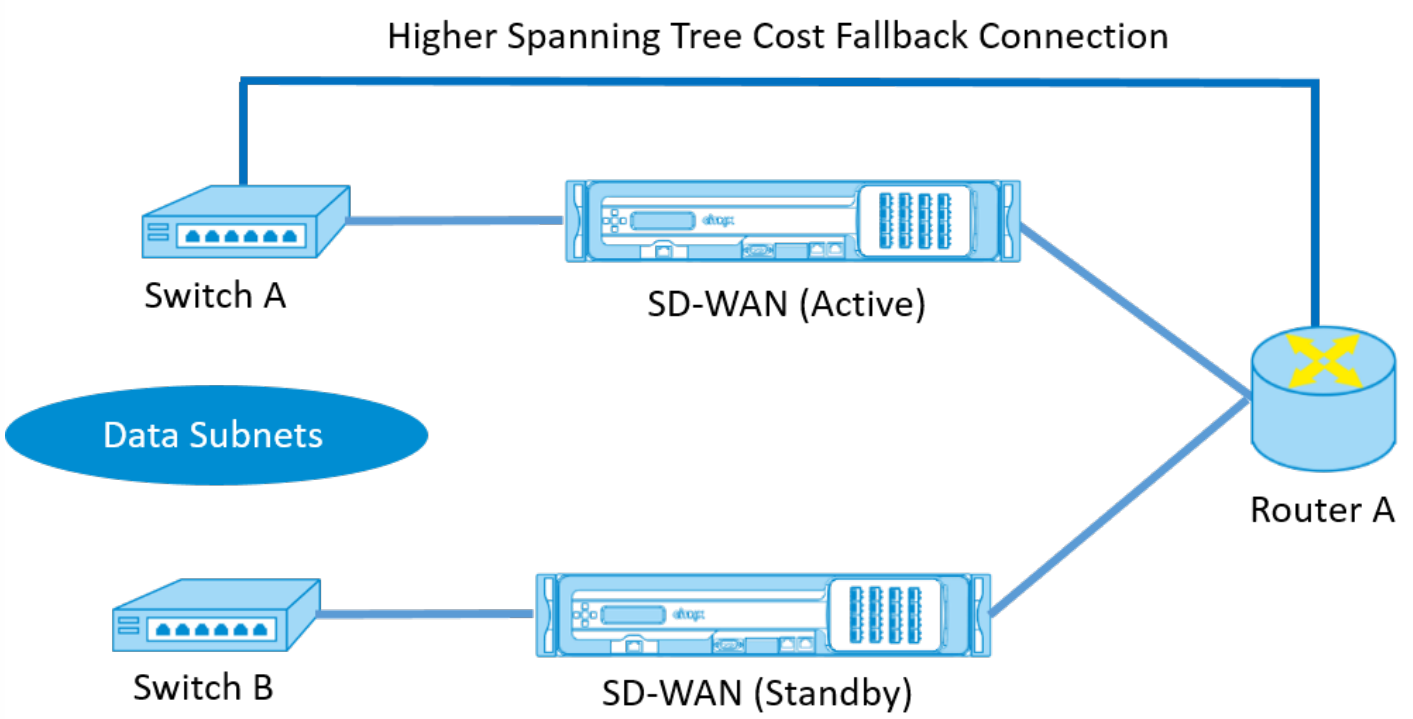

In more complex scenarios, where multiple routers might be using VRRP, non-routable VLANs are recommended to ensure the LAN side switch and routers are reachable at layer 2.

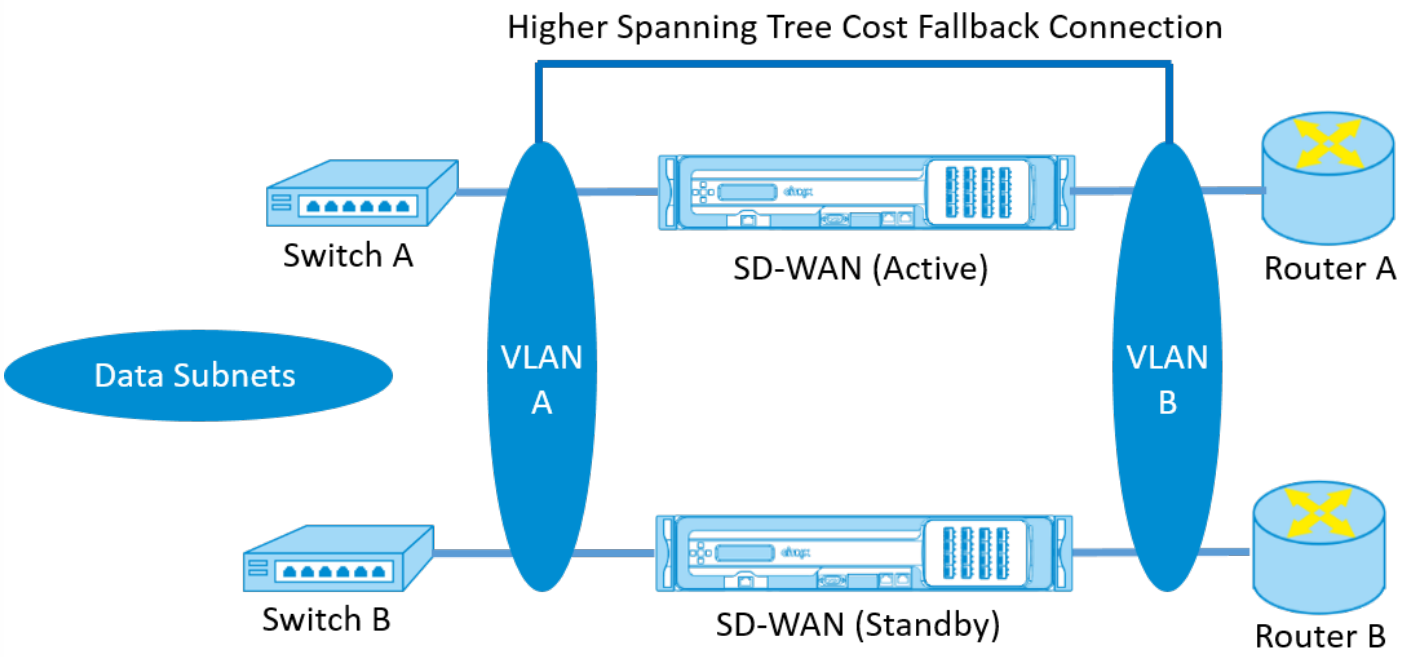

Fail-to-Wire mode:

In fail-to-wire mode, the SD-WAN appliances are inline in the same data path. The bypass interface groups must be in the fail-to-wire mode by using the Standby appliance in a passthrough or bypass state. A direct connection among the two appliances on a seperate port must be configured and used for the high availability interface group.

Note

- High availability switchover in fail-to-wire mode takes longer period, approximately 10–12 seconds because of delay in ports to recover from Fail-to-Wire state.
- When the high availability connection between the appliances fails, both appliances go into Active state and cause a service interruption. This can be mitigated by assigning multiple high availability connections so that there is no single point of failure.
- It is imperative that in high availability Fail-to-Wire Mode, a seperate port be used in the hardware appliance pairs for high availability control exchange mechanism to assist in state convergence.

- Because of a physical state change if the SD-WAN appliances switch over from Active to Standby, failover can cause partial loss of connectivity depending on how long the auto-negotiation takes on the Ethernet ports.
- It is recommended that Fail-to-Wire mode be used on ports that are auto-negotiated, because this increases failover time.

The following illustration shows an example of the Fail-to-Wire deployment.

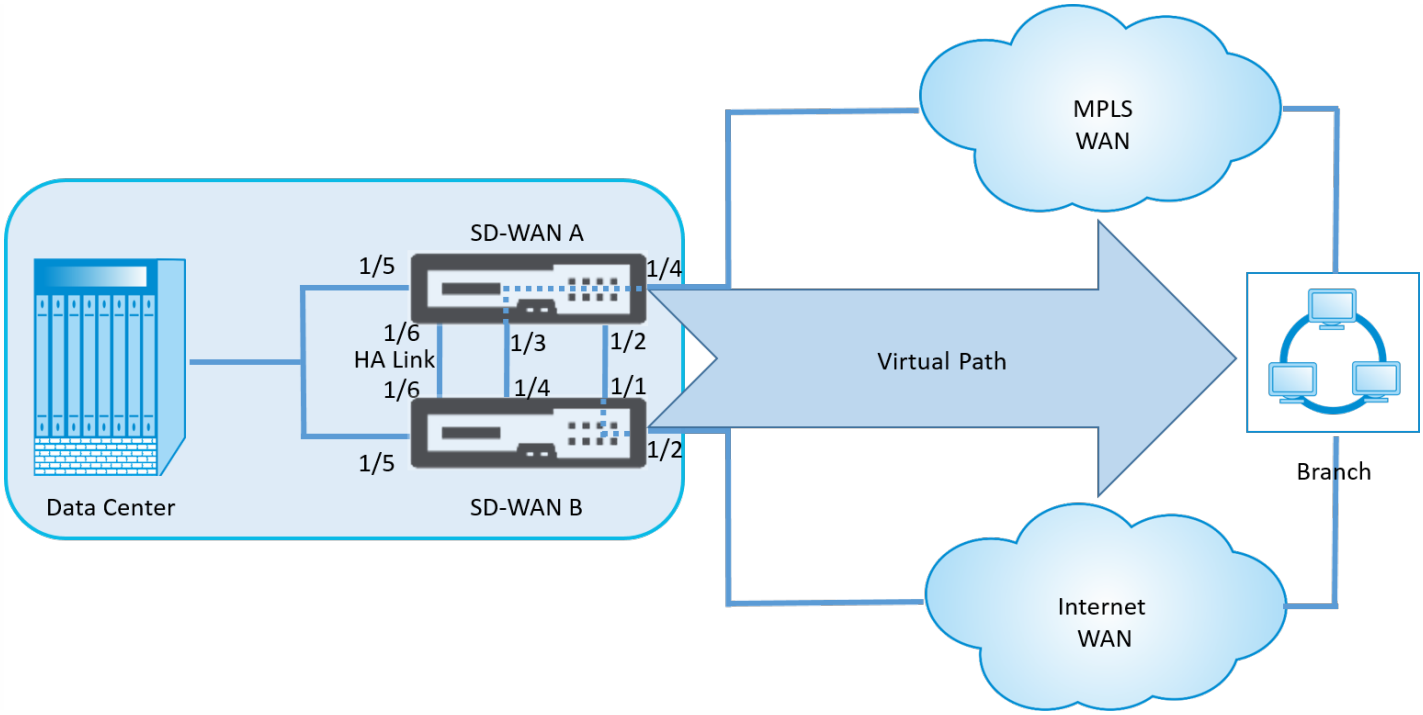

The One-Arm high availability configuration or Parallel Inline high availability configuration is recommended for data centers or Sites that forward a high volume of traffic to minimize disruption during failover.

If minimal loss of service is acceptable during a failover, then Fail-to-Wire high availability mode is a better solution. The Fail-to-Wire high availability mode protects against appliance failure and parallel inline high availability protects against all failures. In all scenarios, high availability is valuable to preserve the continuity of SD-WAN network during a system failure.

Configuring High Availability

To configure high availability:

1. In the Configuration Editor, navigate to **Sites >** site name **> High Availability.** Select **Enable High Availability**.

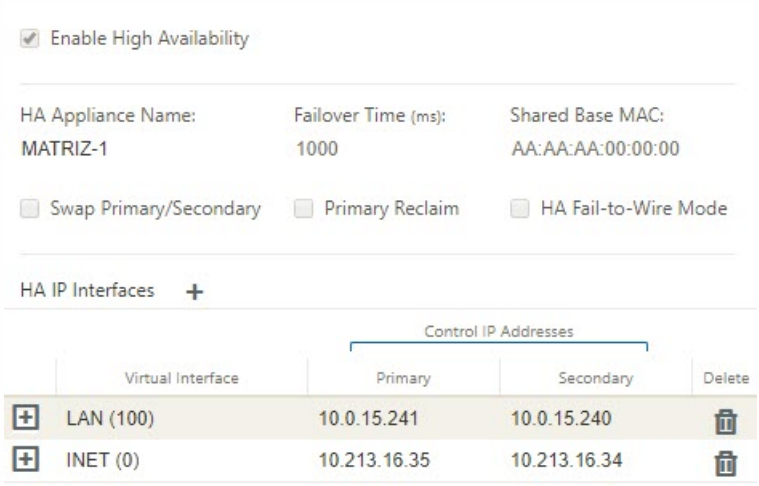

2. Type values for the following parameter:

- **High availability Appliance Name:** This is the name of the high availability (secondary) appliance.
- **Failover Time:** This specifies the wait time (in milliseconds) after contact by using the primary appliance is lost, before the standby appliance becomes active.

- **Shared Base MAC:** This is the shared MAC address for the high availability pair appliances. If a failover occurs, the secondary appliance has the same virtual MAC addresses as the failed primary appliance.
- **Swap Primary/Secondary:** When this is selected, if both appliances in the high availability pair come up simultaneously, the secondary appliance becomes the primary appliance, and takes precedence.
- **Primary Reclaim:** If this is selected, the designated primary appliance reclaims control upon restart after a failover event.
- **HA Fail-to-Wire Mode:** Choose this for Fail-to-wire high availability deployment mode.

> ## Note
>
> For hypervisor and cloud based platforms an extra parameter **Disable Shared Base MAC** is available. Choose this to disable the shared virtual MAC address.

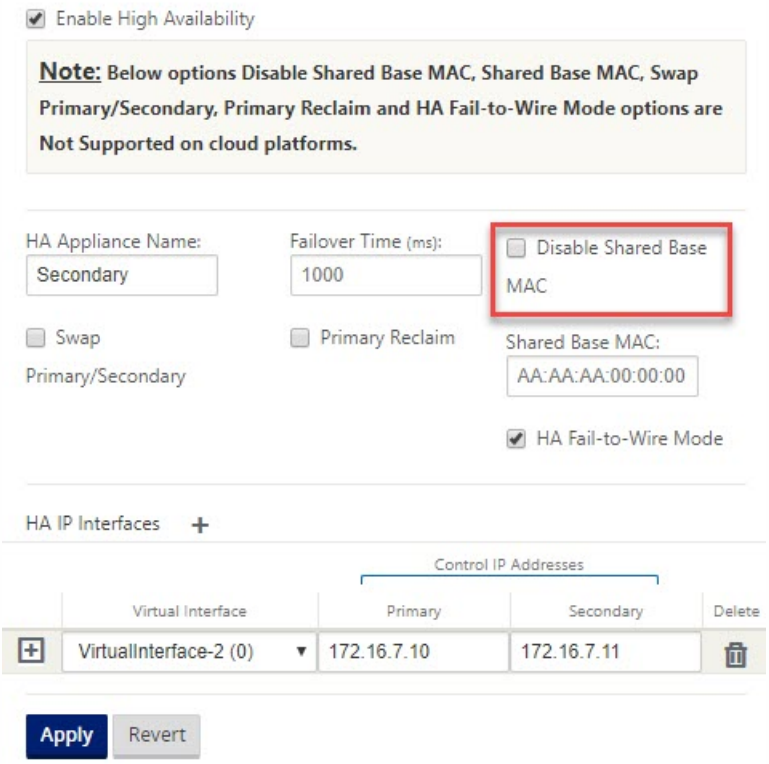

> ## Note
>
> For hypervisor based platforms ensure that the promiscuous mode is enabled on the hypervisors to allow packet sourcing from high availability shared MAC address. When promiscuous mode is not enabled, you can enable Disable **Shared Base MAC**. option.

3. Click + next to **HA IP Interfaces** to configure interface groups. Enter Values for the following parameters:

- **Virtual Interface** – This is the Virtual Interface to be used for communication among the appliances in the high availability pair. This interface monitors the Active appliance for reachability. For One-Arm high availability mode, only one interface group is required.
- **Primary** – This is the unique Virtual IP address for the primary appliance. The secondary appliance uses this for

communication by using the primary appliance.

- **Secondary** – This is the unique Virtual IP address for the secondary appliance. The primary appliance uses this for communication by using the secondary appliance.

> **Note**
>
> For Inline high availability mode, extra interface groups are required for **External Tracking** to monitor the upstream or downstream network infrastructure. For example. Switch port failure, to detect when high availability change state is required.

4. Click **+** to the left of the new **HA IP Interfaces** entry. In the **External Tracking IP Address** field, enter the IP Address of the external device that responds to ARP requests to determine the state of the primary appliance.

5. Choose **Apply**.

Monitoring

To monitor high availability configuration:

Log in to the SD-WAN web management interface for the Active and Standby appliance's for which high availability is implemented. View high availaiblity status under the **Dashboard** tab.

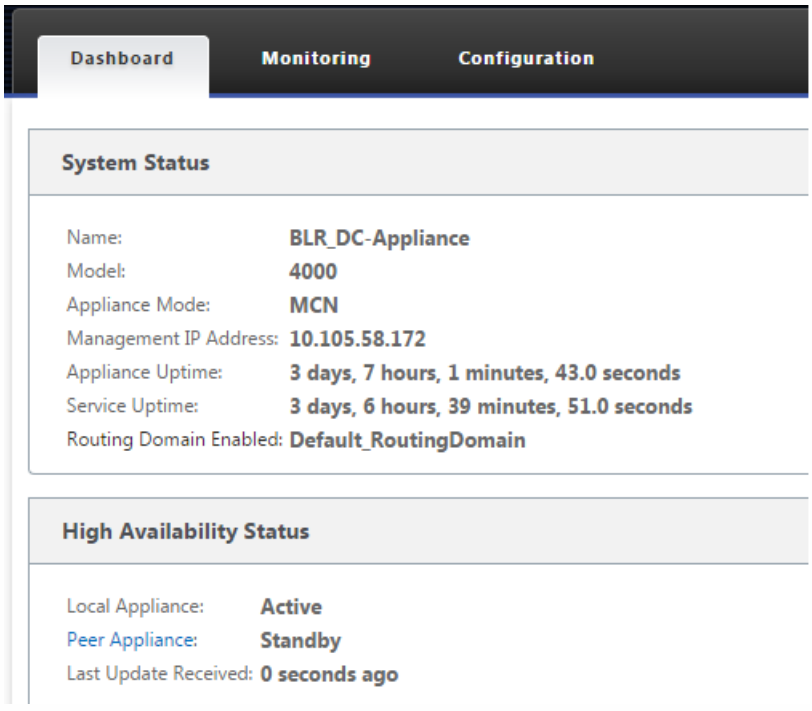

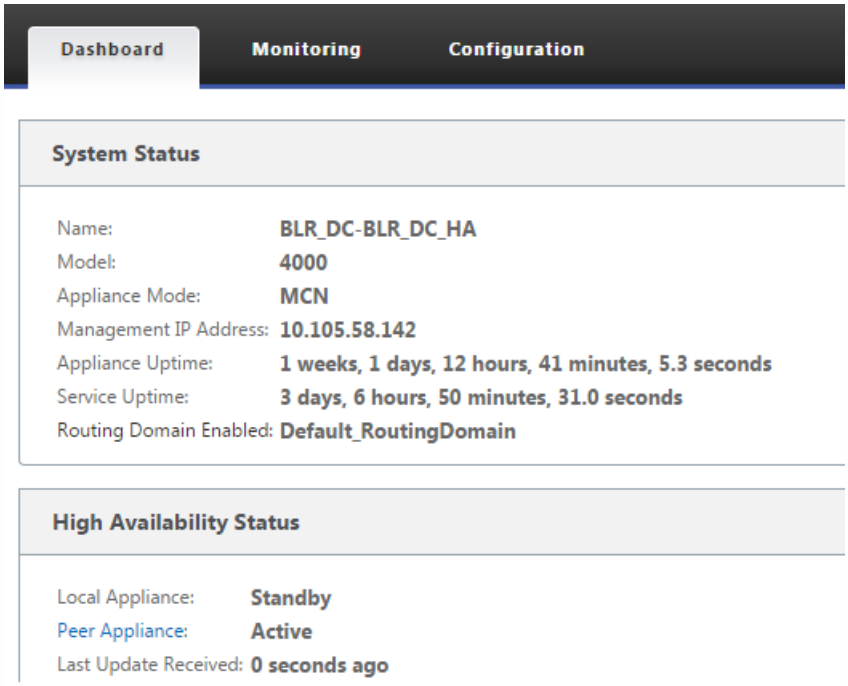

For Network Adapter details of Active and Standby high availability appliances, navigate to **Configuration** > **Appliance Settings** > **Network Adapters** > **Ethernet** tab.

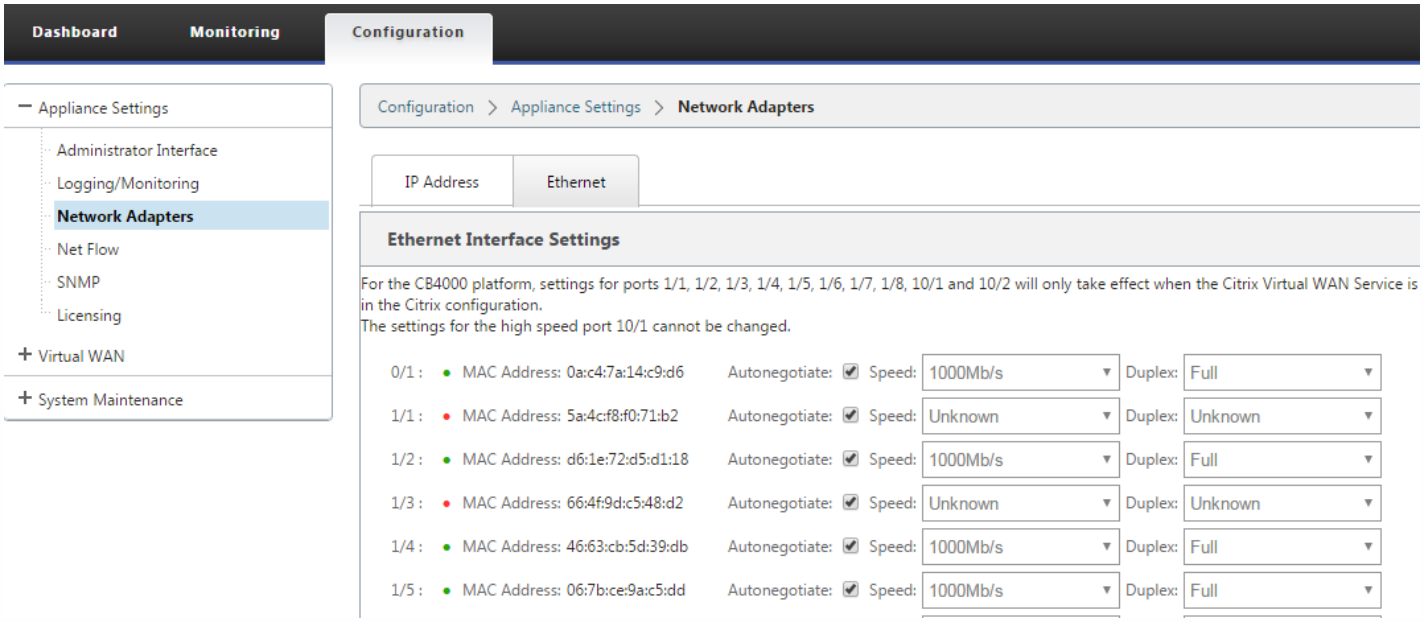

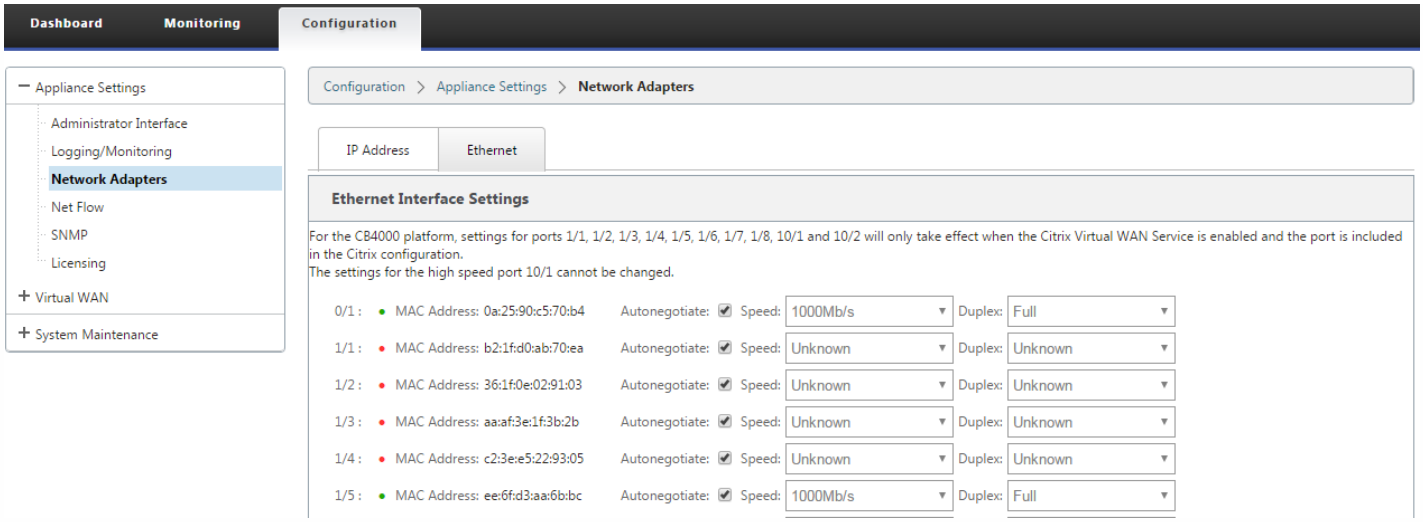

Configuration

Aug 09, 2017

The following topics provide information about how to configure Virtual path service between MCN and branch sites, and enabling WAN optimization.

Configuring Virtual WAN service

Configuring virtual path between MCN and branch sites

Enabling and configuring WAN optimization

Configuring Virtual WAN Service

Aug 09, 2017

The SD-WAN configuration describes and defines the topology of your SD-WAN network. Before you can deploy a SD-WAN network, you must define the Virtual WAN configuration. To do this, use Configuration Editor in the SD-WAN Management Web Interface on the MCN appliance.

Security and Encryption

Enabling encryption for SD-WAN (for the Virtual Paths) is optional. Instructions for configuring this feature are provided in the section, Enabling and Configuring Virtual WAN Security and Encryption (Optional).

When encryption is enabled, SD-WAN uses the Advanced Encryption Standard (AES) to secure traffic across the Virtual Path. Both AES 128 and 256 bit ciphers (key sizes) are supported by the SD-WAN Appliances, and are configurable options. You can select, enable, and configure these and the other encryption options by using the Configuration Editor in the Management Web Interface on the Management Control Node (MCN). You must have administrative access on the MCN to modify the configuration, and to distribute your changes across the SD-WAN network. Once the MCN is secured, the encryption settings and their distribution are also secure.

Authentication between sites functions by means of the Virtual WAN Configuration.
The network configuration has a secret key for each site. For each Virtual Path, the network configuration generates a key by combining the secret keys from the sites at each end of the Virtual Path. The initial key exchange that occurs after a Virtual Path is first set up, is dependent upon the ability to encrypt and decrypt packets by means of that combined key.

Enabling Virtual WAN Service

If this is an initial installation and configuration, as a final step you will need to manually enable the Virtual WAN Service on each SD-WAN appliance in your network. Enabling the service enables and starts the Virtual WAN daemon.

> **Note**
>
> If you are reconfiguring an existing deployment, the MCN automatically enables the service when it distributes the updated Appliance Packages to the client sites. In this case, you can skip this final step.

To manually enable the Virtual WAN Service on an appliance, do the following:

1. Log into the Management Web Interface on the appliance you want to activate.
2. Select **Configuration** tab.
3. In the navigation pane, open the Virtual WAN branch and select **Enable/Disable/Purge Flows**.

If the Virtual WAN Service is currently disabled, this displays the Enable Virtual WAN Service page, as shown below. If the service is already enabled, this displays the Enable/Disable/Purge Flows page.

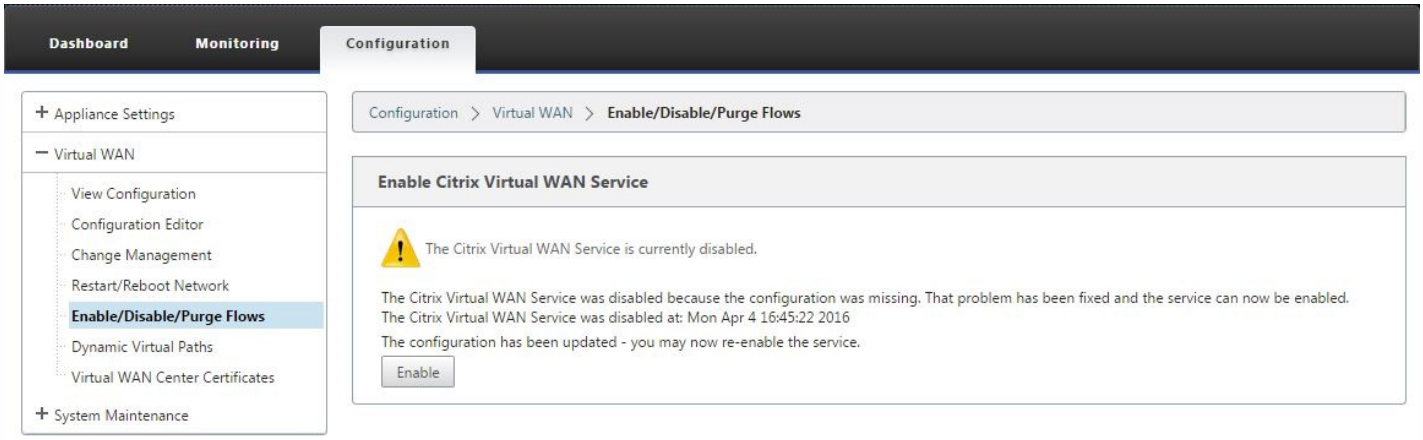

4. Click **Enable**.

   This enables the service, and displays the **Enable/Disable/Purge Flows** page.

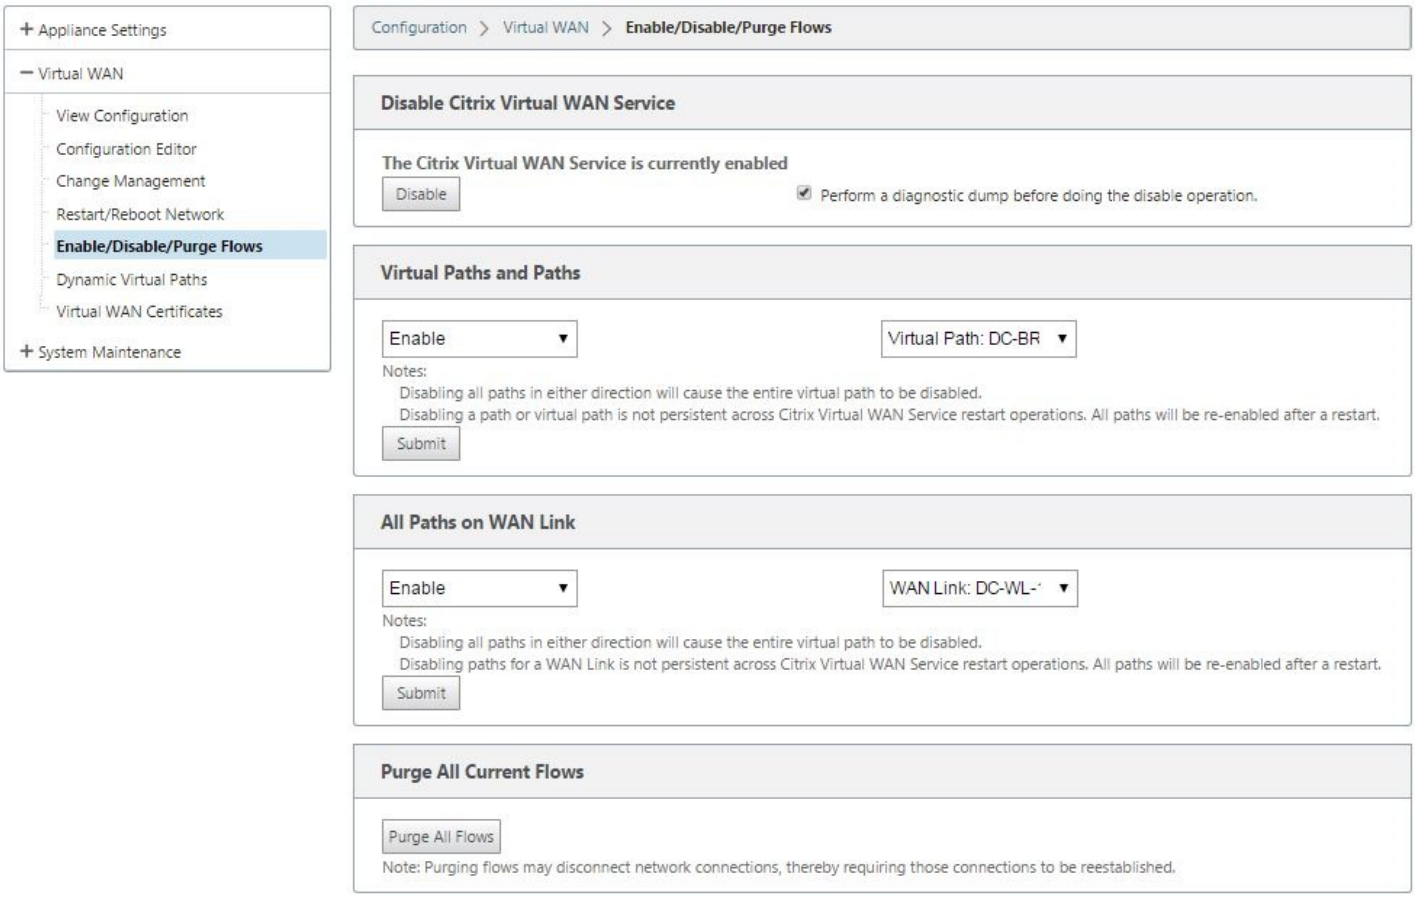

When the Virtual WAN Service is enabled, a status message to that effect displays in the top section of the page.

Note

This page also presents options for enabling/disabling specific paths and Virtual Paths in your network, as well as an option to purge all flows.

This completes the installation and activation of the SD-WAN on the MCN and branch site client appliances. You can now

use the Monitoring pages to verify the activation and diagnose any existing or potential configuration issues.

Configuring the Virtual Path Service Between the MCN and Client Sites

Aug 09, 2017

The next step is to configure the Virtual Path Service between the MCN and each of the client (branch) sites. To do this, you will use the configuration forms and settings available in the **Connections** section configuration tree of the **Configuration Editor**.

To configure the Virtual Path Service between the MCN and a client site, do the following:

1. Continuing in the **Configuration Editor**, click the **Connections** section heading.

   This reveals the **Connections** section configuration tree.

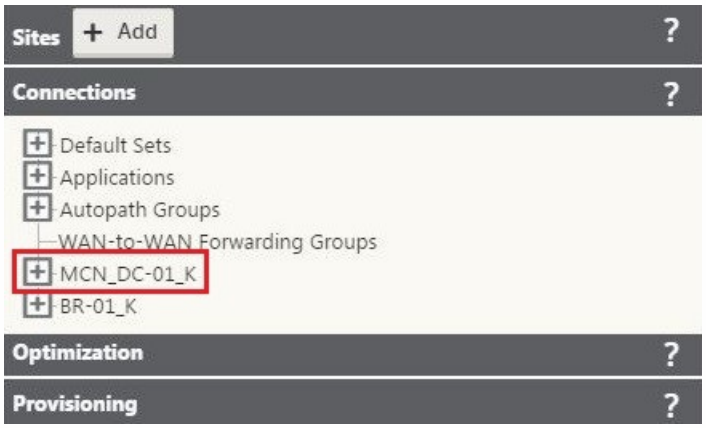

2. Click **+** to the left of the MCN site name in the **Connections** section tree.

   This opens the MCN site branch in the **Connections** configuration tree.

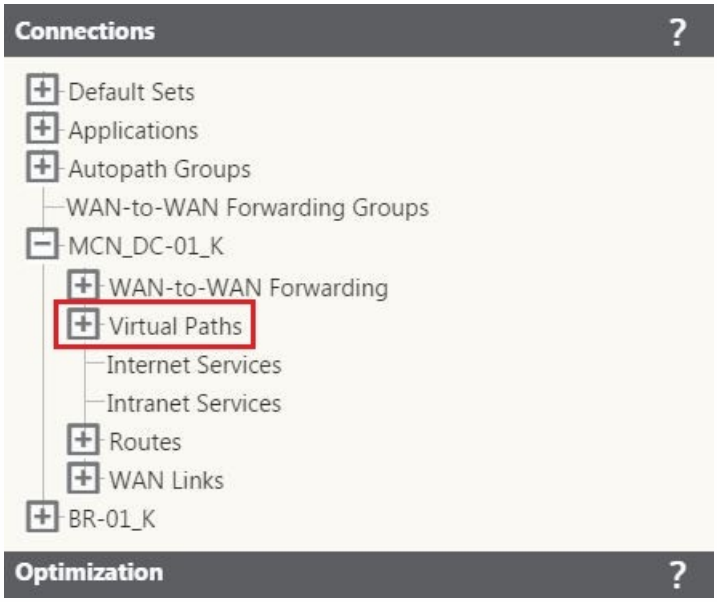

3. Click **+** next to the **Virtual Paths** branch label.

   This opens the **Virtual Paths** configuration section (child branch) for the MCN site branch in the tree. This section

provides settings and forms for configuring the Virtual Path Service between the MCN and each of the Virtual WAN client sites.The below figure shows an example **Virtual Paths** section for an MCN site.

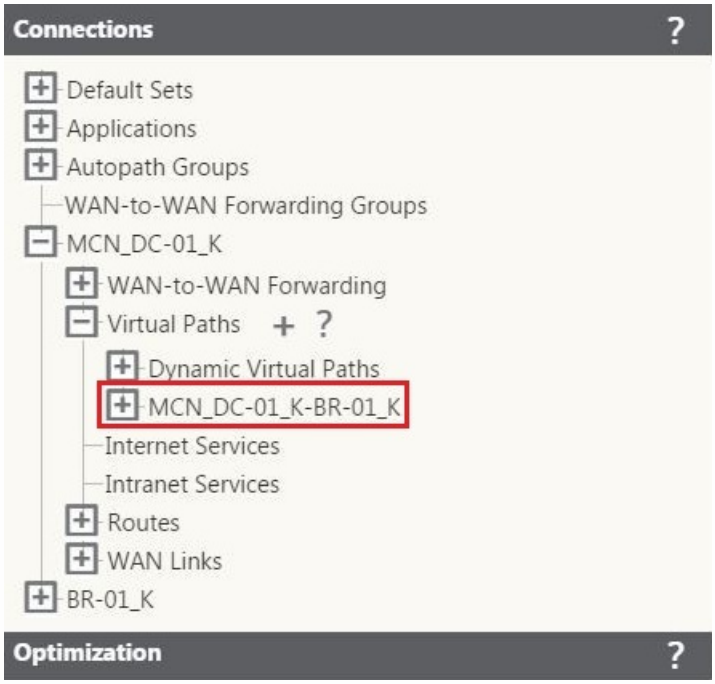

The Virtual Paths section contains the following child branches:

- **Dynamic Virtual Paths** – (Optional) The settings in this section allow you to enable and disable Dynamic Virtual Paths, and set the maximum allowable Dynamic Virtual Paths for the site. Dynamic Virtual Paths are Virtual Paths that are established directly between sites, based on a configured threshold. The threshold is typically based on the amount of traffic occurring between those sites. Dynamic Virtual Paths are operational only after the specified threshold is reached. Dynamic Virtual Paths are not required for normal operation, so configuring this section is optional.

- <**MCN_Site_Name**>_<**Client_Site_Name**> – The system initially automatically adds a static Virtual Path between the MCN and a client site, as this Virtual Path is required. The name for the path uses the following form:

<MCN_Site_Name>_<Client_Site_Name>

Where:

* *MCN_Site_Name* is the name of the MCN for this Virtual WAN.

* *Client_Site_Name* is the name of a client site identified in the current configuration package.

User-configurable default settings are initially applied to the static Virtual Path, as defined in the **Default Sets** > **Virtual Path Default Sets** section of the **Connections** configuration tree. However, you can customize or add to the defined **Default Sets**, and also customize the configuration for a specific site and Virtual Path.

> ## Note
>
> To add more static Virtual Paths for a site, you must do so manually. Instructions for manually adding a static Virtual Path are included in the steps below.

4. Click **+** next to the name of the static Virtual Path in the **Virtual Paths** tree.

This reveals additional child branches (configuration sets) for the static Virtual Path, as follows:

- **Local Site** – This section enables you to view and configure the Virtual Path settings from the perspective of a local site. You can also view, customize, and add **Classes** or **Rules** as required for this specific Virtual Path. You can also add Virtual Paths to the local site, as needed.

- **Remote Site** – This section enables you to view and configure the Virtual Path settings from the perspective of a remote site. As for the Local Site section, you can also view, customize, and add **Class** or **Rules** as required for this specific Virtual Path. You can also add Virtual Paths to the remote site, as needed.

- **Paths** – This section provides settings and forms for configuring the Virtual Paths and Virtual Path Service between the MCN and the client site.

The below figure shows an example MCN static Virtual Path branch and child branches.

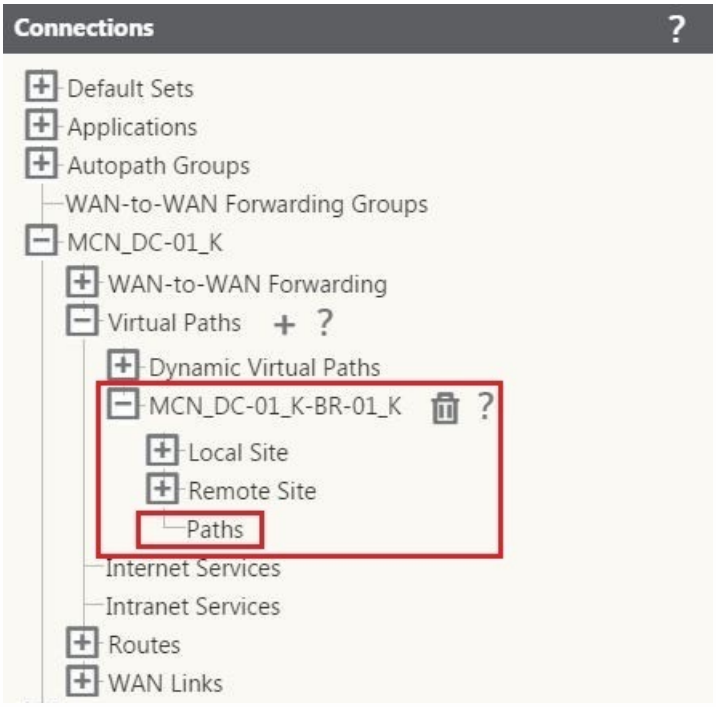

5. Click **+** next to the **Paths** branch label, or click **Paths** if no **+** icon is present.

The system is able to automatically generate Virtual Paths for some WAN Links. If so, there will be a plus sign (**+**) to the left of the **Paths** child branch label. If the **+** icon is not present, click the **Paths** label directly. This reveals the Add (**+**) button for adding a Virtual Path. The Add icon is located to the right of the **Paths** branch label.

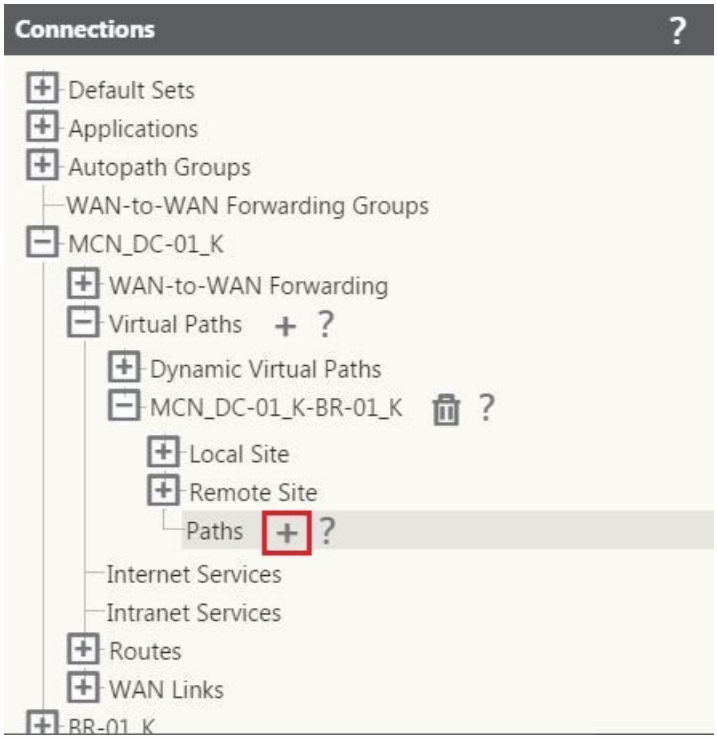

6. Click **+** (Add) to the right of the **Paths** branch label.

   This displays the **Add Path** dialog box (configuration form).

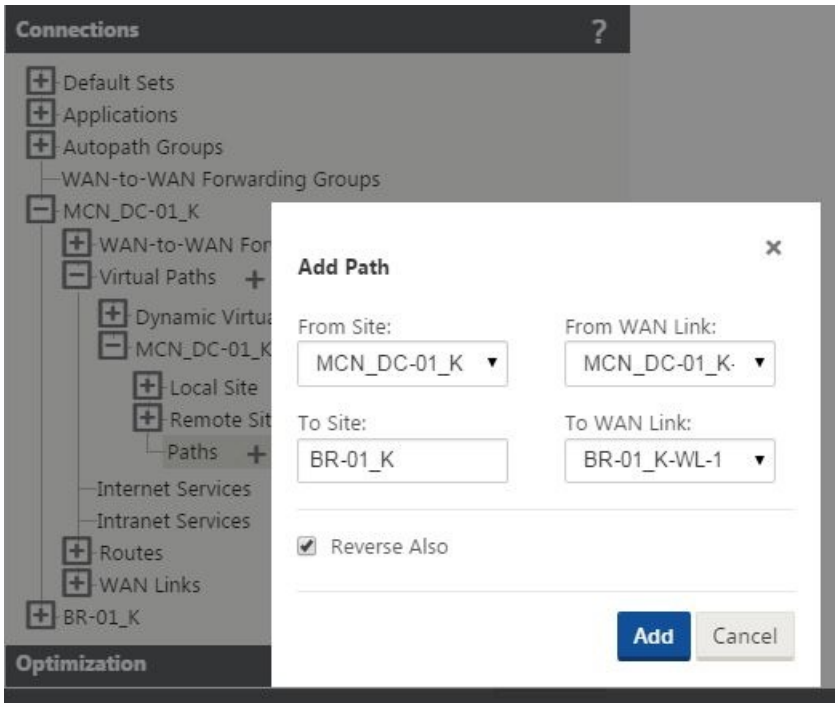

7. Specify the source and destination site information for the new Virtual Path.

   Specify the following from the available drop-down menus:

> ## Note
>
> Depending on how the WAN links are configured for the sites, some fields will be read-only. Fields that are configurable provide a

drop-down menu of the available selections.

- **From Site** – This is the source site for the Virtual Path. For the required static Virtual Path, this is configured as the MCN site by default.

- **From WAN Link** – This is the originating WAN Link for the Virtual Path.

- **To Site** – This is the destination site for the Virtual Path.

- **To WAN Link** – This is the destination WAN link for the Virtual Path.

8. Click **Add**.

This adds the configured Virtual Path to both the MCN and the associated client site in the **Connections** tree. This also automatically opens the **Paths** settings configuration form for the **From Site** for the Virtual Path (in this case, the MCN).

9. Click Edit (pencil icon), to the right of the MCN-to-client Virtual Path label.

This opens the Virtual Path Service configuration form for editing.

10. Configure the settings for the Virtual Path, or accept the defaults.

The **Paths** configuration form contains the following settings:

- **From Site** section:

  * **Site** – This is the source site for the Virtual Path. For the required static Virtual Path, this is configured as the MCN site by default.

  * **WAN Link** – This is the originating WAN Link for the Virtual Path.

- **To Site** section:

  * **Site** – This is the destination site for the Virtual Path.

  * **WAN Link** – This is the destination WAN link for the Virtual Path.

- **Reverse Also** – Select this checkbox to enable Reverse Also for this Virtual Path. If enabled, the system automatically builds a Virtual Path in the opposite direction of the configured path, using the same WAN links as configured for the original path.

- **IP DSCP Tagging** – Select a tag from the drop-down menu. This specifies the DSCP tag to set in the IP header for traffic traveling over this Virtual Path.

- **Enable Encryption** – Select this checkbox to enable encryption of packets sent along this Virtual Path.

- **Bad Loss Sensitive** – Select a setting from the drop-down menu. The options are:

  * **Enable**– (Default) If enabled, paths will be marked **BAD** due to loss, and will incur a path scoring penalty.
  * **Disable** – Disabling **Bad Loss Sensitive** can be useful when the loss of bandwidth is intolerable.

  * **Custom** – Select Custom to specify the percentage of loss over time required to mark a path as BAD. Selecting this option reveals the following additional settings:

    . **Percent Loss (%)** – This specifies the percentage of loss threshold before a path is marked BAD, as measured over the specified time. By default, the percentage is based on the last 200 packets received.

    . **Over Time (ms)** – Specify the time period (in milliseconds) over which to measure packet loss. Select an option between 100 and 2000 from the drop-down menu for this field.

- **Silence Period (ms)** – This specifies the duration (in milliseconds) before the path state transitions from **GOOD** to **BAD**. The default is 150 milliseconds. Select an option between 150 and 1000 from the drop-down menu for this field.

- **Path Probation Period (ms)** – This specifies the wait time (in milliseconds) before a path transitions from BAD to GOOD. Select an option between 500 and 60000 from the drop-down menu for this field. The default is 10000 milliseconds.

- **Instability Sensitive** – Select this checkbox to enable. If enabled, latency penalties due to a path state of **BAD** and other latency spikes are considered in the path scoring algorithm.

- **Tracking IP Address** – Enter a Virtual IP Address on the Virtual Path that can be pinged to determine the state of the path.

- **Reverse Tracking IP Address** – If **Reverse Also** is enabled for the Virtual Path, enter a Virtual IP Address on the path

that can be pinged to determine the state of the reverse path.

11. Click the minus sign (**–**) to the left of the path label to close the settings form.

This reveals that the two new **From Site** and **To Site** Virtual Paths between the MCN and the client site have been added to the tree.

12. Repeat the steps above for each branch you want to connect to the MCN.

13. (Optional) When you finish adding and configuring all of the Virtual Paths to and from the MCN, click the minus sign (–) to the left of the MCN branch label to close the branch.

Next, you have the option of customizing the Virtual Paths configurations for the client sites, as well as adding and configuring additional paths between clients. Instructions are provided in the remaining steps, below.

14. Click + to the left of the client site branch label in the tree.

This opens the client site branch in the **Connections** tree.

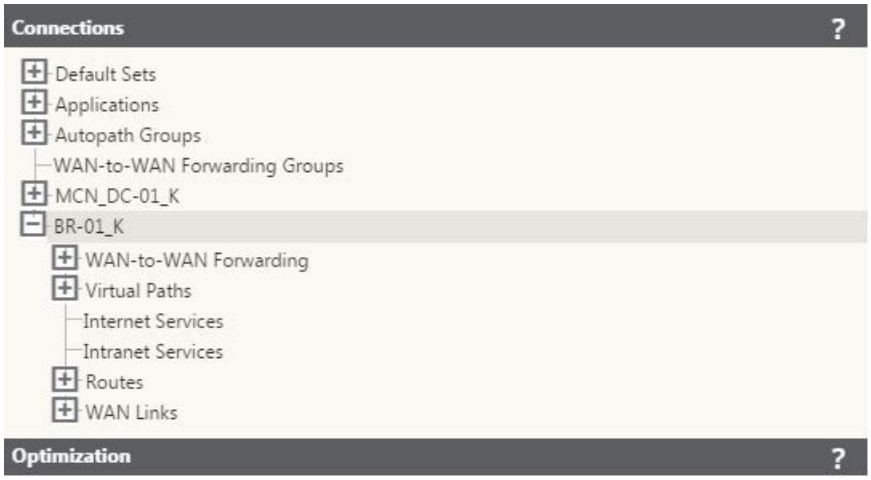

15. Drill down to the **Paths** settings configuration form for any client site Virtual Path you want to configure.

To navigate to the Paths settings form for the client site, do the following:

a. Click **+** to the left of the branch label for the client site.

b. Click **+** to the left of the **Virtual Paths**.

c. Click **+** to the left of the branch label for the Virtual Path you want to configure.

d. Click **+** to the left **Paths**.

The below figures shows an example **Paths** settings form for the new **From Site** path added in the previous steps.

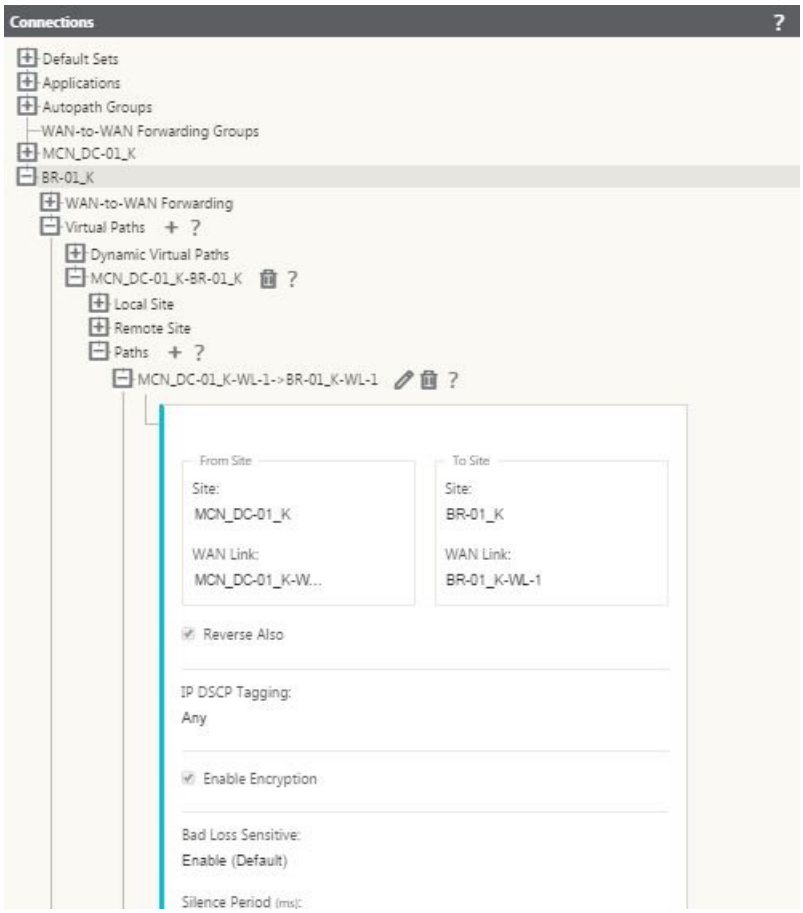

16. Configure the settings for each path you want to customize.

    Follow the same steps as you did to configure the Virtual Paths for the MCN site.

17. (Optional) Click the minus sign (–) to the left of the client site branch label.

    This closes the configured client site branch in the tree.

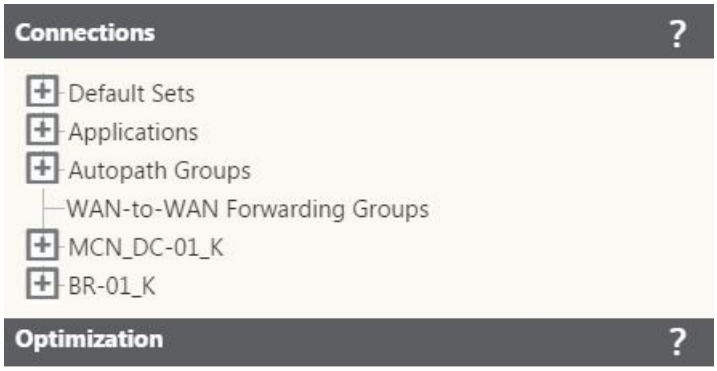

This completes the basic configuration of the Virtual Paths between the client sites and the MCN.

> ## Note
>
> For information on configuring additional settings in the **Connections** or **Provisioning** sections of the **Configuration Editor**, please refer to the Management Web Interface online help for those sections. If you do not want to configure these settings at this time, you can proceed to the appropriate step indicated below.

The next step depends on the SD-WAN Edition license you have activated for your deployment, as follows:

- **SD-WAN Enterprise Edition –** The Enterprise Edition includes the full set of WAN Optimization features. If you want to configure WAN Optimization for your sites, please proceed to the Enabling and Configuring WAN Optimization topic. Otherwise, you can proceed directly to Preparing the SD-WAN Appliance Packages on the MCN.
- **SD-WAN Edition –** This Edition does not include the WAN Optimization features. You can now proceed directly to Preparing the SD-WAN Appliance Packages on the MCN.

Enabling and Configuring WAN Optimization

Aug 09, 2017

The section provides step-by-step instructions for enabling and configuring SD-WAN Enterprise Edition WAN Optimization features for your Virtual WAN. To do this, you will use the **Optimization** section forms in the **Configuration Editor** in the Web Management Interface on the MCN.

> ## Note
>
> You must have a SD-WAN Enterprise Edition license installed to access, enable, configure, and activate WAN Optimization features in your Virtual WAN. SD-WAN Standard Edition does not support these features.

There are two top-level steps for configuring the **Optimization** section sets and parameters. These are as follows, listed in order of dependency:

1. Enable WAN Optimization and customize the **Defaults** configuration, or accept the defaults.

   The **Defaults** configuration is used as the base **Optimization** configuration for all sites eligible for WAN Optimization. The **Defaults** configuration comes pre-configured, and can be customized.

> ## Note
>
> For instructions, see Enabling Optimization and Configuring Default Settings.

2. (Optional) Customize the WAN Optimization configuration for each of the individual branch sites, or accept the **Defaults sets and settings for each.**

By default, the **Defaults** configuration is initially applied to each branch site that is eligible for WAN Optimization. WAN Optimization is supported for 1000-EE and 2000-EE hardware appliances, only. For each supported branch site, you can elect to accept or modify any combination of the **Defaults** sets and settings, or any subset of these. For instructions, see Configuring Optimization for a Branch Site.

To complete these steps, you will use the configuration forms the **Optimization** section of the **Configuration Editor. The Optimization** section is organized as follows:

- **Defaults** – The **Defaults** branch contains the following child branches, which in turn contain one or more forms for configuring their respective sets and settings:

  * **Defaults Features**

  * **Defaults Tuning Settings**

    * **Defaults Application Classifiers (set)**

  * **Defaults Service Classes** (set)

- <**Client Site Name**> – The **Optimization** section configuration tree contains a branch for each client node (branch site) that supports WAN Optimization. If a client node is an unsupported appliance model, the site will not be included in the **Optimization** section configuration tree. Each branch in the tree contains the following child branches, which in turn contain one or more forms for configuring their respective sets and settings:

    * **Defaults Features**

    * **Defaults Tuning Settings**

    * **Defaults Application Classifiers** (set)

    * **Defaults Service Classes** (set)

The below figure shows a simple example of the top and second levels of the **Optimization** section configuration tree. In this example, the branch site **BR-01_K** is included, because the site is configured for a 1000-EE appliance.

The following section provides instructions for enabling WAN Optimization for your Virtual WAN, and configuring the **Defaults** sets and settings.

Enabling Optimization and Configuring Default Settings

Aug 09, 2017

The first step is to enable WAN Optimization in the new configuration package, and configure the **Defaults** sets and settings.

The **Optimization** section **Defaults** sets and settings are categorized as follows:

- **Defaults Features**
- **Defaults Tuning Settings**
- **Defaults Application Classifiers** (set)
- **Defaults Service Classes** (set)

The following sections provide instructions for enabling WAN Optimization and configuring each of these **Defaults** sets and settings.

Enabling Optimization and Configuring the Default Feature Settings

Aug 09, 2017

Enabling WAN Optimization in your Virtual WAN entails the following procedures:

1. Enable WAN Optimization in the **Features** settings of the **Optimization section.**

   Instructions for this part of the process are provided in this section.

2. Configure the **Acceleration** policy setting for each applicable Service Class in the **Service Classes** table.

This procedure occurs further on, after you have completed the rest of the **Optimization** configuration. Instructions are provided in the section, Configuring Optimization Default Service Classes. At this point, WAN Optimization has been enabled in your configuration, but not yet enabled and activated in your Virtual WAN. To enable and activate WAN Optimization in your Virtual WAN, you must complete the Virtual WAN configuration, and then generate, stage, and activate the Virtual WAN Appliance Packages on the eligible sites in your deployment, as outlined in the subsequent chapters of this guide.

To enable WAN Optimization and configure the **Defaults** section **Features** settings, do the following:

1. If necessary, log back into the Management Web Interface, and open the **Configuration Editor.**

   To open the **Configuration Editor**, do the following:

   a. Select the **Configuration** tab at the top of the page to open the **Configuration** navigation tree (left pane).

   b. In the navigation tree, click **+** to the left of the **Virtual WAN** branch to open that branch.

   C. In the **Virtual WAN** branch, select **Configuration Editor.**

2. Open the configuration package you want to modify.

   Click **Open** to display the **Open Configuration Package** dialog box, and select the package from the **Saved Packages** drop-down menu.

   This loads the selected package into the **Configuration Editor** and opens it for editing.

If you have a valid and current license that includes WAN Optimization features, the **Optimization** section will be available in the **Configuration Editor**.

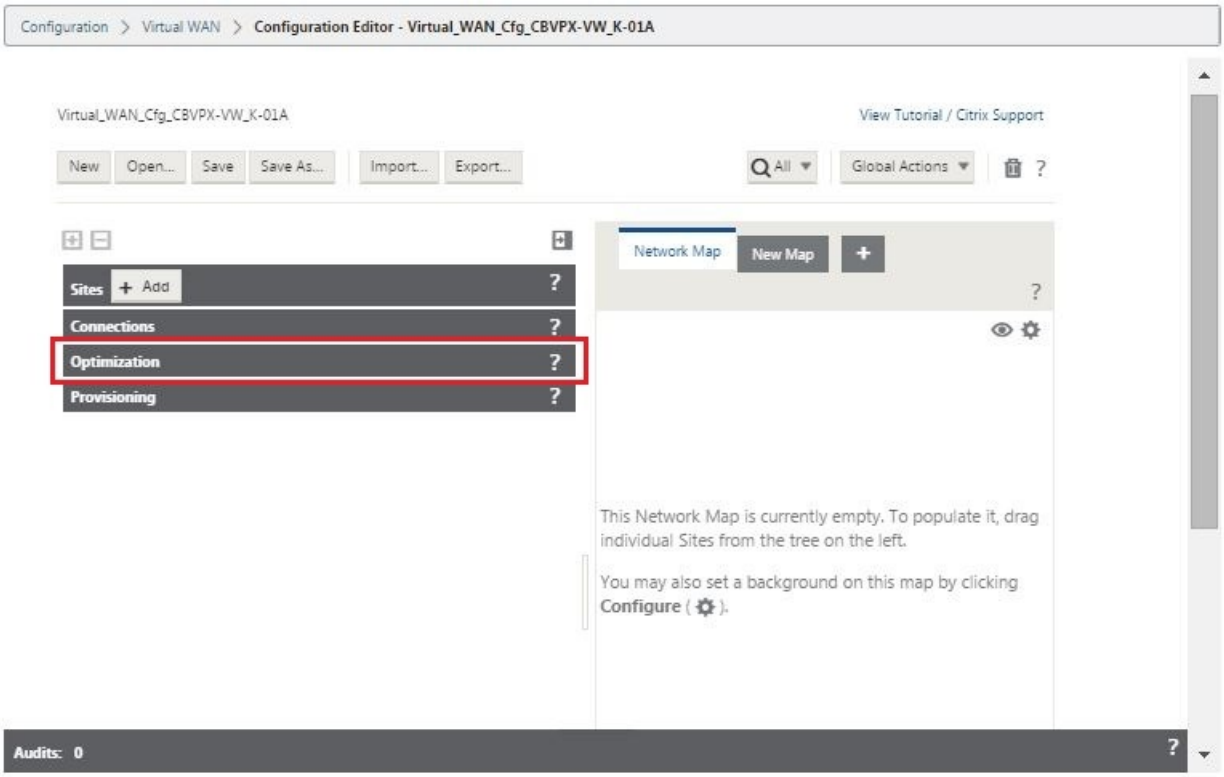

Note

If the **Optimization** section is not available, please check that you have installed a CloudBridge Enterprise Edition license in your Virtual WAN. CloudBridge Virtual WAN Edition does not support WAN Optimization features.

For details and instructions, see the following sections:

- The SD-WAN Editions
- Licensing
- Uploading and Installing the Virtual WAN Software License File

3. Click the **Optimization** section heading.

   This opens the **Configuration Editor Optimization** section tree.

The **Optimization** section tree contains a branch for the **Defaults** settings, and a branch for each eligible client node (branch site) in the current configuration. Optimization is supported for 1000-EE and 2000-EE clients, only. Consequently,

the client node must be a 1000-EE or 2000-EE hardware Virtual WAN Appliance for that site to be included in the tree.

4. Click **+** to the left of the **Defaults** branch label.

   This opens the **Defaults** settings tree.

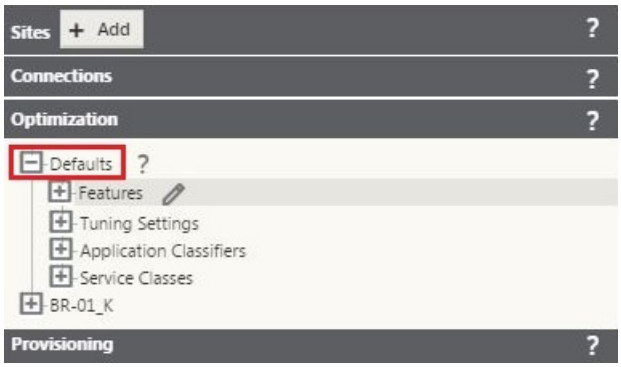

5. Click **+** to the left of the **Features** branch label.

   This opens the default **Features** configuration form.

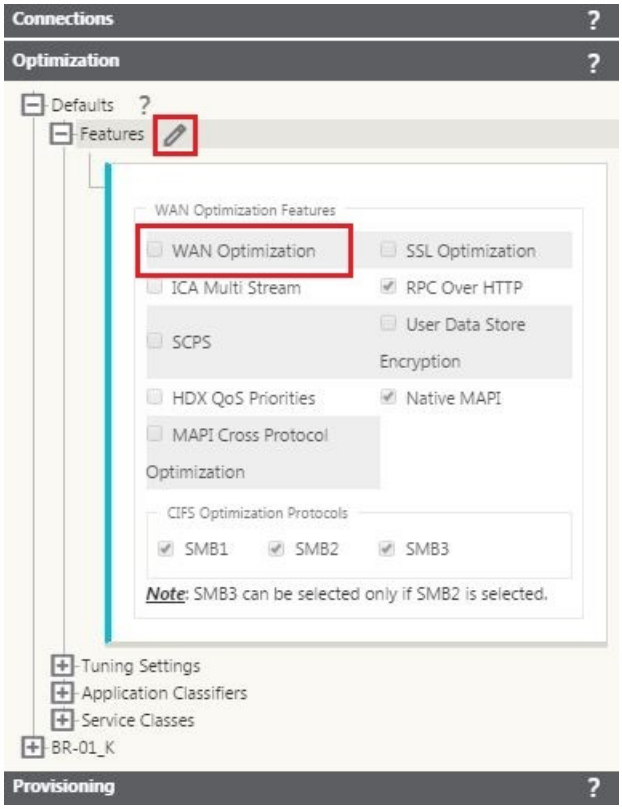

6. Click Edit (pencil icon) to the right of the **Features** branch label to enable editing of the form.

7. Select the **WAN Optimization** checkbox.

The **WAN Optimization** checkbox is in the upper left corner of the **WAN Optimization** Features section at the top of the form. Select the checkbox to select **WAN Optimization** for enabling. This also opens the other options in the form for editing, and reveals the **Apply** and **Revert** buttons.

Note

This selects this feature for enabling, only. WAN Optimization will not be enabled in the **Optimization** section or the configuration package until you click **Apply,** after completing the **Features** configuration. In addition, you must also configure the **Acceleration** setting for each applicable Service Class in the Service Classes table, as instructed further on in the **Optimization** configuration process. (Instructions are provided in the section Configuring Optimization Default Service Classes) Finally, WAN Optimization will not be enabled and activated in your Virtual WAN until you have completed the entire Virtual WAN configuration, and then generated, staged, distributed, and activated the Virtual WAN Appliance Packages on the eligible sites in your Virtual WAN.

The below figure shows the **Defaults** section **Features** configuration form, with WAN Optimization enabled, and the **Apply** and **Revert** buttons revealed.

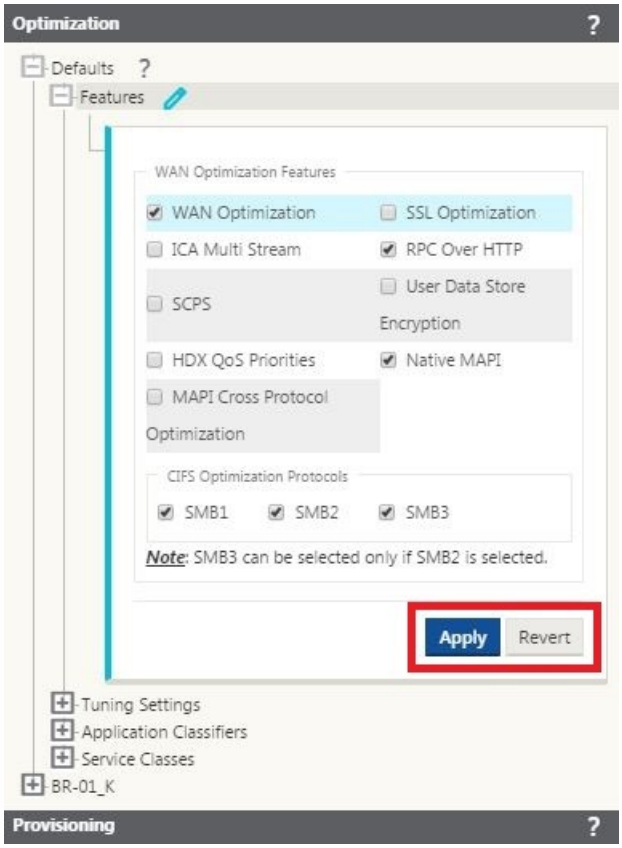

8. Configure the **Features** settings.

Click a checkbox to select or deselect an option. You can accept the default settings pre-selected in the form, or customize the settings.

Note

By default, the settings you configure in the **Defaults** section forms are automatically applied to each branch site included in the tree. However, you can customize the **Optimization** configuration for a specific branch, as outlined in the section, Configuring Optimization for a Branch Site.

The **Features** configuration form contains two sections:

- **WAN Optimization Features**

- **CIFS Optimization Protocols**

The **WAN Optimization Features** settings are as follows:

- **WAN Optimization** – Select this to enable WAN Optimization for this configuration. This also enables compression, deduplication, and TCP Protocol Optimization.

> **Note**
>
> The WAN Optimization option must be selected for the other Optimization section options to be available.

- **ICA Multi Stream** – Select this to enable CloudBridge to negotiate the use of multi-stream for ICA traffic for improved Quality of Service.

- **SCPS** – Select this to enable TCP Protocol optimization for Satellite Links.

- **HDX QoS Priorities** – Select this to enable optimization of ICA traffic based on prioritization of HDX sub-channels.

- **MAPI Cross Protocol O**ptimization – Select this to enable cross-protocol optimization of Microsoft Outlook (MAPI) traffic.

- **SSL Optimization** – Select this to enable optimization for traffic streams with SSL encryption.

- **RPC Over HTTP –** Select this to enable optimization of Microsoft Exchange traffic that uses RPC over HTTP.

- **User Data Store Encryption –** Select this to enable enhanced security of data through the encryption of WAN Optimization compression history.

- **Native MAPI** – Select this to enable optimization of Microsoft Exchange traffic.

The **CIFS Optimization Protocols** options are as follows:

- **SMB1** – Select this to enable Optimization of Windows File Sharing (SMB1)

- **SMB2** – Select this to enable Optimization of Windows File Sharing (SMB2)

- **SMB3** – Select this to enable Optimization of Windows File Sharing (SMB3). You must first select the **SMB2** option before you can select **SMB3**.

9. Click **Apply**.

This enables and adds the selected **Default Features** to the configuration package.

The next step is to configure the **Optimization** default **Tuning Settings**.

Configuring Optimization Default Tuning Settings

Aug 09, 2017

To configure the WAN Optimization default **Tuning Settings**, do the following:

1. Open the **Tuning Settings** configuration form.

   Continuing in the **Defaults** branch of the **Optimization** section of the **Configuration Editor**, click **+** to the left of the **Tuning Settings** branch label.

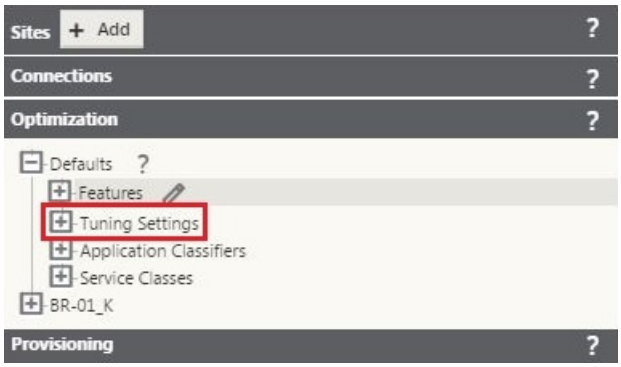

This opens the **Defaults** section **Tuning Settings** configuration form.

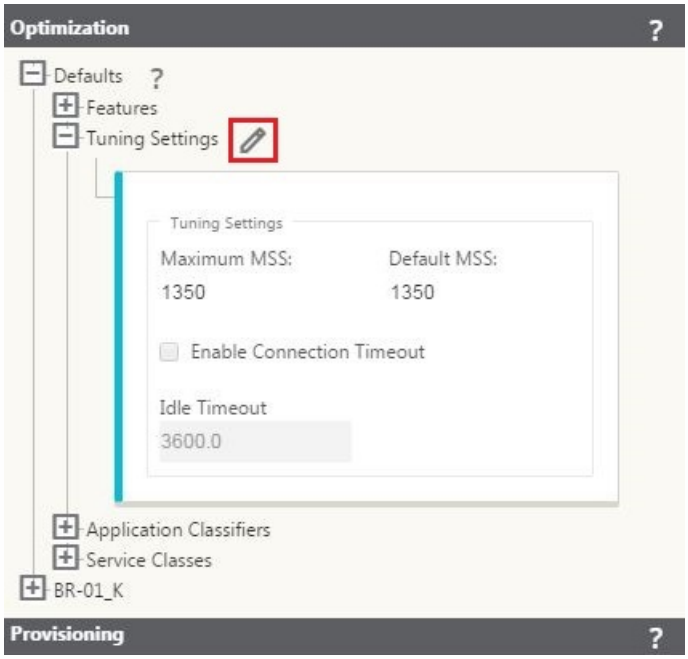

2. Click Edit (pencil icon) to enable editing of the form.
3. Select and configure the **Tuning Settings**.

   The **Tuning Settings** options are as follows:

   - **Maximum MSS** – Enter the maximum size (in bytes) for the Maximum Segment Size (MSS) for a TCP segment.
   - **Default MSS** – Enter the default size (in octets) for the MSS for TCP segments.
   - **Enable Connection Timeout** – Select this to enable automatic termination of a connection when the idle

threshold is exceeded.

- **Idle Timeout** – Enter a threshold value (in seconds) to specify the amount of idle time permitted before an idle connection is terminated. You must first select **Enable Connection Timeout** before this field can be configured.

4. Click **Apply**.

   This applies the modified **Tuning Settings** to the **Defaults** configuration.

   The next step is to configure the default set of WAN Optimization Application Classifiers.

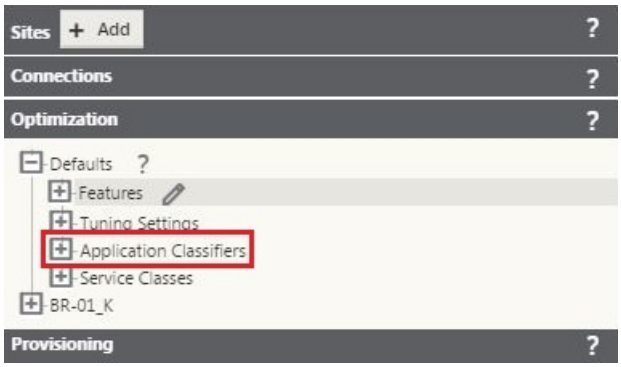

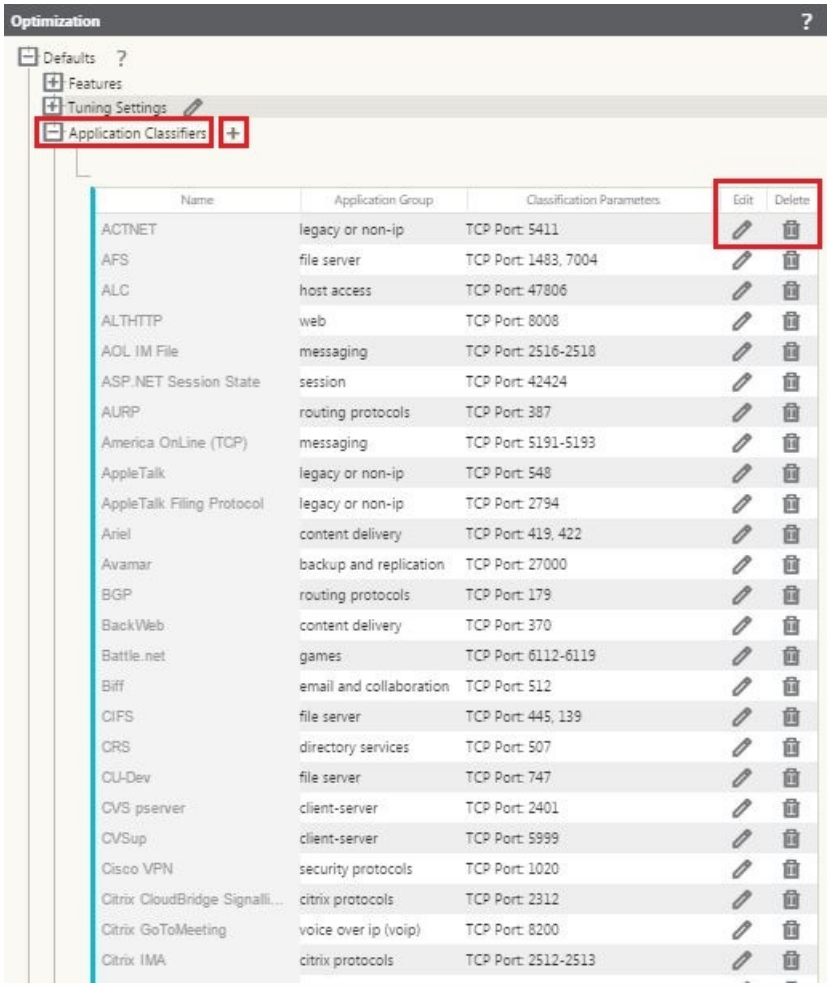

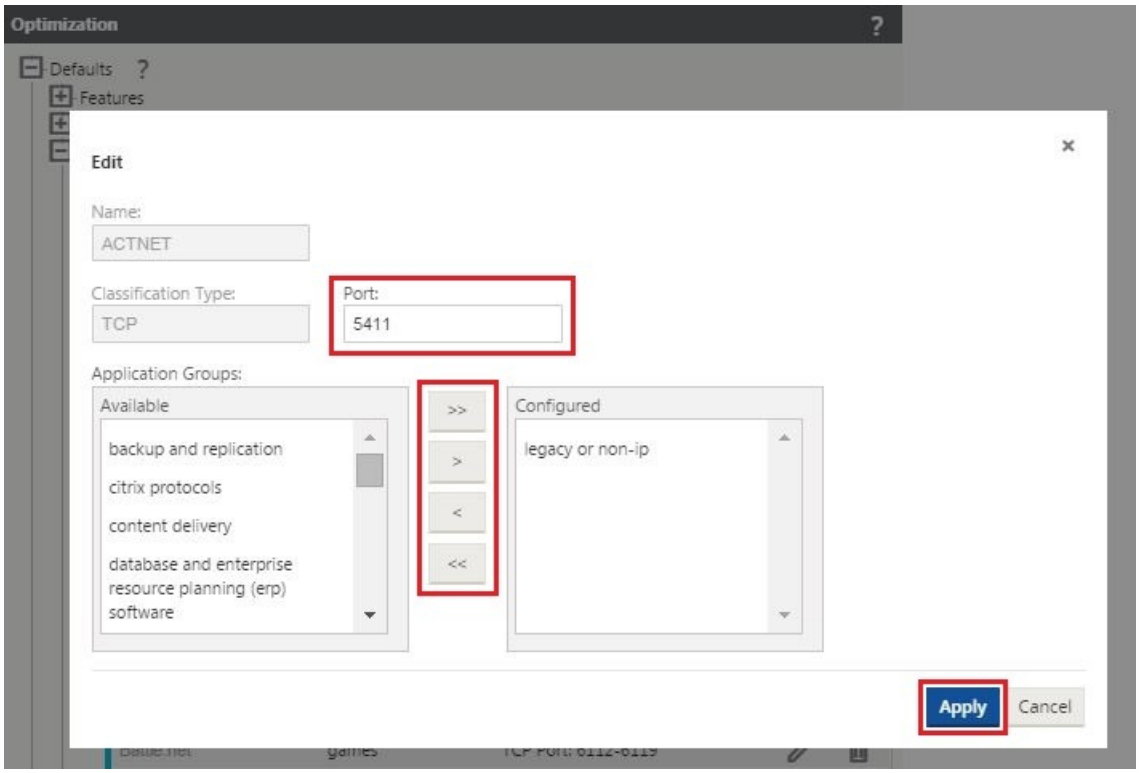

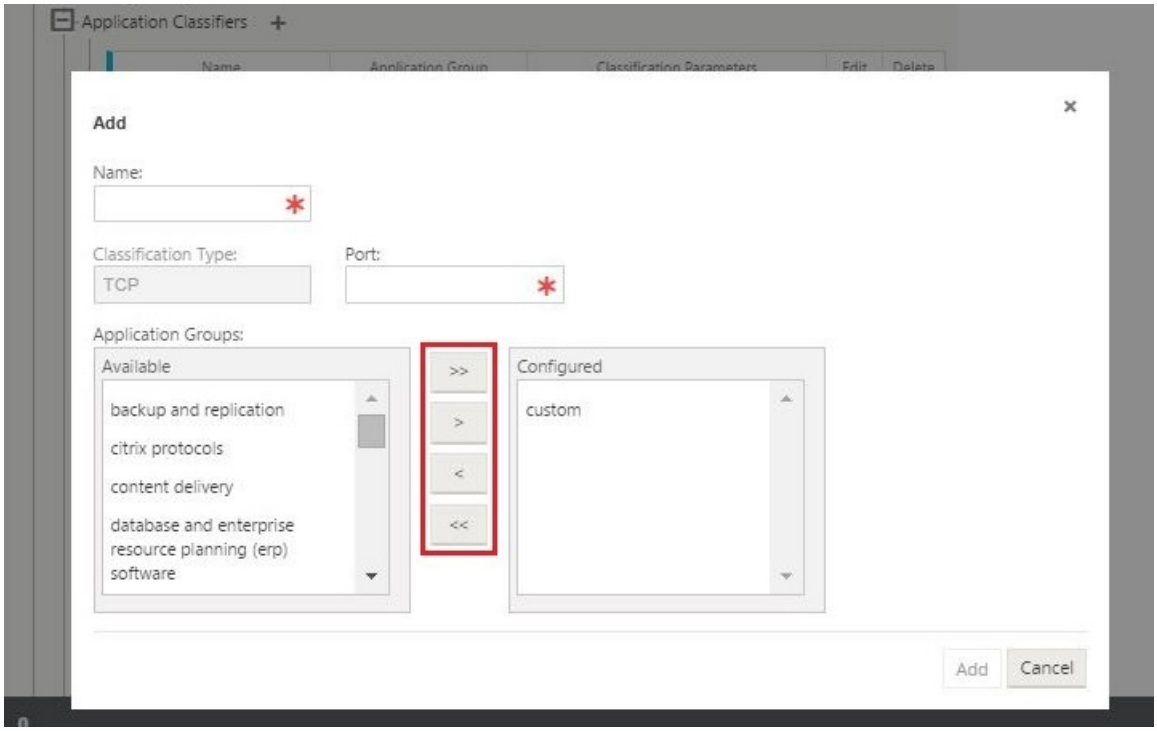

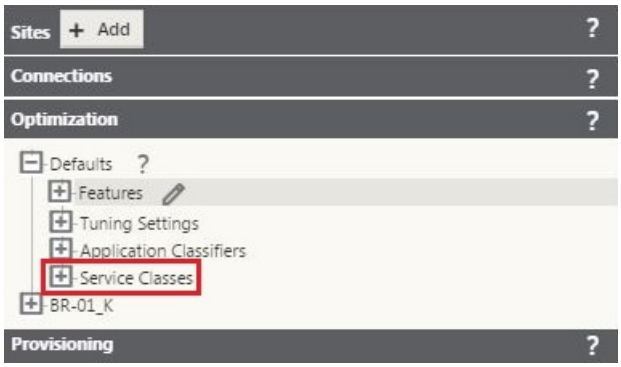

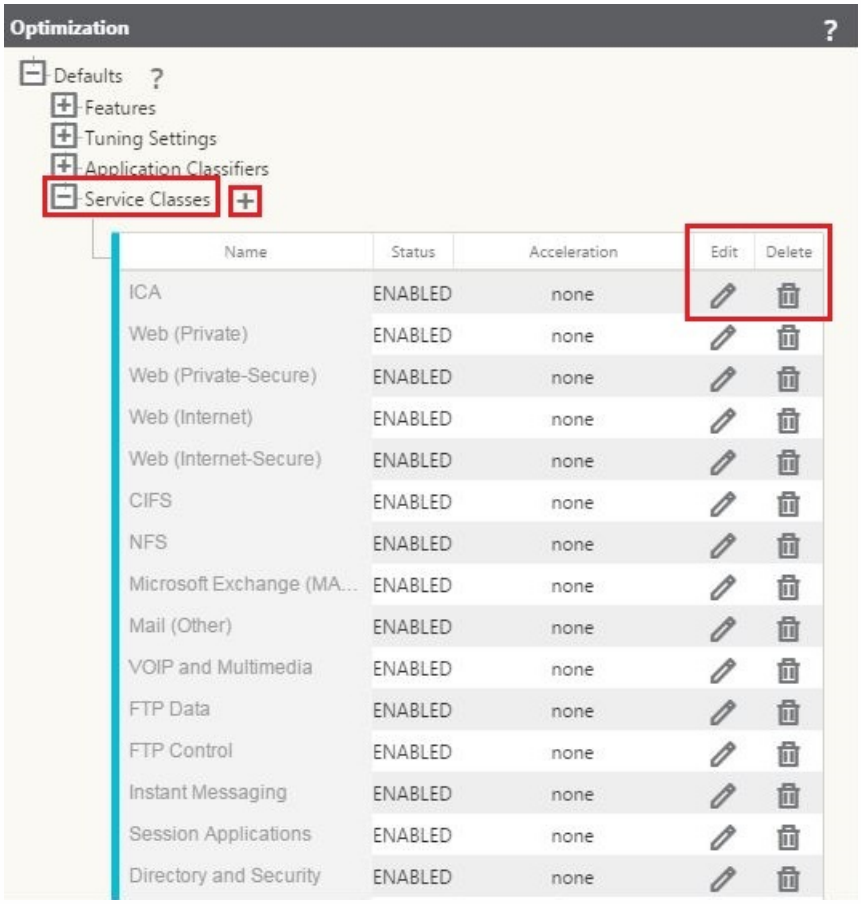

-

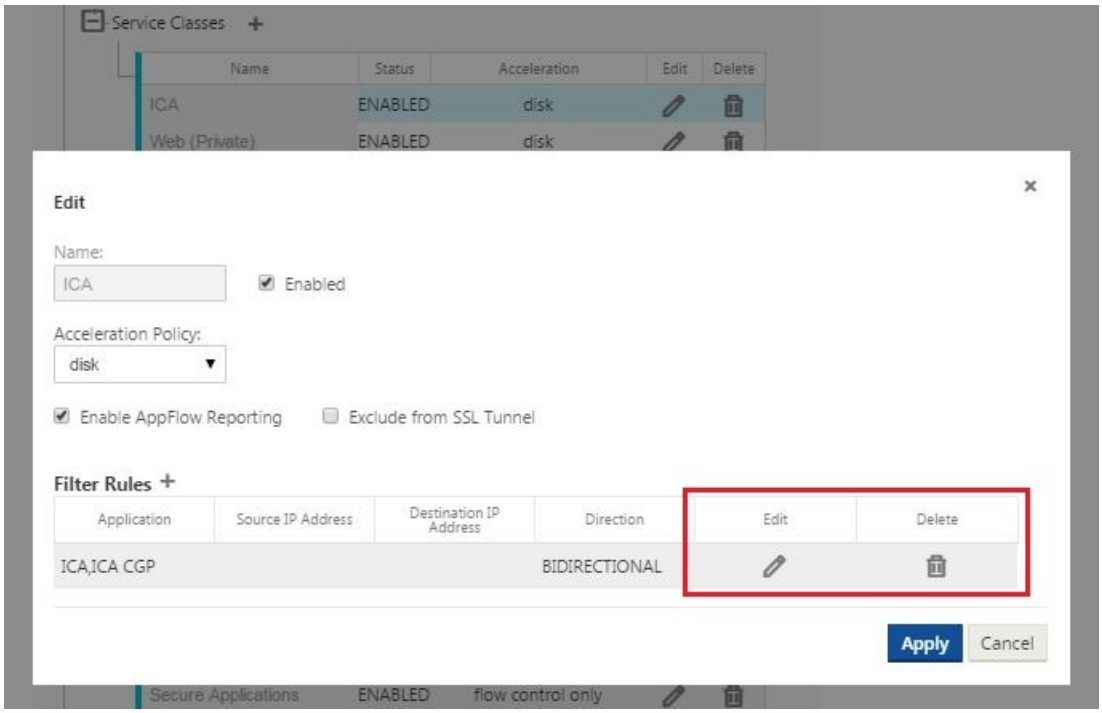

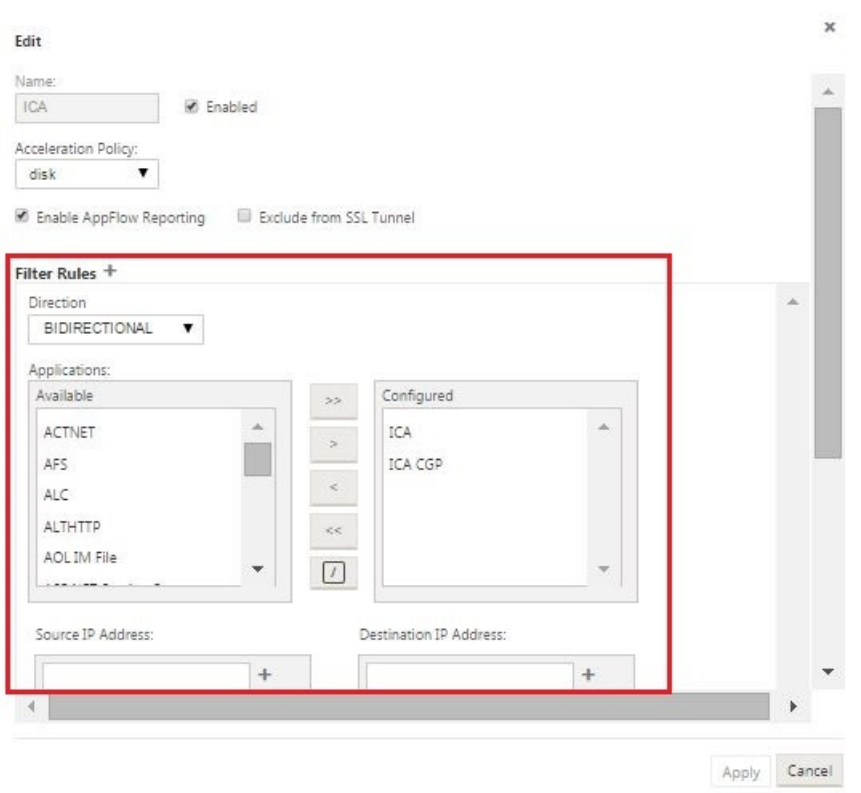

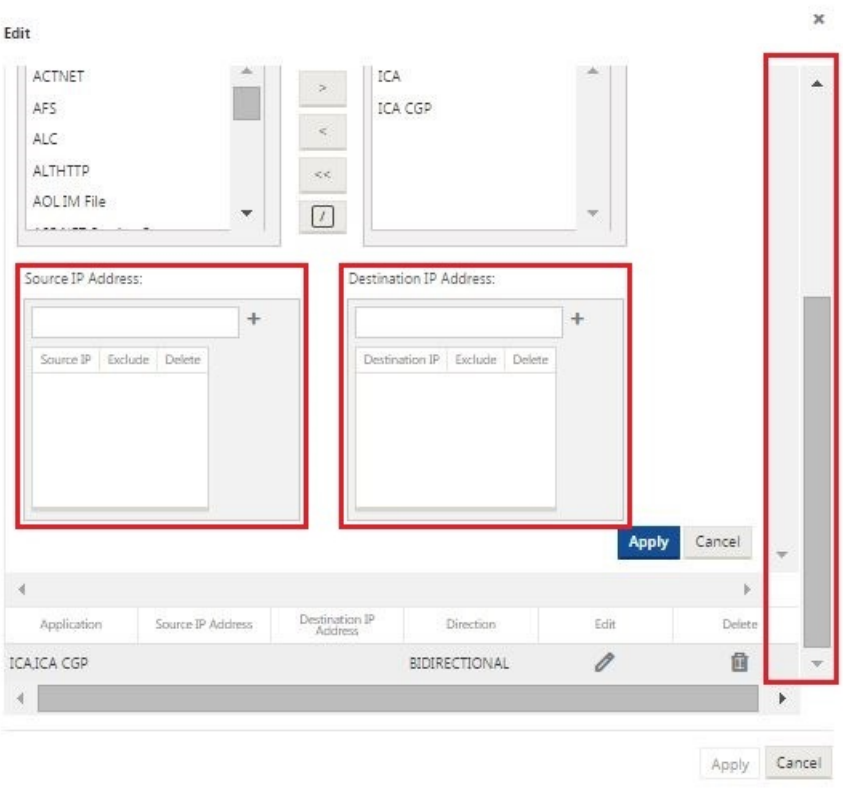

Source IP Address:

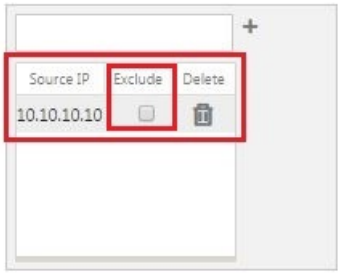

Destination IP Address:

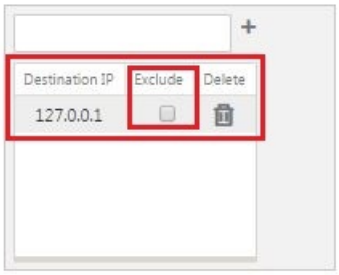

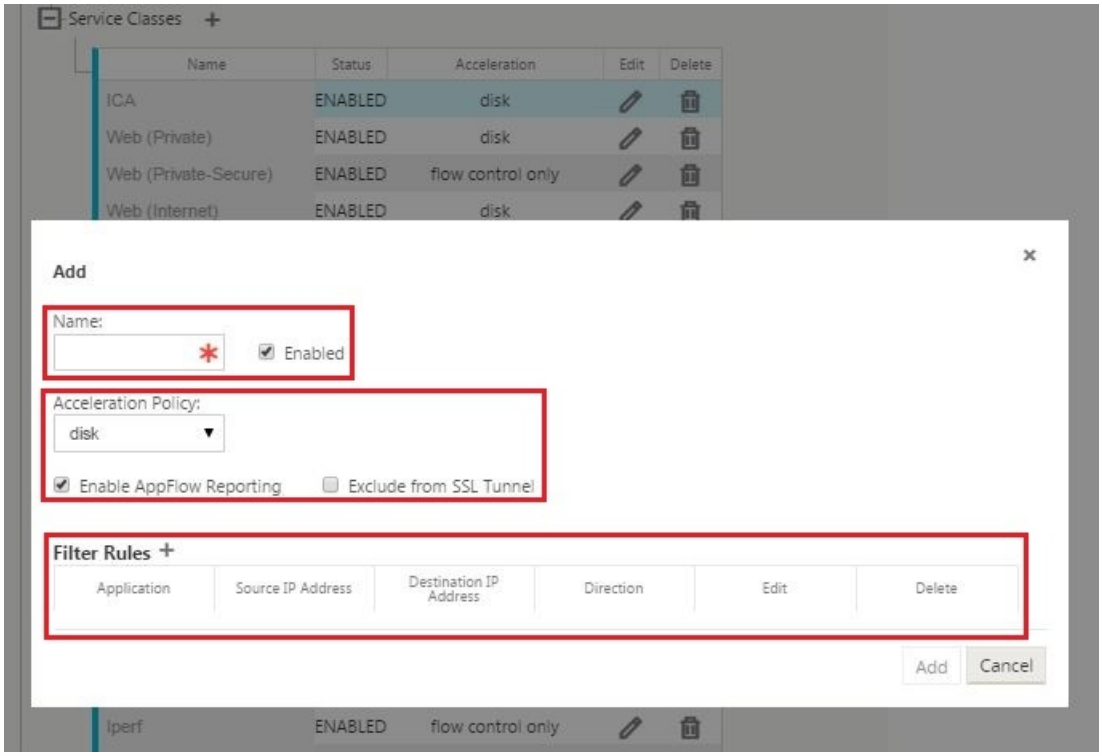

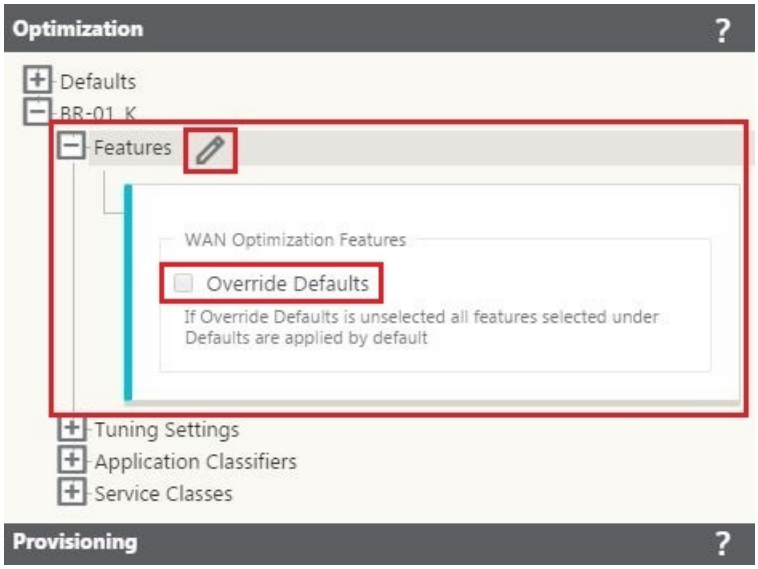

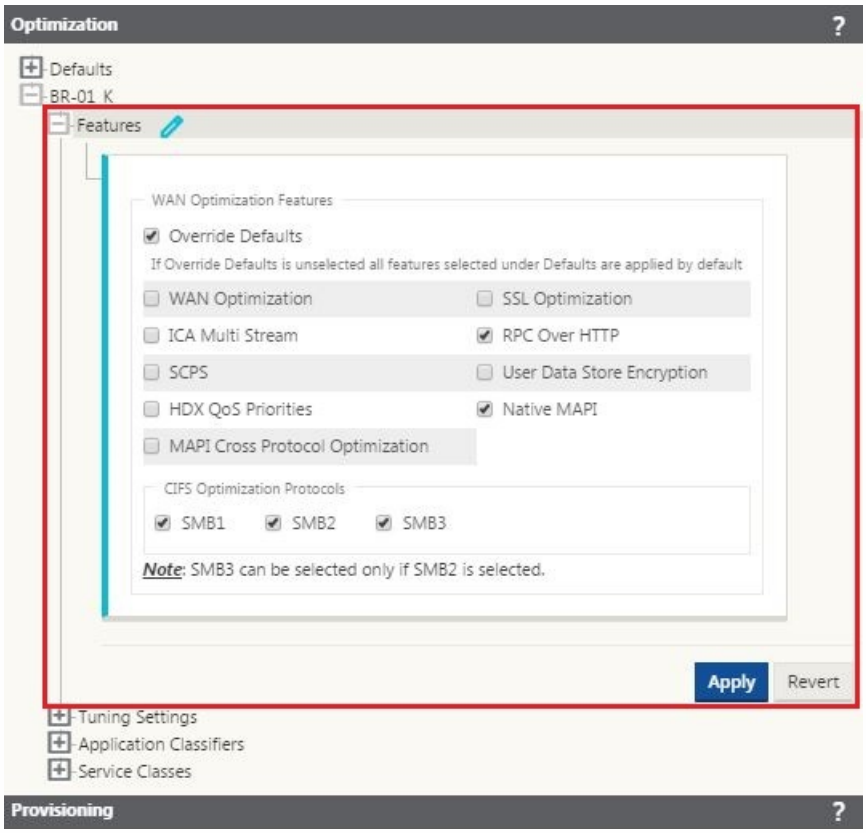

-
-
-
-

-
-
-
-
-

-
-
-
-

Global Virtual WAN Network Encryption

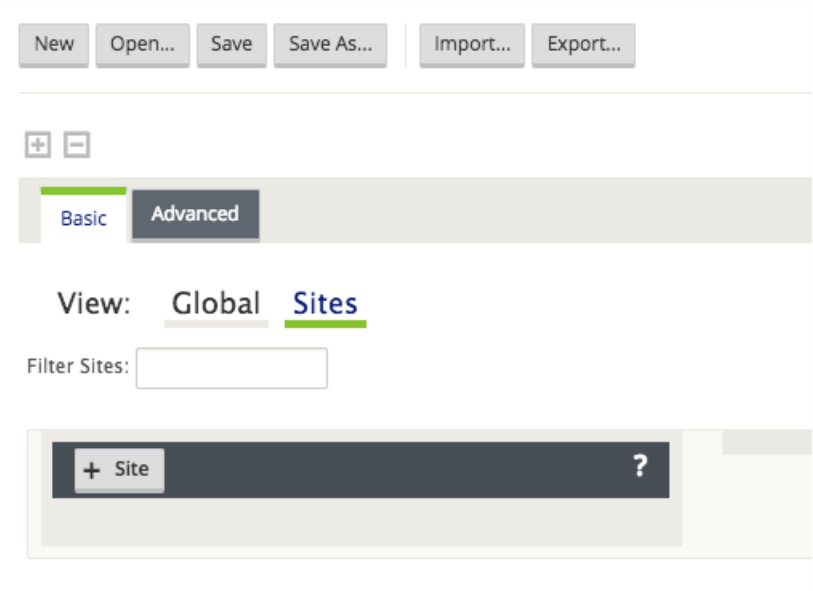

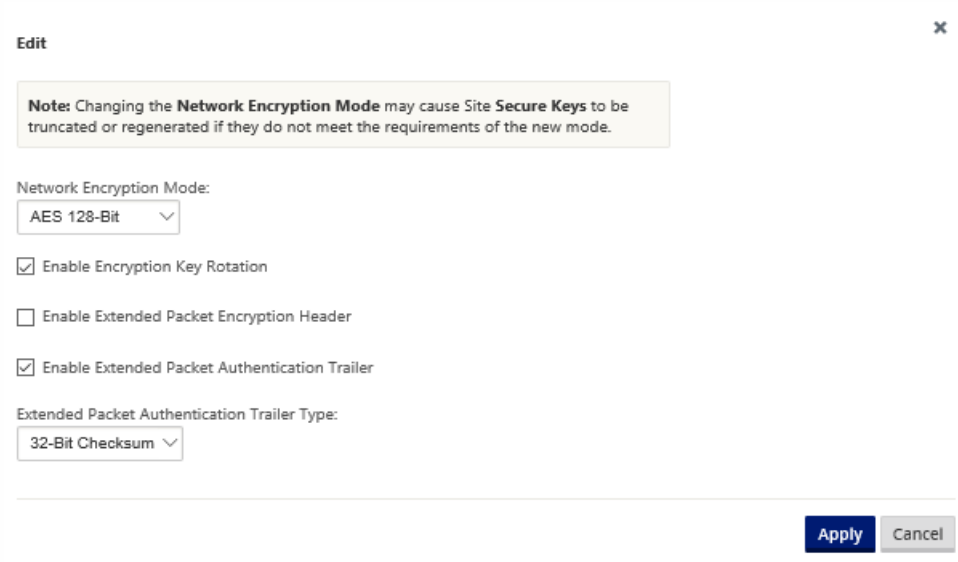

-
-
-
-
-

WAN Link Template

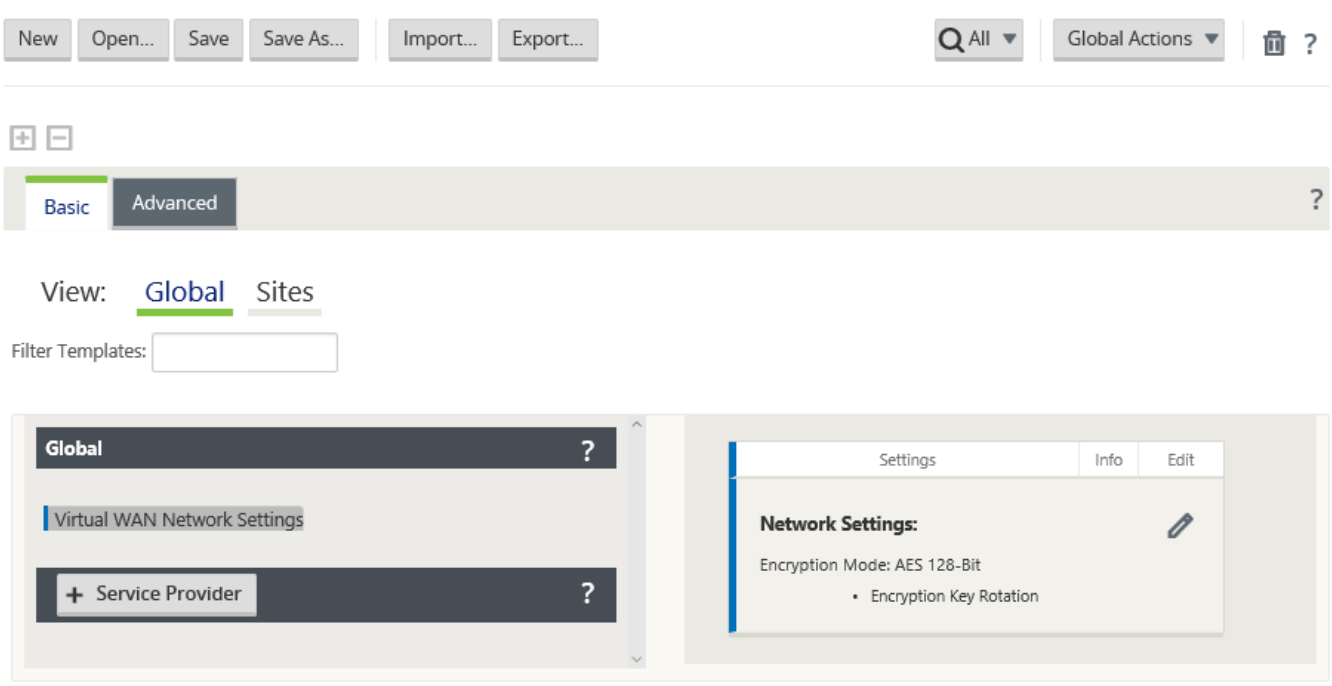

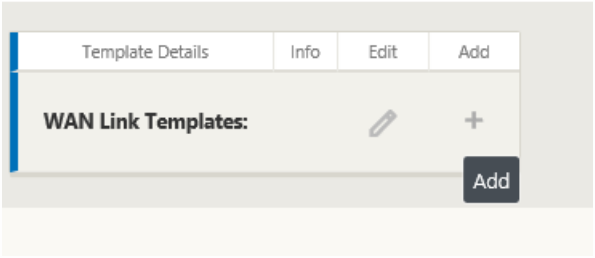

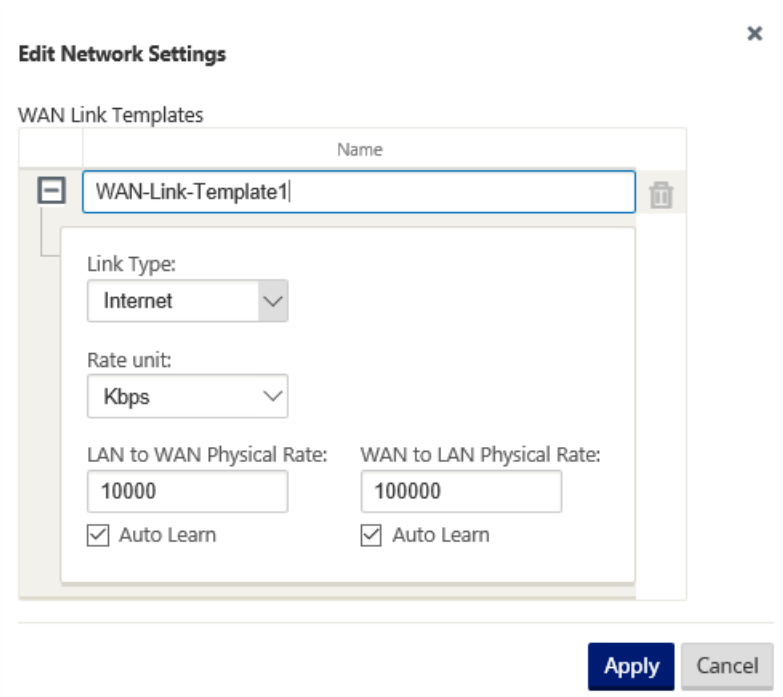

-
-
-

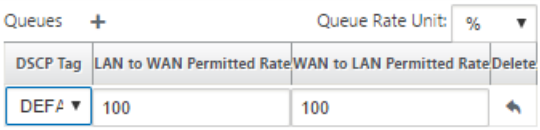

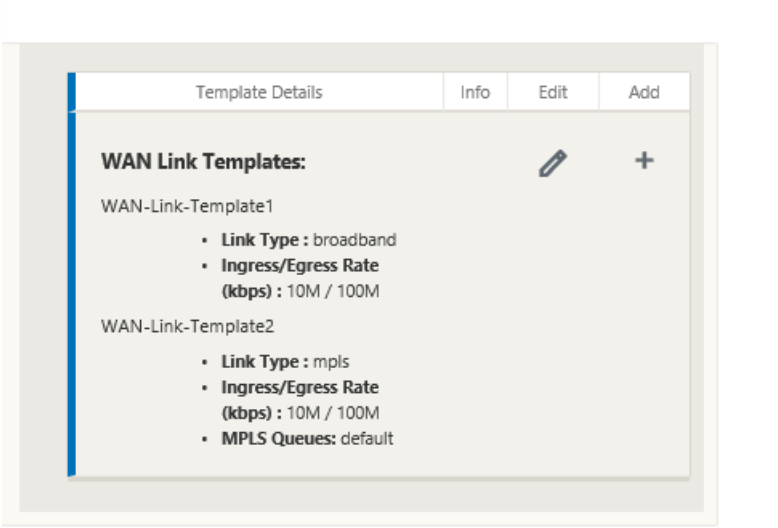

Site Cloning

-
-
-
-

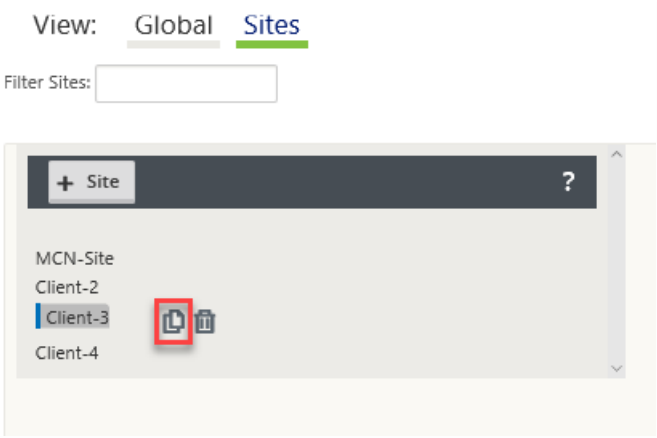

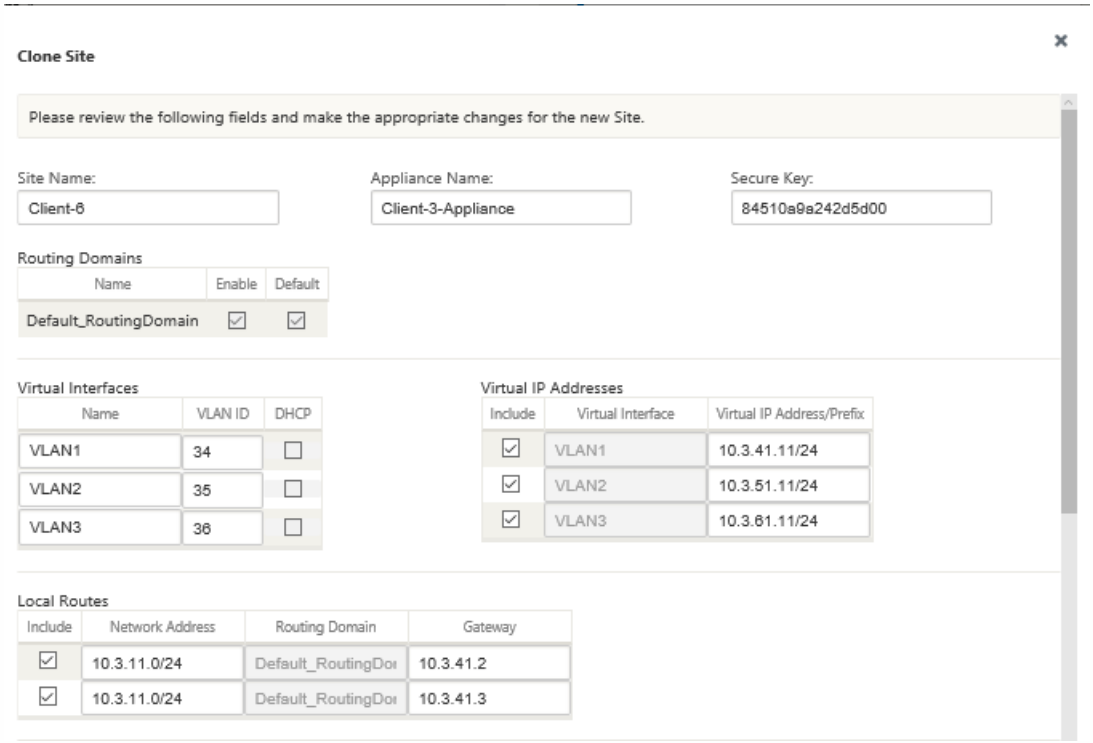

-
-
-
-
  -
  -
  -
  -
  -
  -
  -
  -
-
-
-

NetScaler SD-WAN appliances enable configuring routing protocols providing single point of administration to manage a corporate network, or a branch office network, or a data center network.

To configure routing domain:

1. In the SD-WAN web interface, navigate to **Configuration → Virtual WAN → Configuration Editor**. In the **Configuration Editor**, navigate to **Global → Routing Domains**, click **Add (+)** and enter a Name for your new Routing Domain.

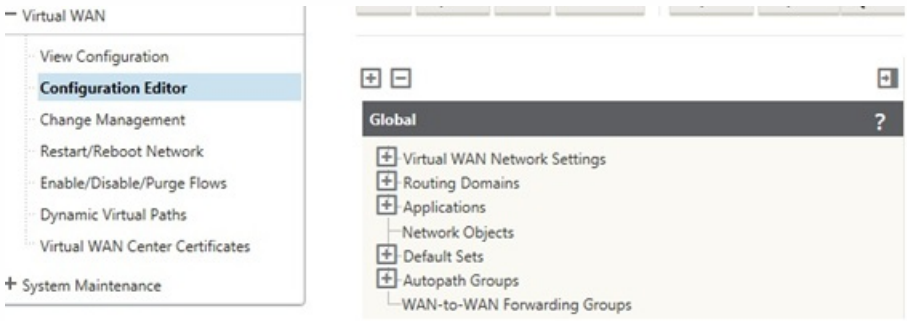

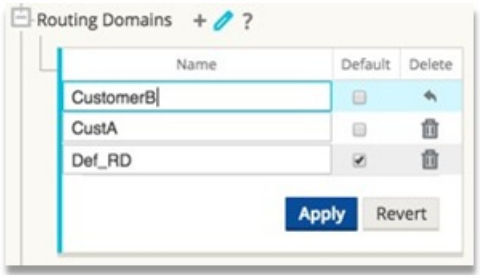

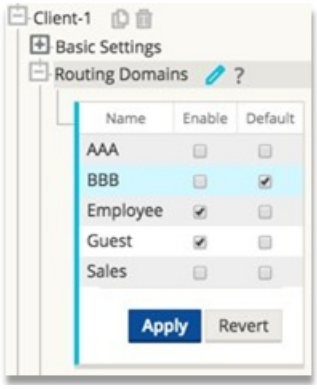

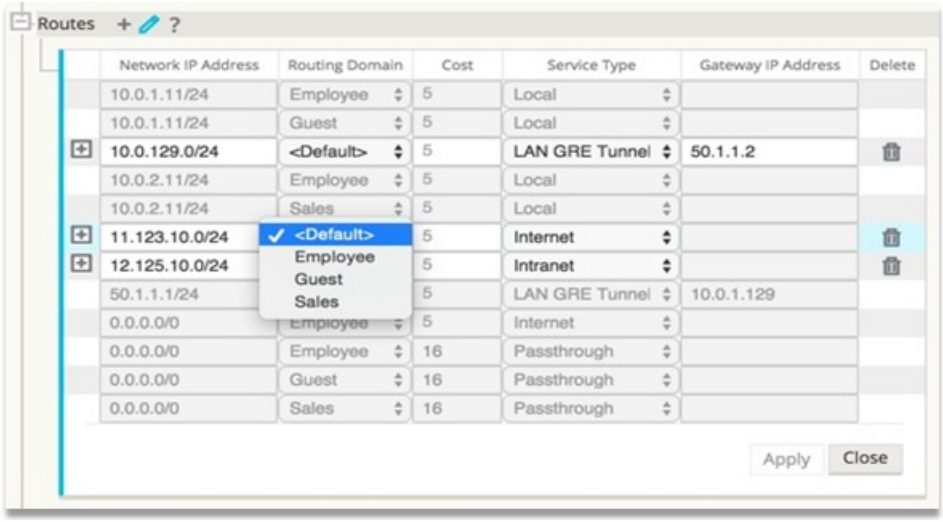

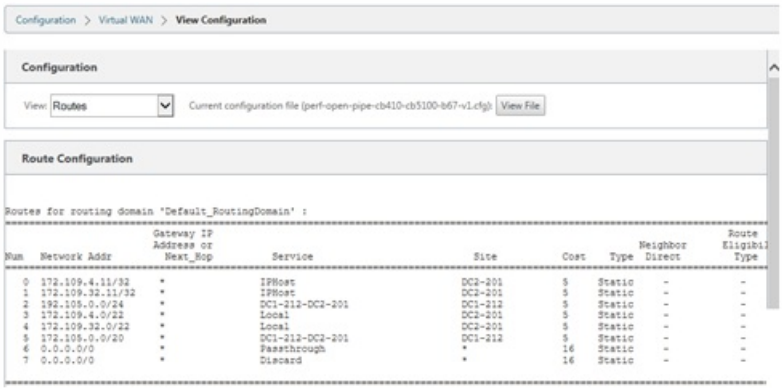

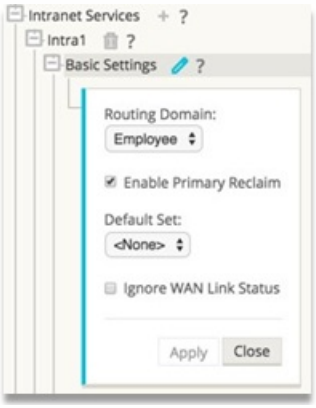

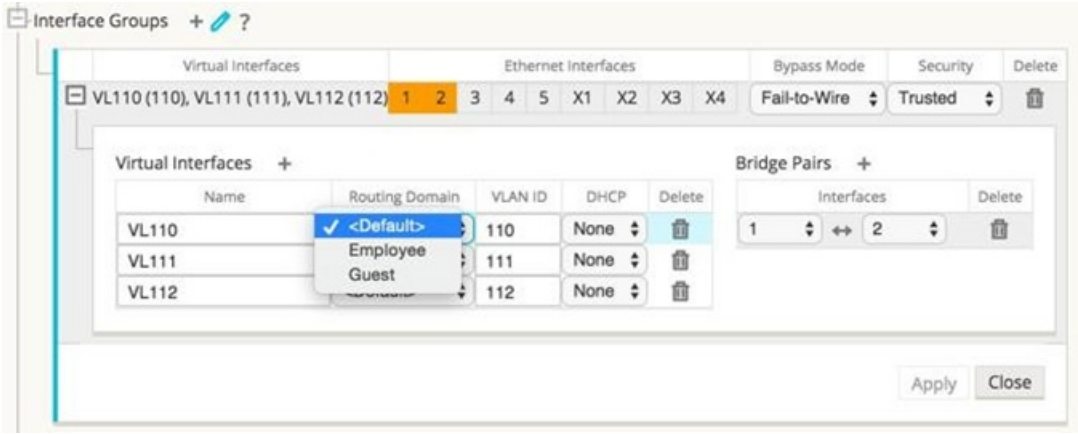

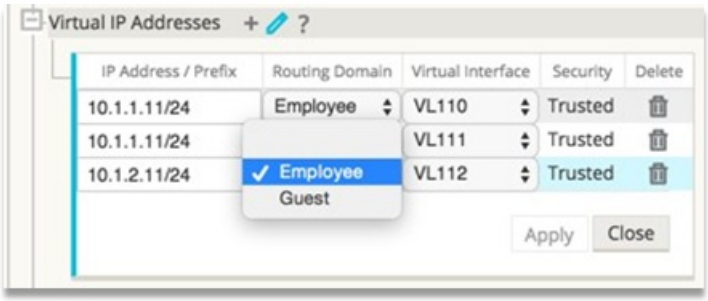

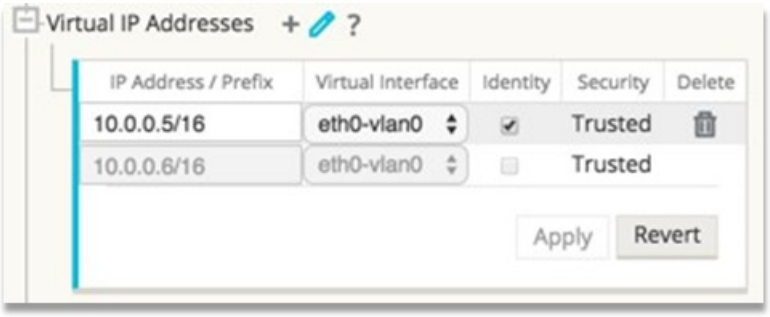

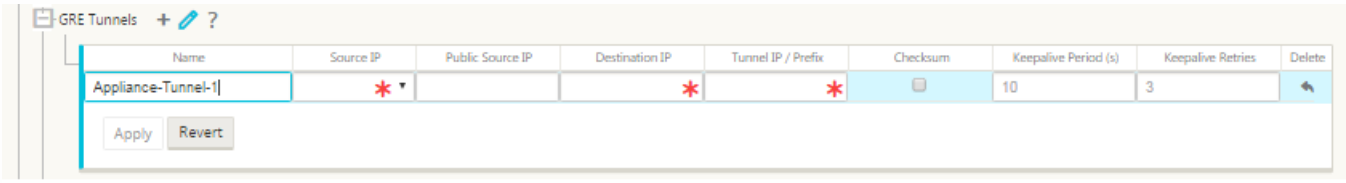

-
-
-

-

-

-

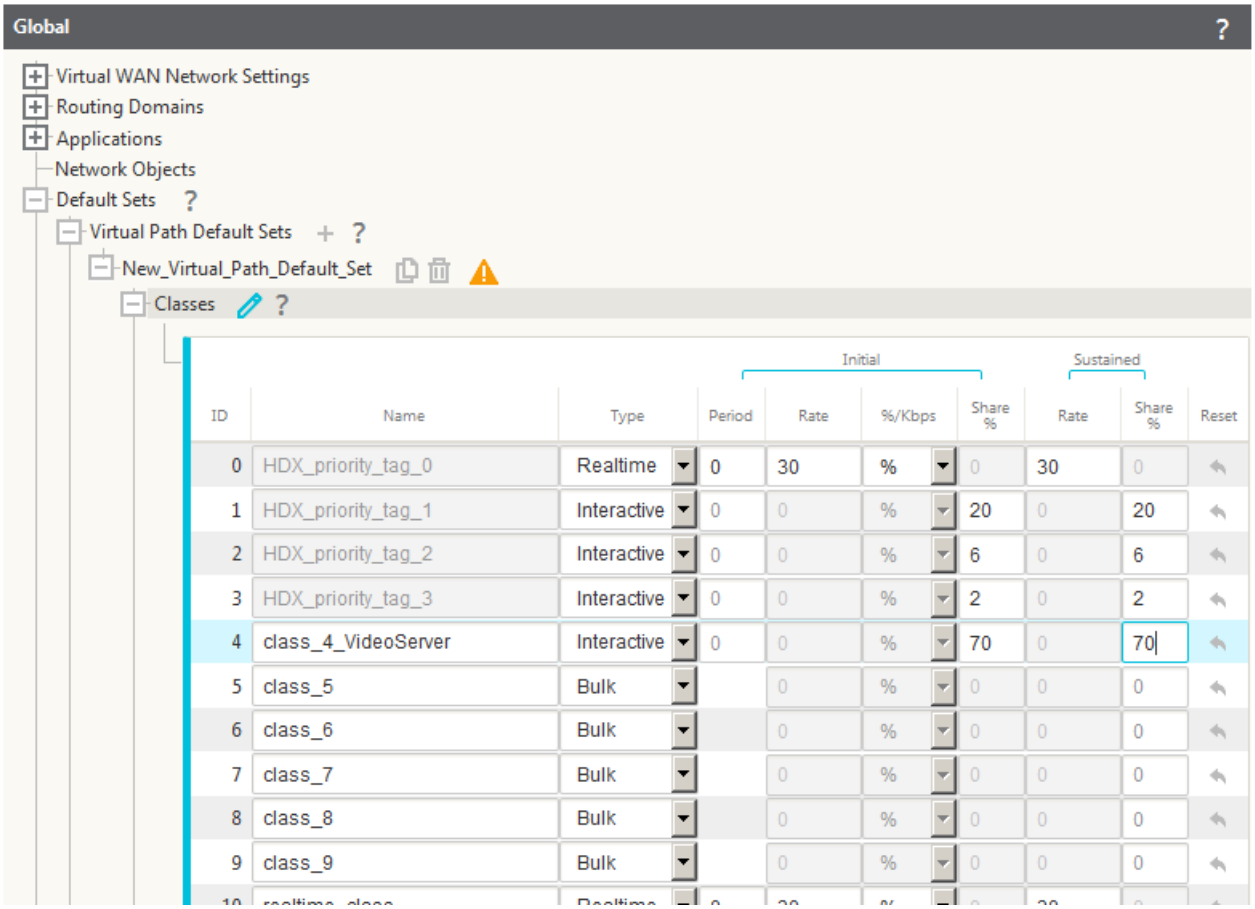

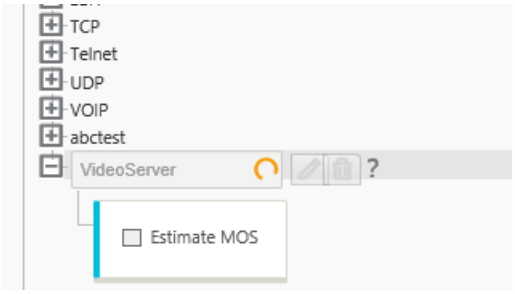

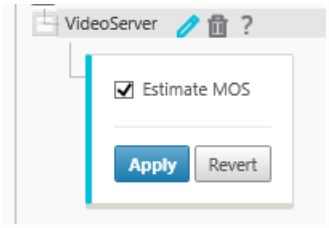

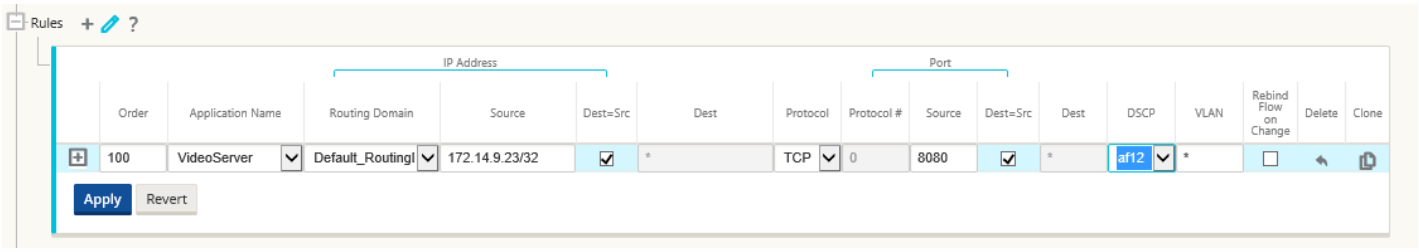

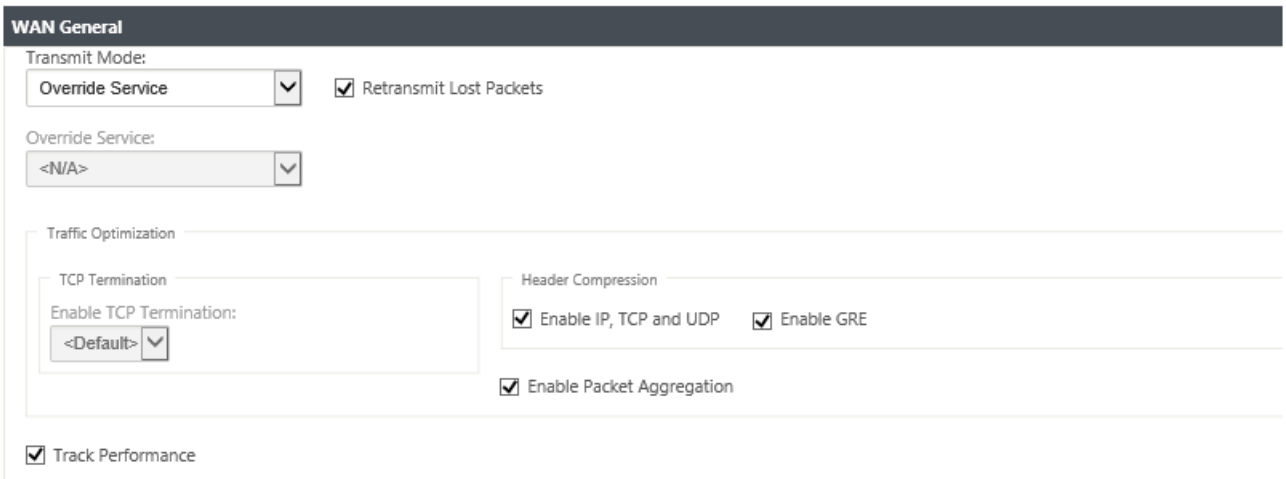

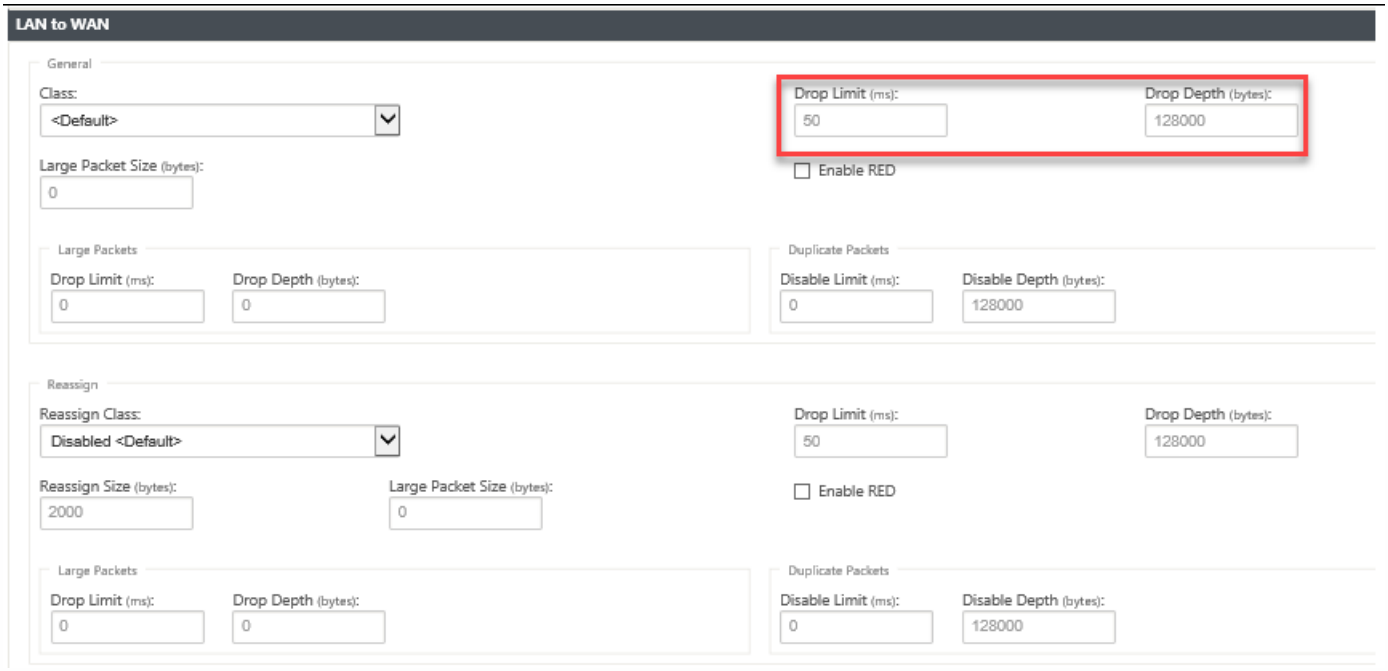

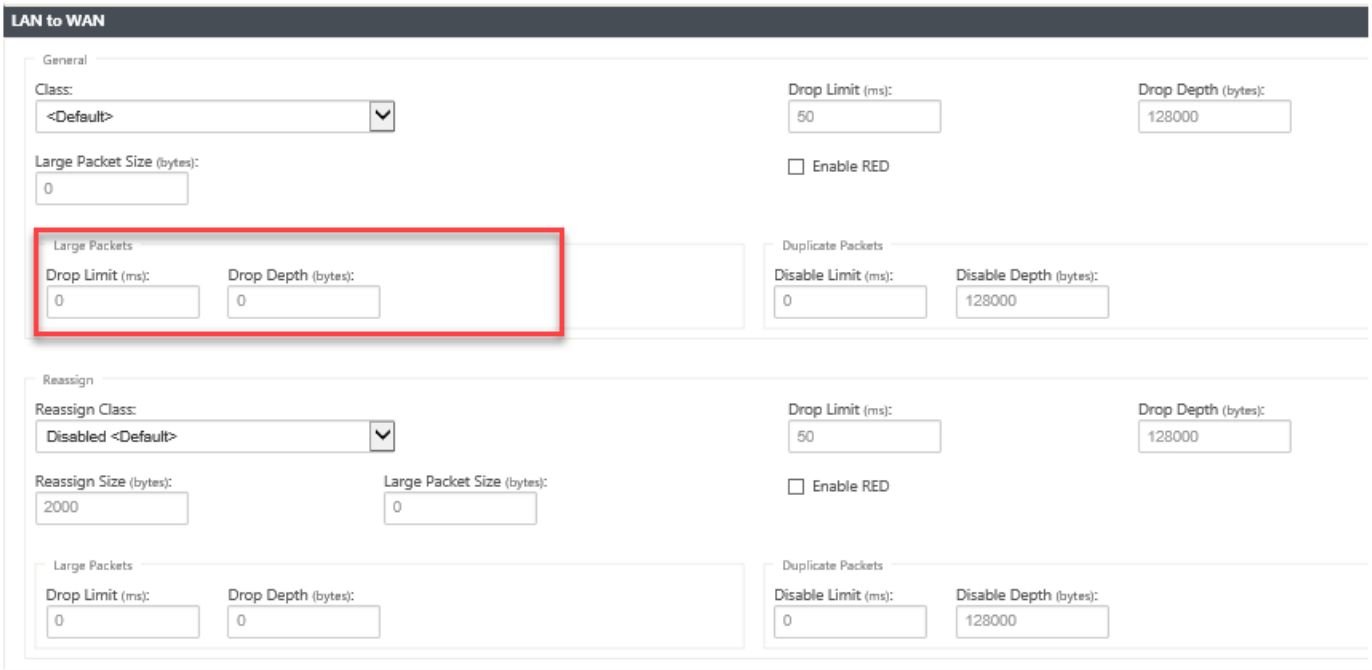

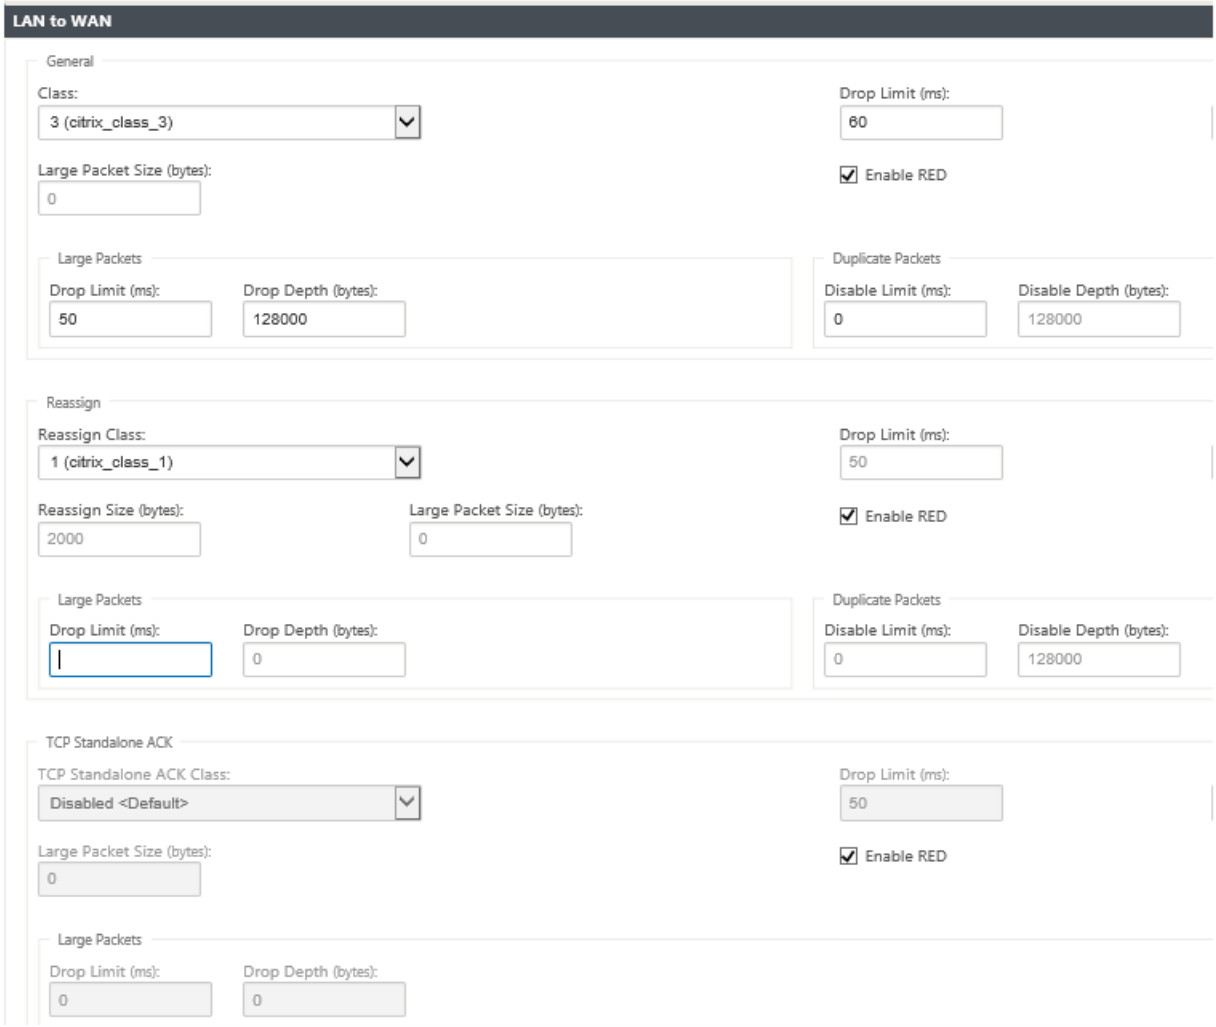

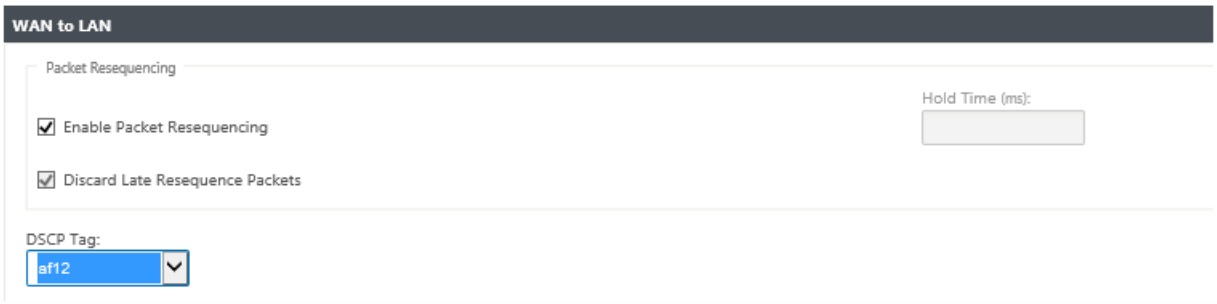

-
-
-
-
-

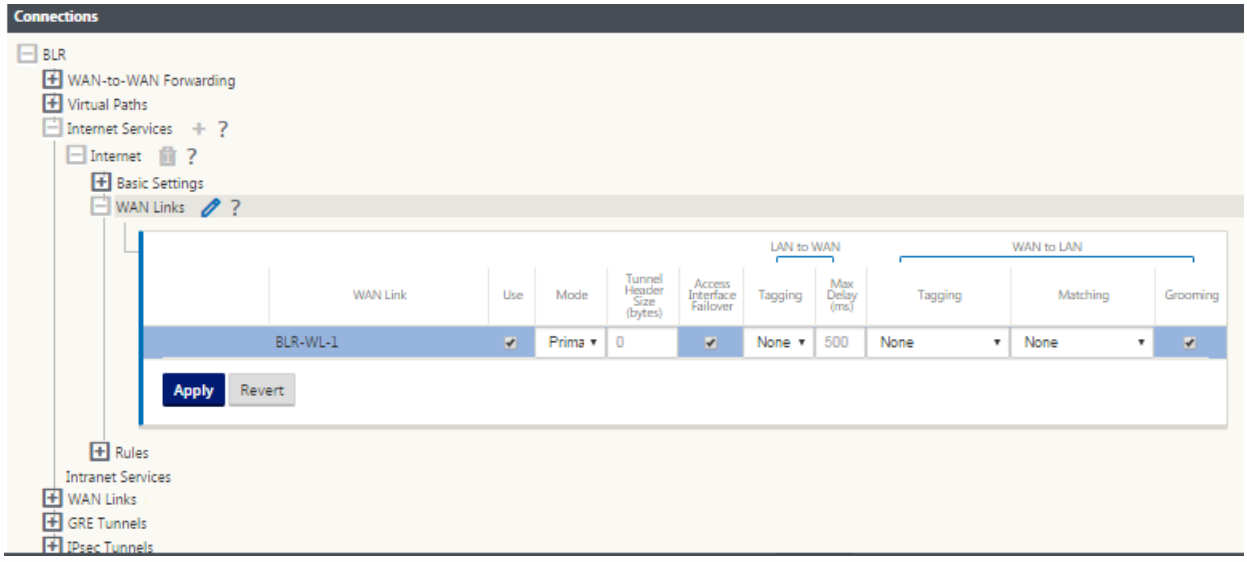

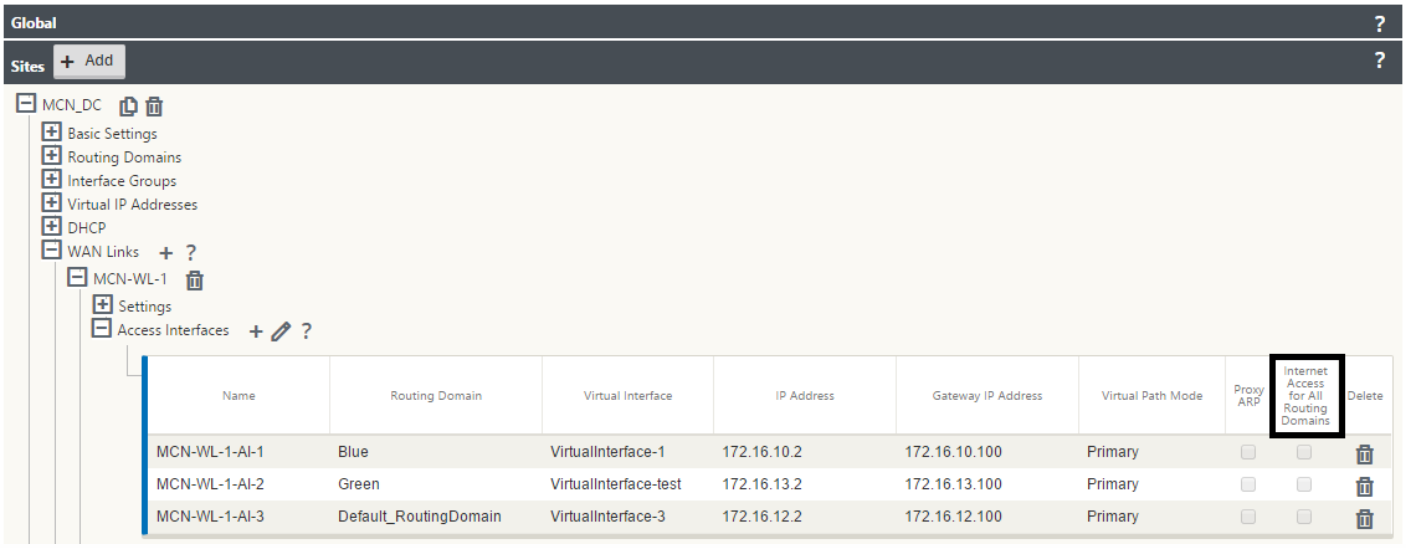

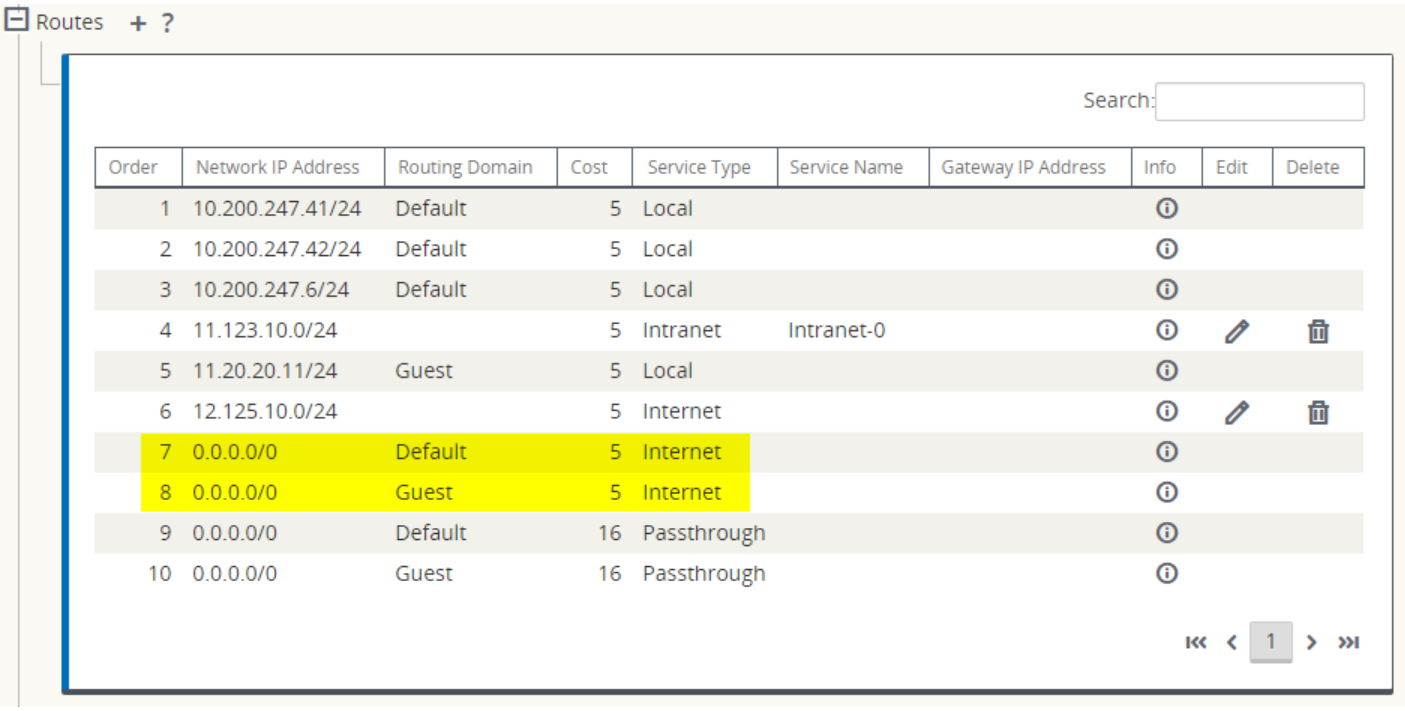

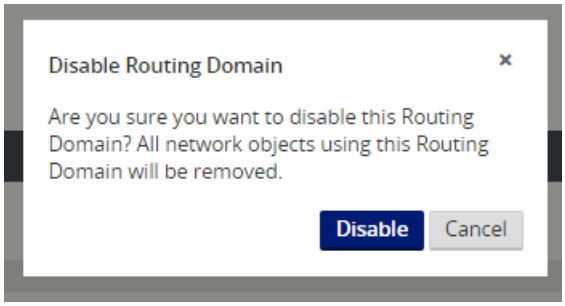

Flows Data

Toggle Columns

Both WAN Ingress and WAN Egress Flows

Routing Domain	Source IP Address	Dest IP Address	Direction	Source Port	Dest Port	IPP	IP DSCP	Hit Count	Service Type	Service Name	LAN GW IP	Age (mS)	Packets	Bytes	PPS	Customer kbps	Conduit Overhead kbps	IPsec Overhead kbps	Rule ID	Class	Class Type	Path	Hdr Compression Saved Bytes	Transmission Type
Guest	11.20.20.20	12.125.10.20	WAN Ingress	8	3335	ICMP	default	62	INTERNET	-	LOCAL	74	62	5208	1.013	0.681	0.000	0.000	202	N/A	N/A	N/A	N/A	N/A
Default	10.200.247.200	12.125.10.20	WAN Ingress	8	16185	ICMP	default	66	INTERNET	-	LOCAL	311	66	5544	1.009	0.678	0.000	0.000	202	N/A	N/A	N/A	N/A	N/A
Guest	12.125.10.20	11.20.20.20	WAN Egress	0	18456	ICMP	default	62	INTERNET	-	LOCAL	94	62	5208	1.013	0.681	0.000	0.000	202	N/A	N/A	N/A	N/A	N/A
Default	12.125.10.20	10.200.247.200	WAN Egress	0	3968	ICMP	default	66	INTERNET	-	LOCAL	328	66	5544	1.008	0.678	0.000	0.000	202	N/A	N/A	N/A	N/A	N/A

Total INGRESS flows displayed: 2 out of 2
Total EGRESS flows displayed: 2 out of 2

Routes for routing domain : Guest

Filter: in Any column Apply

Show 100 entries Showing 1 to 5 of 5 entries First Previous 1 Next Last

Num	Network Addr	Gateway IP Address	Service	Firewall Zone	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
0	11.20.20.0/24	*	Local	Default_LAN_Zone	YES	*	Angelina-CFB	Static	-	-	5	318	YES	N/A	N/A
1	11.10.10.0/24	*	DC-Angelina-CFB	Default_LAN_Zone	YES	*	DC	Static	-	-	5	0	YES	N/A	N/A
2	0.0.0.0/0	*	Internet	Untrusted_Internet_Zon	YES	*	*	Static	-	-	5	159	YES	N/A	N/A
3	0.0.0.0/0	*	Passthrough	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
4	0.0.0.0/0	*	Discard	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

Showing 1 to 5 of 5 entries First Previous 1 Next Last

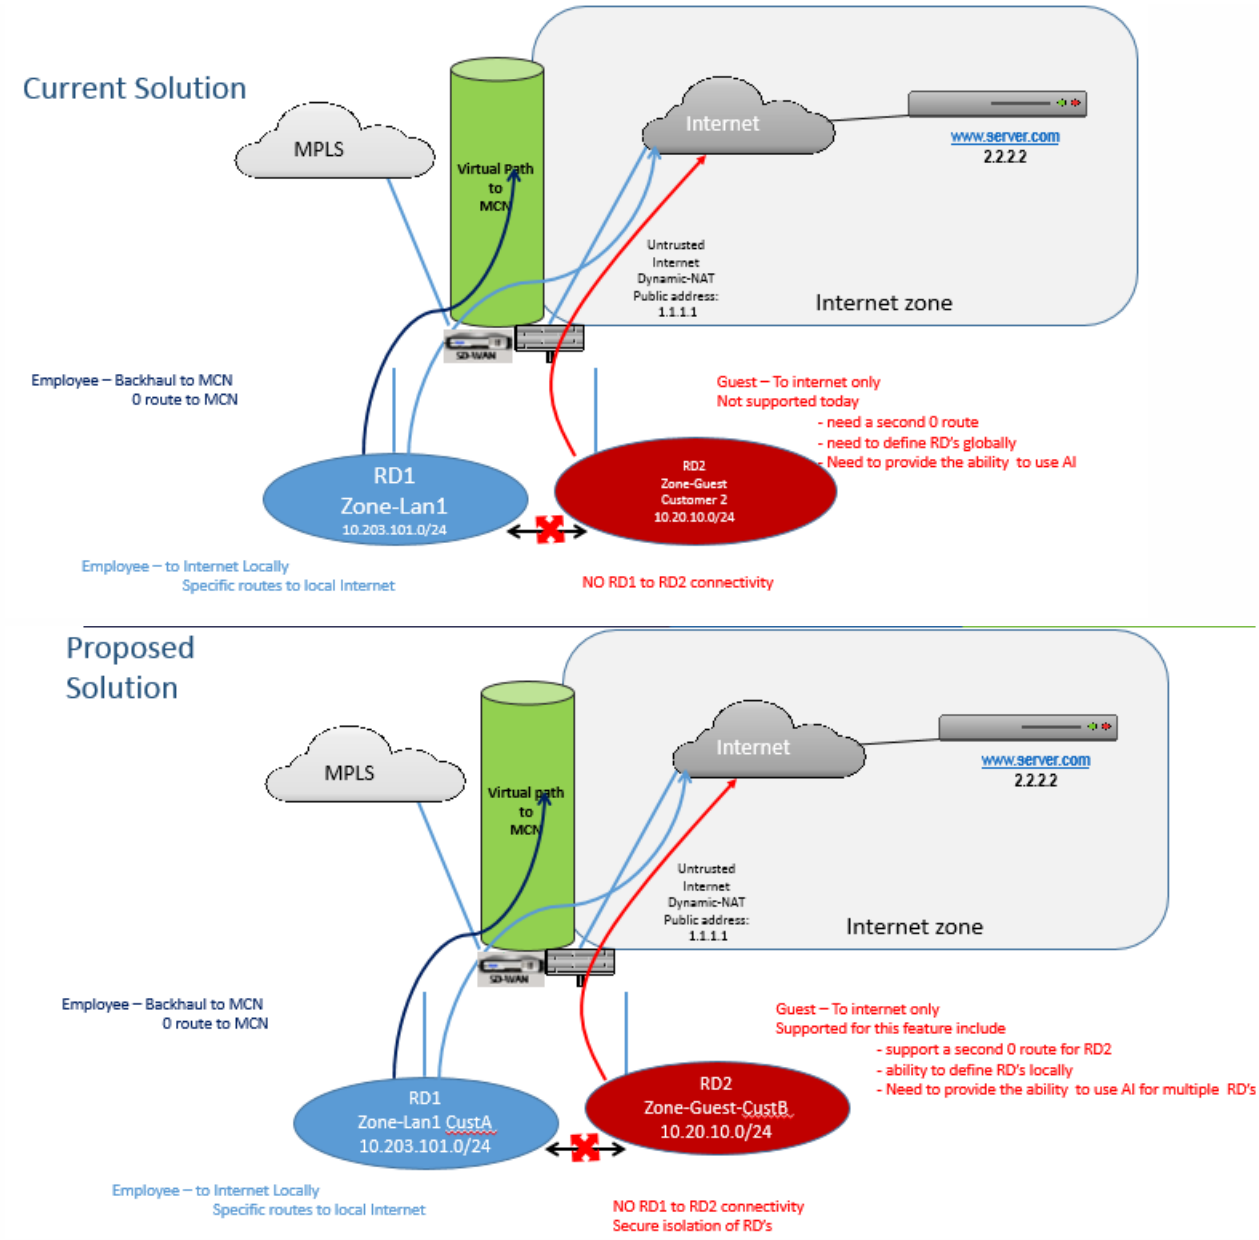

Limitations

-
-
-
-
-

-
-

OSPF Overview

-
-

How To Configure OSPF

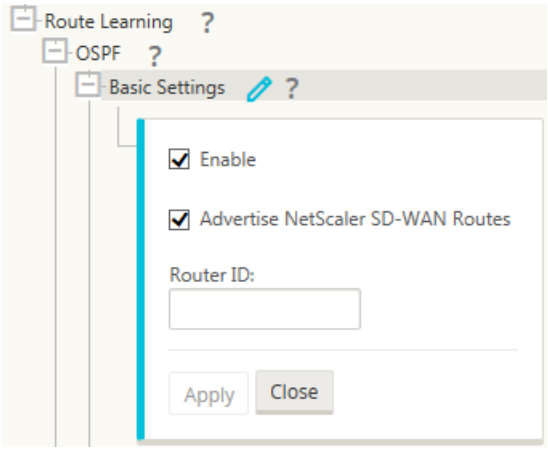

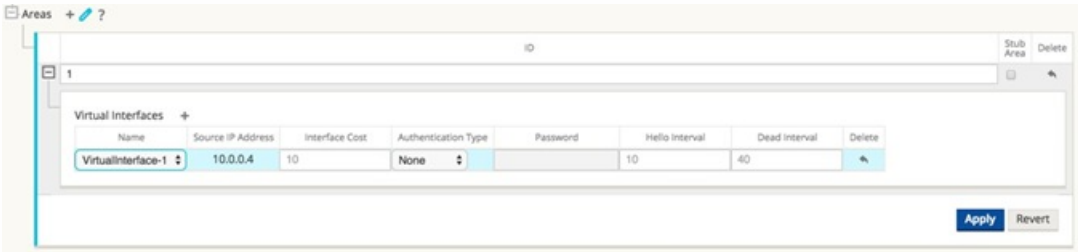

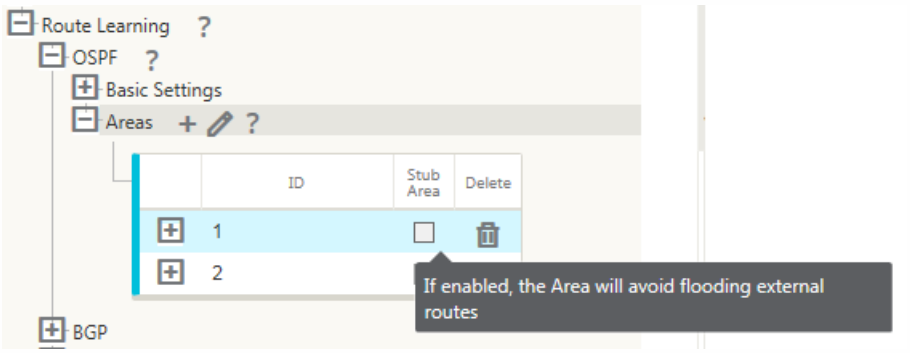

BGP Overview

How To Configure BGP

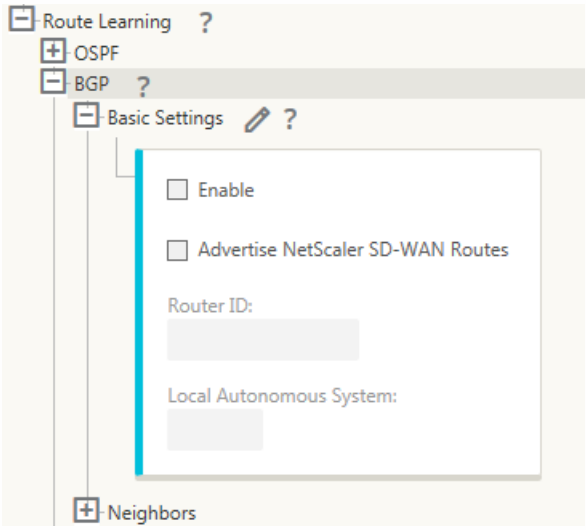

How To Monitor Route Statistics

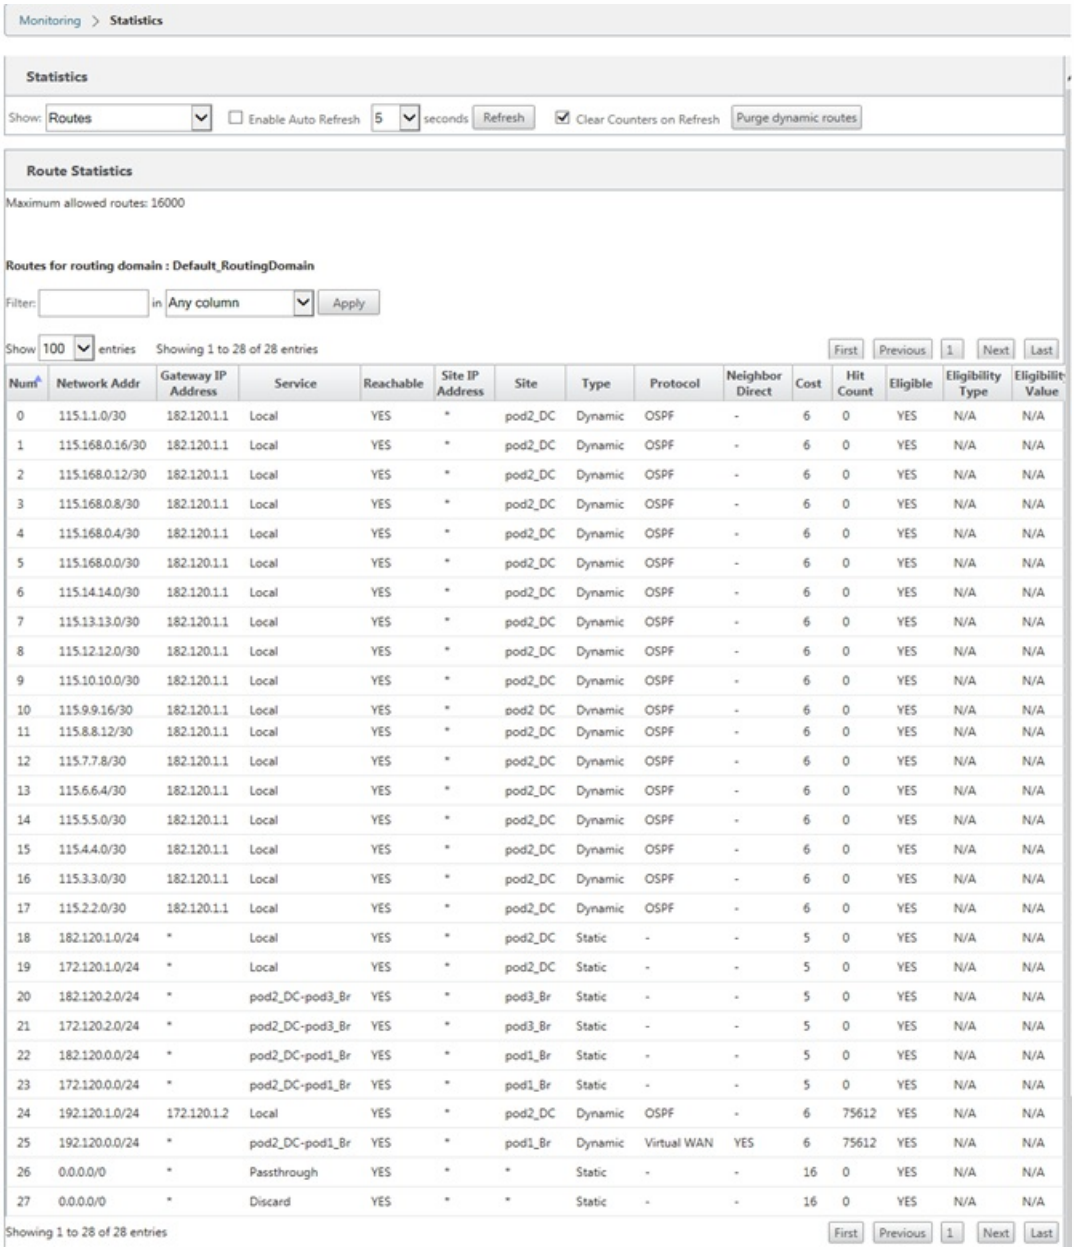

Exterior BGP (eBGP)

-
-
-
-
-
-

LAN Side: Dynamic Route Learning

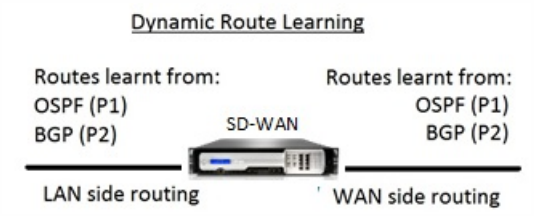

WAN Side: Dynamic Route Sharing

OSPF Deployment Modes

-
-
-

-
-

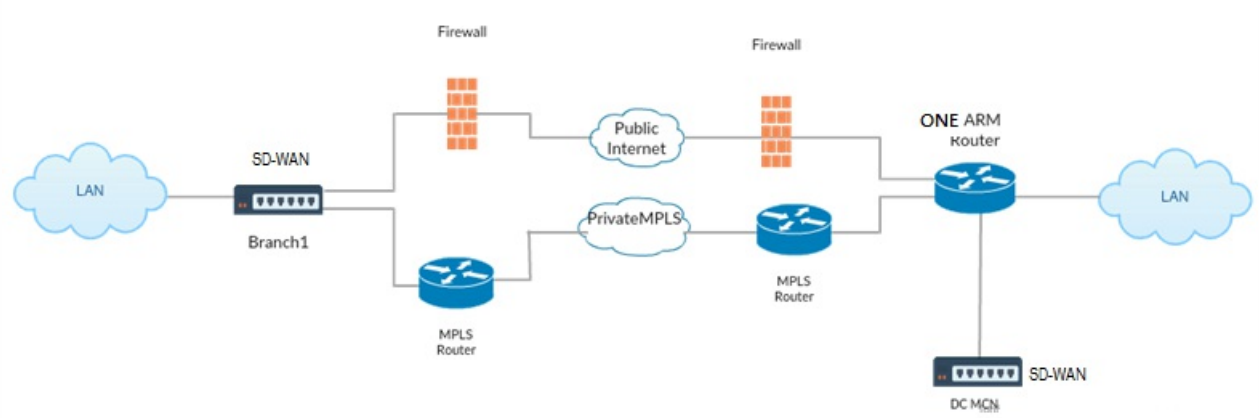

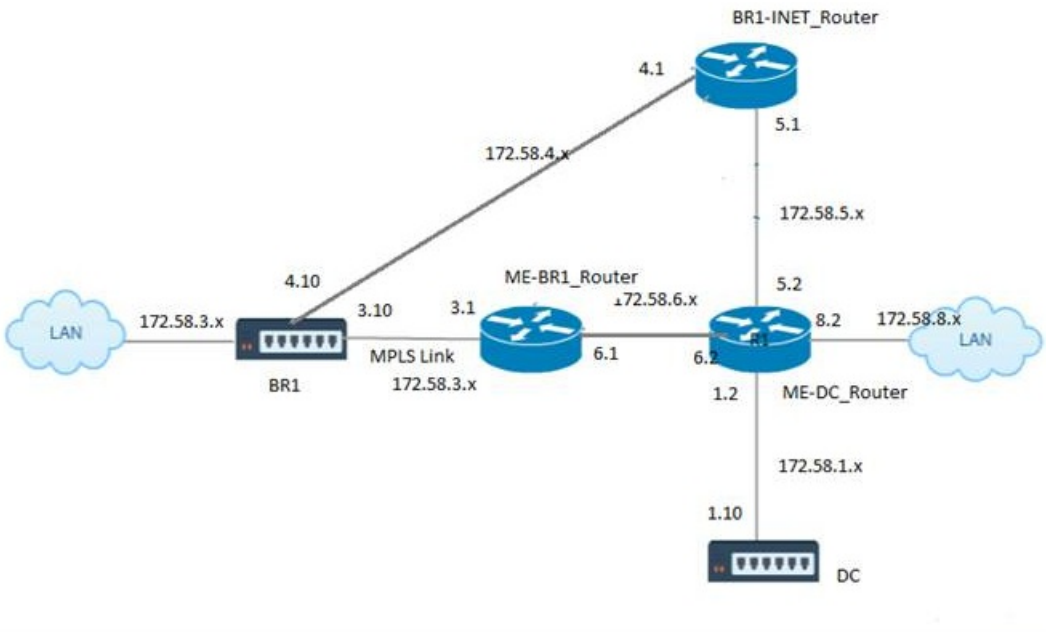

-
-

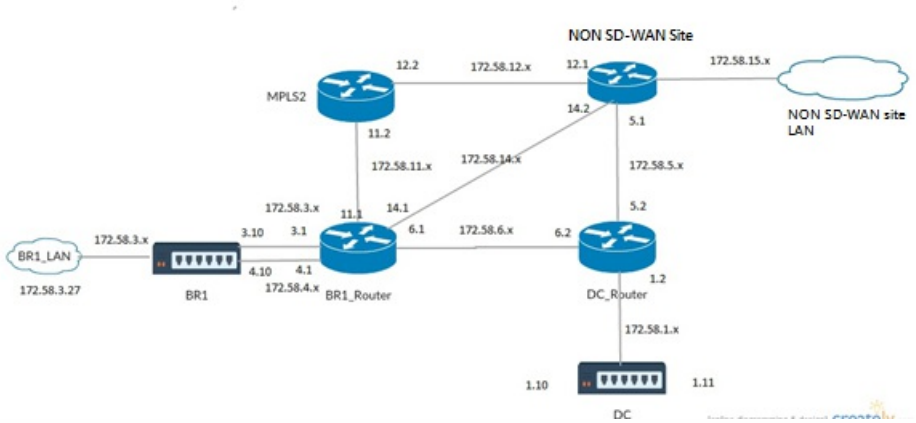

Implementing OSPF with SD-WAN Network in High Availability Setup

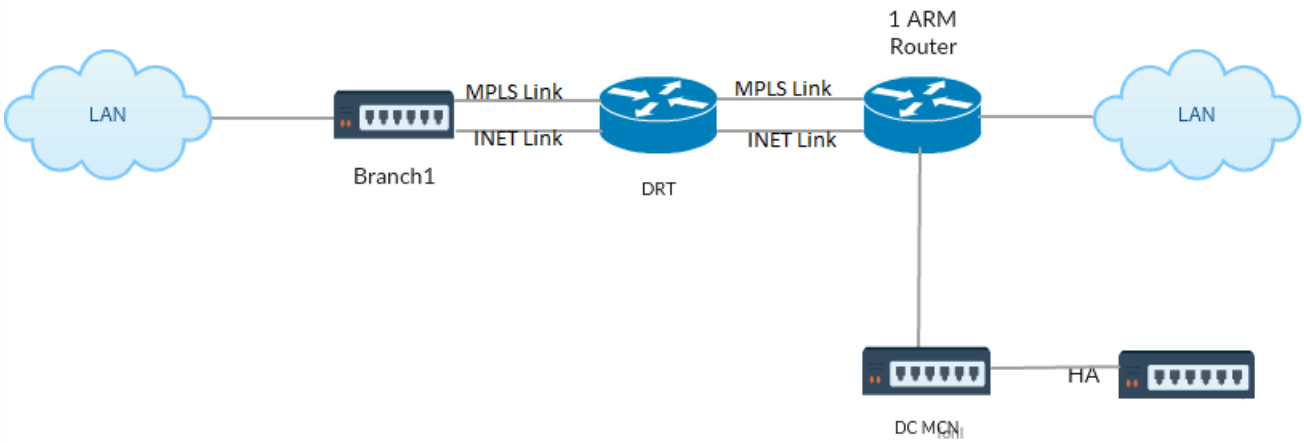

-
-

-

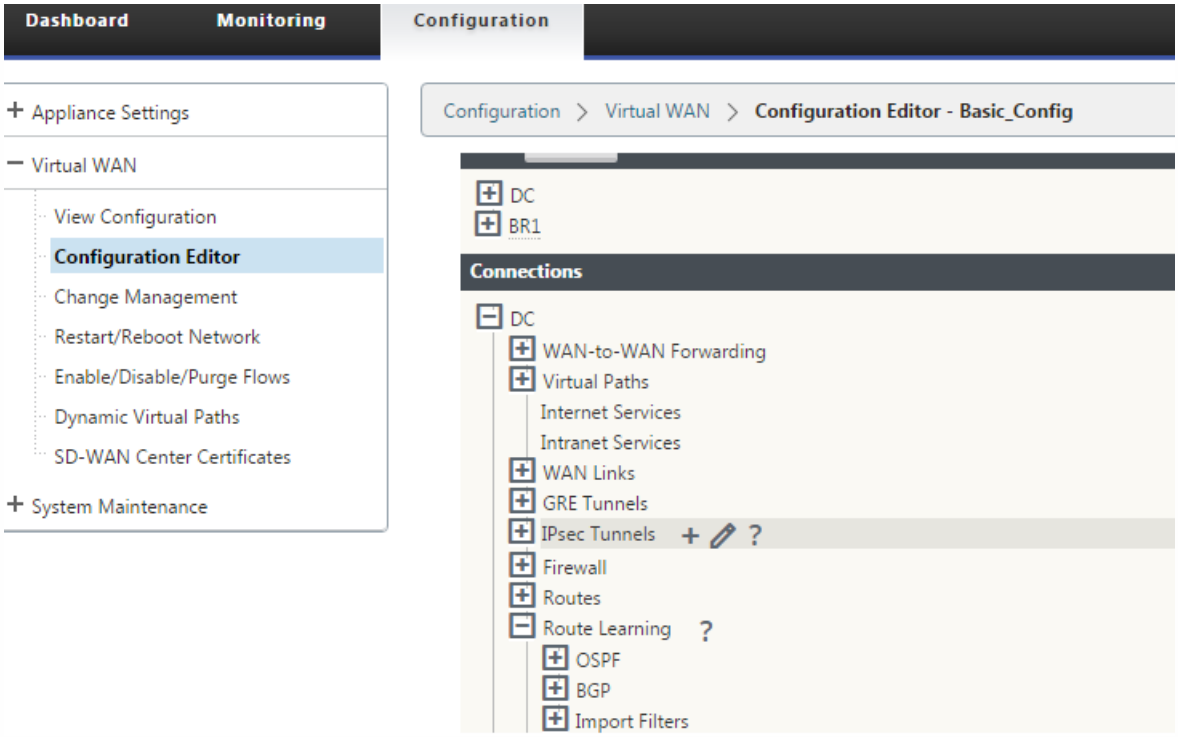

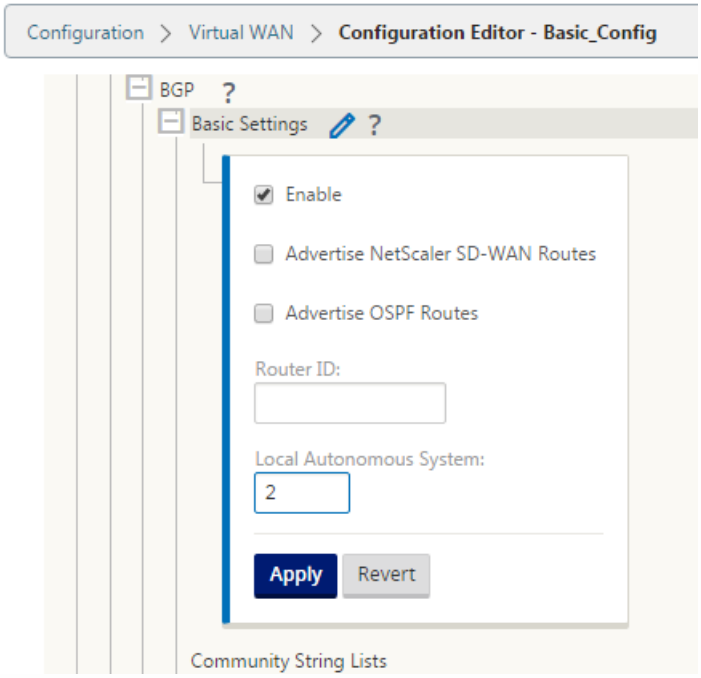

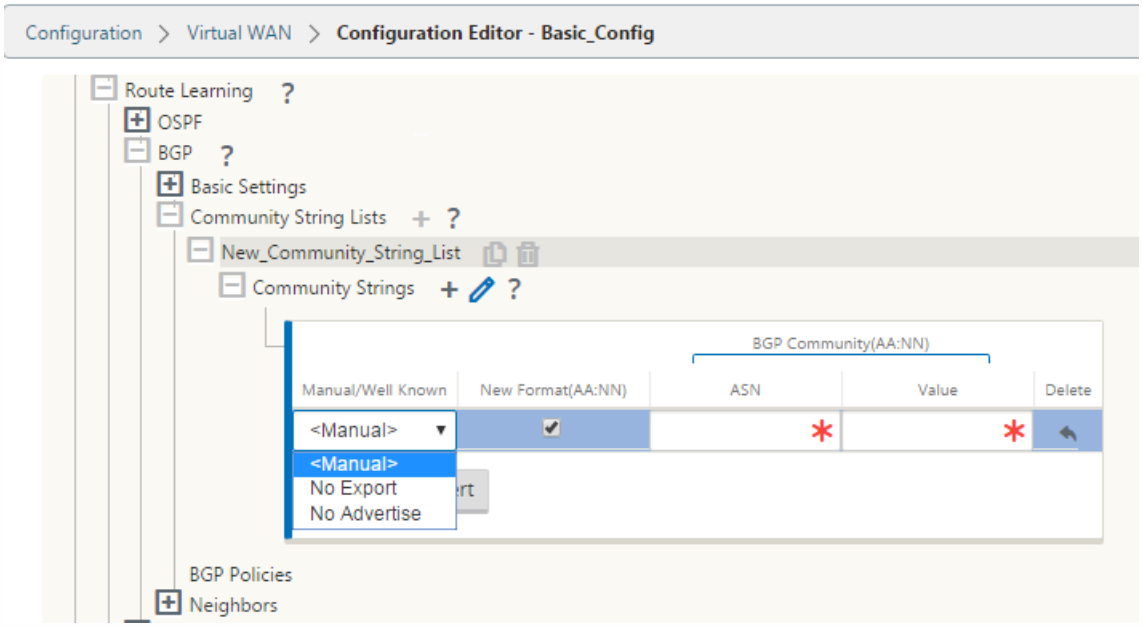

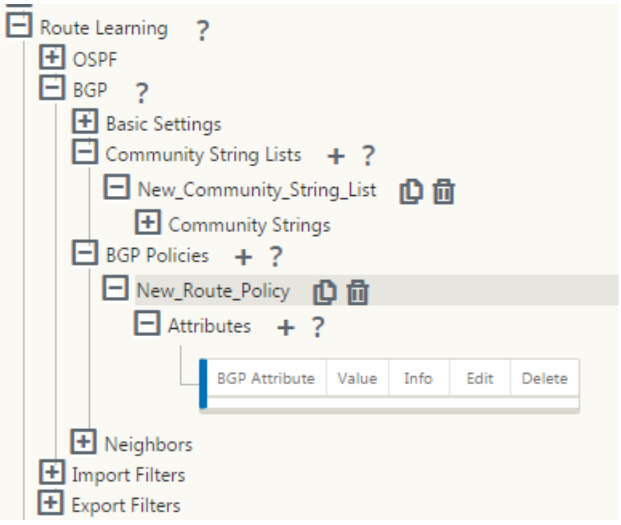

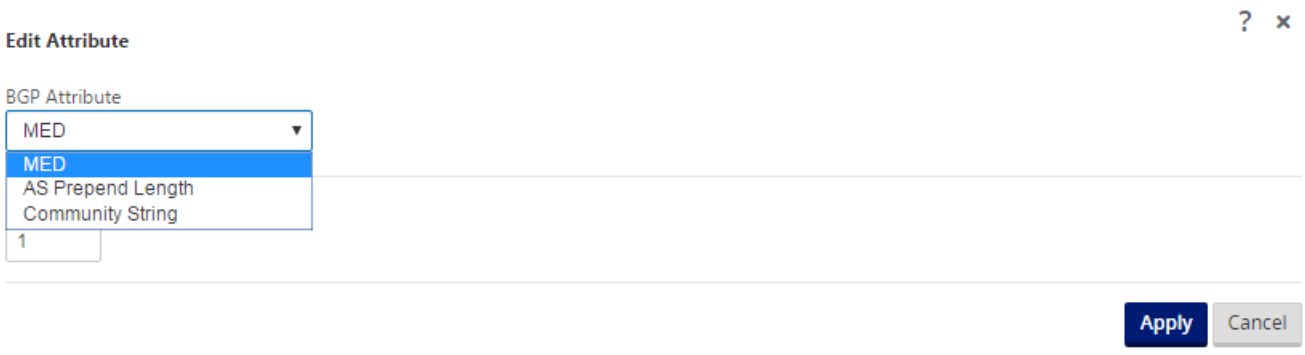

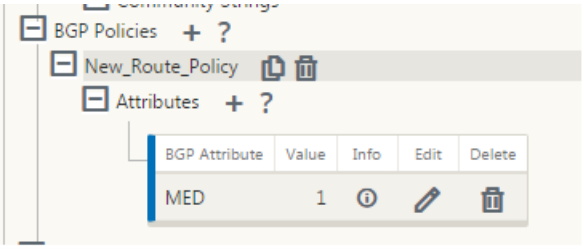

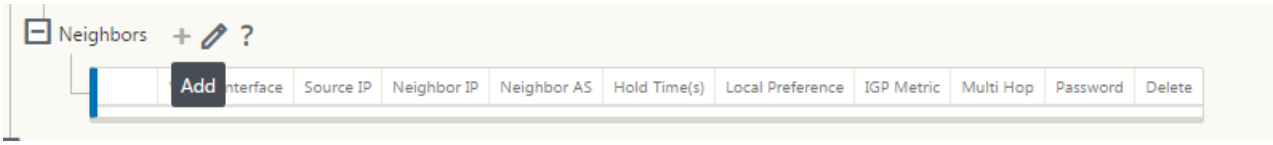

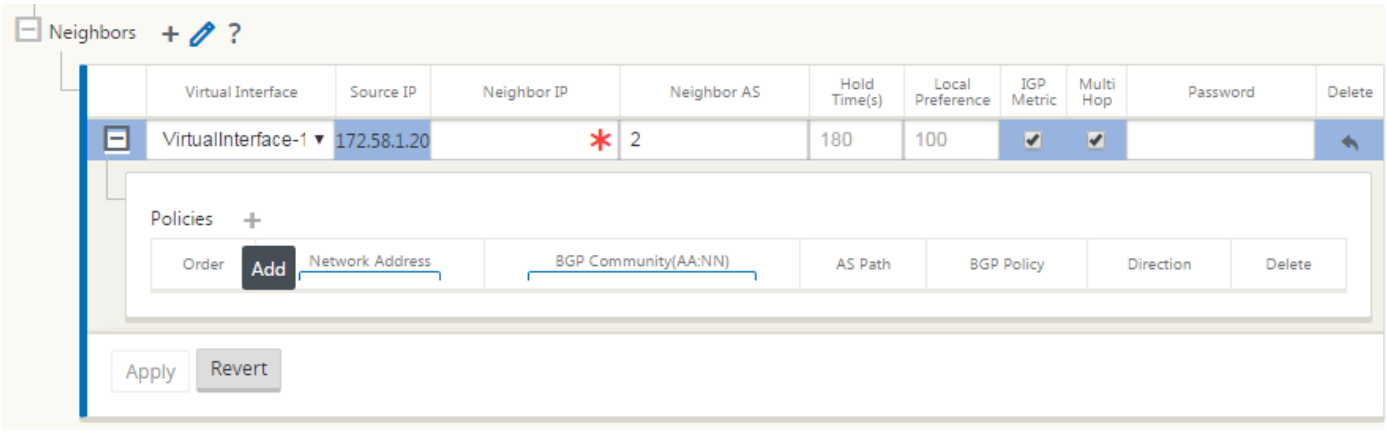

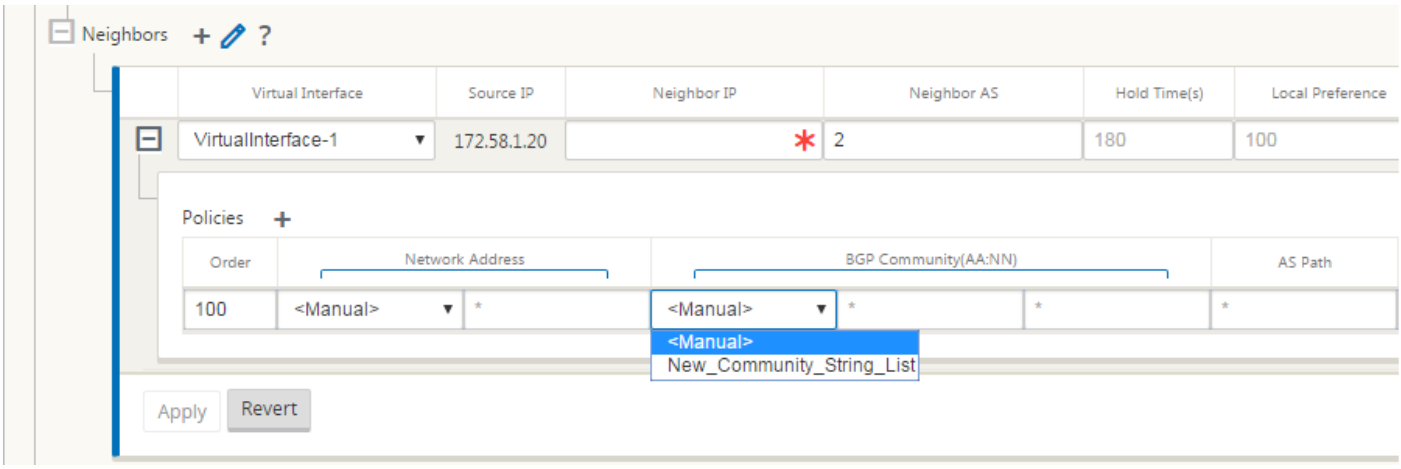

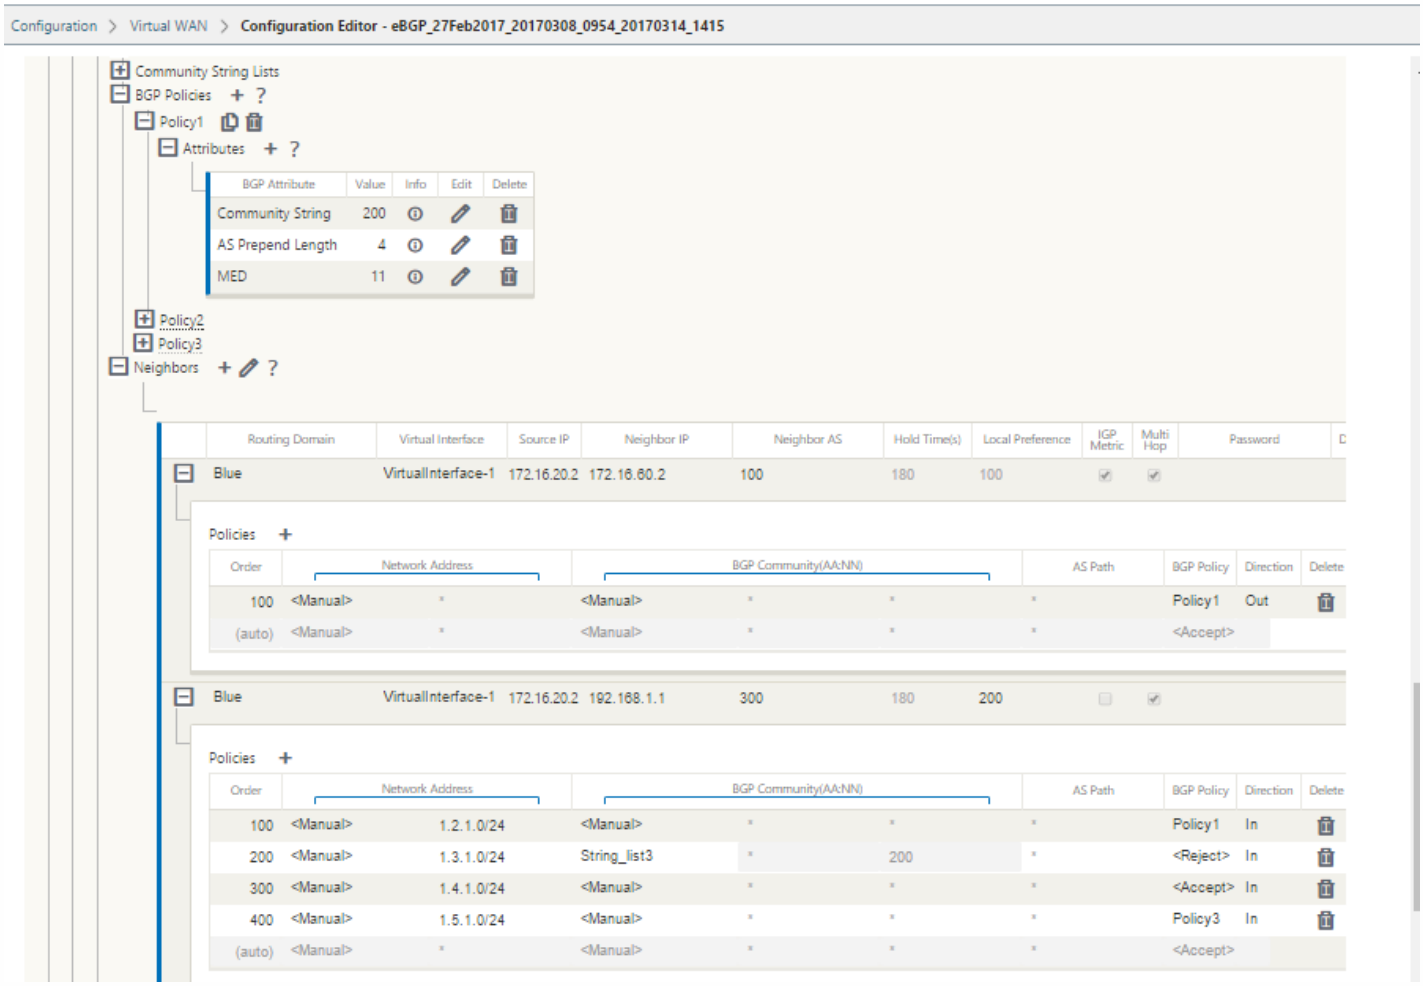

-
-
-
-

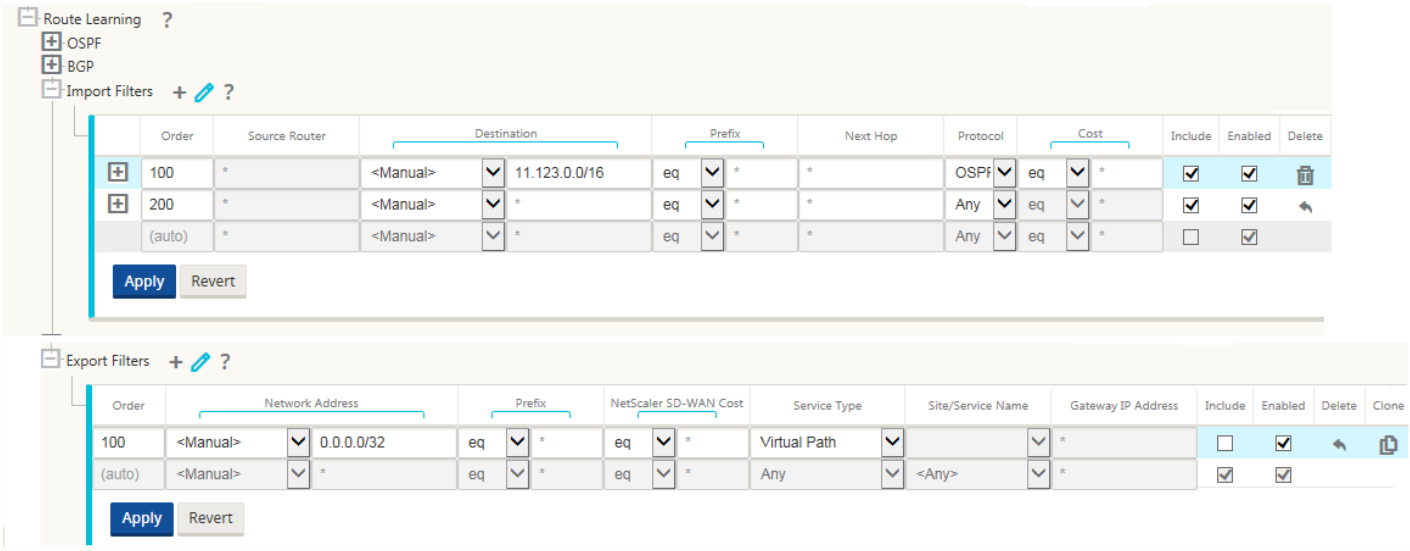

		• • • • • •
		•

		•
		• • • • •
		•
		• • • • • • •
		•
		•
		•
		•

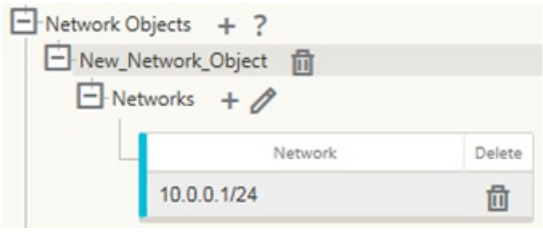

-
-
-
-
-
-
-

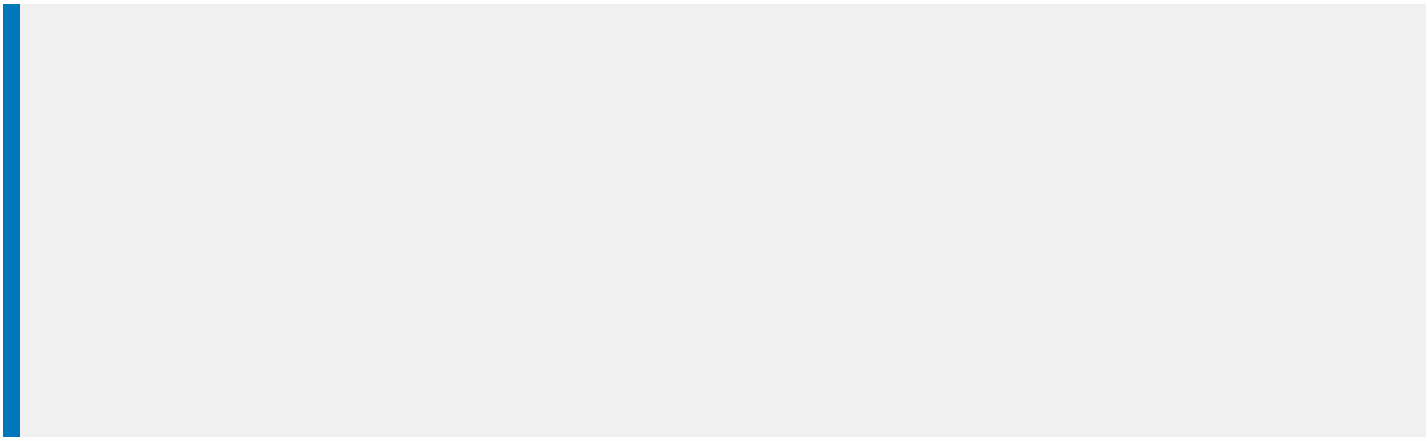

-
-
-
-
-

-
-

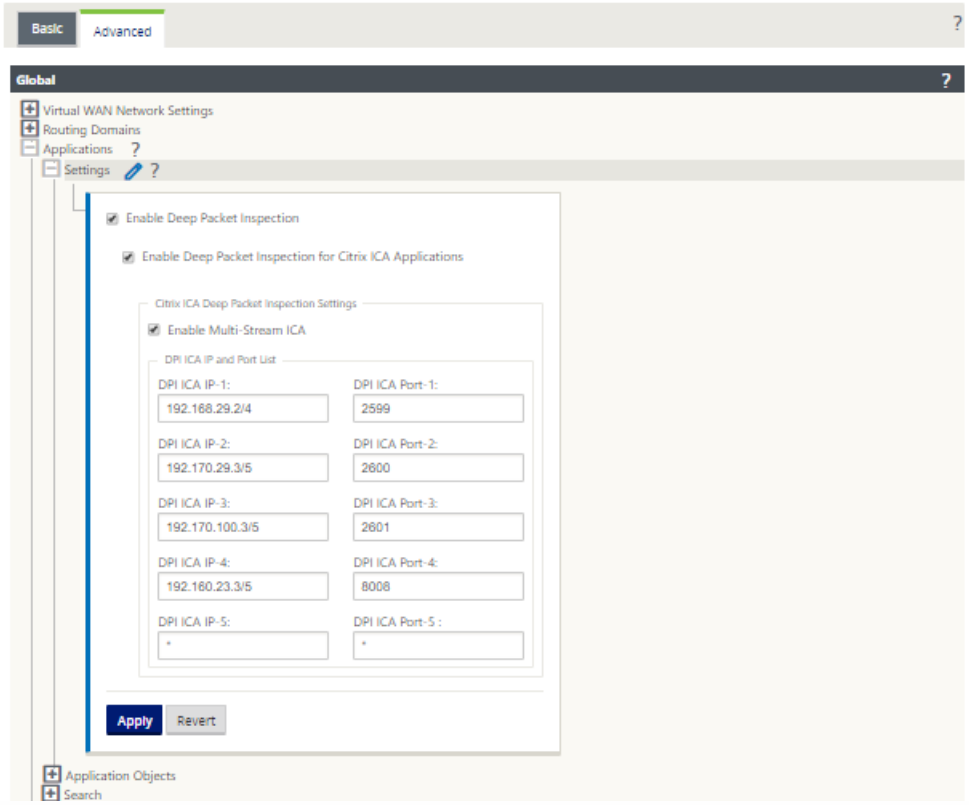

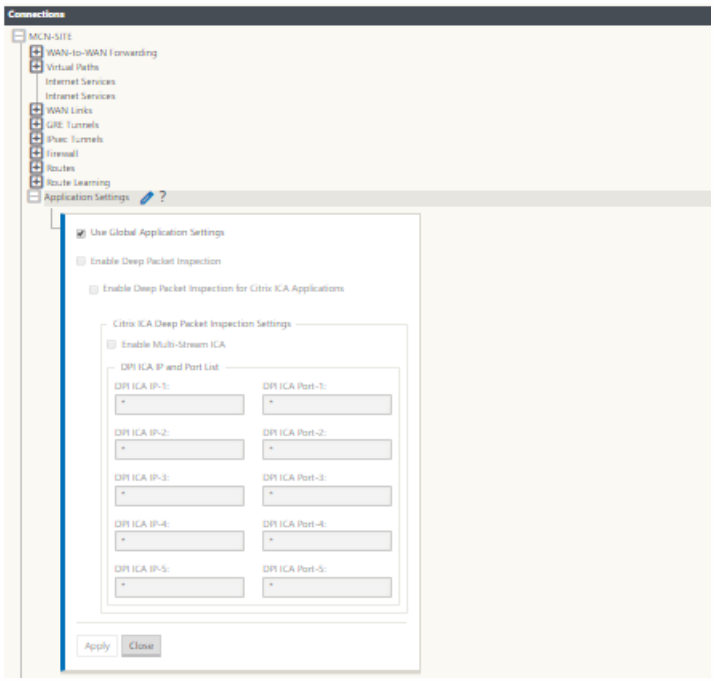

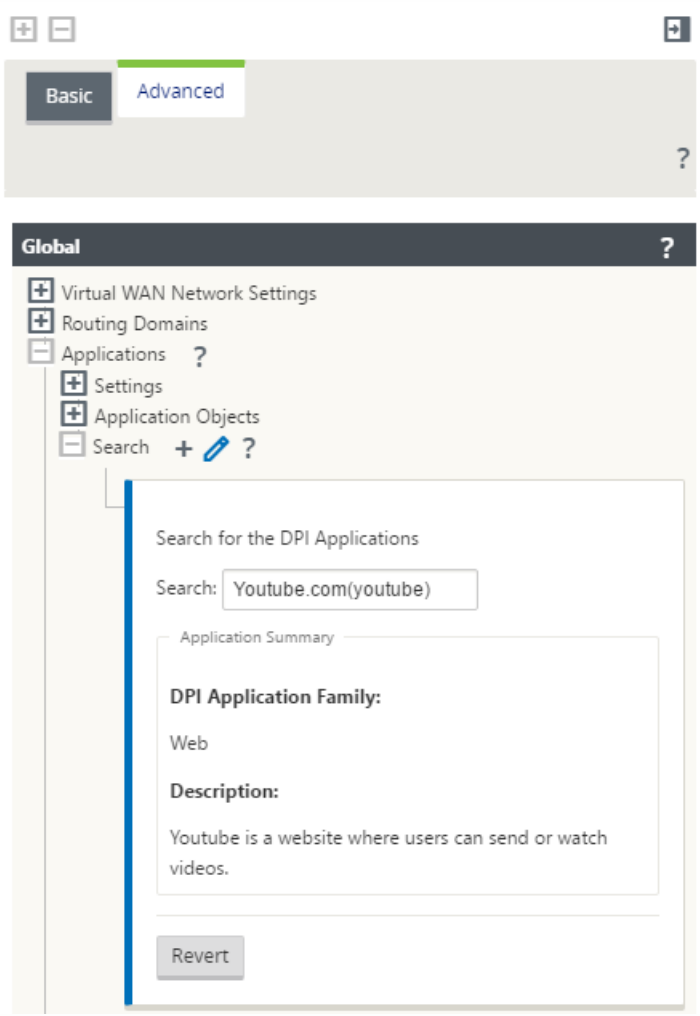

-
-
-

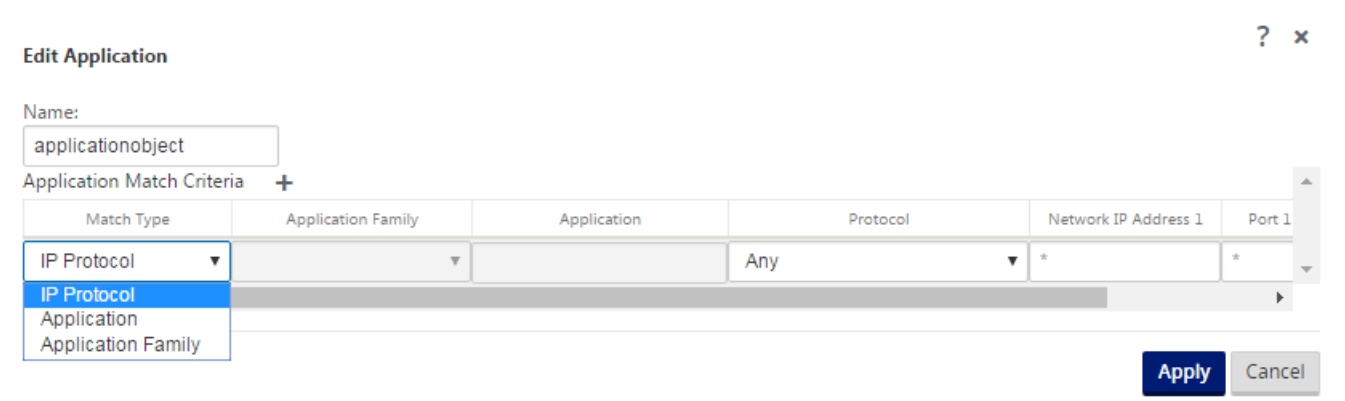

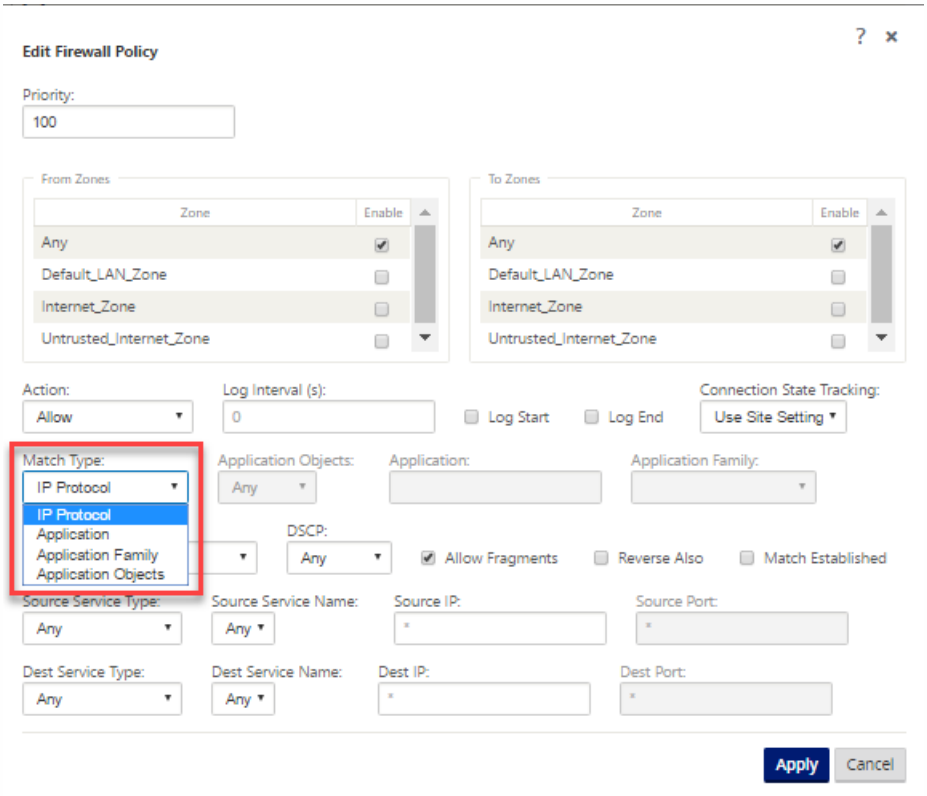

-
-
-

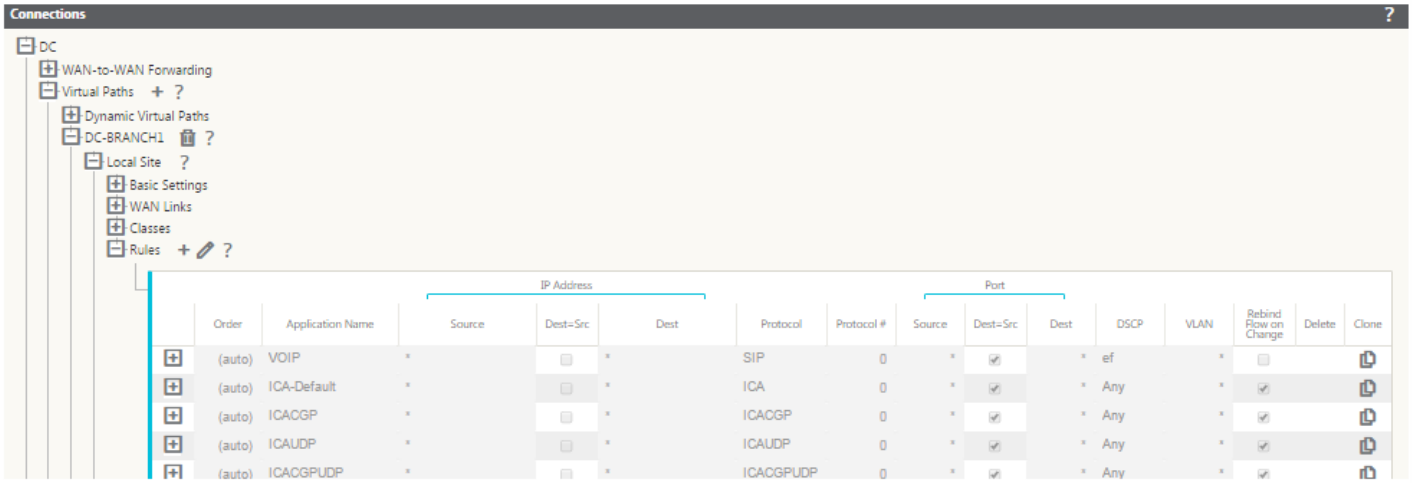

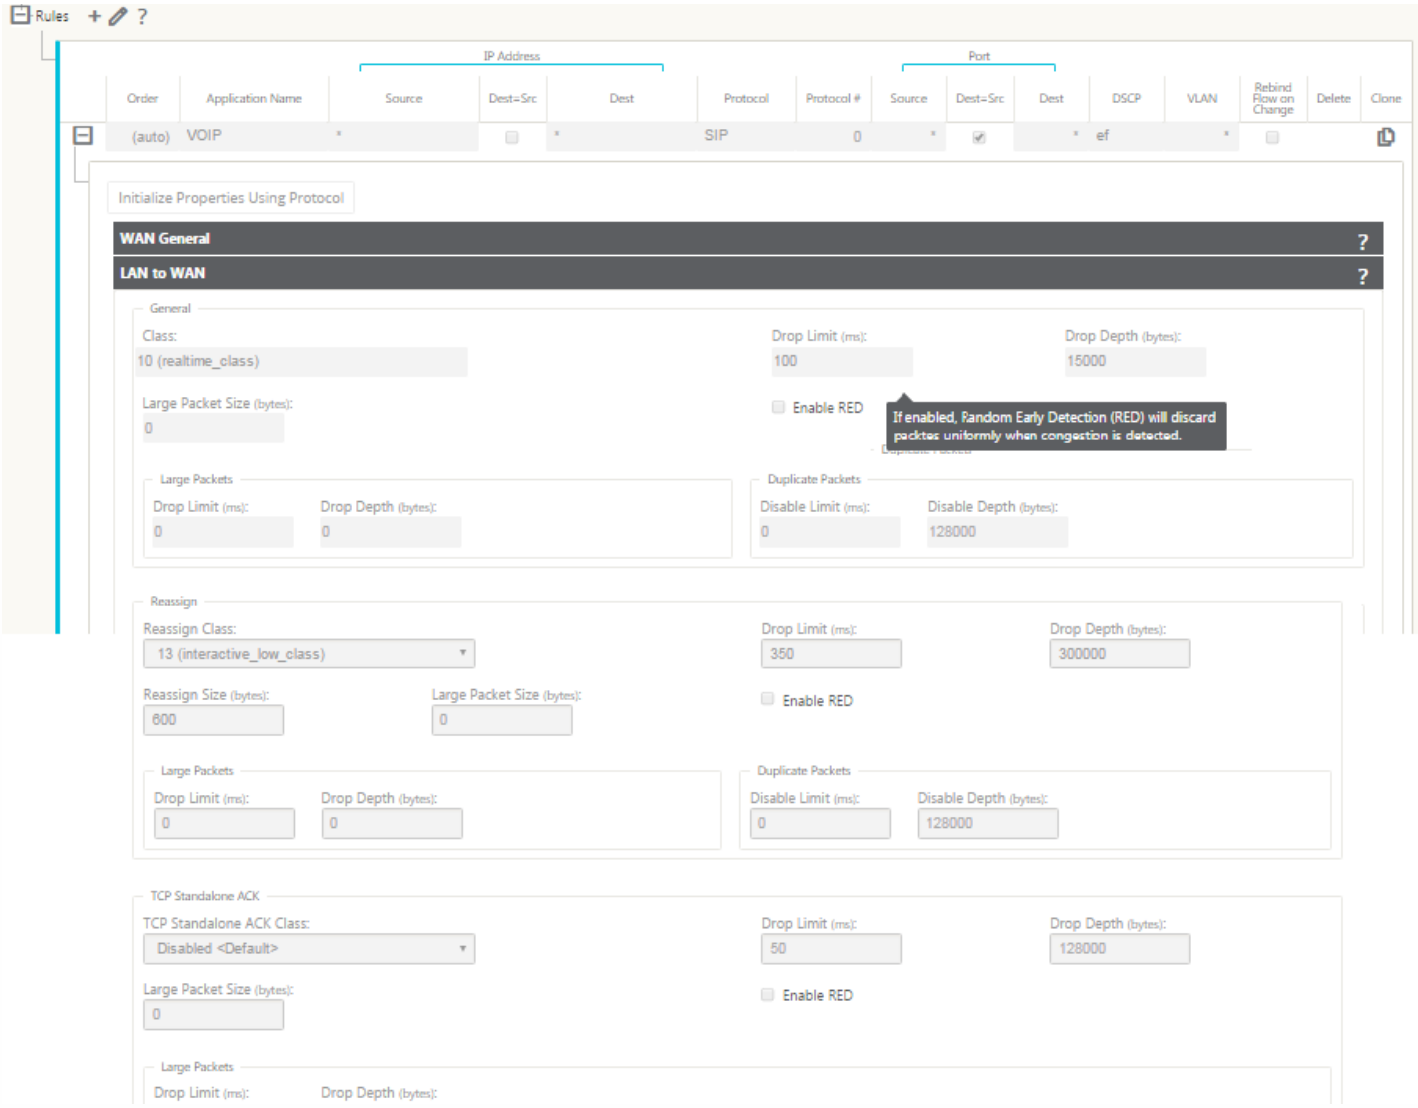

-
-
-

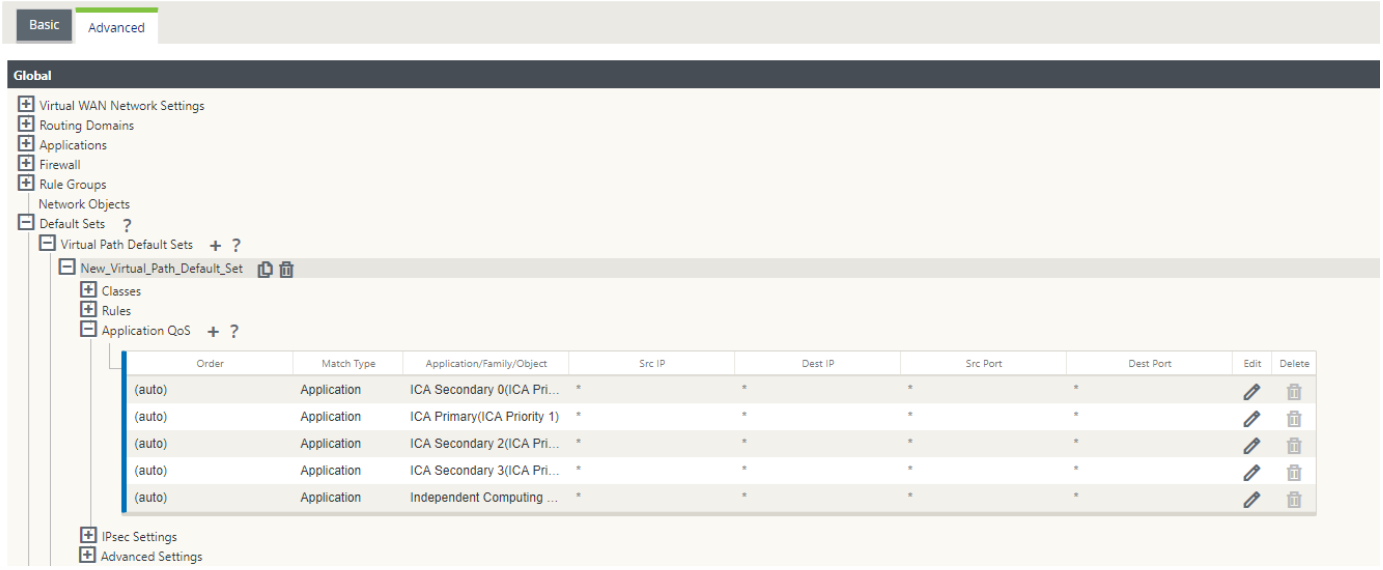

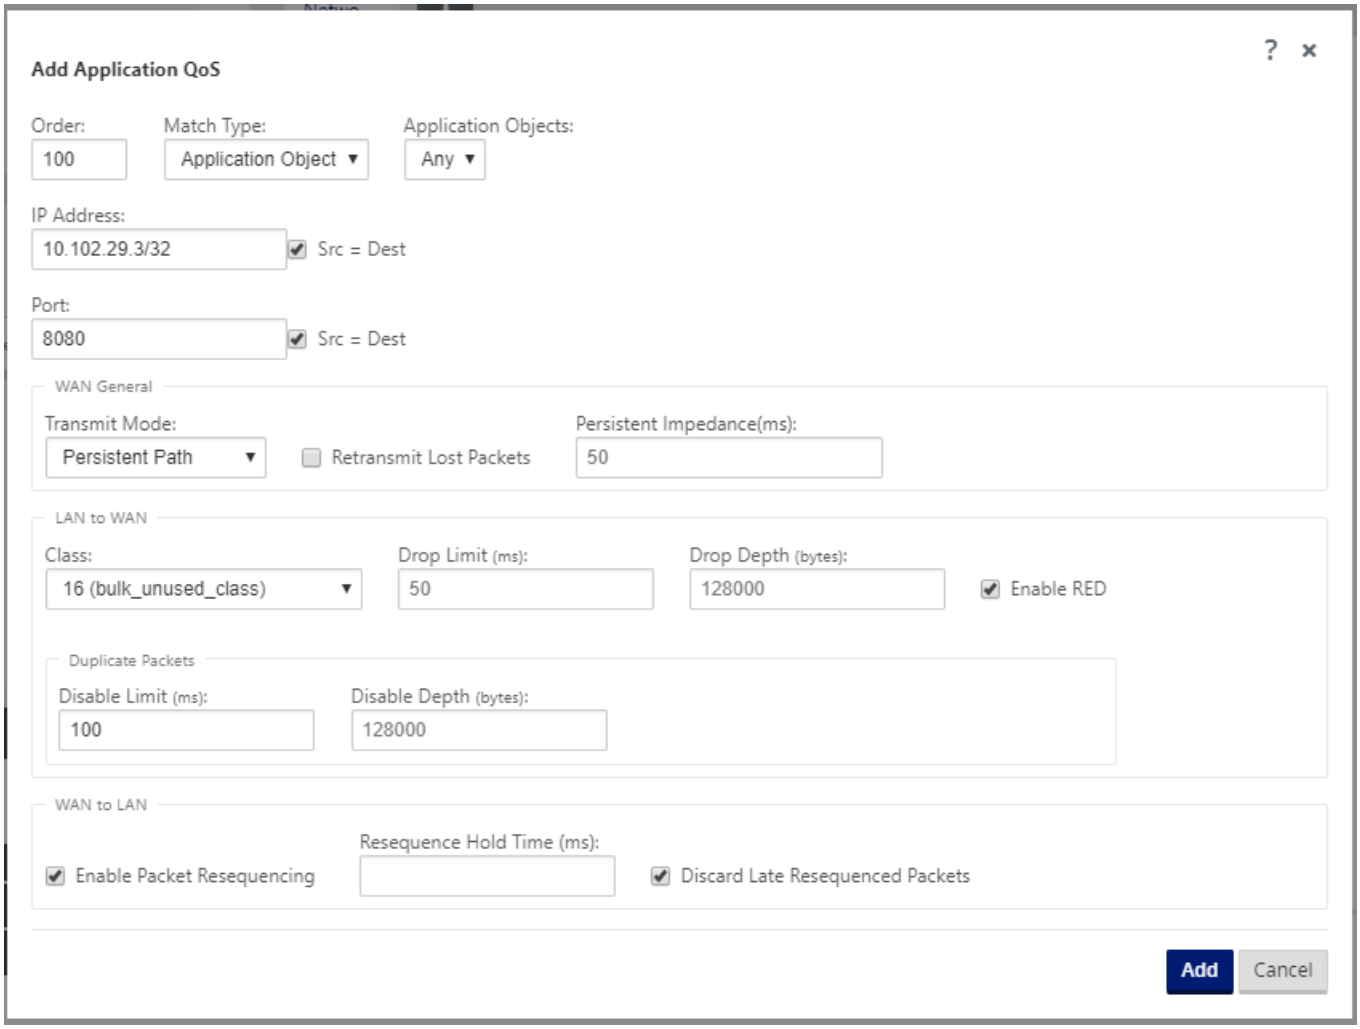

-
-
-
-
-
-
-

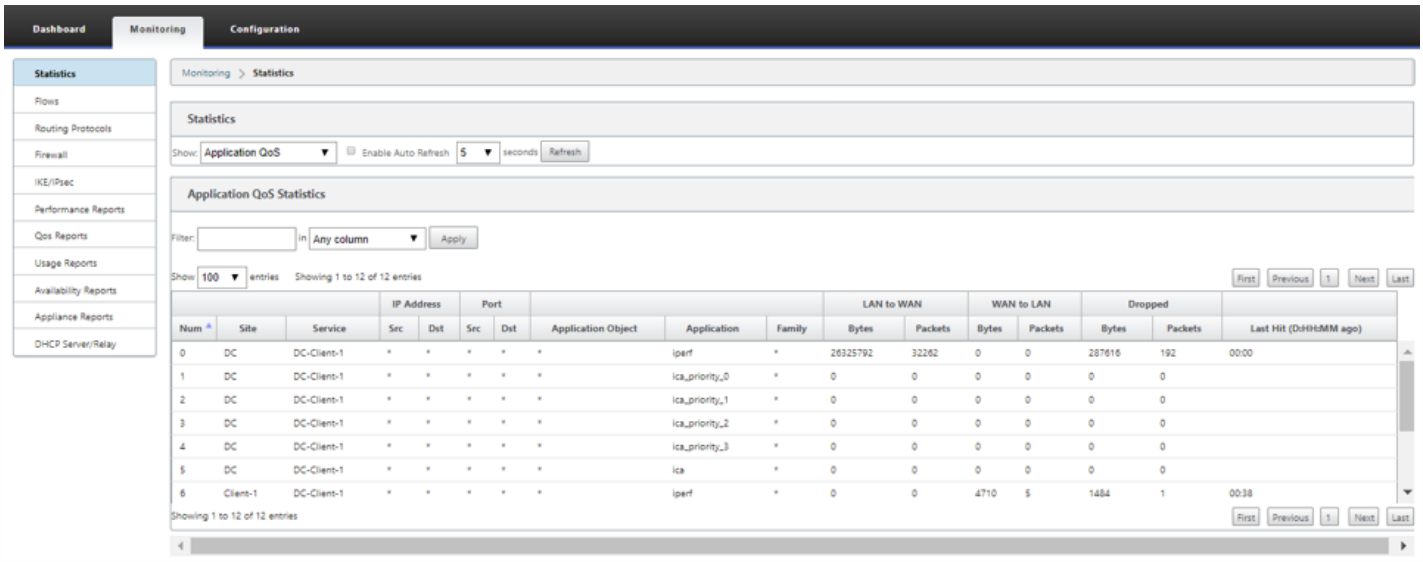

To configure new WAN Link Access Type for Private MPLS

1. In the Configuration Editor, click + (Add) under **Sites** > **[Site Name]** >**WAN Links**, the Add WAN Link pop-up appears.

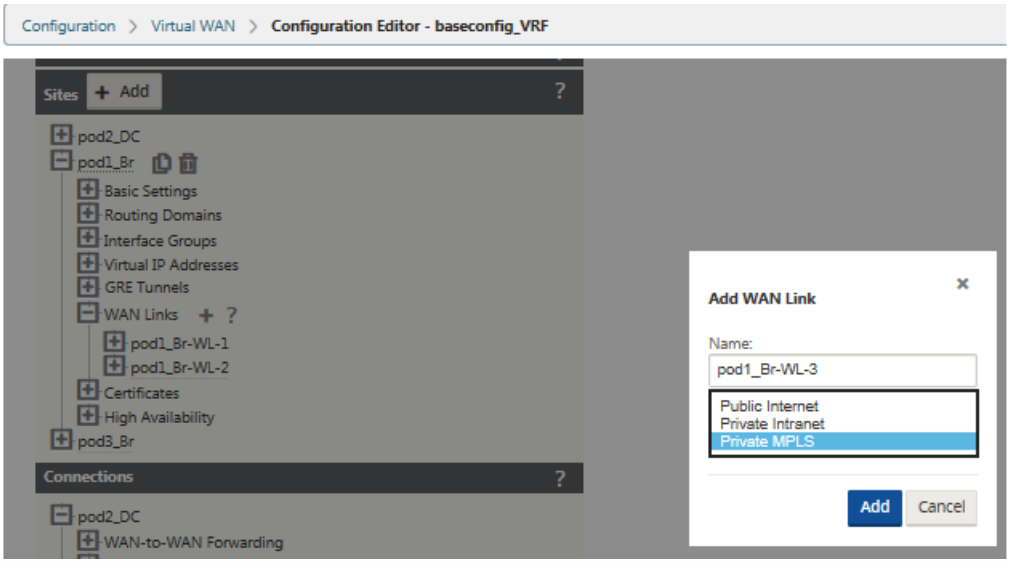

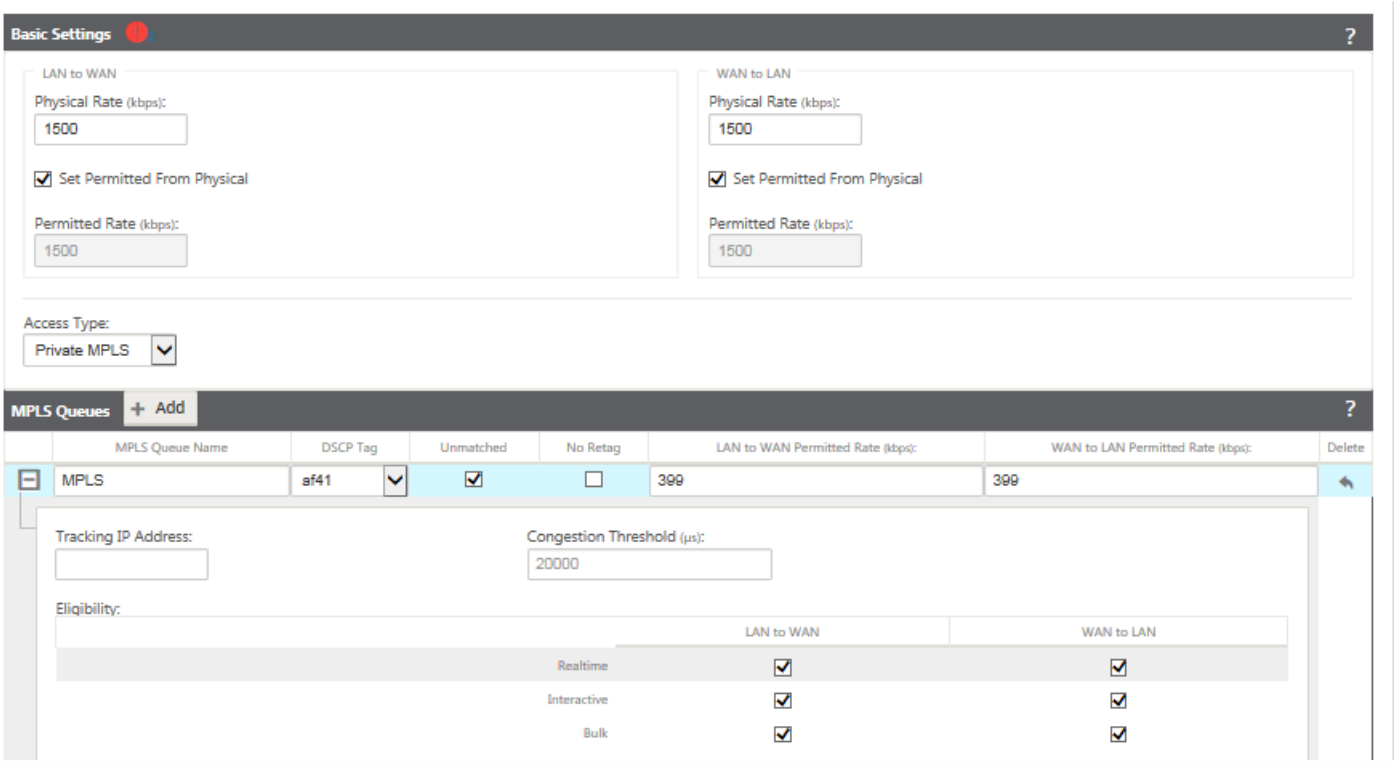

Once the Private MPLS WAN Link with its MPLS Queues is defined, you should assign an Autopath Group for the WAN Link under a specific Virtual Path definition.

To assign autopath group

1. Go to **Connections** > **[Site Name]** > **WAN Links** >[**MPLS WAN Link Name**] > **Virtual Paths** > [**Virtual Path Name**] > [**Local Site**] > **WAN Links** and click **Edit** ().

2 . Click the **Autopath Group** drop-down menu and choose from the available groups. By default, MPLS Queues inherit the Autopath Group assigned to the MPLS WAN Link. You may choose to set the individual MPLS Queues to Inherit the chosen Autopath Group or choose an alternate from the Autopath Group drop-down menu for each MPLS Queue.

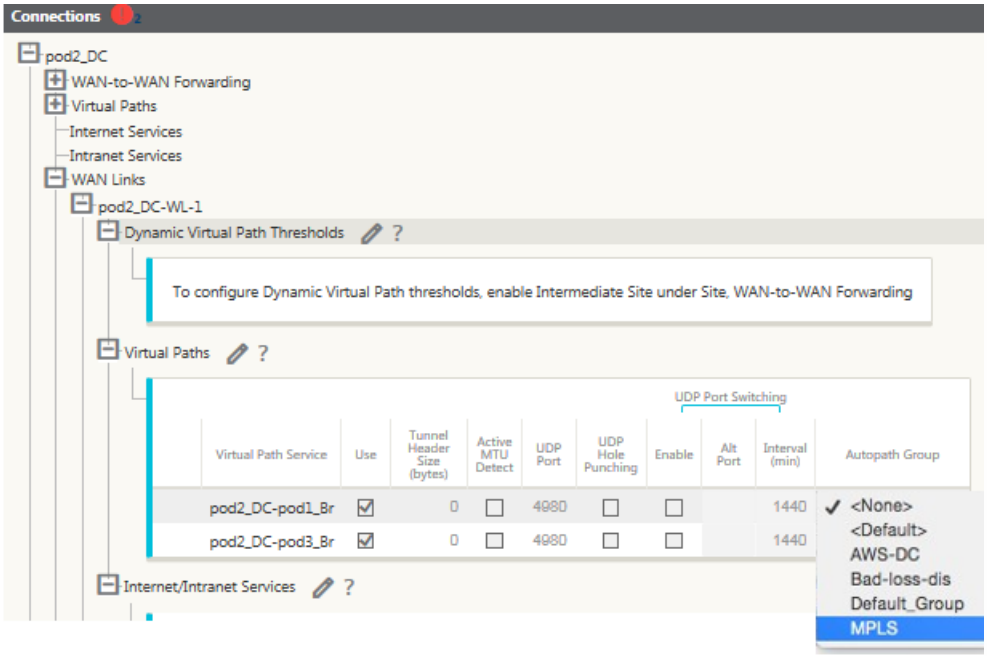

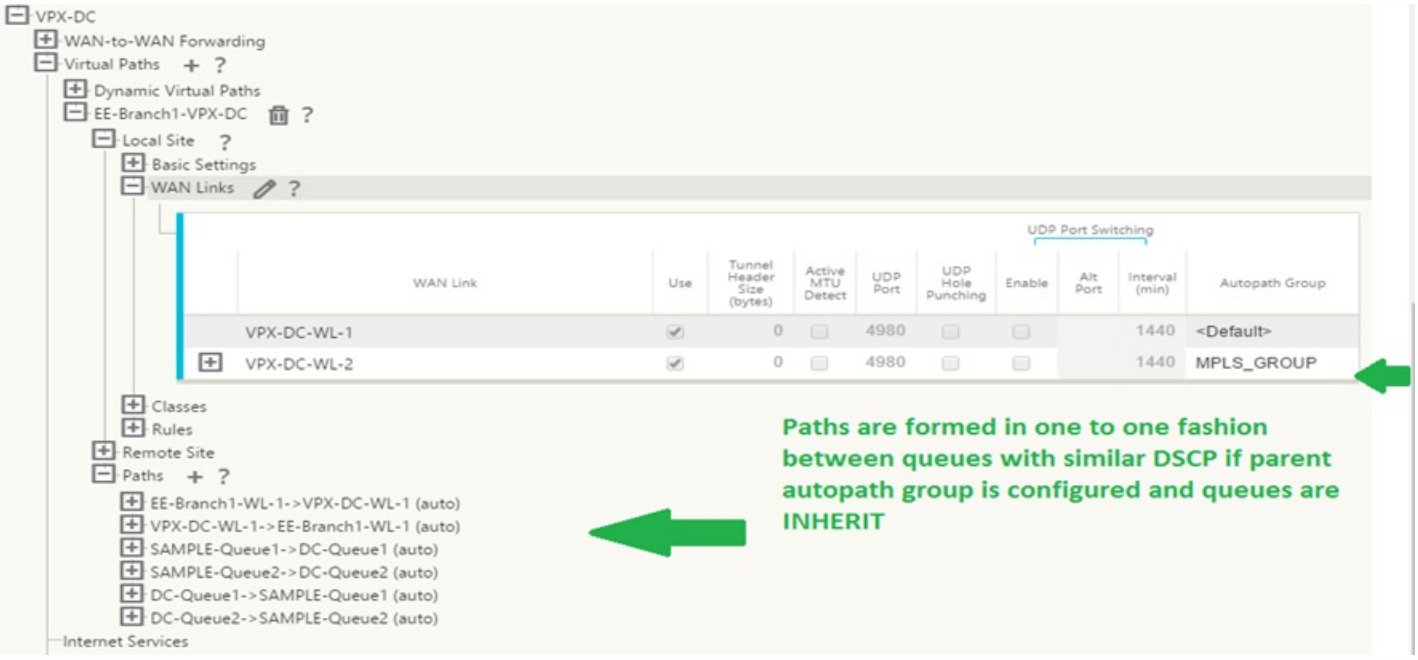

The SD-WAN web interface now allows you to view the permitted rate for WAN Links and WAN Link Usages and whether a WAN Link, Path, or Virtual Path may be in a congested state. In the previous releases, this information was only available in SD-WAN log files and through the CLI. These options are now available in the web interface to assist in troubleshooting.

View Permitted Rate

Permitted Rate is the amount of bandwidth that a particular WAN Link, Virtual Path Service, Intranet Service, or Internet Service is permitted to use at a given point in time. The permitted rate for a WAN Link is static, and is defined explicitly in the SD-WAN configuration. The permitted rate for a Virtual Path Service, Intranet Service, or Internet Service will fluctuate over time, in response to congestion, user demand, and Fair Shares, but will always be greater than or equal to the Minimum Reserved Bandwidth for the Service.

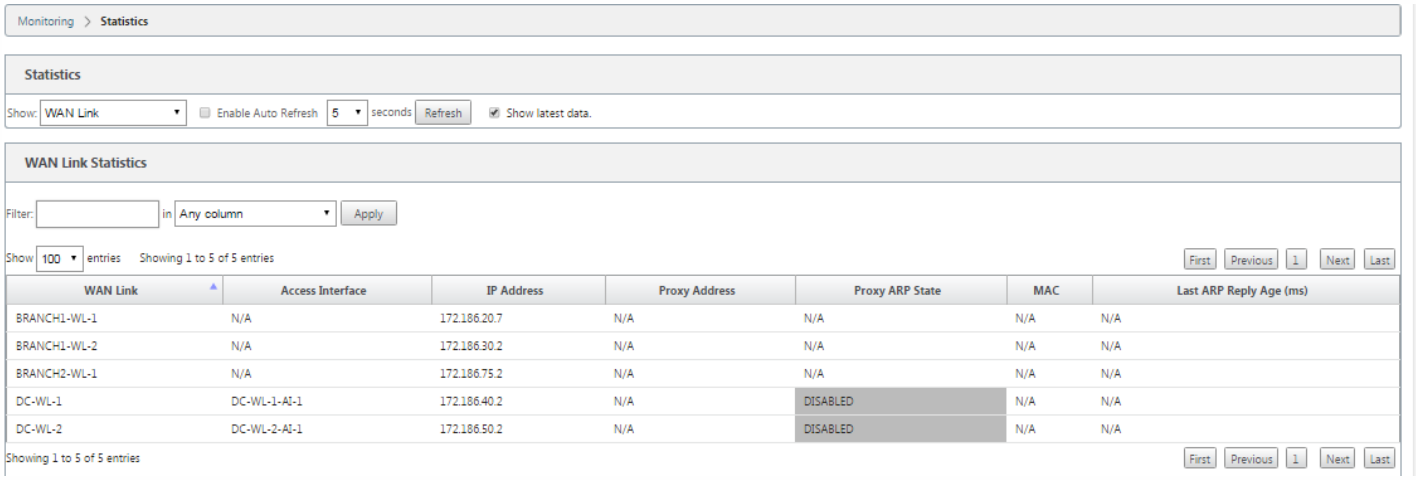

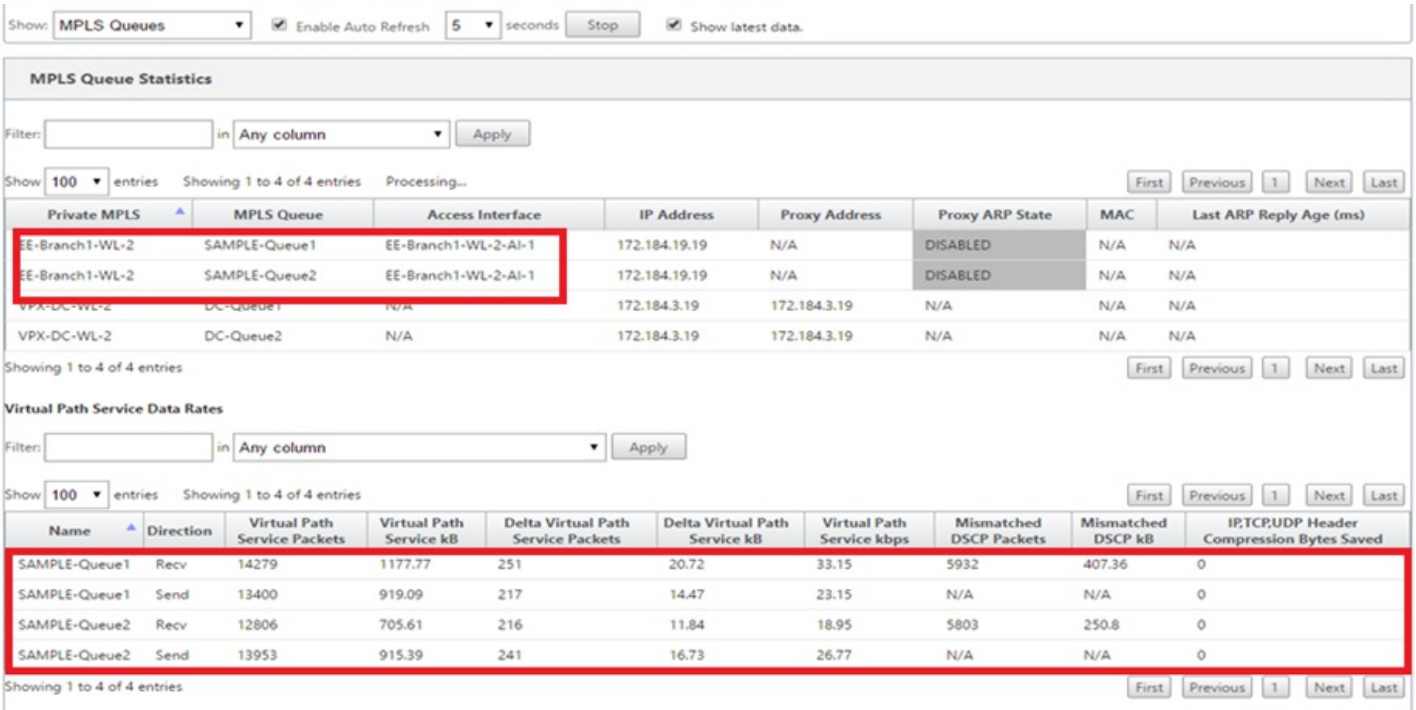

-
-

-
-
-
-
-

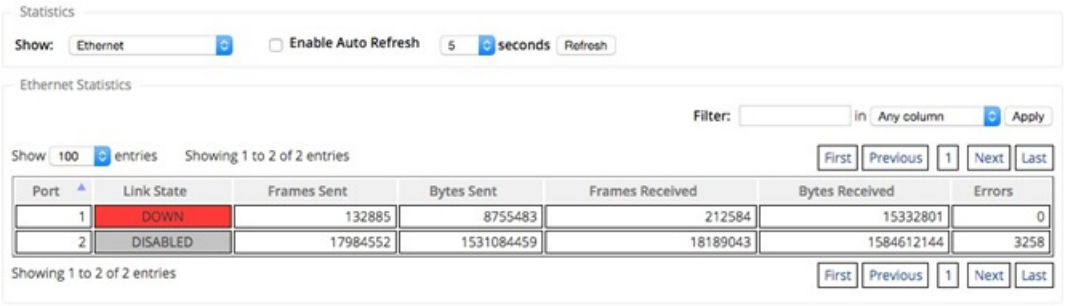

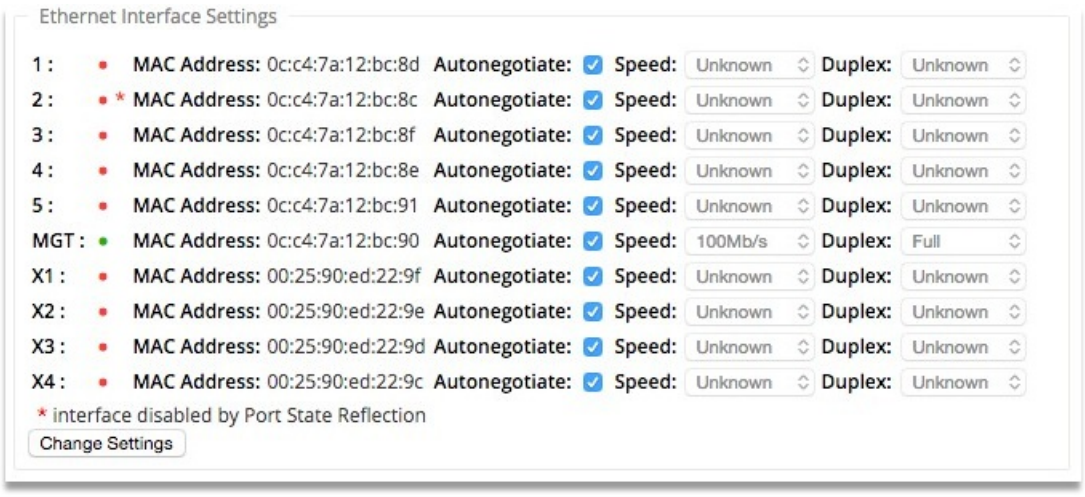

-
-
-
-

-
  -
  -
-
-
-
-
-
-
-
-

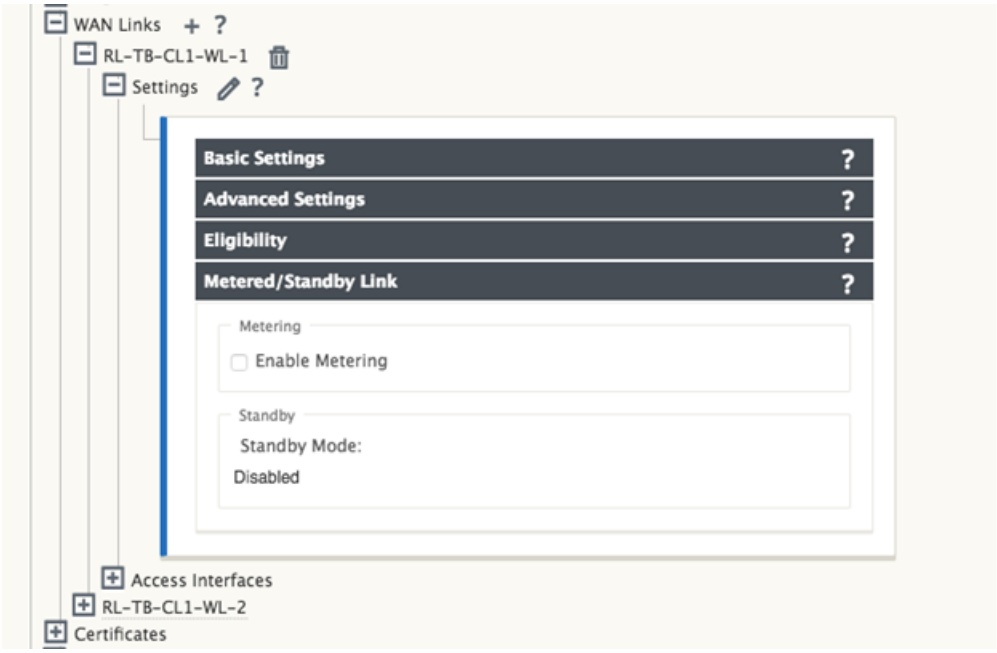

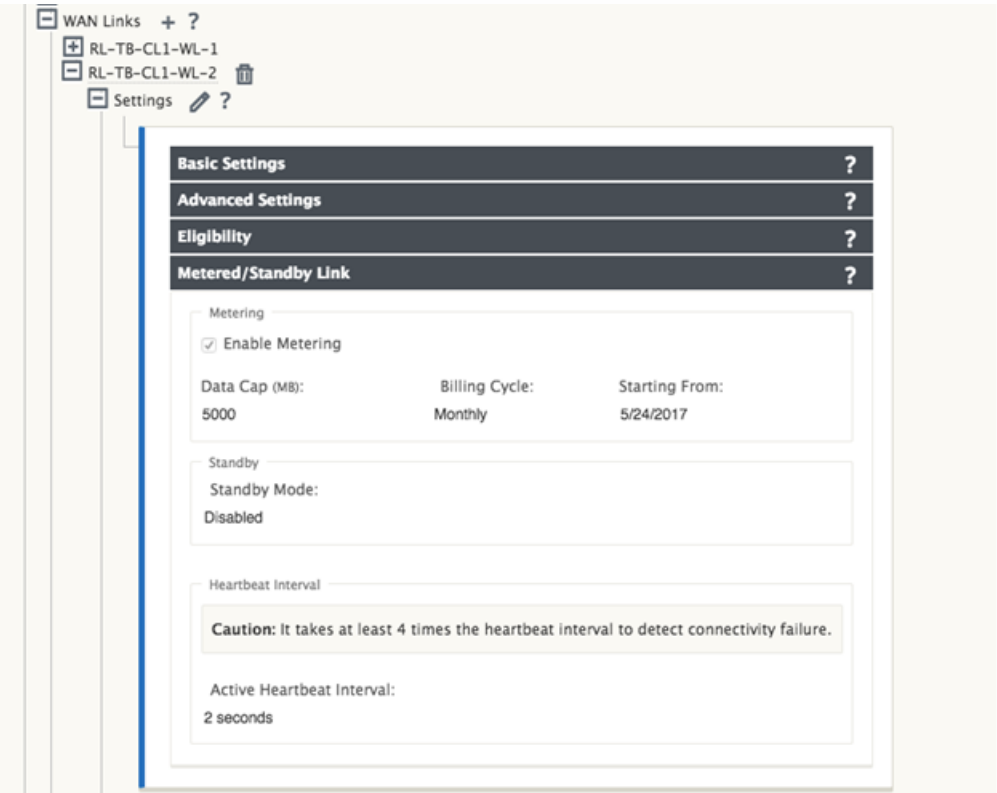

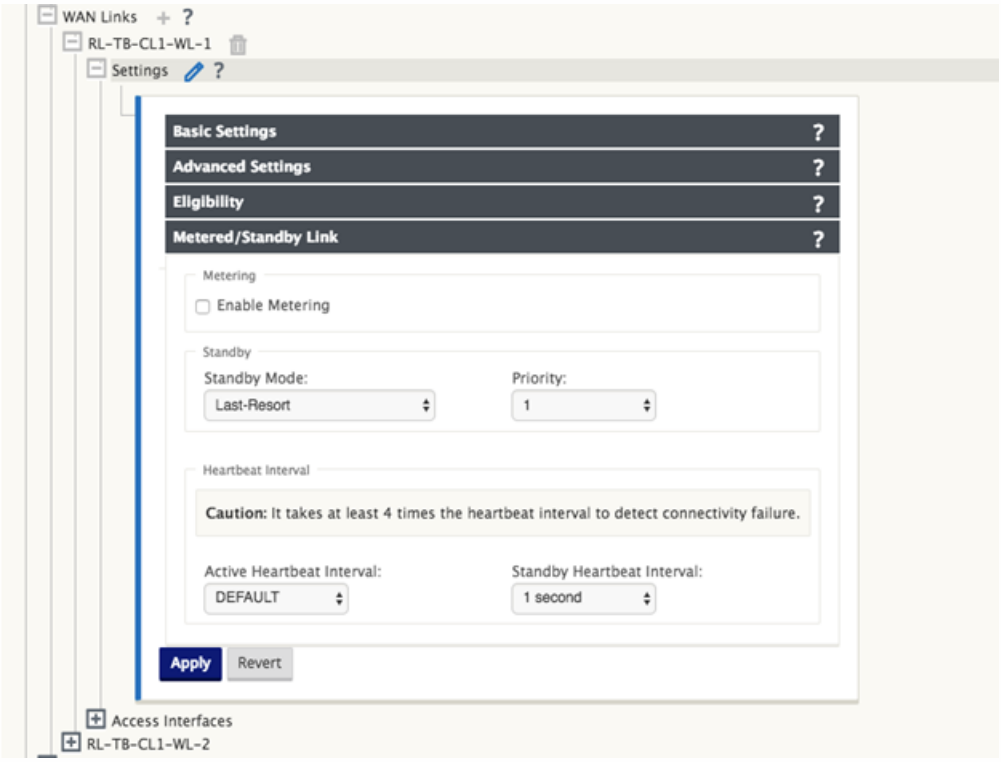

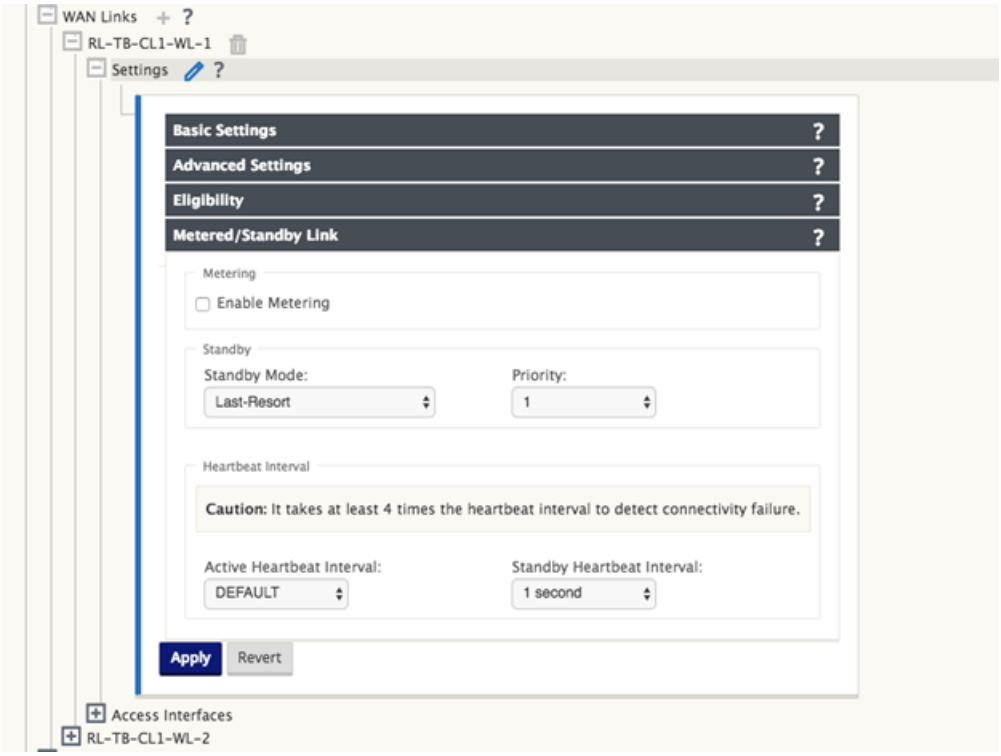

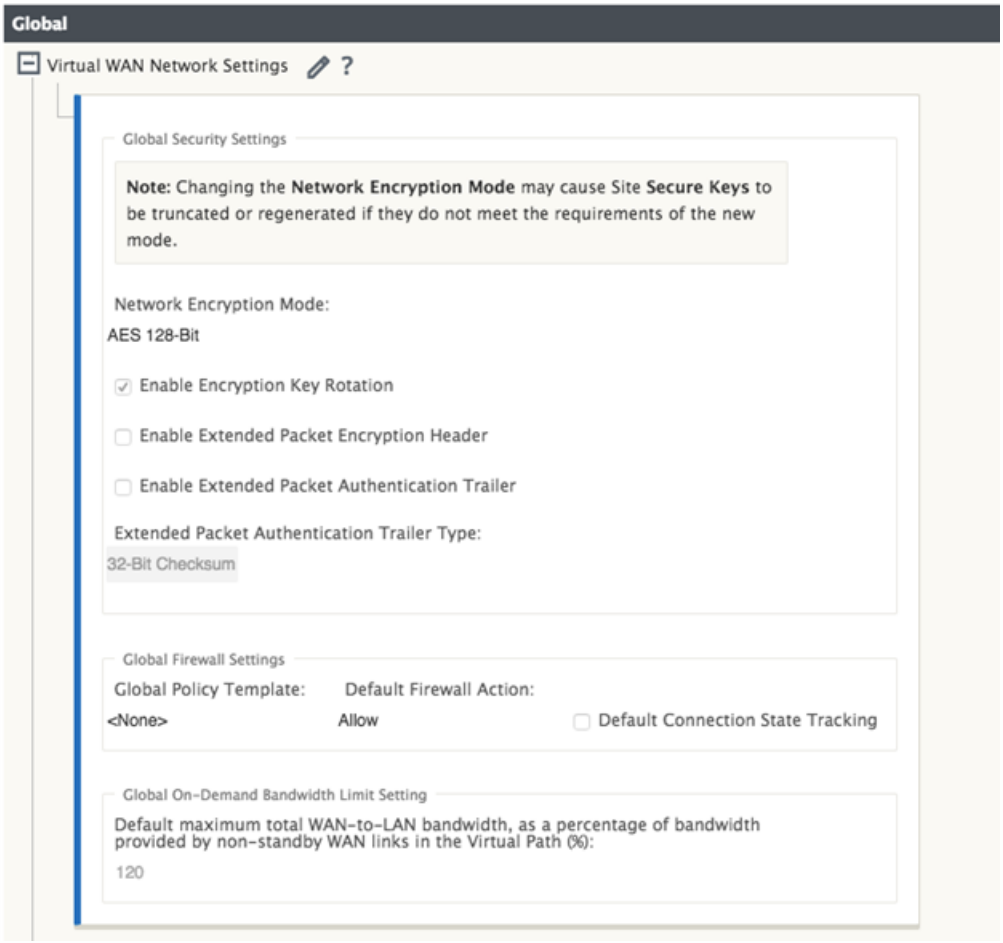

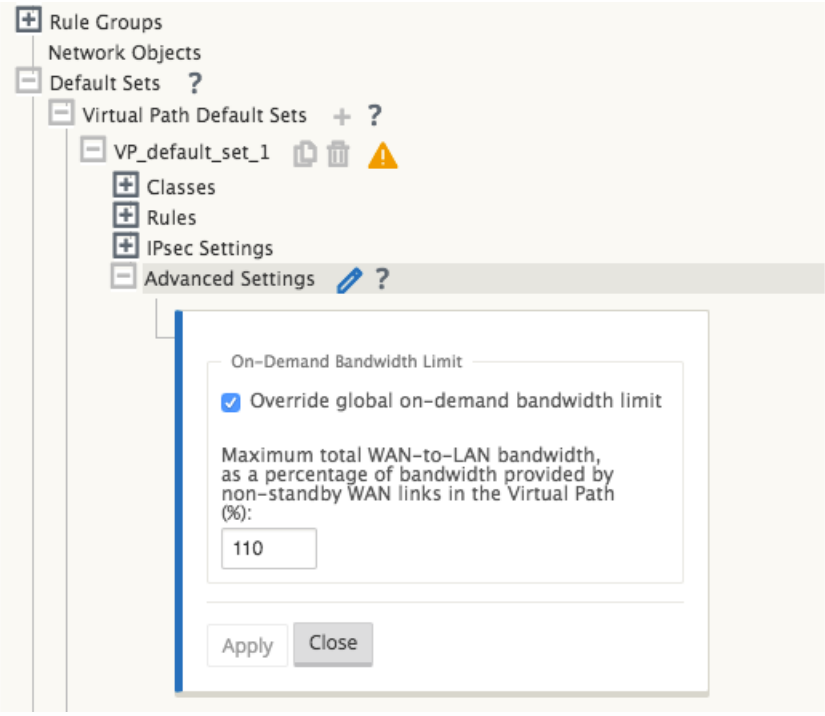

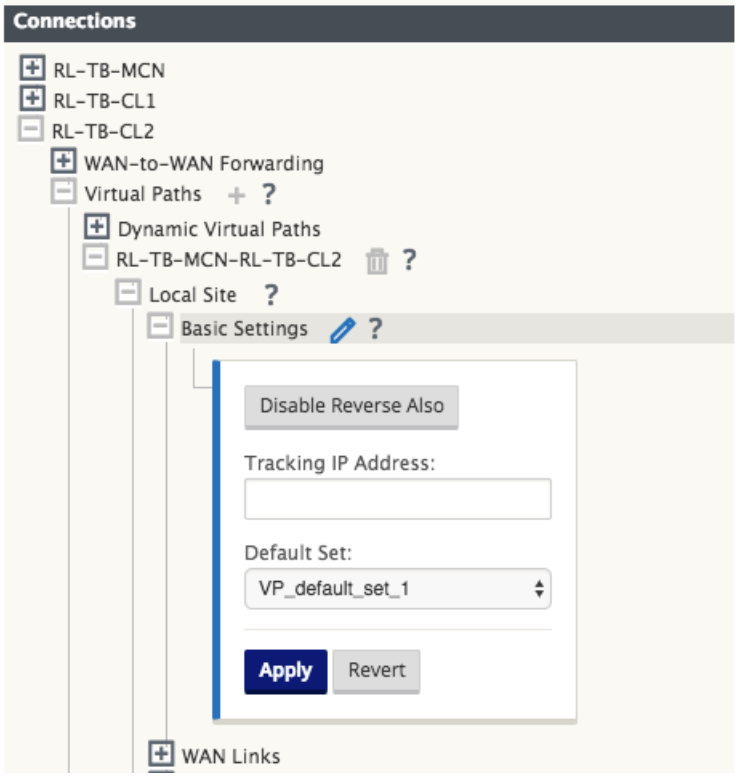

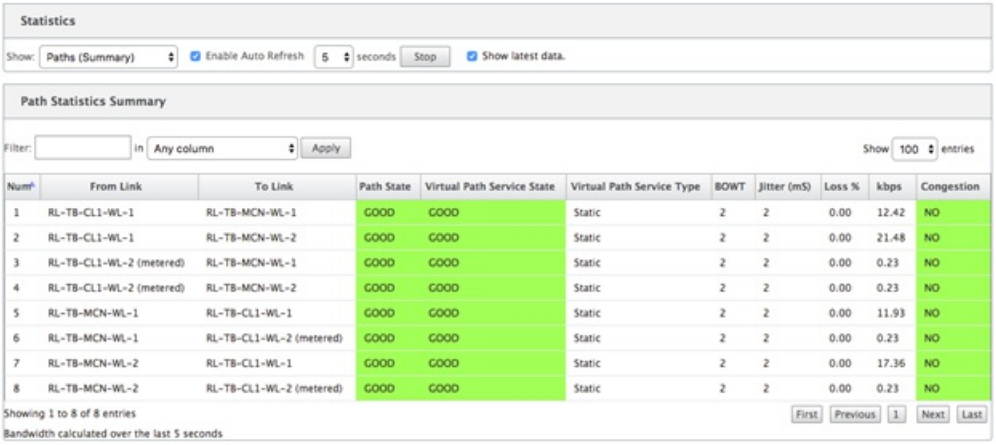

Statistics

Show: Paths (Summary) Enable Auto Refresh 5 seconds Stop Show latest data.

Path Statistics Summary

Filter: in Any column Apply Show 100 entries

Num	From Link	To Link	Path State	Virtual Path Service State	Virtual Path Service Type	BOWT	Jitter (mS)	Loss %	kbps	Congestion
1	RL-TB-CL1-WL-1	RL-TB-MCN-WL-1	GOOD	GOOD	Static	2	2	0.00	12.42	NO
2	RL-TB-CL1-WL-1	RL-TB-MCN-WL-2	GOOD	GOOD	Static	2	2	0.00	21.48	NO
3	RL-TB-CL1-WL-2 (metered)	RL-TB-MCN-WL-1	GOOD	GOOD	Static	2	2	0.00	0.23	NO
4	RL-TB-CL1-WL-2 (metered)	RL-TB-MCN-WL-2	GOOD	GOOD	Static	2	2	0.00	0.23	NO
5	RL-TB-MCN-WL-1	RL-TB-CL1-WL-1	GOOD	GOOD	Static	2	2	0.00	11.93	NO
6	RL-TB-MCN-WL-1	RL-TB-CL1-WL-2 (metered)	GOOD	GOOD	Static	2	2	0.00	0.23	NO
7	RL-TB-MCN-WL-2	RL-TB-CL1-WL-1	GOOD	GOOD	Static	2	2	0.00	17.36	NO
8	RL-TB-MCN-WL-2	RL-TB-CL1-WL-2 (metered)	GOOD	GOOD	Static	2	2	0.00	0.23	NO

Showing 1 to 8 of 8 entries First Previous 1 Next Last

Bandwidth calculated over the last 5 seconds

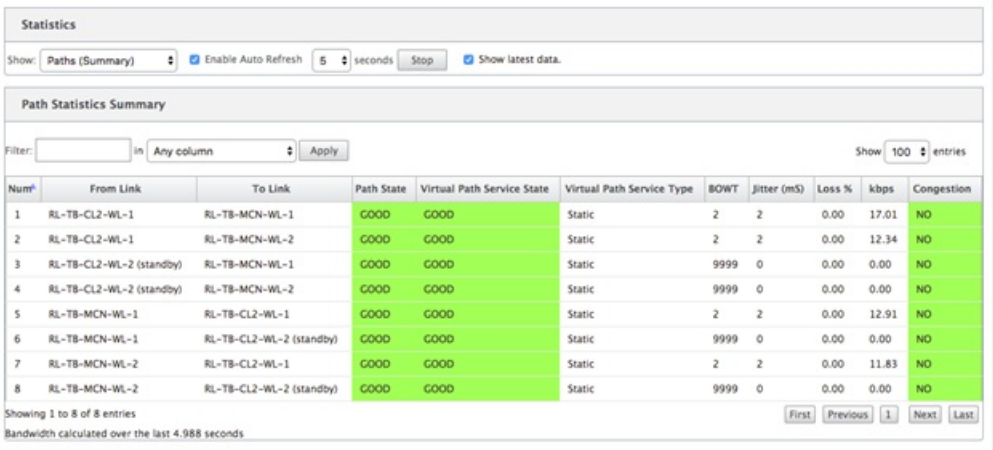

Statistics

Show: Paths (Summary) Enable Auto Refresh 5 seconds Stop Show latest data.

Path Statistics Summary

Filter: in Any column Apply Show 100 entries

Num	From Link	To Link	Path State	Virtual Path Service State	Virtual Path Service Type	BOWT	Jitter (mS)	Loss %	kbps	Congestion
1	RL-TB-CL2-WL-1	RL-TB-MCN-WL-1	GOOD	GOOD	Static	2	2	0.00	17.01	NO
2	RL-TB-CL2-WL-1	RL-TB-MCN-WL-2	GOOD	GOOD	Static	2	2	0.00	12.34	NO
3	RL-TB-CL2-WL-2 (standby)	RL-TB-MCN-WL-1	GOOD	GOOD	Static	9999	0	0.00	0.00	NO
4	RL-TB-CL2-WL-2 (standby)	RL-TB-MCN-WL-2	GOOD	GOOD	Static	9999	0	0.00	0.00	NO
5	RL-TB-MCN-WL-1	RL-TB-CL2-WL-1	GOOD	GOOD	Static	2	2	0.00	12.91	NO
6	RL-TB-MCN-WL-1	RL-TB-CL2-WL-2 (standby)	GOOD	GOOD	Static	9999	0	0.00	0.00	NO
7	RL-TB-MCN-WL-2	RL-TB-CL2-WL-1	GOOD	GOOD	Static	2	2	0.00	11.83	NO
8	RL-TB-MCN-WL-2	RL-TB-CL2-WL-2 (standby)	GOOD	GOOD	Static	9999	0	0.00	0.00	NO

Showing 1 to 8 of 8 entries First Previous 1 Next Last

Bandwidth calculated over the last 4.988 seconds

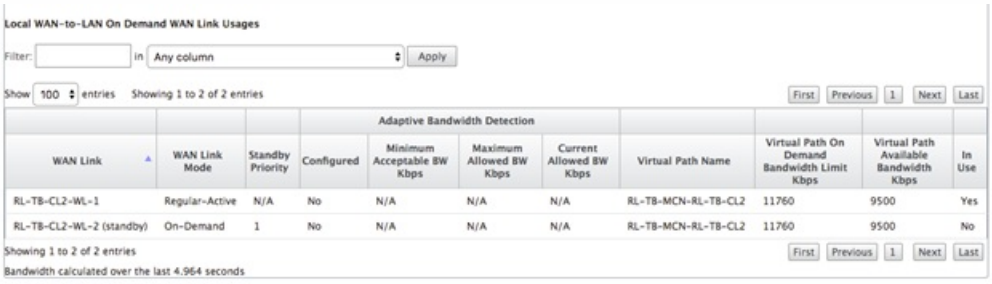

Local WAN-to-LAN On Demand WAN Link Usages

Filter: in Any column Apply

Show 100 entries Showing 1 to 2 of 2 entries First Previous 1 Next Last

WAN Link	WAN Link Mode	Standby Priority	Adaptive Bandwidth Detection				Virtual Path Name	Virtual Path On Demand Bandwidth Limit Kbps	Virtual Path Available Bandwidth Kbps	In Use
			Configured	Minimum Acceptable BW Kbps	Maximum Allowed BW Kbps	Current Allowed BW Kbps				
RL-TB-CL2-WL-1	Regular-Active	N/A	No	N/A	N/A	N/A	RL-TB-MCN-RL-TB-CL2	11760	9500	Yes
RL-TB-CL2-WL-2 (standby)	On-Demand	1	No	N/A	N/A	N/A	RL-TB-MCN-RL-TB-CL2	11760	9500	No

Showing 1 to 2 of 2 entries First Previous 1 Next Last

Bandwidth calculated over the last 4.964 seconds

System Status

Name:	RL-TB-CL1
Model:	VPX
Appliance Mode:	Client
Serial Number:	c4ec4b39-04db-2633-5ed4-38b3ad3f52d2
Management IP Address:	10.200.32.236
Appliance Uptime:	2 weeks, 3 days, 20 hours, 7 minutes, 22.6 seconds
Service Uptime:	2 hours, 1 minutes, 52.0 seconds
Routing Domain Enabled:	Default_RoutingDomain

Local Versions

Configuration Created On:	Fri May 26 11:42:15 2017
Software Version:	9.3.0.43.594998
Built On:	May 26 2017 at 03:42:20
Hardware Version:	VPX
OS Partition Version:	4.6

Virtual Path Service Status

Virtual Path RL-TB-MCN-RL-TB-CL1 Uptime: 2 hours, 55.0 seconds.

WAN Link Metering

WAN Link Name:	RL-TB-CL1-WL-2
Data Usage:	6921.72 MBs of 5000 MBs
Usage(in %):	138
Billing Cycle:	MONTHLY
Starting From:	05/20/2017
Days Elapsed:	7 days of 31 days

17654	1	RL-TB-CL1-WL-2	WAN LINK	2017-05-24 10:22:58	USAGE_3	WARNING	Total usage 1.84 GBytes used (91% of limit 2.00 GBytes) in 1 of 31 days in this billing cycle since 00:00:00 05/24/2017
17653	1	RL-TB-CL1-WL-2	WAN LINK	2017-05-24 10:17:58	USAGE_2	WARNING	Total usage 1.52 GBytes used (75% of limit 2.00 GBytes) in 1 of 31 days in this billing cycle since 00:00:00 05/24/2017
17652	1	RL-TB-CL1-WL-2	WAN LINK	2017-05-24 10:09:58	USAGE_1	WARNING	Total usage 1.00 GBytes used (50% of limit 2.00 GBytes) in 1 of 31 days in this billing cycle since 00:00:00 05/24/2017

24640	3	RL-TB-MCN-WL-2->RL-TB-CL2-WL-2	PATH	2017-05-26 10:18:32	STANDBY	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Backup Path RL-TB-MCN-WL-2->RL-TB-CL2-WL-2 has become standby
24639	1	RL-TB-MCN-WL-1->RL-TB-CL2-WL-2	PATH	2017-05-26 10:18:32	STANDBY	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Backup Path RL-TB-MCN-WL-1->RL-TB-CL2-WL-2 has become standby
24638	1	RL-TB-CL2-WL-1->RL-TB-MCN-WL-2	PATH	2017-05-26 10:18:27	GOOD	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Path RL-TB-CL2-WL-1->RL-TB-MCN-WL-2 state has changed from BAD to GOOD because notified by peer.
24637	2	RL-TB-MCN-WL-2->RL-TB-CL2-WL-1	PATH	2017-05-26 10:18:27	GOOD	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Path RL-TB-MCN-WL-2->RL-TB-CL2-WL-1 state has changed from BAD to GOOD .
24636	2	RL-TB-MCN:RL-TB-CL2	VIRTUAL PATH	2017-05-26 10:18:27	GOOD	NOTICE	The state of Virtual Path RL-TB-MCN-RL-TB-CL2 has changed from BAD to GOOD
24635	0	RL-TB-CL2-WL-1->RL-TB-MCN-WL-1	PATH	2017-05-26 10:18:27	GOOD	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Path RL-TB-CL2-WL-1->RL-TB-MCN-WL-1 state has changed from BAD to GOOD because notified by peer.
24634	0	RL-TB-MCN-WL-1->RL-TB-CL2-WL-1	PATH	2017-05-26 10:18:27	GOOD	NOTICE	Virtual Path RL-TB-MCN-RL-TB-CL2 Path RL-TB-MCN-WL-1->RL-TB-CL2-WL-1 state has changed from BAD to GOOD .
24633	3	RL-TB-MCN-WL-2->RL-TB-CL2-WL-2	PATH	2017-05-26 10:18:27	ACTIVE	ERROR	Virtual Path RL-TB-MCN-RL-TB-CL2 Backup Path RL-TB-MCN-WL-2->RL-TB-CL2-WL-2 has become active
24632	1	RL-TB-MCN-WL-1->RL-TB-CL2-WL-2	PATH	2017-05-26 10:18:27	ACTIVE	ERROR	Virtual Path RL-TB-MCN-RL-TB-CL2 Backup Path RL-TB-MCN-WL-1->RL-TB-CL2-WL-2 has become active

Path Configuration

Paths on virtual path 1 'RL-TB-MCN-RL-TB-CL1':

Path ID	From Link	To Link	Primary Src IP Address	Primary Dst IP Address	Secondary Src IP Address	Secondary Dst IP Address	Src Port	Dst Port	Al Sr
0	RL-TB-MCN-WL-1	RL-TB-CL1-WL-1	172.16.124.2	172.16.152.2	-	-	4980	4980	
3	RL-TB-MCN-WL-2	RL-TB-CL1-WL-2	192.168.15.2	192.168.16.194	-	-	4980	4980	
1	RL-TB-MCN-WL-1	RL-TB-CL1-WL-2	172.16.124.2	192.168.16.194	-	-	4980	4980	
2	RL-TB-MCN-WL-2	RL-TB-CL1-WL-1	192.168.15.2	172.16.152.2	-	-	4980	4980	
0	RL-TB-CL1-WL-1	RL-TB-MCN-WL-1	172.16.152.2	172.16.124.2	-	-	4980	4980	
3	RL-TB-CL1-WL-2	RL-TB-MCN-WL-2	192.168.16.194	192.168.15.2	-	-	4980	4980	
1	RL-TB-CL1-WL-1	RL-TB-MCN-WL-2	172.16.152.2	192.168.15.2	-	-	4980	4980	
2	RL-TB-CL1-WL-2	RL-TB-MCN-WL-1	192.168.16.194	172.16.124.2	-	-	4980	4980	

From Link	To Link	Realtime Eligible	Interactive Eligible	Bulk Eligible	Path Group	Standby Heartbeat Interval(ms)	Active Heartbeat Interval(ms)
RL-TB-MCN-WL-1	RL-TB-CL1-WL-1	YES	YES	YES	0	n/a	n/a
RL-TB-MCN-WL-2	RL-TB-CL1-WL-2	YES	YES	YES	0	n/a	2000
RL-TB-MCN-WL-1	RL-TB-CL1-WL-2	YES	YES	YES	0	n/a	2000
RL-TB-MCN-WL-2	RL-TB-CL1-WL-1	YES	YES	YES	0	n/a	n/a
RL-TB-CL1-WL-1	RL-TB-MCN-WL-1	YES	YES	YES	0	n/a	n/a
RL-TB-CL1-WL-2	RL-TB-MCN-WL-2	YES	YES	YES	0	n/a	2000
RL-TB-CL1-WL-1	RL-TB-MCN-WL-2	YES	YES	YES	0	n/a	n/a
RL-TB-CL1-WL-2	RL-TB-MCN-WL-1	YES	YES	YES	0	n/a	2000

Path Configuration

Paths on virtual path 2 'RL-TB-MCN-RL-TB-CL2':

Path ID	From Link	To Link	Primary Src IP Address	Primary Dst IP Address	Secondary Src IP Address	Secondary Dst IP Address	Src Port	Dst Port	Alternate Src Port	Alternate Dst Port	IP DSCP
0	RL-TB-CL2-WL-1	RL-TB-MCN-WL-1	172.16.156.2	172.16.124.2	-	-	4980	4980	-	-	*
3	RL-TB-CL2-WL-2	RL-TB-MCN-WL-2	192.168.17.2	192.168.15.2	-	-	4980	4980	-	-	*
1	RL-TB-CL2-WL-1	RL-TB-MCN-WL-2	172.16.156.2	192.168.15.2	-	-	4980	4980	-	-	*
2	RL-TB-CL2-WL-2	RL-TB-MCN-WL-1	192.168.17.2	172.16.124.2	-	-	4980	4980	-	-	*
0	RL-TB-MCN-WL-1	RL-TB-CL2-WL-1	172.16.124.2	172.16.156.2	-	-	4980	4980	-	-	*
3	RL-TB-MCN-WL-2	RL-TB-CL2-WL-2	192.168.15.2	192.168.17.2	-	-	4980	4980	-	-	*
1	RL-TB-MCN-WL-1	RL-TB-CL2-WL-2	172.16.124.2	192.168.17.2	-	-	4980	4980	-	-	*
2	RL-TB-MCN-WL-2	RL-TB-CL2-WL-1	192.168.15.2	172.16.156.2	-	-	4980	4980	-	-	*

From Link	To Link	Realtime Eligible	Interactive Eligible	Bulk Eligible	Path Group	Standby Heartbeat Interval(ms)	Active Heartbeat Interval(ms)
RL-TB-CL2-WL-1	RL-TB-MCN-WL-1	YES	YES	YES	0	n/a	n/a
RL-TB-CL2-WL-2	RL-TB-MCN-WL-2	YES	YES	YES	1	0	1000
RL-TB-CL2-WL-1	RL-TB-MCN-WL-2	YES	YES	YES	0	n/a	n/a
RL-TB-CL2-WL-2	RL-TB-MCN-WL-1	YES	YES	YES	1	0	1000
RL-TB-MCN-WL-1	RL-TB-CL2-WL-1	YES	YES	YES	0	n/a	n/a
RL-TB-MCN-WL-2	RL-TB-CL2-WL-2	YES	YES	YES	1	0	1000
RL-TB-MCN-WL-1	RL-TB-CL2-WL-2	YES	YES	YES	1	0	1000
RL-TB-MCN-WL-2	RL-TB-CL2-WL-1	YES	YES	YES	0	n/a	n/a

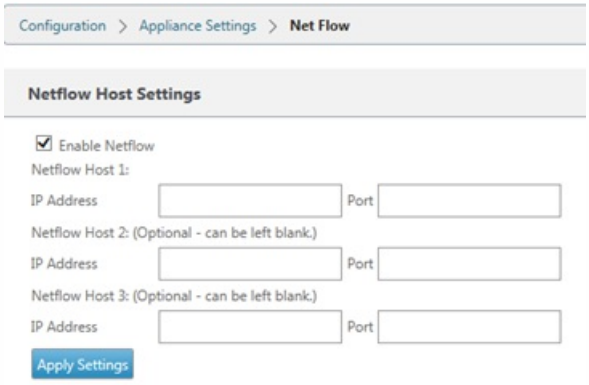

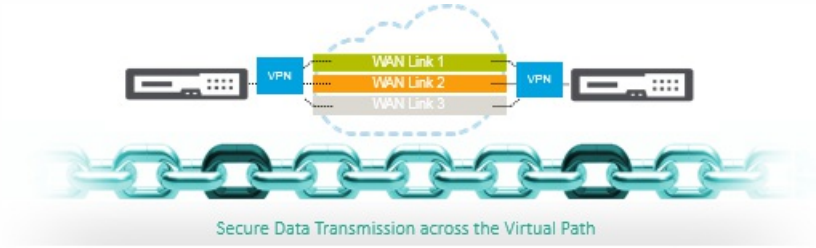

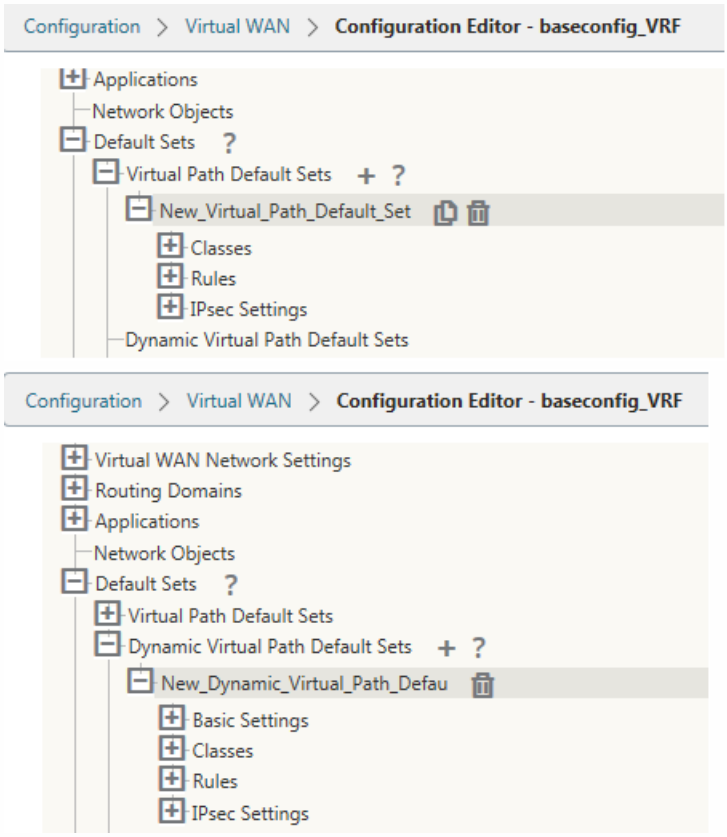

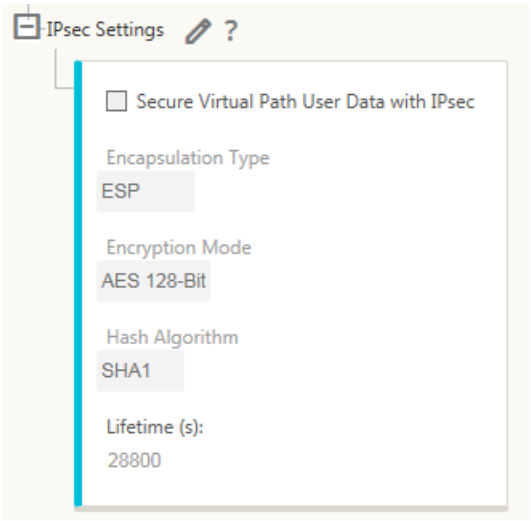

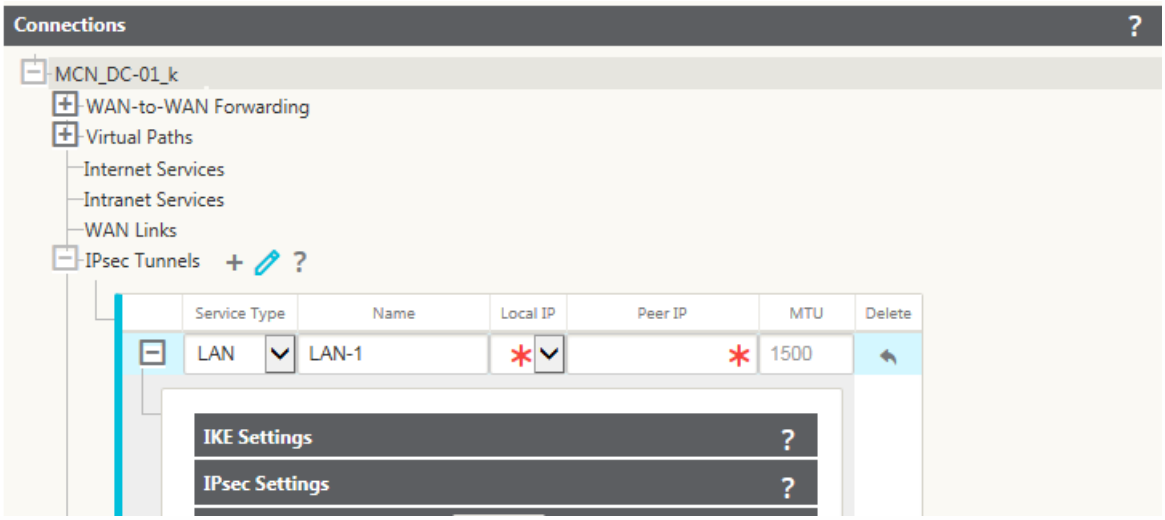

		• •
		•

		•
		•
		•
		• •
		• •
		• •
		• •
		•
		• • •
		• • •
		• • •
		•

		•
		• • •
		•
		• • •

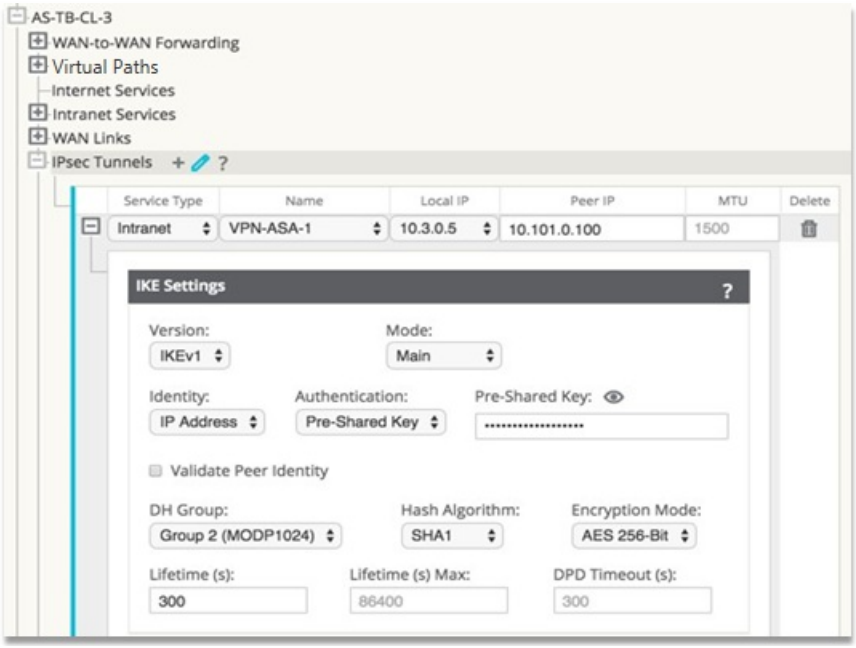

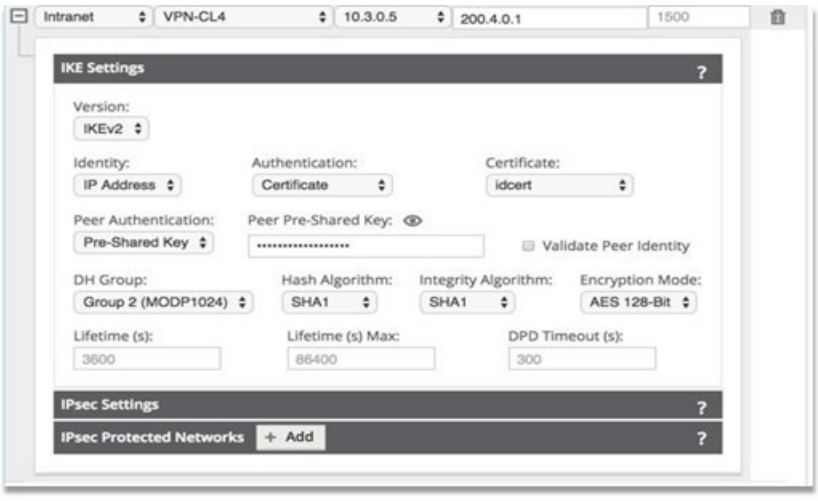

		• • • •
		• • •
		• • •
		•
		•
		•
		•
		• • •
		•

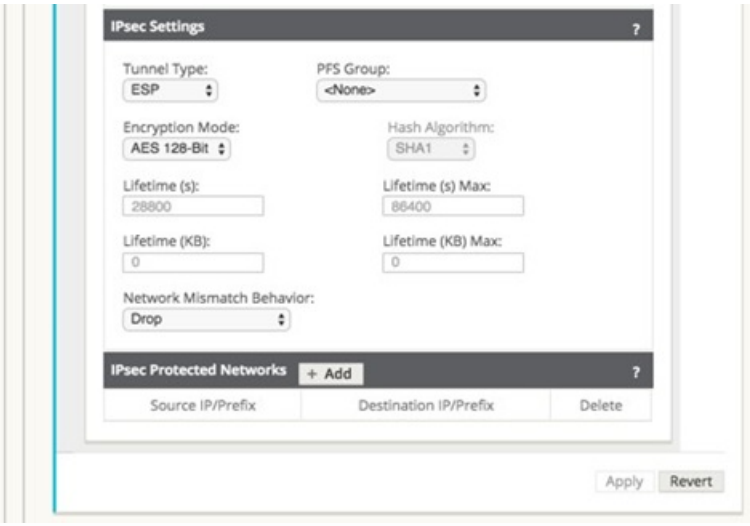

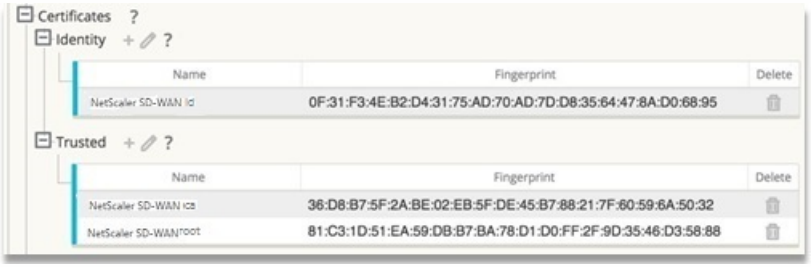

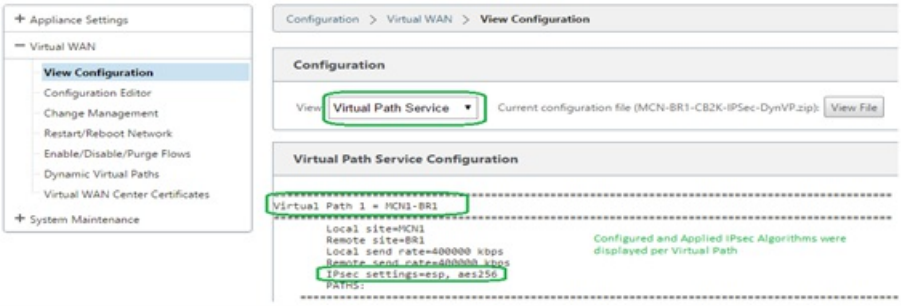

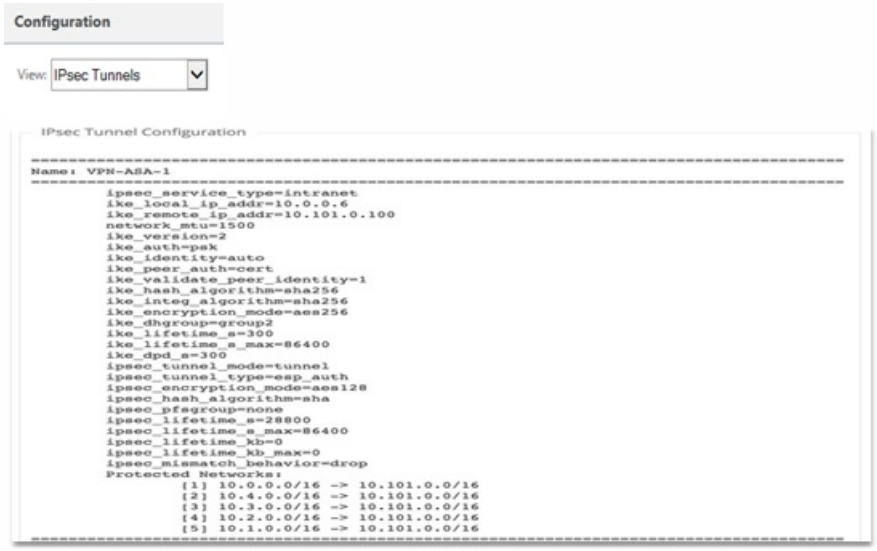

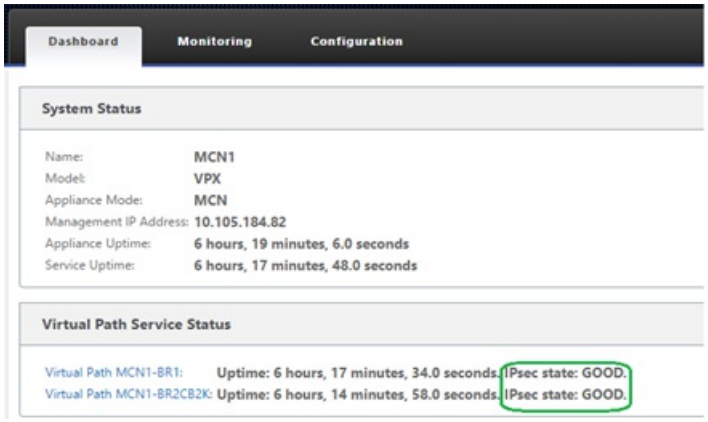

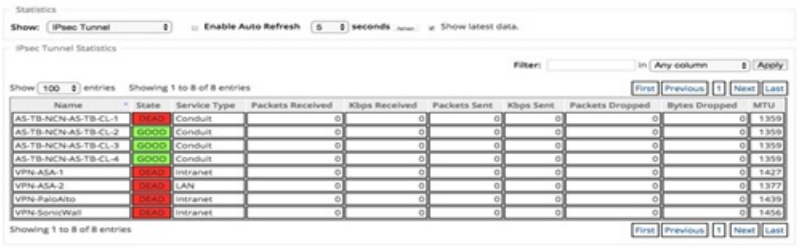

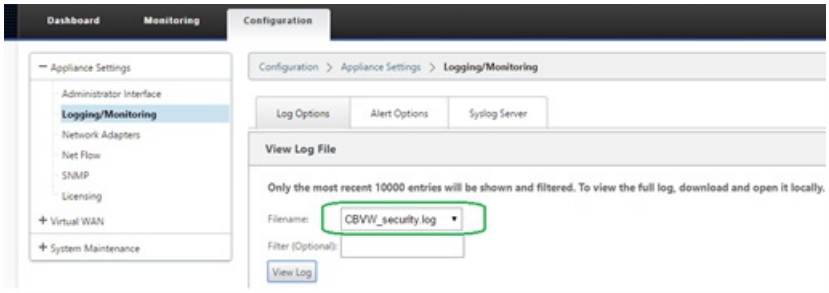

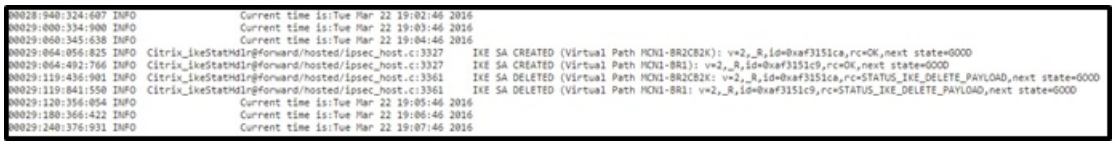

Monitoring > **Statistics**

Statistics

Show: Routes | Enable Auto Refresh | 5 seconds | Refresh | Clear Counters on Refresh | Purge dynamic routes

Route Statistics

Maximum allowed routes: 16000

Routes for routing domain : Default_RoutingDomain

Filter: in Any column Apply

Show 100 entries Showing 1 to 13 of 13 entries First Previous 1 Next Last

Num	Network Addr	Gateway IP Address	Service	Firewall Zone	Reachable	Site IP Address	Site	Type	Protocol	Neighbor Direct	Cost	Hit Count	Eligible	Eligibility Type	Eligibility Value
0	172.186.120.0/24	172.186.40.1	Local	Default_LAN_Zone	YES	*	DC	Static	-	-	5	11369	YES	N/A	N/A
1	172.186.50.0/24	*	Local	Default_LAN_Zone	YES	*	DC	Static	-	-	5	0	YES	N/A	N/A
2	172.186.40.0/24	*	Local	Default_LAN_Zone	YES	*	DC	Static	-	-	5	11389	YES	N/A	N/A
3	172.186.75.0/24	*	DC-BRANCH2	Default_LAN_Zone	YES	*	BRANCH2	Static	-	-	5	0	YES	N/A	N/A
4	172.186.30.0/24	*	DC-BRANCH1	Default_LAN_Zone	YES	*	BRANCH1	Static	-	-	5	0	YES	N/A	N/A
5	172.186.20.0/24	*	DC-BRANCH1	Default_LAN_Zone	YES	*	BRANCH1	Static	-	-	5	0	YES	N/A	N/A
6	172.186.160.0/24	172.186.40.1	Local	Default_LAN_Zone	YES	*	DC	Dynamic	BGP	-	6	0	YES	N/A	N/A
7	155.155.155.0/24	172.186.40.1	Local	Default_LAN_Zone	YES	*	DC	Dynamic	BGP	-	6	0	YES	N/A	N/A
8	172.186.30.0/24	*	New_Intranet_Service-1	Default_LAN_Zone	YES	*	DC	Static	-	-	15	0	YES	N/A	N/A
9	172.186.20.0/24	*	New_Intranet_Service-1	Default_LAN_Zone	YES	*	DC	Static	-	-	15	0	YES	N/A	N/A
10	16.16.0.0/16	172.186.40.1	Local	Default_LAN_Zone	YES	*	DC	Dynamic	BGP	-	6	0	YES	N/A	N/A
11	0.0.0.0/0	*	Passthrough	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A
12	0.0.0.0/0	*	Discard	Any	YES	*	*	Static	-	-	16	0	YES	N/A	N/A

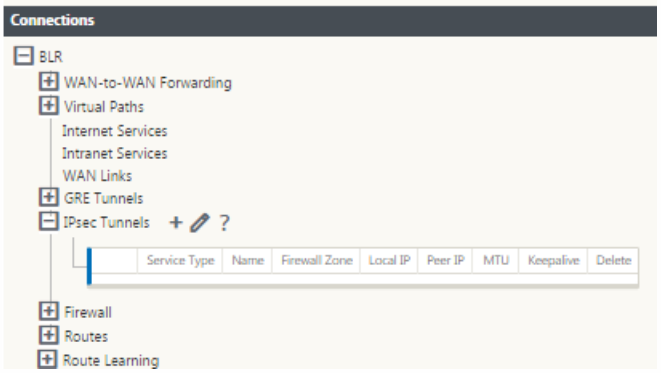

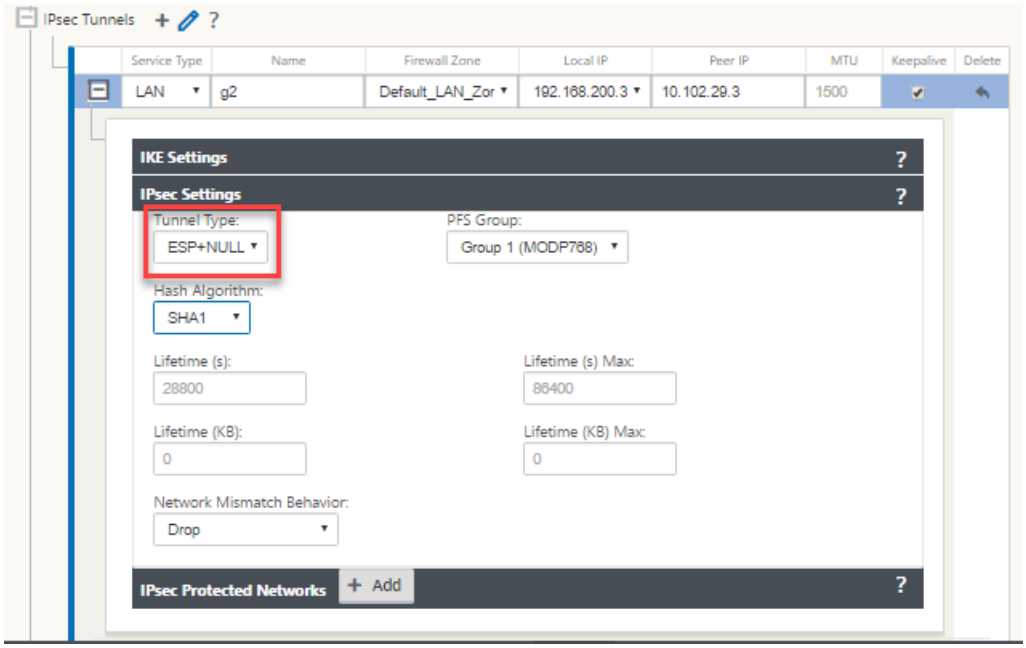

-
-
-
-
-
-
-
-
-
-
-

-
-
-
-

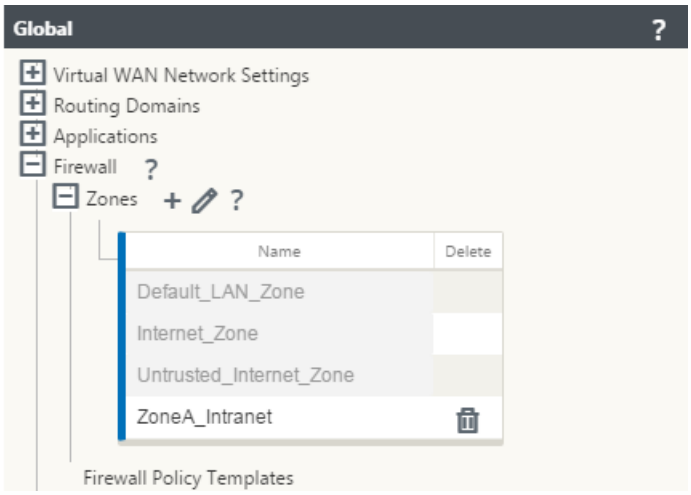

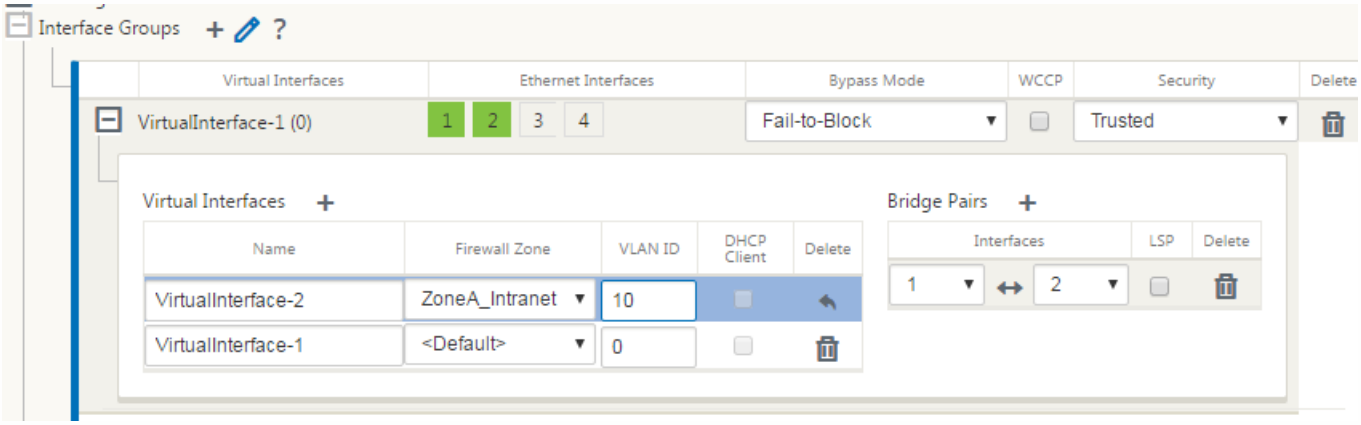

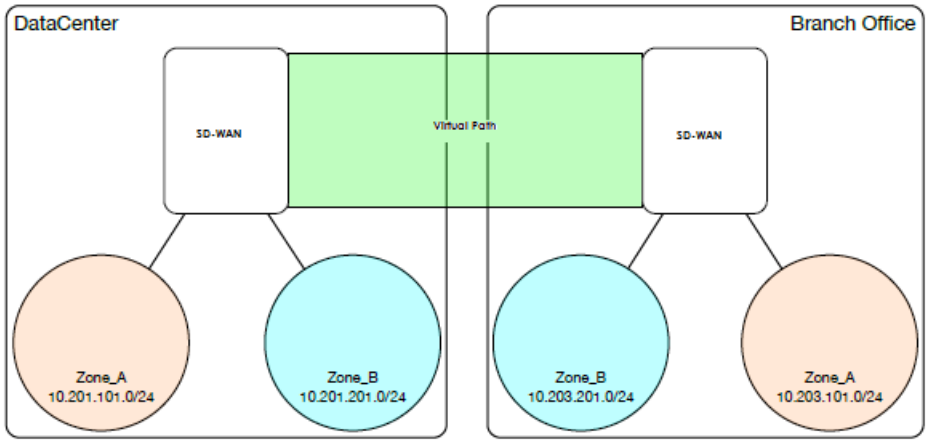

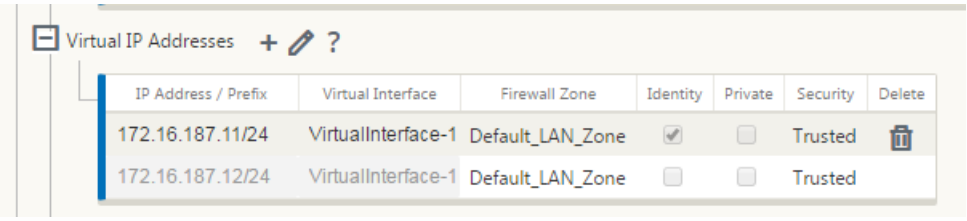

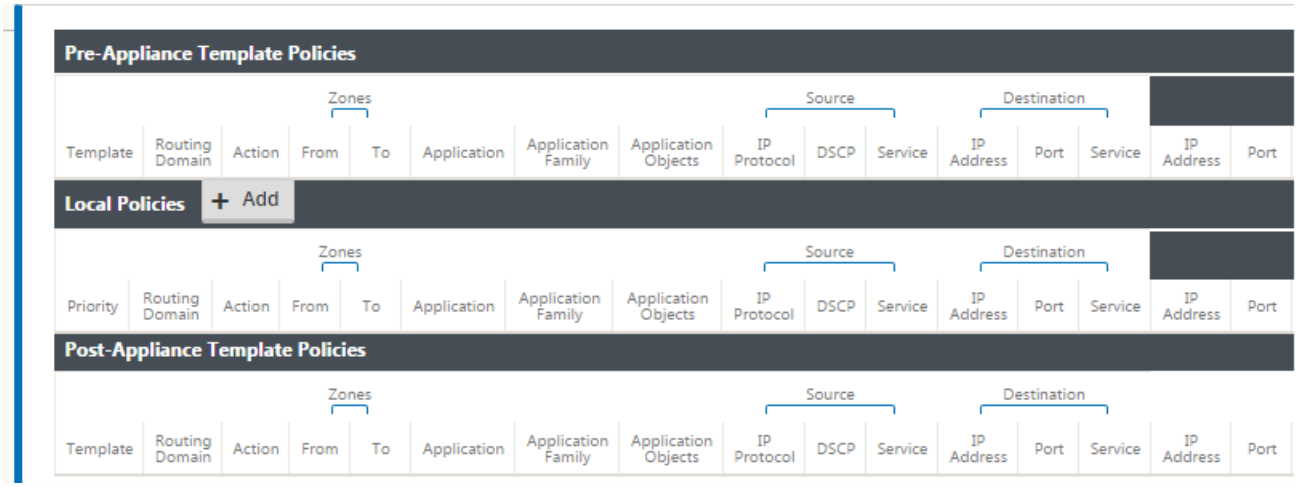

Filter Policy Evaluation Order

Policy definitions - Global and Local (site)

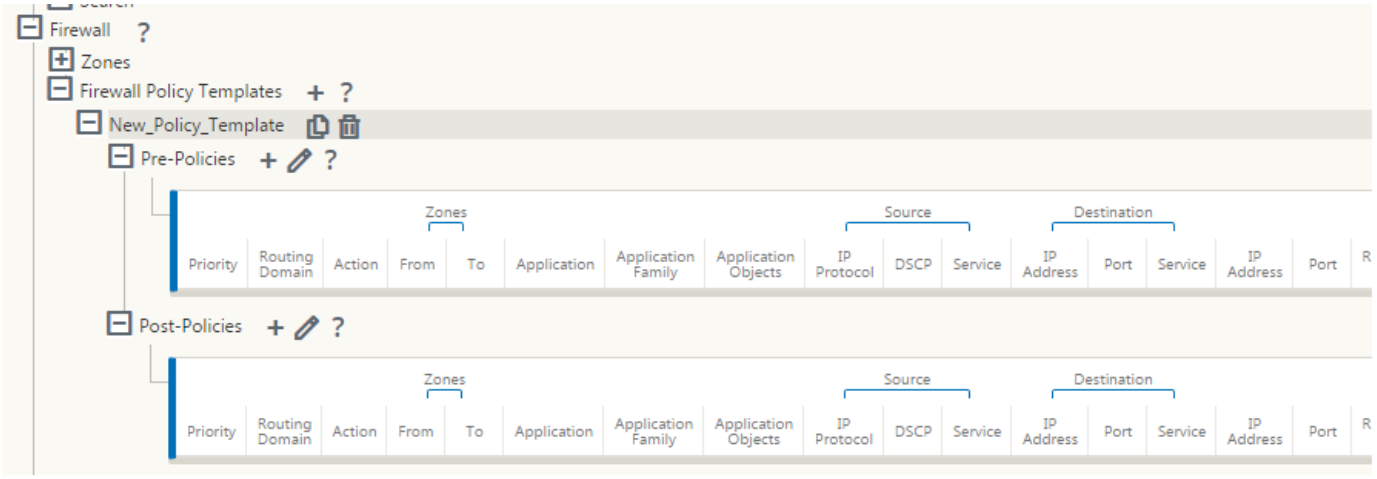

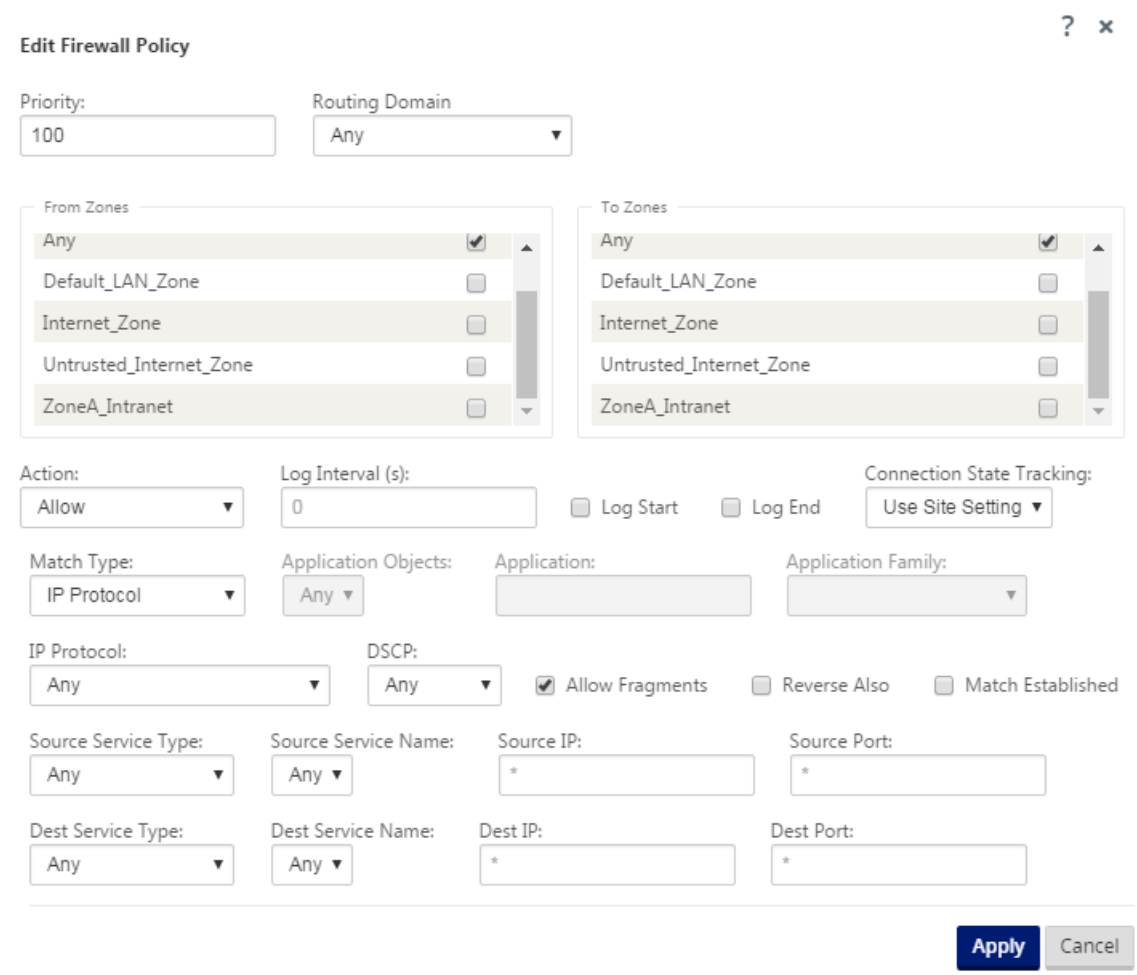

Policy Attributes

State Table for The Track Option

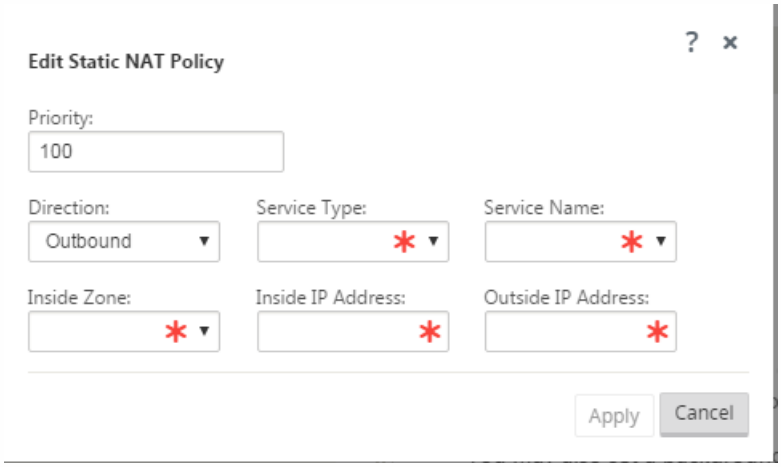

Configuration Options

-
-

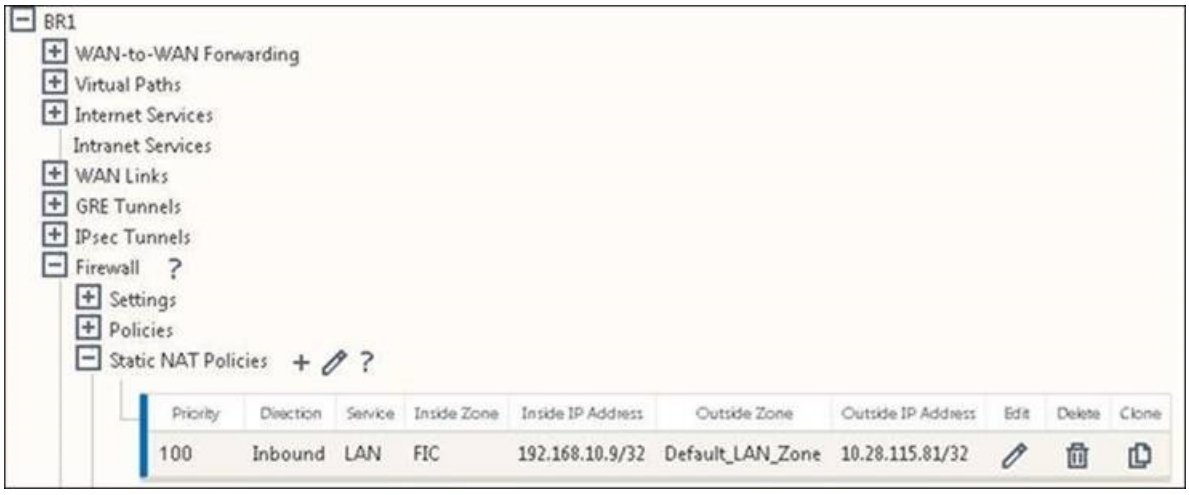

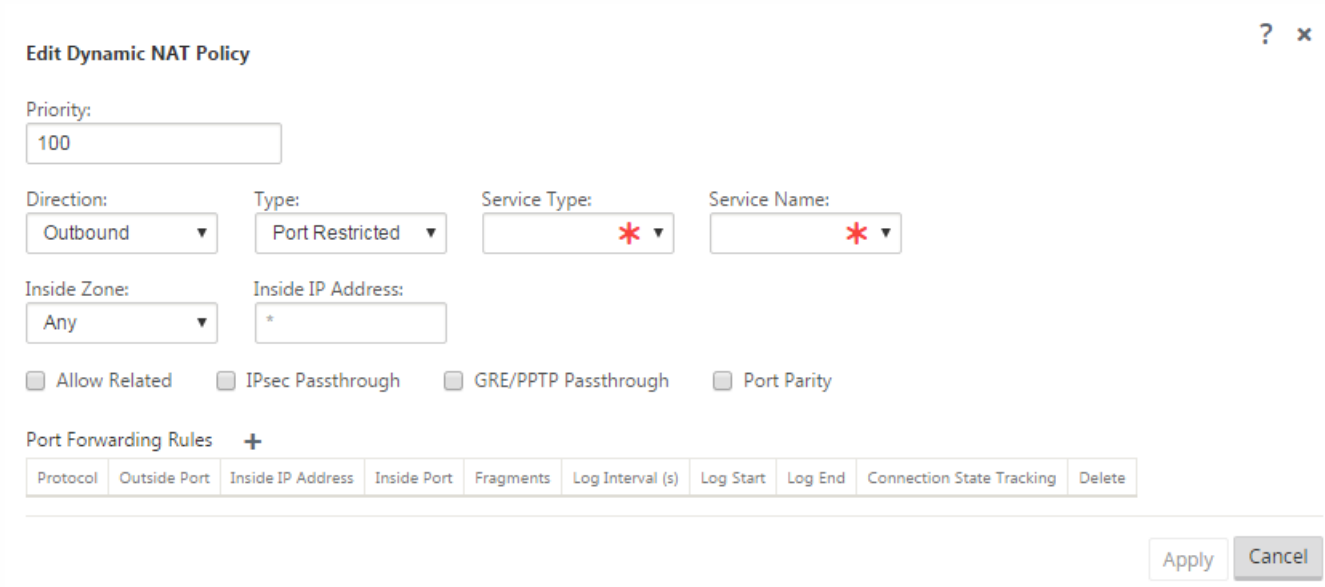

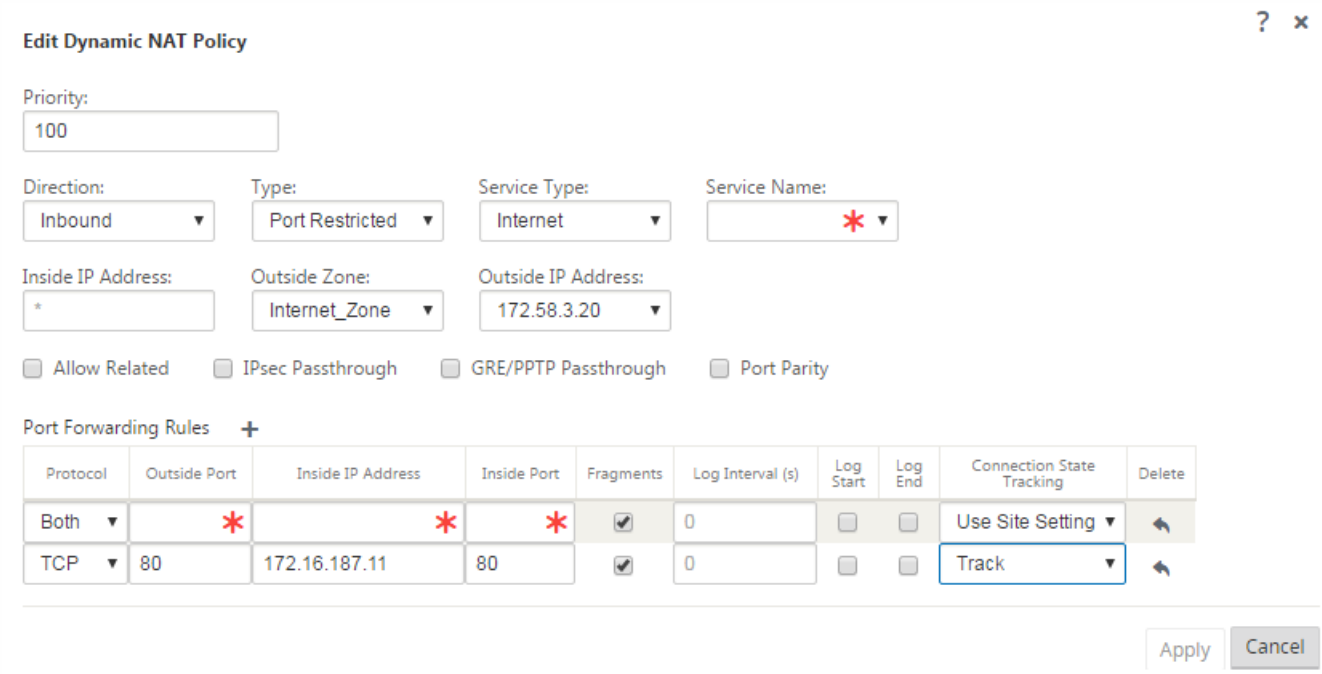

Configuration Options

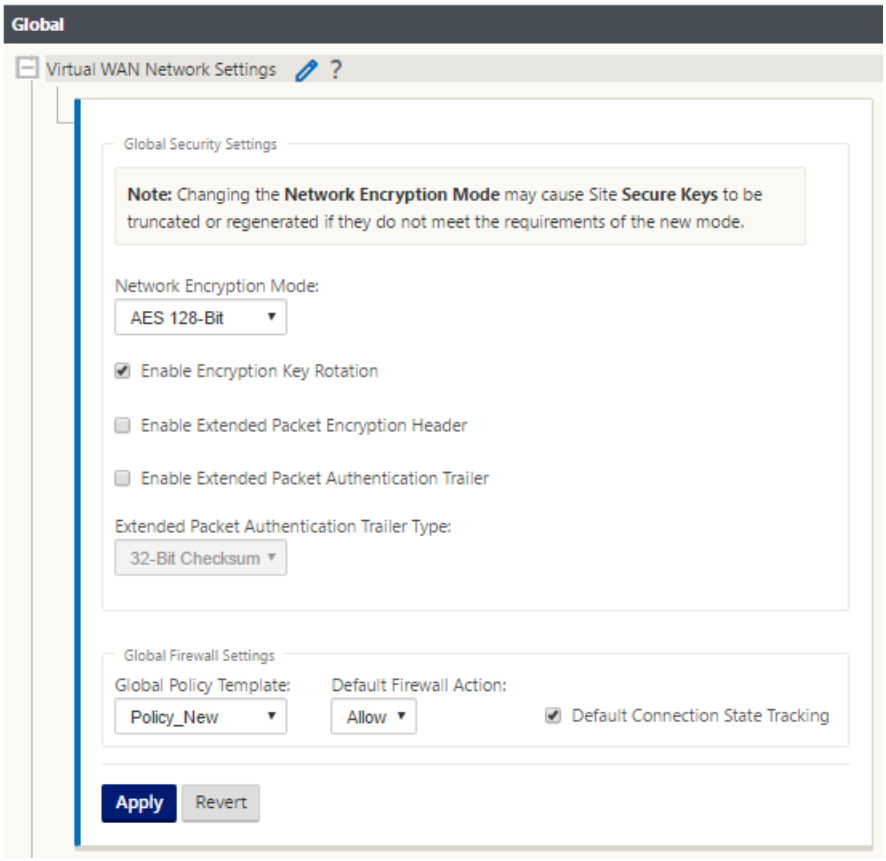

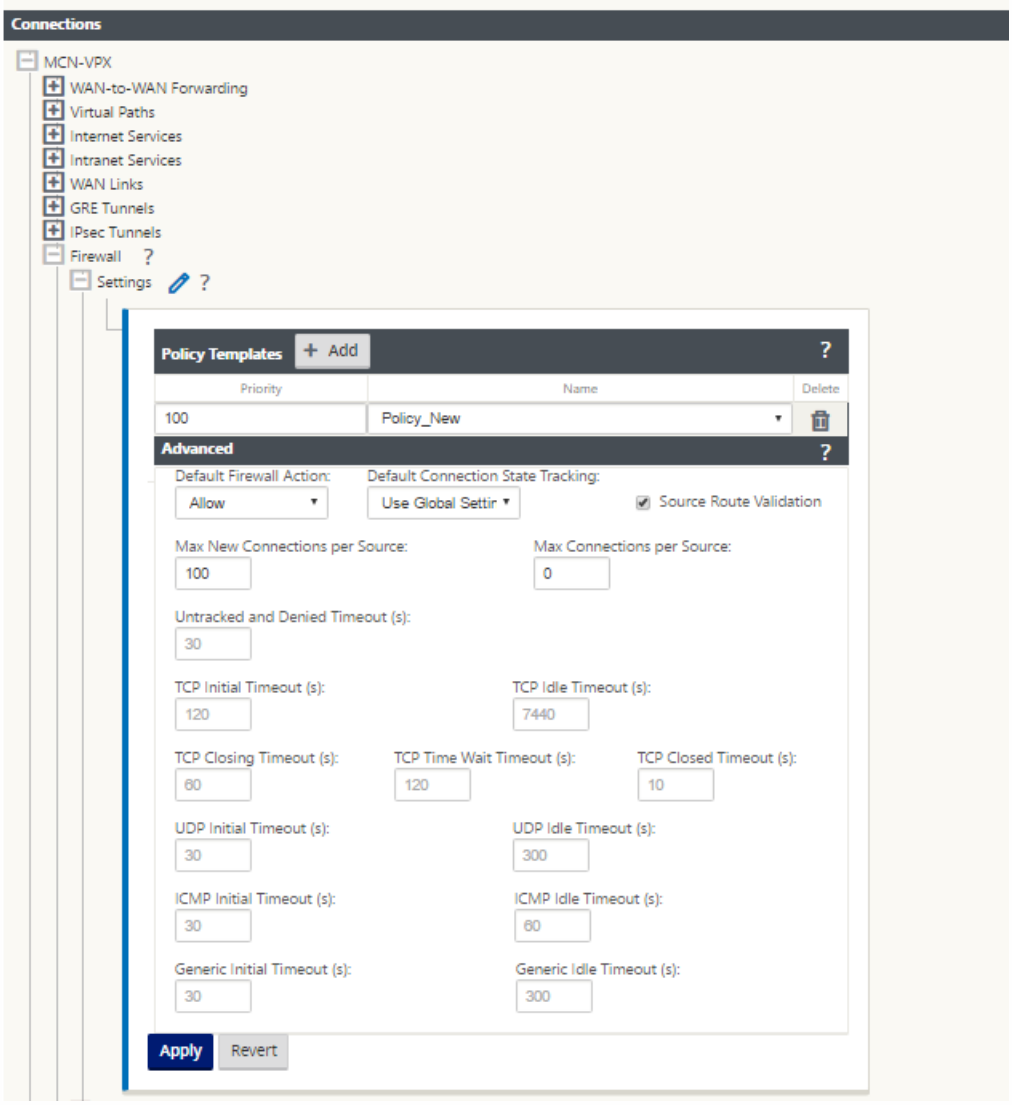

How to Configure Multicast Groups in SD-WAN GUI

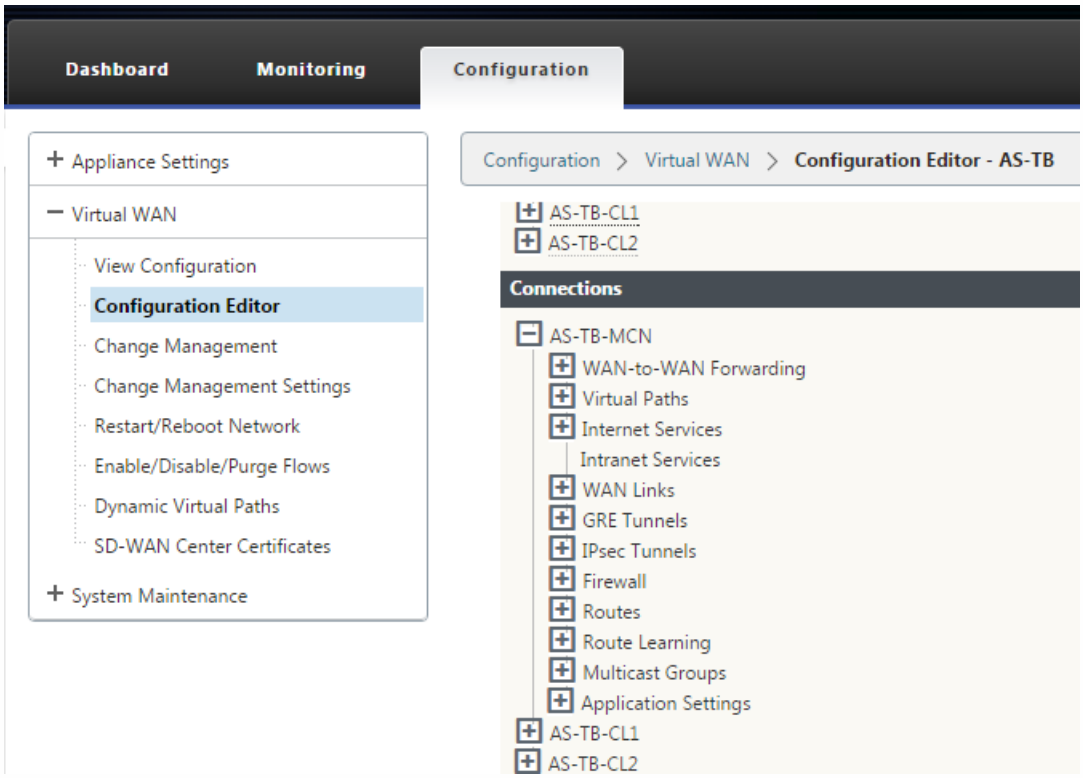

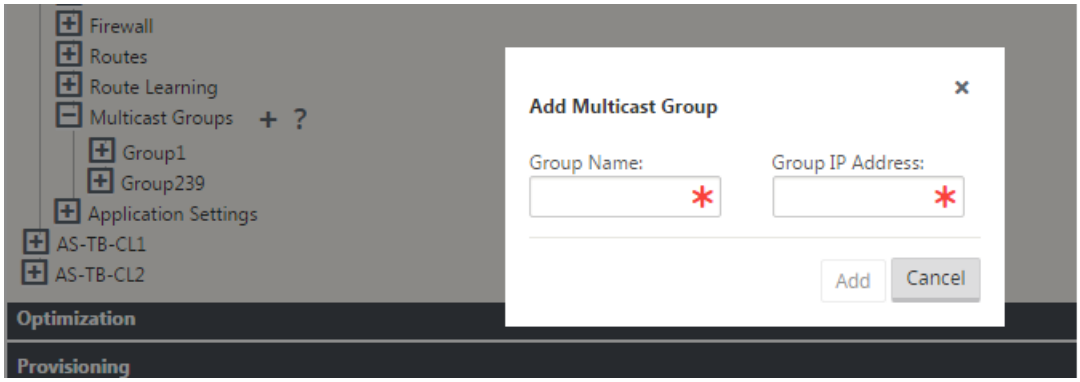

-
-
-

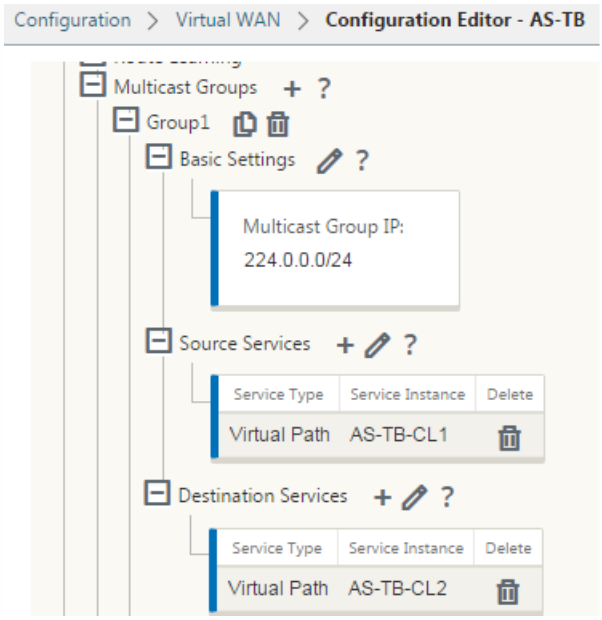

Monitoring Multicast Traffic

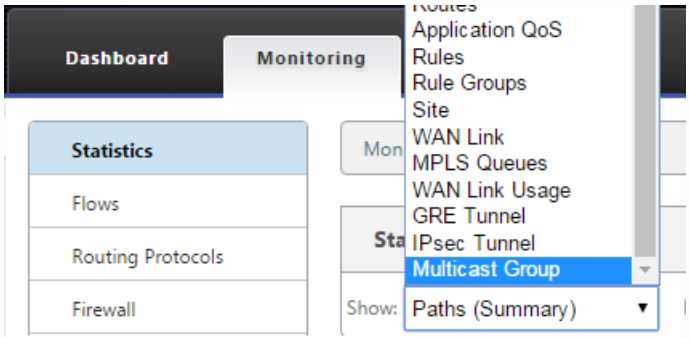

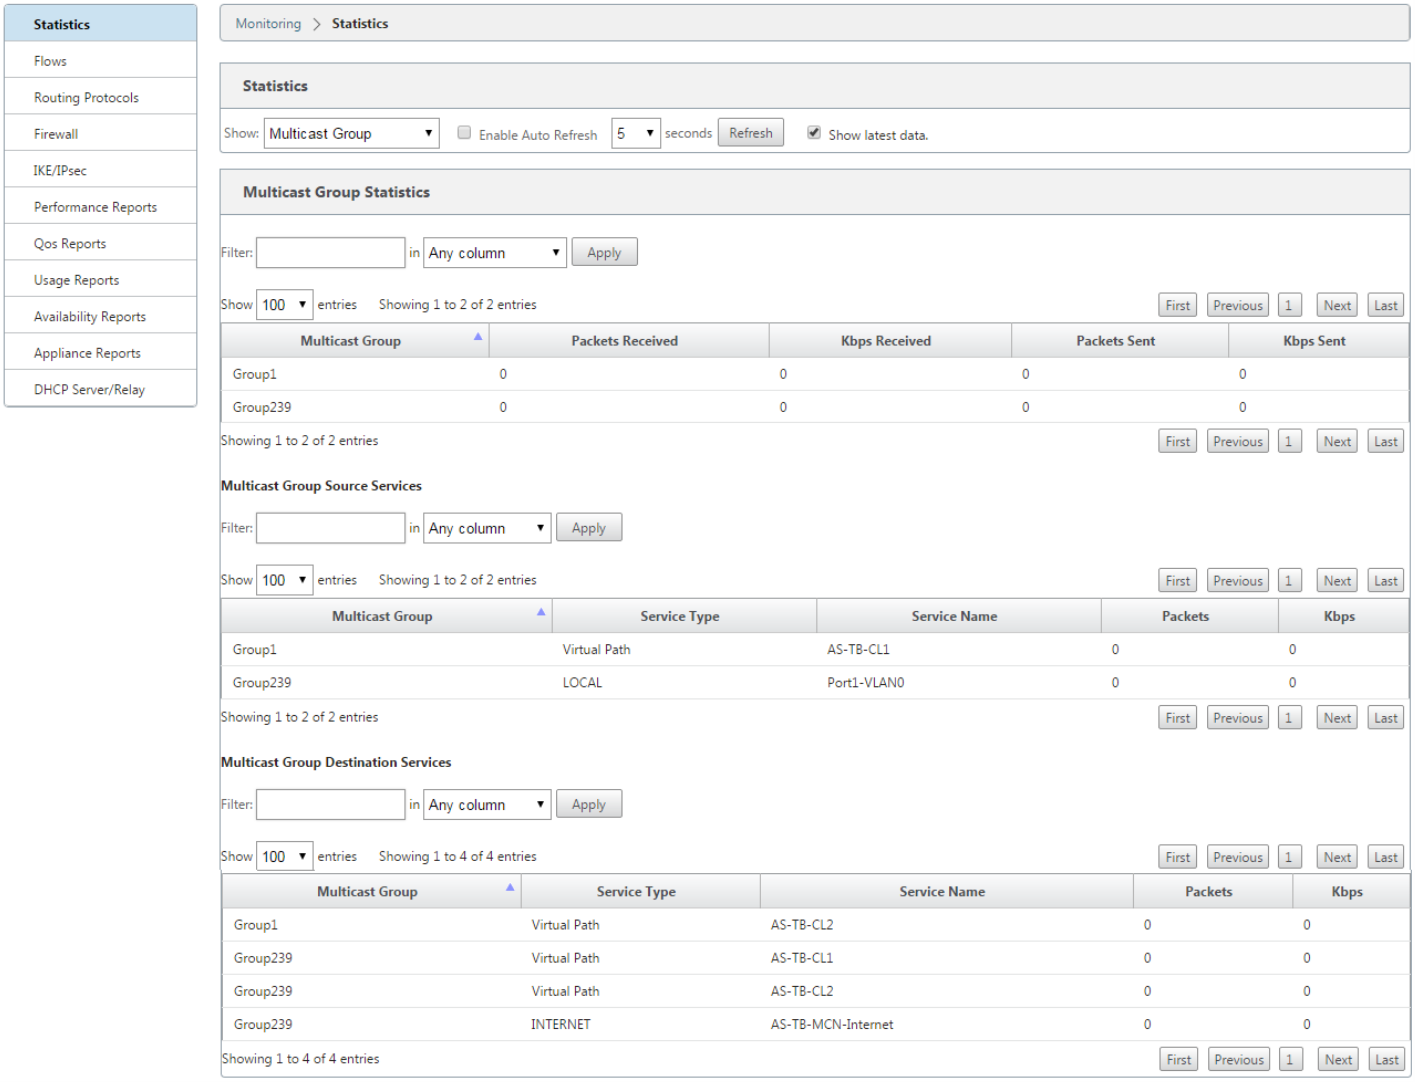

-
-
  -
  -
-
-

-
-

Topology

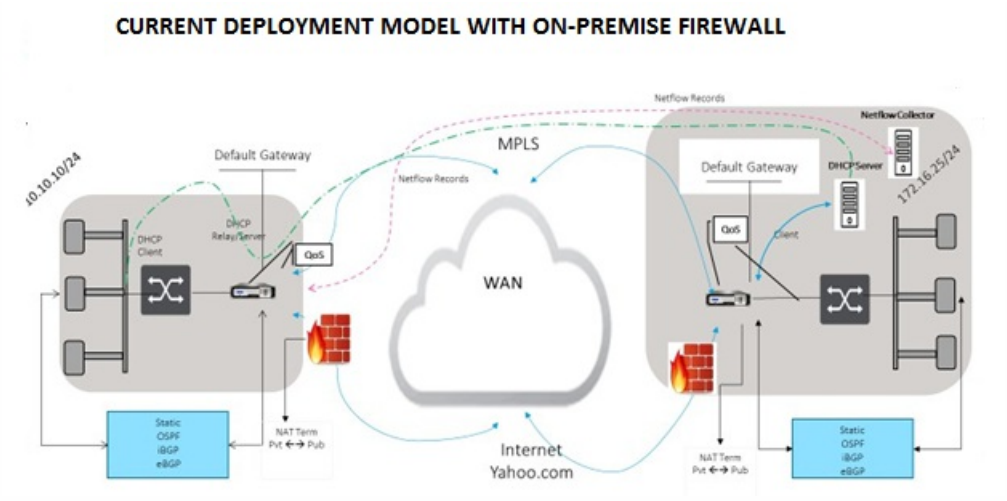

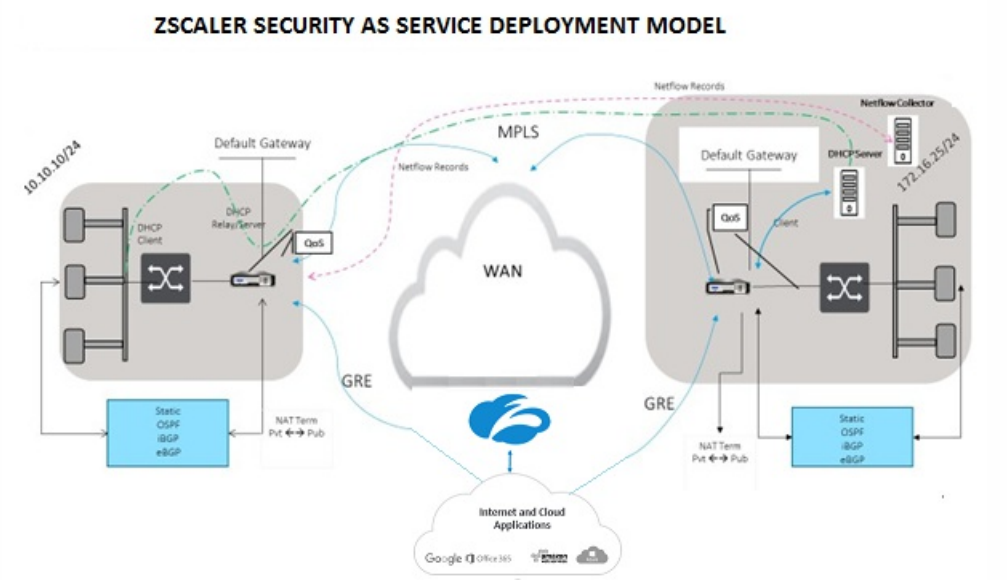

Internet traffic forwarding to Zscaler through GRE and IPsec Tunnels:

1. Log into the Zscaler help portal at: https://help.zscaler.com/submit-ticket.
2. Raise a ticket and provide the static public IP address, which is used as the GRE tunnel or IPsec tunnel source IP address.

Zscaler uses the source IP address value to identify the customer IP address. This value must be a static public IP address. Zscaler responds with two ZEN IP addresses to which to redirect traffic. GRE keep-alive messages can be used to determine the health of the tunnels.

Sample IP addresses:

Configuring an Internet Service

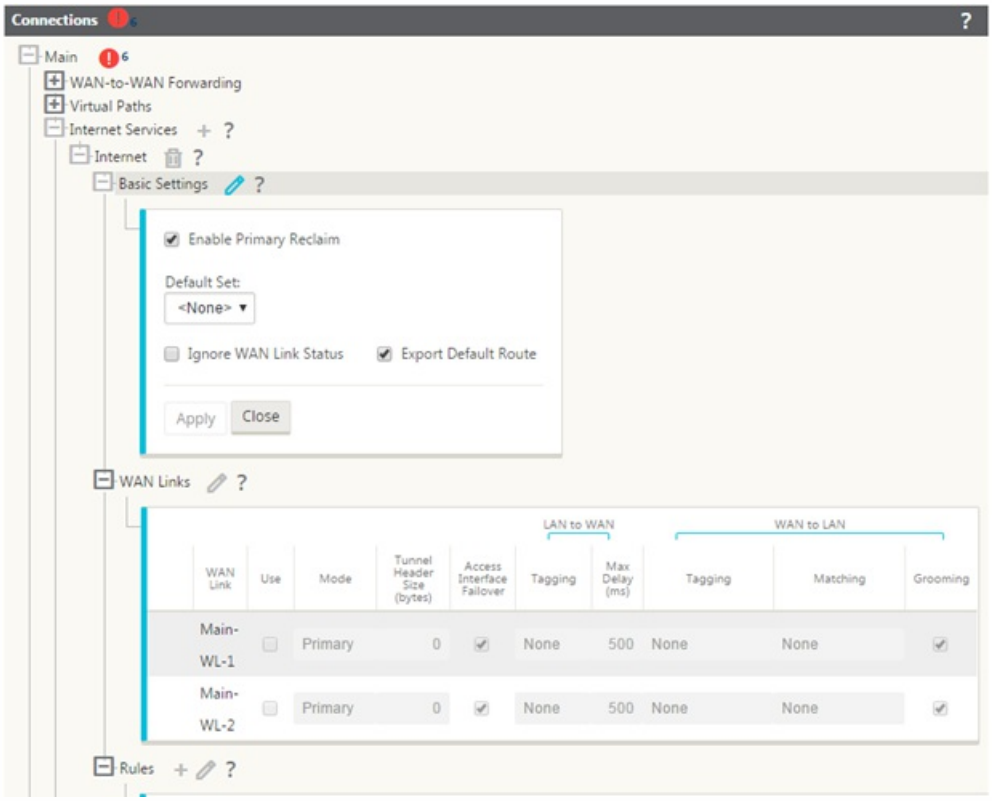

Configure GRE Tunnel

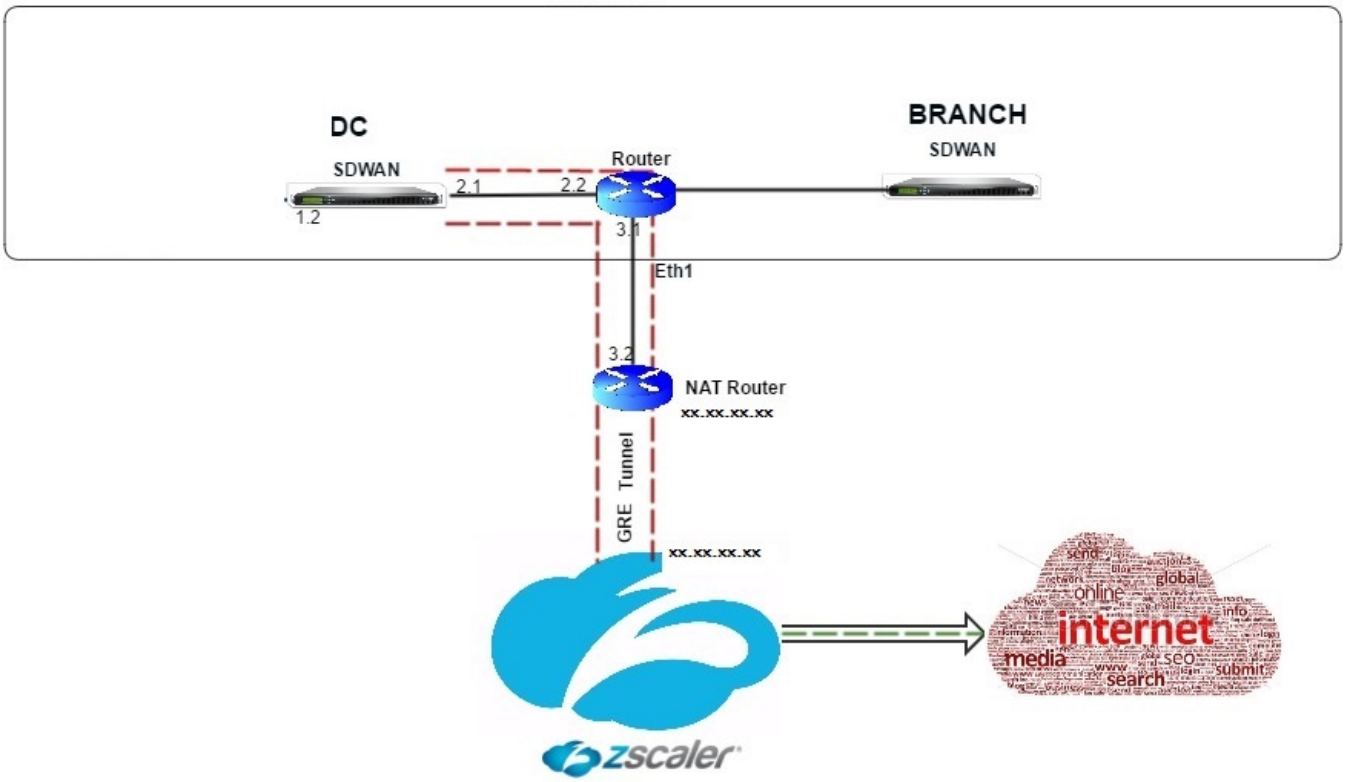

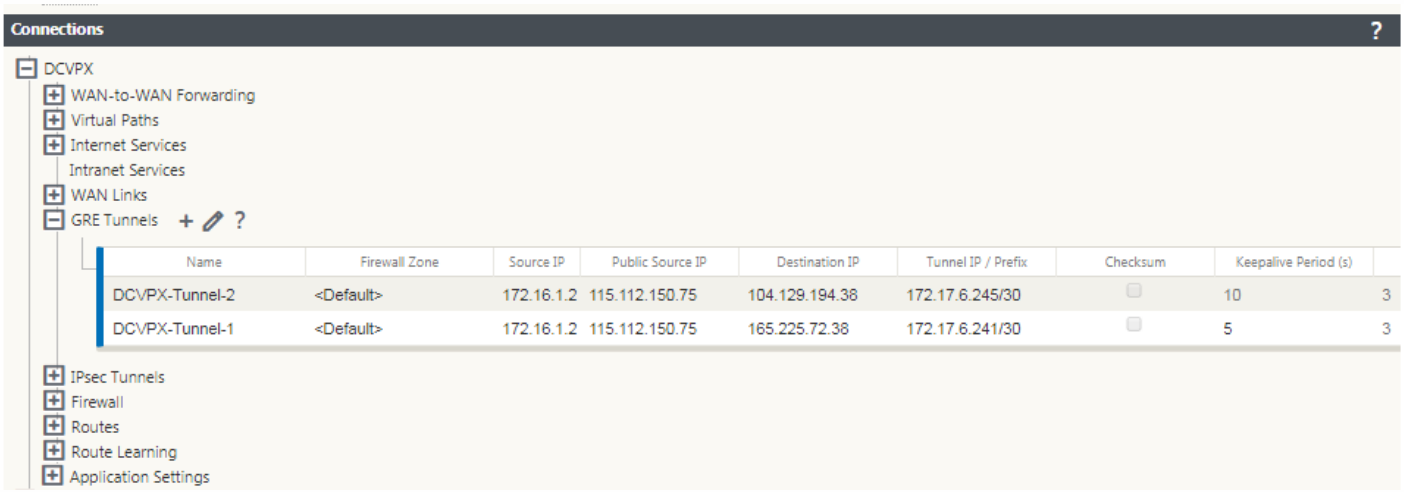

Configure Routes for GRE Tunnels

-
-

-
-

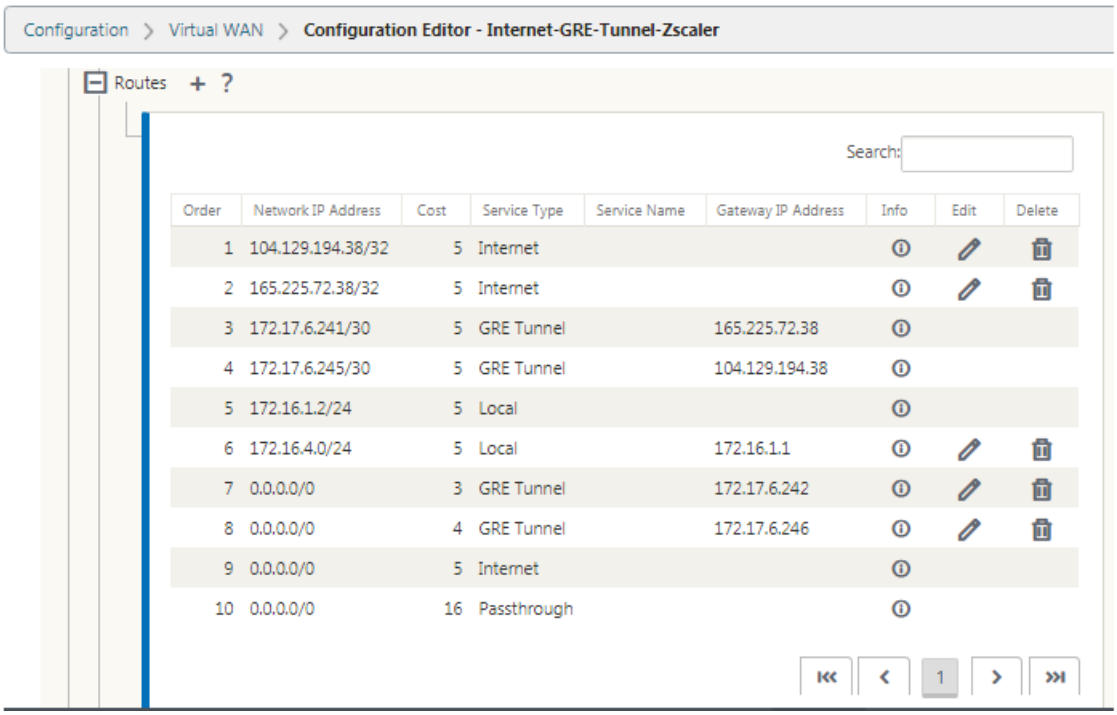

Limitations

-
-

Configure IPsec Tunnels

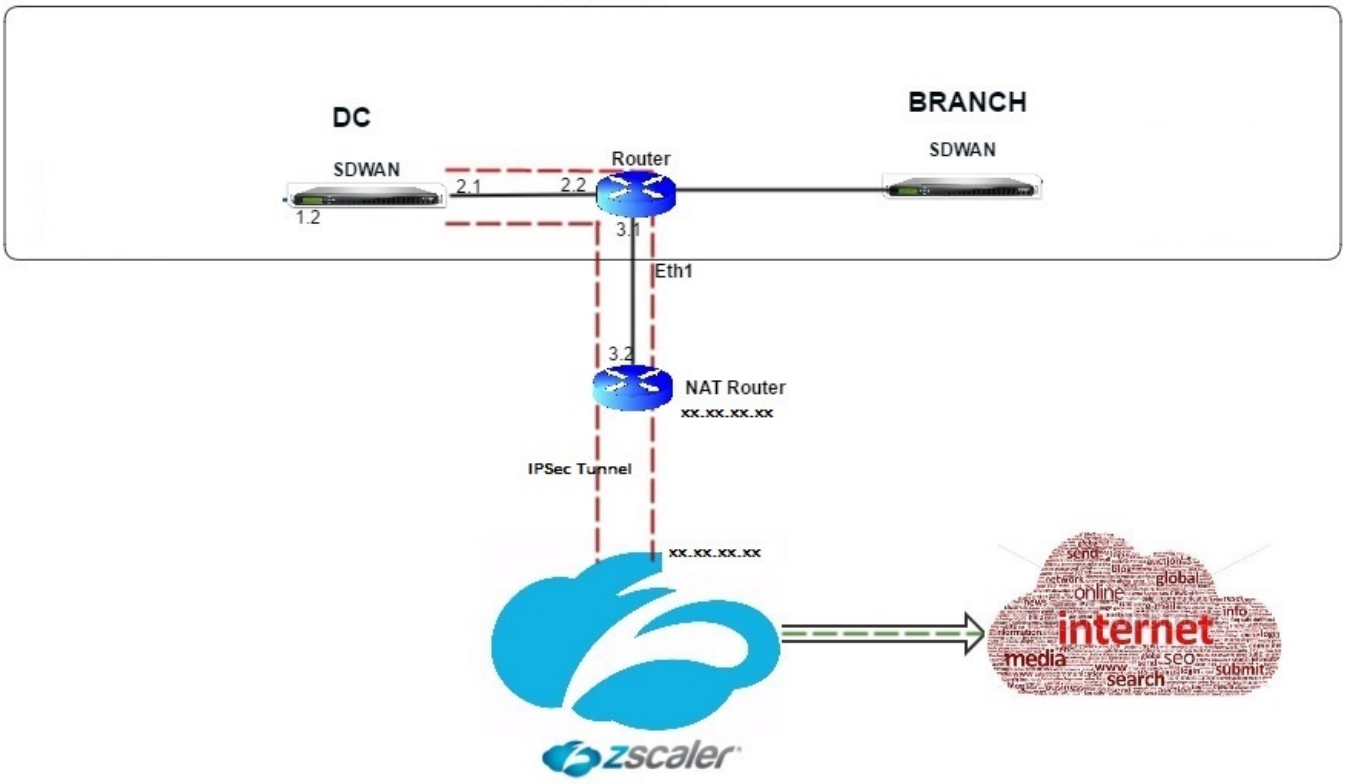

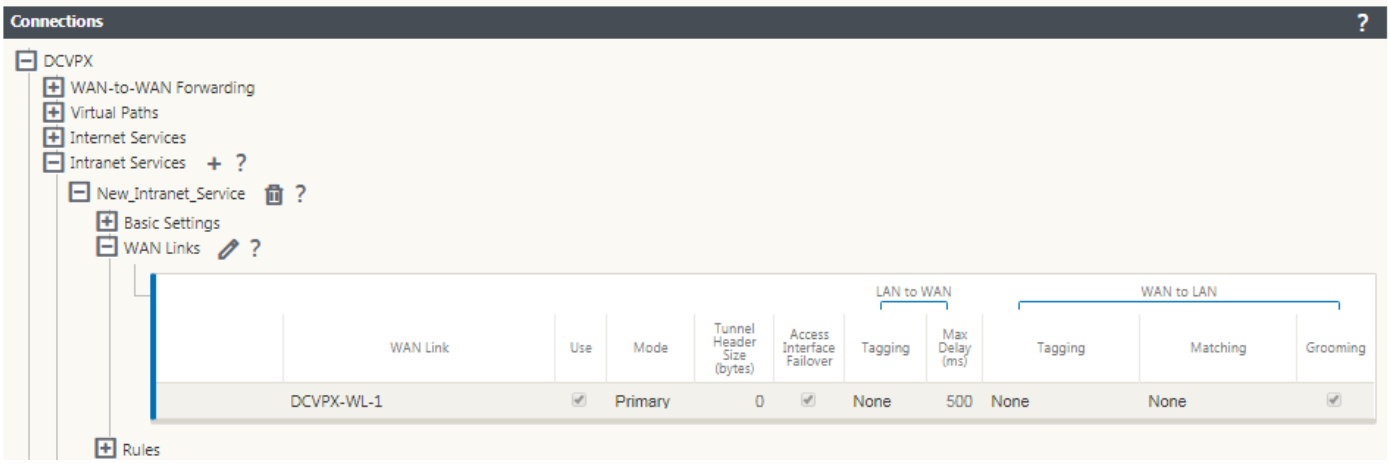

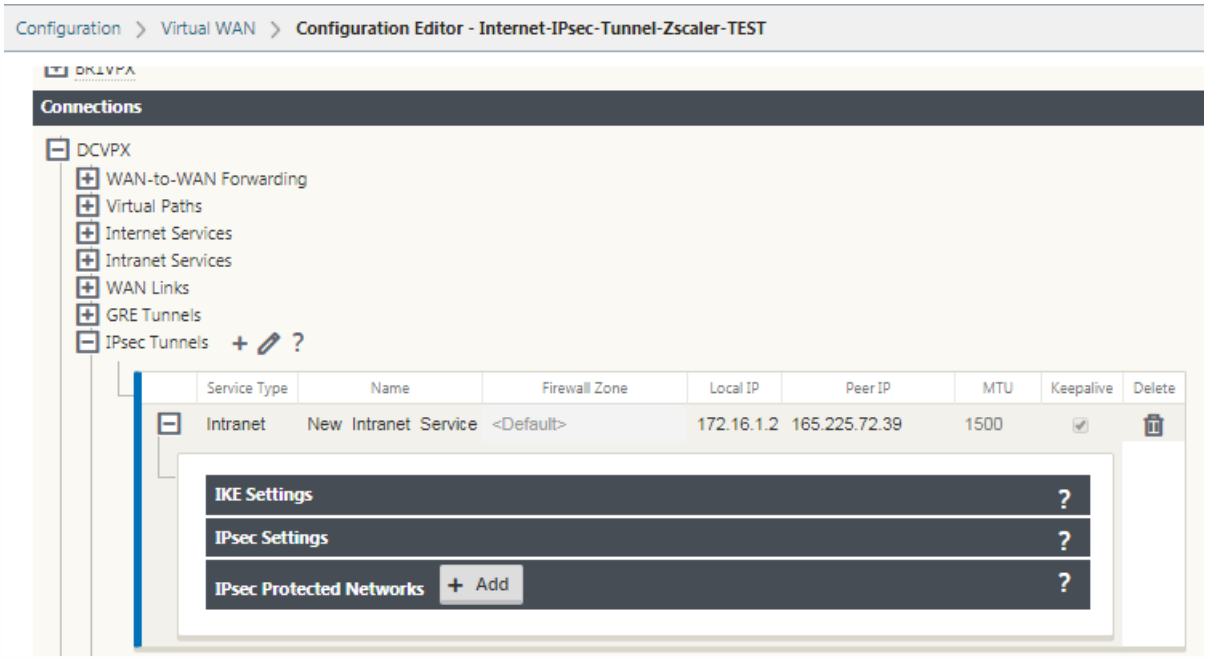

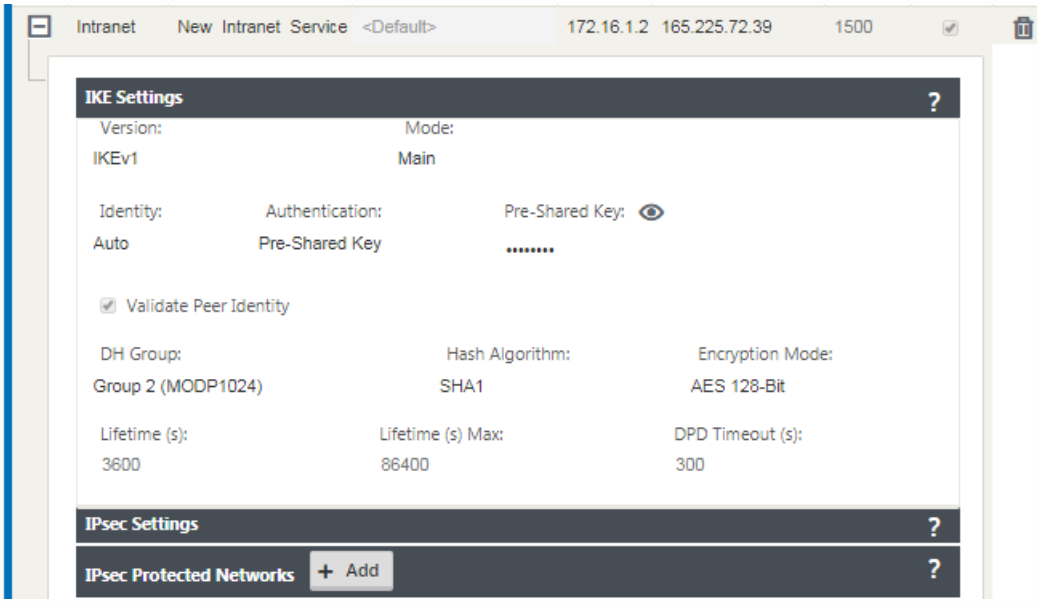

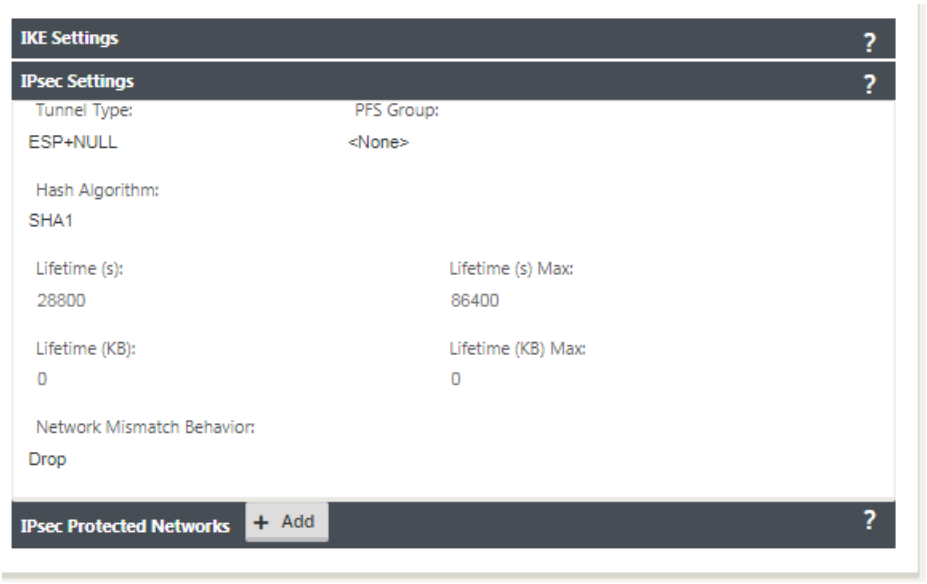

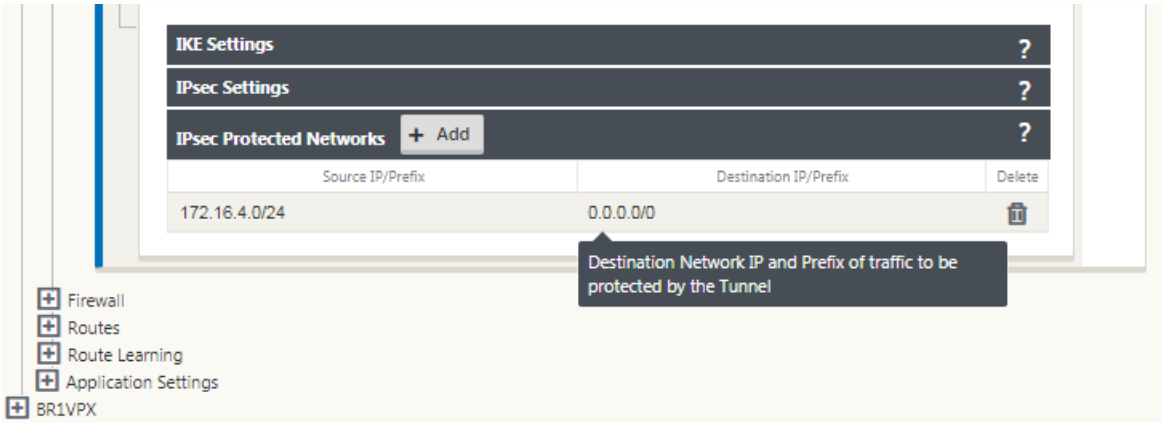

Configure Routes for IPsec Tunnels

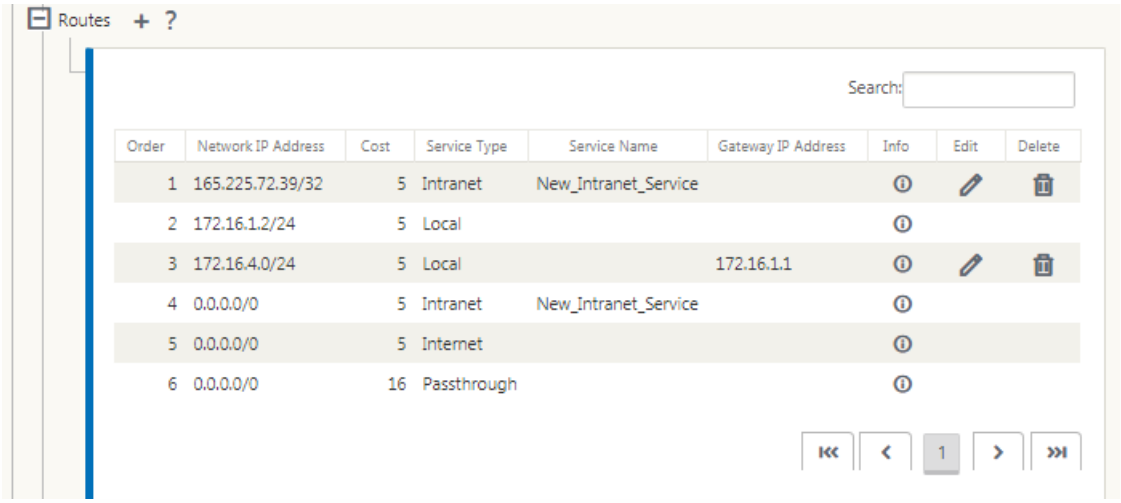

-
-
-

-
-

-

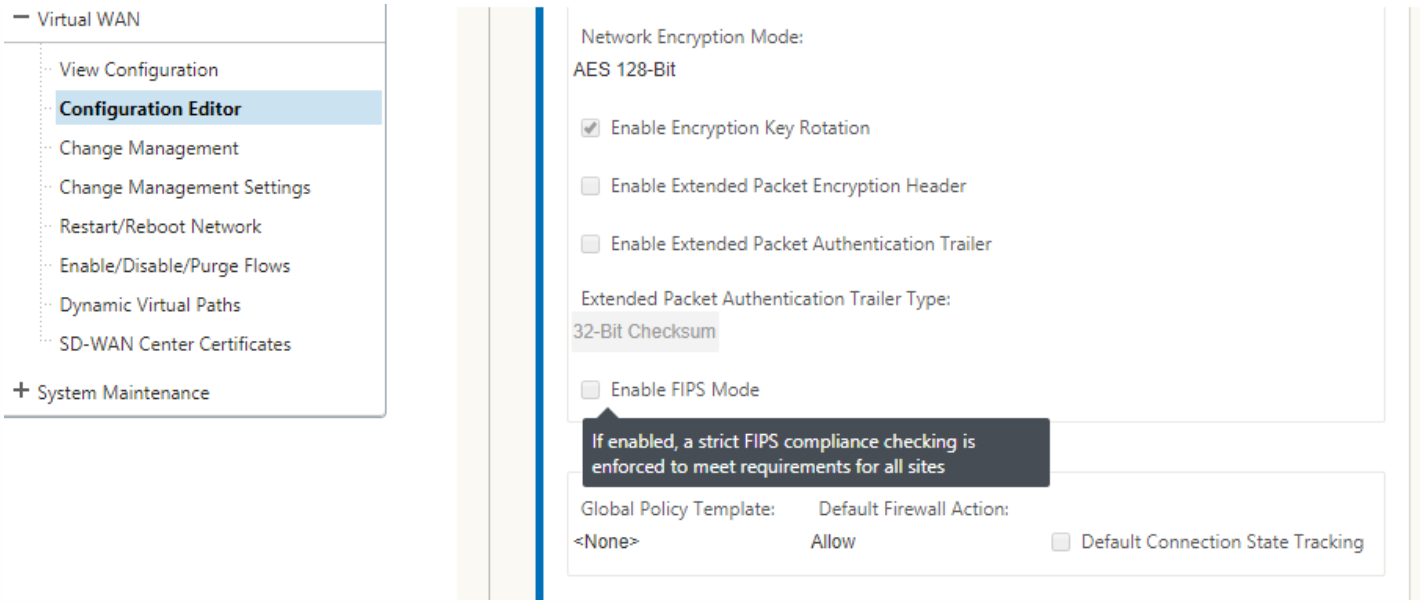

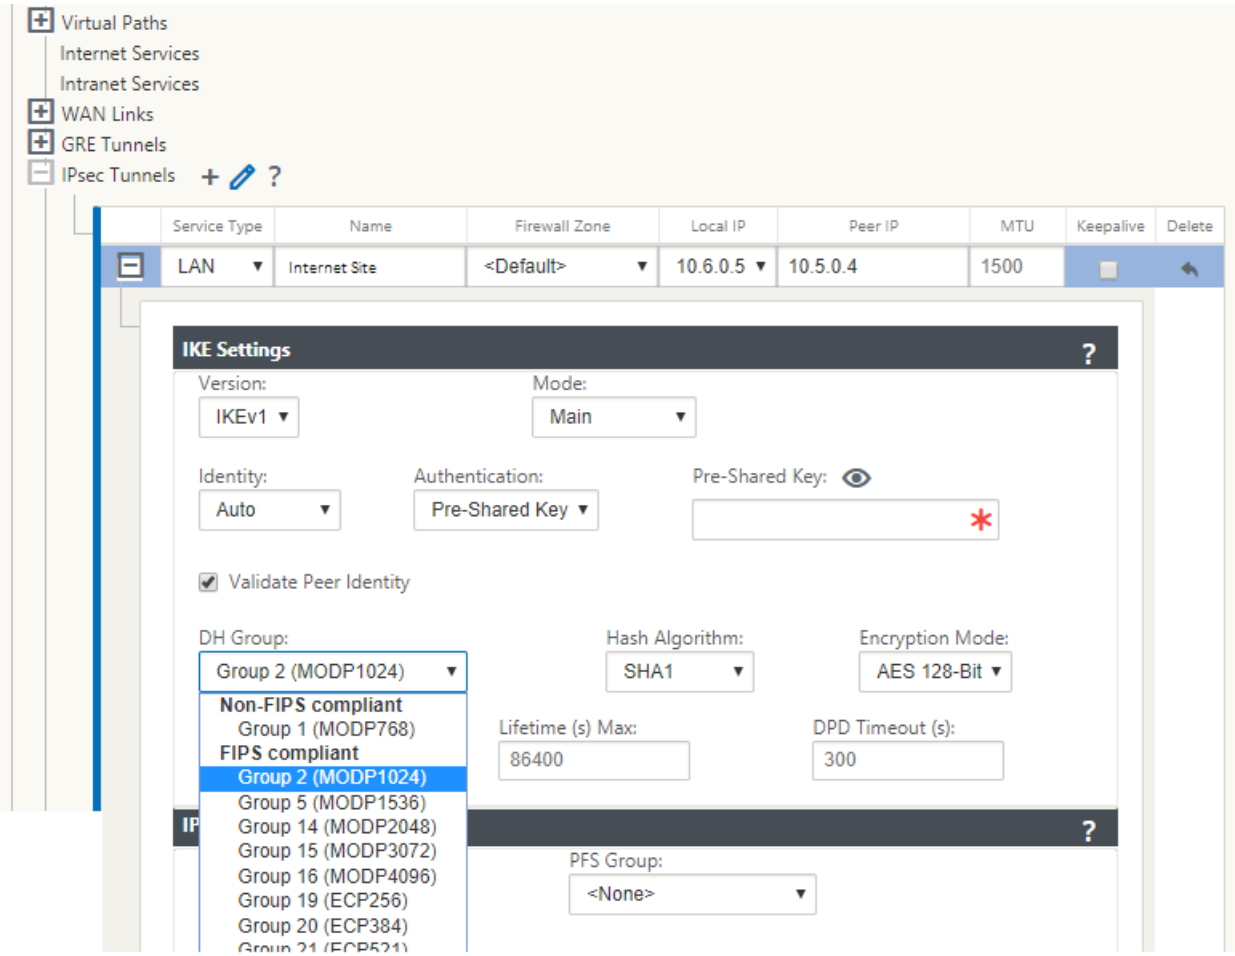

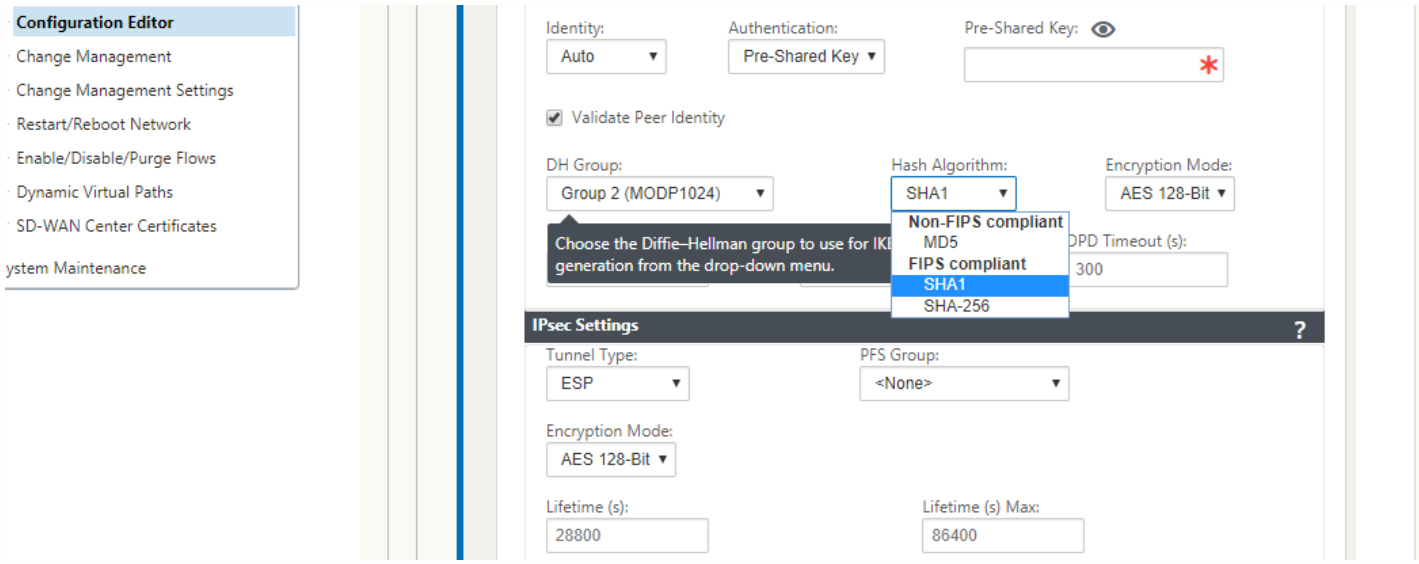

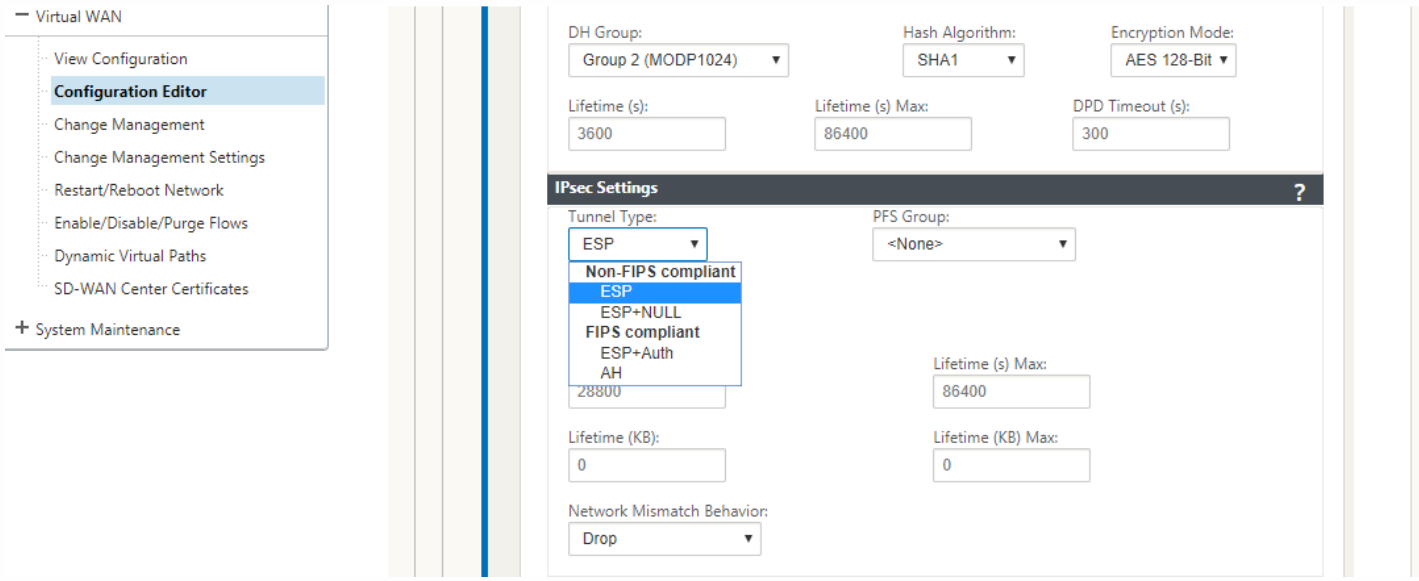

-
-
-
-

-
-
-

-
-
-
-
-

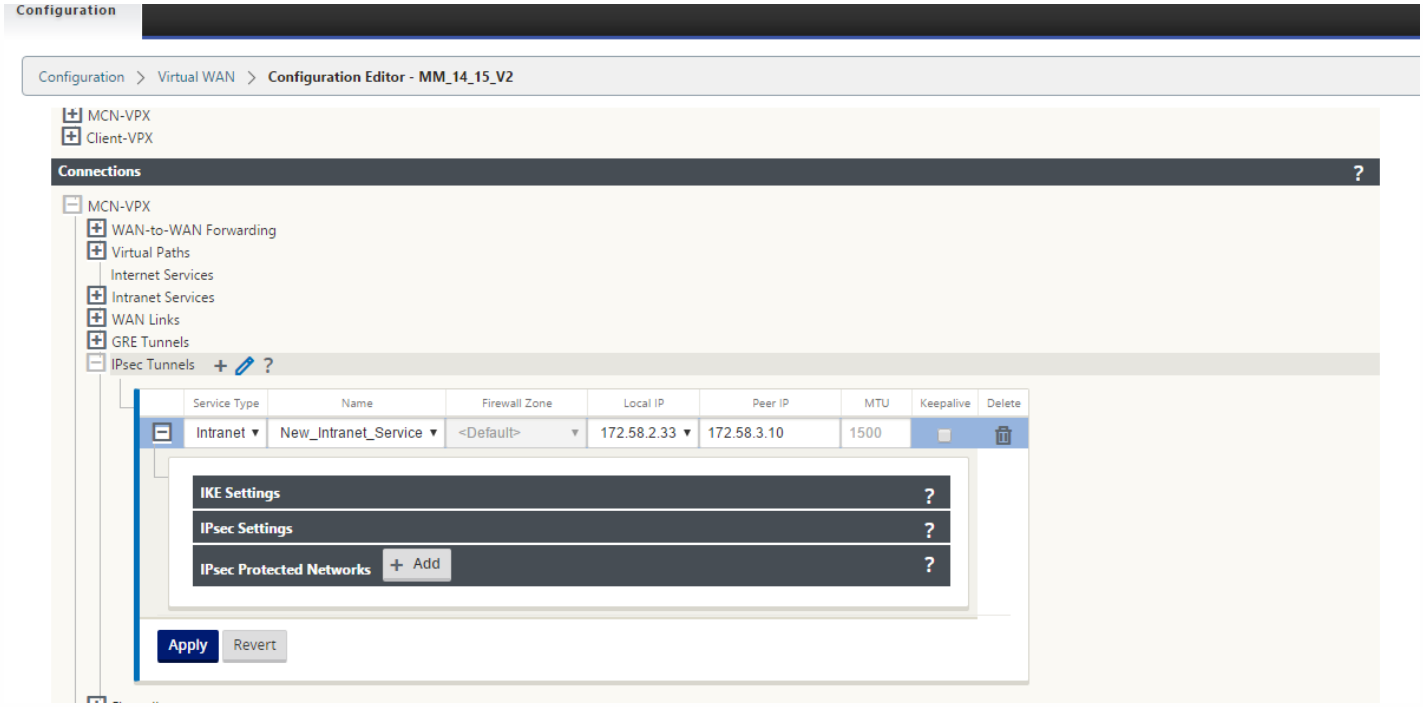

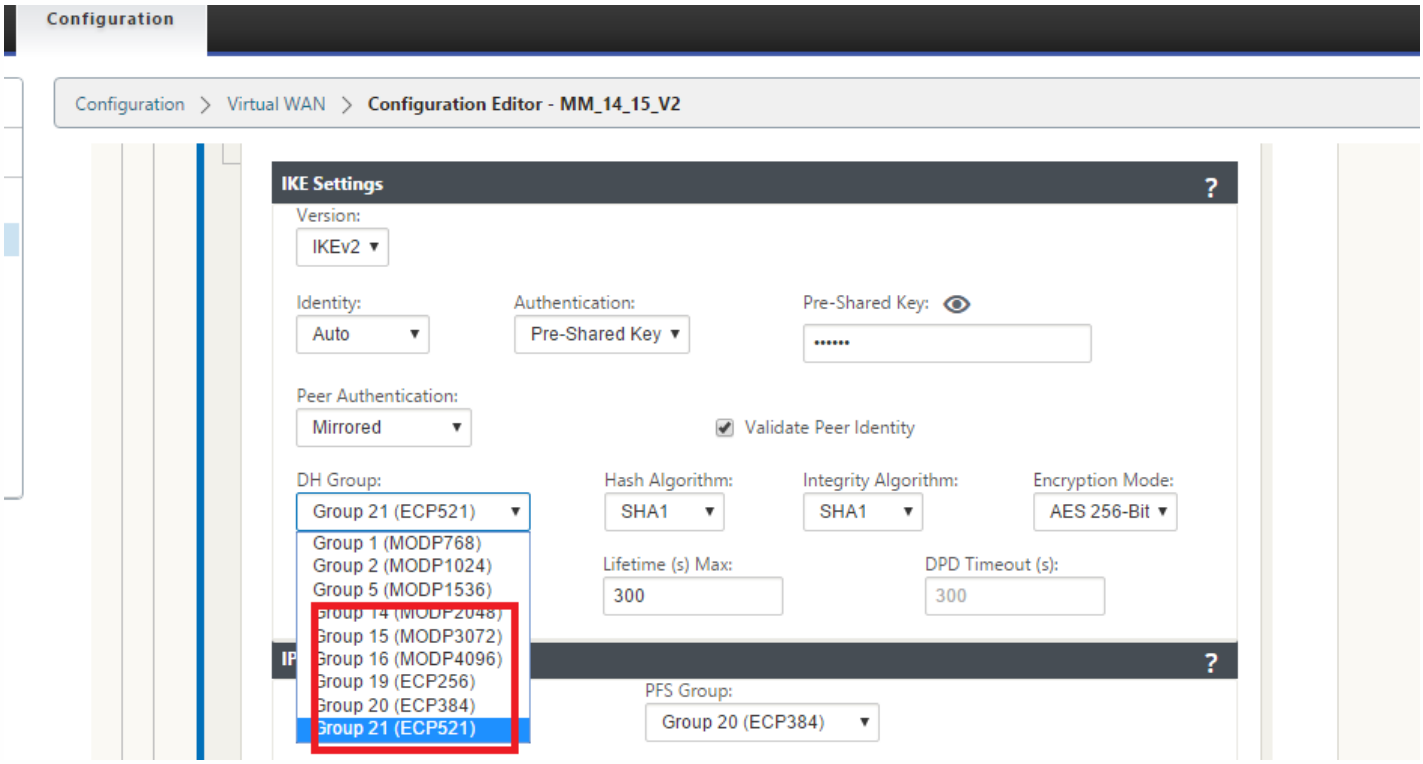

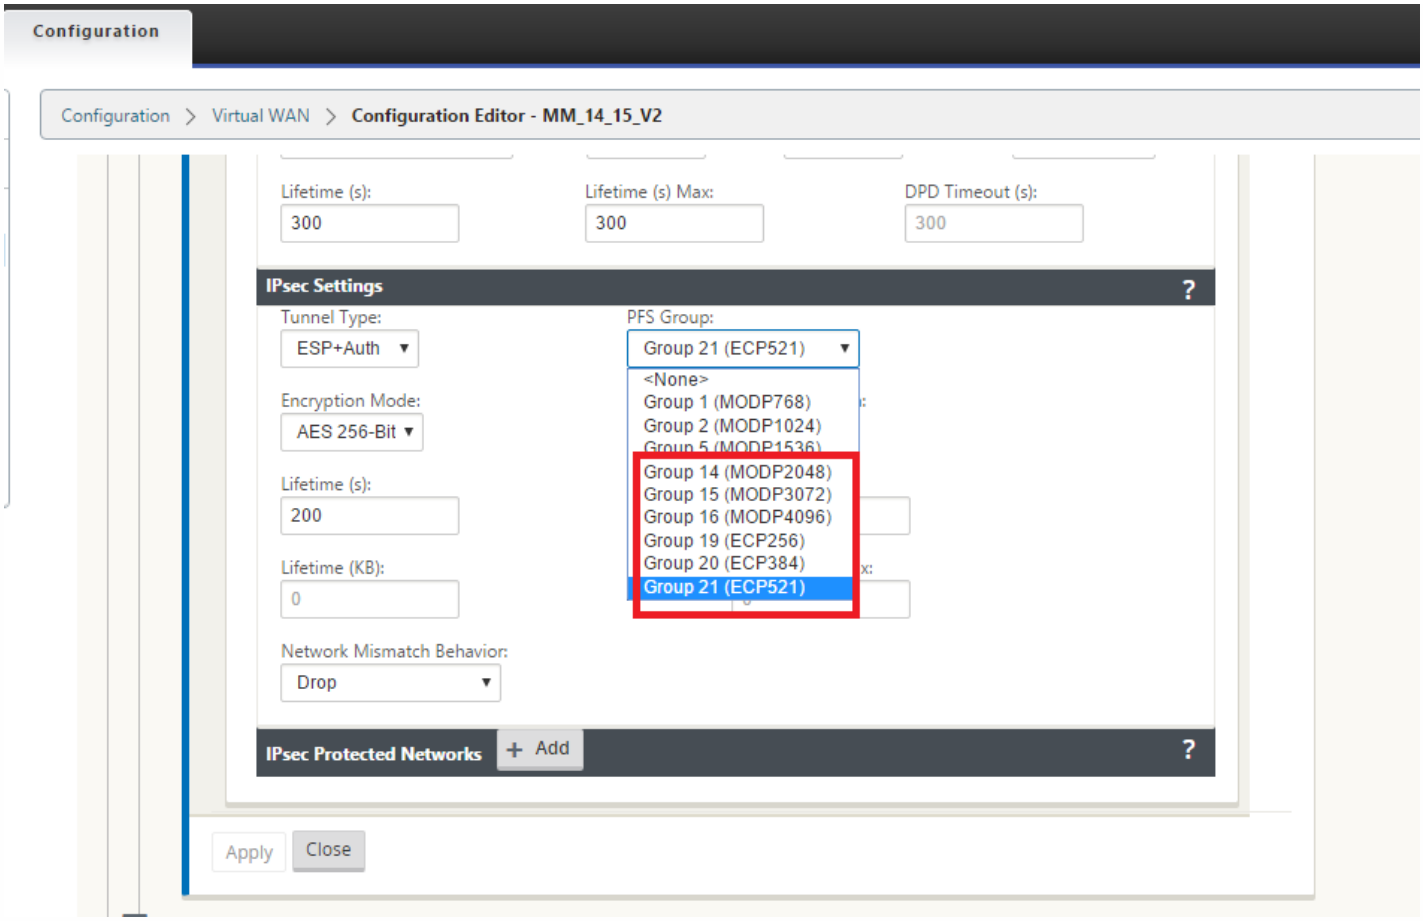

-
-
-
-
-
-

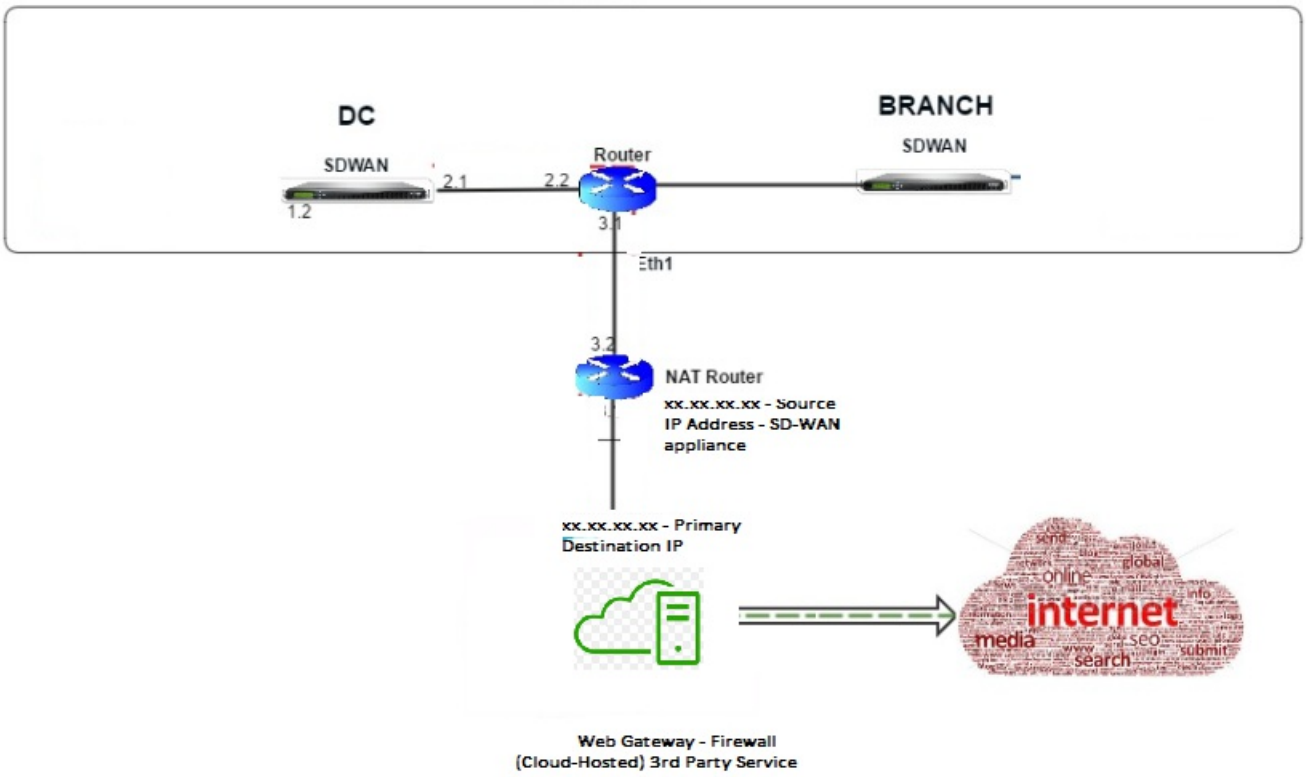

Configuring Destination NAT (DNAT)

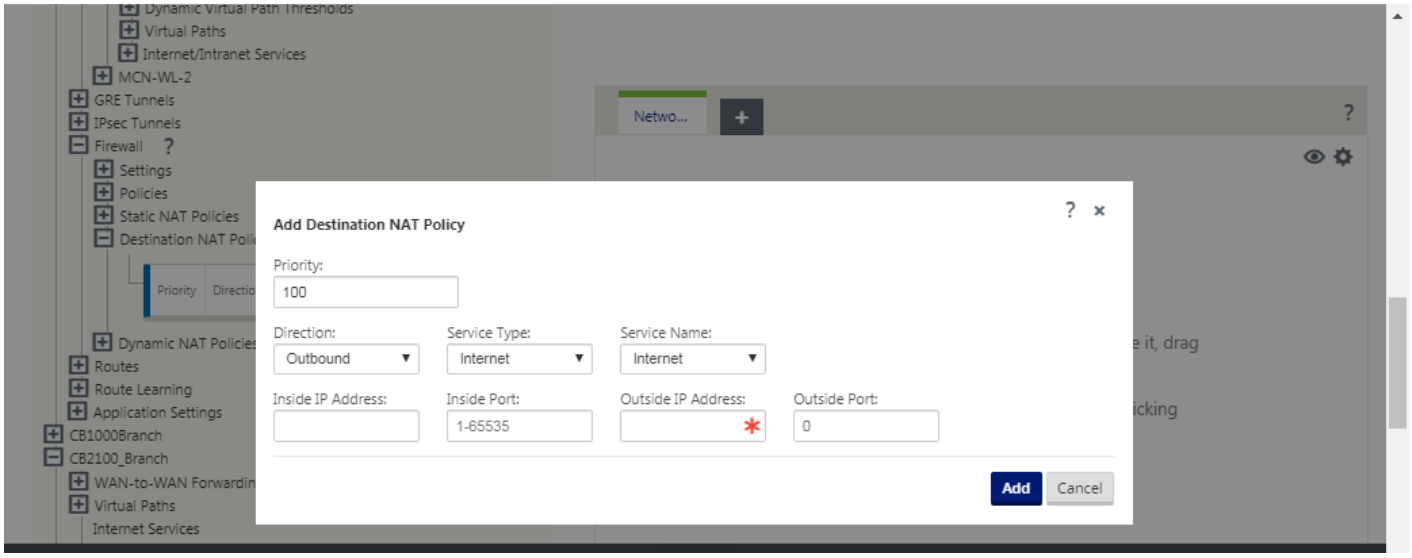

Monitoring a Destination NAT Policy (Firewall)

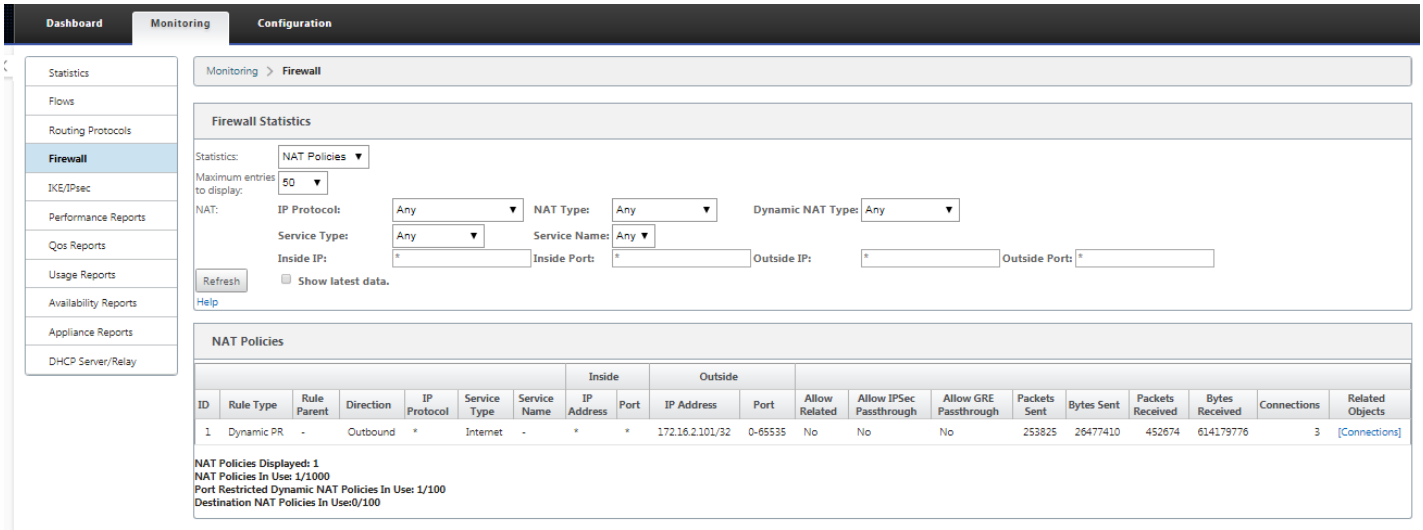

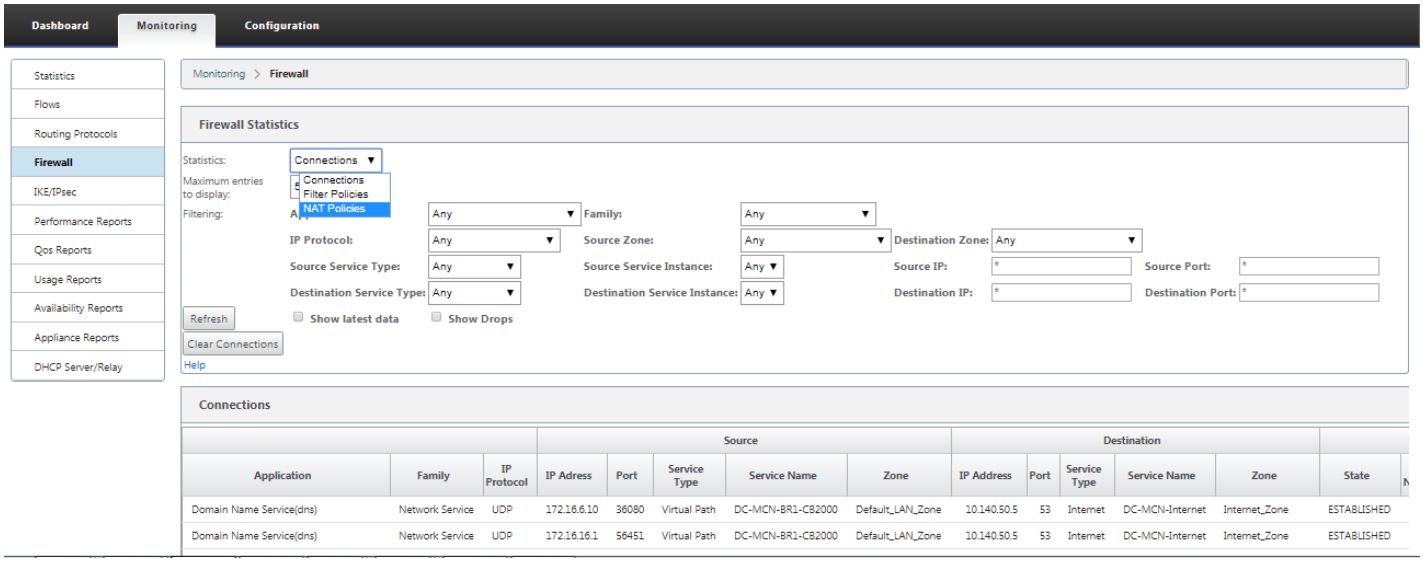

-
-
-

Internet Traffic Control

Direct Internet Breakout at Branch with Integrated Firewall

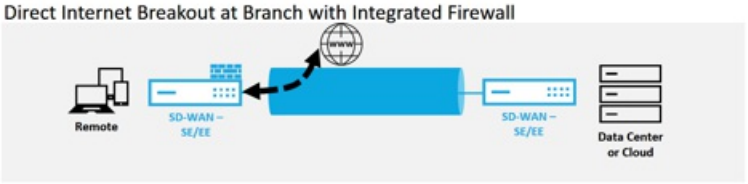

Direct Internet Breakout at Branch with forwarding to Secure Web Gateway

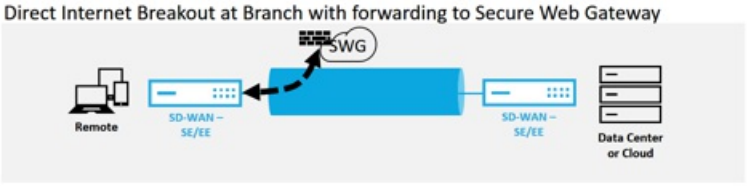

Backhaul Internet to Data Center MCN

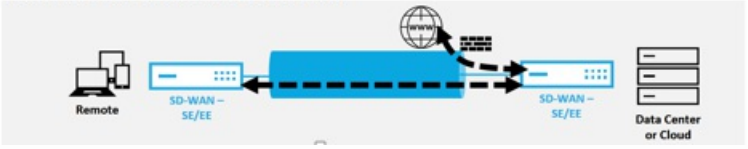

Direct Internet Breakout at Branch with Integrated Firewall

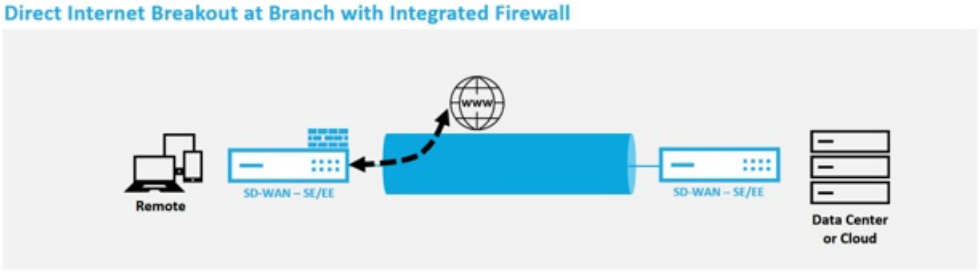

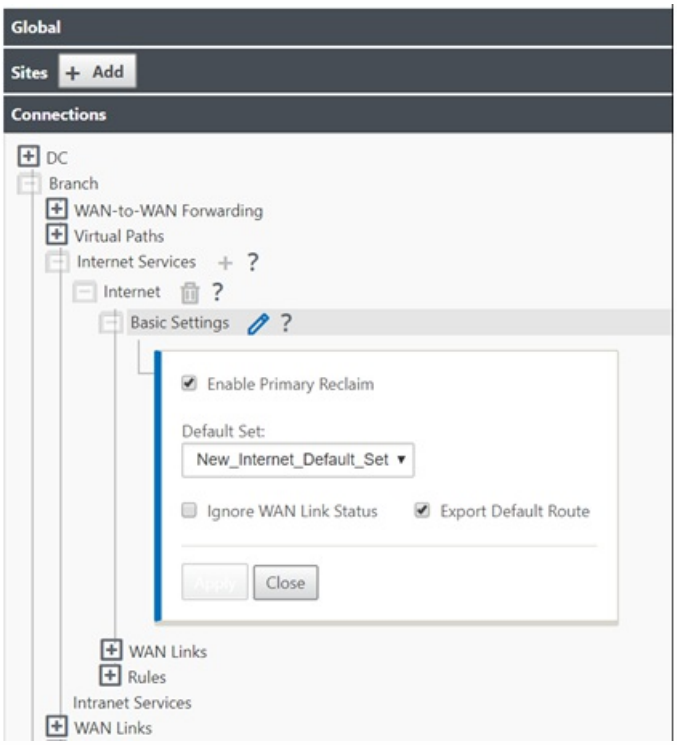

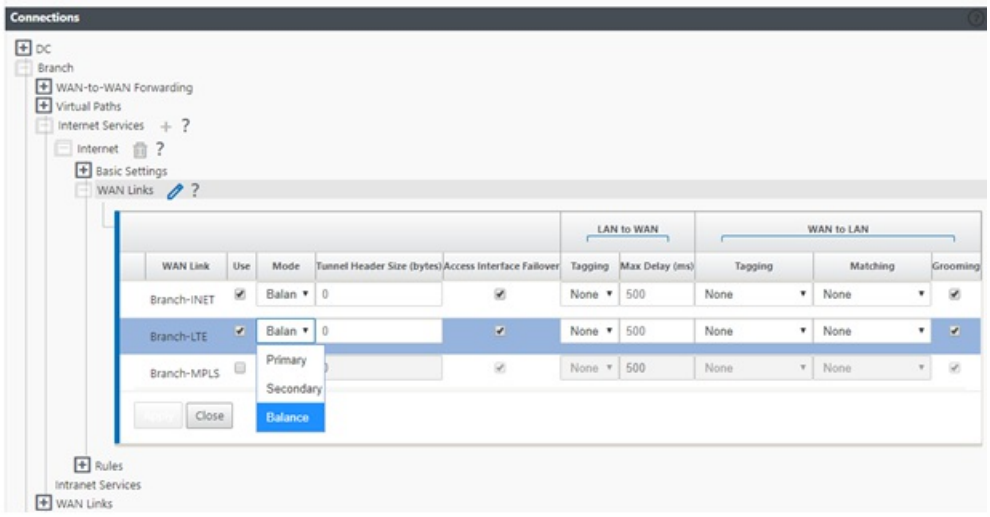

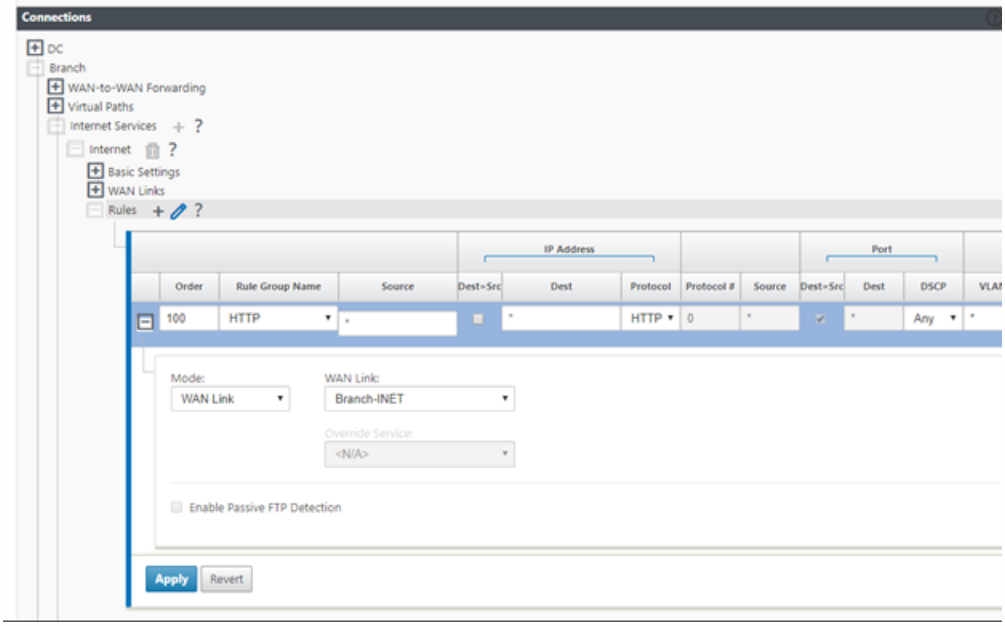

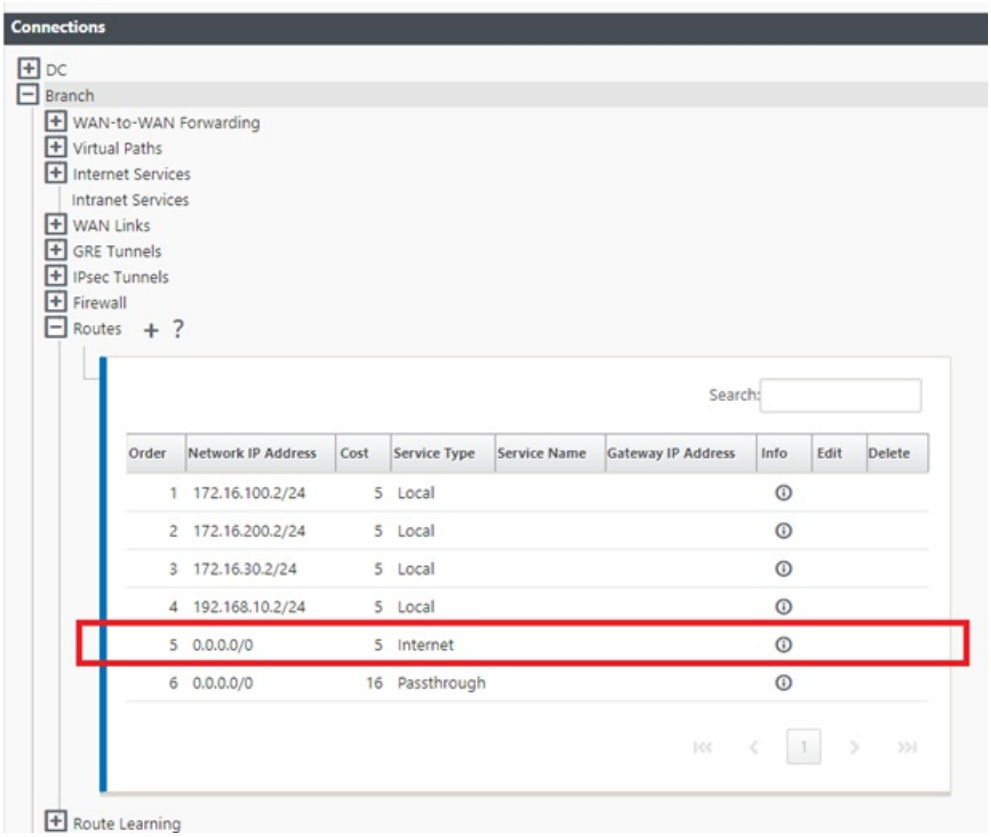

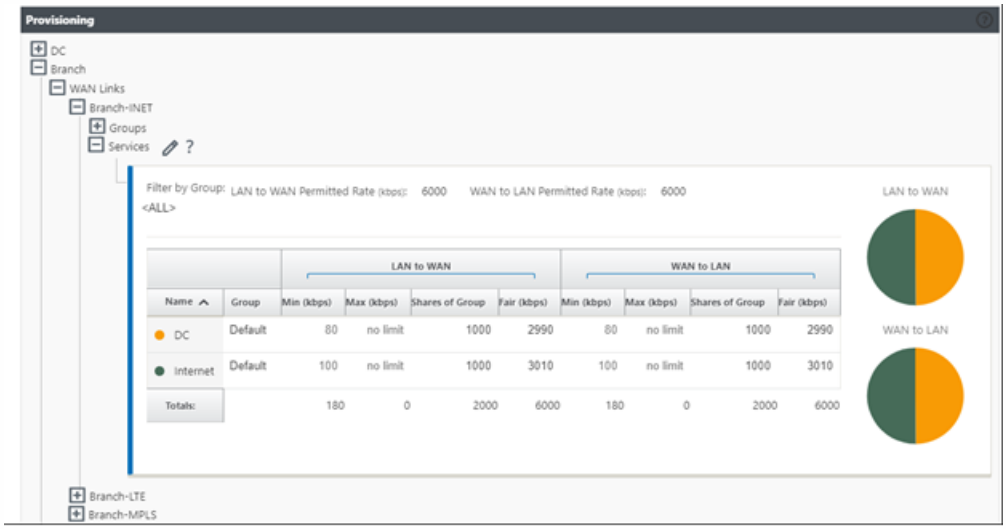

-

Direct Internet Breakout at Branch with forwarding to Secure Web Gateway

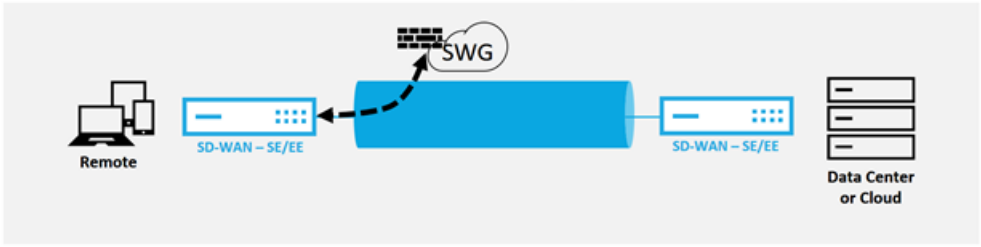

-
-

-
-
-

Backhaul Internet to Data Center MCN

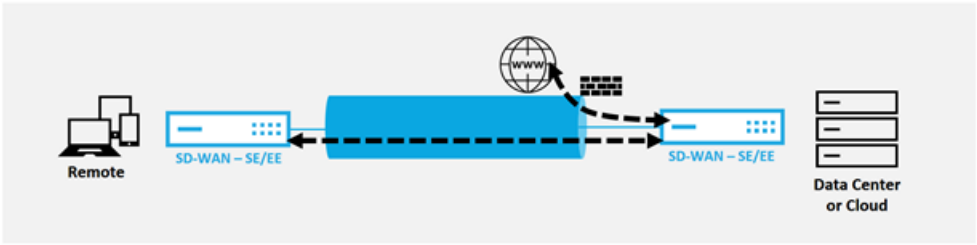

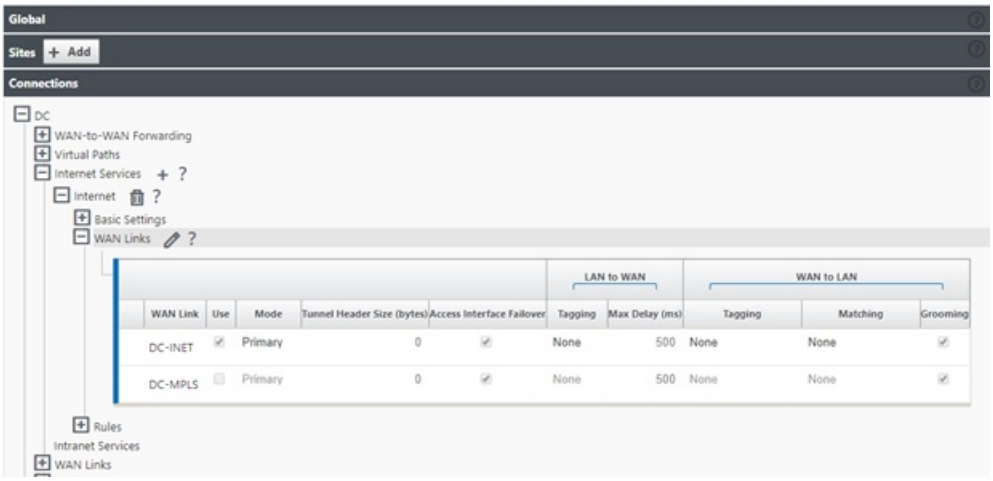

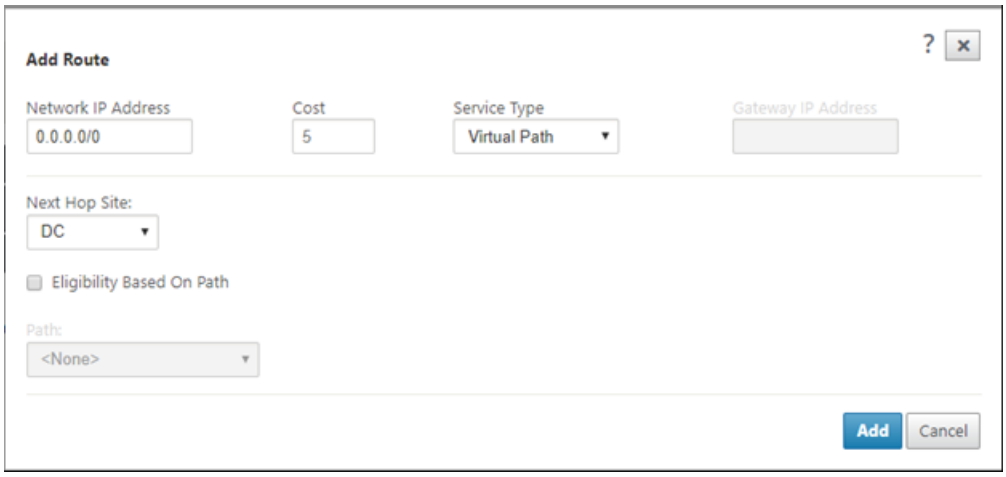

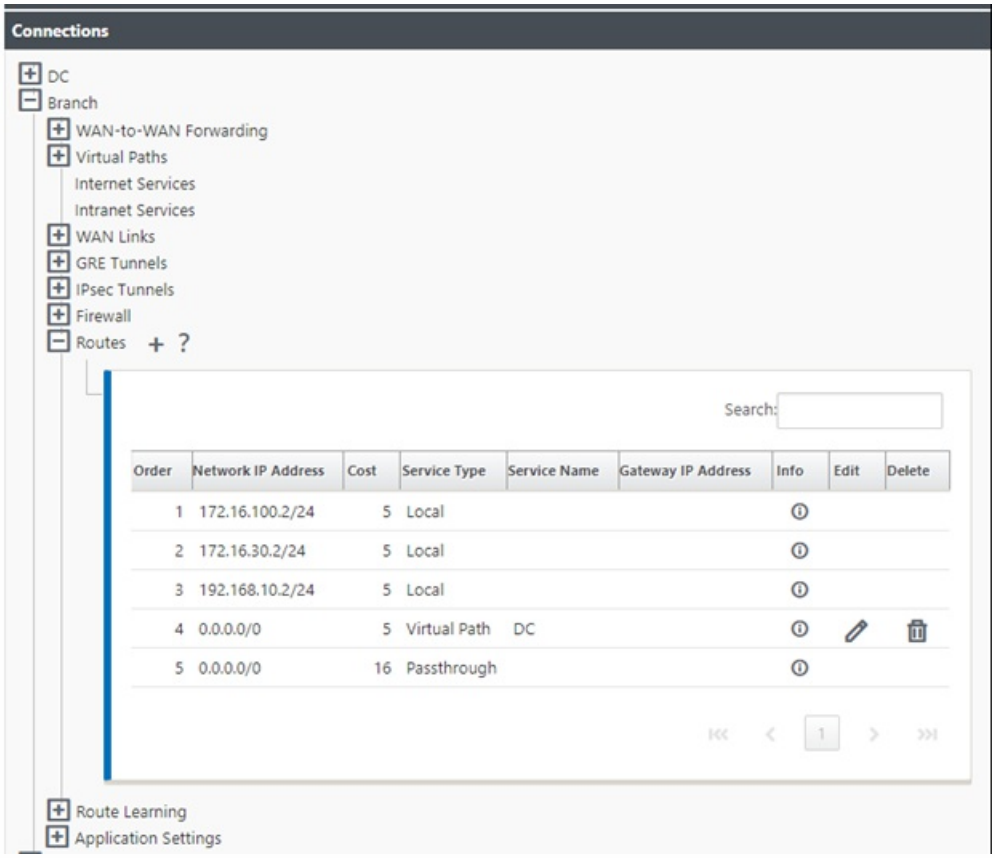

-
-

How To Enable DHCP Server

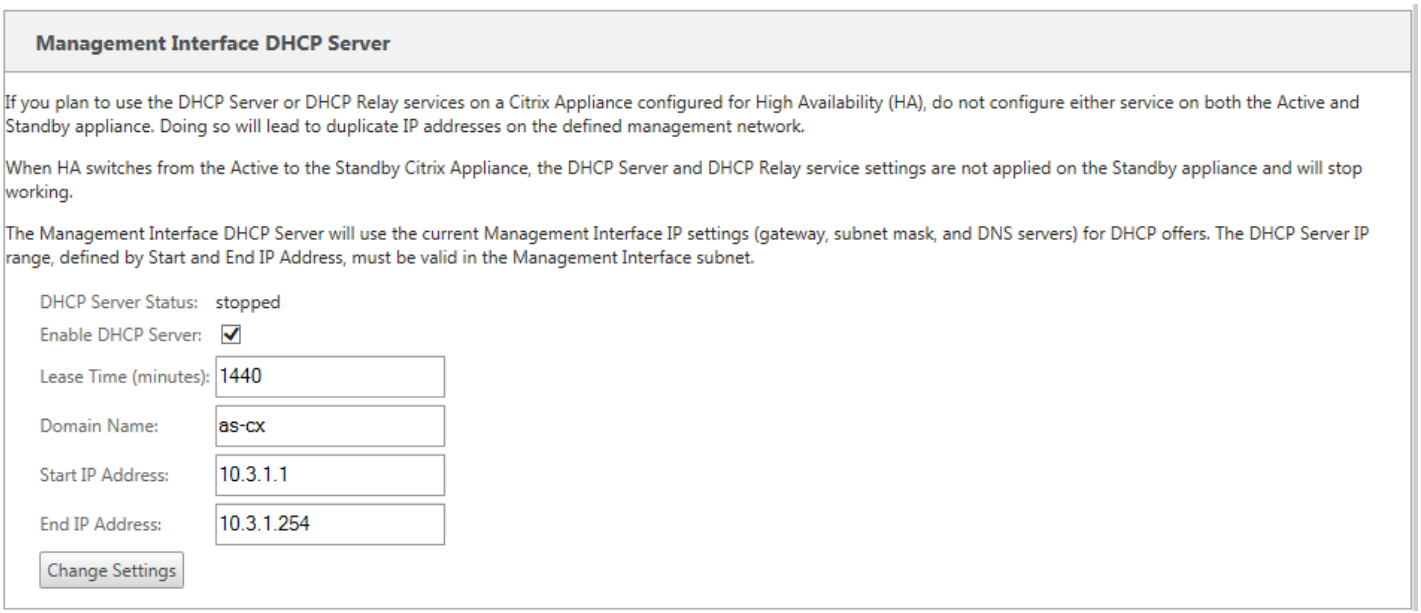

How To Enable DHCP Relay

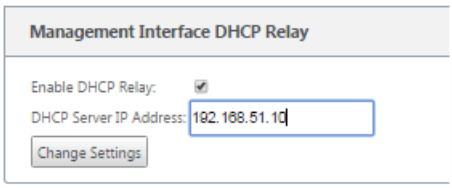

How To Monitor DHCP Client WAN Links

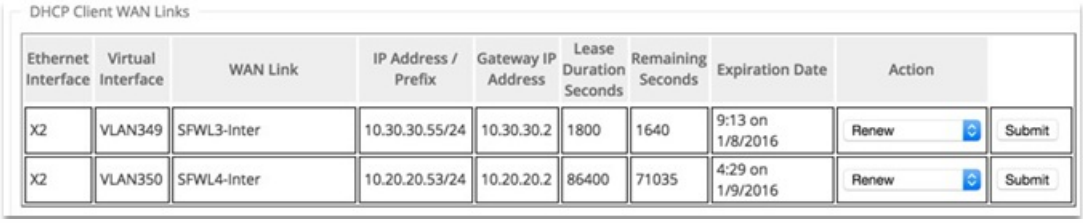

DHCP Management

Aug 09, 2017

Devices on the same network as the SD-WAN appliance's LAN/WAN interface can now use the SD-WAN DHCP Relay & DHCP Server features to provide those devices with their IP configuration. These features help to simplify the client site network by reducing the amount of equipment necessary.

- Reduce equipment at client site
- Replace router at client site (Easy deployment of edge router services)
- Simplify the client site network
- Configuration of Router without CLI commands
- Use of Dynamic Host Configuration Protocol (DHCP) on Internet Protocol (IP) networks to request IP addresses and networking parameters automatically reducing manual configuration needs by network admin on simple client sites

DHCP Relay

Network administrators can now use the DHCP Relay service of the SD-WAN appliance to relay requests and replies between local DHCP clients and a remote DHCP Server. This allows local hosts to acquire dynamic IP addresses from the remote DHCP Server.

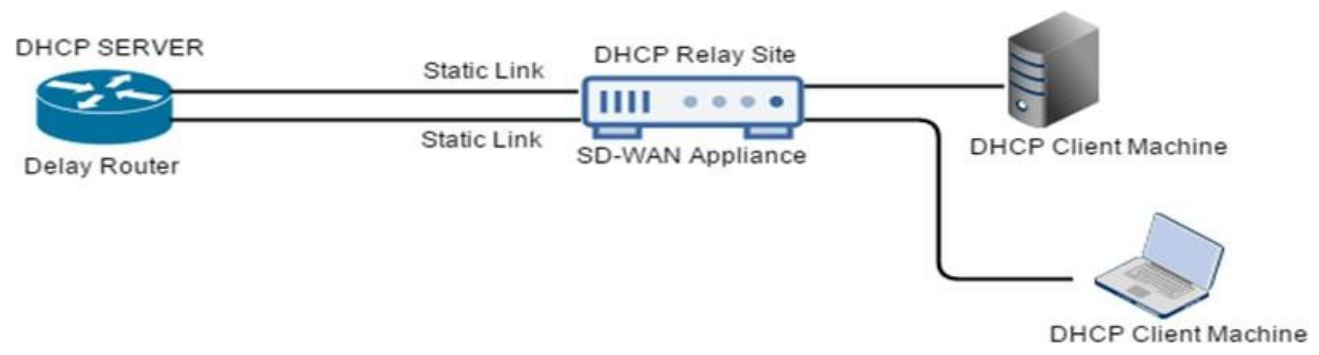

To configure DHCP Relay:

1. In the SD-WAN web management interface, navigate to **Configuration Editor** > **Sites** > [Site Name] > **DHCP** > **Relays**. Expand Relays, and then specify the server IP address.
2. Select the **Virtual Interface** to be used.
3. Configure a static route to reach the DHCP Server.

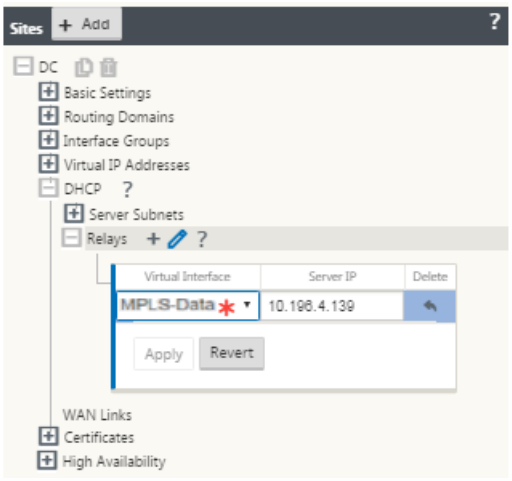

DHCP Server

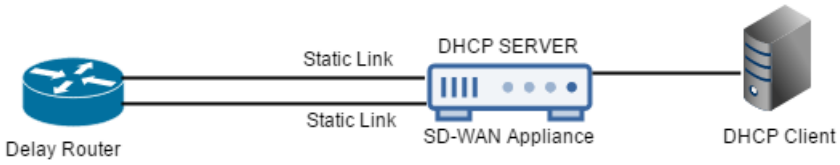

Network administrators can use the DHCP Server feature on data ports of an SD-WAN appliance to allow local hosts to acquire dynamic or static IP addressing directly from the SD-WAN appliance.

To configure DHCP Server:

1. In the SD-WAN web management interface, navigate to **Configuration Editor** > **Sites** > [Site Name] > **DHCP** > **Server Subnets**.

2. Select the Virtual Interface to be used and specify the range of IP addresses allowed to be dynamically assigned to local hosts.

3. Optionally, specify settings for configuring the hosts, such as gateway IP address, DNS address, and an Option Set (described below).

The Hosts option of this drop-down gives users an option to manually associate specific IP addresses with specific hosts through the host MAC addresses.

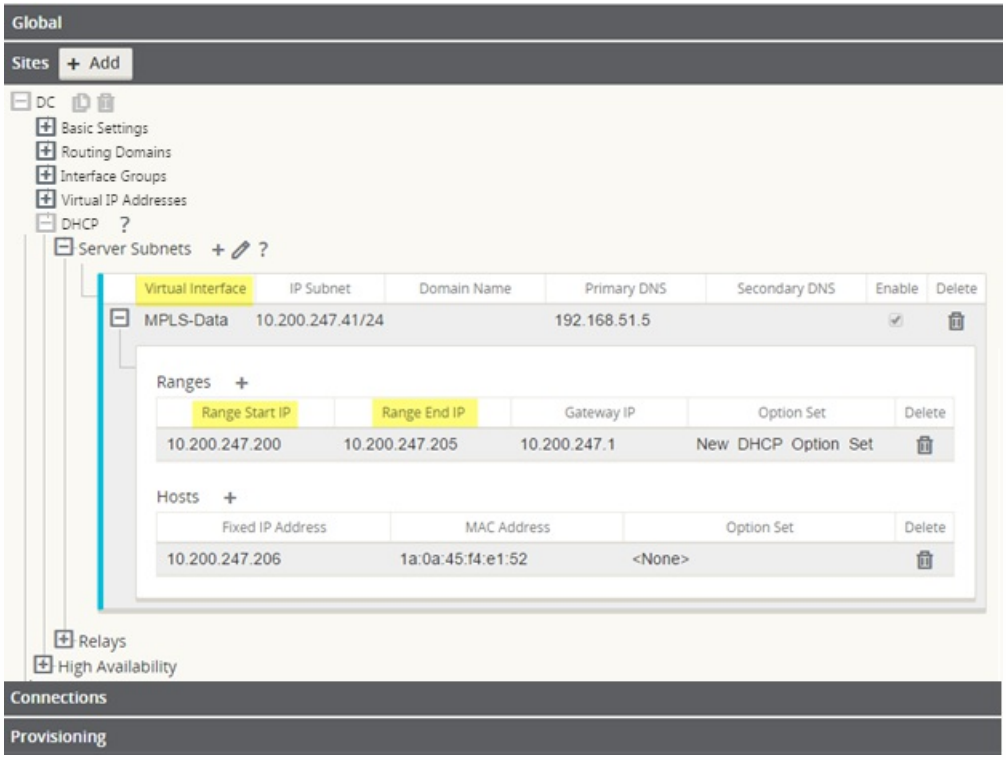

Note

The following feature is optional, not required.

DHCP Option Sets are groups of DHCP settings that can be applied to individual IP address ranges. To create DHCP Option Sets, navigate to the **Global** section of the configuration and expand **DHCP Options Sets**. Enter the settings that you would like to include in the set, and then click **Apply**.

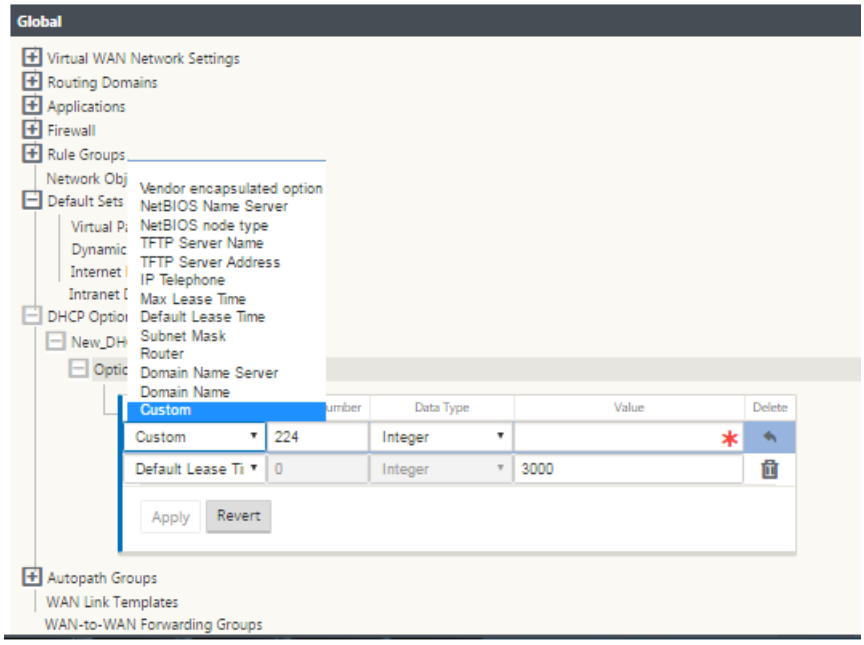

Your DHCP Option Set must then be assigned to a DHCP range. Do this in the Sites section where the IP address range was defined. The DHCP Option Set defined globally takes precedence over local configuration for same parameters

configured for the DHCP Server pool.

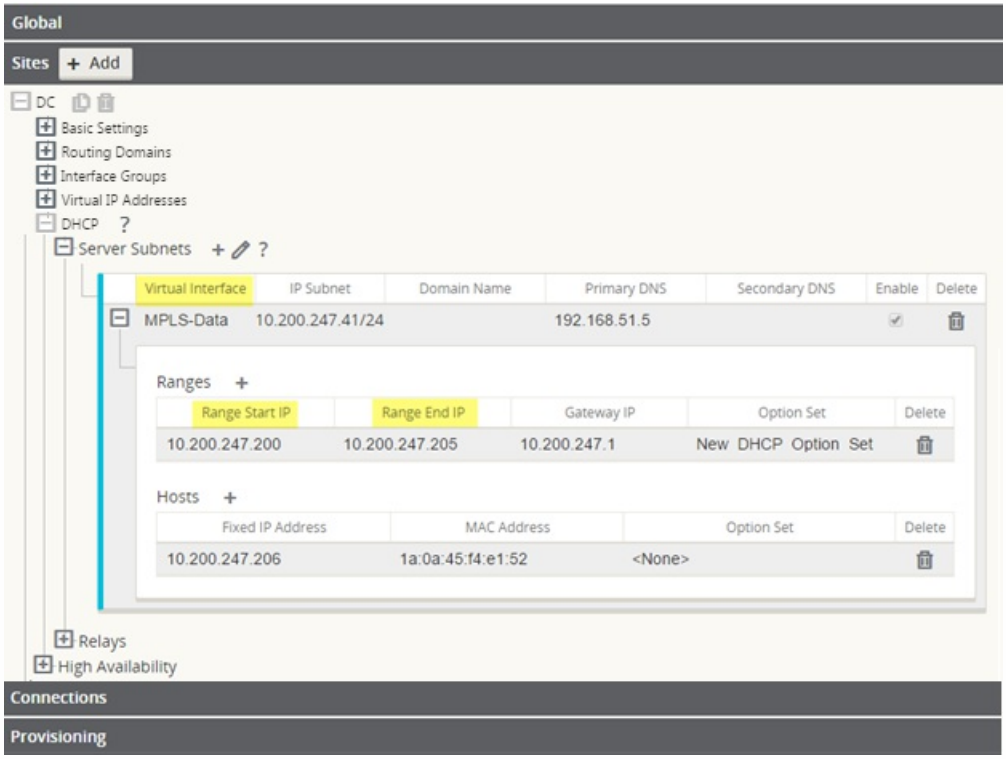

To view a list of Clients from the DHCP Server Database, in the web management interface, navigate to **Monitor** > **DHCP Server**/**Relay**.

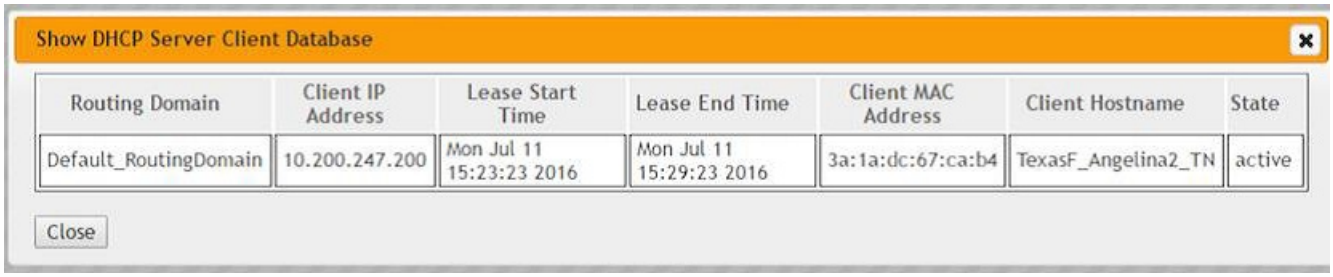

DHCP Client for Data Port (WAN Link IP Address Learning)

Aug 09, 2017

NetScaler SD-WAN appliances support WAN Link IP address learning through DHCP Clients. This functionality reduces the amount of manual configuration to deploy SD-WAN appliances and reduces ISP costs by eliminating the need to purchase static IP addresses. SD-WAN appliances can obtain dynamic IP addresses for WAN Links on untrusted interfaces eliminating the need for an intermediary WAN router to perform this function.

Note

- DHCP Client can only be configured for untrusted non-bridged interfaces configured as Client Nodes.
- DHCP Client for Data Port can be enabled only on non-MCN sites.
- One-Arm or Policy Based Routing (PBR) deployment is not supported on the site with DHCP Client configuration.
- DHCP events are logged from the client's perspective only and no DHCP server logs are generated.

How To Configure DCHP Client

To configure DHCP for an untrusted virtual interface:

1. In the **Configuration Editor**, choose **Client** from the **DHCP** drop-down menu under **[Site Name] → Interface Groups → Virtual Interfaces**.

Note

The physical interface in the interface group should be a non-bridged pair on a single interface.

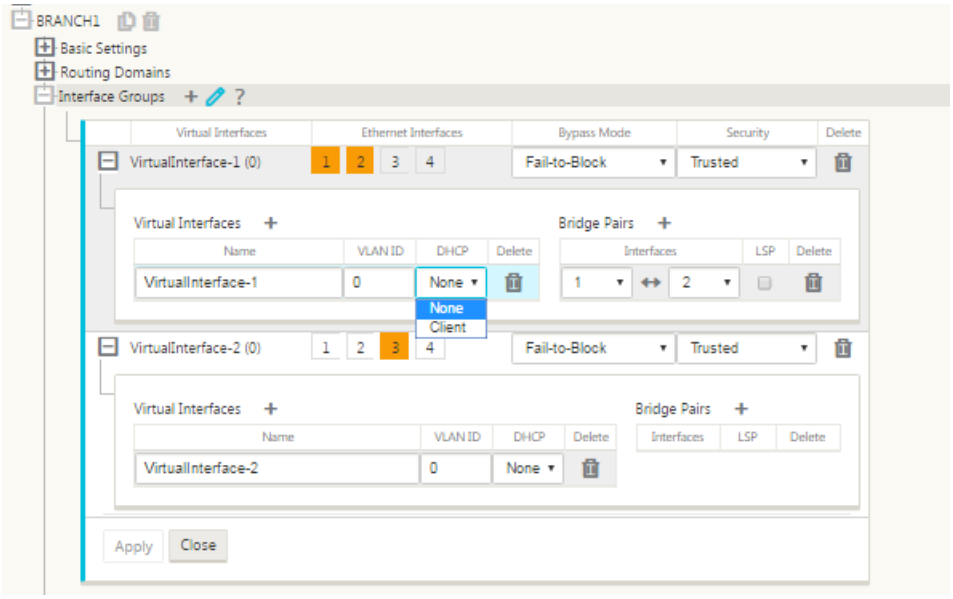

2. Navigate to **WAN Links → [WAN Link Name] → Settings → Basic Settings**.

3. Click the **Autodetect Public IP** checkbox to enable the MCN to detect the Public IP Address used by the Client. This is required when DHCP Client mode is configured for the WAN Link.

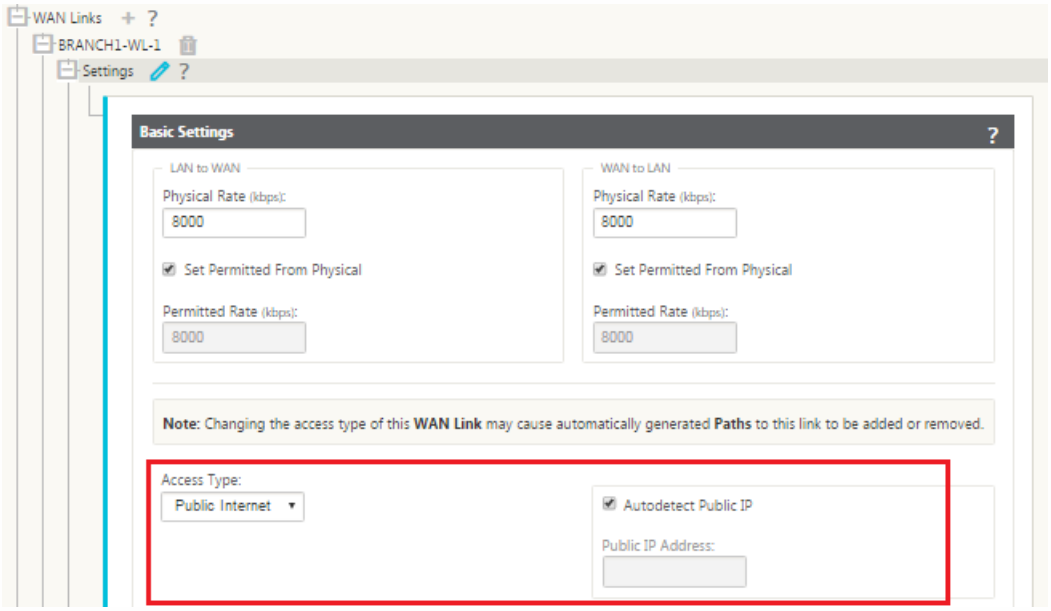

Adaptive Bandwidth Detection

Aug 09, 2017

This feature is applicable to networks with VSAT, LOS, Microwave, 3G/4G/LTE WAN Links, for which the available bandwidth varies based on weather and atmosphere conditions, location, and line of site obstructions. It allows the NetScaler SD-WAN appliances to adjust the bandwidth rate on the WAN Link dynamically based on a defined bandwidth range (minimum and maximum WAN link rate) to use the maximum amount of available bandwidth without marking the paths BAD.

- Greater bandwidth reliability (Over VSAT, Microwave, 3G/4G, and LTE)
- Greater predictability of adaptive bandwidth over user configured settings

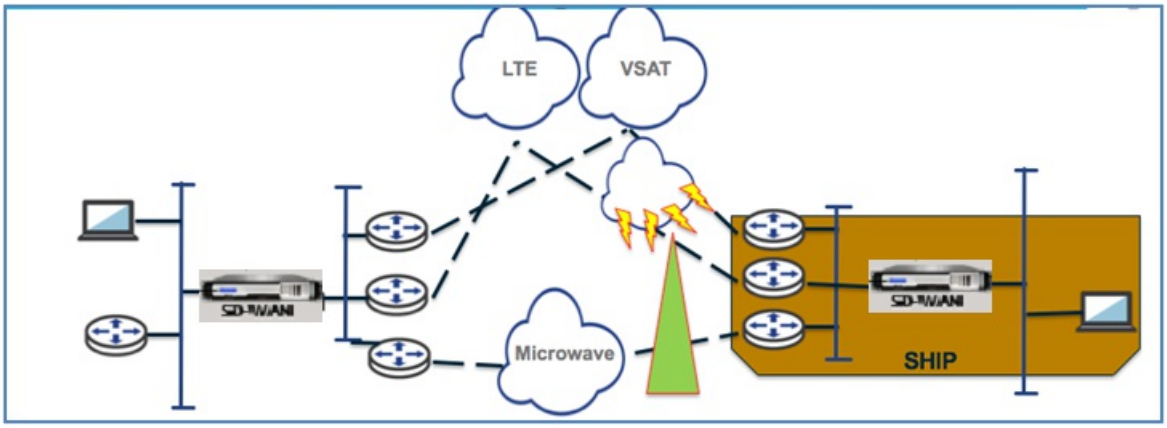

To enable adaptive bandwidth detection:

This feature needs Bad loss sensitivity option to be enabled (default/custom) as a prerequisite. You can enable it under **Global** >**Autopath Groups** > [**Autopath Group Name**] > **Bad Loss Sensitive**.

1. Enable **Adaptive Bandwidth Detection** under **Global** >**Autopath Groups** > [**Autopath Group Name**] > **Bad Loss Sensitive**.

2. Navigate to **Configuration Editor** > **Sites** > [**Site Name**] > **WAN Links** > [**WAN Link Name**] > **Settings** > **Advanced Settings**.

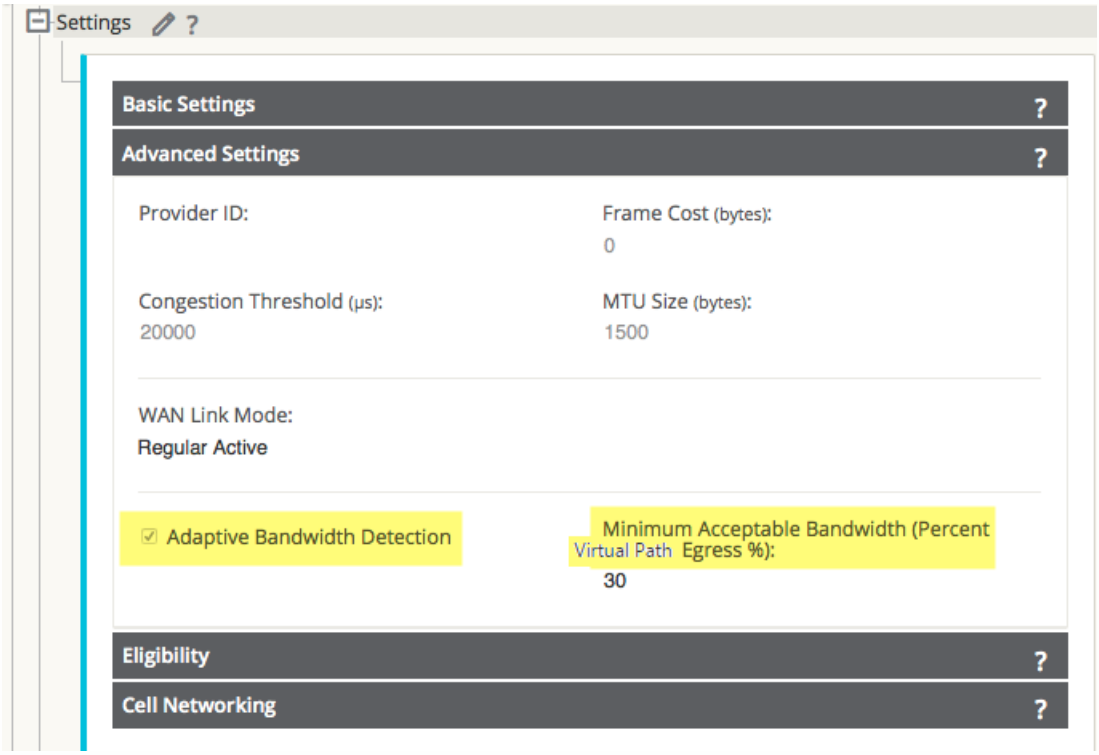

3. Check the **Adaptive Bandwidth Detection** box and enter a value in the **Minimum Acceptable Bandwidth** field.

4. View the **Usage and Permitted Rates** table by navigating to **Monitor** > **Statistics** > **WAN Link Usage** > **Usage** and **Permitted Rates**.

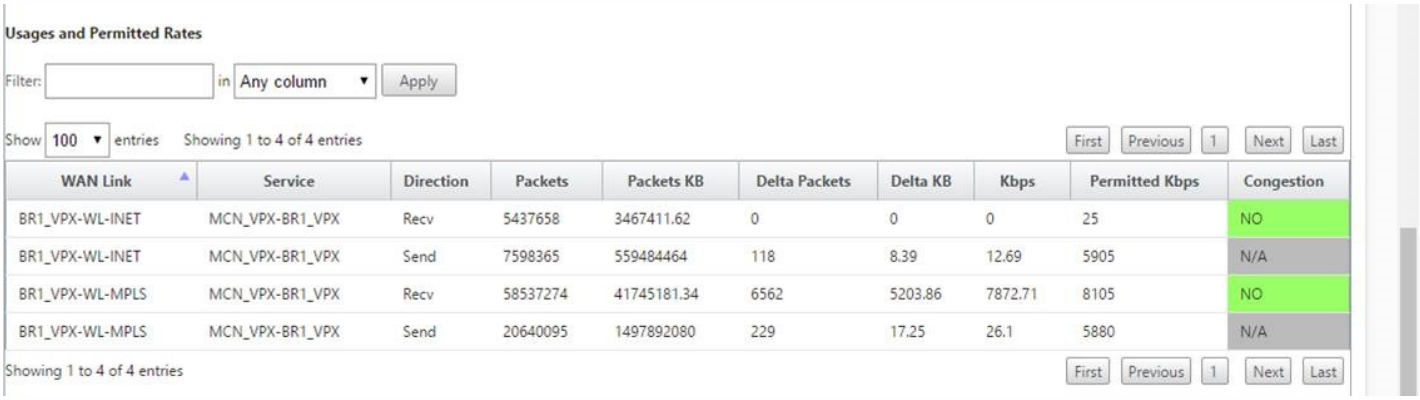

Usages and Permitted Rates

Filter: in Any column Apply

Show 100 entries Showing 1 to 4 of 4 entries First Previous 1 Next Last

WAN Link	Service	Direction	Packets	Packets KB	Delta Packets	Delta KB	Kbps	Permitted Kbps	Congestion
BR1_VPX-WL-INET	MCN_VPX-BR1_VPX	Recv	5437658	3467411.62	0	0	0	25	NO
BR1_VPX-WL-INET	MCN_VPX-BR1_VPX	Send	7598365	559484464	118	8.39	12.69	5905	N/A
BR1_VPX-WL-MPLS	MCN_VPX-BR1_VPX	Recv	58537274	41745181.34	6562	5203.86	7872.71	8105	NO
BR1_VPX-WL-MPLS	MCN_VPX-BR1_VPX	Send	20640095	1497892080	229	17.25	26.1	5880	N/A

Showing 1 to 4 of 4 entries First Previous 1 Next Last

Active Bandwidth Testing

Aug 09, 2017

Active Bandwidth Testing enables you the ability to issue an instant path bandwidth test through public internet WAN link, or to schedule public internet WAN link bandwidth testing to be completed at specific times on a recurring basis. This feature is useful for demonstrating how much bandwidth is available between two locations during new and existing installations, also for testing paths to determine the outcome of setting and confirmation changes, such as adjusting DSCP tag settings or bandwidth Permitted Rates.

To use the active bandwidth testing feature:

1. Navigate to **System Maintenance** > **Diagnostics** > **Path Bandwidth**.
2. Select the desired **Path** and click **Test**.

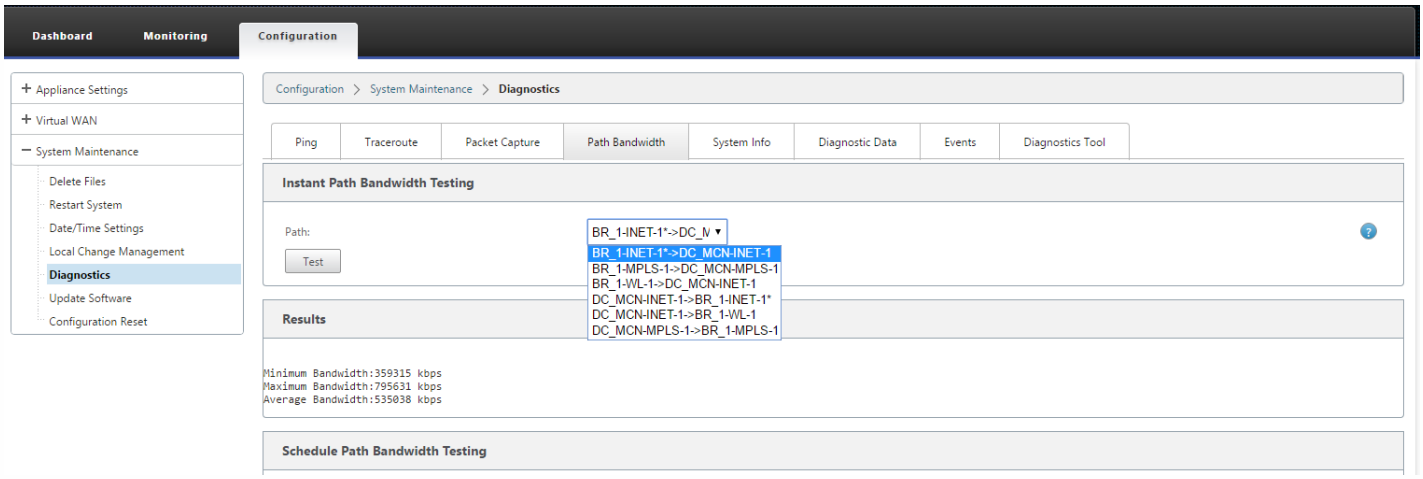

The output displays average bandwidth used as value to set as the permitted rate for the WAN Link minimum and maximum bandwidth results of the test. Along with the ability to test the bandwidth, you can now change the configuration file to use the learned bandwidth. This is accomplished through the Auto Learn option is under **Site** > [Site Name] > **WAN Links** > [WAN Link Name] > **Settings** and if enabled, the system will use the learned bandwidth.

You can also schedule recurring tests of path bandwidth in weekly, daily, or hourly intervals.

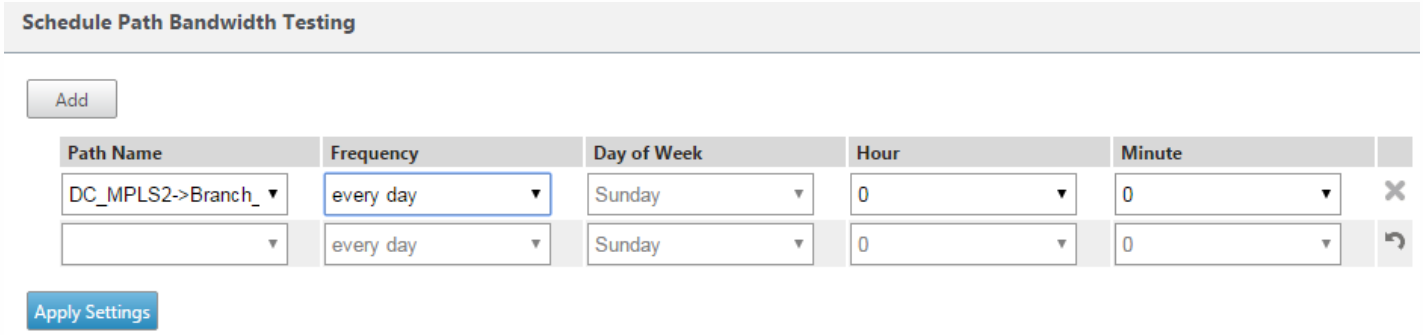

Note

A history of the path bandwidth testing results is displayed at the bottom of this page and results are archived every 7 days.

Schedule Path Bandwidth Testing

Add

Path Name	Frequency	Day of Week	Hour	Minute

Apply Settings

History Path Bandwidth Testing Result

Show 50 entries Showing 1 to 14 of 14 entries Search

First Previous 1 Next Last

Num	From Link	To Link	Test Time	Min Bandwidth (kbps)	Max Bandwidth (kbps)	Avg Bandwidth (kbps)
1	BR_1-INET-1*	DC_MCN-INET-1	3/29/2017, 1:29:54 AM	363140	780616	525927
2	BR_1-INET-1*	DC_MCN-INET-1	3/29/2017, 1:30:00 AM	281995	573073	430345
3	BR_1-INET-1*	DC_MCN-INET-1	3/29/2017, 1:30:06 AM	317568	636640	480818
4	BR_1-MPLS-1	DC_MCN-MPLS-1	3/29/2017, 1:34:00 AM	440056	1083357	725514
5	DC_MCN-INET-1	BR_1-INET-1*	3/29/2017, 1:34:10 AM	506768	786784	638673
6	DC_MCN-INET-1	BR_1-INET-1*	3/29/2017, 1:34:18 AM	462584	1388712	669232
7	DC_MCN-INET-1	BR_1-WL-1	3/29/2017, 1:34:27 AM	380679	727895	533286
8	DC_MCN-MPLS-1	BR_1-MPLS-1	3/29/2017, 1:35:12 AM	26823	35495	30578
9	DC_MCN-INET-1	BR_1-INET-1*	3/29/2017, 1:36:09 AM	350097	733929	591542
10	DC_MCN-INET-1	BR_1-INET-1*	3/29/2017, 1:36:47 AM	476024	789756	639048
11	DC_MCN-INET-1	BR_1-WL-1	3/29/2017, 1:36:56 AM	446292	777674	608533

Diagnostic Tool

Aug 09, 2017

A new inbuilt traffic generator is introduced in SD-WAN 9.2. This diagnostic tool is used to generate test traffic which allows you to troubleshoot network issues that may results in:

- Frequent change in path state from Good to Bad.
- Poor application performance.
- Higher packet loss

Most often, these problems arise due to rate limiting configured on firewall / router, incorrect bandwidth settings, low link speed, priority queue set by network provider and so on. The diagnostic tool allows you to identify the root cause of such issues and troubleshoot it.

The diagnostic tool removes the dependency on third party tools such as iPerf which has to be manually installed on the Data Center and Branch hosts. It provides more control over the type of diagnostic traffic sent, the direction in which the diagnostic traffic flows, and the path on which the diagnostic traffic flows.

The diagnostic tool allows to generate the following two types of traffic:

- Control: Generates traffic with no QOS/scheduling applied to the packets. As a result, the packets are sent over the path selected in the UI, even if the path is not the best at the time. This traffic is used to test specific paths and helps to identify ISP related issues. You can also use this to determine the bandwidth of the selected path.
- Data: Simulates the traffic generated from the host with SD-WAN traffic processing. Since QoS/scheduling is applied to the packets, the packets are sent over the best path available at that time. Traffic will be sent over multiple paths if load balancing is enabled. This traffic is used to troubleshoot QoS/scheduler related issues.

> ## Note
>
> To run a diagnostic test on a path, you have to start the test on the appliances at both ends of the path. Start the diagnostic test as a server on one appliance and as a client on the other appliance.

To use diagnostics tool:

1. On both the appliances, click **Configuration** > **System Maintenance** > **Diagnostics** > **Diagnostics Tool**.

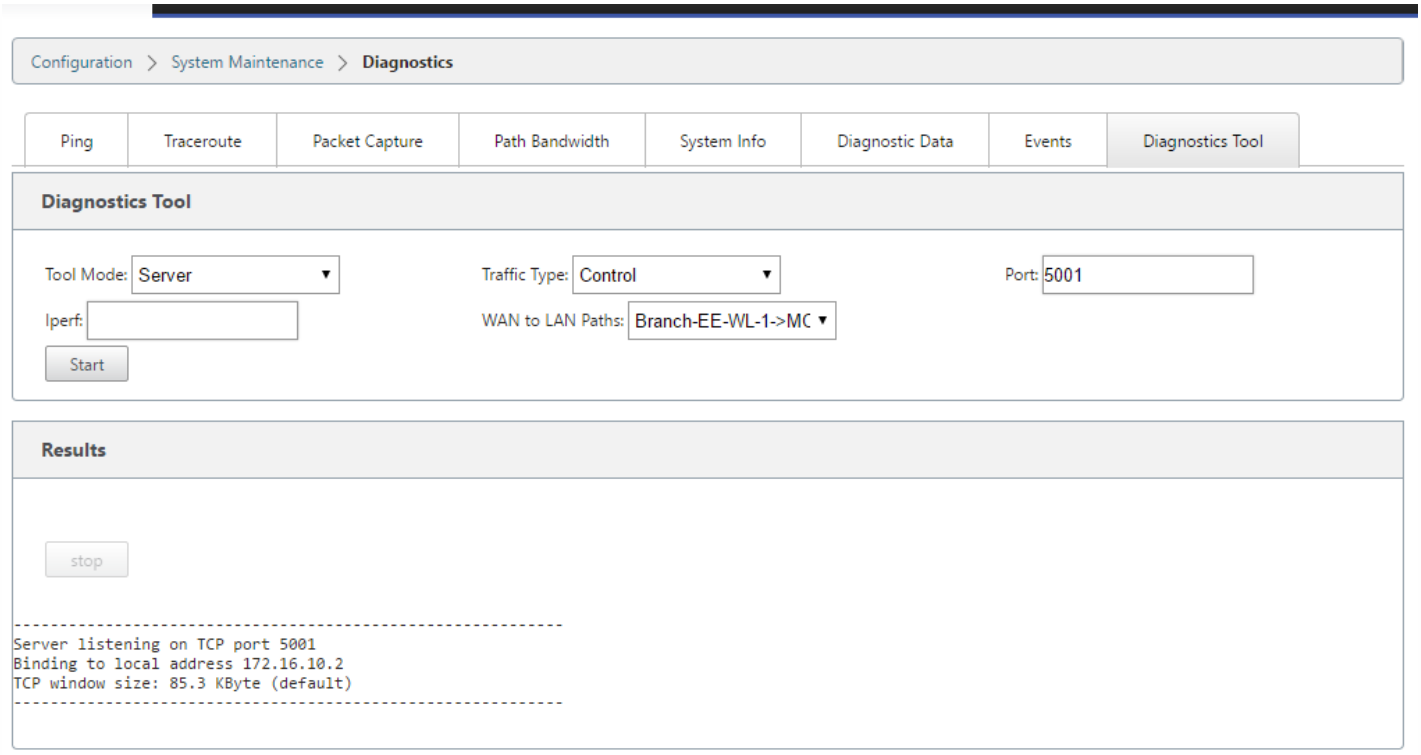

2. In the **Tool Mode** field, select **Server** on one appliance and select **Client** on the other appliance.

3. In the **Traffic Type** field, select the type of diagnostic traffic, either **Control** or **Data**. Select the same traffic type on both the appliances.

4. In the **Port** field, specify the **TCP / UDP** port number on which the diagnostic traffic will be sent. Specify the same port number on both the appliances.

5. In the **Iperf** field, specify IPERF command line options, if any.

Note

You need not specify the following IPERF command line options:

- -c: Client mode option is added by the diagnostic tool.
- -s: Server mode option is added by the diagnostic tool.
- -B: Binding IPERF to specific IP/interface is done by the diagnostic tool depending on the path selected.
- -p: Port number is provided in the diagnostics tool.

6. Select the path on which you want to send the diagnostic traffic. Select the same path on both the appliances.

7. Click **Start** on both the appliances.

Monitoring Your Virtual WAN

Aug 09, 2017

Viewing Basic Information for an Appliance

Use a browser to connect to the Management Web Interface of the appliance you want to monitor, and click the **Dashboard** tab to display basic information for that appliance.

The **Dashboard** page displays the following basic information for the local appliance:

System Status:

- **Name** – This is the name you assigned to the appliance when you added it to the system.
- **Model** – This is the Virtual WAN appliance model number.
- **Appliance Mode** – This indicates whether this appliance has been configured as the primary or secondary MCN, or as a client appliance.
- **Management IP Address –** This is the Management IP Address for the appliance.
- **Appliance Uptime –** This specifies the duration for which the appliance has been running since the last reboot.
- **Service Uptime** – This specifies the duration for which the Virtual WAN Service has been running since the last restart.

Virtual Path Service Status:

- **Virtual Path [site name]** – This displays the current status of all the Virtual Paths associated with this appliance. If the Virtual WAN Service is enabled, this section is included on the page. If the Virtual WAN Service is disabled, an Alert icon (goldenrod delta) and Alert message to that effect displays in place of this section.

Local Version Information:

- **Software version –** This is the version of the CloudBridge Virtual Path software package currently activated on the appliance.
- **Build on** – This is the build date for the product version currently running on the local appliance.
- **Hardware version** – This is the hardware model number and version of the appliance.
- **OS Partition Version –** This is the version of the OS partition currently active on the appliance.

  The below figure shows a sample Dashboard page for the MCN, and MCN Appliance information.

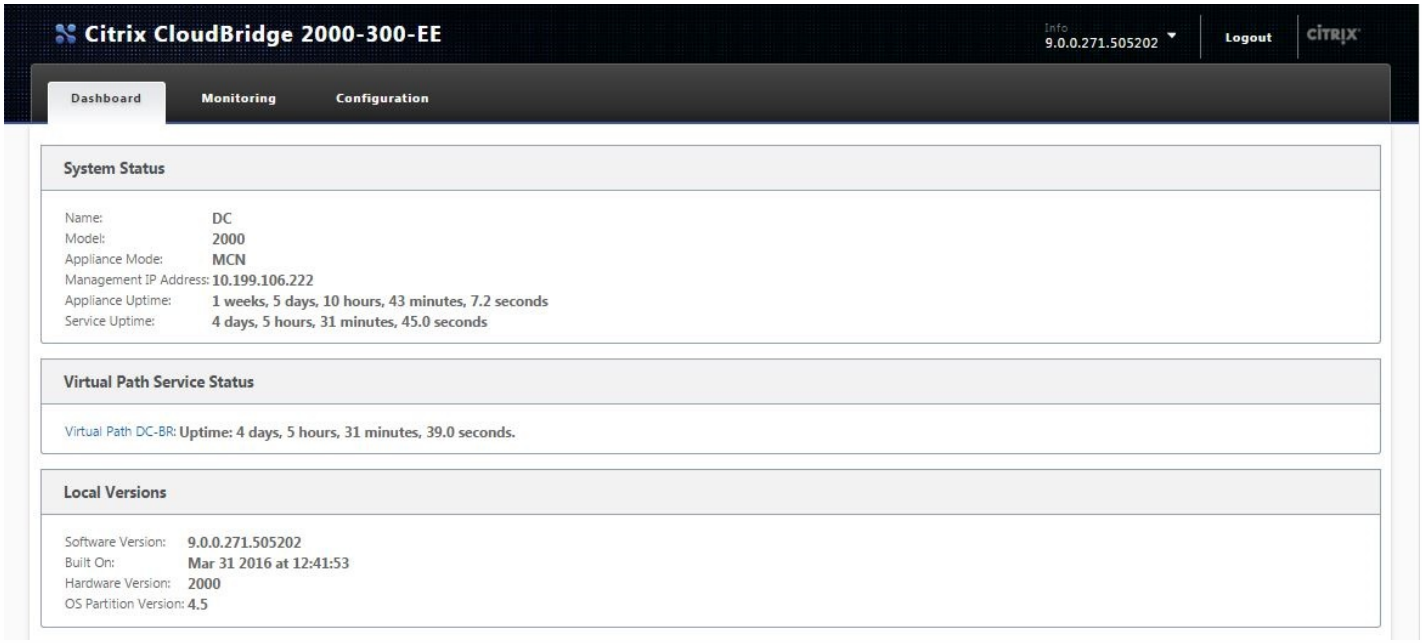

The below figure shows a sample Dashboard page and information for a client appliance.

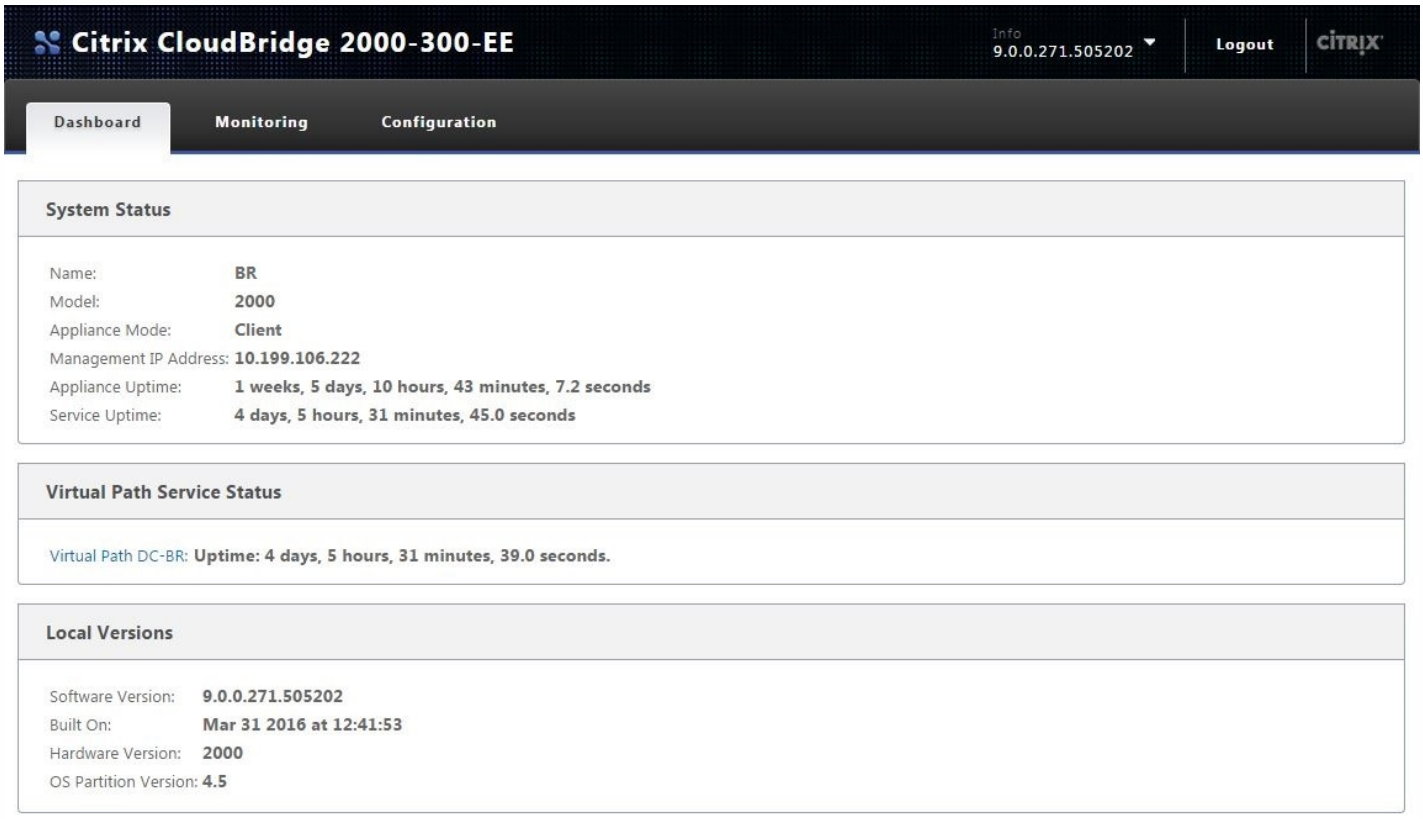

Viewing Statistical Information

Aug 09, 2017

This section provides basic instructions for viewing Virtual WAN statistics information.

1. Log onto the Management Web Interface for the MCN.
2. Select the **Monitoring** tab.

This opens the **Monitoring** navigation tree in the left pane. By default, this also displays the **Statistics** page with **Paths** preselected in the **Show** field. This contains a detailed table of path statistics.

> ## Note
>
> If you navigate to another **Monitoring** page (for example, **Flows**), you can return to this page by selecting **Statistics** in the **Monitoring** navigation tree (left pane).

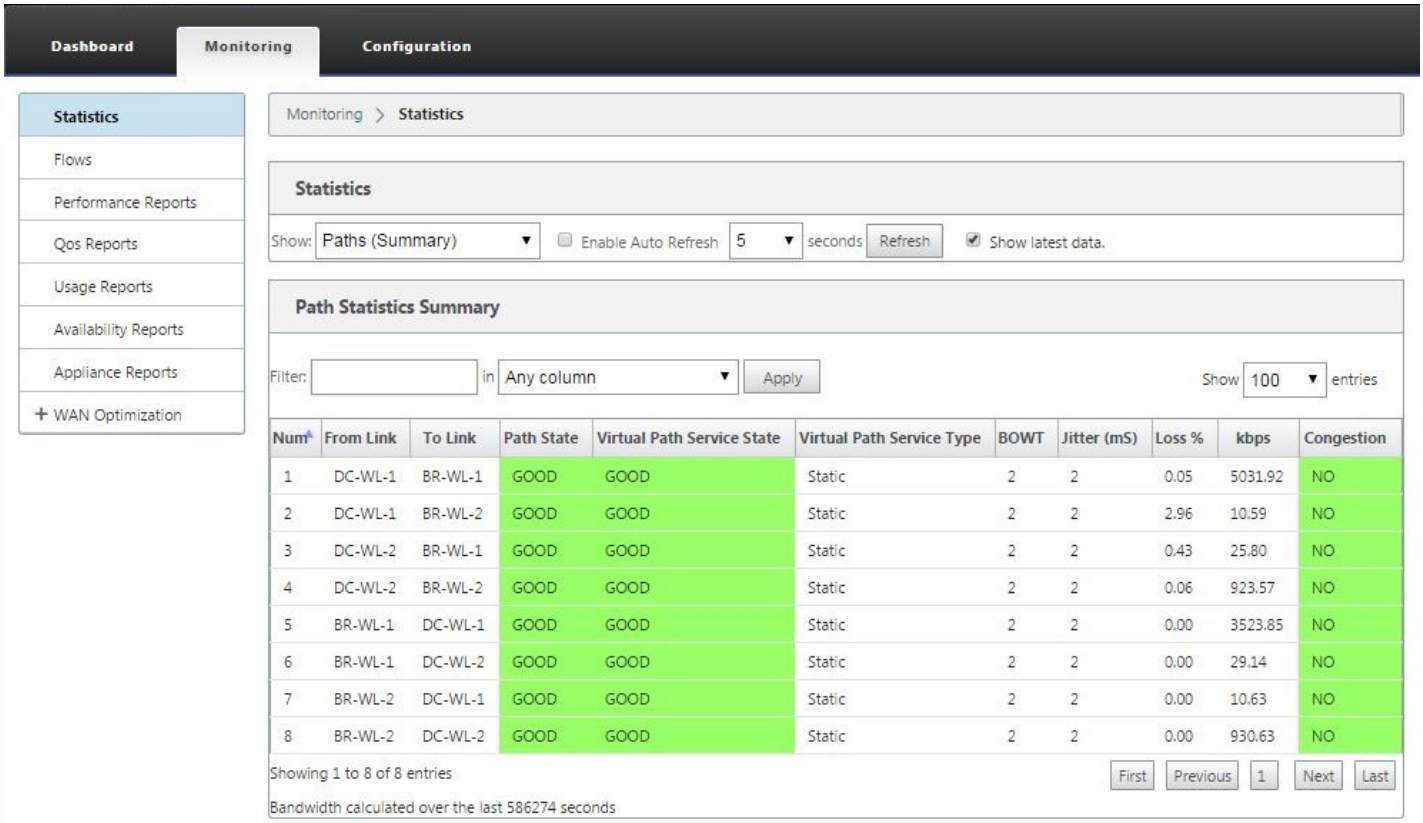

3. Open the **Show** drop-down menu next to the **Show** field.

In addition to the **Paths** statistics, the **Show** menu also offers several additional options for filtering and viewing statistical information.

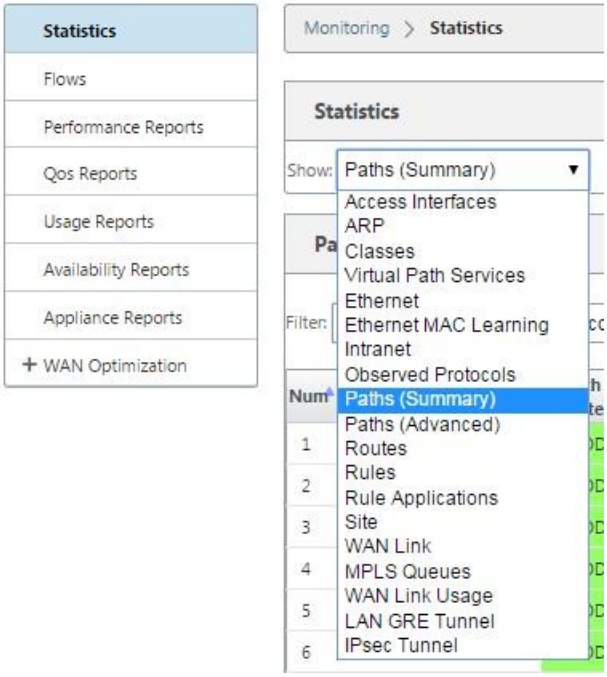

4. Select a filter from the **Show** menu to view a table of statistical information for that topic.

Viewing Flow Information

Aug 09, 2017

This section provides basic instructions for viewing Virtual WAN flow information.

To view flow information, do the following:

1. Log onto the Management Web Interface for the MCN, and select the **Monitoring** tab.

   This opens the **Monitoring** navigation tree in the left pane.

2. Select the **Flows** branch in the navigation tree.

   This displays the **Flows** page with **LAN** to **WAN** preselected in the **Flow Type** field.

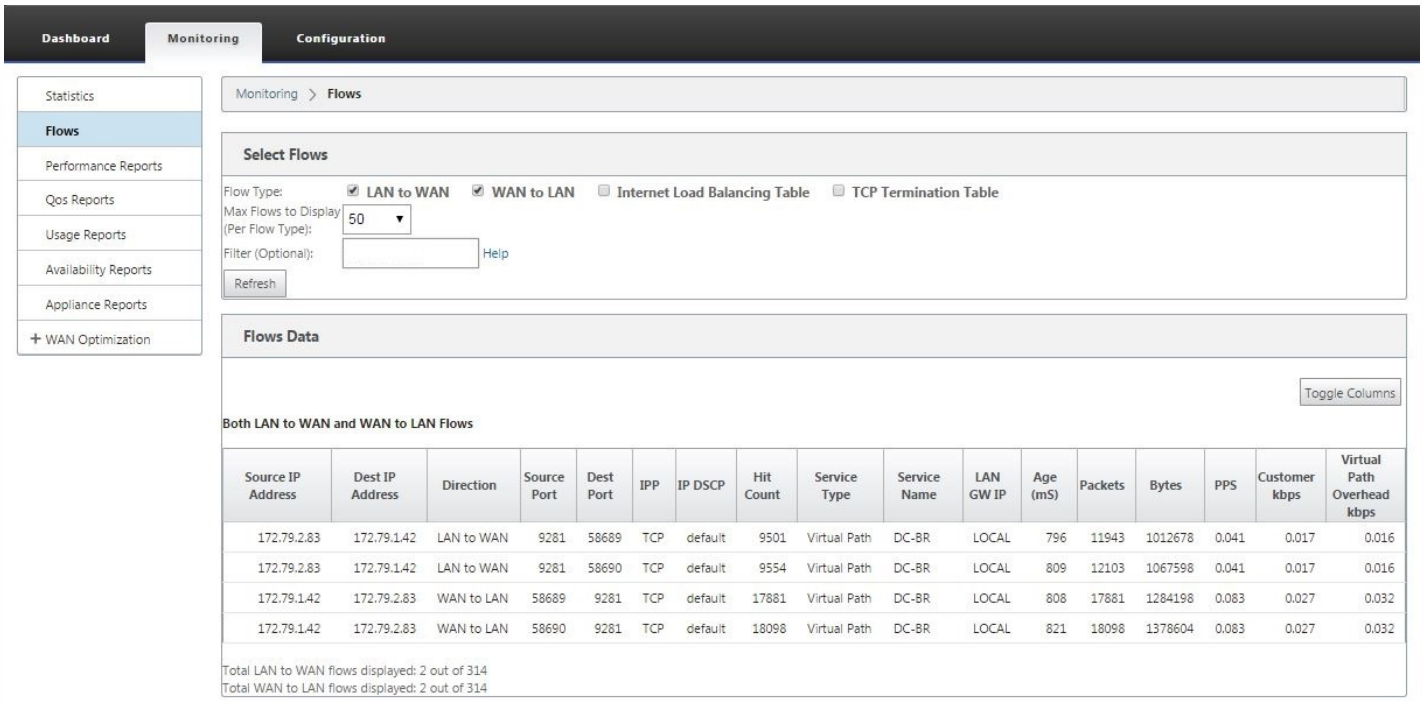

3. Select the **Flow Type**.

The **Flow Type** field is located in the **Select Flows** section at the top of the **Flows** page. Next to the **Flow Type** field is a row of checkbox options for selecting the flow information you want to view. You can check one or more boxes to filter the information to be displayed.

4. Select the **Max Flows to Display** from the drop-down menu next to that field.

   This determines the number of entries to display in the **Flows** table. The options are: **50**, **100**, **1000**.

5. (Optional) Enter search text in the **Filter** field.

   This filters the table results so that only entries containing the search text display in the table.

Tip

To see detailed instructions for using filters to refine **Flow** table results, click **Help** to the right of the **Filter** field. To close the help display, click **Refresh** in the bottom left corner of the **Select Flows** section.

6. Click **Refresh** to display the filter results.

The below figure shows a sample **Flows** page filtered display with all flow types selected.

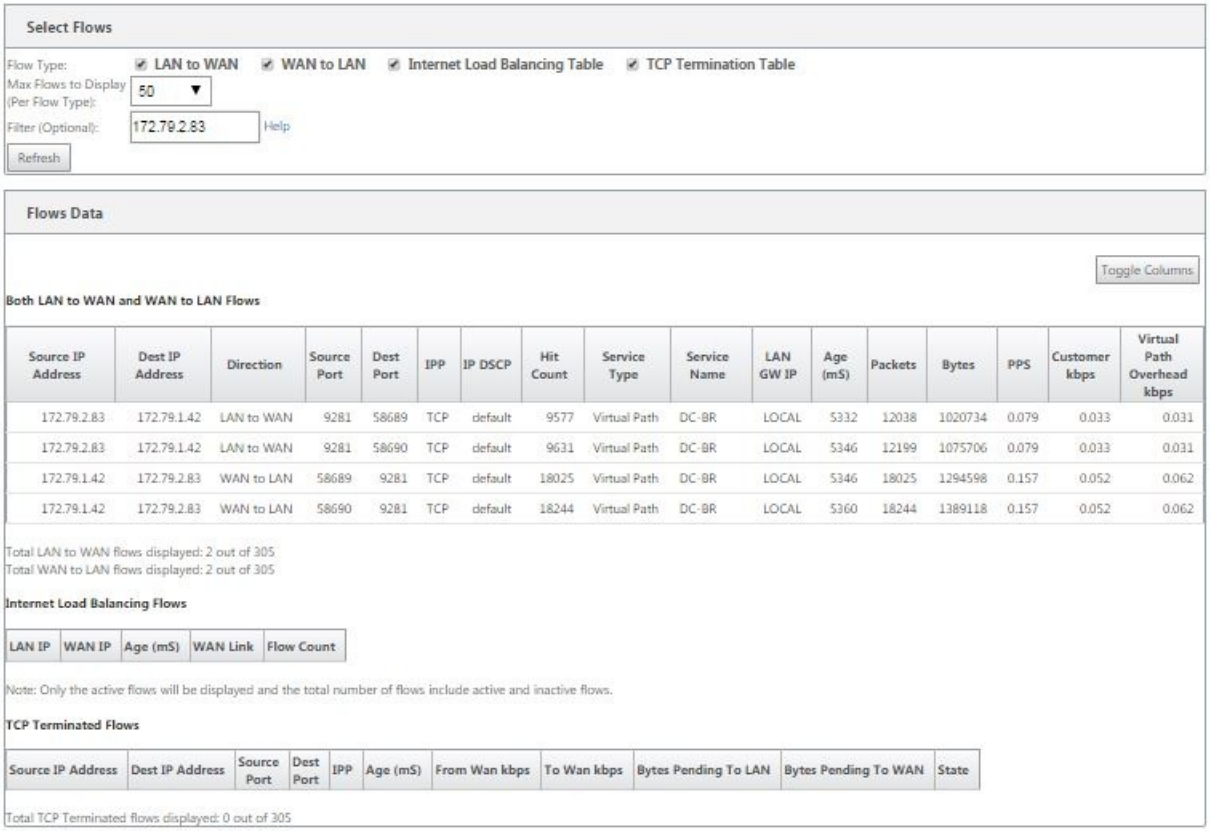

7. (Optional) Select the columns to include in the table.

Do the following:

a. Click **Toggle** Columns.

The **Toggle Columns** button is just above the top right corner of the **Flows** table. This reveals any deselected columns, and opens a checkbox above each column for selecting or deselecting that column. Deselected columns display greyed out, as shown in the below figure.

Note

By default, all of the columns are selected, which can cause the table to be truncated in the display, obscuring the **Toggle Columns** button. If so, a horizontal scroll bar displays beneath the table. Slide the scroll bar to the right to view the truncated section of the table and reveal the **Toggle Columns** button. If the scroll bar is not available, try resizing the width of your browser window until the scroll bar is revealed.

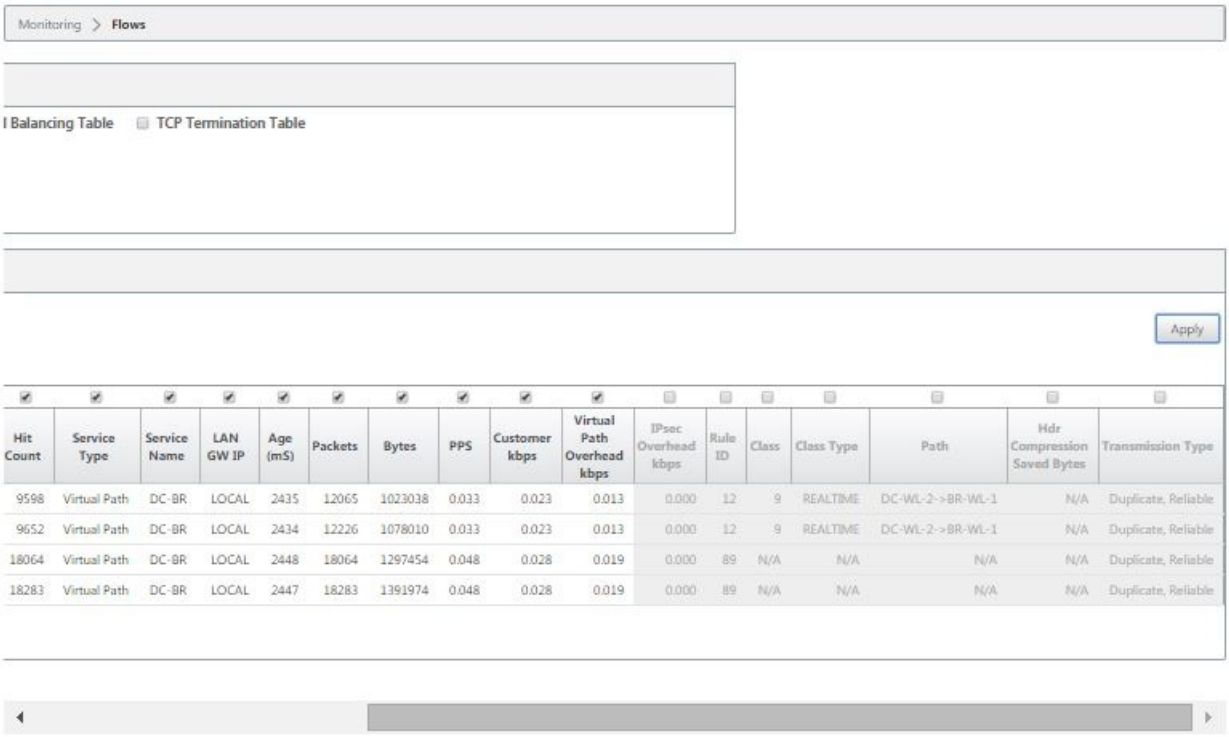

b. Click a checkbox to select or deselect a column.

c. Click **Apply** (above the top right corner of the table).

This dismisses the selection options, and refreshes the table to include only the selected columns.

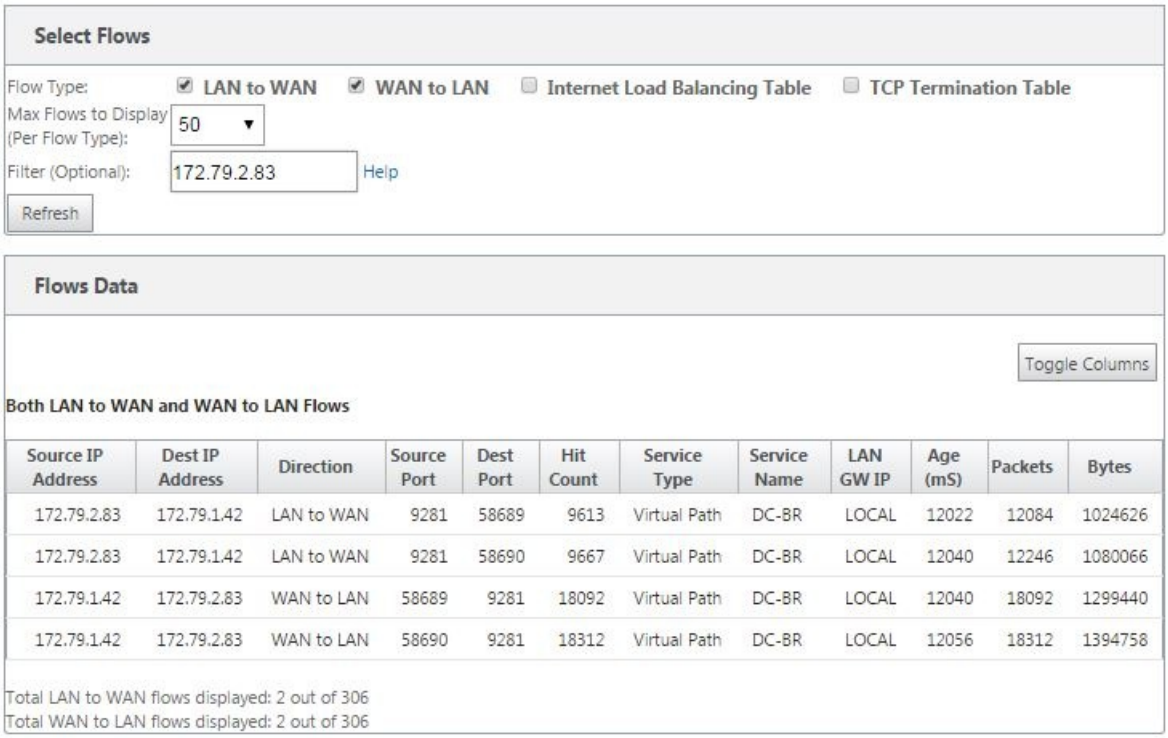

Viewing Reports

Aug 09, 2017

This section provides basic instructions for generating and viewing CloudBridge Virtual WAN reports about the local appliance using the Management Web Interface.

> **Note**
>
> Reports generated on the Management Web Interface apply to the local appliance, only. To generate and view reports for the Virtual WAN, use the Virtual WAN Center Web Interface.

To generate and view CloudBridge Virtual WAN reports, do the following:

1. Log onto the Management Web Interface for the MCN, and select the **Monitoring** tab.

   This opens the **Monitoring** navigation tree in the left pane.

2. Select a report type from the navigation tree.

   The report types are listed as branches in the navigation tree, just below the **Flows** branch.

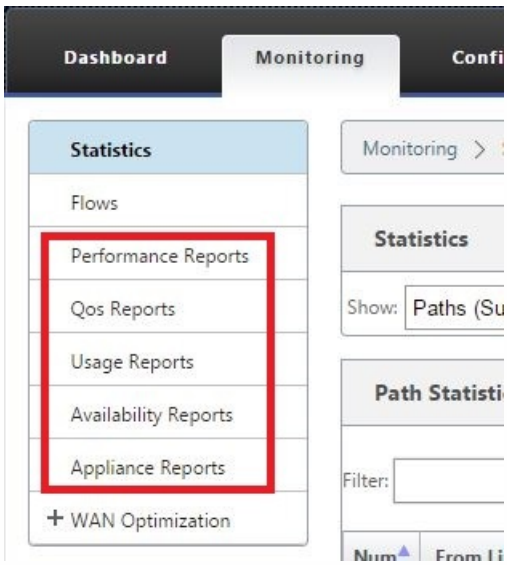

   The available report types are as follows:

   - **Performance Reports**
   - **QoS Reports**
   - **Usage Reports**
   - **Availability Reports**
   - **Appliance Reports**

3. Select the report options.

In addition to the various types of reports, for each report type there are numerous options and filters for refining report results.

Viewing Firewall Statistics

Aug 09, 2017

Once you have configured firewall and NAT policies you can view the statistics of the connections, firewall policies and NAT policies as reports. You can filter the reports using the various filtering parameters.

For information on configuring firewall and NAT policies, see Stateful Firewall and NAT Support.

Connections

You can check the statistics for Applications for the Firewall Policy. This enables you to see all connections that match to the selected Application, where they are coming from, where they are going to, and how much traffic they are generating. You can see how the firewall policies are acting on the traffic for each Application.

You can filter the connections statistics using the following parameters:

- Application - The application used as filter criteria for the connection.
- Family - The application family the used as filter criteria for the connection.
- IP Protocol - The IP protocol used by the connection.
- Source Zone - The zone from which the connection originated.
- Destination Zone - The zone from which responding traffic originates.
- Source Service Type - The service from which the connection originated.
- Source Service Instance - The instance of the service from which the connection originated.
- Source IP - The IP address from which the connection originated, input in dotted decimal notation with an optional subnet mask.
- Source Port - The port or range of ports from which the connection originated. A single port or a range of ports using the "-" character is accepted.
- Destination Service Type - The service from which responding traffic originates.
- Destination Service Instance - The instance of the service from which responding traffic originates.
- Destination IP - The IP address of the responding device, input in dotted decimal notation with an optional subnet mask.
- Destination Port - The port or range of ports used by the responding device. A single port or a range of ports using the "-" character is accepted.

Filter Policies

Policies enable you to specify actions for traffic flows. Group of firewall filters are created using Firewall Policy Templates and can be applied to all sites in the network or only to specific sites.

You can view statistics report for all the filter policies and filter it using the following parameters.

- Application object - The Application object used as a filter criteria in the firewall policy.
- Application - The application used as a filter criteria in the firewall policy
- Family - The application family used as filter criteria in the firewall policy.
- IP Protocol - The IP protocol that the filter policy matches.
- DSCP: The DSCP tag that the filter policy matches.
- Filter Policy Action - The action taken by the policy when a packet matches the filter.
- Source Service Type - The service from which the connection originated.
- Source Service Name - The instance of the service from which the connection originated.
- Source IP - The IP address from which the connection originated, input in dotted decimal notation with an optional

subnet mask.

- Source Port - The port or range of ports from which the connection originated. A single port or a range of ports using the "-" character is accepted.
- Destination Service Type - The service to which responding traffic is destined.
- Destination Service Name - When applicable, the service to which responding traffic is destined.
- Destination IP - The IP address of the responding device, input in dotted decimal notation with an optional subnet mask.
- Destination Port - The port or range of ports used by the responding device. A single port or a range of ports using the "-" character is accepted.
- Source Zone - The origination zone matched by the filter policy.
- Destination Zone - The responding zone matched by the filter policy.

NAT Policies

You can view the statistics of all the Network Address Translation (NAT) policies and filter the report using the following parameters.

- IP Protocol - The IP protocol that the NAT policy matches.
- NAT Type - The type of NAT in use by the NAT policy.
- Dynamic NAT Type - The type of Dynamic NAT in use by the NAT policy.
- Service Type - The service type used by the NAT policy.
- Service Name - The instance of the service used by the NAT policy.
- Inside IP - The inside IP address, input in dotted decimal notation with an optional subnet mask.
- Inside Port- The inside port range used by the NAT policy. A single port or a range of ports using the "-" character is accepted.
- Outside IP - The outside IP address, input in dotted decimal notation with an optional subnet mask.
- Outside Port - The outside port range used by the NAT policy. A single port or a range of ports using the "-" character is accepted.

To view Firewall Statistics:

1. Navigate to **Monitoring** > **Firewall.**
2. In the Statistics field select, **Connections**, **Filter Policies** or **NAT Policies** as required.
3. Set the filtering criteria as require.

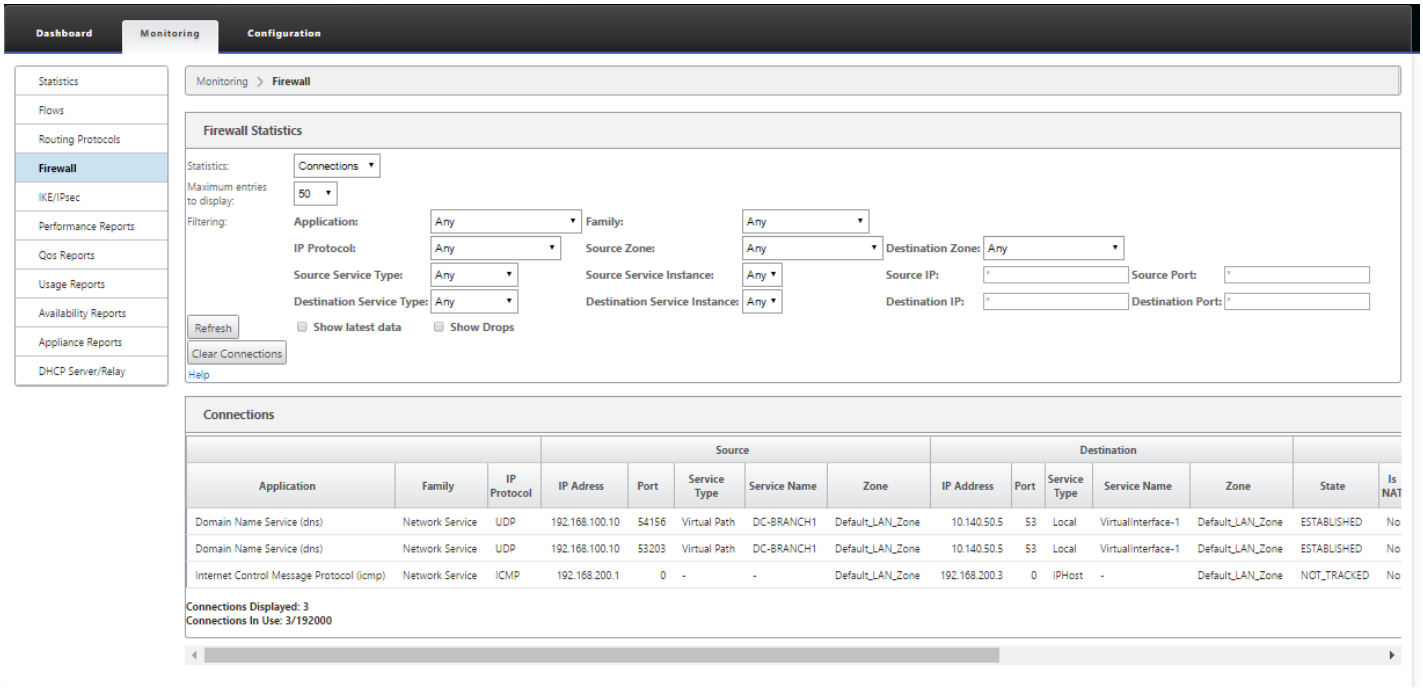

4. Click **Refresh**.

Auto Secure Peering and Manual Secure Peering

Aug 09, 2017

Enterprise Edition appliance can be installed at the data center and has the capability to initiate auto or manual secure peering, create SSL profile and associate service class, and join the appliance to a Windows Domain Controller for allowing users/administrator to make use of the extended rich feature of standalone WANOP appliance.

Following are the deployment modes supported for Auto Secure Peering and Manual Secure Peering:

Auto Secure Peering Deployments

1. To perform auto secure peering to an EE appliance from a standalone WANOP / SDWAN SE/WANOP on the DC site.

   Steps to initiate this deployment:

   - WANOP DC appliance is in LISTEN ON mode (2312/Any non-standard port) and Branch EE is in CONNECT-TO mode.
   - WANOP DC initiates automatic secure peering to an EE appliance which installs the Private CA Certs and CERT KEY Pairs and configure CONNECT-TO on the EE appliance with WANOPs LISTEN-ON IP.

2. To perform Auto-secure peering initiated from EE appliance at DC site and Branch site EE appliance.

   Steps to initiate this deployment:

   - EE DC appliance is in LISTEN ON mode (on port 443). Branch EE is in CONNECT-TO mode.
   - EE DC appliance initiates automatic secure peering to an EE Branch appliance which installs the Private CA Certs and CERT KEY Pairs and configures CONNECT-TO on the EE Branch appliance with DC EE's LISTEN-ON IP.
   - LISTEN-ON IP for EE is in the interface IP associated to the routing domain for which "Redirect to WANOP" is enabled.

3. Auto Secure Peering initiated from EE Appliance at DC site and Branch with WANOP/ SDWAN SE appliance.

   Steps to initiate this deployment:

   - EE DC appliance is in LISTEN ON mode (on port 443). Branch WANOP / SDWAN SE is in CONNECT-TO mode.
   - EE DC appliance initiates automatic secure peering to Branch WANOP / SDWAN SE appliance which installs the Private CA Certs and CERT KEY Pairs and configures CONNECT-TO on the EE appliance with DC EE's LISTEN-ON IP.

Manual Secure Peering Deployments

4. Manual Secure Peering initiated from EE appliance at DC site to Branch EE Appliance.

   Steps to initiate this deployment:

   - EE DC appliance is in LISTEN ON mode (on port 443). Branch EE is in CONNECT-TO mode.
   - LISTEN-ON IP for EE is in the interface IP associated to the routing domain for which "Redirect to WANOP"

is enabled.

- Manually upload CA and Cert Key pair certificates obtained from authentic source of certificate authority.

5. Manual Secure Peering initiated from EE appliance at DC site to Branch WANOP/SDWAN-SE Appliance.

Steps to initiate this deployment:

- EE DC appliance is in LISTEN ON mode (on port 443). Branch WANOP / SDWAN SE is in CONNECT-TO mode.
- LISTEN-ON IP for EE is in the interface IP associated to the routing domain for which “Redirect to WANOP” is enabled
- Manually upload CA and Cert Key pair certificates obtained from authentic source of certificate authority.

Auto Secure Peering to an EE appliance from a Standalone WANOP / SDWAN SE/WANOP on the DC site

Aug 09, 2017

Configuration

To perform auto secure peering an EE appliance from a standalone WANOP on the DC Side:

1. On a standalone WANOP appliance at the data center, click **Secure** in the **Secure Peering** pane of the **Secure Acceleration** page.

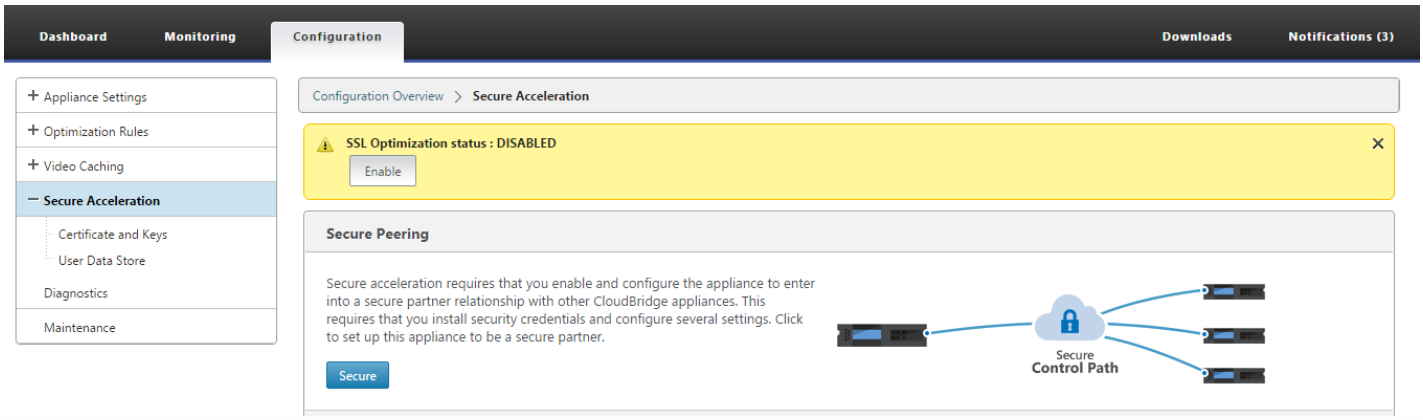

2. Configure the keystore settings by providing the **keystore password** or by disabling the keystore.

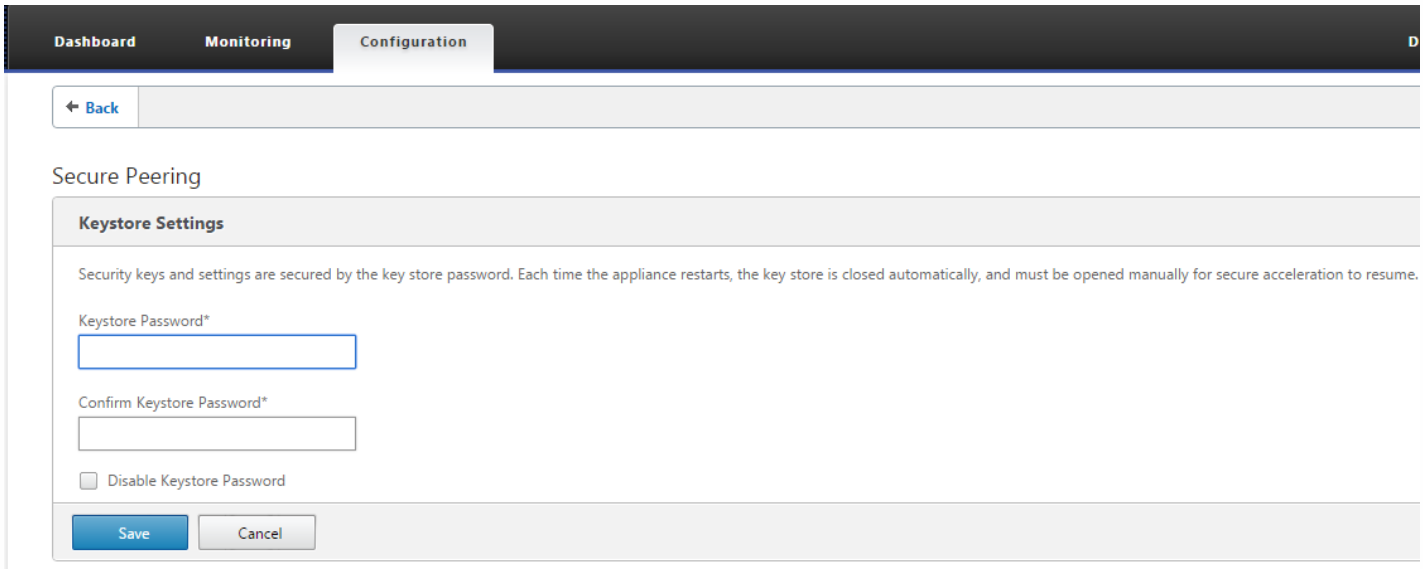

3. **Enable Secure Peering** by selecting **Private CA** to perform AUTOMATIC SECURE PEERING.

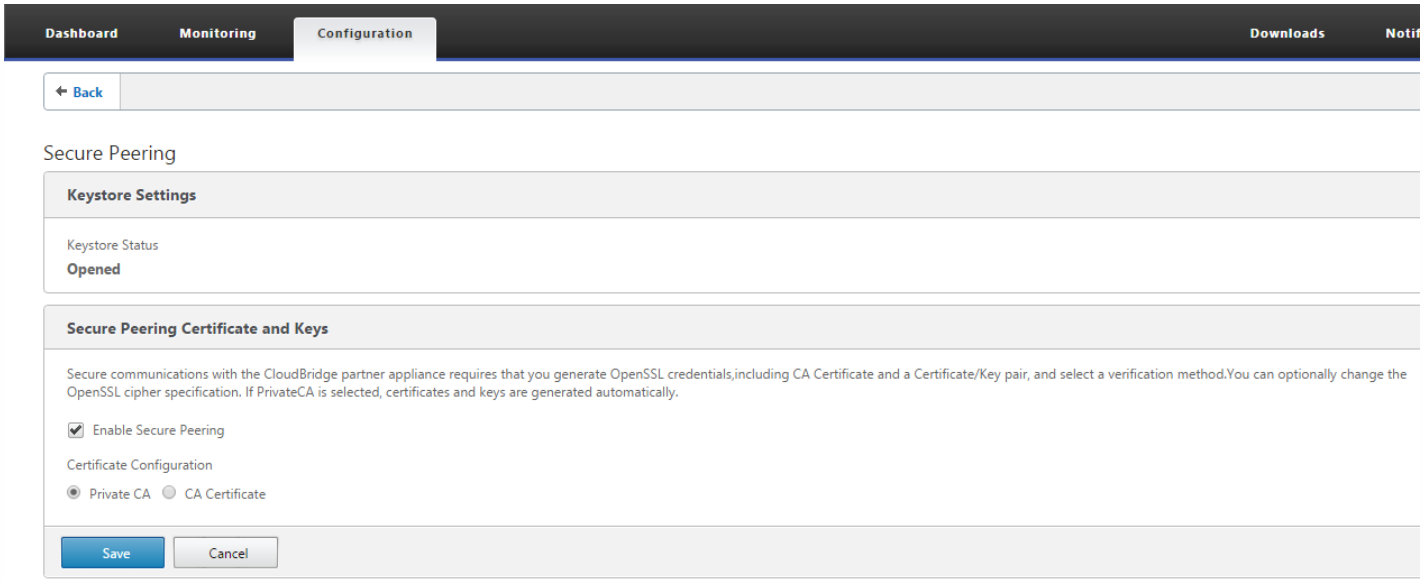

4. The appliance level CA certificate and private Certificate and Key will be generated on the local WANOP and a table to add a REMOTE PEER TO Perform AUTO secure peering with is displayed.

5. Click on the '**+**' icon and a popup window to add IP address with username and password is dispalyed. After successful authentication with the remote IP with credentials provided, a request is sent to the remote machine that installs CA Certificate and the Private certificate and key for itself locally (on the remote machine).

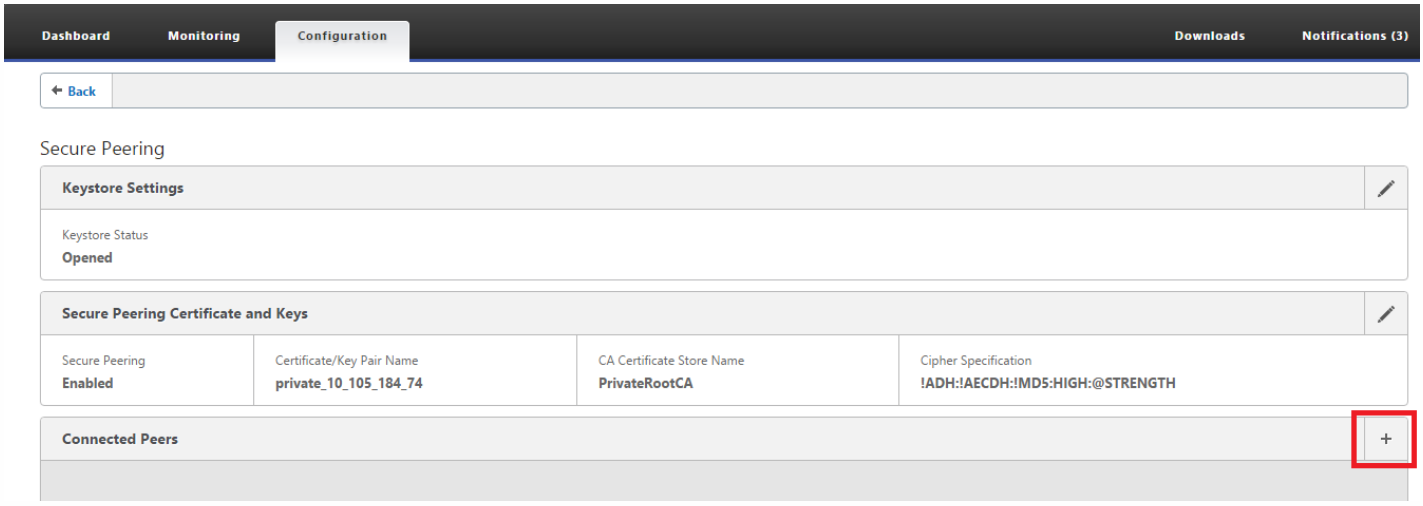

Note

- IP Address – IP Address of remote ENTERPRISE EDITION APPLIANCE MANAGEMENT IP
- Username – Username of remote ENTERPRISE EDITION APPLIANCE
- Password – Password of remote ENTERPRISE EDITION APPLIANCE

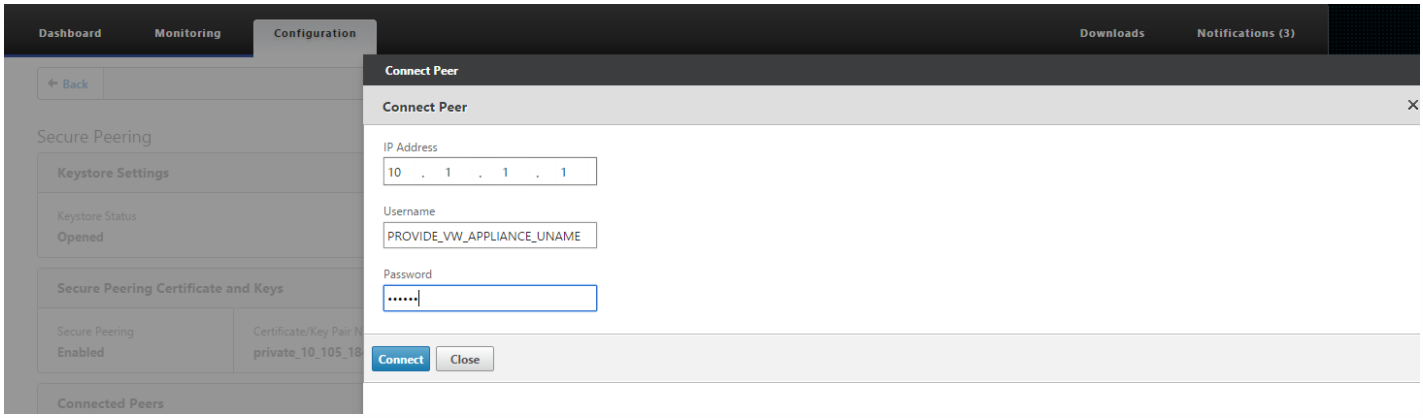

After Successful Authentication, you will see Secure Peering as TRUE and the partner IP address as one of the Virtual IP address of the remote Enterprise Edition Appliance.

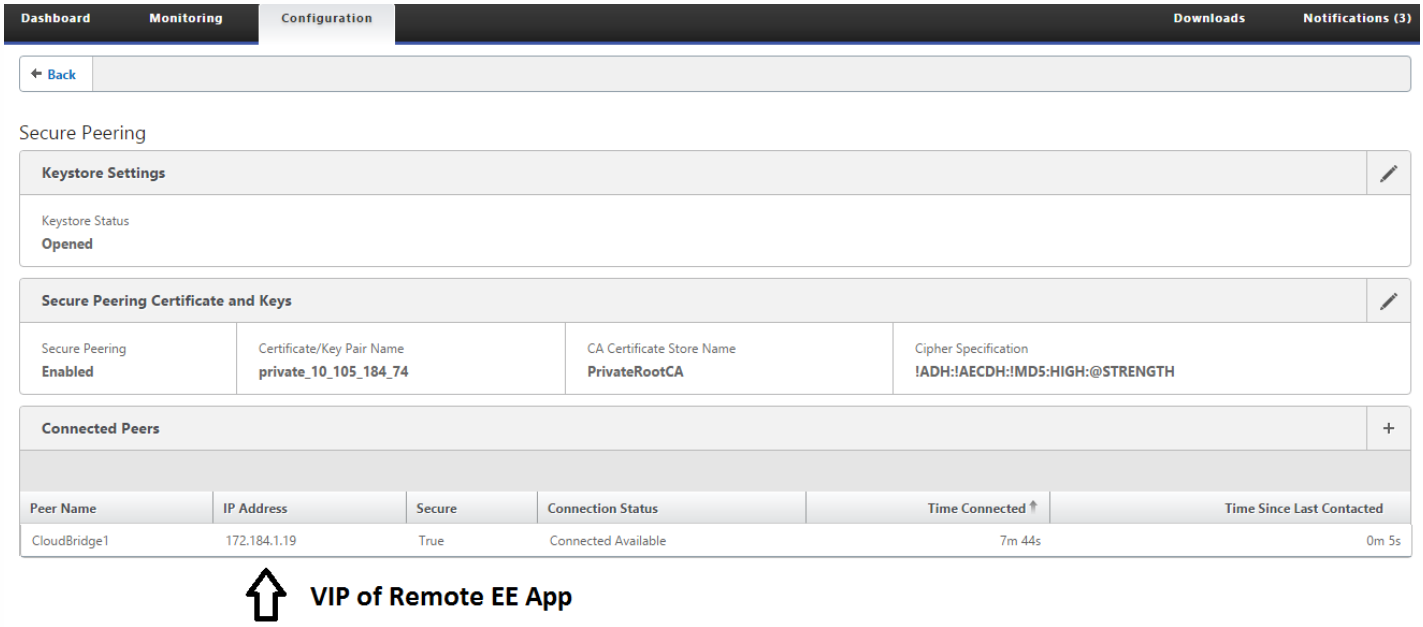

Monitoring

View Secure Partner Information on the Enterprise Edition Appliance under **WANOPTIMIZATION** > **Partners** in the **Monitoring** page.

a. Data Store Encryption can be performed on the Enterprise Edition appliance through feature enablement from the MCN under Optimization node for an Enterprise Edition appliance.

b. For an Enterprise Edition appliance, secure peering will always be enabled.

1. To validate if the **Private CA** and **Private Certificate Key** pair is generated successfully, review the information below:

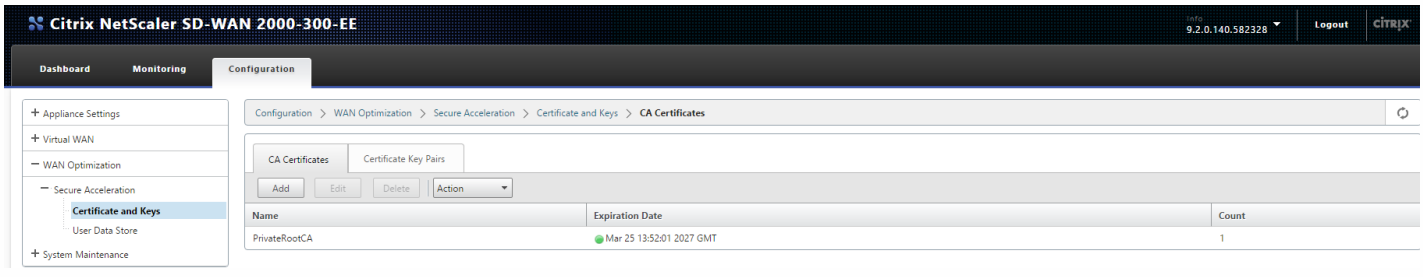

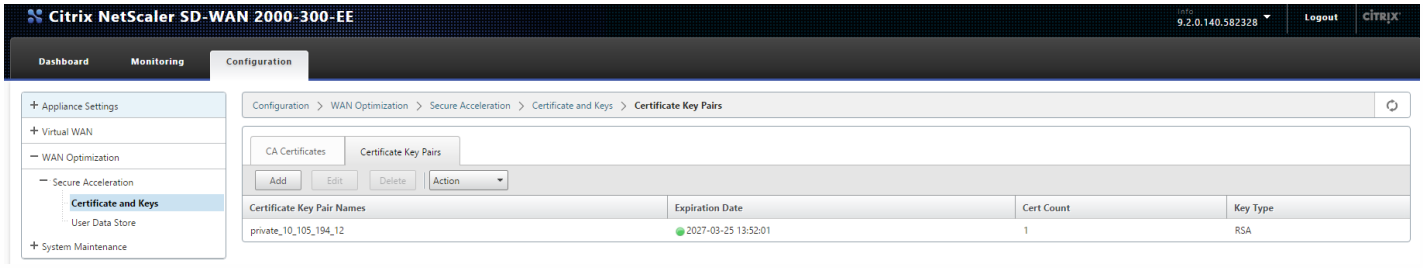

2. View **Secure Partner Information** on the Enterprise Edition appliance under **Monitoring** > **WAN Optimization** > **Partners** page.

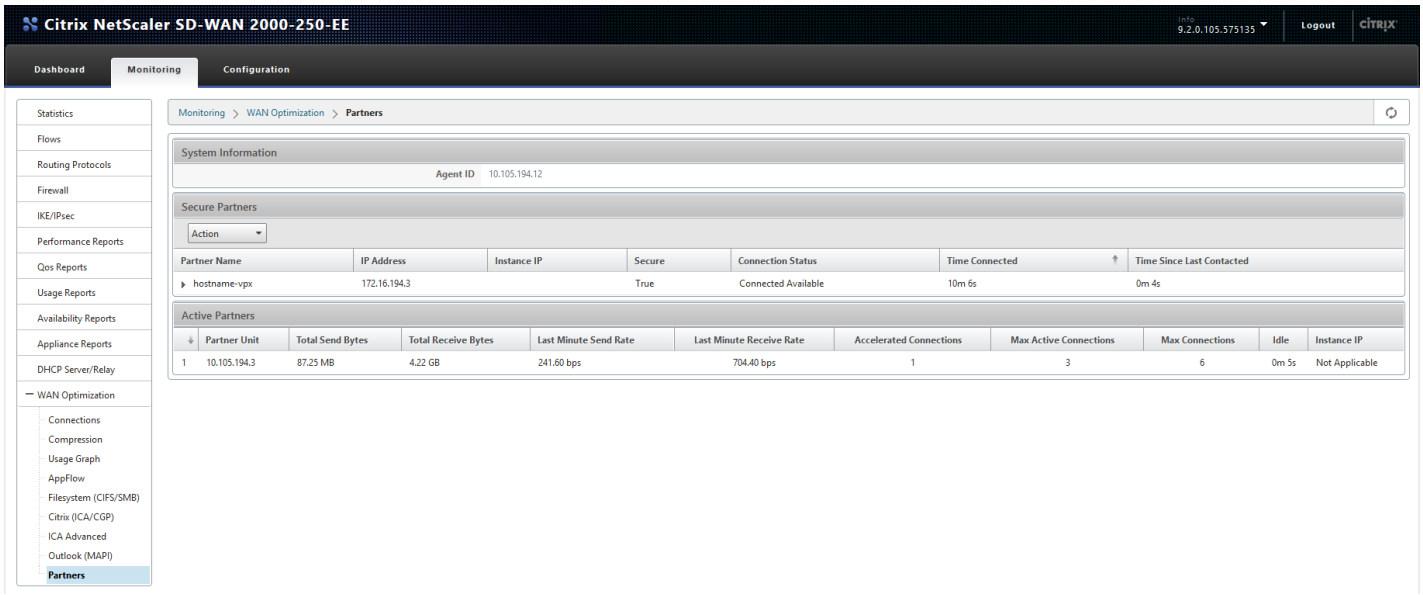

3. On partner appliance, **View Secure Partner Information** of the Enterprise Edition appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partner**s page.

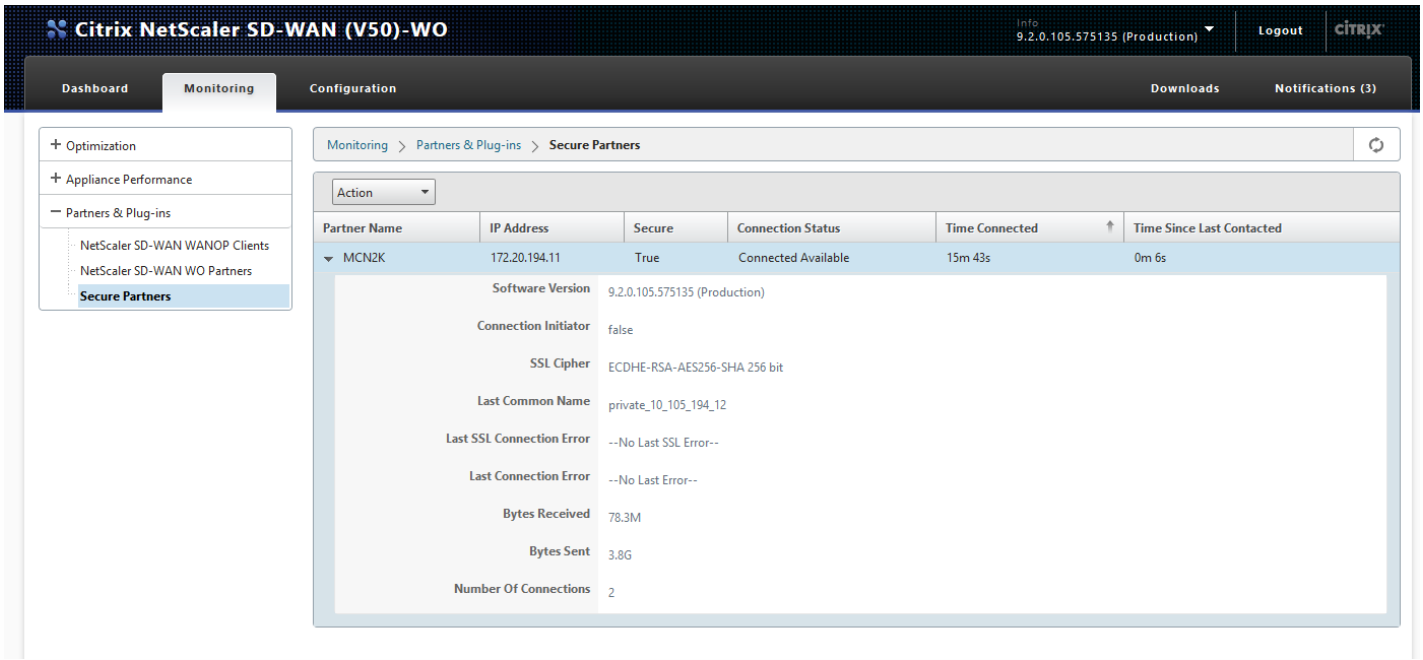

Troubleshooting

1. View **Secure Partner Success** / **Failure Information** on the Enterprise Edition appliance under **Monitoring** > **WAN**

Optimization > **Partners** > **Secure Partners** page.

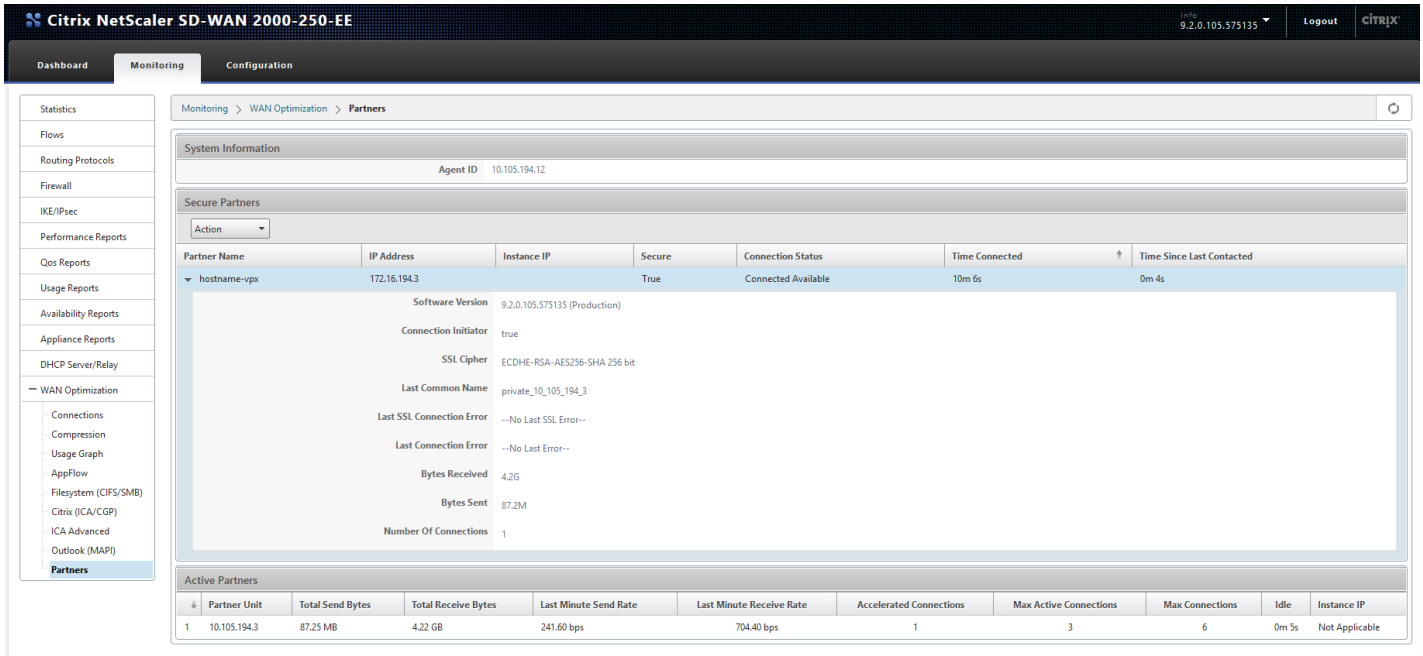

2. On partner appliance, view Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partners** page.

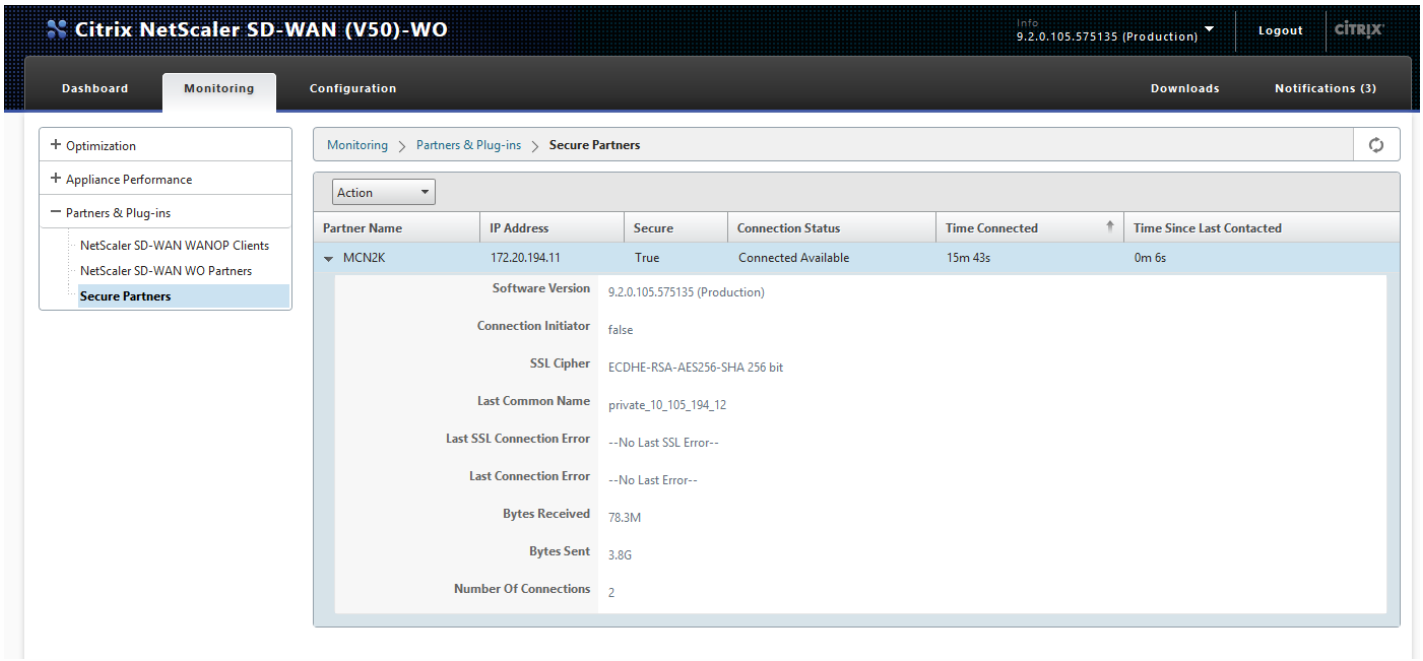

3. On partner appliance, view Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **Appliance Performance** > **Logging** page.

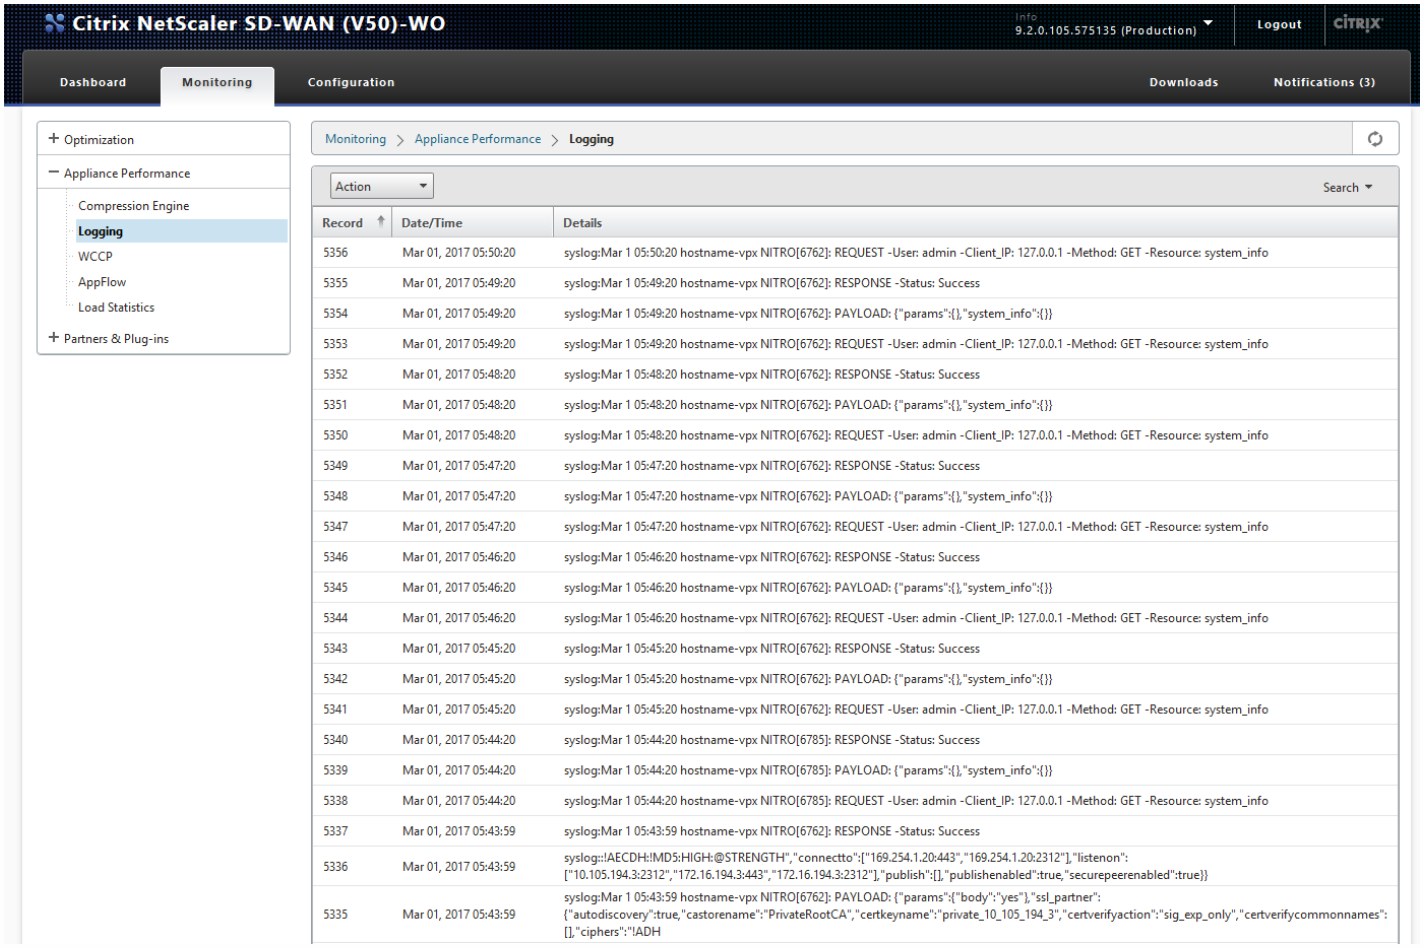

Auto Secure Peering Initiated from EE Appliance at DC Site and Branch Site EE Appliance

Aug 09, 2017

Configuration

To configure auto secure peering on a new Enterprise Edition appliance at DC:

1. In the SD-WAN web GUI, navigate to **Configuration** > **WAN Optimization** > **Secure Acceleration** > **Secure Peering**.

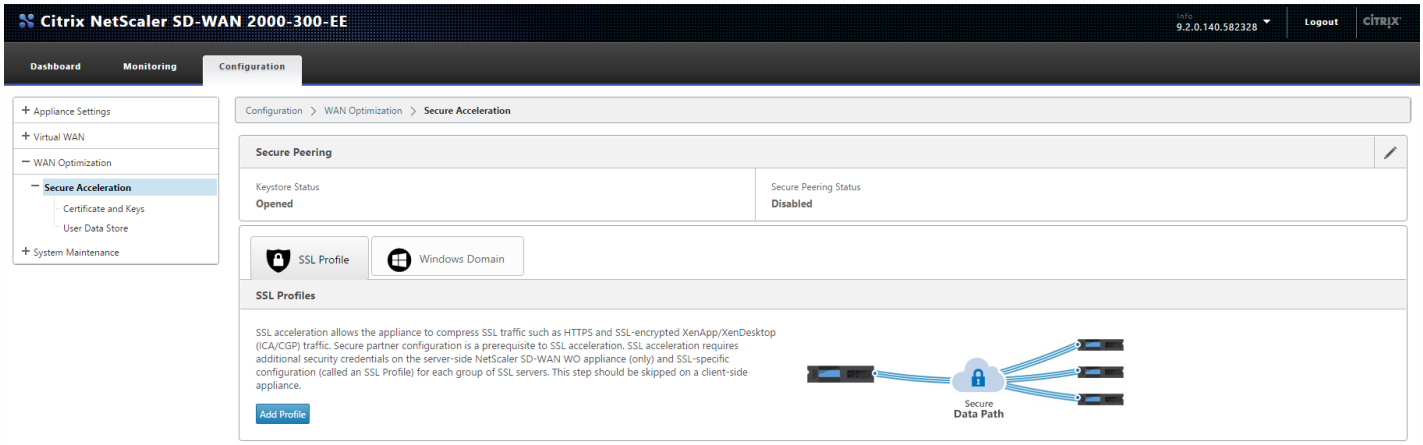

2. Configure keystore by providing the keystore password or by disabling keystore.

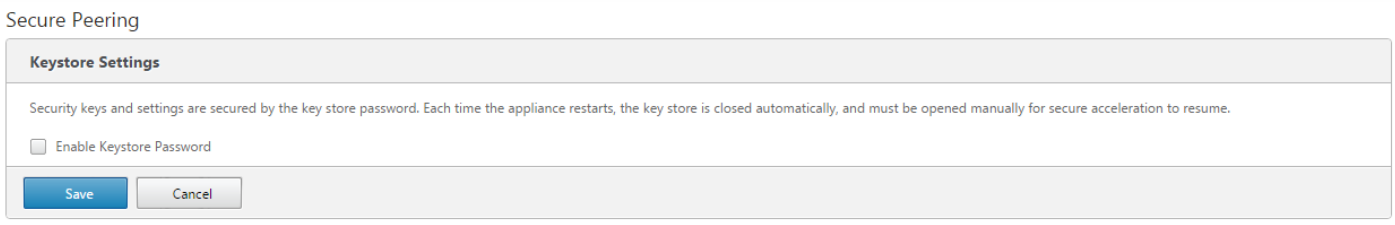

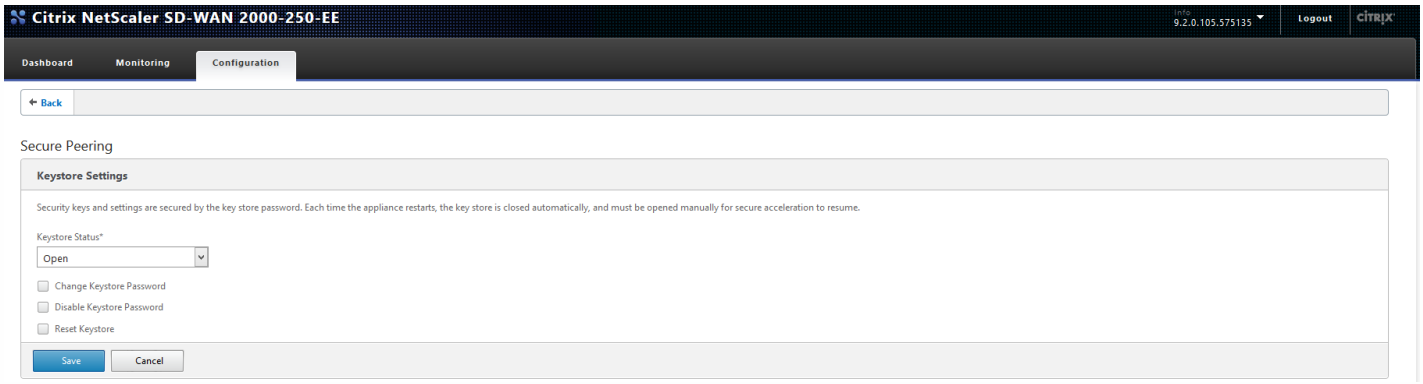

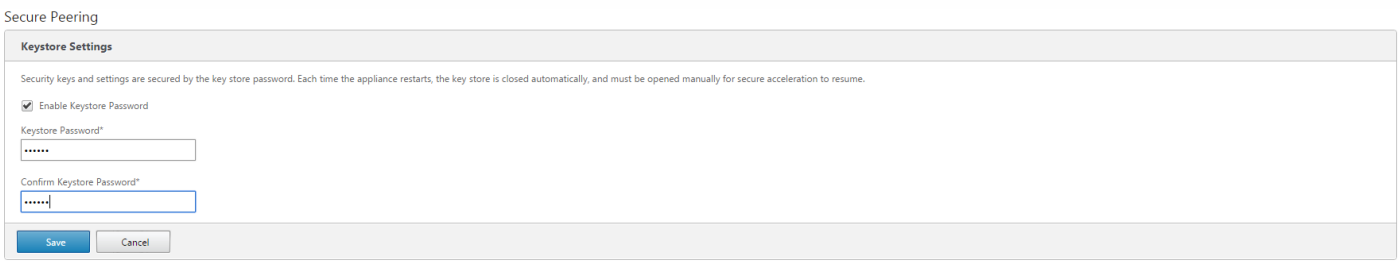

3. Enable **Secure Peering** by selecting **Private CA** to perform AUTOMATIC SECURE PEERING.

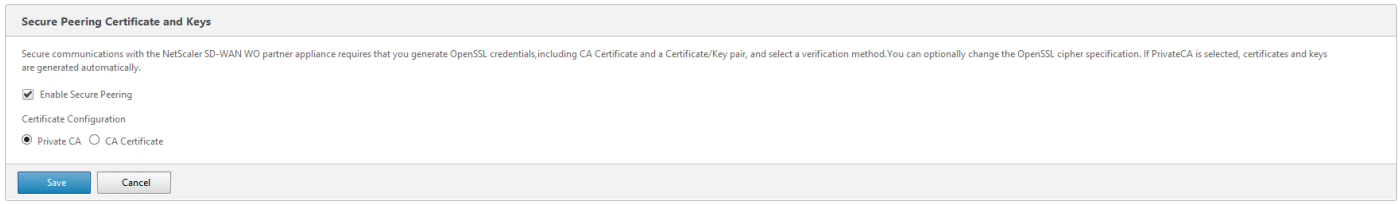

4. Click on the '+' icon and to add IP with username and password. After successful authentication with the remote IP and credentials provided, a request is sent to the remote machine that will install CA Certificate and the Private cert and key for itself locally on the remote machine.

> ## Note
>
> IP Address – IP Address of remote EE Appliance MANAGEMENT IP
>
> Username – Username of remote EE Appliance
>
> Password – Password of remote EE Appliance

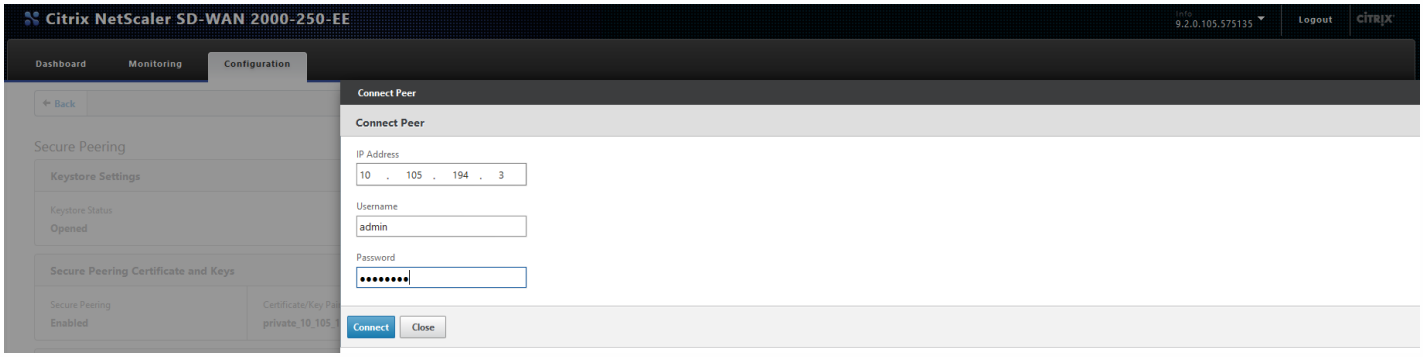

Monitoring

1. To validate if the Private CA and Private Certificate Key pair is generated successfully, review the information displayed below.

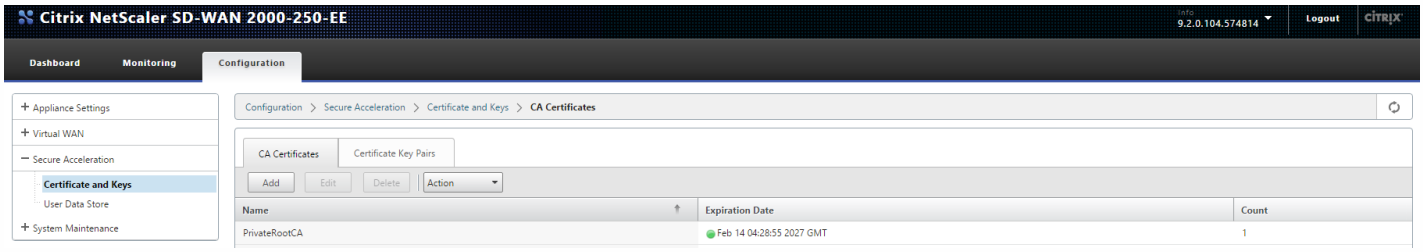

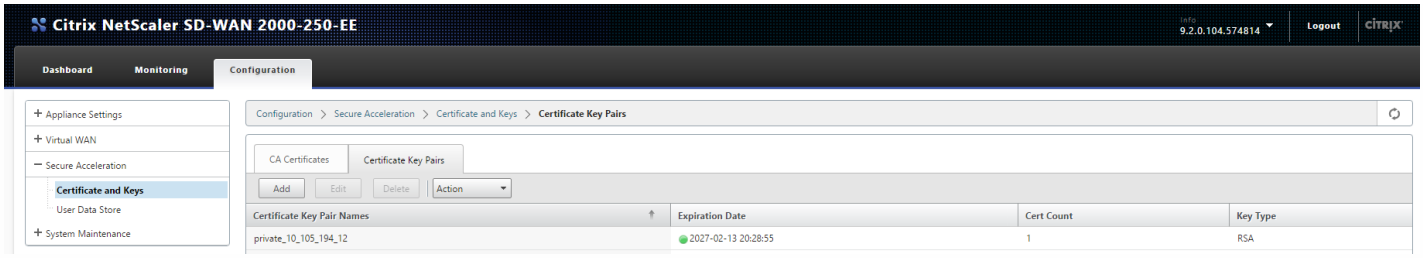

2. View **Secure Partner Information** on the Enterprise Edition Appliance under **Monitoring** > **WAN Optimization** > **Partners** page.

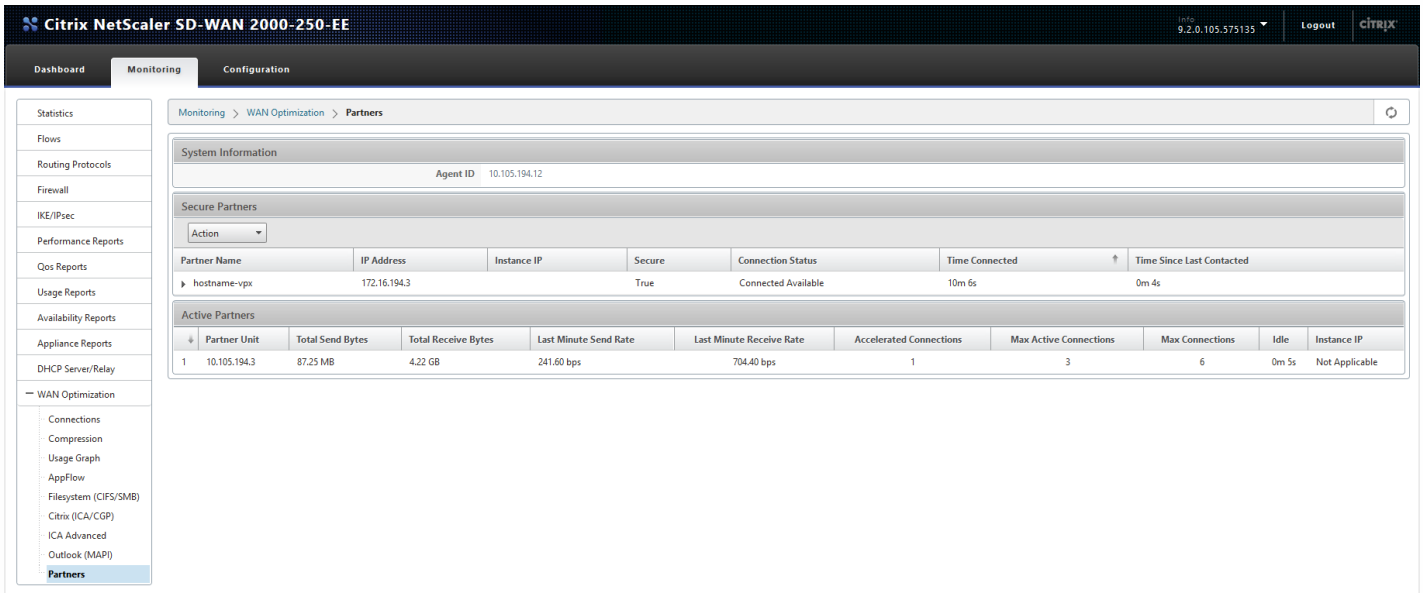

3. On partner appliance, view Secure Partner Information on the Enterprise Edition Appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partners** page.

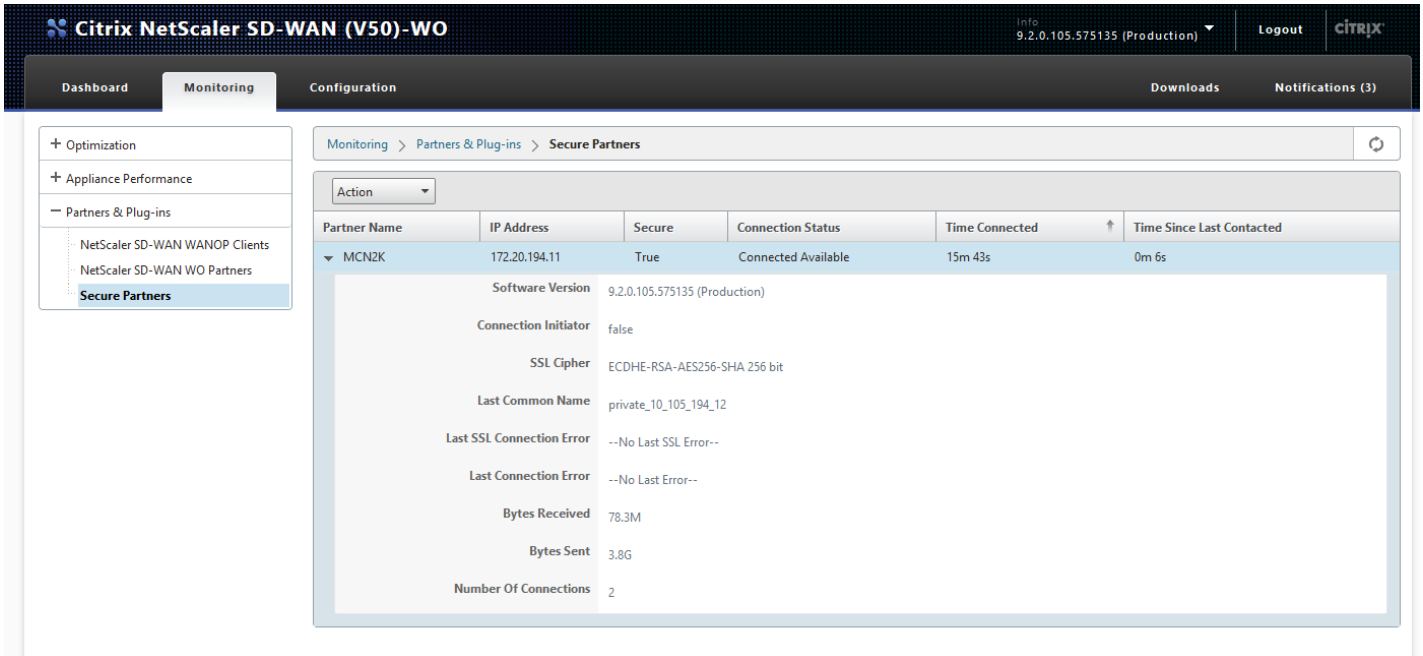

Troubleshooting

1. View Secure Partner Success / Failure Information on the Enterprise Edition Appliance under **Monitoring** > **WAN Optimization** > **Partners** > **Secure Partners** page.

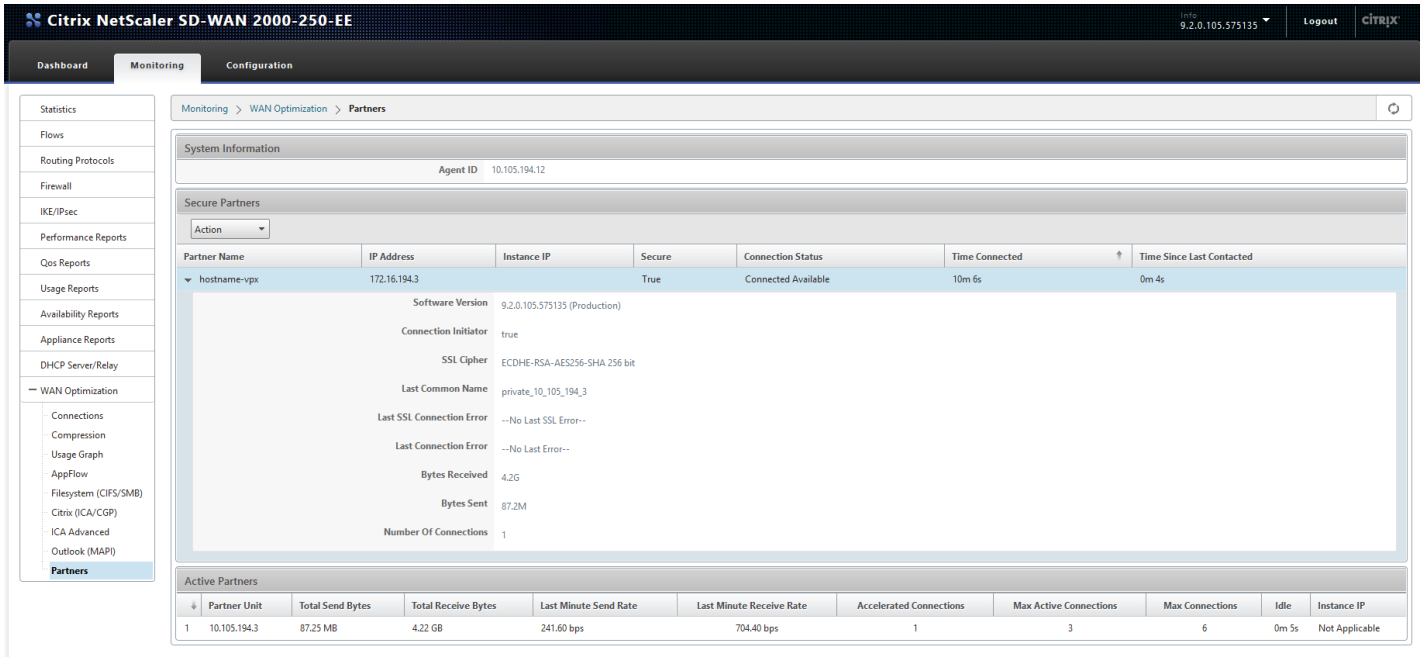

2. On partner Appliance, view Secure Partner Information on the Enterprise Edition Appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partners** page.

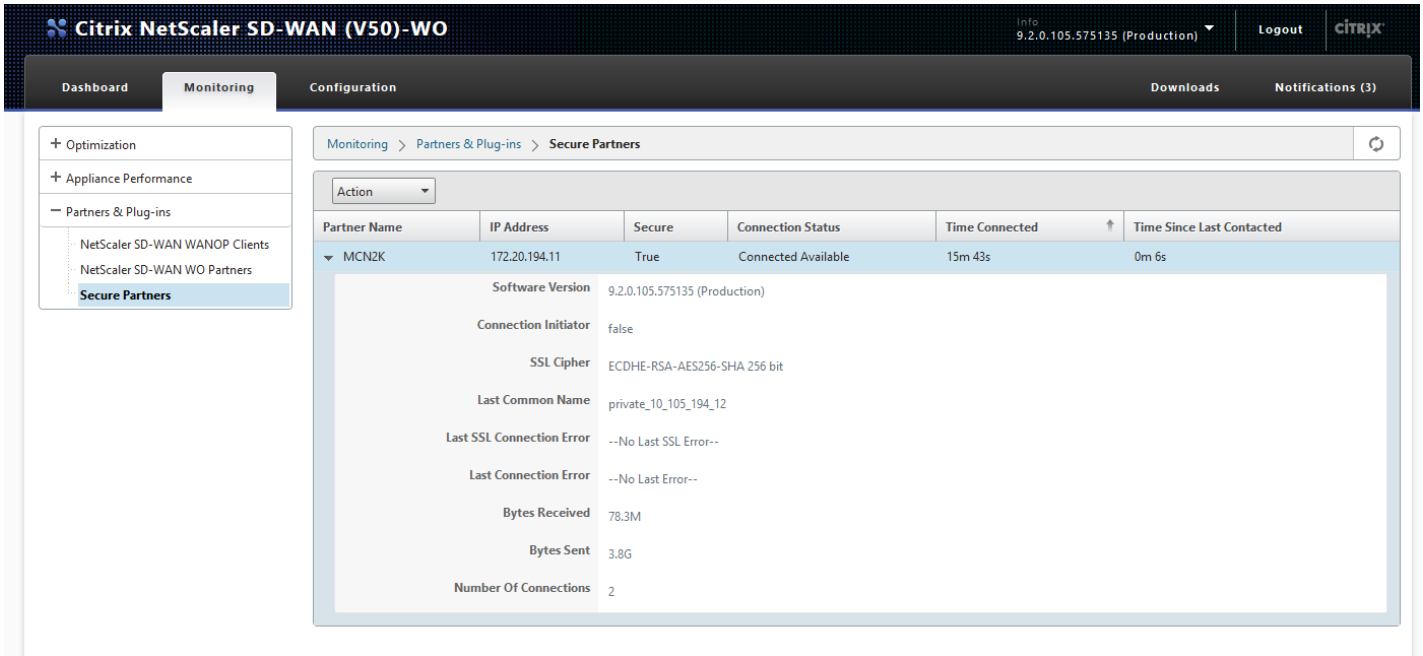

3. On partner Appliance, view Secure Partner Information on the Enterprise Edition Appliance under **Monitoring** > **Appliance Performance** > **Logging** page.

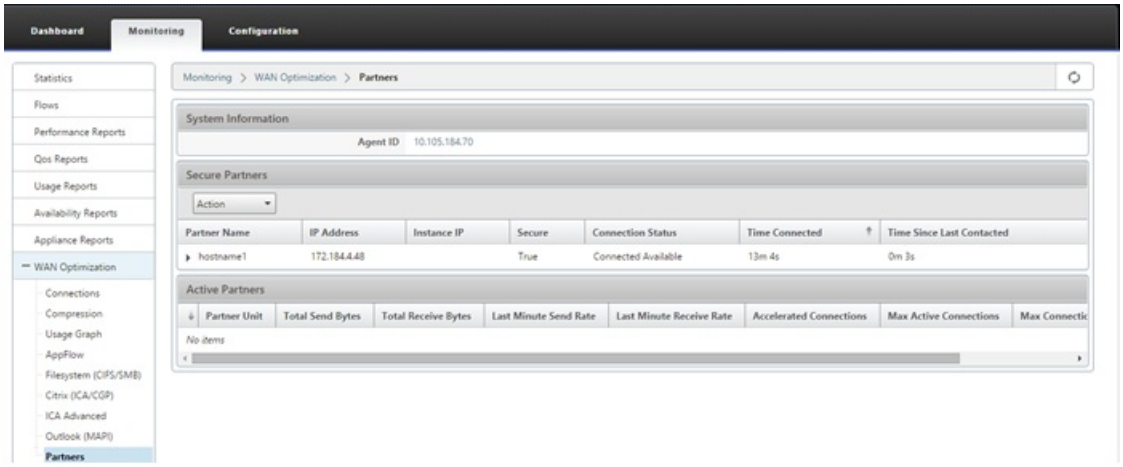

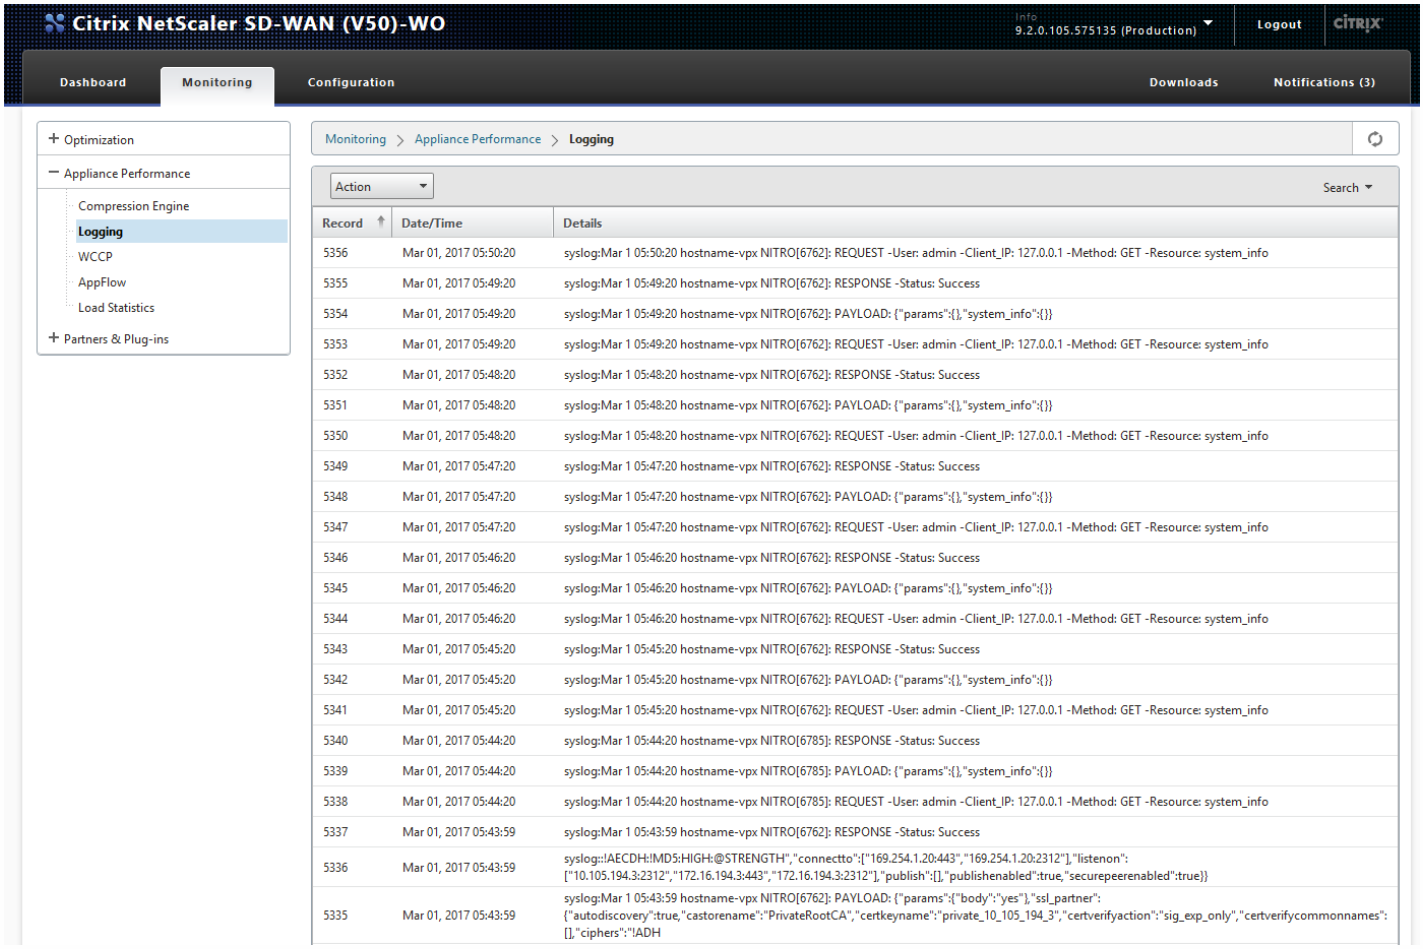

Auto Secure Peering Initiated from EE Appliance at DC Site and Branch with WANOP/SE Appliance

Aug 09, 2017

Configuration

To configure a new Enterprise Edition appliance with auto secure peering at the DC site:

1. In the SD-WAN web GUI, navigate to **Configuration** > **WAN Optimization** > **Secure Acceleration** > **Secure Peering**.

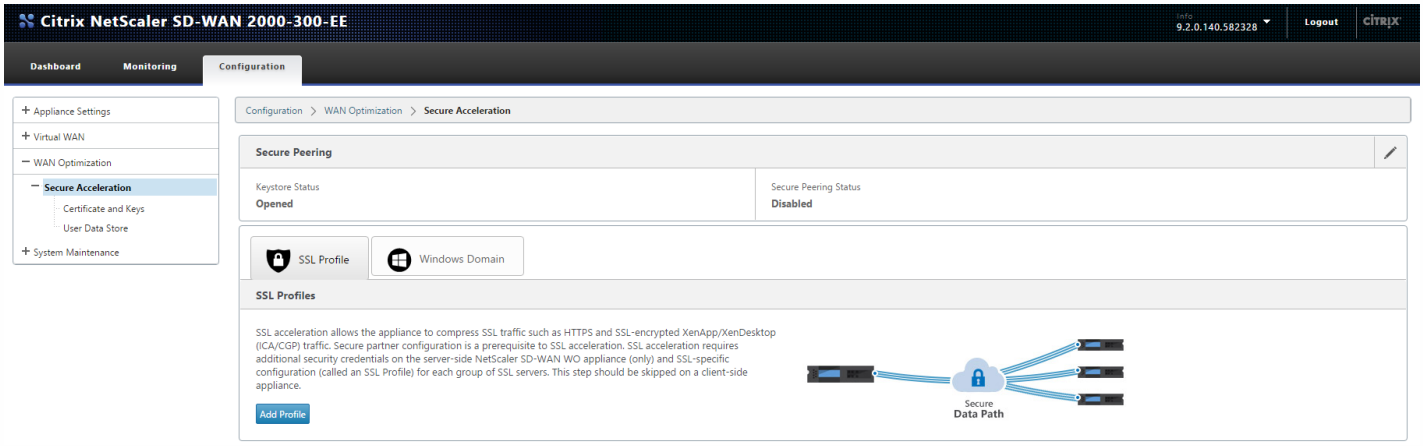

2. Configure keystore by providing the keystore password or by disabling the keystore.

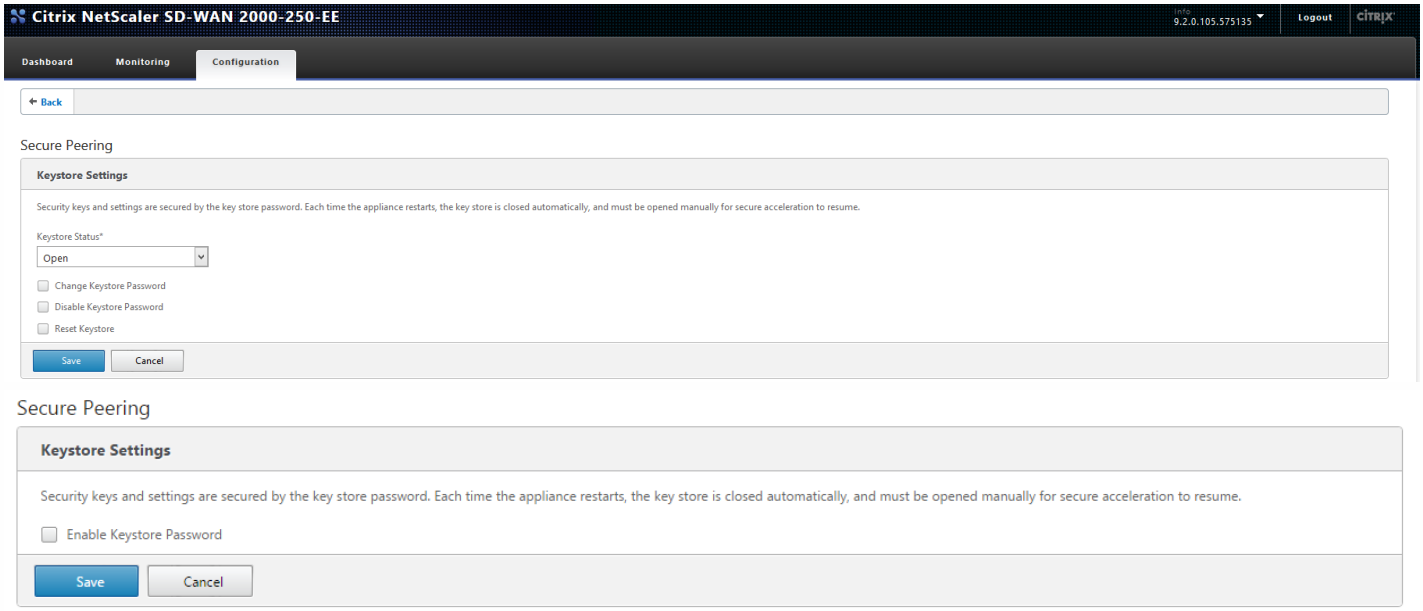

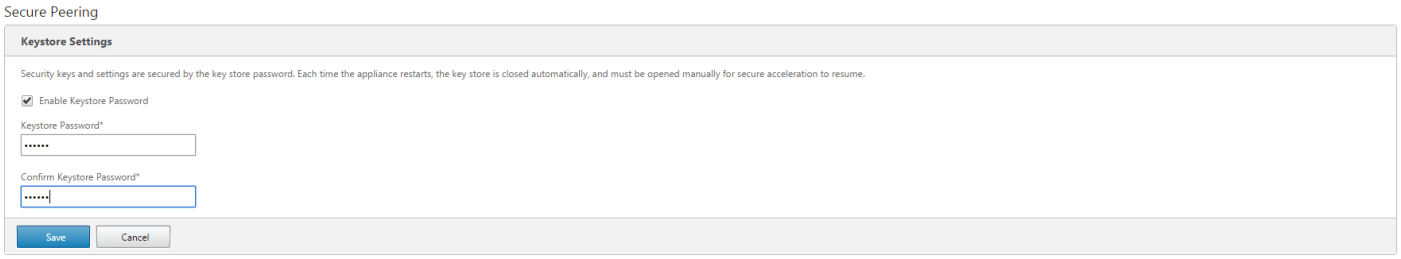

3. Enable **Secure Peering** by selecting **Private CA** to perform AUTOMATIC SECURE PEERING.

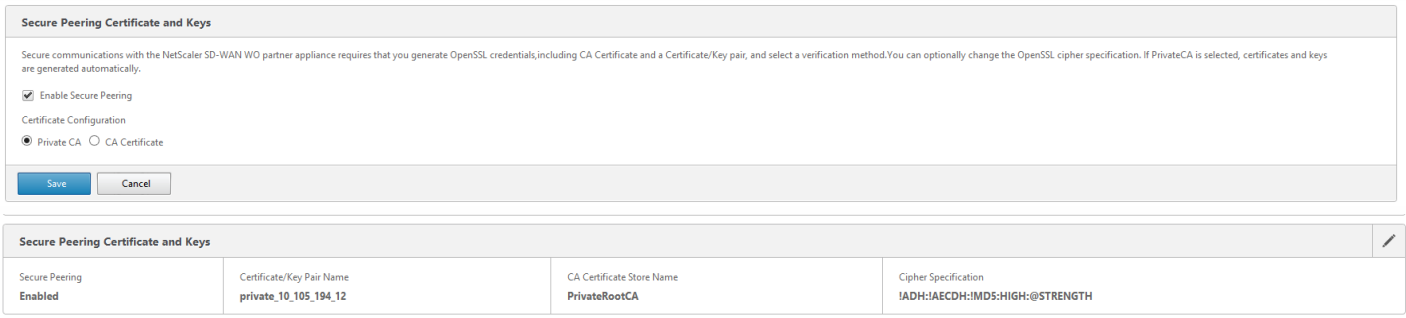

4. Click on the '+' icon and to add IP with username and password. After successful authentication with the remote IP and credentials provided, a request is sent to the remote machine that will install CA Certificate and the Private cert and key for itself locally on the remote machine.

Note

IP Address – IP Address of remote WANOP Standalone or Standard Edition Appliance MANAGEMENT IP.

Username – Username of remote WANOP Standalone or Standard Edition Appliance.

Password – Password of remote WANOP Standalone or Standard Edition Appliance.

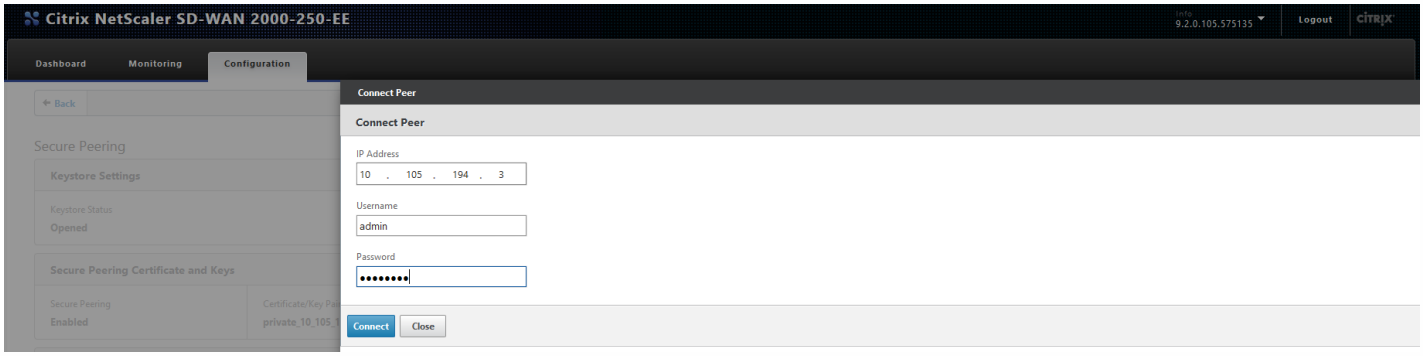

5. After Successful Authentication, you can view Secure Peering as TRUE and the partner IP as one of the Virtual IP of the remote WANOP Standalone appliance.

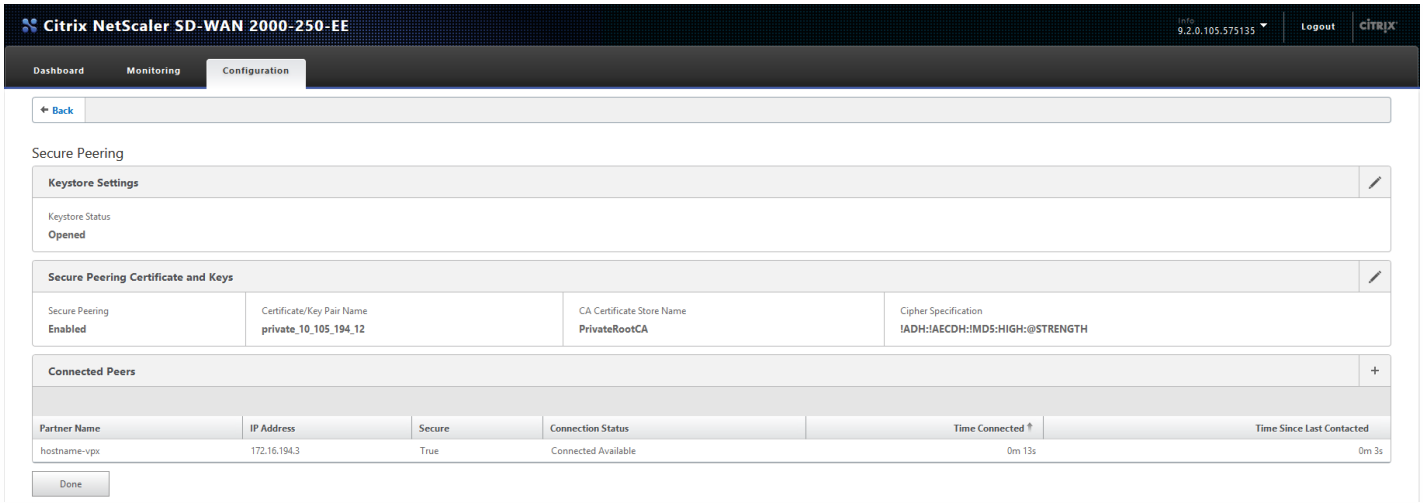

Monitoring

1. To validate if the Private CA and Private Certificate Key pair is generated successfully, review the information below.

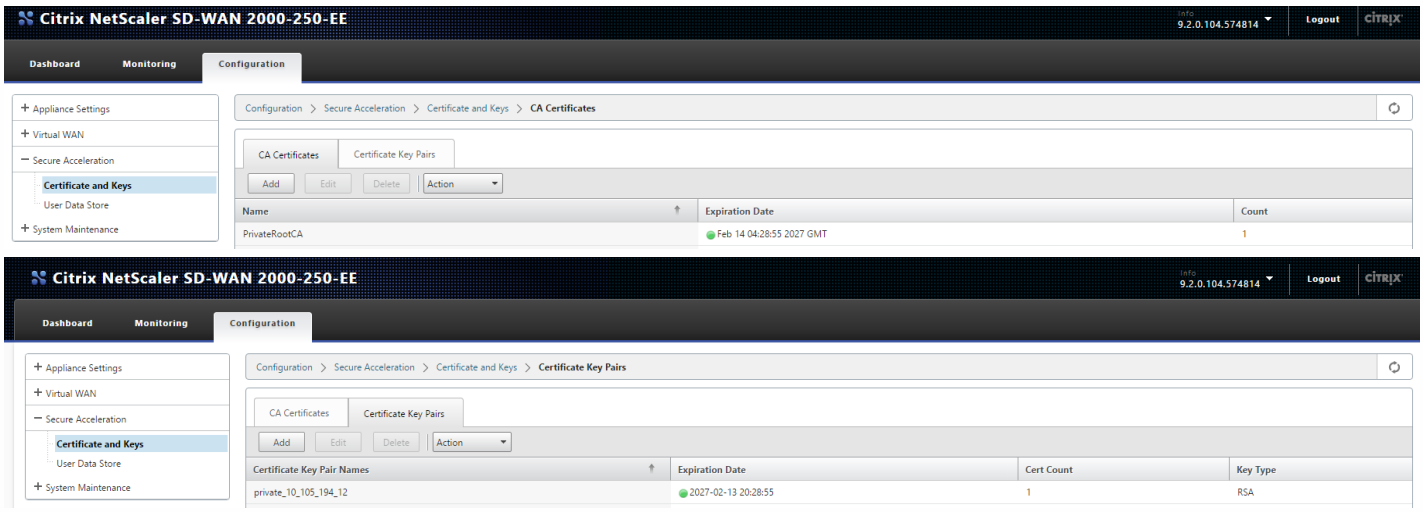

2. View Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **WAN Optimization** > **Partners** page.

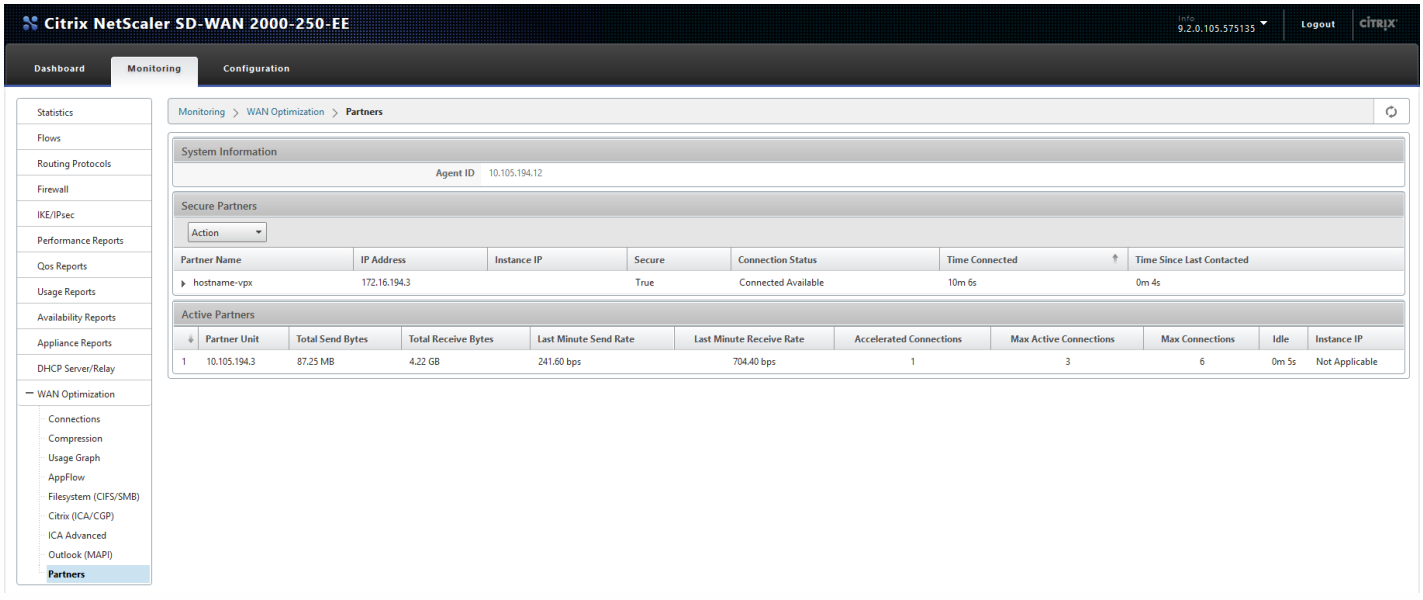

3. On partner appliance, View Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partners** page.

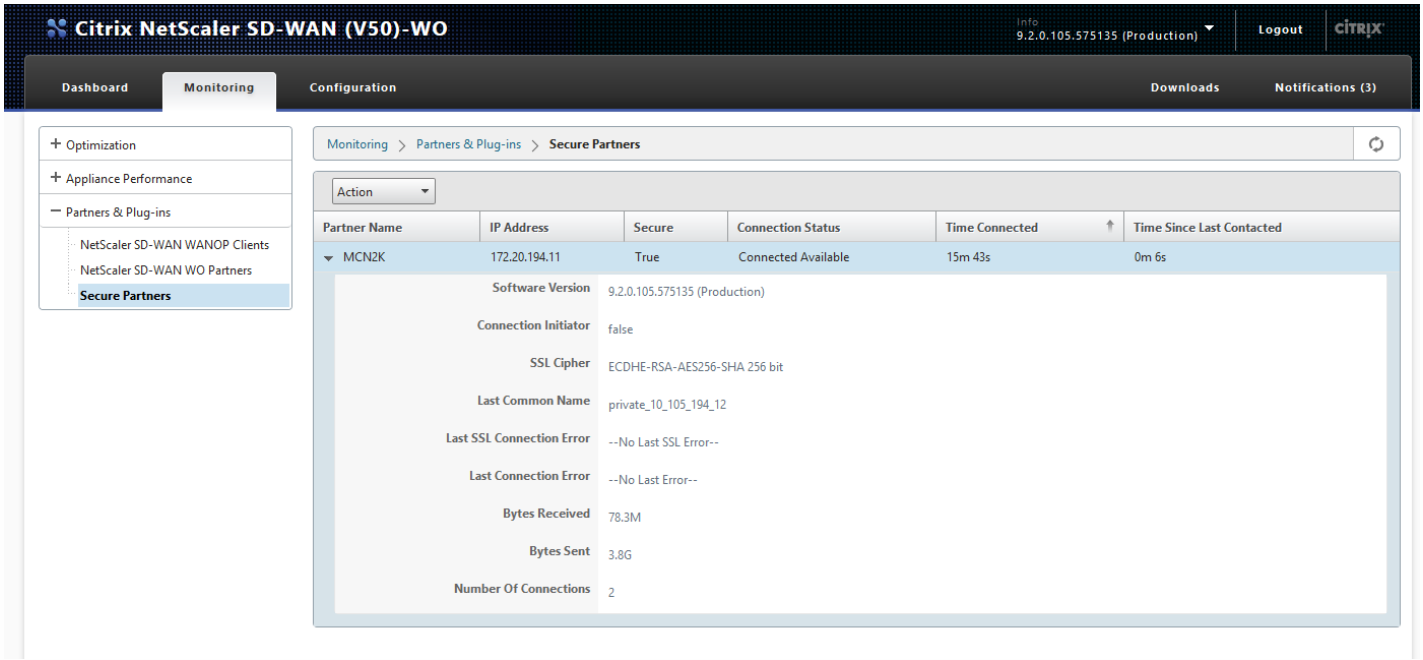

Troubleshooting

1. View Secure Partner Success / Failure Information on the Enterprise Edition appliance under **Monitoring** > **WAN Optimization** > **Partners** > **Secure Partners** page.

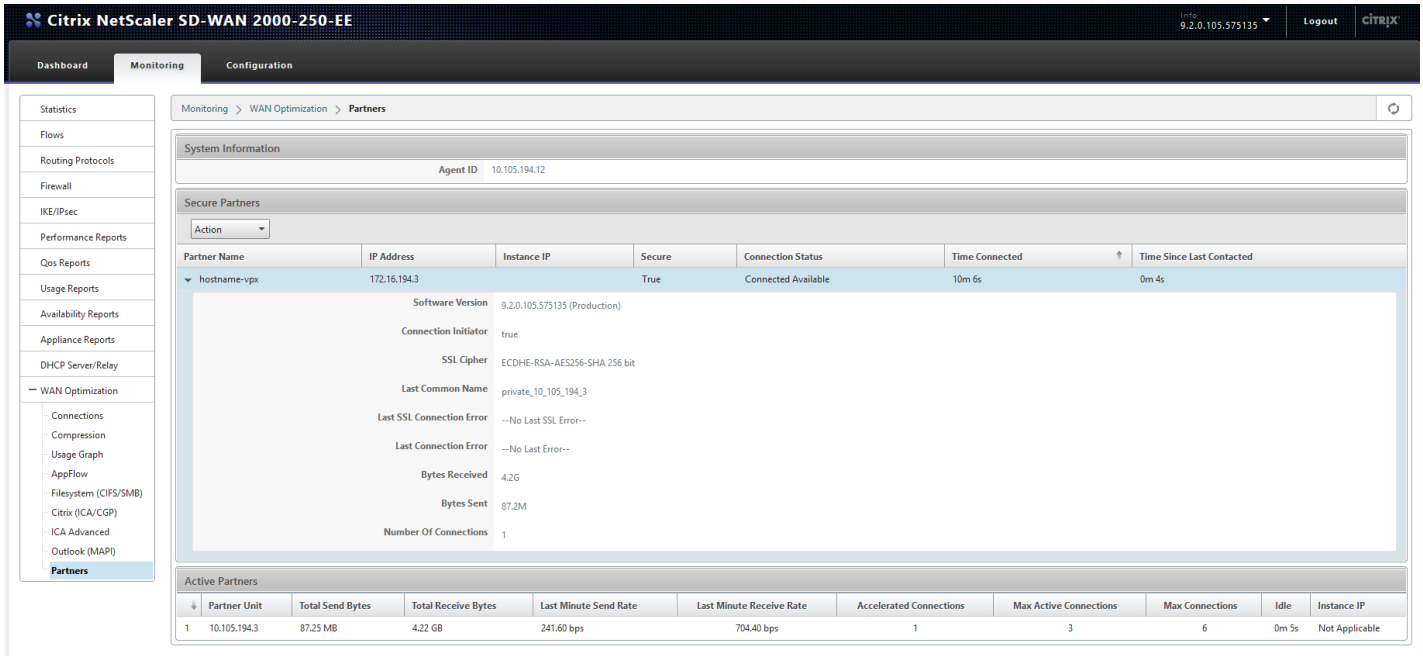

2. On partner appliance, view **Secure Partner Information** on the Enterprise Edition appliance under **Monitoring** > **Partners & Plug-ins** > **Secure Partners** page.

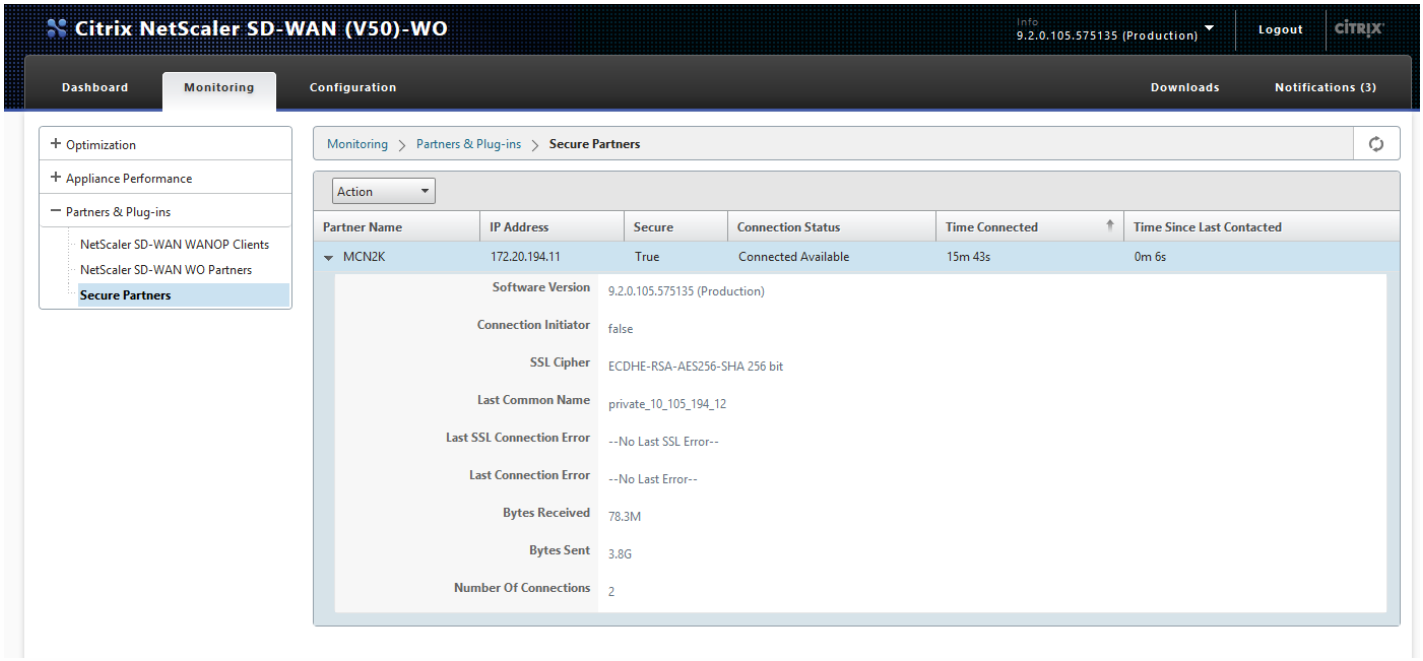

3. On partner appliance, view **Secure Partner Information** on the Enterprise Edition appliance under **Monitoring** > **Appliance Performance** > **Logging** page.

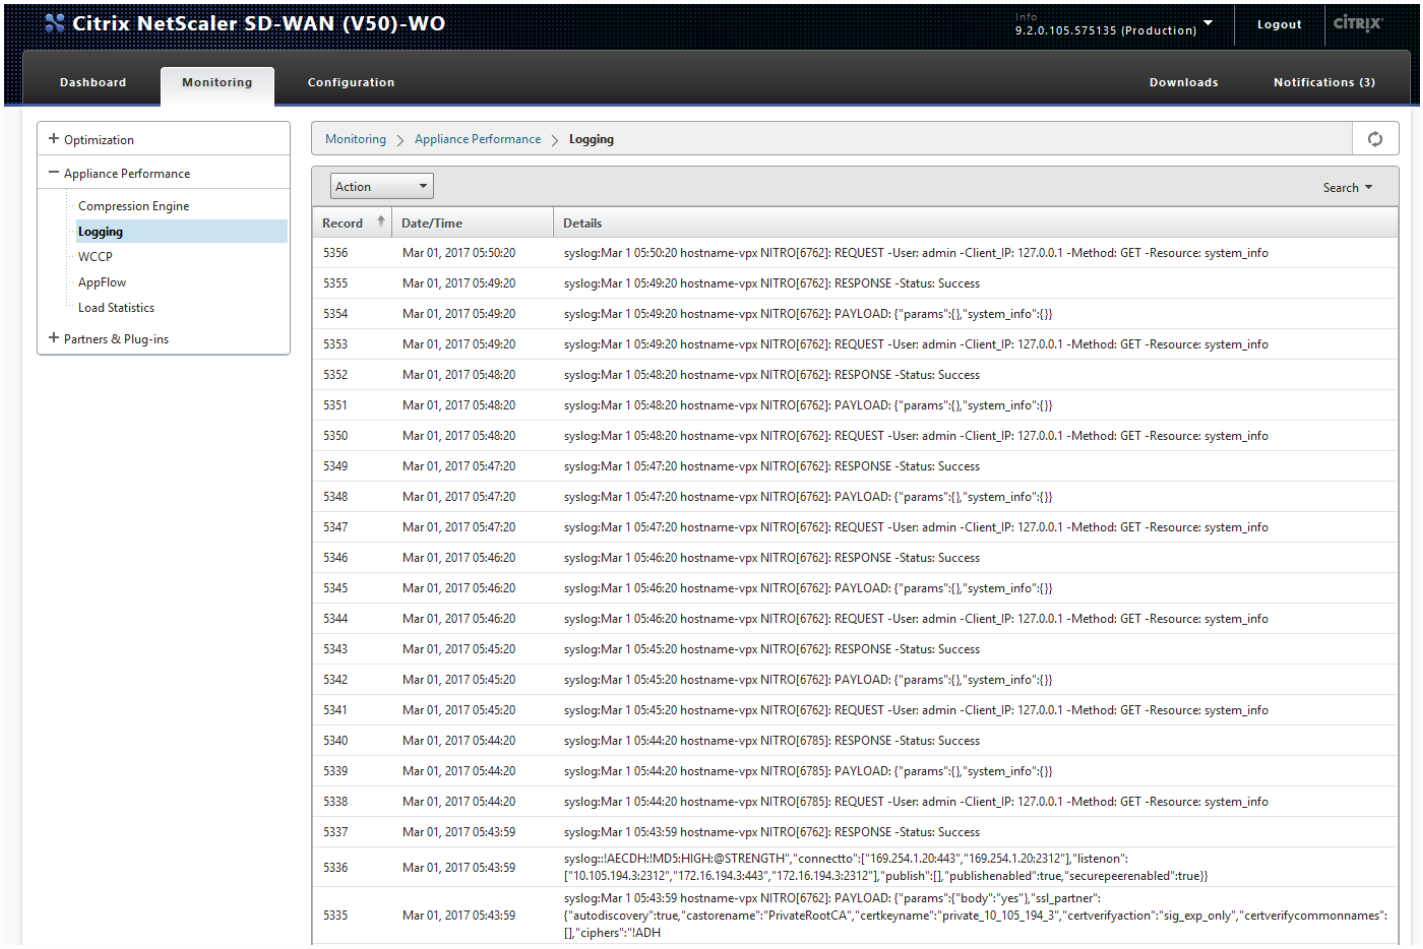

Manual Secure Peering Initiated from EE Appliance at DC Site and Branch EE Appliance

Aug 09, 2017

This deployment configures DC site EE appliance in LISTEN ON mode and Branch site EE appliance in CONNECT TO mode.

Configuration

To configure auto secure peering initiated from an EE appliance at DC site and EE appliance at Branch site:

1. Upload **CA Certificate** and **CA Key Certificate** obtained from authentic certificate and provide to SD-WAN as shown below.

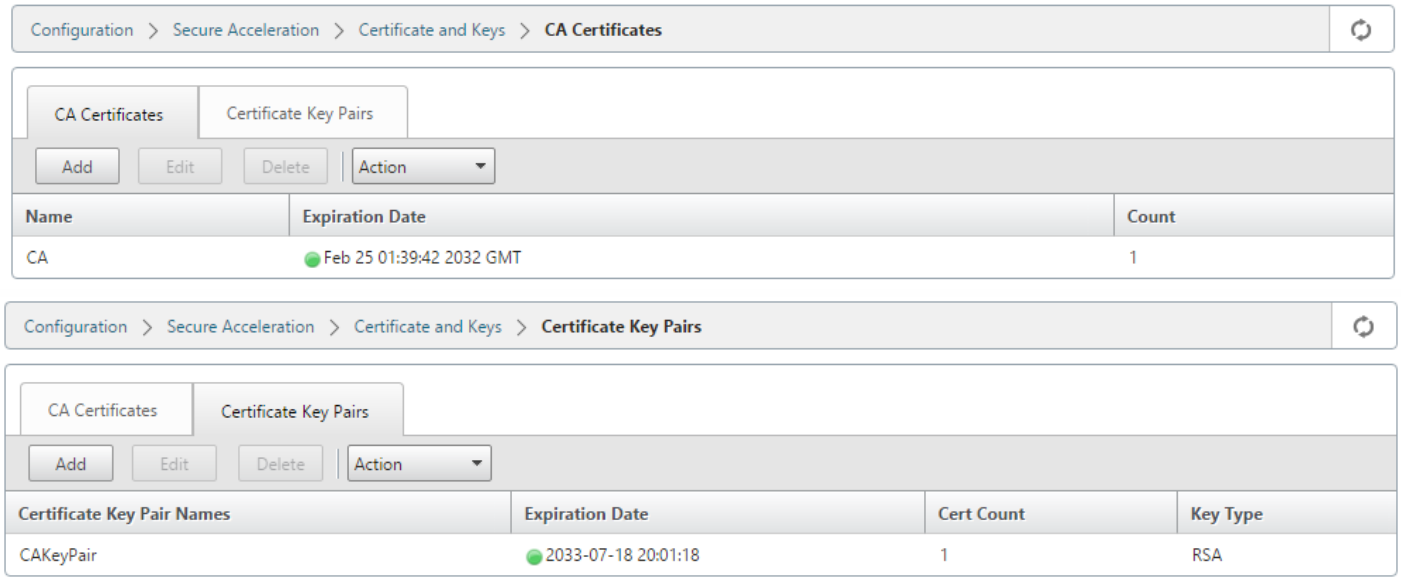

2. On a new EE appliance at the DC site, in the SD-WAN web GUI, go to **Configuration** > **Secure Acceleration** > **Secure Peering**.

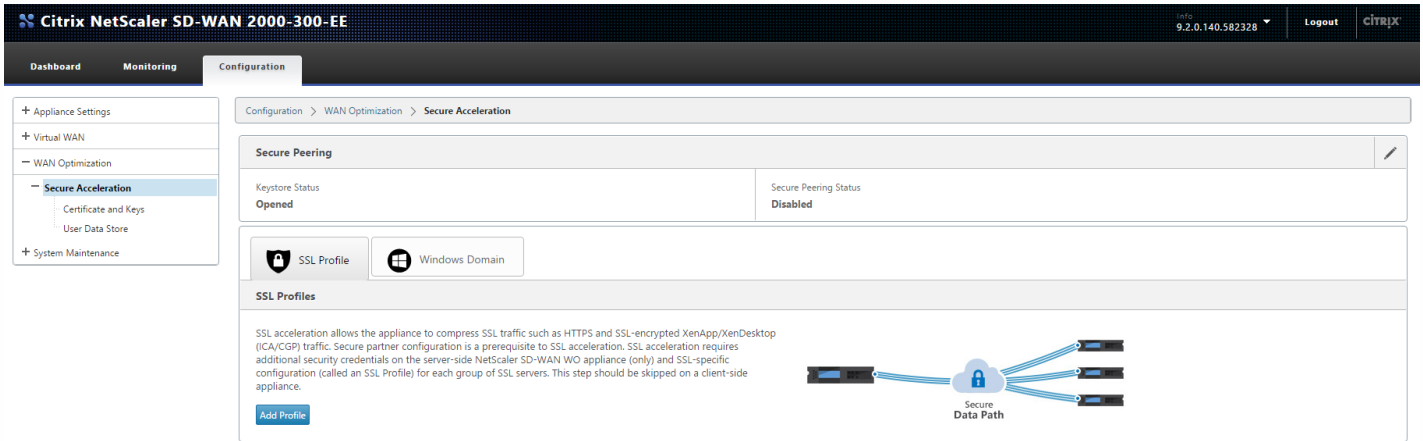

3. Configure keystore by providing the keystore password or by disabling the keystore.

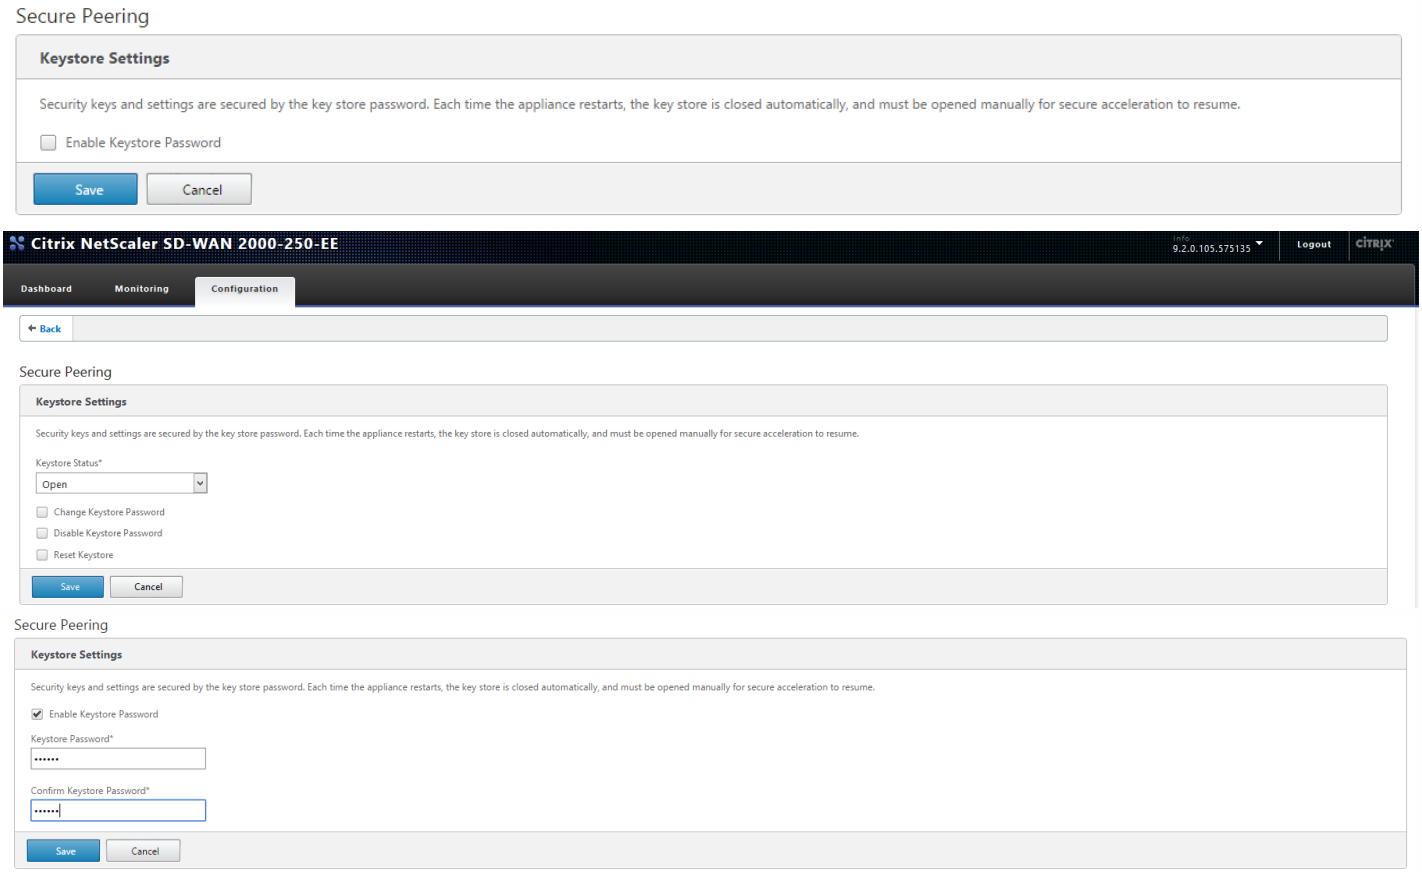

4. Enable secure peering by selecting **CA Certificate** radio button and providing uploaded CA and CA Key pair certificates appropriately as shown below.

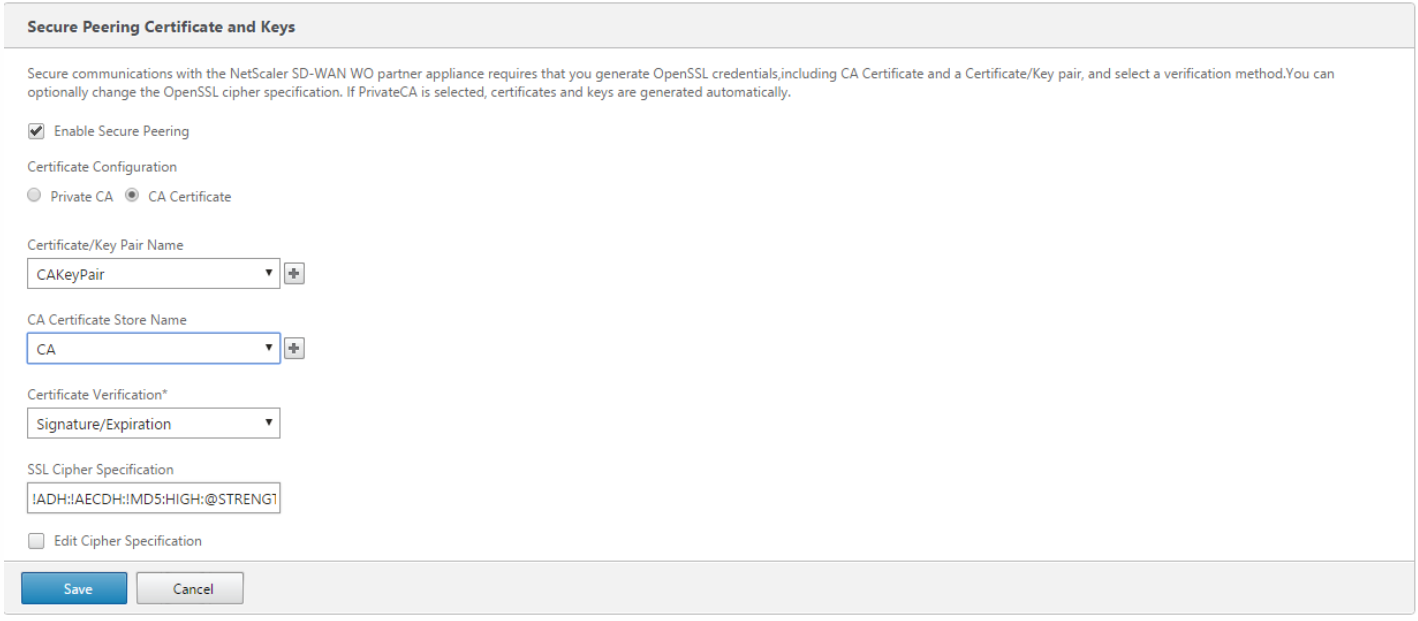

5. Provide Remote machine's Virtual IP along with Port 443 as shown below.

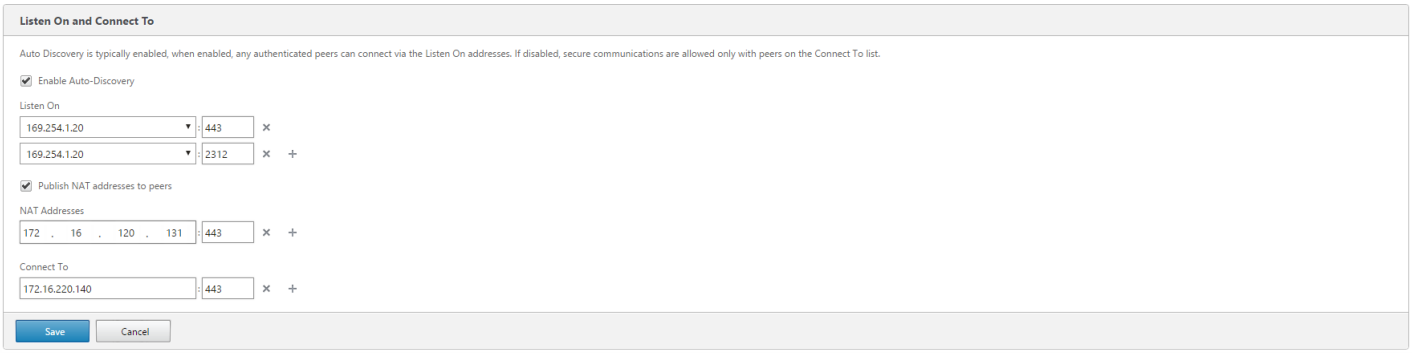

Monitoring

1. To validate if the **Private CA** and **Private Certificate Key** pair is generated successfully, review the information below.

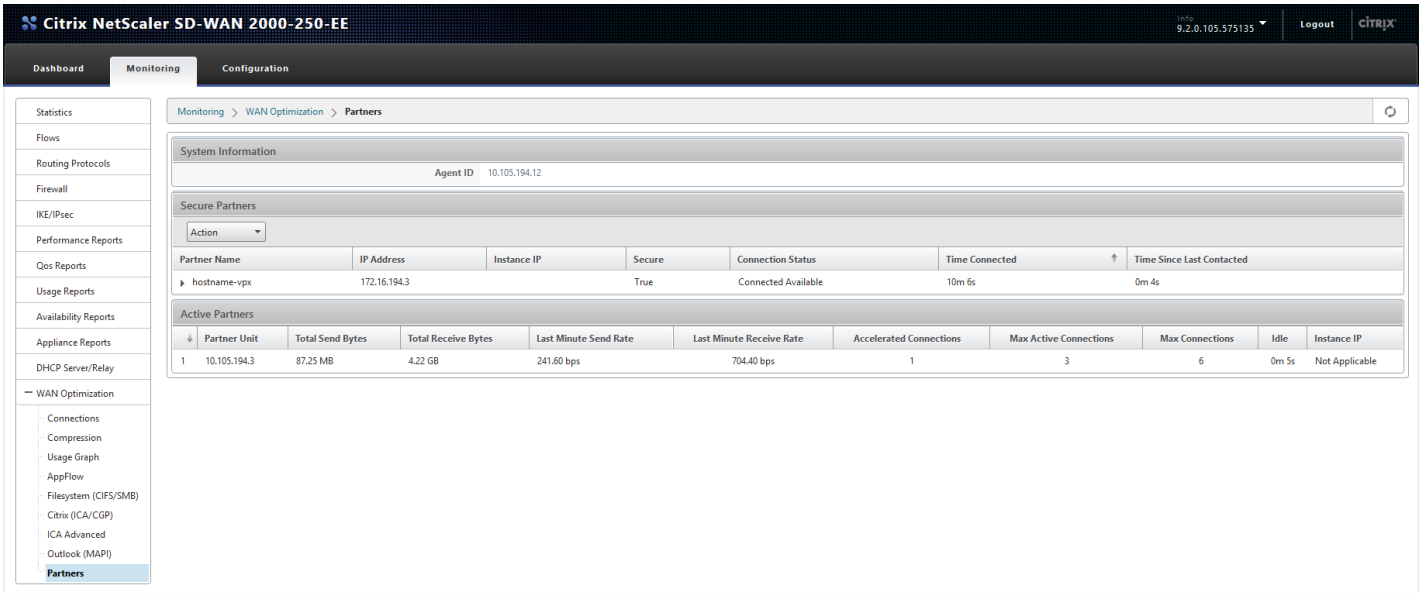

2. On partner appliance, **View Secure Partner Information** on the Enterprise Edition appliance under **Monitoring** > **Partners** > **Secure Partners** page.

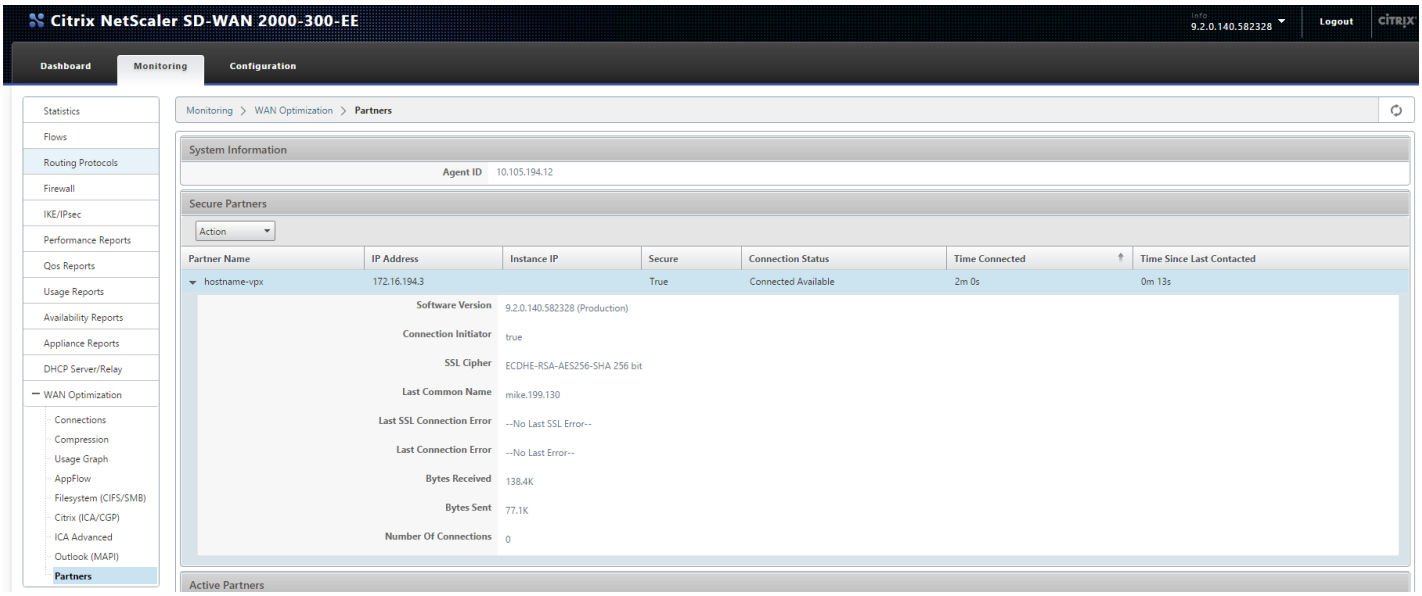

Troubleshooting

1. View **Secure Partner Success / Failure** Information on the Enterprise Edition Appliance under **Monitoring** > **WAN Optimization** > **Partners** > **Secure Partners** page.

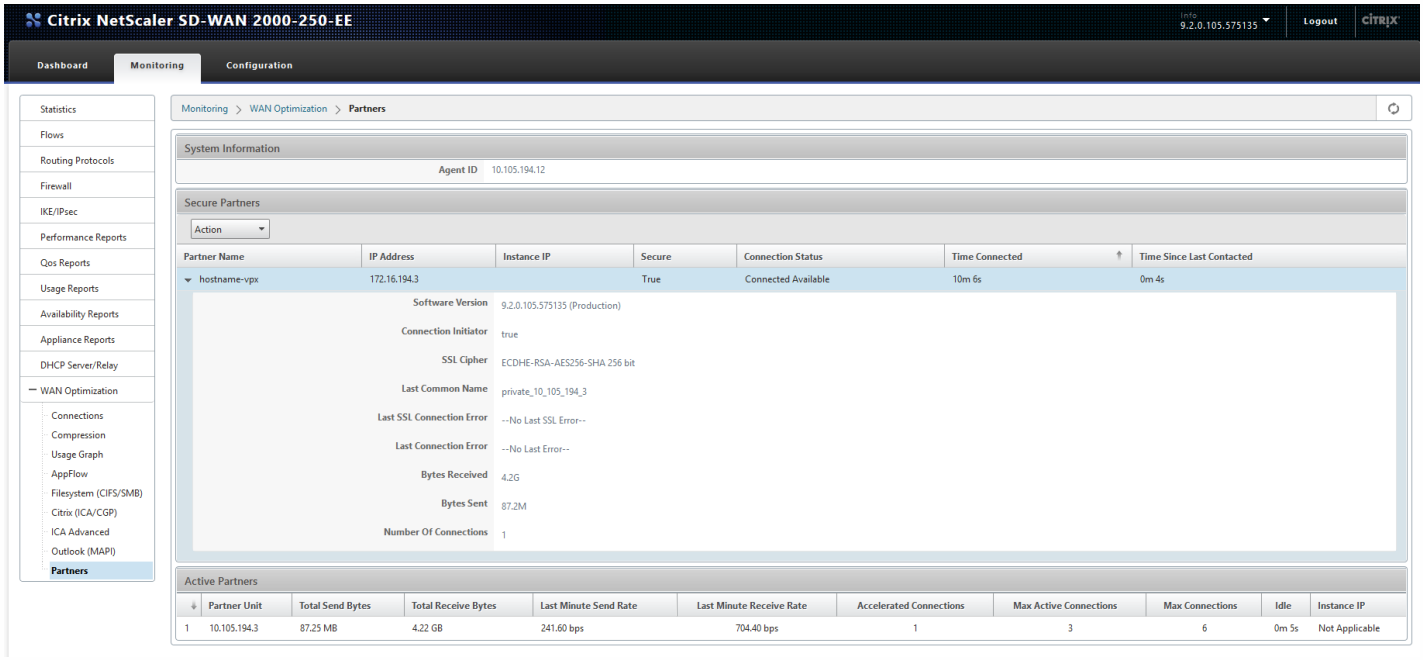

Manual Secure Peering initiated from EE appliance at DC site to Branch WANOP/SDWAN-SE Appliance

Aug 09, 2017

Configuration

1. Upload **CA Certificate** and **CA Key Certificate** obtained from authentic certificate and provide to SD-WAN as shown below.

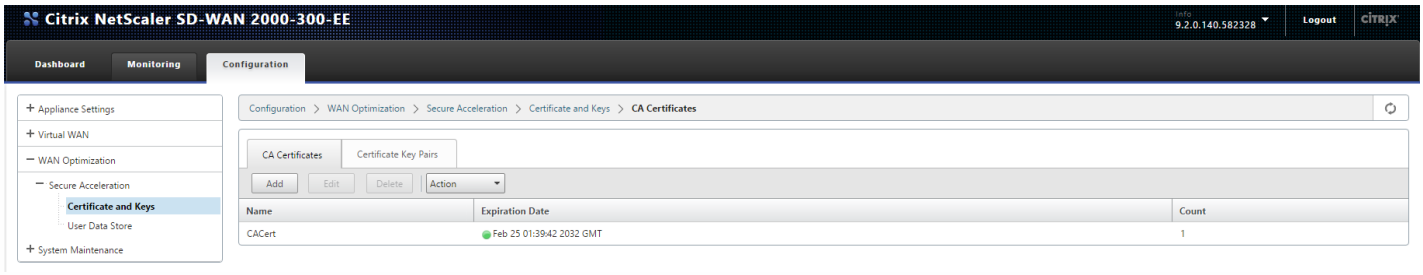

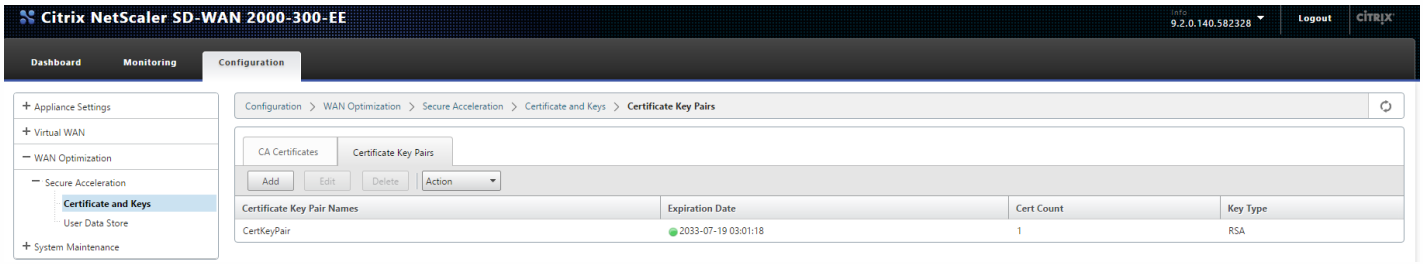

2. On a new EE appliance at the DC site, in the SD-WAN web GUI, go to **Configuration** > **Secure Acceleration** > **Secure Peering**.

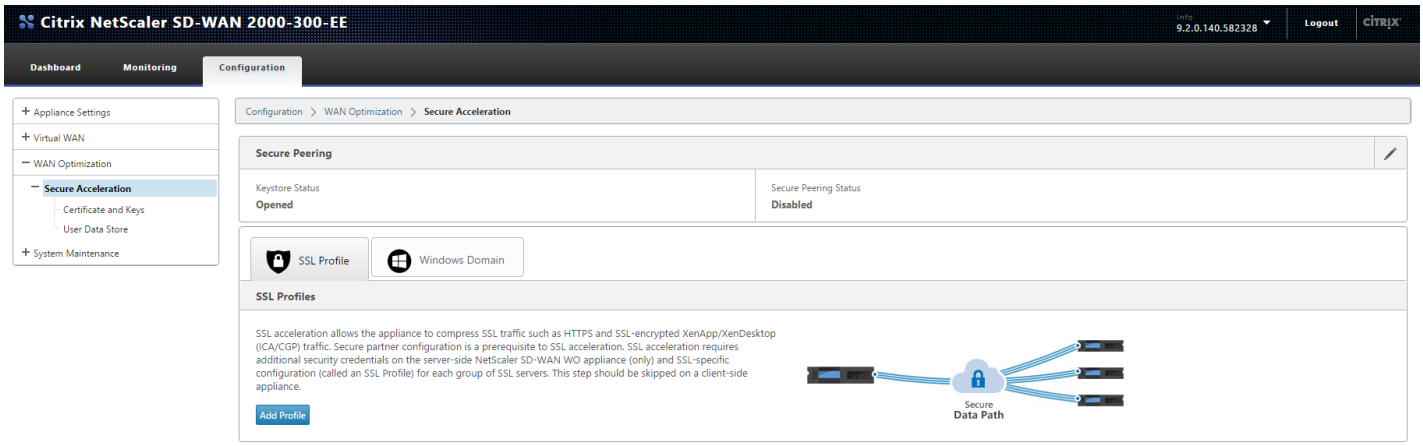

3. Enable the keystore by providing the k**eystore password** or disable the keystore.

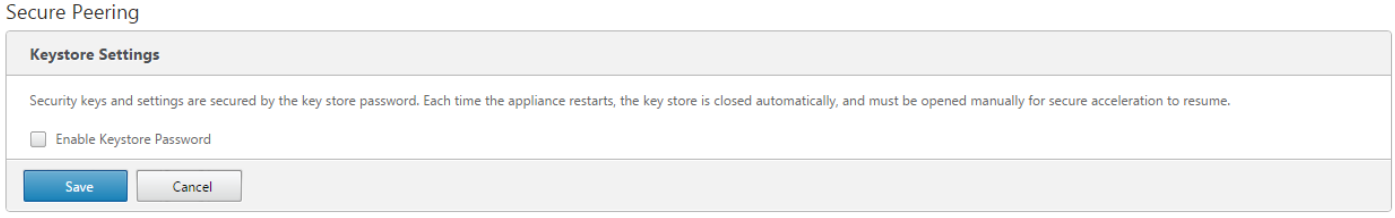

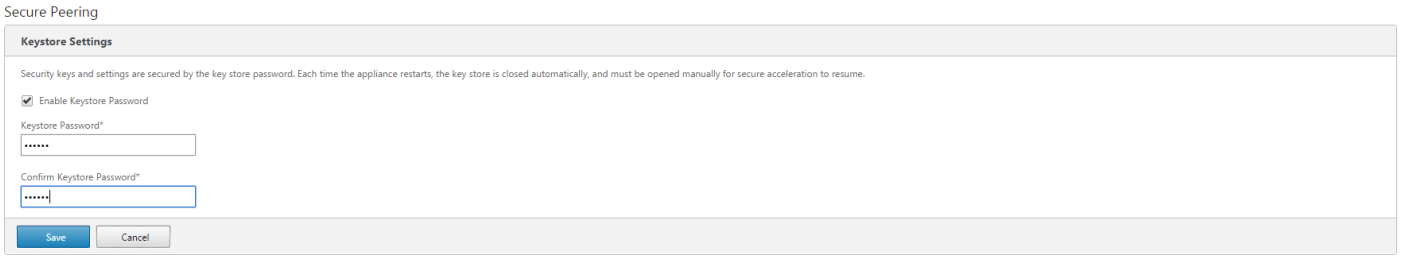

4. Enable secure peering by selecting **CA Certificate** radio button and providing uploaded CA and CA Key pair certificates appropriately as shown below.

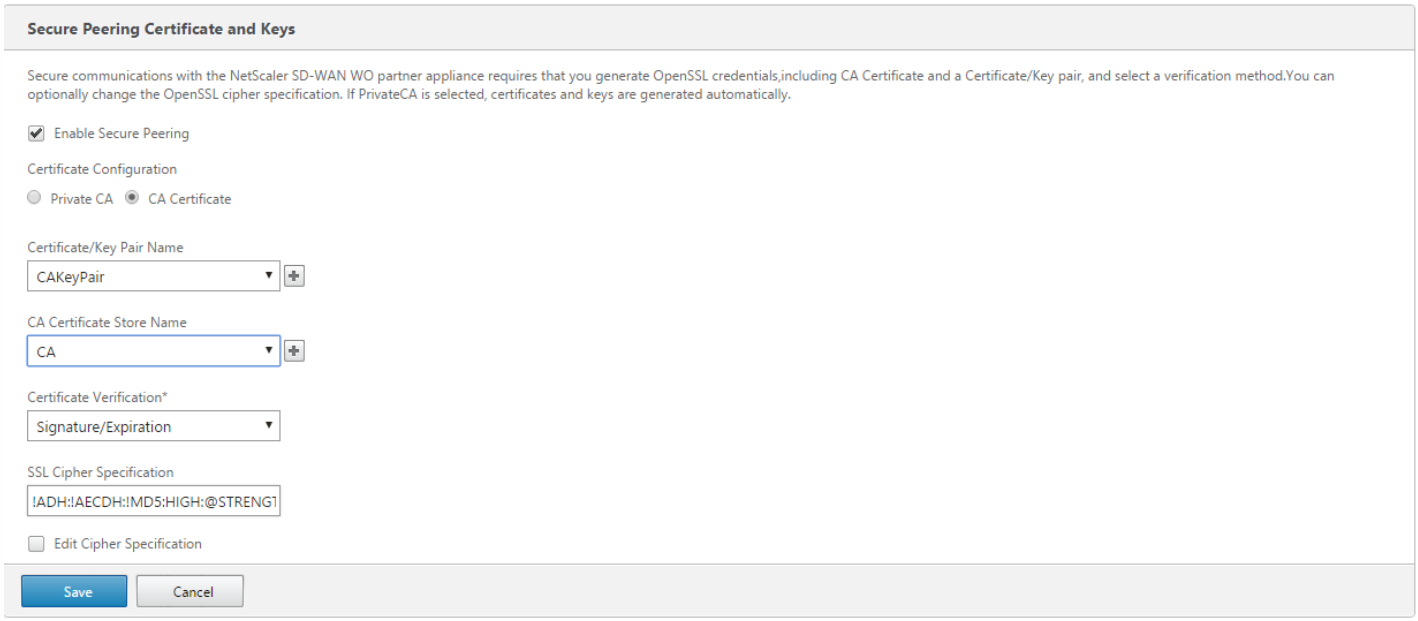

5. Provide Remote machine's Virtual IP along with Port 443 as shown below.

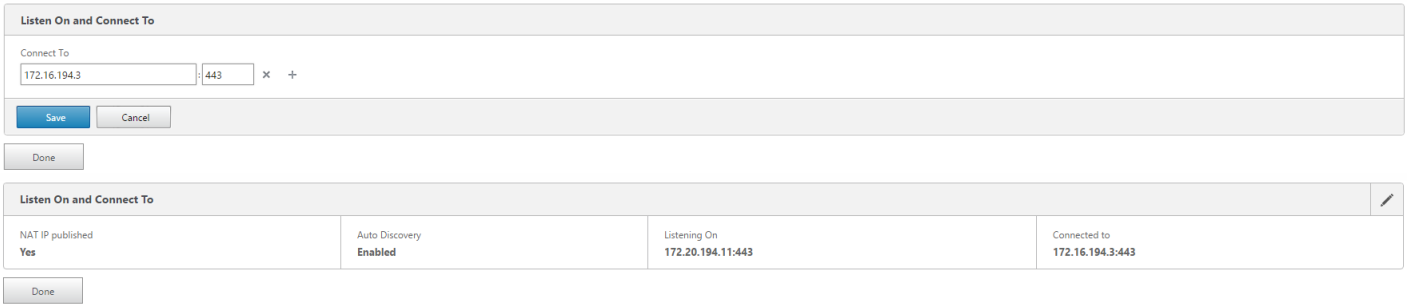

Monitoring

1. View Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **WAN Optimization** > **Partners** page.

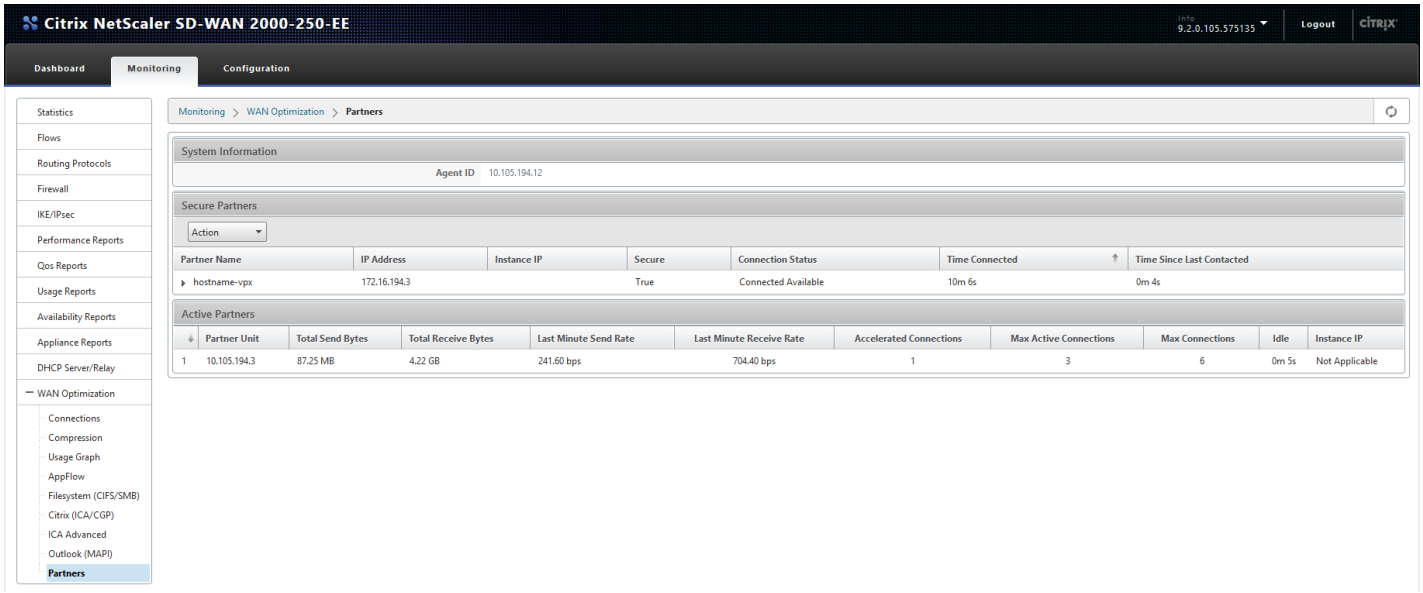

2. On partner appliance, View Secure Partner Information on the Enterprise Edition appliance under **Monitoring** > **Partners** > **Secure Partners** page.

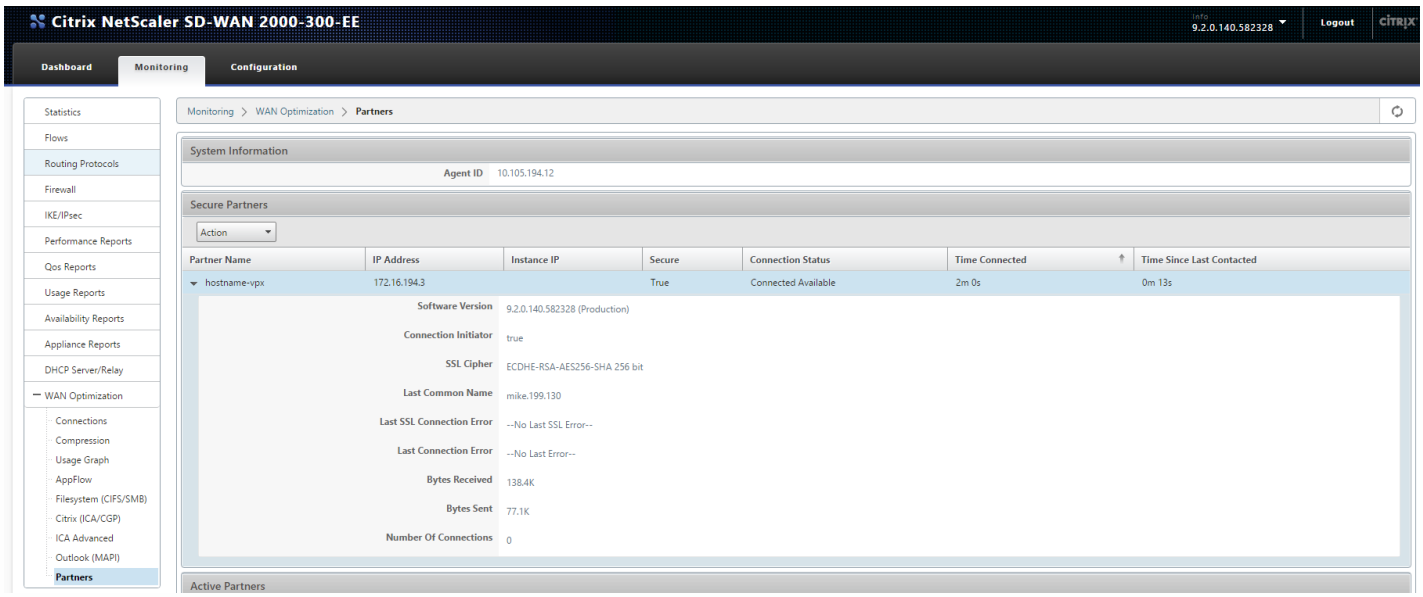

Troubleshooting

1. View **Secure Partner Success** / **Failure Information** on the Enterprise Edition Appliance under **Monitoring** > **WAN Optimization** > **Partners** > **Secure Partners** page.

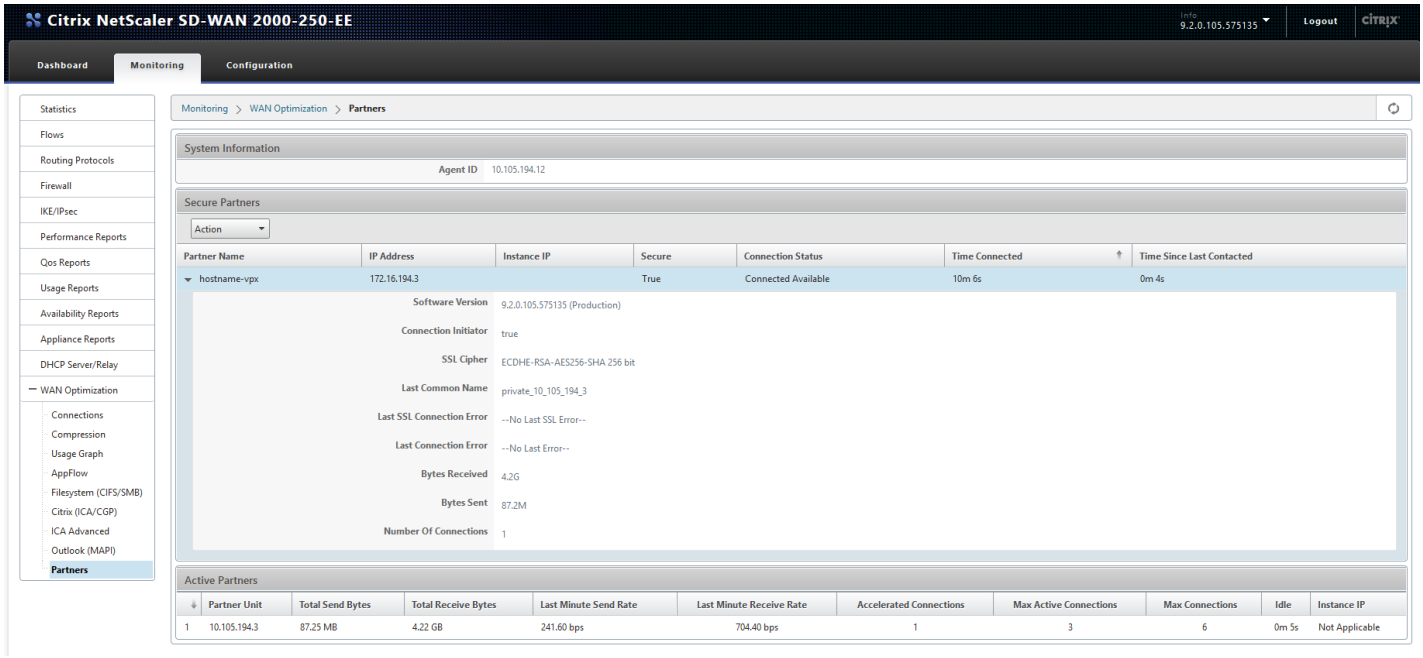

2. On partner appliance, view **Secure Partner Information** on the Enterprise Edition appliance under **Monitoring** > **Appliance Performance** > **Logging** page.

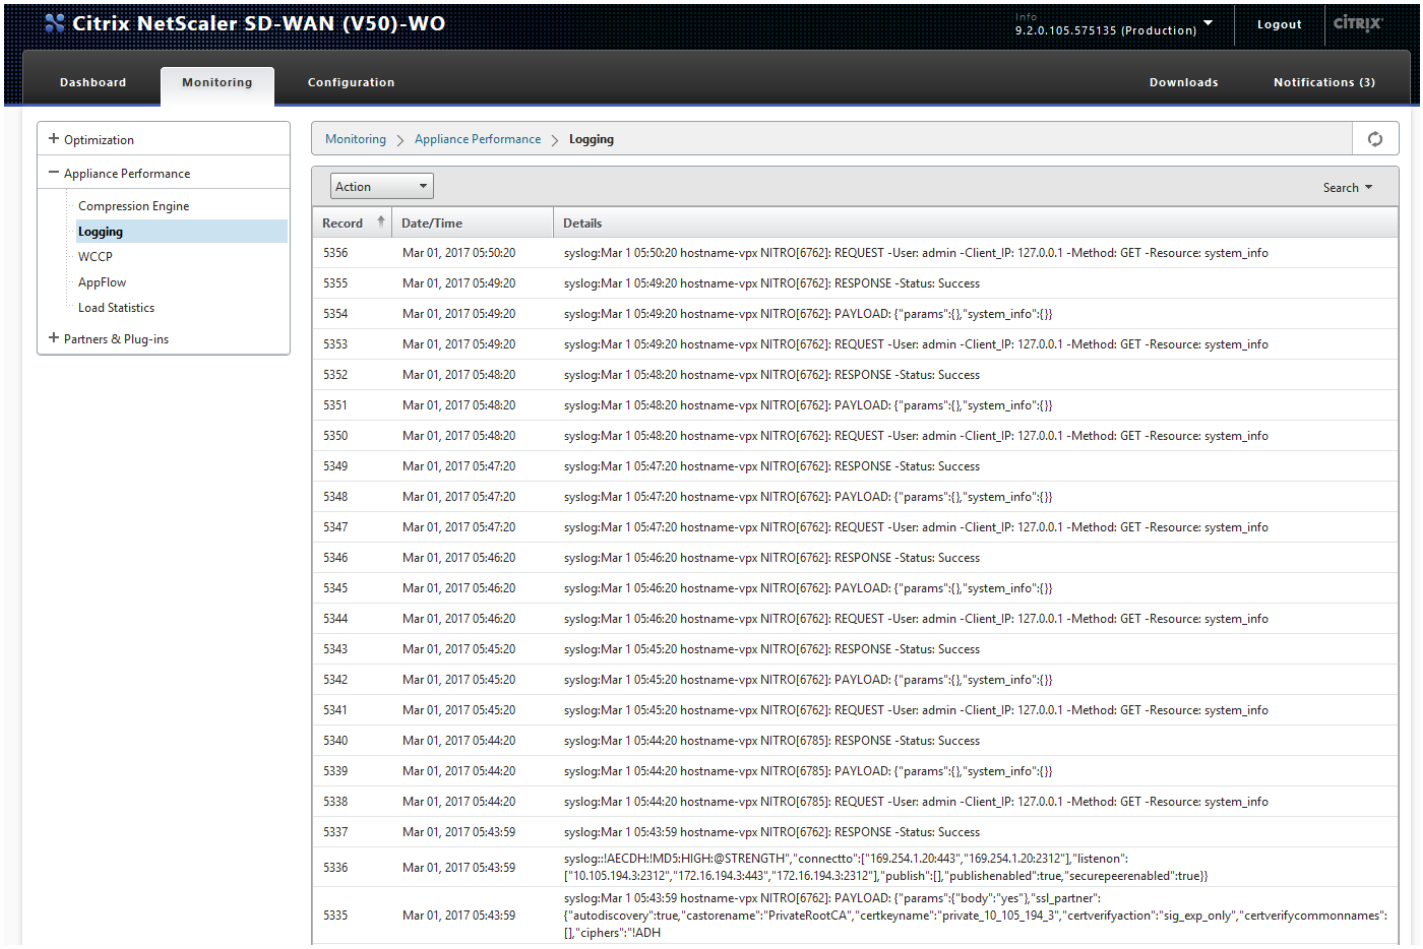

Domain Join and Delegate User Creation

Aug 09, 2017

Domain Join

To configure new EE appliance at the DC to windows domain:

1. Go to Windows Domain in SD-WAN web GUI, navigate to **Configuration** > **Secure Acceleration** > and click **Join Windows Domain**.

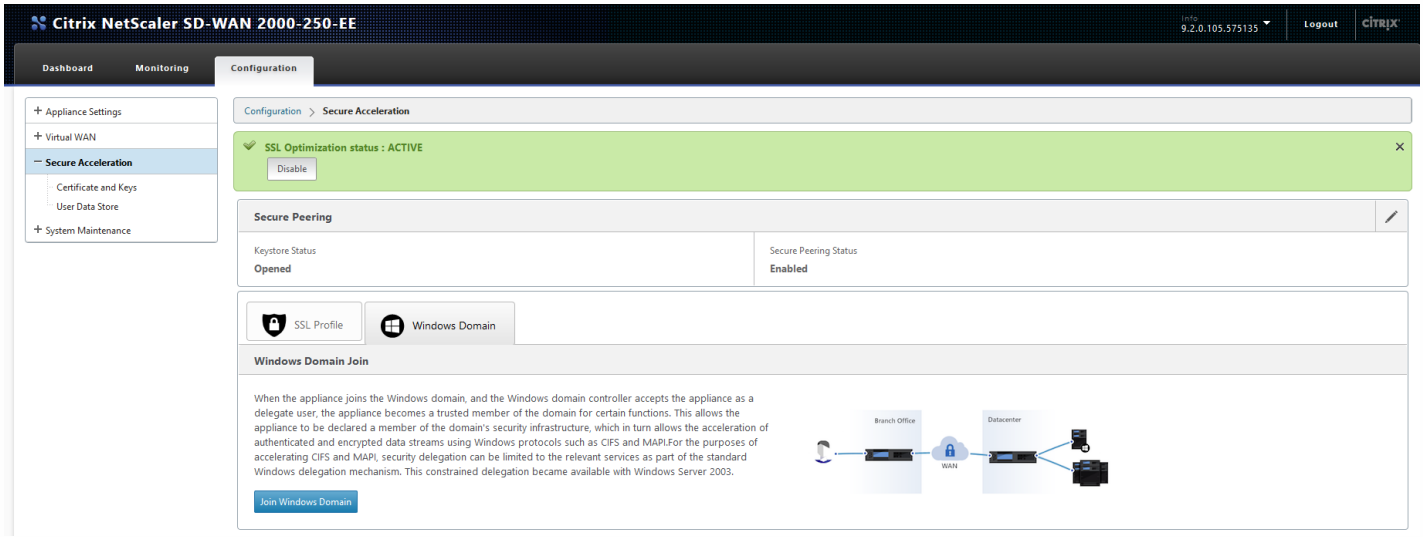

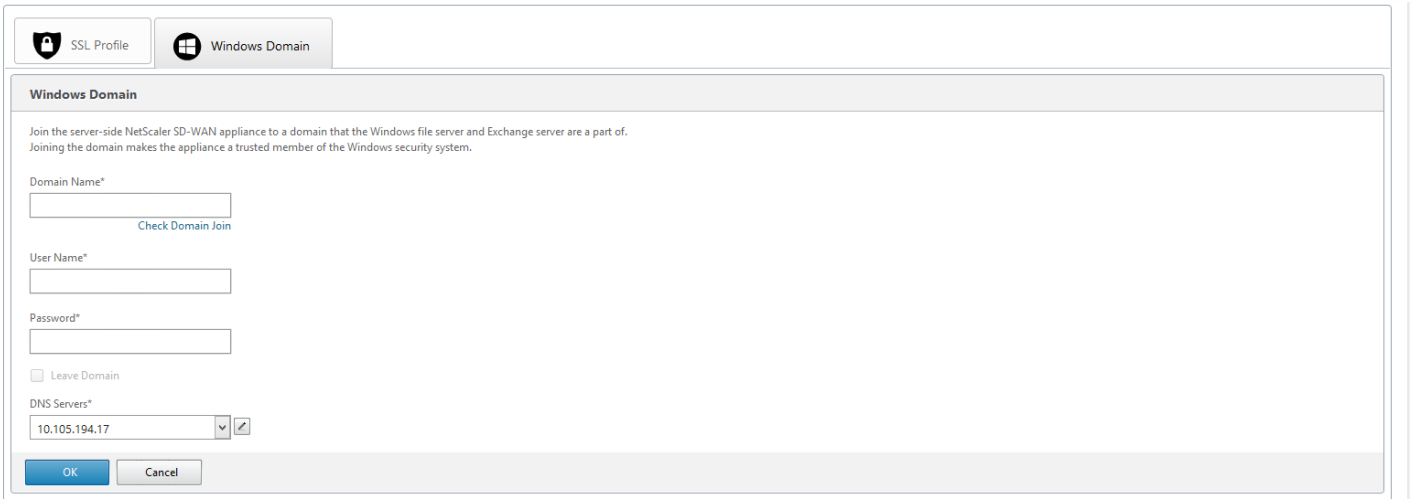

2. Provide **Windows domain name** and perform **Domain Join** pre-checks.

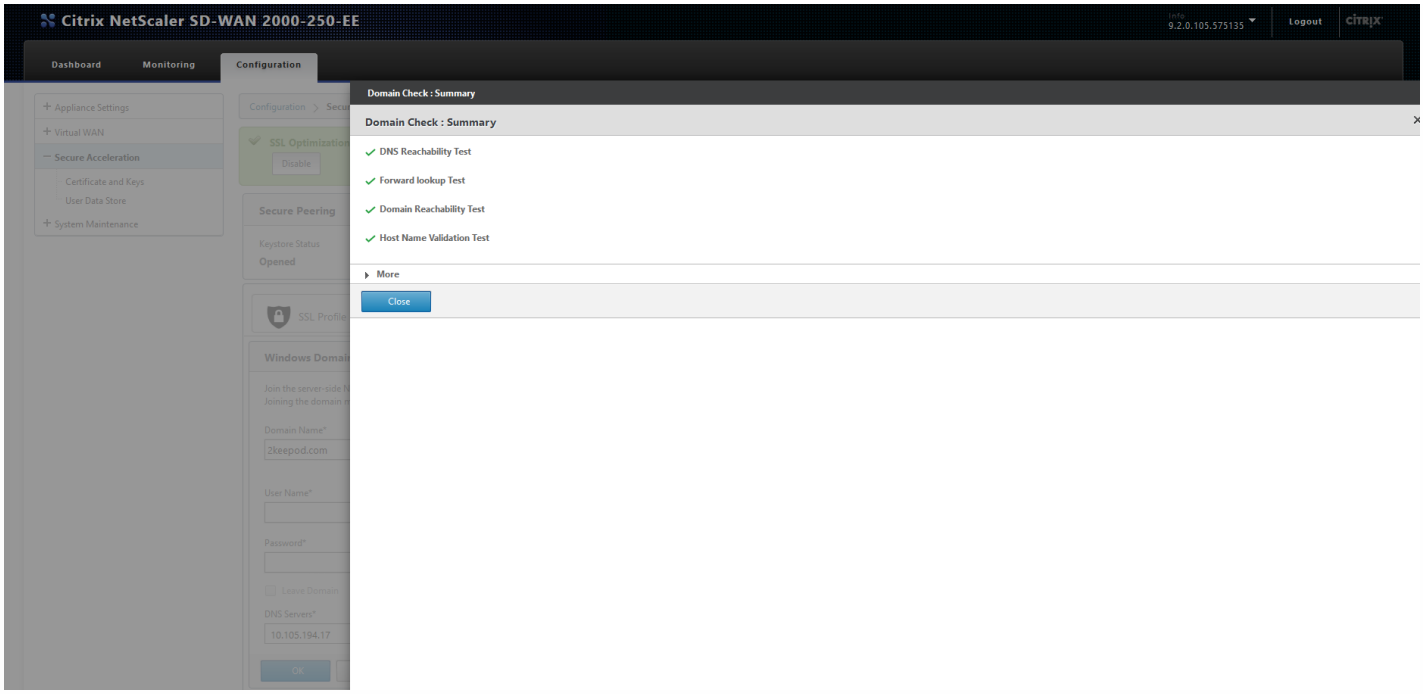

3. After pre-check summary shows as successful, enter domain controller's credentials.

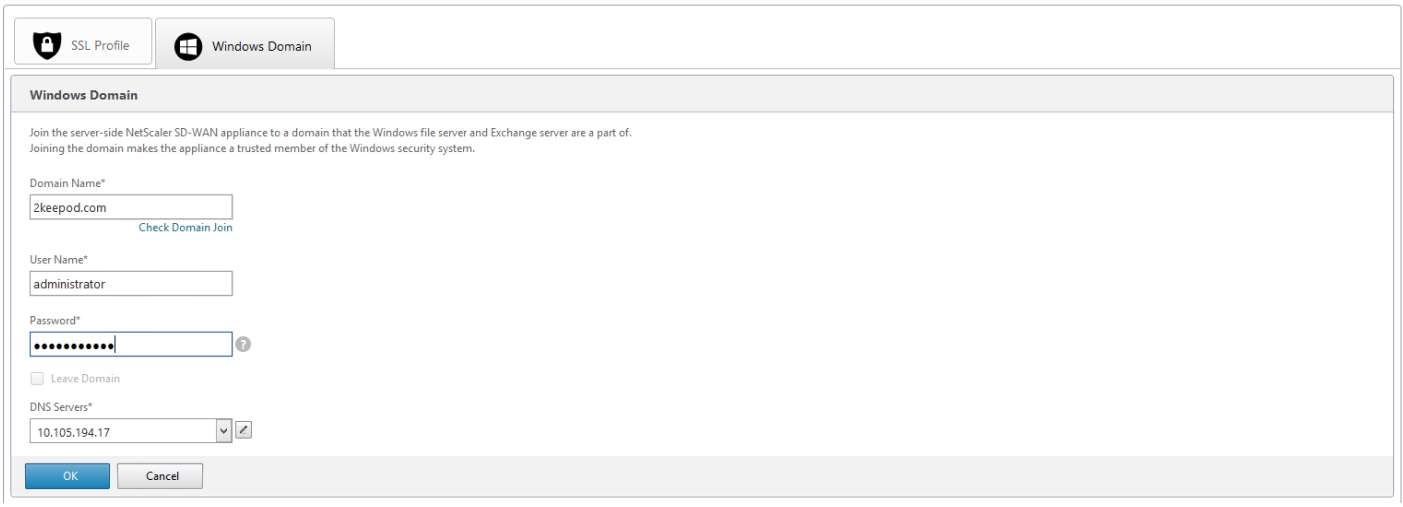

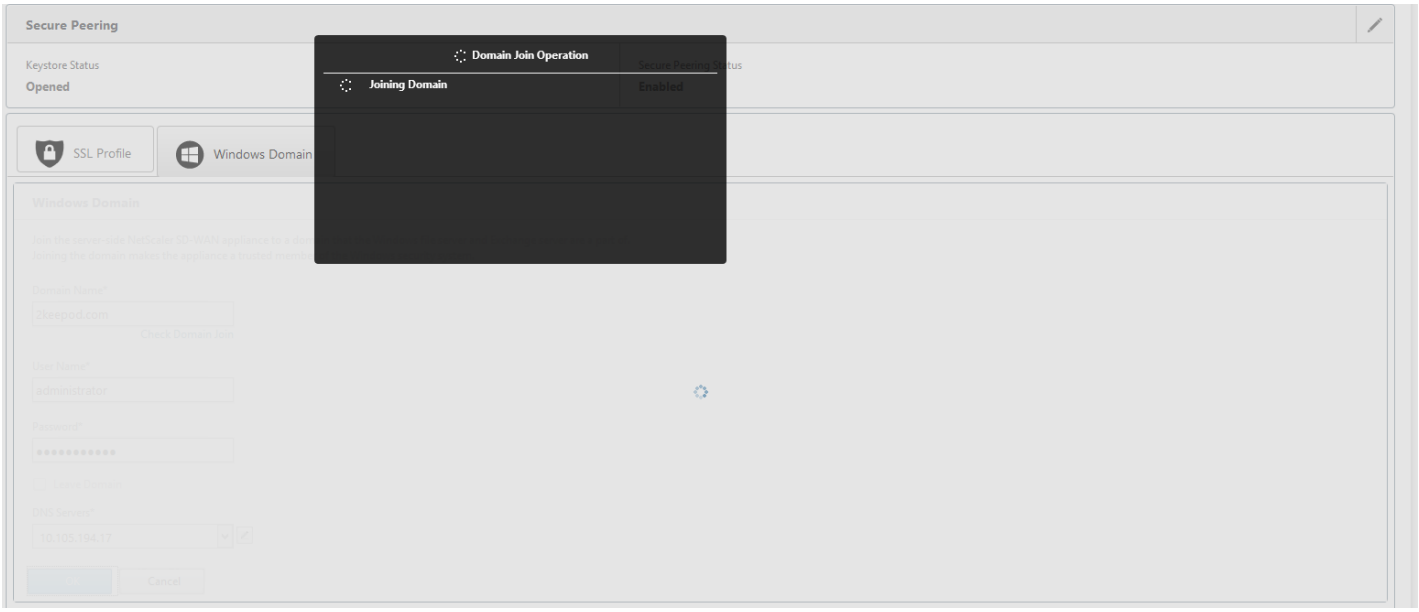

4. On successful domain join, you get the following output.

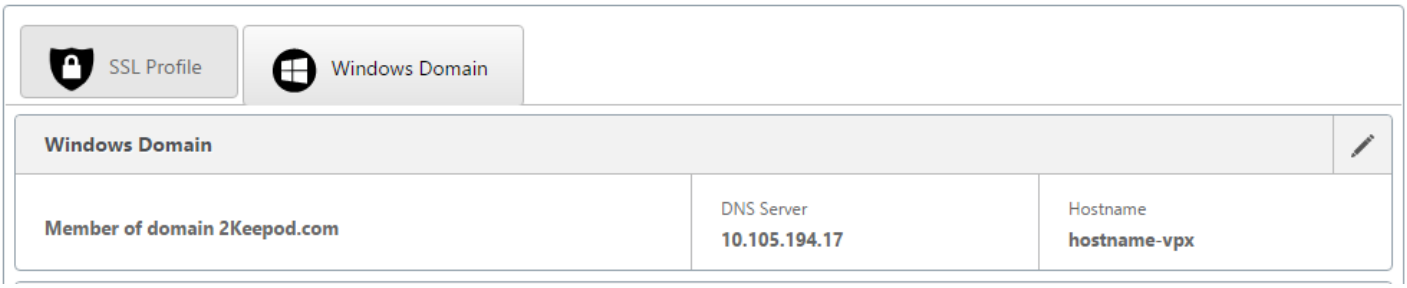

Delegate User

1. Add delegate user to delegate the services as shown below.

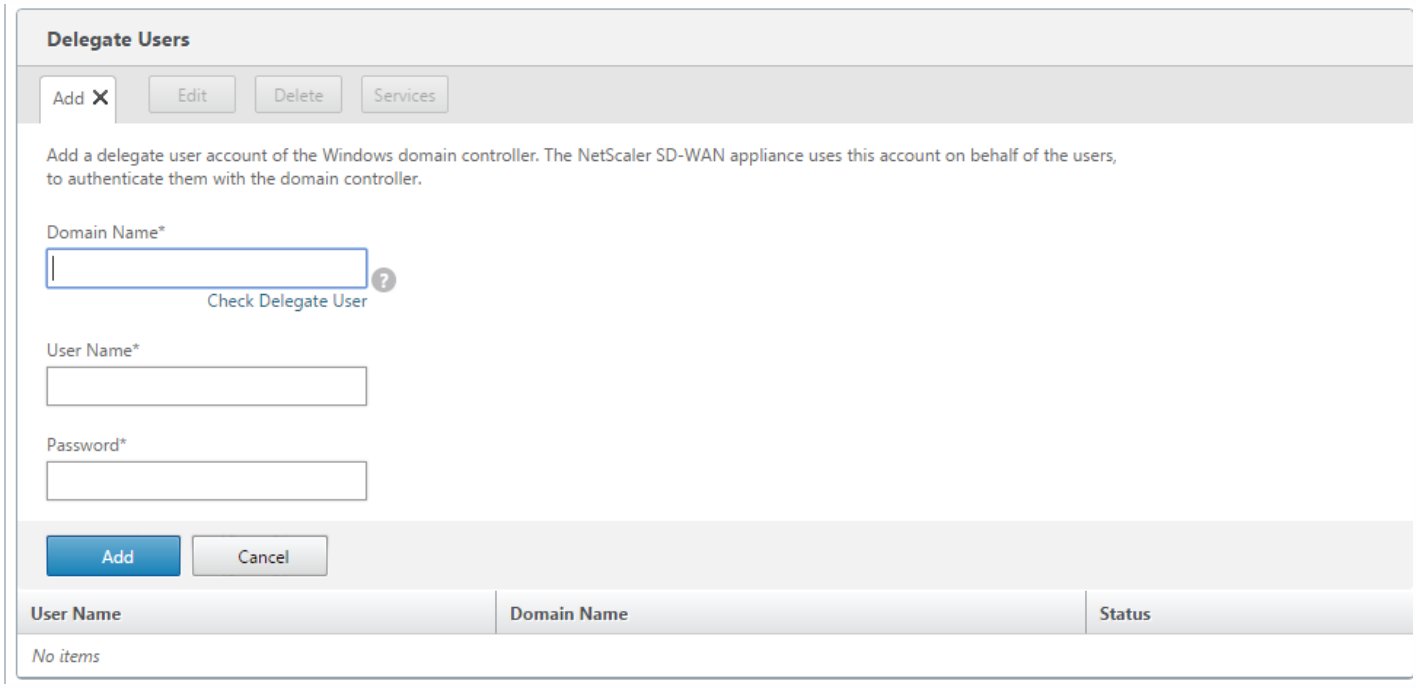

2. Provide correct domain Name and perform delegate user pre-check.

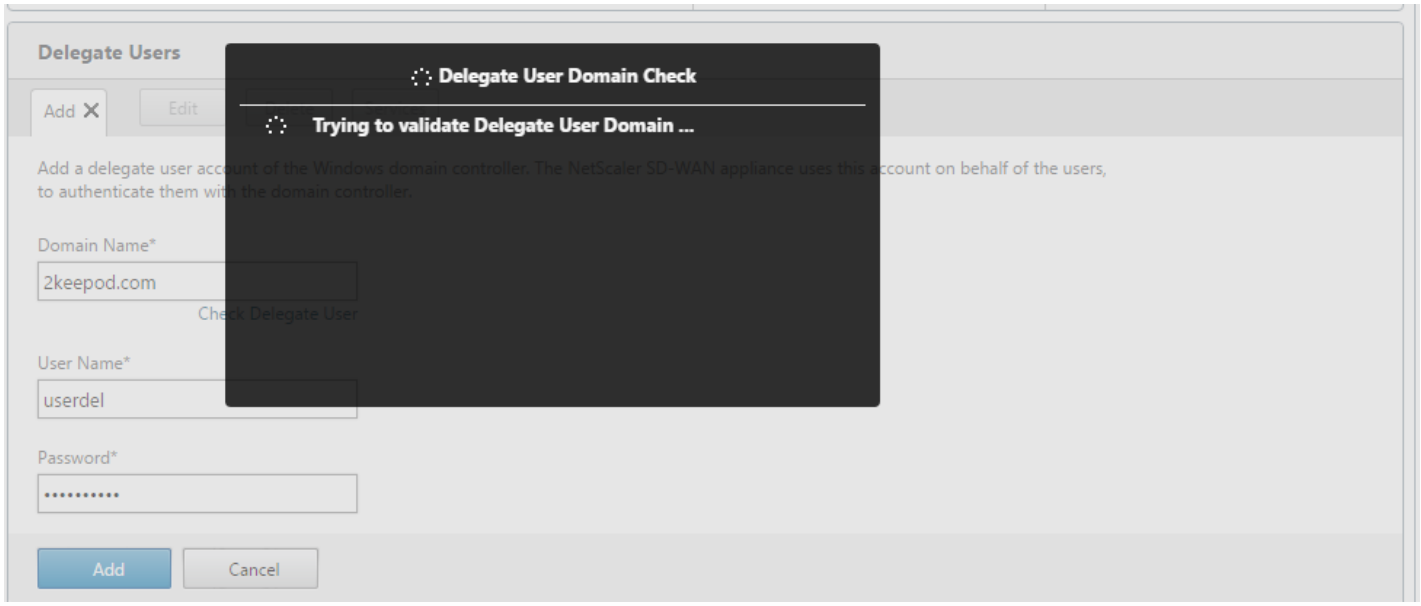

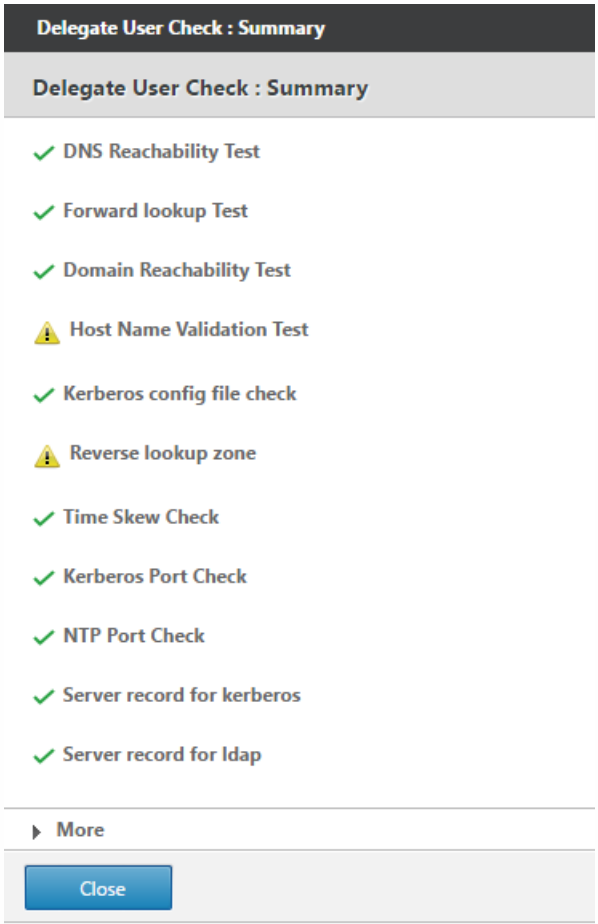

3. After delegate user pre-checks are successful, provide valid credentials of the delegate user.

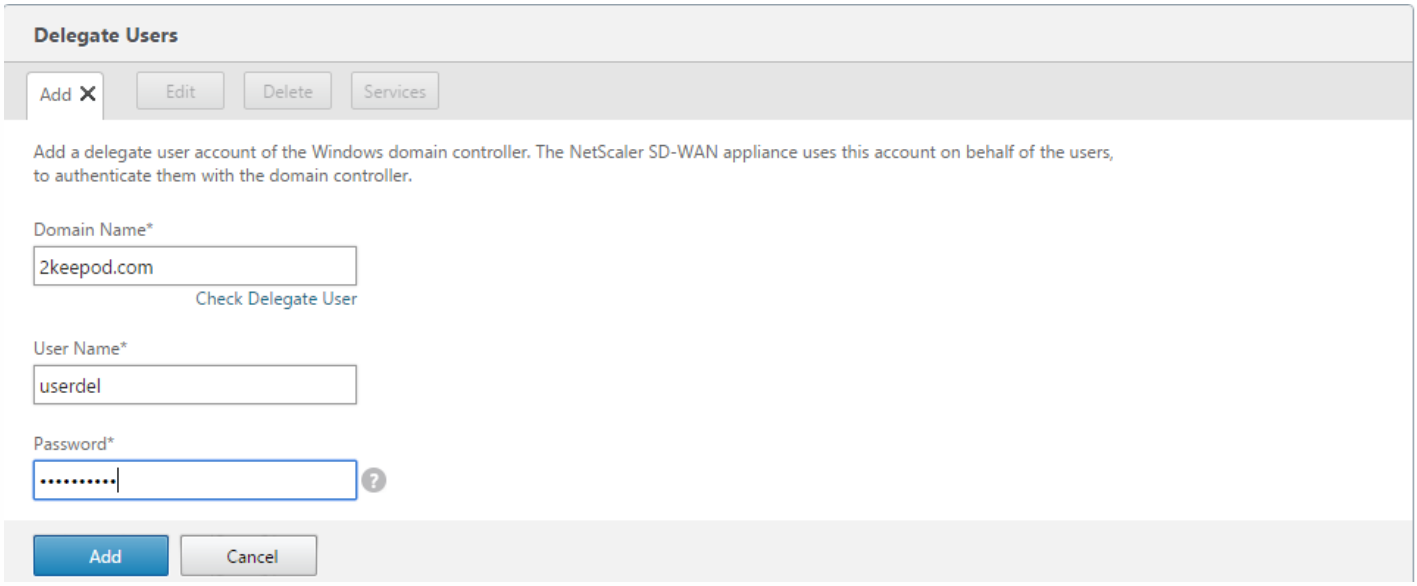

4. After delegate user is added successfully to SD-WAN, you will see a success message.

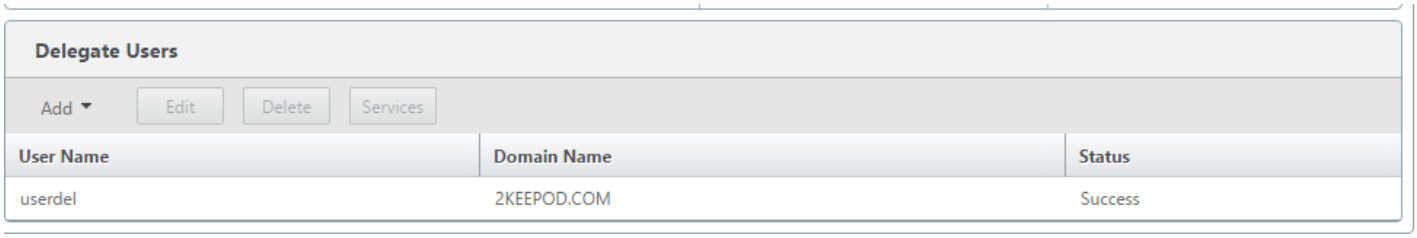

5. To check what all services are delegated by the delegate user, point to the user and select services.

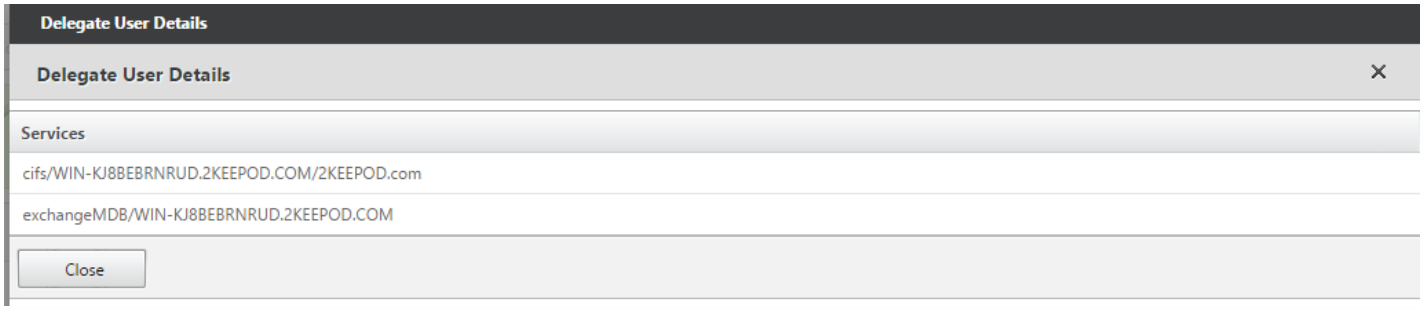

SNMPv3 Polling and Trap Capability

Aug 09, 2017

NetScaler SD-WAN supports only a single user account for each SNMPv3 capability. This restriction provides the following advantages:

- Ensuring SNMPv3 compliance for network devices
- Verification of SNMPv3 capability
- Easy configuration of SNMPv3

To configure SNMPv3 Polling and Traps, navigate to the SNMPv3 section of the **Integrate** > **Configure Events** and **Alerts** page and fill in the fields as required.

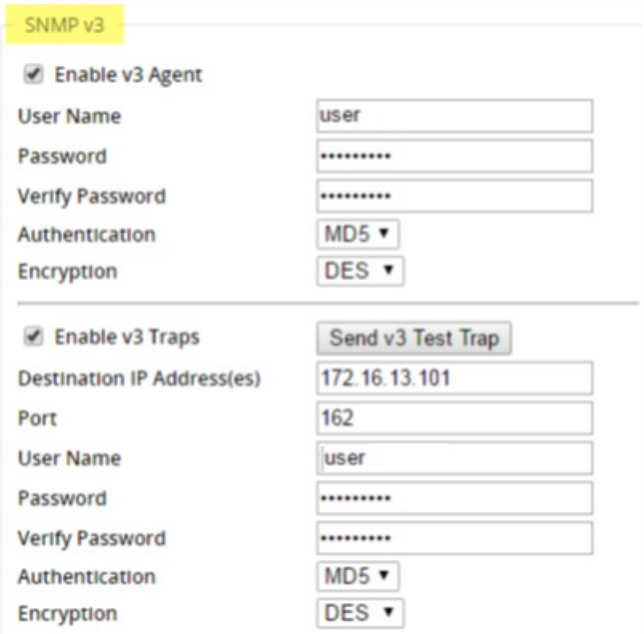

Zero Touch Deployment

Jan 17, 2018

> **Note**
>
> The Zero Touch Deployment service is supported only on select NetScaler SD-WAN appliances:
>
> - NetScaler SD-WAN 210 Standard Edition
> - NetScaler SD-WAN 410 Standard Edition
> - NetScaler SD-WAN 2100 Standard Edition
> - NetScaler SD-WAN 1000 Standard Edition (reimage required)
> - NetScaler SD-WAN 1000 Enterprise Edition (reimage required)
> - NetScaler SD-WAN 2000 Standard Edition (reimage required)
> - NetScaler SD-WAN 2000 Enterprise Edition (reimage required)
> - NetScaler SD-WAN AWS VPX instance

Zero Touch Deployment (ZTD) Cloud Service is a Citrix operated and managed cloud-based service which allows discovery of new appliances in the NetScaler SD-WAN network, primarily focused on streamlining the deployment process for NetScaler SD-WAN at branch or cloud service office locations. The ZTD Cloud Service is publicly accessible from any point in a network via public Internet access. The ZTD Cloud Service is accessed over Secure Socket Layer (SSL) Protocol.

The ZTD Cloud Services securely communicates with backend Citrix services hosting stored identification of Citrix customers who have purchased Zero Touch capable devices (e.g. NetScaler SD-WAN 410-SE, 2100-SE). The backend services are in place to authenticate any Zero Touch Deployment request, properly validating association between the Customer Account and the Serial Numbers of NetScaler SD-WAN appliances.

ZTD High-Level Architecture and Workflow

Data Center Site:

NetScaler SD-WAN Administrator – A user with Administration rights of the NetScaler SD-WAN environment with the following primary responsibilities:

- Configuration creation using NetScaler SD-WAN Center Network Configuration tool, or import of configuration from the Master Control Node (MCN) SD-WAN appliance
- Citrix Cloud Login to initiate the Zero Touch Deployment Service for new site node deployment.

> **Note**

If your SD-WAN Center is connected to the internet through a proxy server, you have to configure the proxy server settings on the SD-WAN Center. For more information, see How to Configure Proxy Server Settings for Zero Touch Deployment.

Network Administrator – A user responsible for Enterprise network management (DHCP, DNS, internet, firewall, etc.)

- If required, configure firewalls for outbound communication to FQDN ***sdwanzt.citrixnetworkapi.net*** from SD-WAN Center.

Remote Site:

Onsite Installer – A local contact or hired installer for on-site activity with the following primary responsibilities:

- Physically unpack the NetScaler SD-WAN appliance
- Reimage non-ZTD ready appliances

  - Required for: NetScaler SD-WAN 1000-SE, 2000-SE, 1000-EE, 2000-EE

  - Not required for: NetScaler SD-WAN 410-SE, 2100-SE

- Power cable the appliance
- Cable the appliance for internet connectivity on the Management interface (e.g. MGMT, or 0/1)
- Cable the appliance for WAN link connectivity on the Data interfaces (e.g. apA.WAN, apB.WAN, apC.WAN, 0/2, 0/3, 0/5, etc)

Note

The interface layout will be different each model, so please reference the documentation for identification of data and management ports.

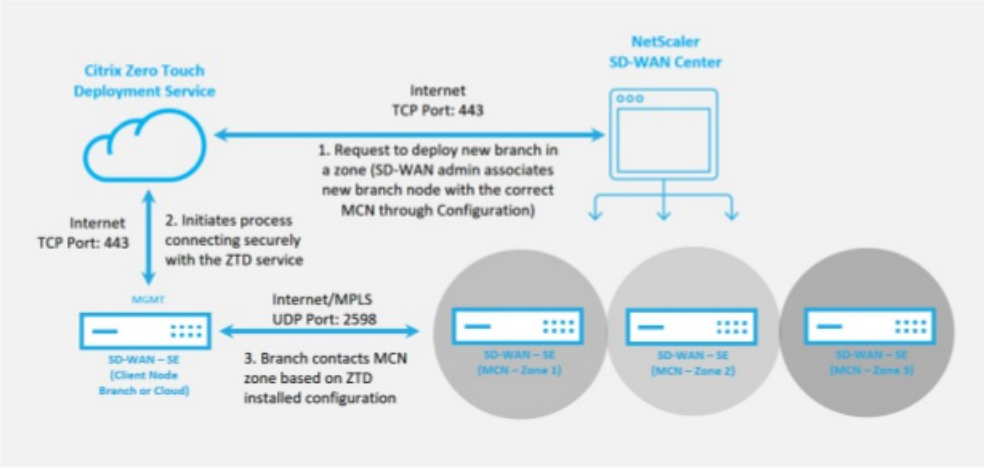

The following prerequisites are required before starting any Zero Touch Deployment service:

- Actively running NetScaler SD-WAN promoted to Master Control Node (MCN).
- Actively running NetScaler SD-WAN Center with connectivity to the MCN through Virtual Path.
- Citrix Cloud Login credentials created on https://onboarding.cloud.com (reference the instruction below on account

creation).

- Management network connectivity (SD-WAN Center and SD-WAN Appliance) to the Internet on port 443, either directly or through a proxy server.
- (optional) At least one actively running NetScaler SD-WAN appliance operating at a branch office in Client Mode with valid Virtual Path connectivity to MCN to help validate successful path establishment across the existing underlay network.

The last prerequisite is not a requirement, but allows the NetScaler SD-WAN Administrator to validate that the underlay network will successfully allow Virtual Paths to be successfully established as soon as the Zero Touch Deployment is complete with any newly added site. Primarily, this validates that the appropriate Firewall and Route policies are in place to either NAT traffic accordingly or confirm ability for UDP port 4980 can successfully penetrate the network to reach the MCN.

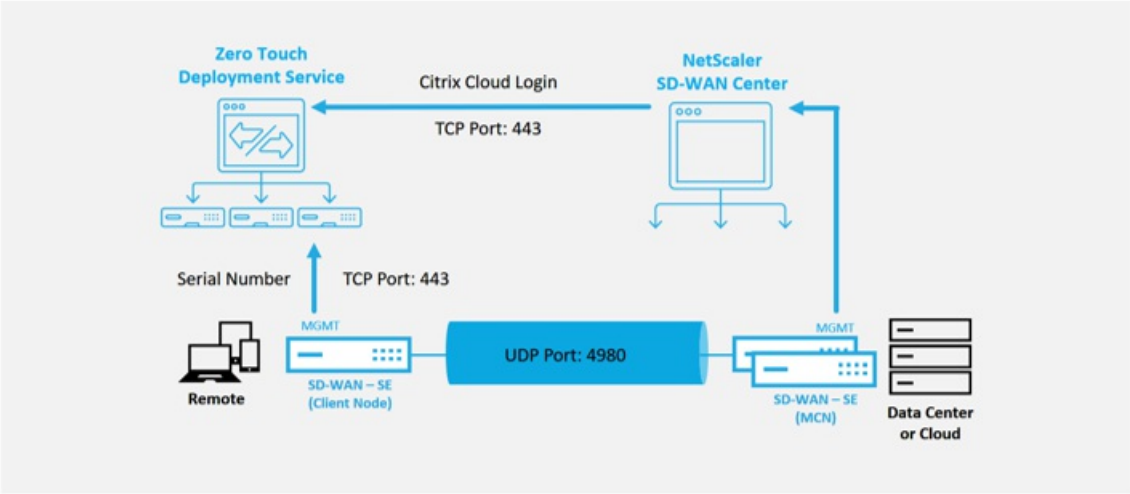

Zero Touch Deployment Service Overview

The Zero Touch Deployment Service works in tandem with the NetScaler SD-WAN Center to provide an easier deployment of branch office SD-WAN appliances. SD-WAN Center is configured and used as the central management tool for the SD-WAN Standard and Enterprise Edition appliances. To utilize the Zero Touch Deployment Service (or ZTD Cloud Service), an Administrator must begin by deploying the first NetScaler SD-WAN device in the environment, then configure and deploy the SD-WAN Center as the central point of management. When the SD-WAN Center, release 9.1 or later, is installed with connectivity to the public internet on port 443, SD-WAN Center will automatically call home to the Cloud Service and install necessary components to unlock the Zero Touch Deployment features and to make the Zero Touch Deployment option available in the GUI of SD-WAN Center. Zero Touch Deployment is not available by default in the SD-WAN Center software. This is purposely designed to make sure the proper preliminary components on the underlay network are present before allowing an Administrator to initiate any on-site activity involving Zero Touch Deployment.

After a working SD-WAN environment is up and running registration into the Zero Touch Deployment Service is accomplished through creating a Citrix Cloud account login. With SD-WAN Center able to communicate with the ZTD service, the GUI will expose the Zero Touch Deployment options under the Configuration tab. Logging into the Zero Touch Service authenticates the Customer ID associated with the particular NetScaler SD-WAN environment and registers the SD-WAN Center, in addition to unlocking the account for further authentication of ZTD appliance deployments.

Using the Network Configuration tool in NetScaler SD-WAN Center, the SD-WAN Administrator will then need to utilize the templates or clone site capability to build out the SD-WAN Configuration to add new sites. The new configuration will be used by the SD-WAN Center to initiate the deployment of ZTD for the newly added sites. When the SD-WAN Administrator initiates a site for deployment using the ZTD process, he or she will have the option to pre-authenticate the appliance to be used for ZTD by pre-populating the serial number, and initiating email communication to on-site installer to begin on-site activity.

The Onsite Installer will receive email communication that the site is ready for Zero Touch Deployment and can begin the installation procedure of powering on and cabling the appliance for DHCP IP address assignment and internet access on the MGMT port. Also, cabling in any LAN and WAN ports. Everything else will be automated by the ZTD Service and progress can be monitored by the utilizing the activation URL. In the event the remote node to be installed is a cloud instance, opening up the activation URL will begin the workflow to automatically install the instance in the designated cloud environment, no action is needed by a local installer.

The Zero Touch Deployment Cloud Service will automate the following actions:

Download and Update the ZTD Agent if new features are available on the branch appliance.

- Authenticate the branch appliance by validating the serial number
- Authenticate that the SD-WAN Administrator accepted the site for ZTD using the SD-WAN Center
- Pull the configuration file specific for the targeted appliance from the SD-WAN Center
- Push the configuration file specific for the targeted appliance to the branch appliance
- Install the configuration file on the branch appliance
- Push any missing SD-WAN software components or required updates to the branch appliance
- Push a temporary 10Mbps license file for confirmation of Virtual Path establishment to the branch appliance
- Enable the SD-WAN Service on the branch appliance

Additional steps are required of the SD-WAN Administrator to install a permanent license file on the appliance.

Zero Touch Deployment Service Procedure

The following procedure detail the steps required to successfully deploy a new site using the Zero Touch Deployment Service. It is recommended to have a running MCN and one client node already working with proper communication to NetScaler SD-WAN Center, as well as established Virtual Paths confirming connectivity across the underlay network.

The following steps are required of the SD-WAN Administrator to initiate the deployment of zero touch:

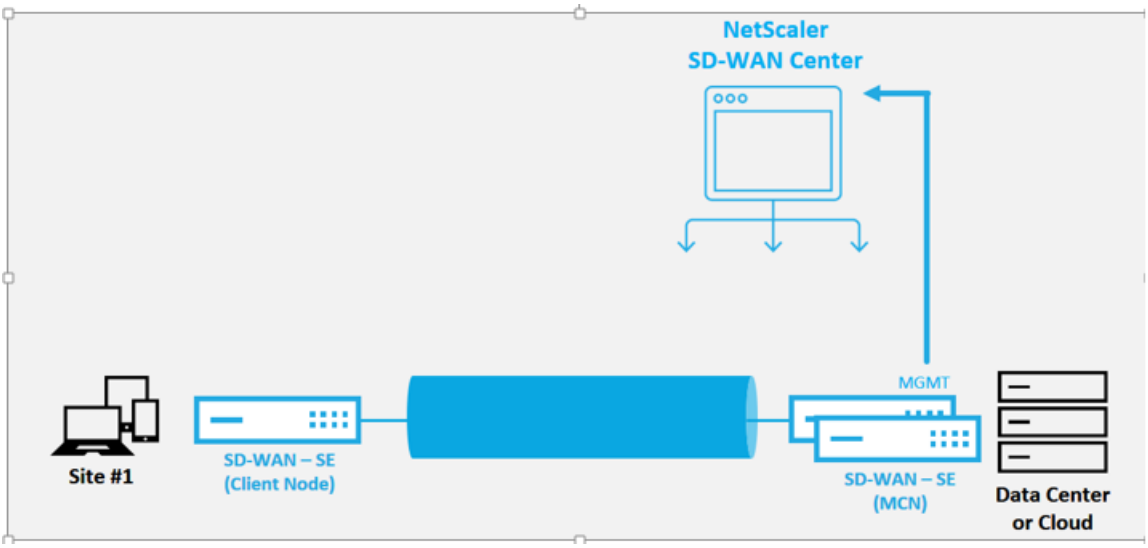

How to Configure Zero Touch Deployment Service

The SD-WAN Center has the functionality to accept requests from newly connected appliances to join the SD-WAN Enterprise network. The request is forwarded to the web interface through the zero touch deployment service. Once the appliance connects to the service, configuration and software upgrade packages are downloaded.

Configuration worflow:

- Access **SD-WAN Center** > **Create New site configuration** or Import existing configuration and save it.
- Login to Citrix Workspace Cloud to enable ZTD service. The Zero Touch Deployment menu option is now displayed in the SD-WAN center web management interface.
- In SD-WAN Center, navigate to **Configuration** > **Zero Touch Deployment** > **Deploy New Site**.
- Select an appliance, click **Enable** and click **Deploy**.
- Installer receives activation email > Enter the serial number > **Activate** > Appliance is deployed successfully.

To configure Zero Touch Deployment service:

1. Install SD-WAN Center with enabled Zero Touch Deployment capabilities:

    a) Install NetScaler SD-WAN Center with DHCP assigned IP address.

    b) Verify SD-WAN Center is assignment a proper management IP address and network DNS address with connectivity to the public internet across the management network.

    c) Upgrade the NetScaler SD-WAN Center to the latest 9.2 firmware.

    d) With proper internet connectivity, the SD-WAN Center will call home to the Zero Touch Deployment (ZTD) Cloud Service and automatically download and install any firmware updates specific to ZTD, if this call home procedure fails the following Zero Touch Deployment option will not be available in the GUI.

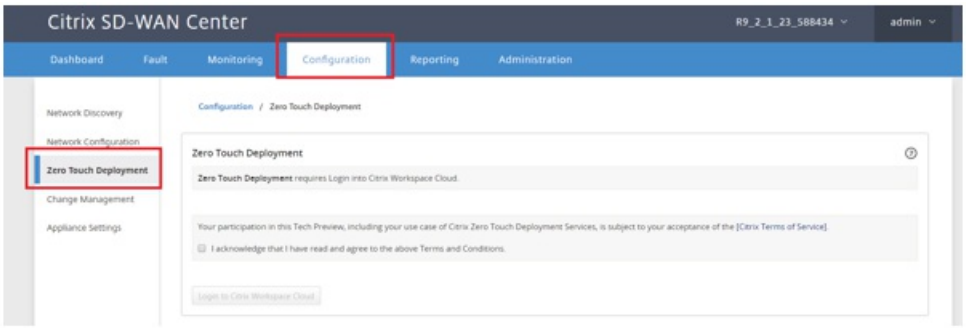

e) Read the Terms and Conditions, and then select "**I acknowledge that I have read and agree to the above Terms and Conditions**."

f) Click the "**Login to Citrix Workspace Cloud**" button if a Citrix Cloud account has leady been created.

g) Login into the Citrix Cloud account, and upon receiving the following message of successful login, **PLEASE DO NOT CLOSE THIS WINDOW UP, THE PROCESS REQUIRES ANOTHER ~20 SECONDS FOR THE SD-WAN CENTER GUI TO BE REFRESHED.** The window should close on its own when it is complete.

h) To create a Cloud Login account follow the below procedure:

- Open a web browser to https://onboarding.cloud.com

- Click on the link for "**Wait, I have a Citrix.com account**".

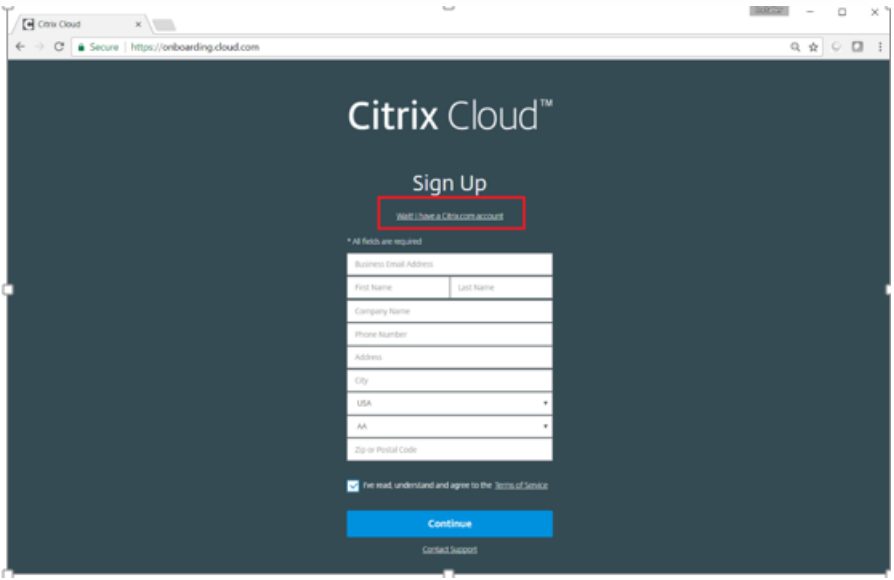

i) Sign in with an existing Citrix account.

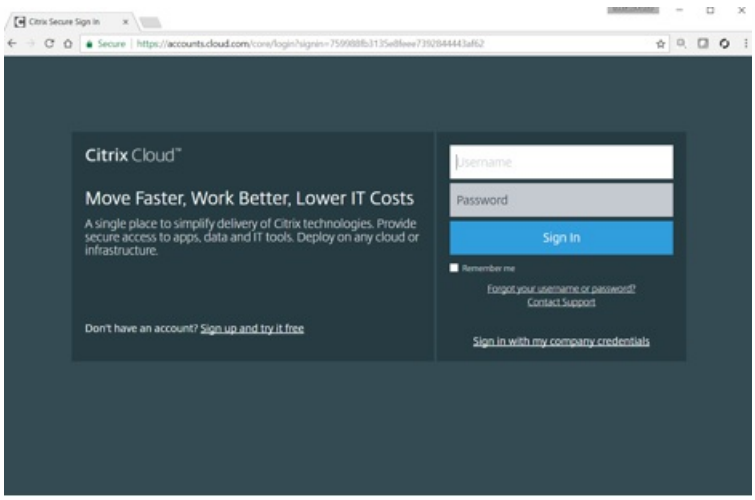

j) Once logged into SD-WAN Center Zero Touch Deployment page, you may notice no sites are available for ZTD deployment, this could be because of the following reasons:

- The active configuration has not been selected from the Configuration drop-down menu
- All the sites for the current active configuration have already been deployed
- The configuration was not built using the SD-WAN Center, but rather the Configuration Editor available on the MCN
- Sites were not built in the configuration referencing zero touch capable appliances (e.g. 410-SE, 2100-SE, Cloud VPX)

2. Update the configuration to add a **new remote** site with a **ZTD capable SD-WAN appliance** using SD-WAN Center Network Configuration.

If the SD-WAN configuration was not built using the SD-WAN Center Network Configuration, import the active configuration from the MCN and begin modifying the configuration using SD-WAN Center. For Zero Touch Deployment capability, the SD-WAN Administrator must build the configuration using SD-WAN Center. The following procedure should be used to add a new site targeted for zero touch deployment.

a) Design the new site for SD-WAN appliance deployment by first outlining the details of the new site (i.e. Appliance Model, Interface Groups usage, Virtual IP Addresses, WAN Link(s) with bandwidth and their respective Gateways).

Important

You may notice any site node that has VPX selected as the model will also be listed, but currently ZTD support is only available for the AWS VPX instance.

Note

- Make sure that you are using a support web browser for Citrix SD-WAN Center
- Make sure the web browser is not blocking any pop-up windows during the Citrix Workspace Login

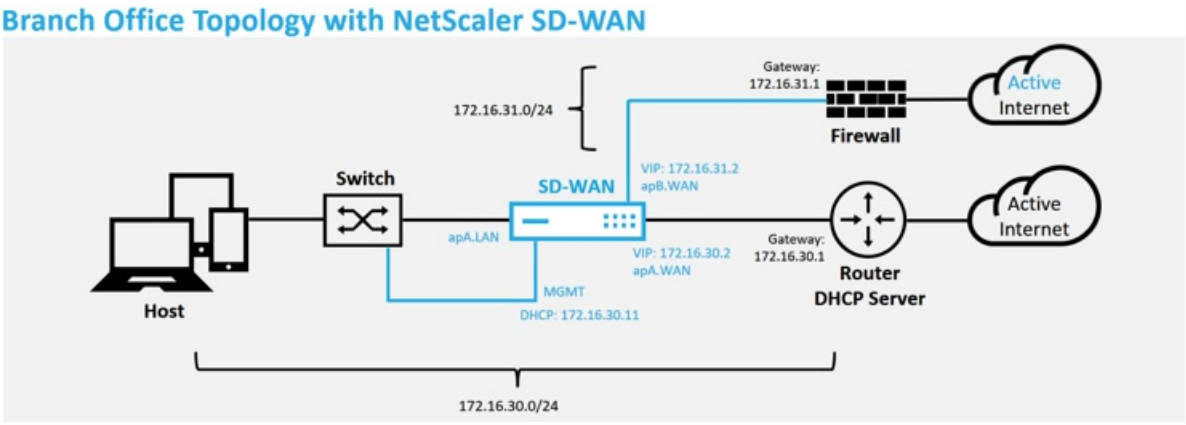

This is an example deployment of a branch office site, the NetScaler SD-WAN appliance is deployed physically in path of the existing MPLS WAN link across a 172.16.30.0/24 network, and leveraging an existing backup link by enabling it into an active state and terminating that second WAN link directly into the NetScaler SD-WAN appliance on a different subnet 172.16.31.0/24.

Note

The NetScaler SD-WAN appliances automatedly assign a default IP address of 192.168.100.1/16. With DHCP enabled by default, the DHCP Server in the network may provide the appliance a second IP address in a subnet that overlaps the default. This can possibly result in a routing issue on the appliance where the appliance may fail to connect to the ZTD Cloud Service. It is recommended to configure the DHCP server to assign IP addresses outside of the range of 192.168.0.0/16.

There are various different deployment modes available for NetScaler SD-WAN product placement in a network. In the above example, SD-WAN is being deployed as an overlay on top of existing networking infrastructure. For new sites, SD-WAN Administrators may choose to deploy the NetScaler SD-WAN in Edge or Gateway Mode deployment, eliminating the need for a WAN edge router and firewall, and consolidating the network needs of edge routing and firewall onto the NetScaler SD-WAN solution.

a) Open the SD-WAN Center web management interface and navigate to the **Configuration** > **Network Configuration** page.

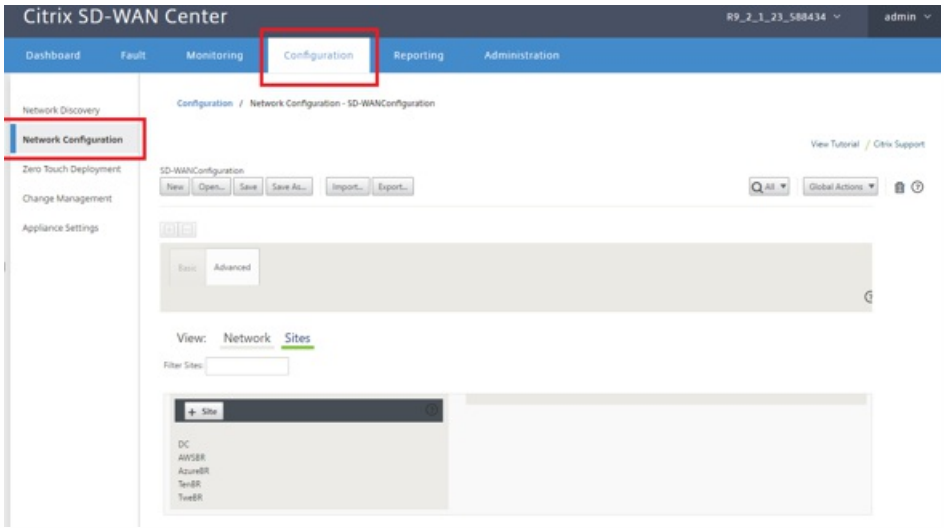

b) Make sure a working configuration is already in place, or import the configuration from the MCN.

c) Navigate to the Advanced tab to create a new site.

d) Open the Sites tile to display the currently configurated sites.

e) Quickly built the configuration for the new site by utilizing the clone feature of any existing site.

f) Populate all the required fields from the topology designed for this new branch site

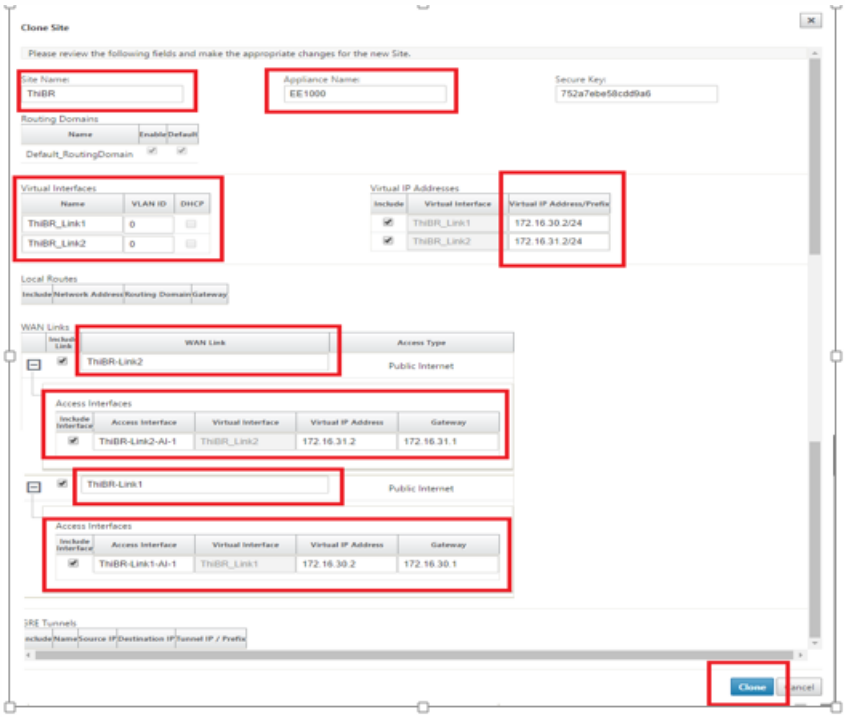

g) After cloning a new site, navigate to the site's **Basic Settings**, and verify that the Model of SD-WAN is correctly selected which would support the zero touch service.

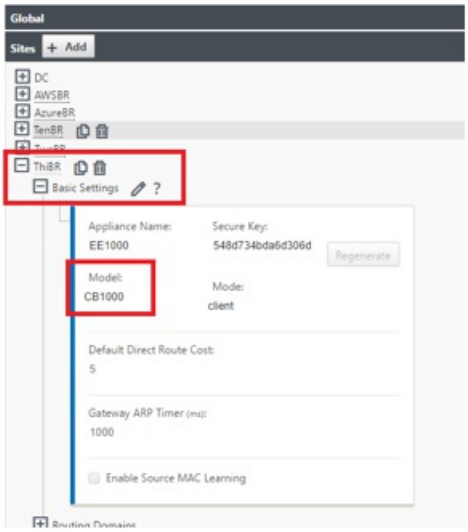

h) The SD-WAN model for the site can be updated, but do be aware that the Interface Groups may have to be redefined since the updated appliance may have a new interface layout then what was used to clone.

i) Save the new configuration on SD-WAN Center, and use the export to the "**Change Management inbox**" option to push the configuration using Change Management.

j) Follow the Change Management procedure to properly stage the new configuration, which makes the existing SD-WAN devices aware of the new site to be deployed via zero touch, you will need to utilize the "Ignore Incomplete" option to skip attempting to push the configuration to the new site that still needs to go through the ZTD workflow.

3. Navigate back to the SD-WAN Center Zero Touch Deployment page, and with the new active configuration running, the new site will be available for deployment.

a) In the Zero Touch Deployment page, under the **Deploy New Site** tab, select the running network configuration file

b) After the running configuration file is selected, the list of all the branch sites with undeployed NetScaler SD-WAN devices that are supported for zero touch will be displayed

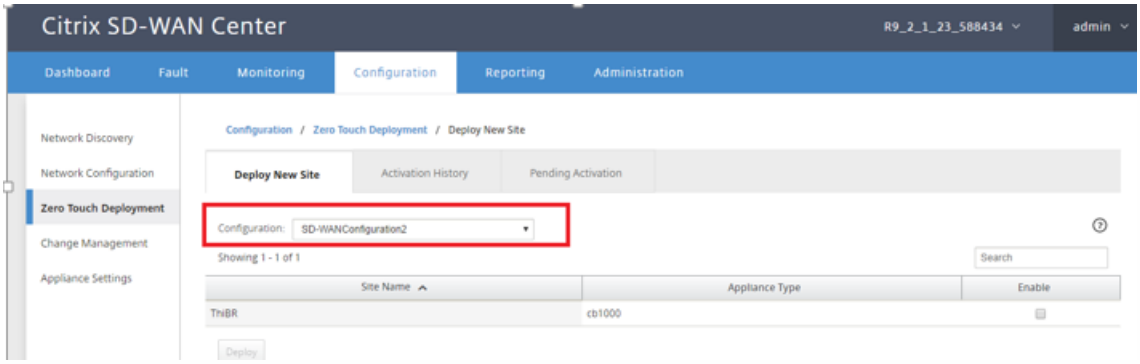

c). Select the branch sites you want to configure for Zero Touch service, click **Enable**, and then **Deploy**.

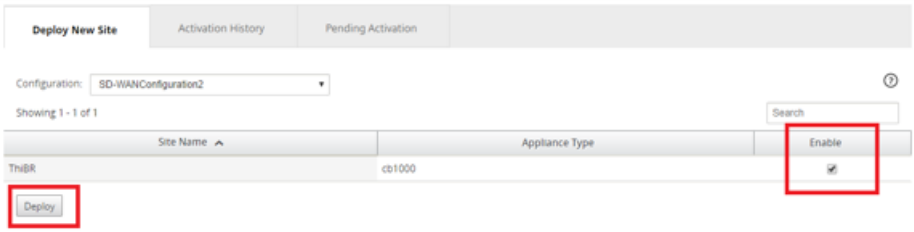

d) A Deploy New Site pop-up window will appear, where the Admin can provide the Serial Number, branch site Street Address, Installer Email address, and Additional Notes, if required.

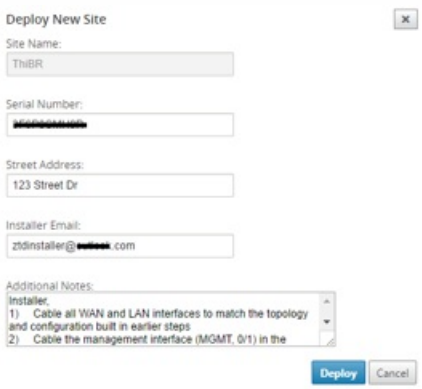

Note

The Serial Number entry field is optional and depending if it is populated or not, will result in a change in on-site activity the Installer is responsible for.

- If Serial Number field is populated – The installer in not required to enter serial number into the activation URL generated with the deploy site command

- If Serial Number field is left black – The installer will be responsible for entering in the correct serial number of the appliance into the activation URL generated with the deploy site command

e) After clicking the **Deploy** button, a message will appear indicating that “The Site configuration has been deployed”.

f) This action triggers the SD-WAN Center, which was previously registered with the ZTD Cloud Service, to share the configuration of this particular site to be temporality stored in the ZTD Cloud Service.

g) Navigate to the Pending Activation tab to confirm that the branch site information populated successfully and was put into a pending installer activity status.

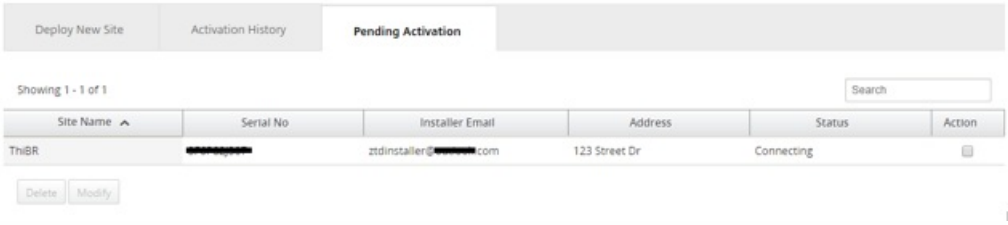

Note

A zero touch deployment in the Pending Activation state can optionally be chosen to Delete or Modify if information is seen to be incorrect. If a Site is deleted from the pending activation page, it will become available to be deployed in the Deploy New Site tab page. Once you choose to delete the branch site from Pending activation, the activation link send to the installer will become invalid.

If the Serial Number field was not populated by the SD-WAN Administrator, the Status Field will indicate "Waiting for Installer" instead of "Connecting".

4. The next series of activities will be conducted by the On-site Installer.

a) The Installer will need to check the mailbox of the email address the SD-WAN Administrator used when deploying the site.

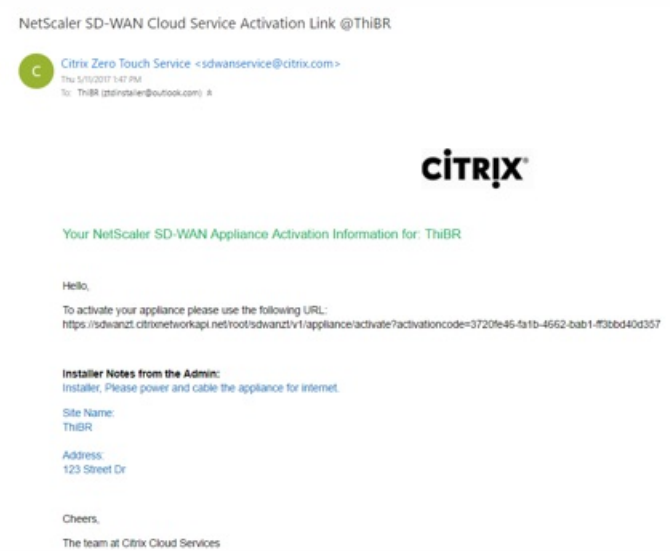

b) Open the zero touch deployment Activation URL in an internet browser window (e.g. https://sdwanzt.citrixnetworkapi.net).

c) If the SD-WAN Administrator did not pre-populate the serial number in the deploy site step, then the Installer would be responsible for locating the serial number on the physical appliance and entering the serial number manually into the activation URL, then click the **Activate** button.

d) If the Admin pre-populating the Serial Number information, the Activation URL will have already progressed to the next step.

e) The installer must physically be on-site to perform the following actions:

- Cable all WAN and LAN interfaces to match the topology and configuration built in earlier steps.
- Cable the management interface (MGMT, 0/1) in the segment of the network that will provide DHCP IP address and connectivity to the Internet with DNS and FQDN to IP address resolution.
- Power cable the SD-WAN appliance.
- Turn on the power switch of the appliance.

> ## Note
>
> Most appliances will automatically power on as soon as the power cable is attached. Some appliance may have to be powered on using the power switch on the front of the appliance, others may have the power switch on the rear of the appliance. Some power switches require holding the power button until the unit powers up.

5. The next series of steps are automated with the help of the Zero Touch Deployment service, but requires that the following pre-requisites are available.

- The branch appliance should be powered up
- DHCP must be available in the existing network to assign management and DNS IP address
- Any DHCP assigned IP address will require connectivity to the internet with ability to resolve FQDNs
- IP assignment can be configured manually, as long as the other pre-requisites are meet

a) The appliance obtains an IP address from the networks DHCP Server, in this example topology this is achieved through the bypassed data interfaces of a factory default state appliance.

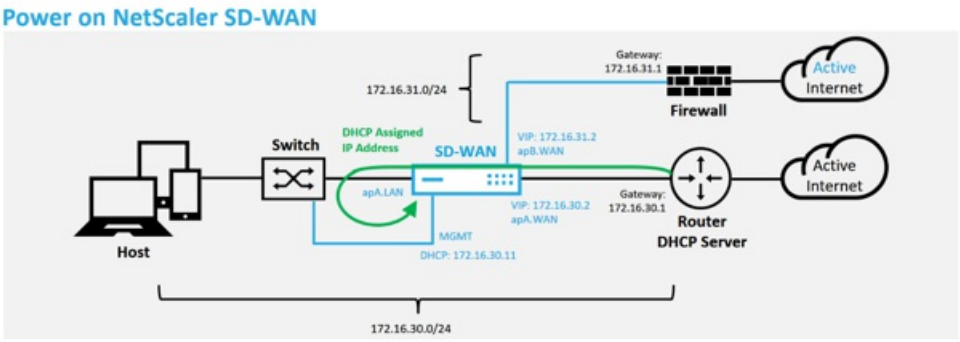

b) As the appliance obtains the web management and DNS IP addresses from the underlay network DHCP Server, the appliance will call home to the Zero Touch Deployment Service and download any ZTD related software updates.

c) With successful connectivity to the ZTD Cloud Service, the deployment process will automatically perform the following:

- Download the Configuration File that was stored earlier by the SD-WAN Center
- Applying the Configuration to the local appliance
- Download and Install a temporary 10 MB license file
- Download and Install any software updates if needed
- Activate the SD-WAN Service

d) Further confirmation can be done in the SD-WAN Center web management interface, the Zero Touch Deployment menu will display successfully activated appliances in the **Activation History** tab.

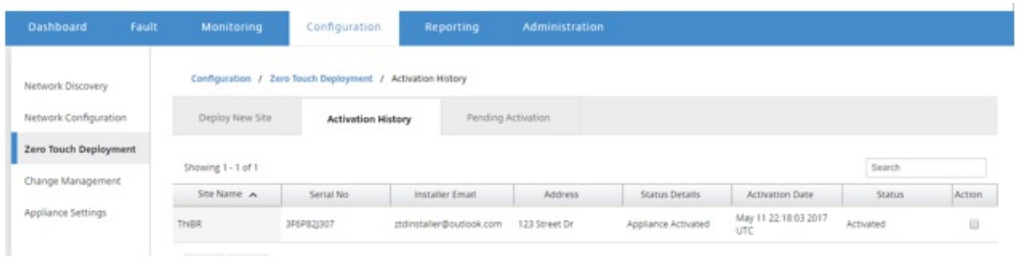

e) The Virtual Paths may not immediately show in a connected state, this is because the MCN may not trust the configuration handed down from the ZTD Cloud Service, and will report “Configuration version mismatch” in the MCN Dashboard.

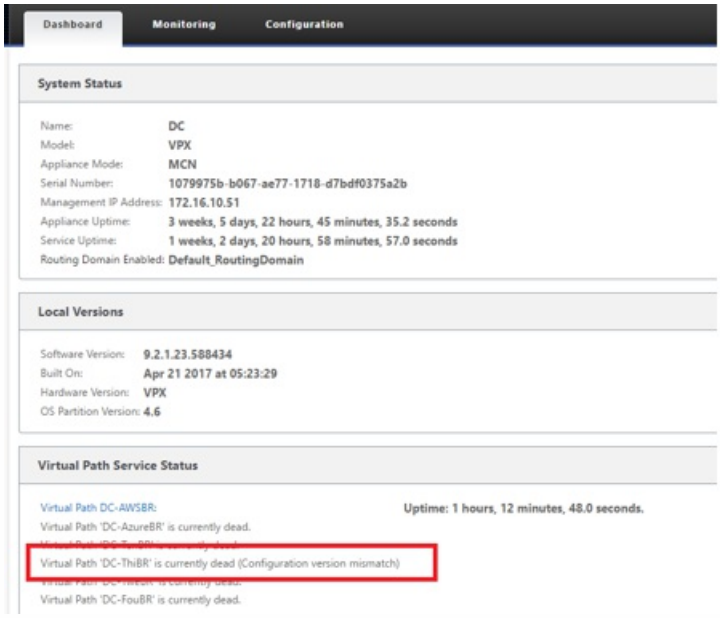

f) The configuration will automatically be redelivered to the newly installed branch office appliance, the status of this can be monitoring on the **MCN** > **Configuration** > **Virtual WAN** > **Change Management** page (this process can take several minutes to complete).

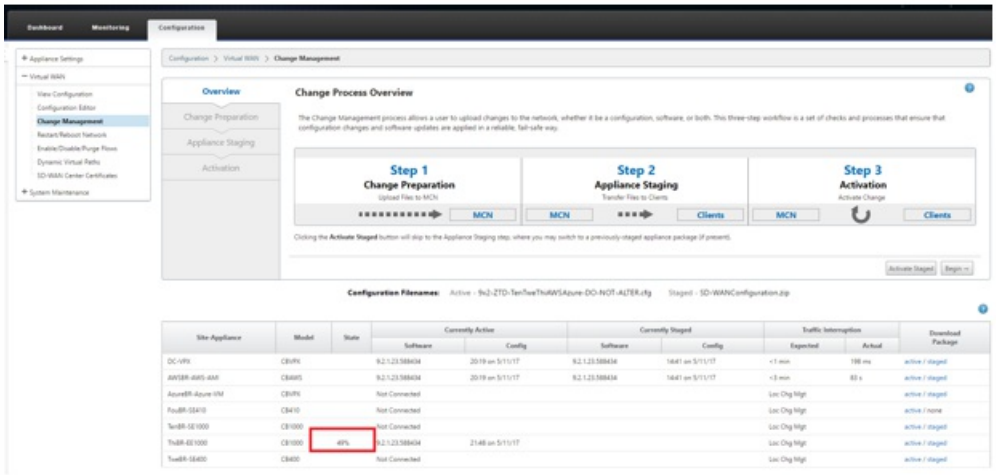

g) The SD-WAN Administrator can monitor the head-end MCN web management page for the established Virtual Paths of the remote site.

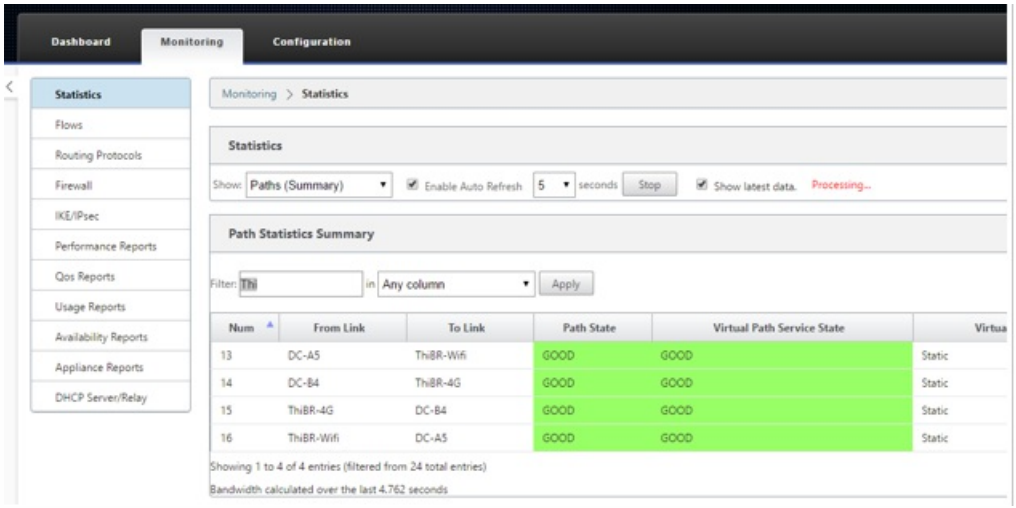

h) SD-WAN Center can also be utilized to identify the DHCP assigned IP address of the on-site appliance from the **Configuration** > **Network Discovery** > **Inventory and Status** page.

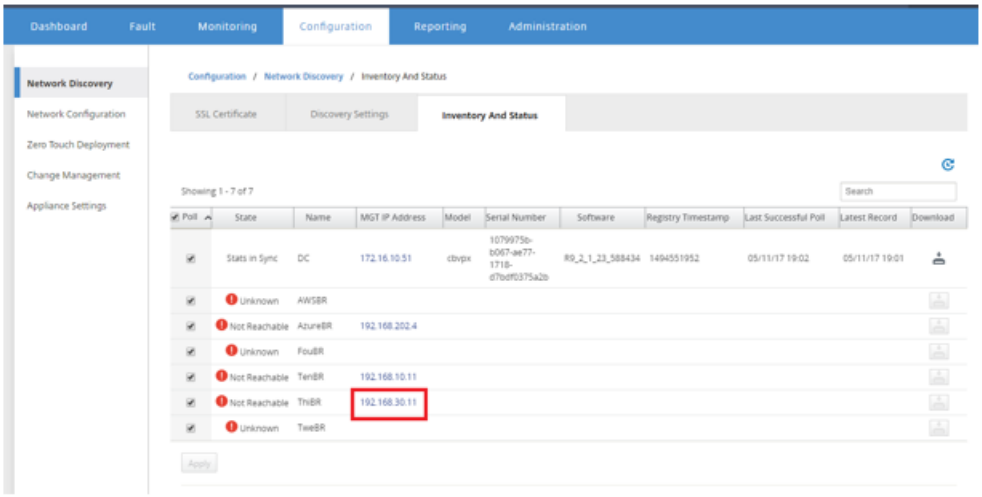

i) At this point the SD-WAN Network Administrator can gain web management access to on-site appliance utilizing the SD-WAN overlay network.

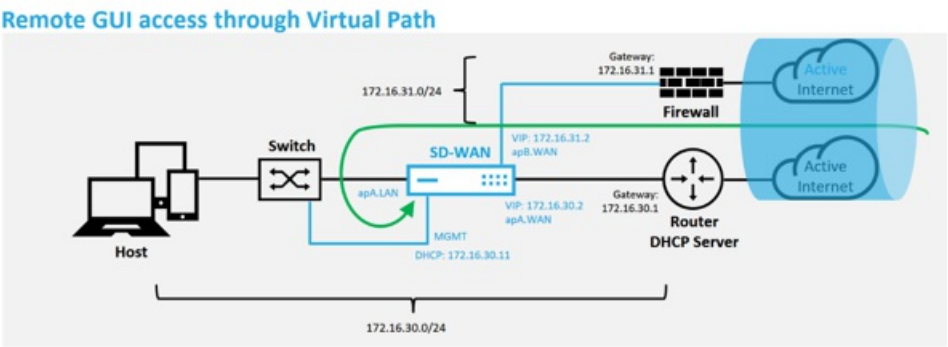

j) Web management access to the remote site appliance will indicate that the appliance has been installed with a temporary Grace License at 10Mbps, which enables the ability for the Virtual Path Service Status to report as active.

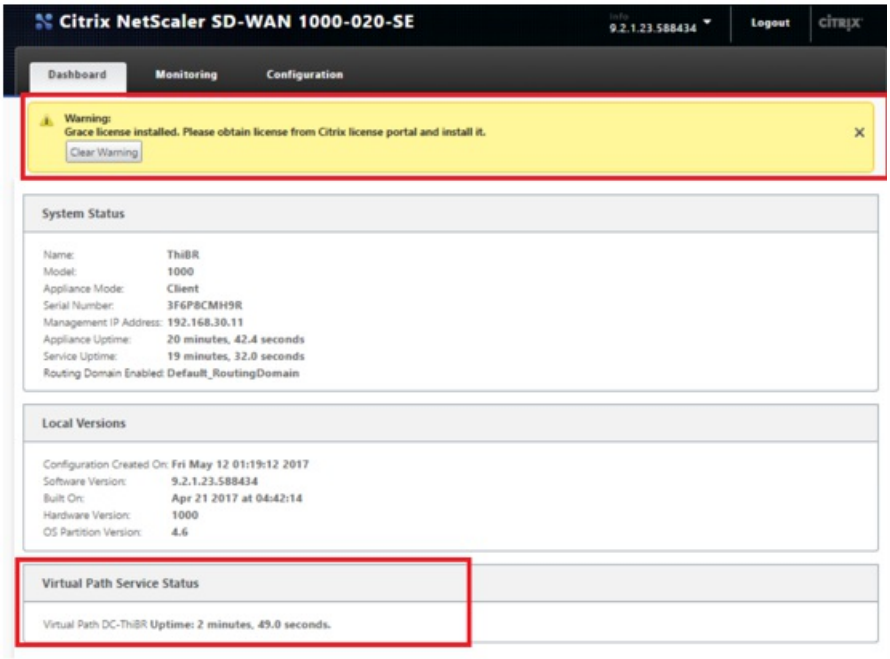

k) The appliance configuration can be validated using the **Configuration** > **Virtual WAN** > **View Configuration** page.

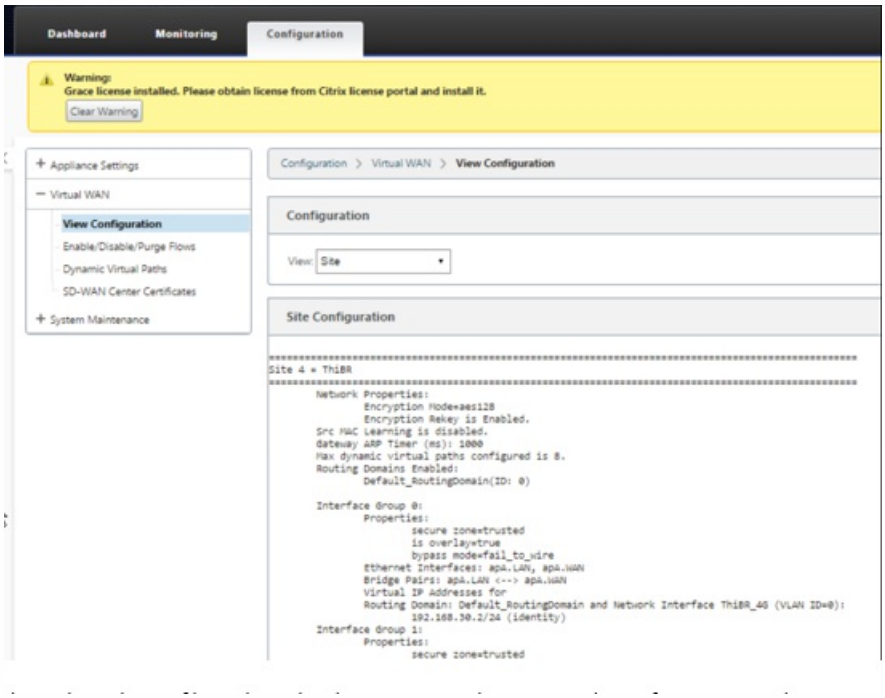

l) The appliance license file can be updated to a permanent license using the **Configuration** > **Appliance Settings** > **Licensing** page.

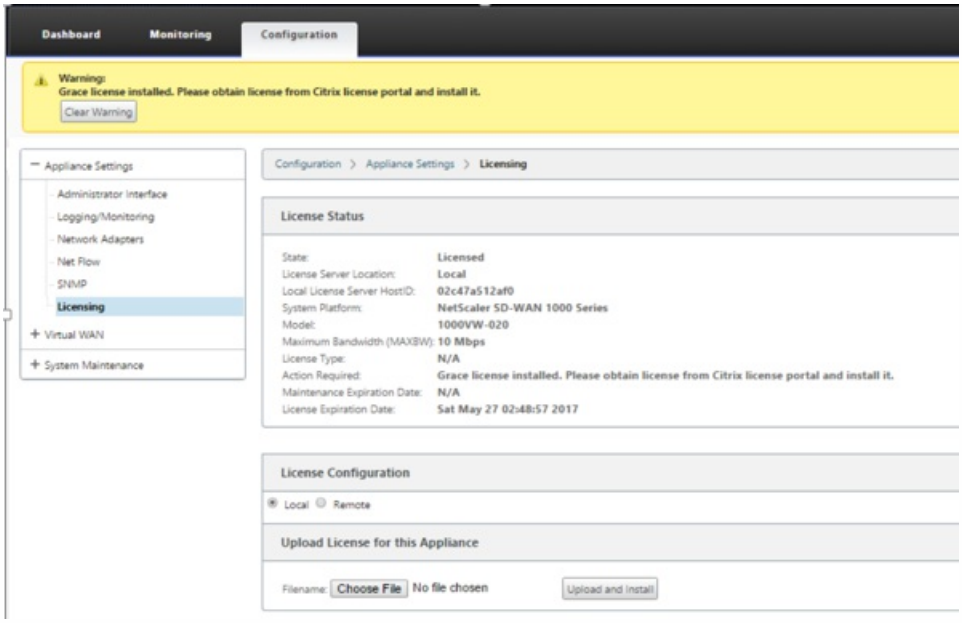

m) After uploading and installing the permanent license file, the Grace License warning banner is will disappear, and during the license install process no loss in connectivity to the remote site will occur (zero pings are dropped).

Appliance Deployment with Zero Touch

Sep 20, 2017

For instructions about how to deploy an SD-WAN appliance with Zero Touch Service, see the topic; How to Configure Zero Touch Deployment Service.

AWS Deployment with Zero Touch

Oct 25, 2017

With NetScaler SD-WAN release 9.3, zero touch deployment capabilities have extended to Cloud instances. The procedure to deploy zero touch deployment process four cloud instances is slightly different from appliance deployment for zero touch service.

1. Update the configuration to add a new remote site with a ZTD capable SD-WAN cloud device using SD-WAN Center Network Configuration.

If the SD-WAN configuration was not built using the SD-WAN Center Network Configuration, import the active configuration from the MCN and begin modifying the configuration using SD-WAN Center. For Zero Touch Deployment capability, the SD-WAN Administrator must build the configuration using SD-WAN Center. The following procedure should be used to add a new cloud node targeted for zero touch deployment.

a) Design the new site for SD-WAN cloud deployment by first outlining the details of the new site (i.e. VPX size, Interface Groups usage, Virtual IP Addresses, WAN Link(s) with bandwidth and their respective Gateways).

> ## Note
>
> - Cloud deployed SD-WAN instances must be deployed in Edge/Gateway mode.
> - The template for the cloud instance is limited to three interfaces; Management, LAN, and WAN (in that order).
> - The available cloud templates for SD-WAN VPX are currently hard-set to obtain the #.#.#.#.11 IP address of the available subnets in the VPC .

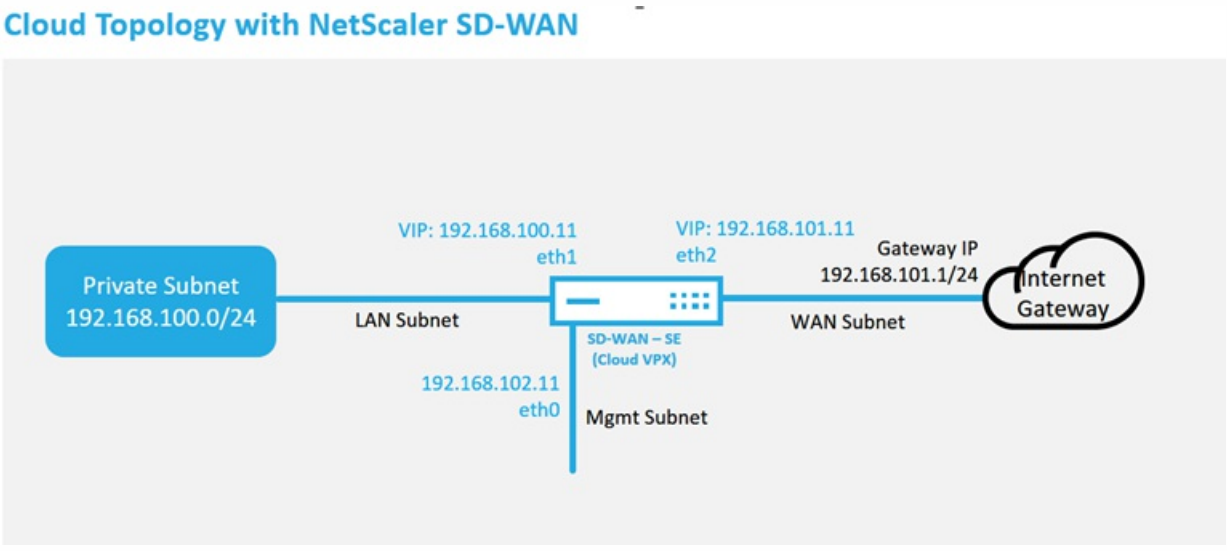

This is an example deployment of a SD-WAN cloud deployed site, the NetScaler SD-WAN device is deployed as the edge device servicing a single Internet WAN link in this cloud network. Remote sites will be able to leverage multiple distinct Internet WAN links connecting into this same Internet Gateway for the cloud, providing resiliency and aggregated bandwidth connectivity from any SD-WAN deploy site to the cloud infrastructure. This provides cost effective and highly reliable connectivity to the cloud.

b) Open the SD-WAN Center web management interface and navigate to the **Configuration** > **Network Configuration** page.

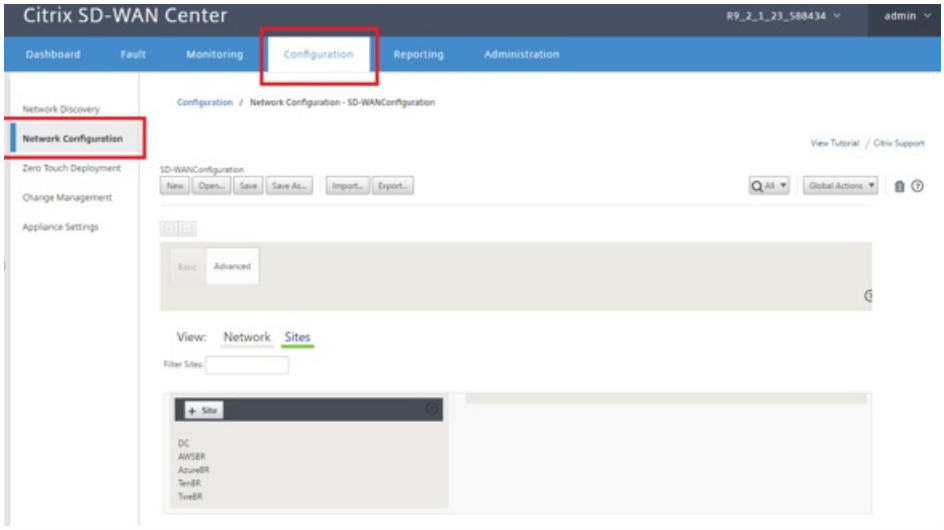

c) Make sure a working configuration is already in place, or import the configuration from the MCN.

d) Navigate to the Basic tab to create a new site.

e) Open the Sites tile to display the currently configured sites.

f) Quickly built the configuration for the new cloud site by utilizing the clone feature of any existing site, or manually build a new site.

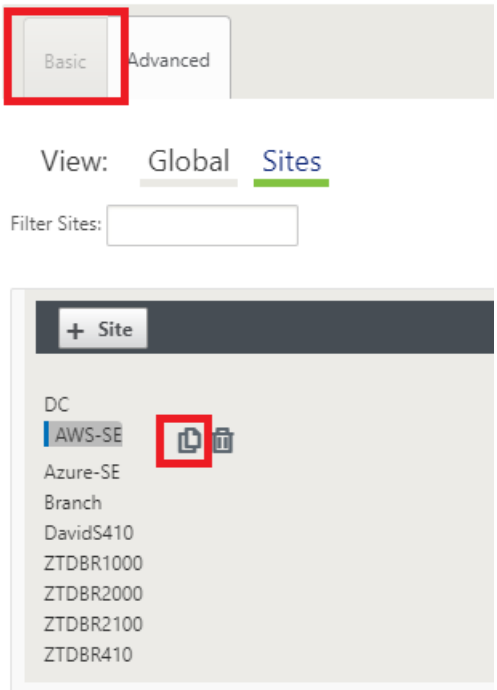

g) Populate all the required fields from the topology designed earlier for this new cloud site

Keep in mind that the template available for cloud ZTD deployments are hard-set to utilize the #.#.#.11 IP address for the Mgmt, LAN, and WAN subnets. If the configuration is not set to match the expected .11 IP host address for each interface, then the device will not be able to properly establish ARP to the cloud environment gateways and IP connectivity to the Virtual Path of the MCN.

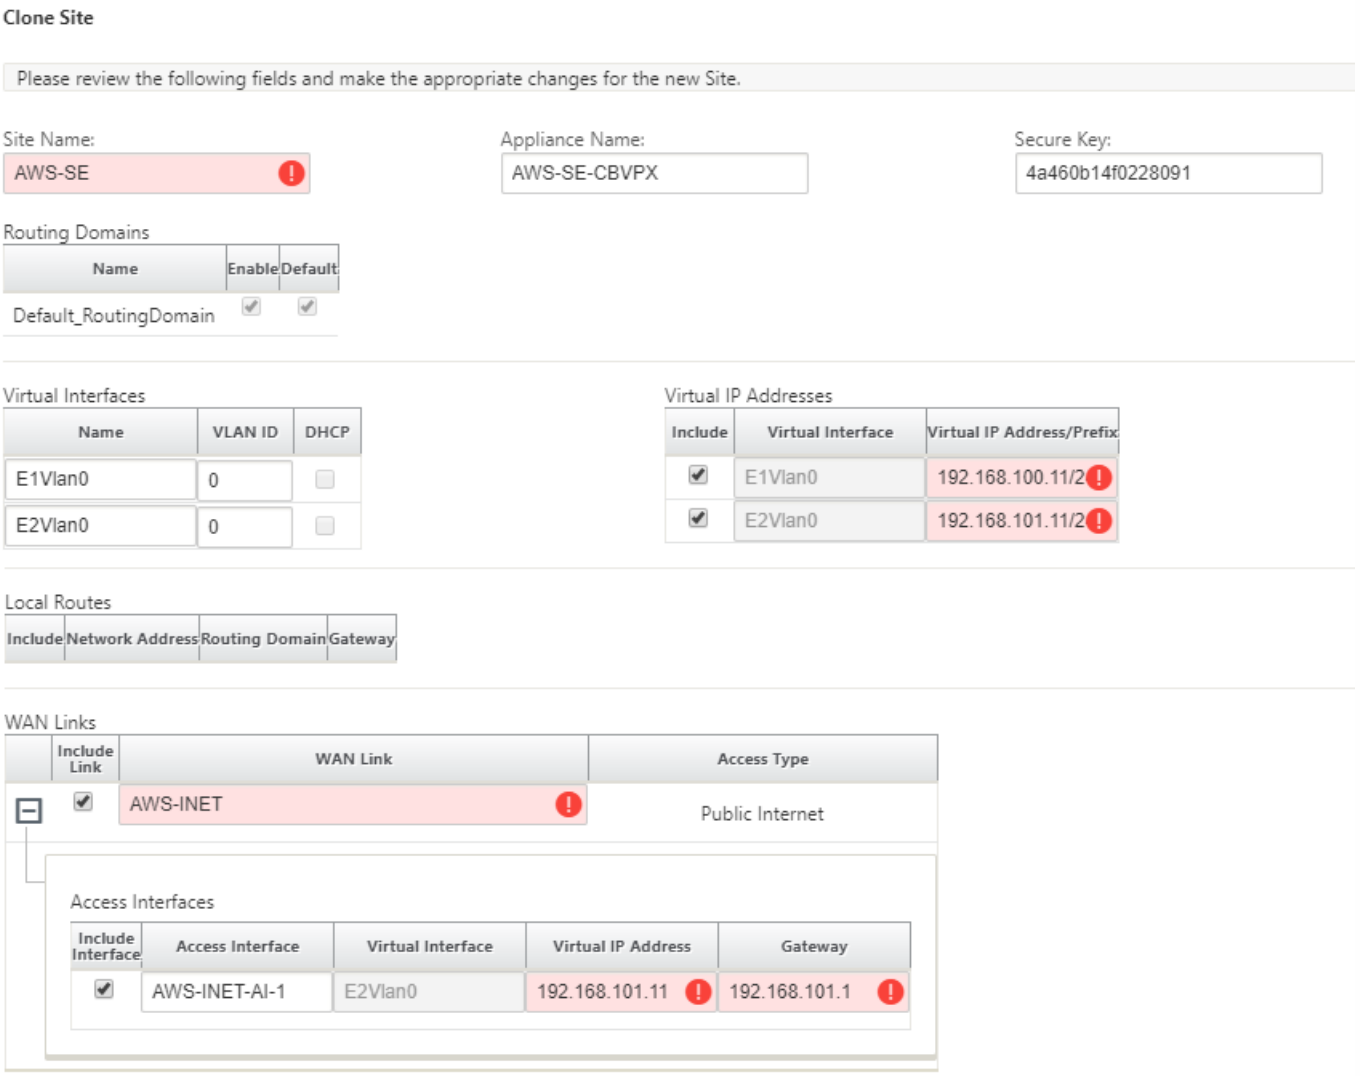

h) After cloning a new site, navigate to the site's **Basic Settings**, and verify that the Model of SD-WAN is correctly selected which would support the zero touch service.

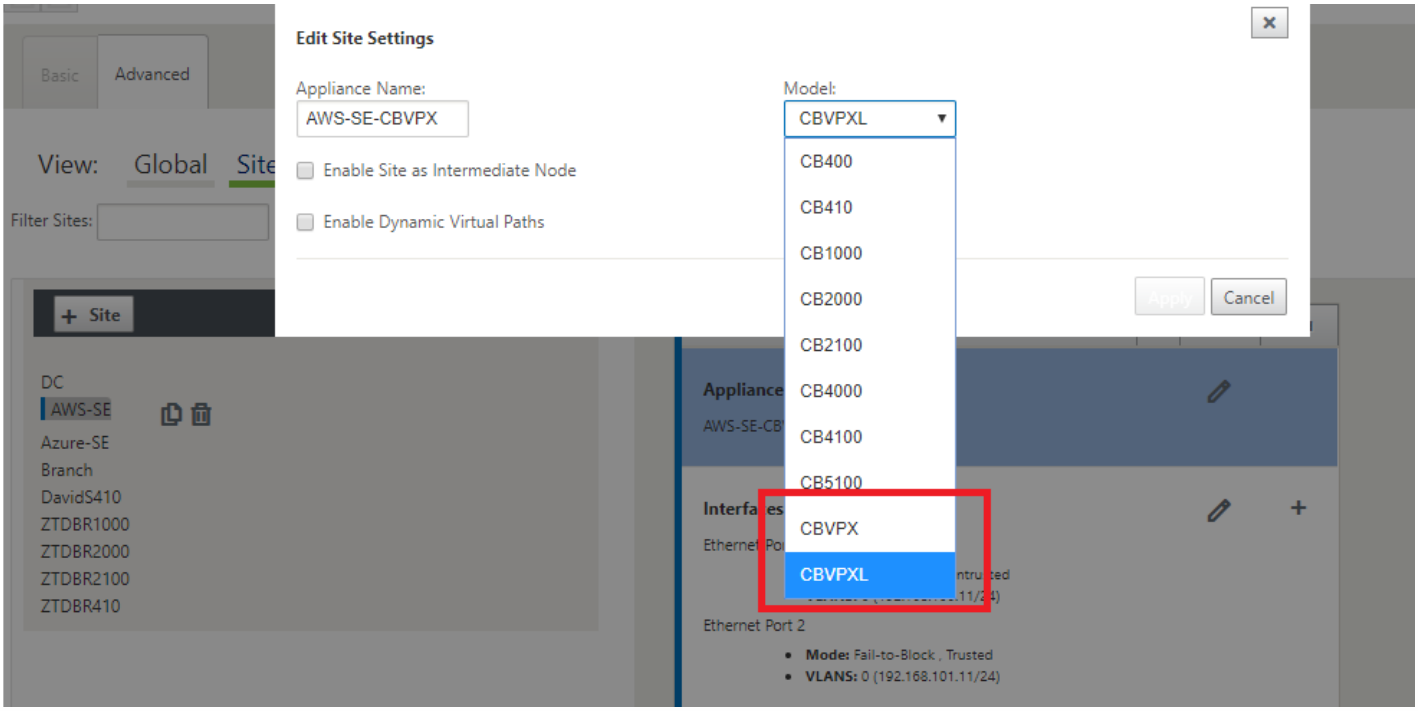

i) Save the new configuration on SD-WAN Center, and use the export to the "**Change Management inbox**" option to push the configuration using Change Management.

j) Follow the Change Management procedure to properly stage the new configuration, which makes the existing SD-WAN devices aware of the new site to be deployed via zero touch, you will need to utilize the "*Ignore Incomplete*" option to skip attempting to push the configuration to the new site that still needs to go through the ZTD workflow.

3. Navigate back to the SD-WAN Center Zero Touch Deployment page, and with the new active configuration running, the new site will be available for deployment.

a) In the Zero Touch Deployment page, under the **Deploy New Site** tab, select the running network configuration file.

b) After the running configuration file is selected, the list of all the branch sites with undeployed NetScaler SD-WAN devices that are supported for zero touch will be displayed.

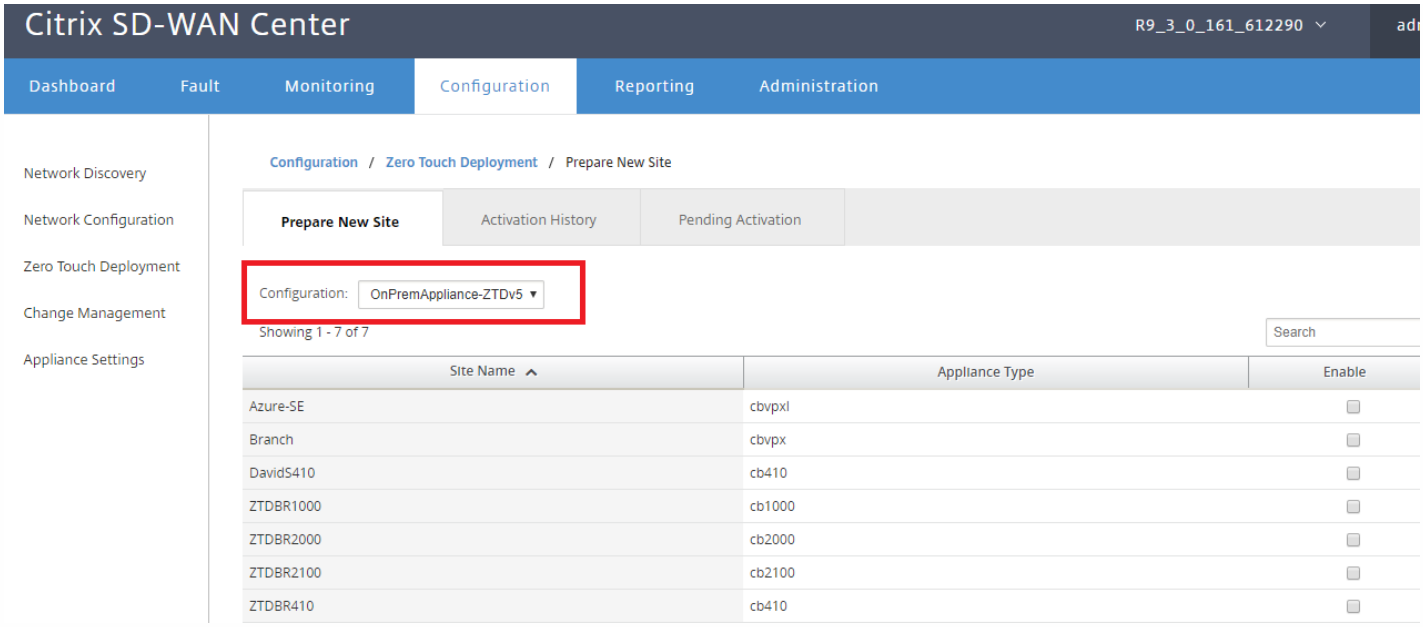

c). Select the target cloud site you want to deploy using the Zero Touch service, click **Enable**, and then **Provision and Deploy**.

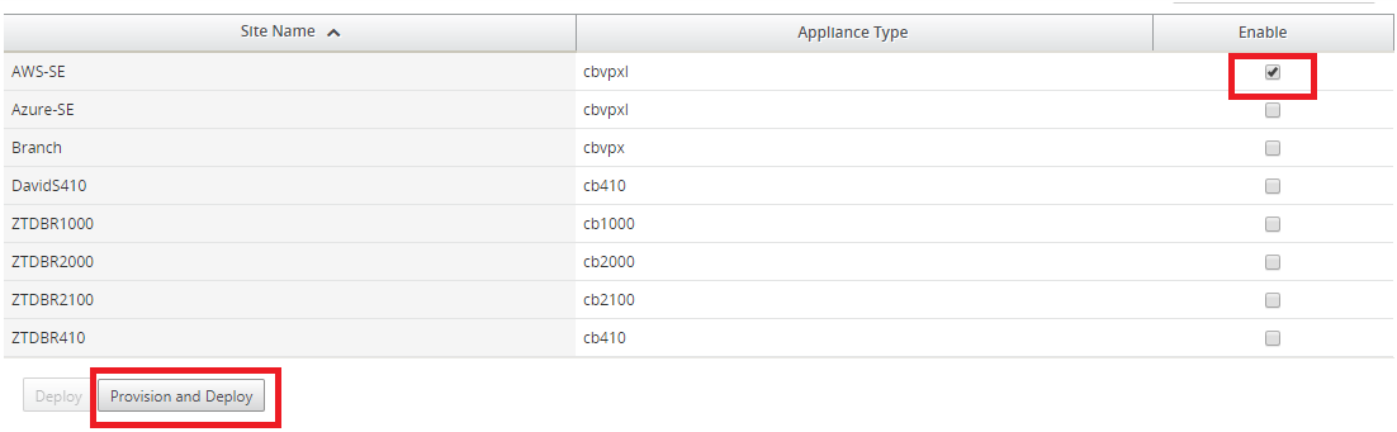

d) A pop-up window will appear, where the NetScaler SD-WAN Admin can initiate the deployment for Zero Touch.

Populate an email address where the activation URL can be delivered, and select the **Provision Type** for the desired Cloud.

e) After clicking **Next**, Select the appropriate Region, Instance size, populate the SSH Key name and Role ARN fields appropriately.

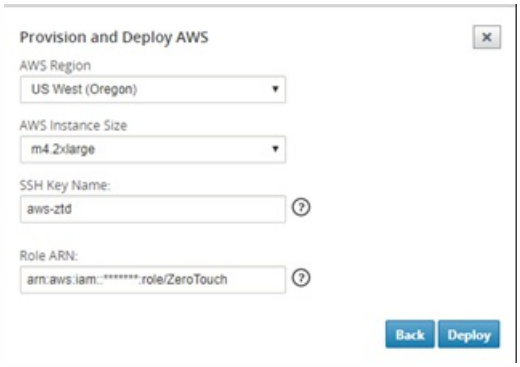

Note

Make use of the help links for guidance on how to setup the SSH Key and Role ARN on the Cloud account. Also make sure the select region matches what is available on the account and that the selected Instance Size matches VPX or VPXL as the selected model in the SD-WAN configuration.

f) Click **Deploy**, triggering the SD-WAN Center, which was previously registered with the ZTD Cloud Service, to share the configuration of this site to be temporality stored in the ZTD Cloud Service.

g) Navigate to the **Pending Activation** tab to confirm that the site information populated successfully and was put into a provisioning status.

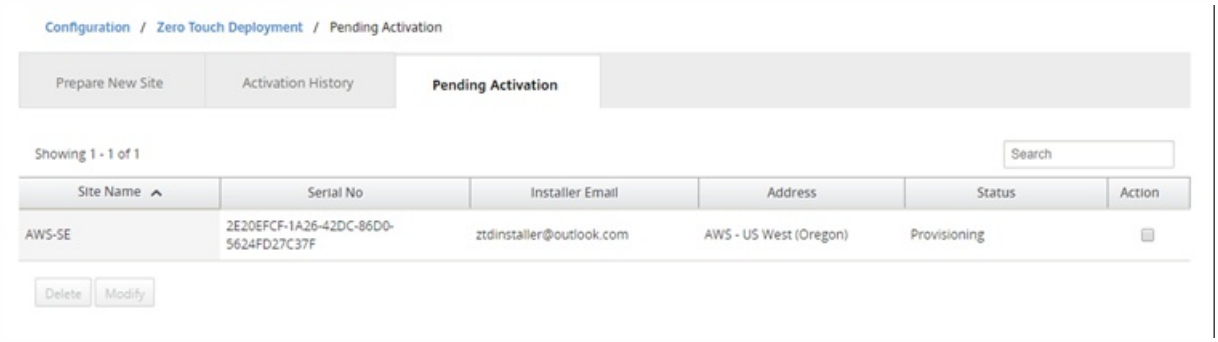

4. Initiate the Zero Touch Deployment process as the Cloud Admin.

a) The Installer will need to check the mailbox of the email address the SD-WAN Administrator used when deploying the site.

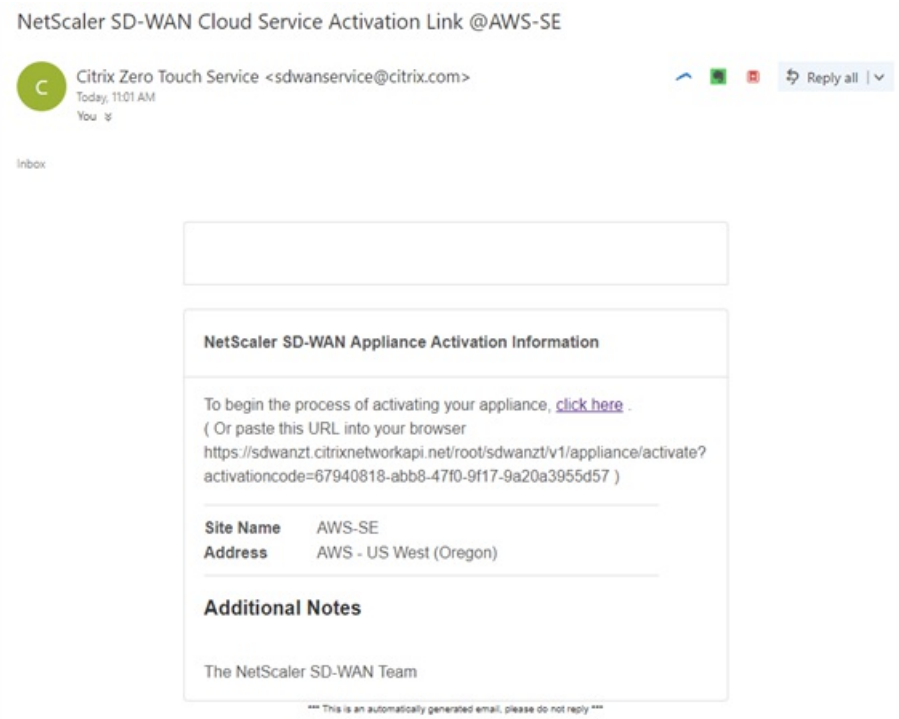

b) Open the activation URL found in the email in an internet browser window (example; *https://sdwanzt.citrixnetworkapi.net*).

c) If the SSH Key and Role ARN are properly inputted, the Zero Touch Deployment Service will immediately start provisioning the SD-WAN instance, otherwise connections errors will immediately be displayed.

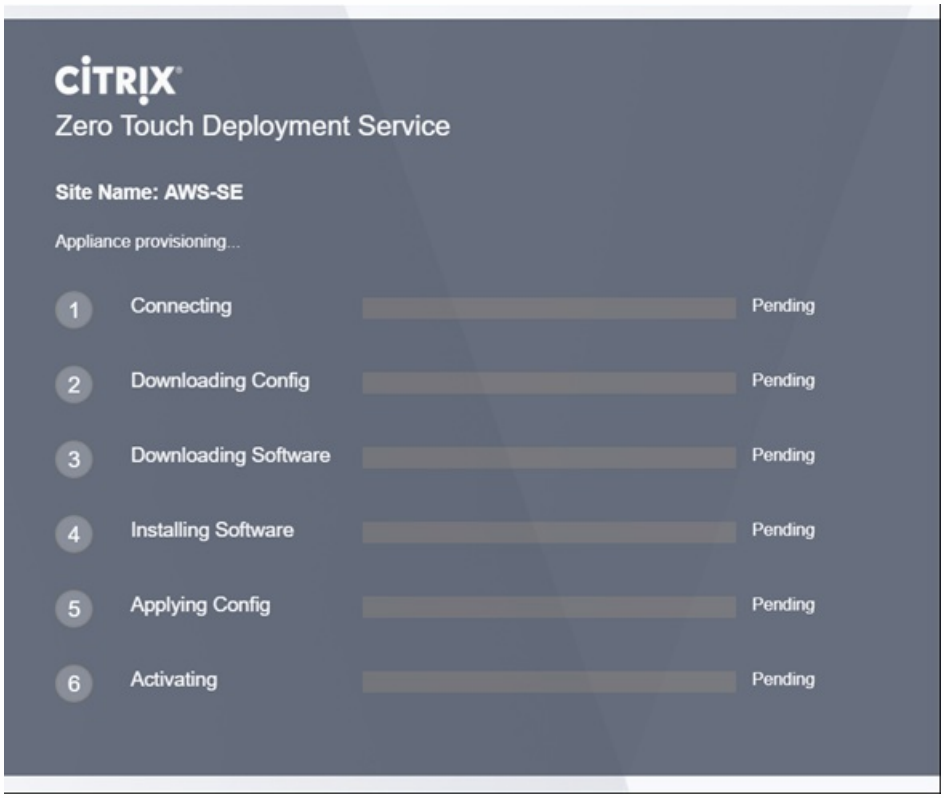

d) For additional troubleshooting on the AWS console, the Cloud Formation service can be utilized to catch any events that occur during the provisioning process.

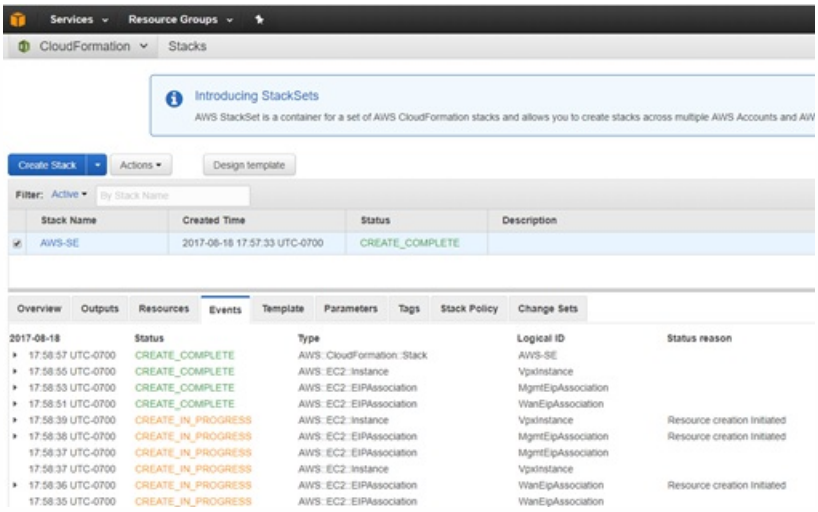

e) Allow the provisioning process ~8-10 minutes and activation another ~3-5 minutes to fully complete.

f) With successful connectivity of the SD-WAN cloud instance to the ZTD Cloud Service, the service will automatically perform the following:

- Download the site-specific Configuration File that was stored earlier by the SD-WAN Center
- Applying the Configuration to the local instance
- Download and Install a temporary 10 MB license file
- Download and Install any software updates if needed
- Activate the SD-WAN Service

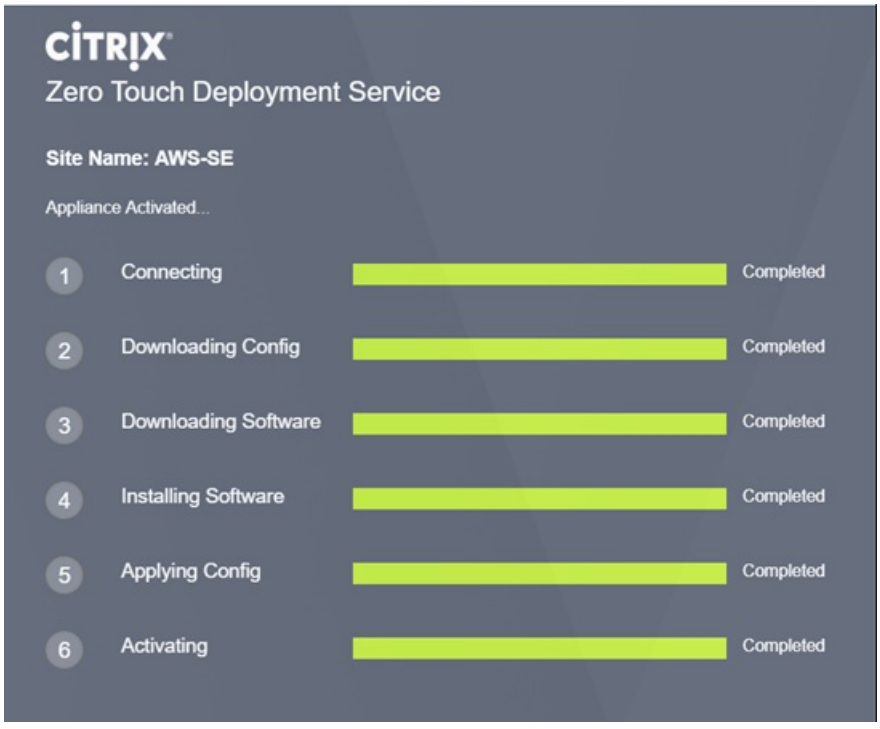

g) Further confirmation can be done in the SD-WAN Center web management interface; the Zero Touch Deployment menu will display successfully activated appliances in the **Activation History** tab.

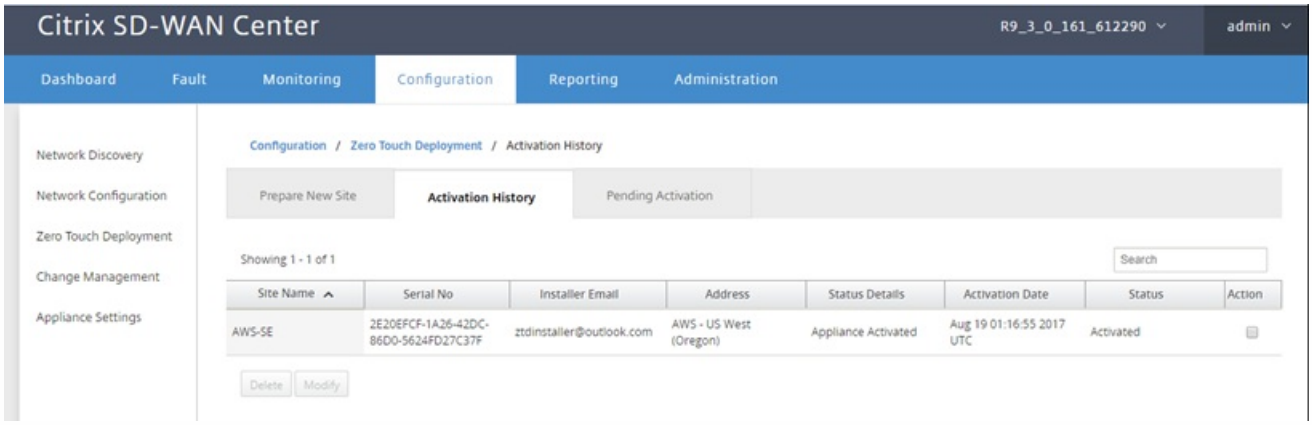

h) The Virtual Paths may not immediately show in a connected state, this is because the MCN may not trust the configuration handed down from the ZTD Cloud Service, and will report "*Configuration version mismatch*" in the MCN Dashboard.

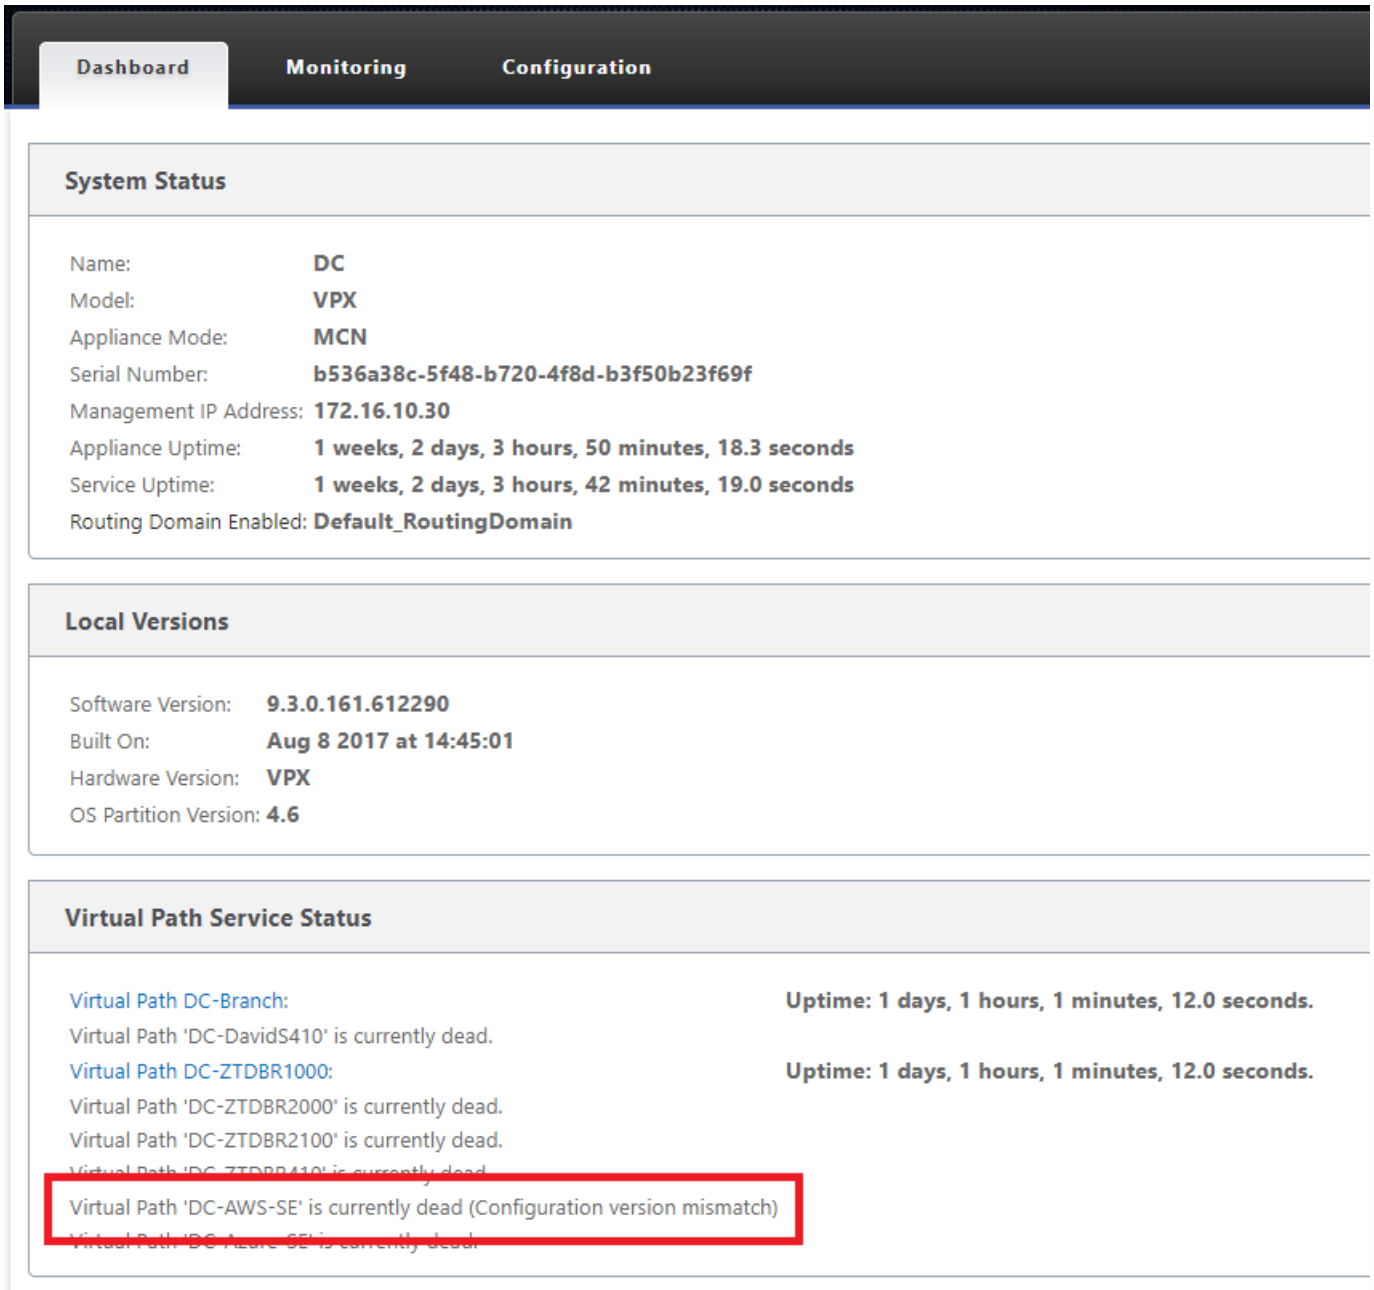

i) The configuration will automatically be redelivered to the newly installed branch office appliance, the status of this can be monitoring on the **MCN** > **Configuration** > **Virtual WAN**> **Change Management** page (depending on the connectivity, this process can take several minutes to complete).

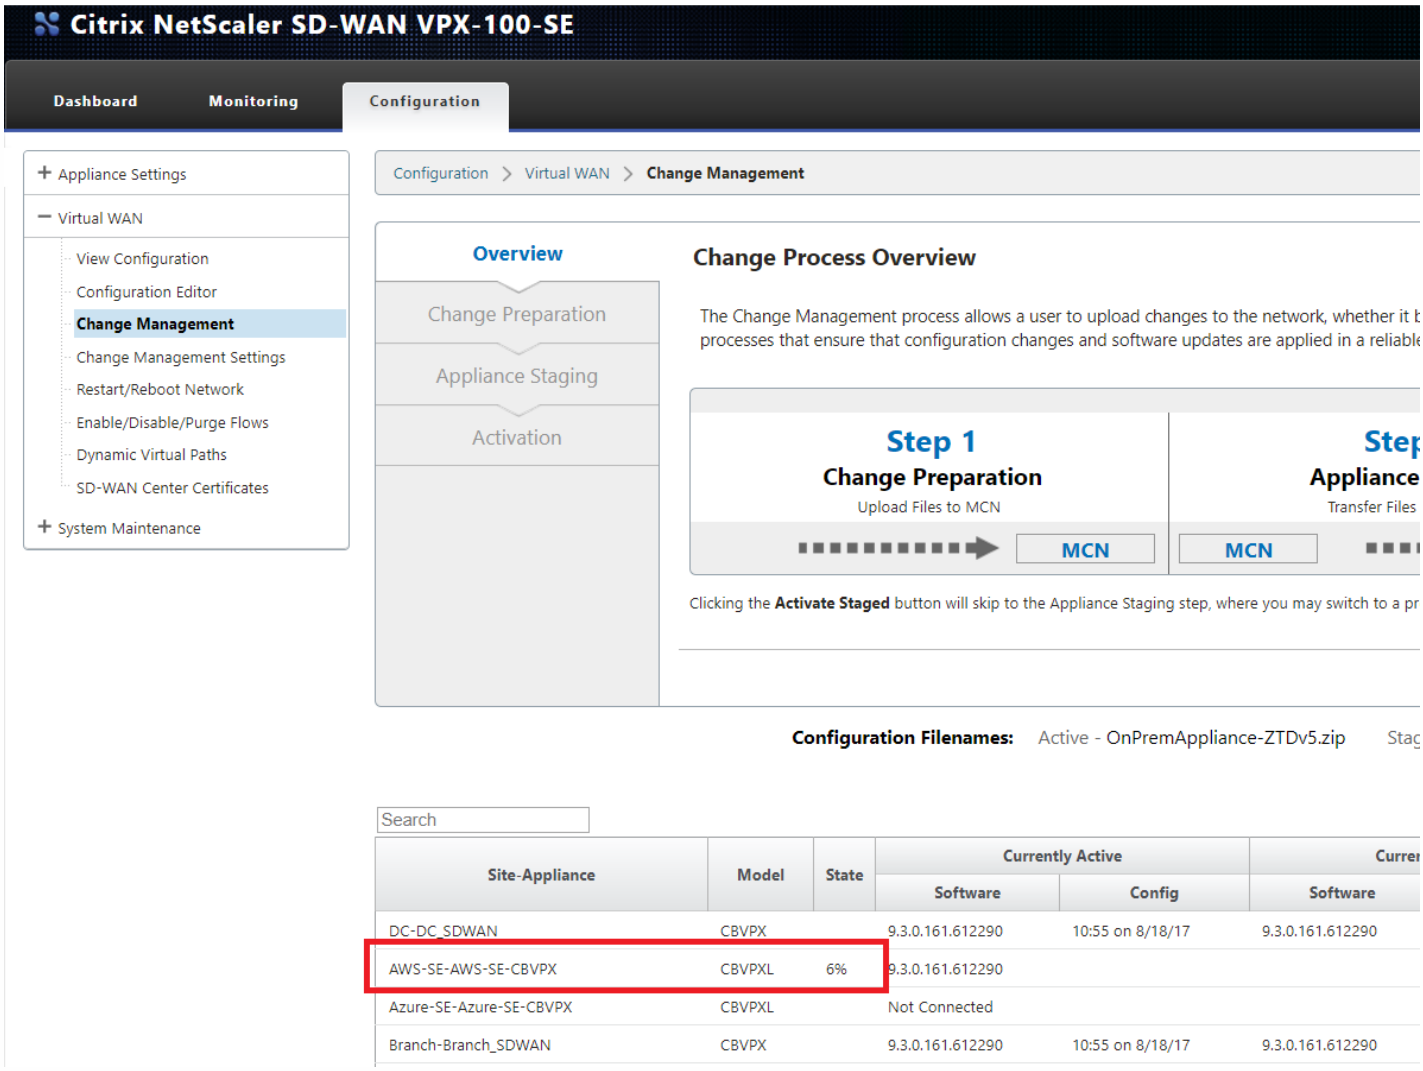

j) The SD-WAN Administrator can monitor the head-end MCN web management page for the established Virtual Paths of the newly added cloud site.

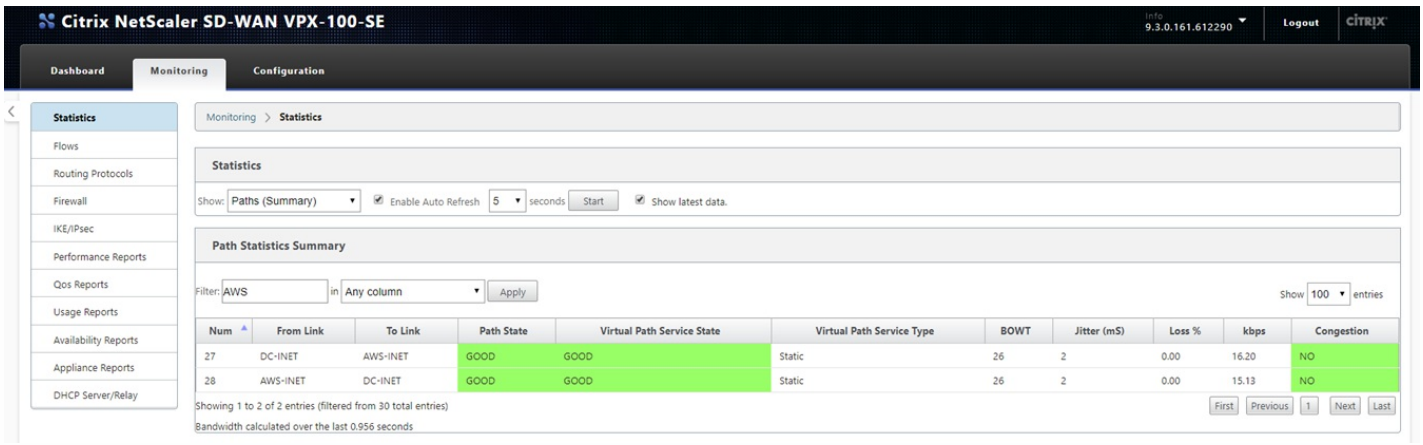

k) If troubleshooting is required, open the SD-WAN instances user interface using the public IP assigned by the cloud environment during provisioning, and utilize the ARP table in the **Monitoring** > **Statistics** page to identify any issues

connecting to the expected gateways, or utilize the trace route and packet capture options in diagnostics.

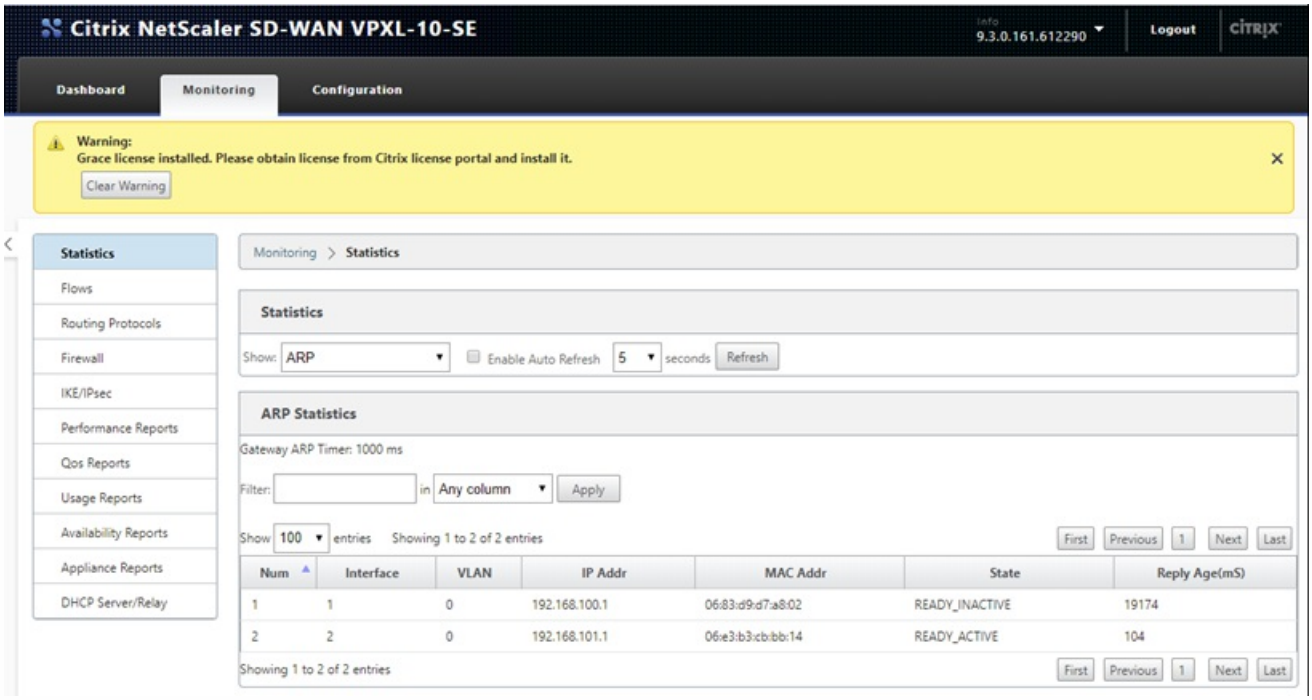

Azure Deployment with Zero Touch

Oct 23, 2017

With NetScaler SD-WAN release 9.3, zero touch deployment capabilities have extended to Cloud instances. The procedure to deploy zero touch deployment process for cloud instances is slightly different from appliance deployment for zero touch service.

Updating the configuration to add a new remote site with a ZTD capable SD-WAN cloud device using SD-WAN Center Network Configuration

If the SD-WAN configuration was not built using the SD-WAN Center Network Configuration, import the active configuration from the MCN and begin modifying the configuration using SD- WAN Center. For Zero Touch Deployment capability, the SD-WAN Administrator must build the configuration using SD-WAN Center. The following procedure should be used to add a new cloud node targeted for zero touch deployment.

a) Design the new site for SD-WAN cloud deployment by first outlining the details of the new site (i.e. VPX size, Interface Groups usage, Virtual IP Addresses, WAN Link(s) with bandwidth and their respective Gateways).

Note

- Cloud deployed SD-WAN instances must be deployed in Edge/Gateway mode.
- The template for the cloud instance is limited to three interfaces; Management, LAN, and WAN (in that order).
- The available Azure cloud templates for SD-WAN VPX are currently hard-set to obtain the 10.9.4.106 IP for the WAN, 10.9.3.106 IP for the LAN, and 10.9.0.16 IP for the Management address. The SD-WAN configuration for the Azure node targeted for Zero Touch must match this layout.
- The Azure site name in the configuration must be all lowercase with no special characters (e.g. ztdazure).

Azure Cloud Topology with NetScaler SD-WAN

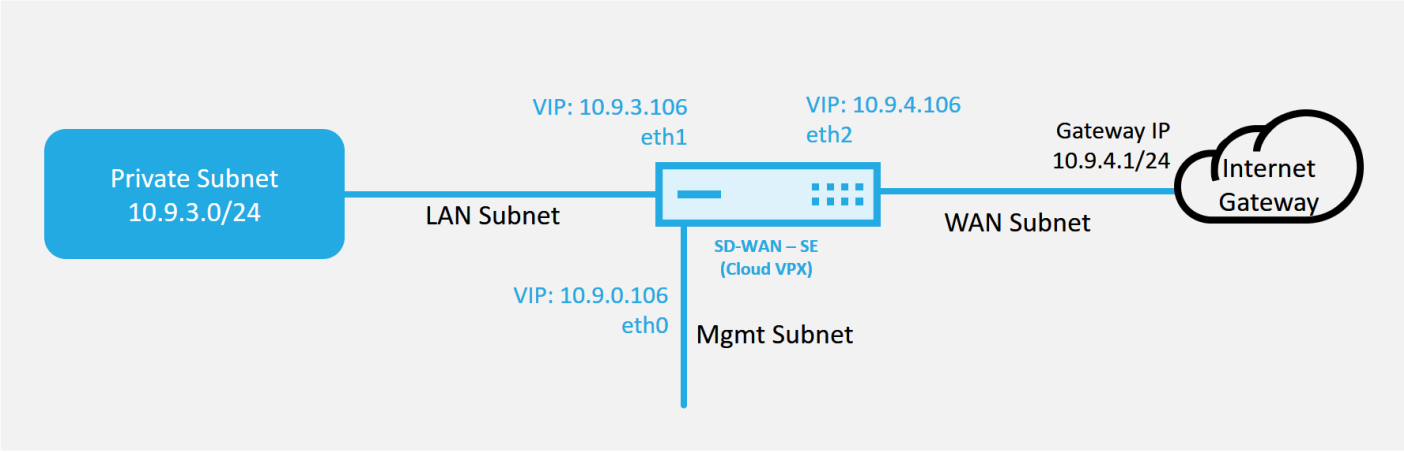

This is an example deployment of a SD-WAN cloud deployed site, the NetScaler SD-WAN device is deployed as the edge device servicing a single Internet WAN link in this cloud network. Remote sites will be able to leverage multiple distinct Internet WAN links connecting into this same Internet Gateway for the cloud, providing resiliency and aggregated bandwidth connectivity from any SD-WAN deploy site to the cloud infrastructure. This provides cost effective and highly reliable connectivity to the cloud.

b) Open the SD-WAN Center web management interface and navigate to the **Configuration** > **Network Configuration** page.

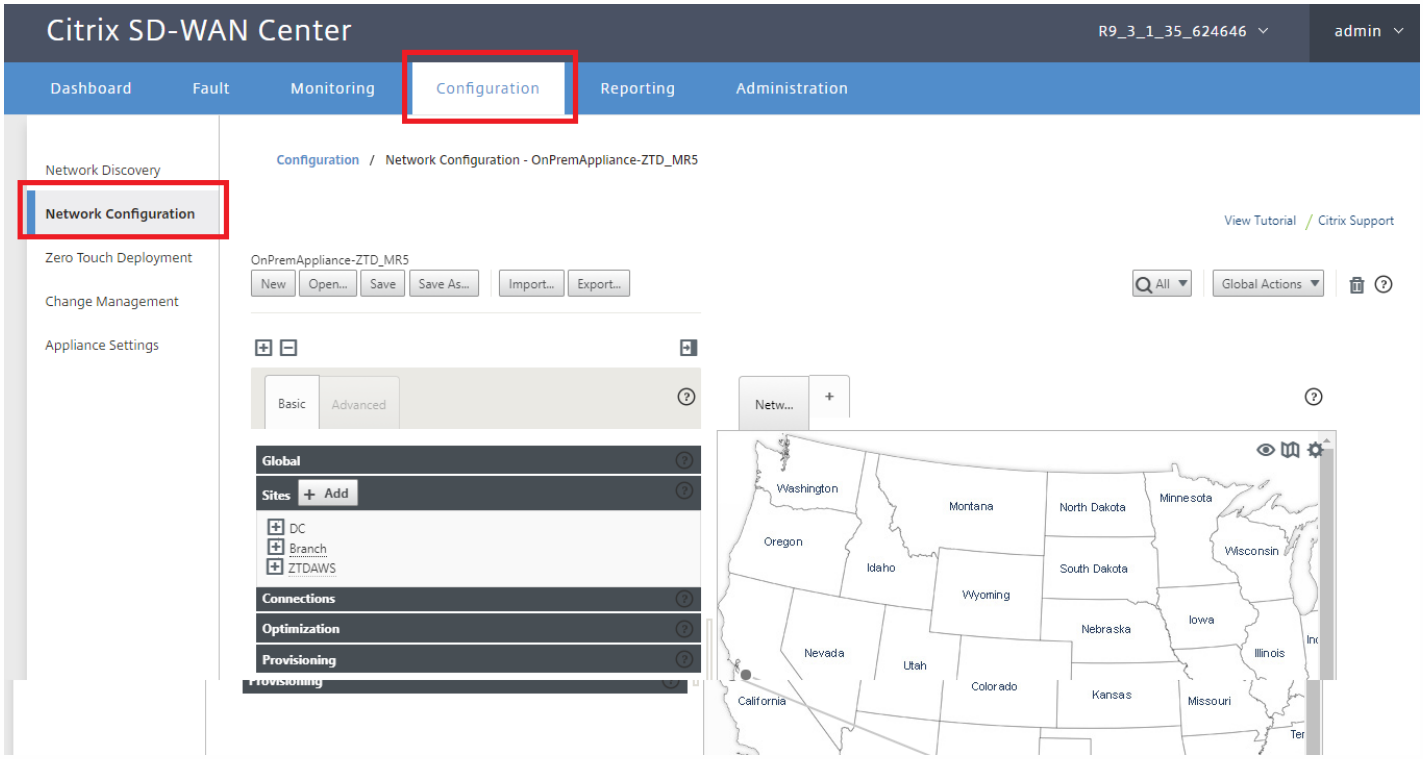

c) Make sure a working configuration is already in place, or import the configuration from the MCN.

d) Navigate to the Basic tab to create a new site.

e) Open the Sites tile to display the currently configured sites.

f) Quickly built the configuration for the new cloud site by utilizing the clone feature of any existing site, or manually build a new site.

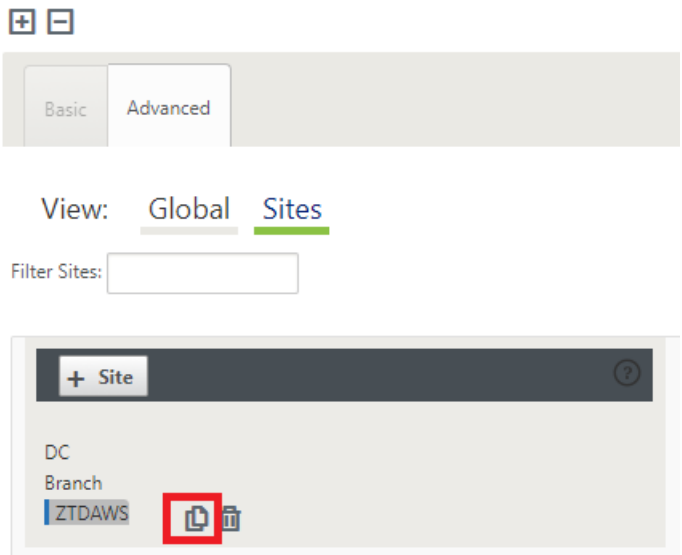

g) Populate all the required fields from the topology designed earlier for this new cloud site.

Keep in mind that the template available for Azure cloud ZTD deployments is currently hard-set to obtain the 10.9.4.106 IP

for the WAN, 10.9.3.106 IP for the LAN, and 10.9.0.16 IP for the Management address. If the configuration is not set to match the expected VIP address for each interface, then the device will not be able to properly establish ARP to the cloud environment gateways and IP connectivity to the Virtual Path of the MCN.

It is import that the site name be compliant with what Azure expects. The site name must be in all lower case, at least 6 characters, with no special characters, it must confirm to the following regular expression **^[a-z][a-z0-9-]{1,61}[a-z0-9]$.**

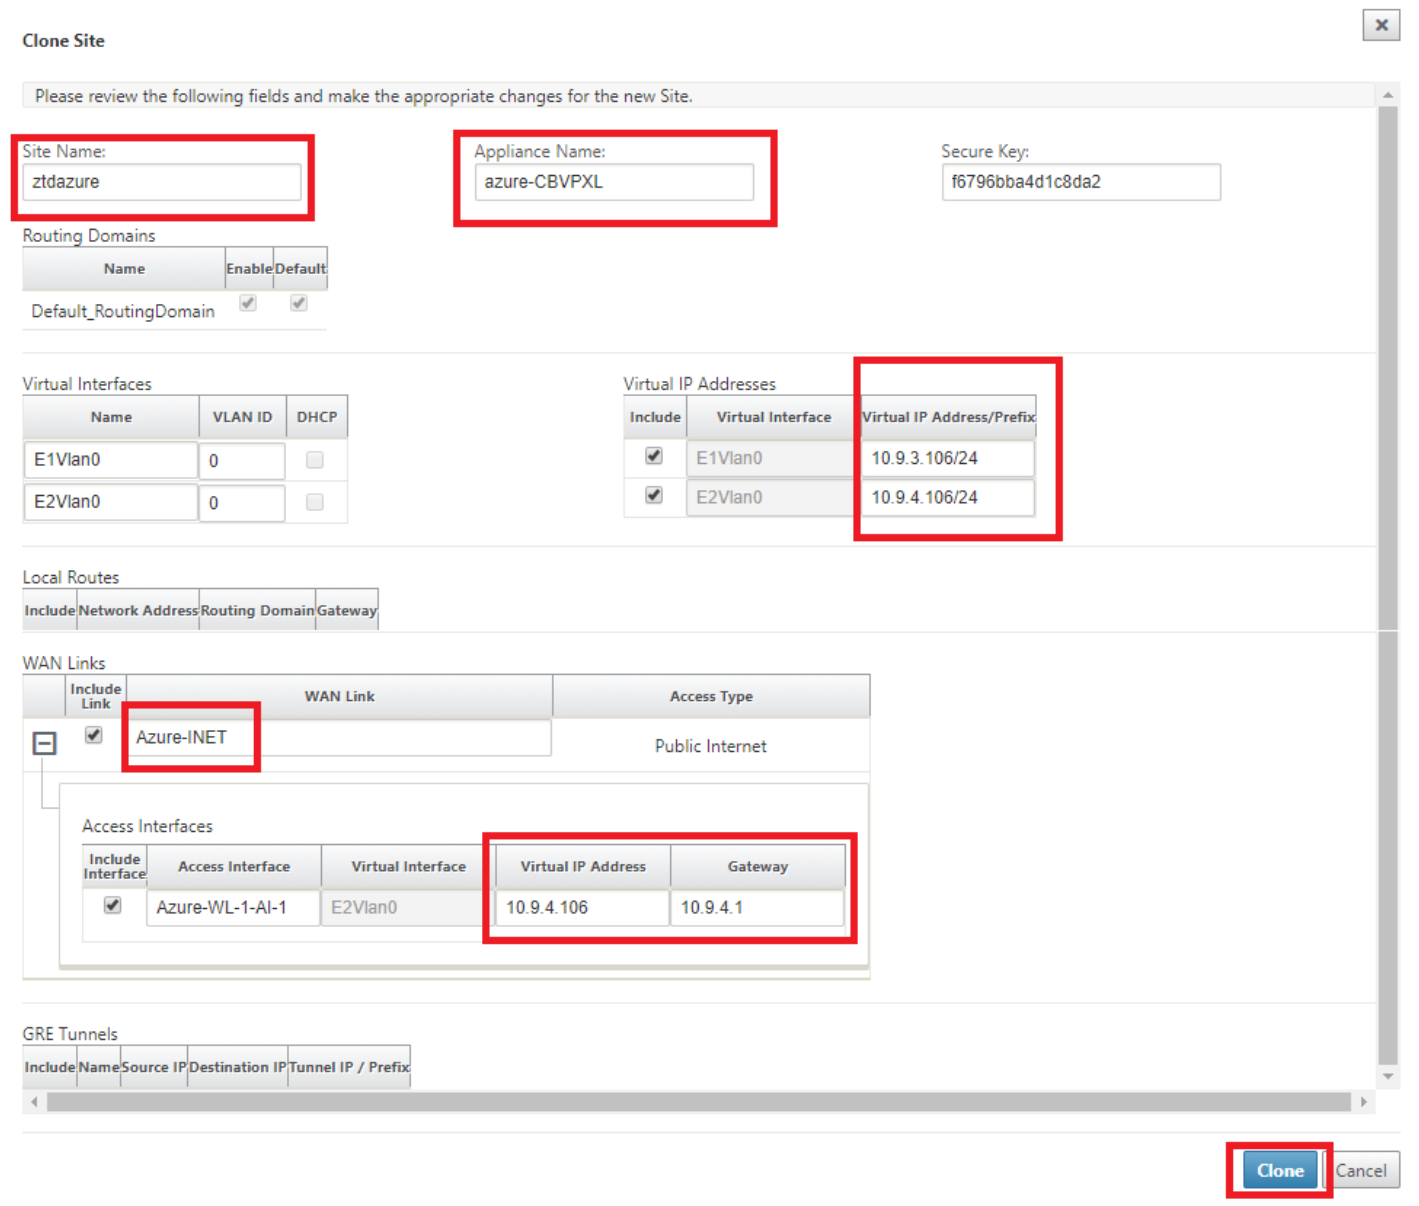

h) After cloning a new site, navigate to the site's **Basic Settings**, and verify that the Model of SD-WAN is correctly selected which would support the zero touch service.

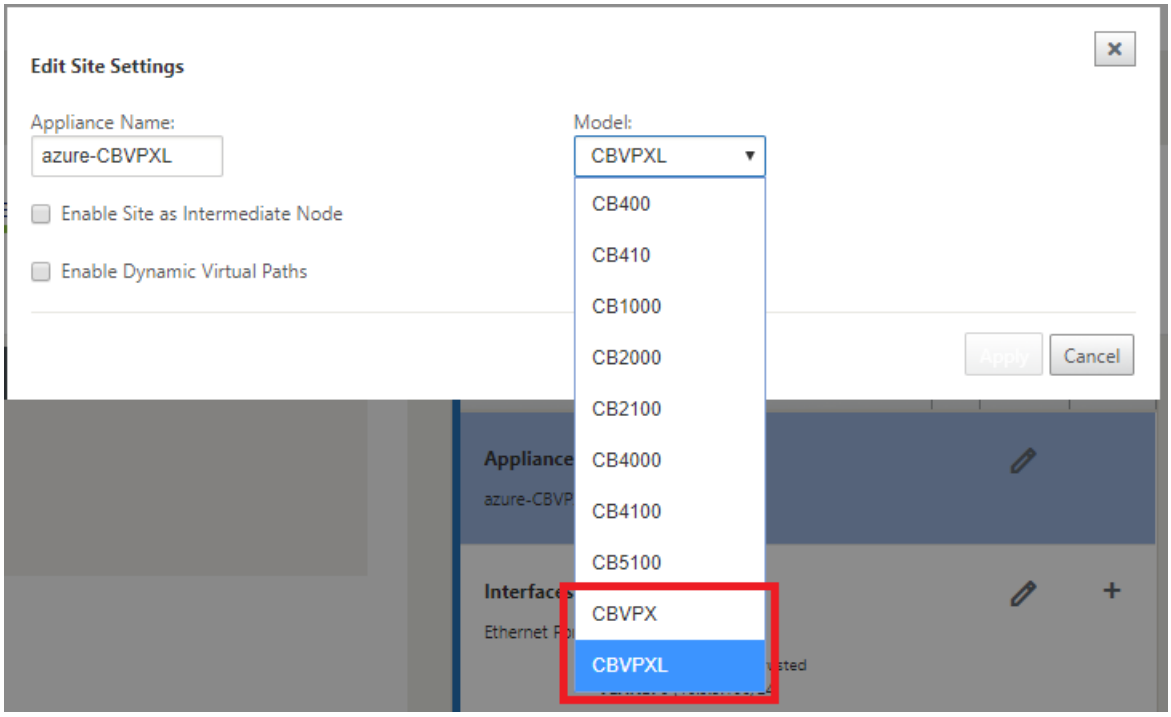

i) Save the new configuration on SD-WAN Center, and use the export to the "**Change Management inbox**" option to push the configuration using Change Management.

j) Follow the Change Management procedure to properly stage the new configuration, which makes the existing SD-WAN devices aware of the new site to be deployed via zero touch, you will need to utilize the "*Ignore Incomplete*" option to skip attempting to push the configuration to the new site that still needs to go through the ZTD workflow.

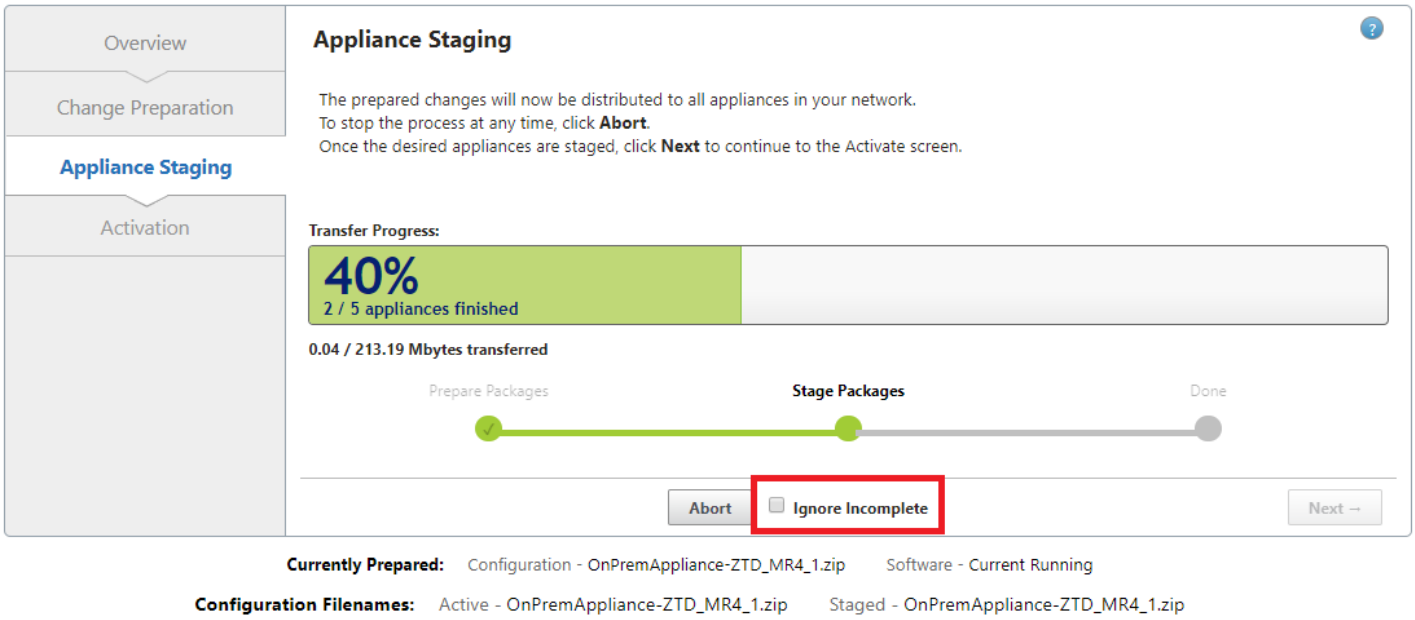

Navigate to the SD-WAN Center's Zero Touch Deployment page, and with the new active configuration running, the new site will be available for SD-WAN Center Provision and Deploy Azure (Step 1 of 2)

a) In the Zero Touch Deployment page, login with your Citrix account credentials. Under the **Deploy New Site** tab, select the running network configuration file.

b) After the running configuration file is selected, the list of all the branch sites with ZTD capable NetScaler SD-WAN

devices will be displayed.

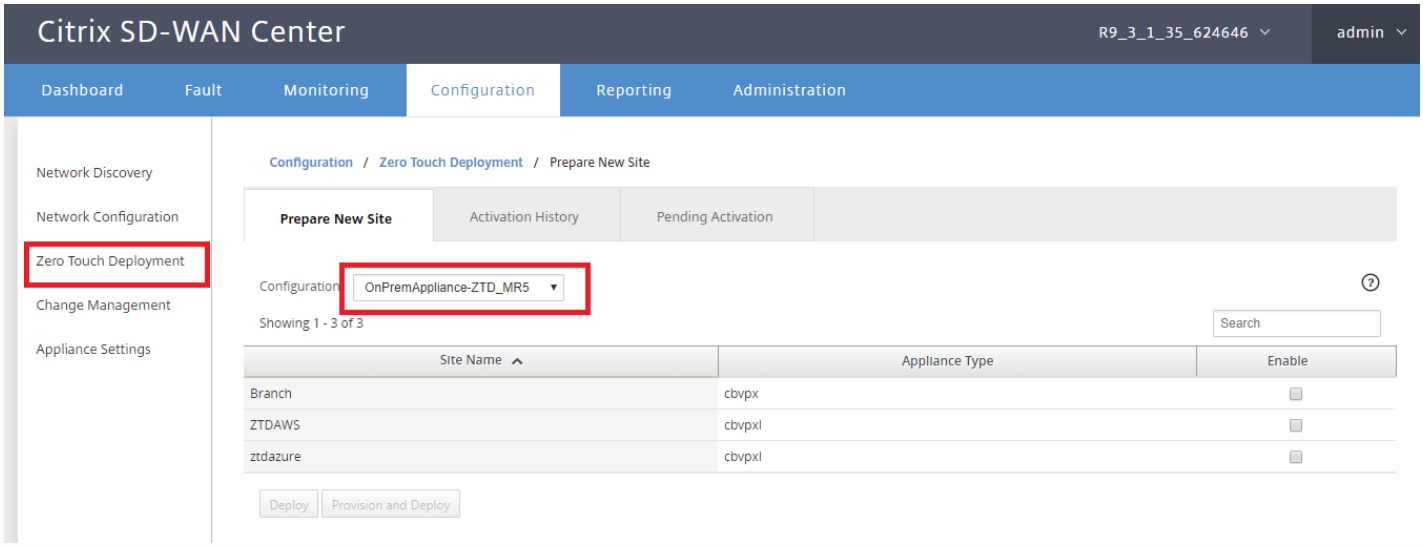

c) Select the target cloud site you want to deploy using the Zero Touch service, click **Enable**, and then **Provision and Deploy**.

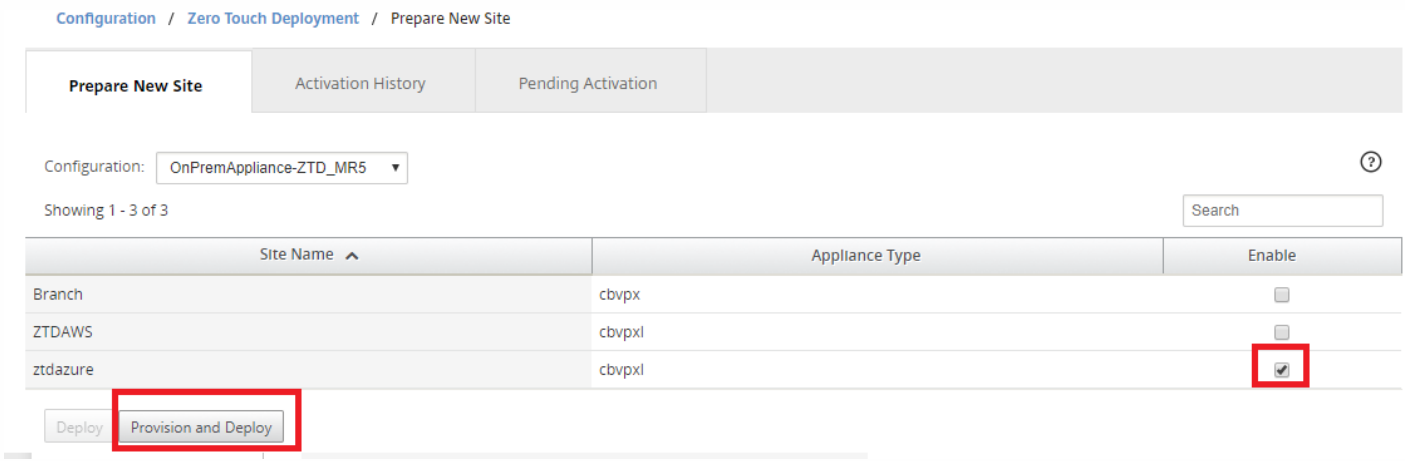

d) A pop-up window will appear, where the NetScaler SD-WAN Admin can initiate the deployment for Zero Touch. Validate that the site name complys with the requirements on Azure (lowercase with no special characters). Populate an email address where the activation URL can be delivered, and select Azure as the **Provision Type** for the desired Cloud, before clicking **Next**.

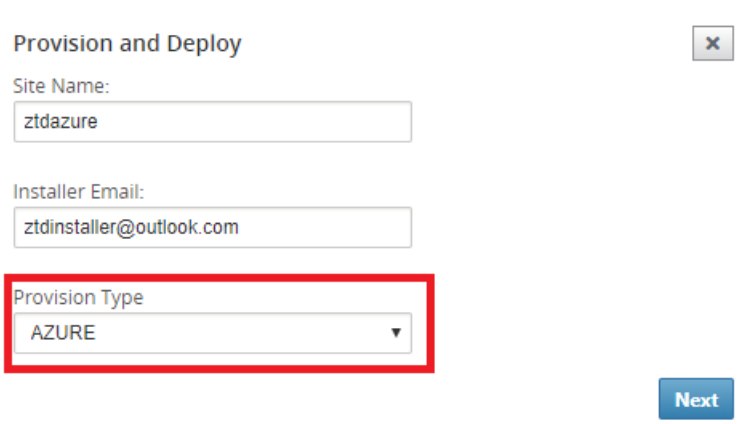

e) After clicking **Next**, the Provision and Deploy Azure (step 1of 2) window will require input of obtained from the Azure

account.

Copy and paste each required field after obtaining the information from your Azure account. The steps below outline how to obtain the required Subscription ID, Application ID, Secret Key, and Tenant ID from your Azure account, then proceed by clicking **Next**.

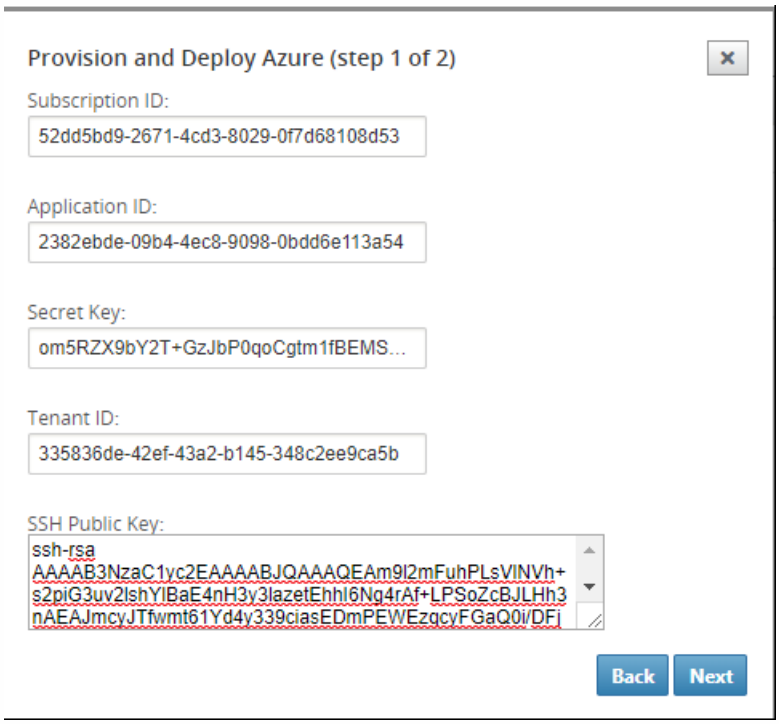

1. On the Azure account, we can identify the required ***Subscription ID*** by navigating to “More Services” and select **Subscriptions.**

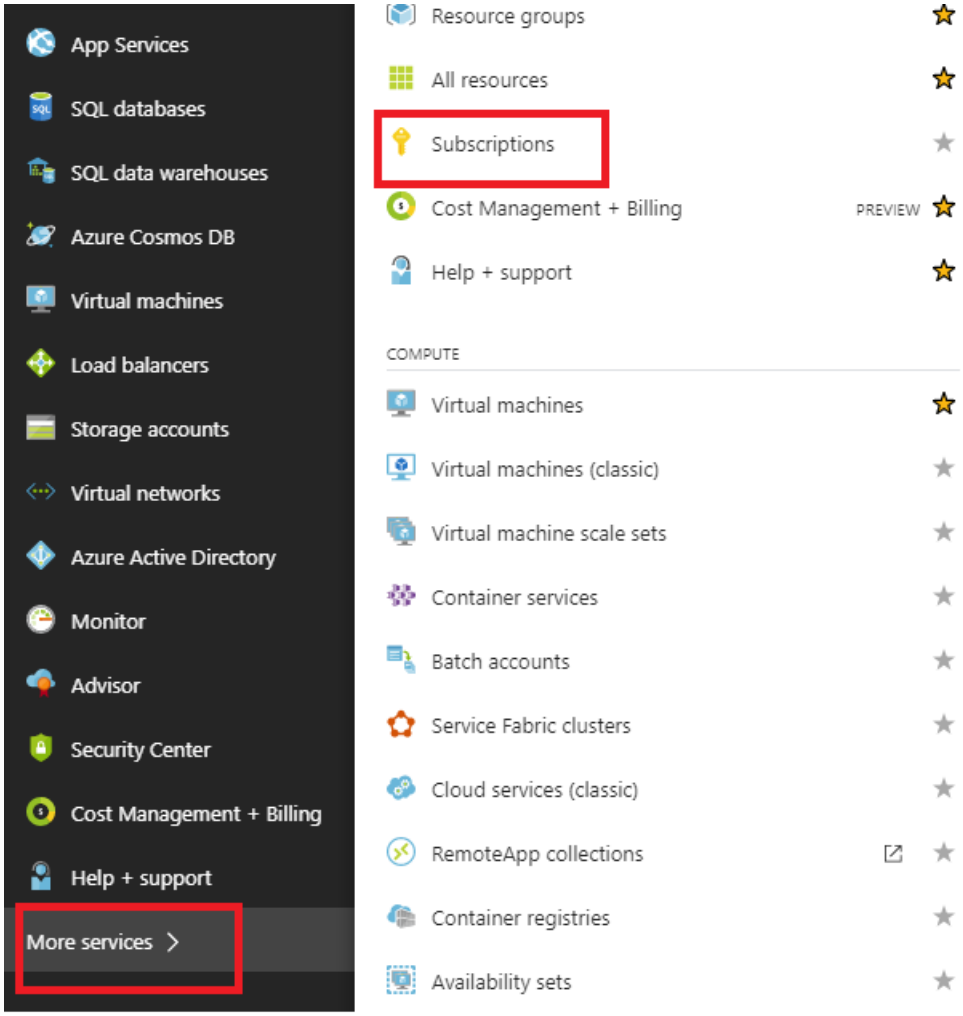

2. To identify the required ***Application ID***, navigate to Azure Active Directory, Application registrations, and click N**ew application registration.**

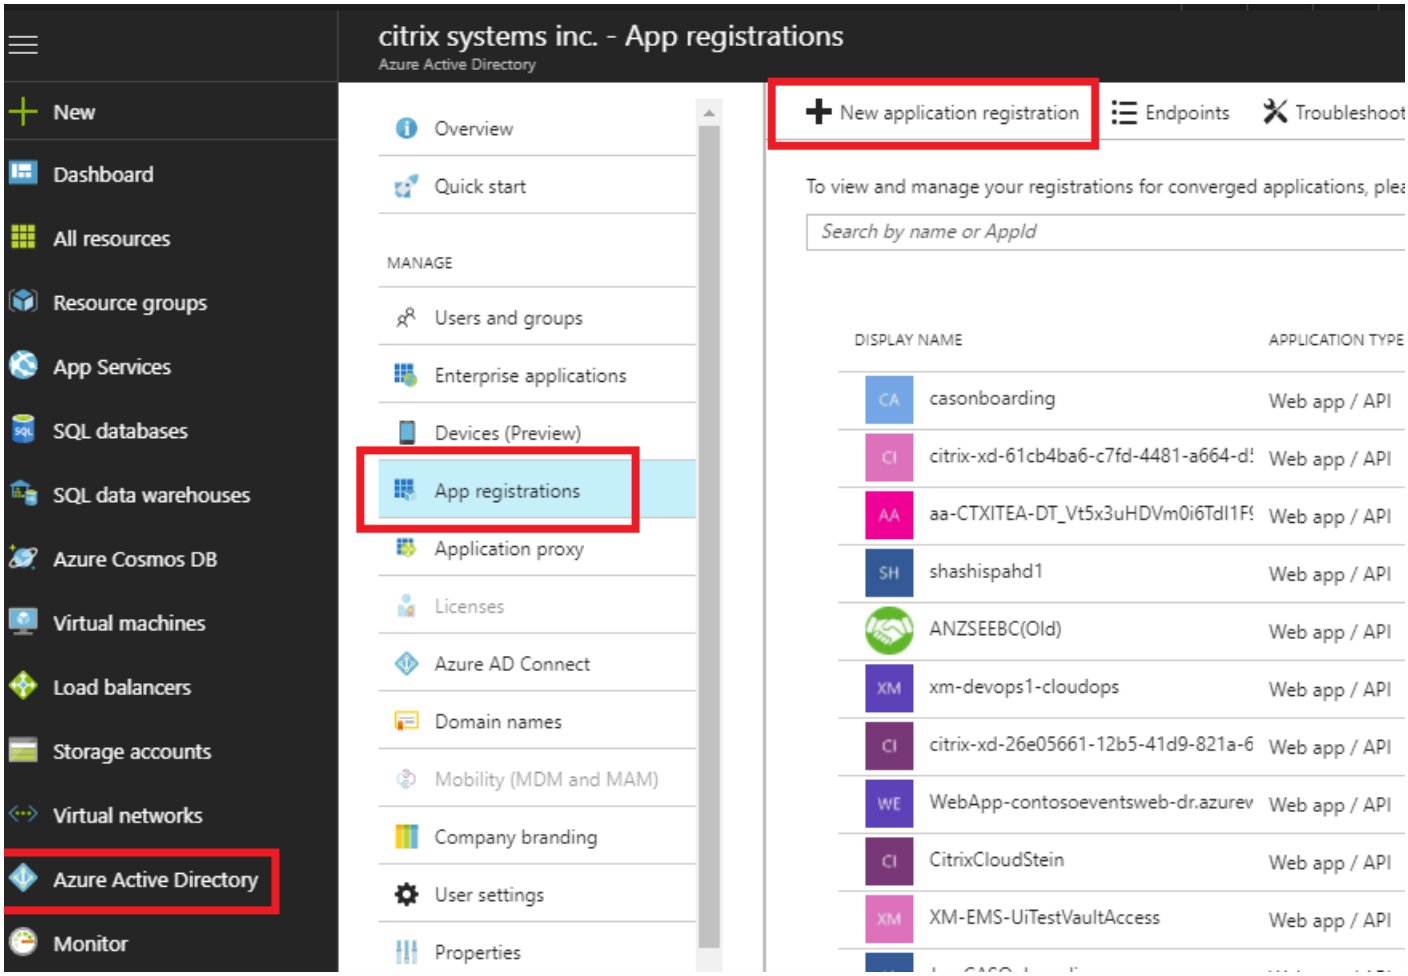

3. In the app registration create menu, enter a Name and a Sign-on URL (this can be any URL, the only requirement is that it must be valid), then click **Create**.

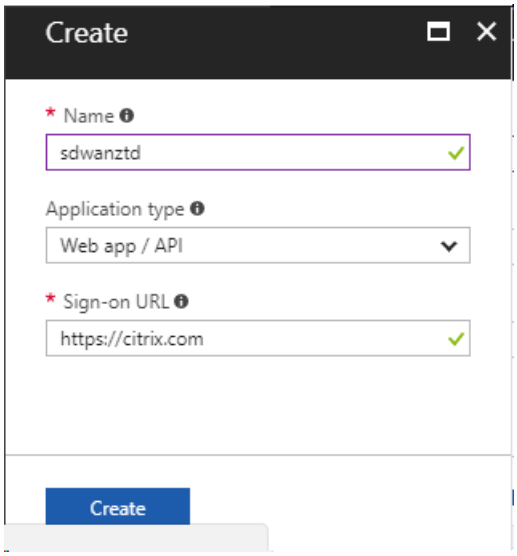

4. Search for and open the newly created Registered App, and note the Application ID.

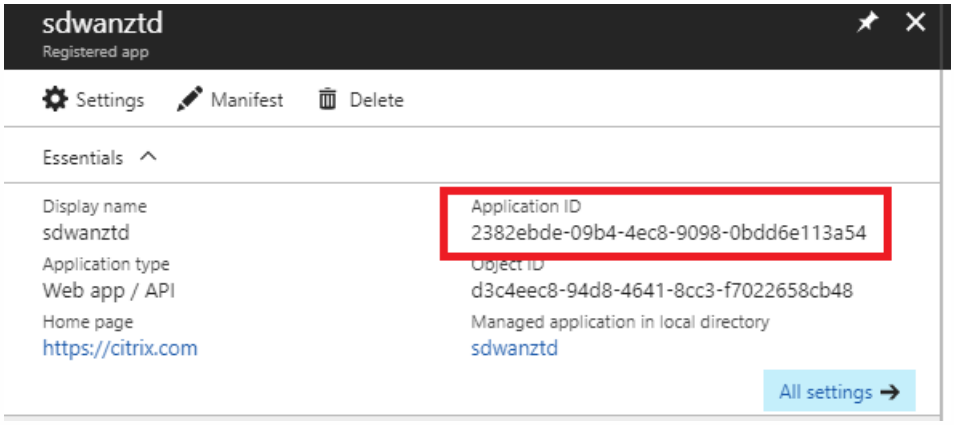

5. Again open the newly created Registration App, and to identify the required *Security Key*, under API Access, select **Required permissions,** to allow a third party to provision and instance. Then select **Add**.

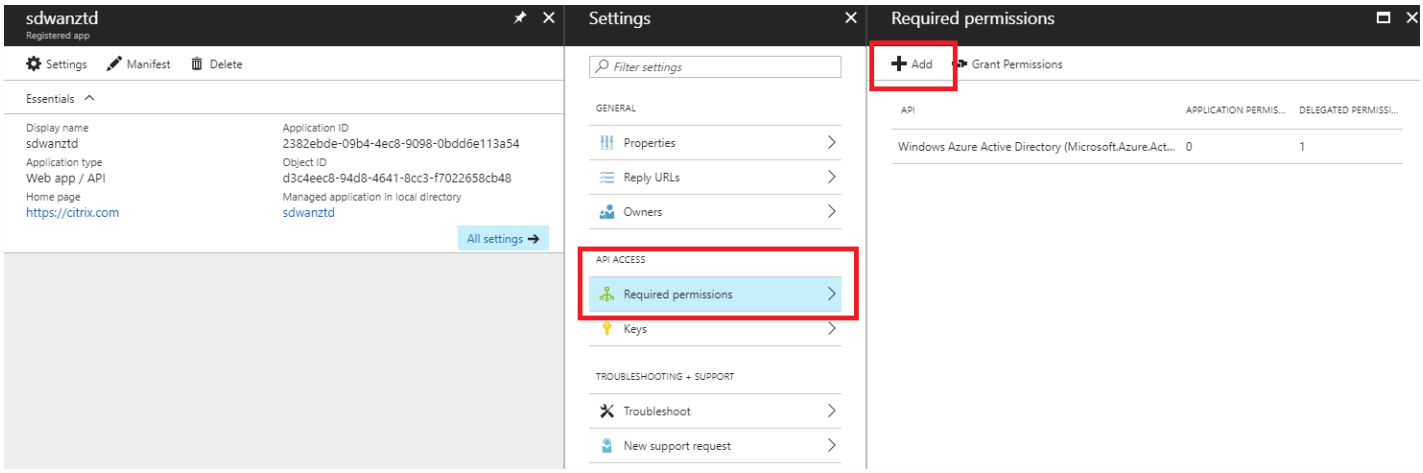

6. When adding the Required permissions, **Select an API,** then highlight **Windows Azure Service Management API.**

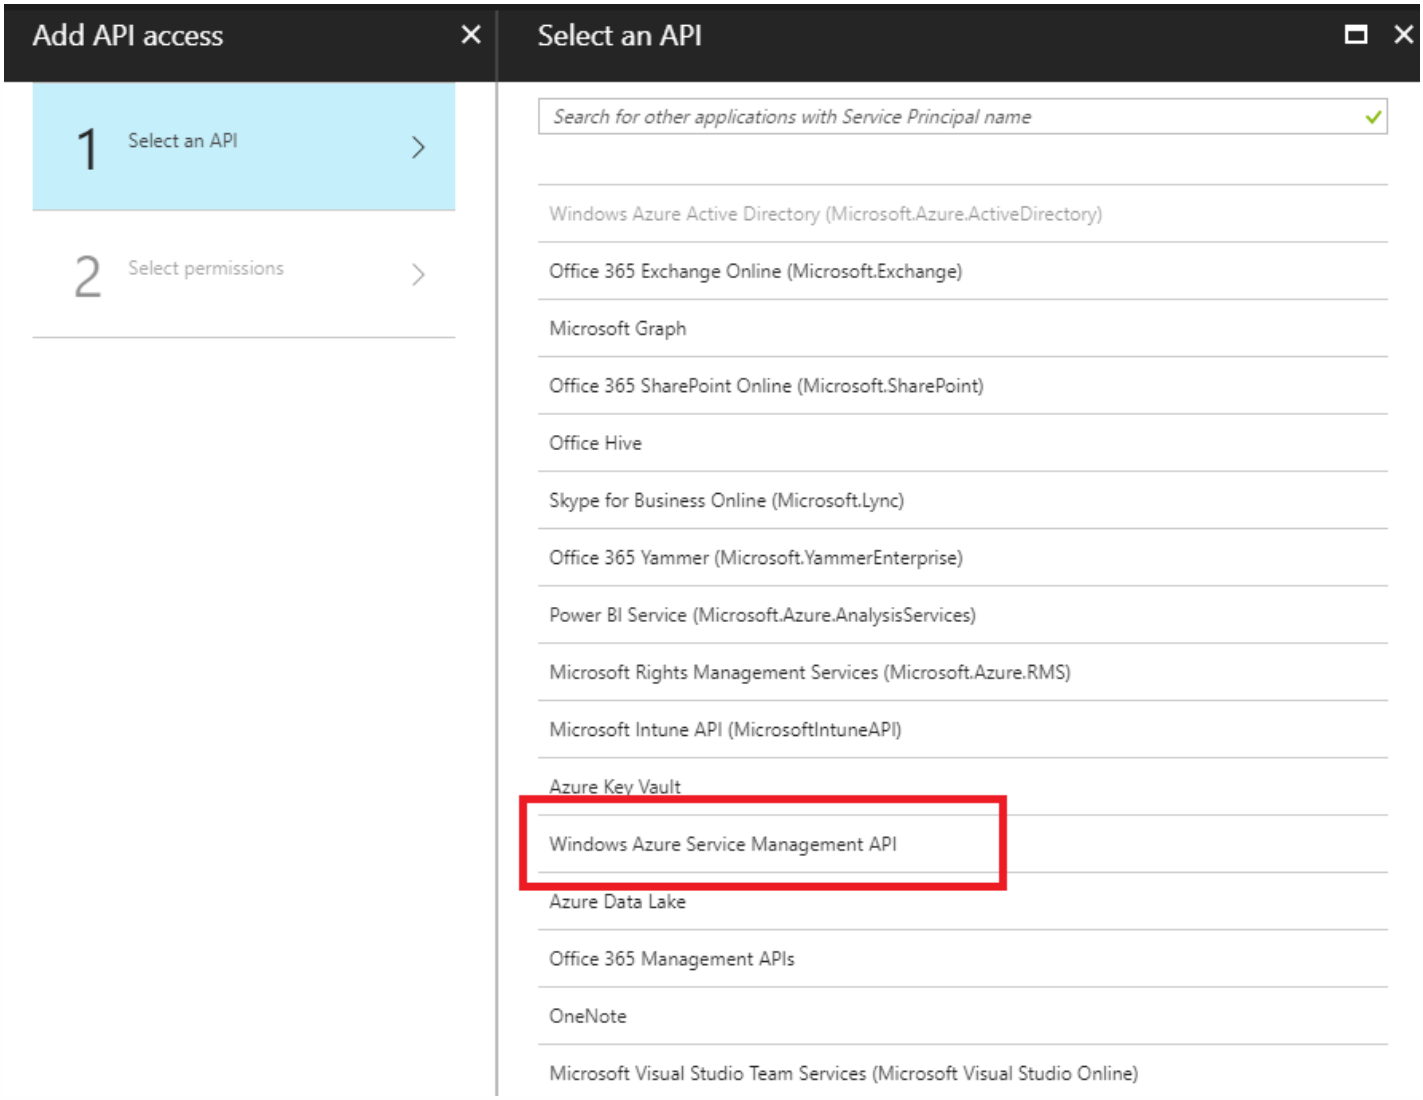

7. Enable **Delegate Permissions** to provision instances, then click **Select** and **Done**.

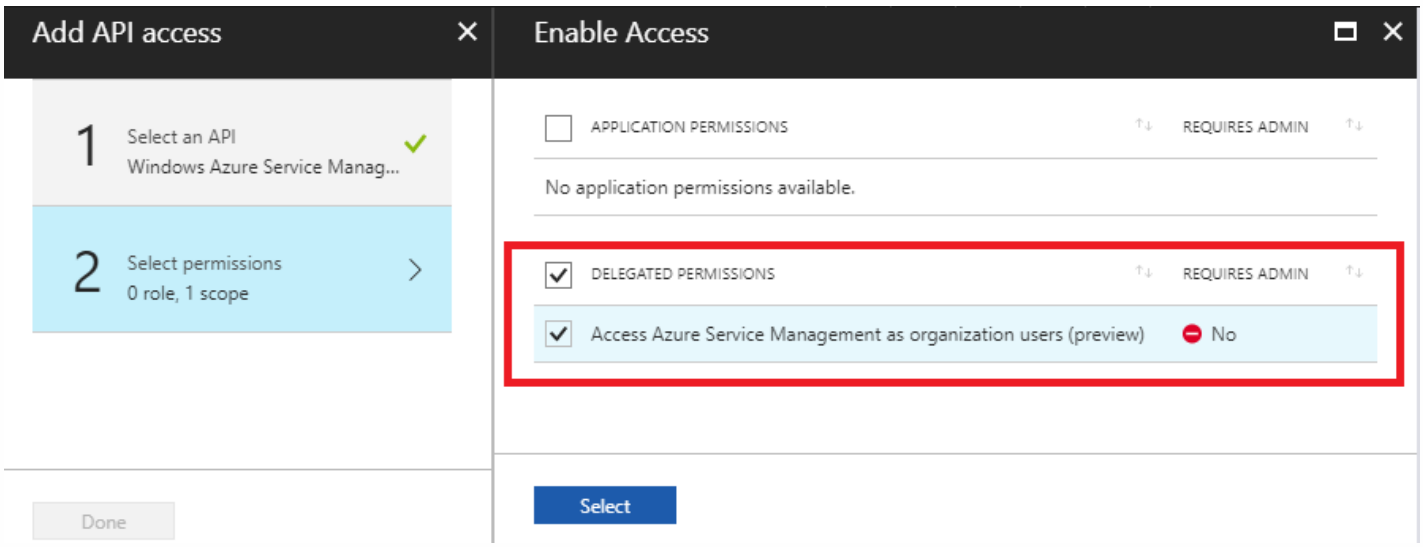

8. For this Registered App, under API Access, select **Keys,** and create a secret **key description** and the desired **duration** for the key to be valid. Then click **Save which** will produce a **secret key** (the key is only required for the provisioning process, it can be deleted after the instance is made available).

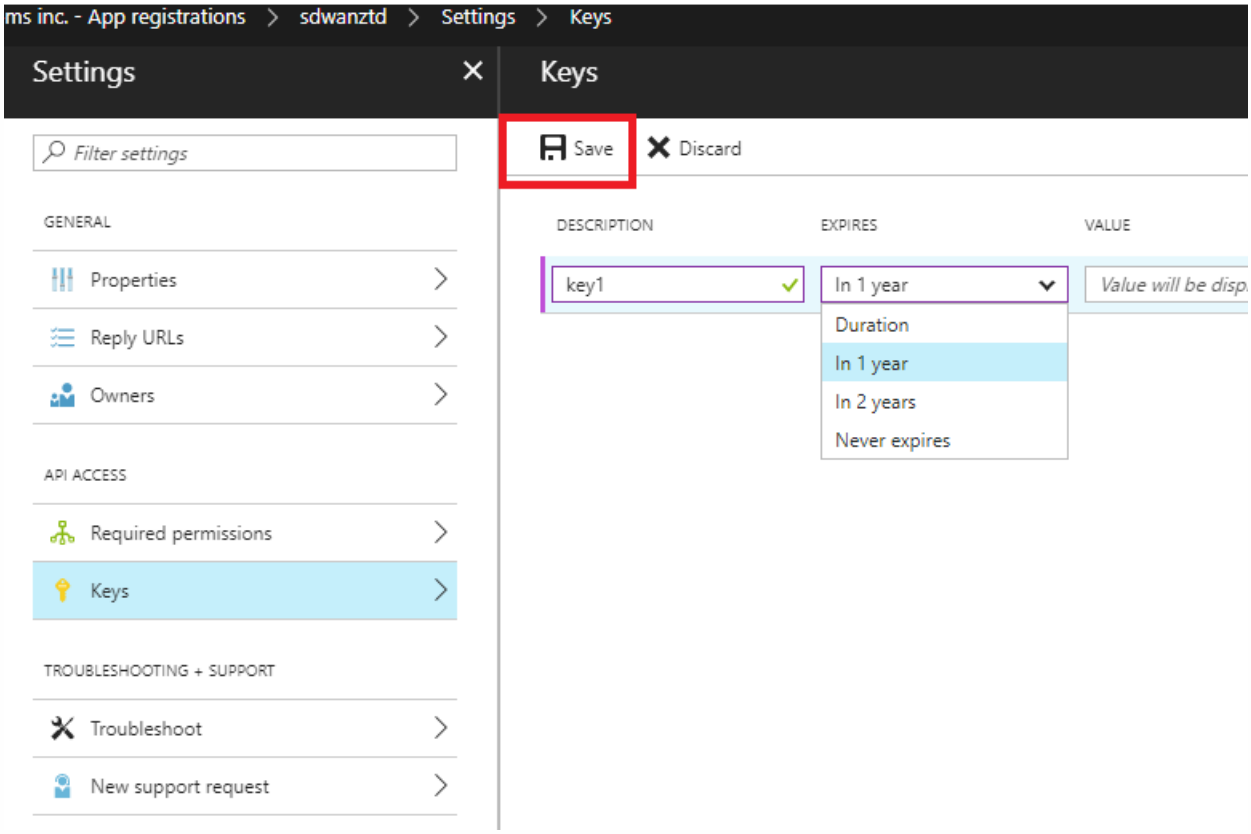

9. Copy and save the secret key (note you will not be able to retrieve this later).

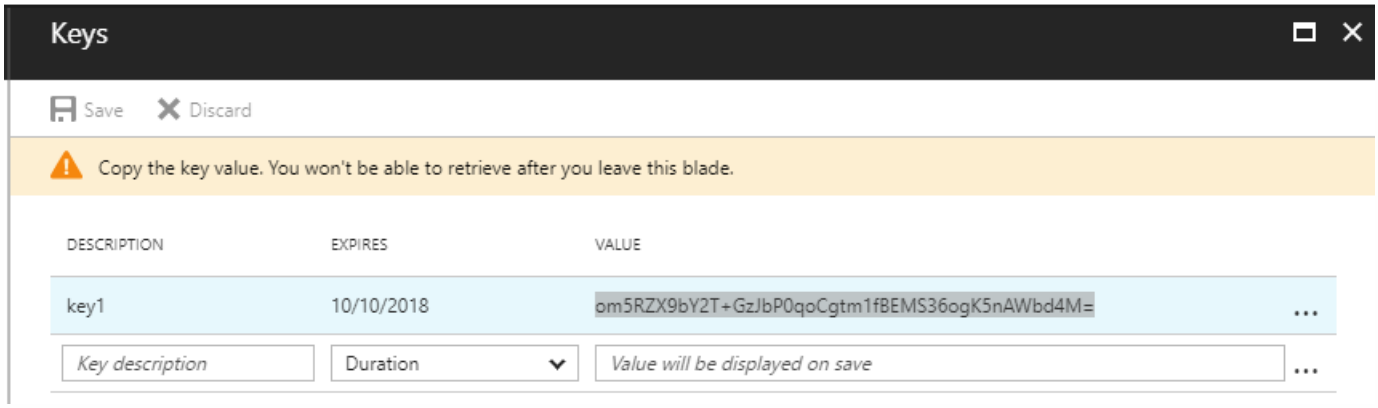

10. To identify the required ***Tenant ID***, navigate back to the App registration pane, and select **Endpoints**.

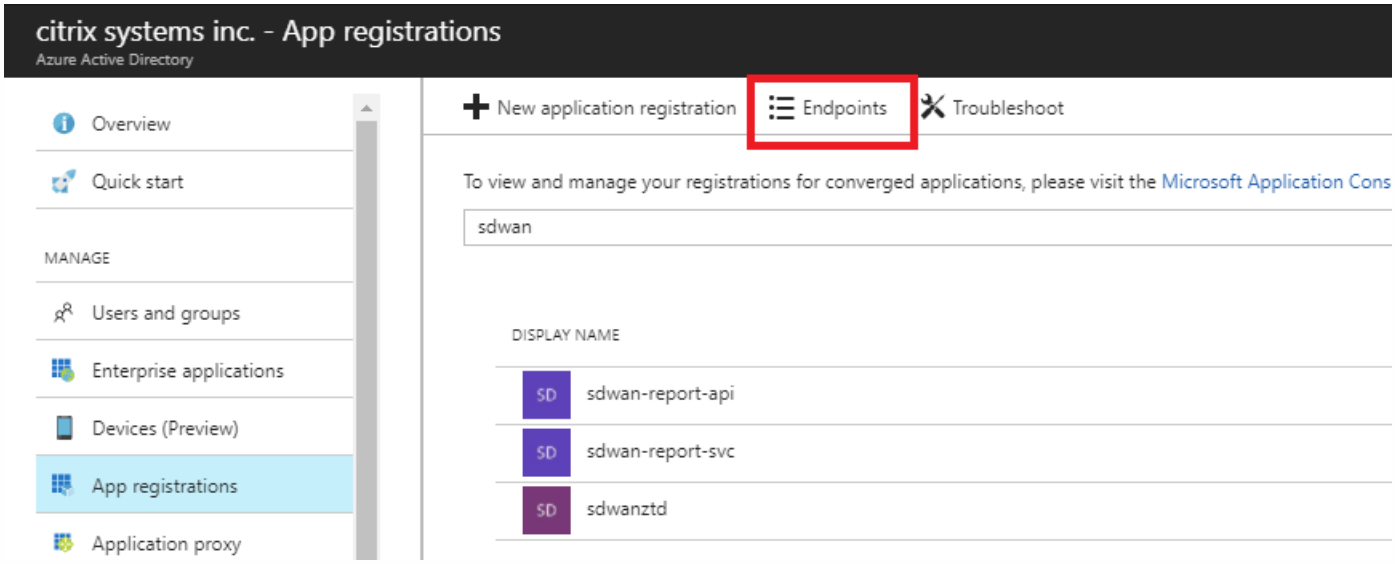

11. Copy the **Federation Metadata Document**, to identify your Tenant ID (note the Tenant ID is 36-character string located between the "online.com/" and the "/federation" in the URL).

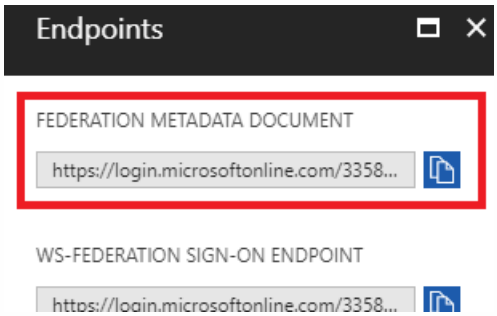

12. The last item required is the ***SSH Public Key***. This can be created using Putty Key Generator or ssh-keygen and will be utilized for authentication, eliminating the need for passwords to log in. The SSH public key can be copied (including the heading ssh-rsa and trailing rsa-key strings). This public key will be shared through SD-WAN Center input to the Citrix Zero Touch Deployment Service.

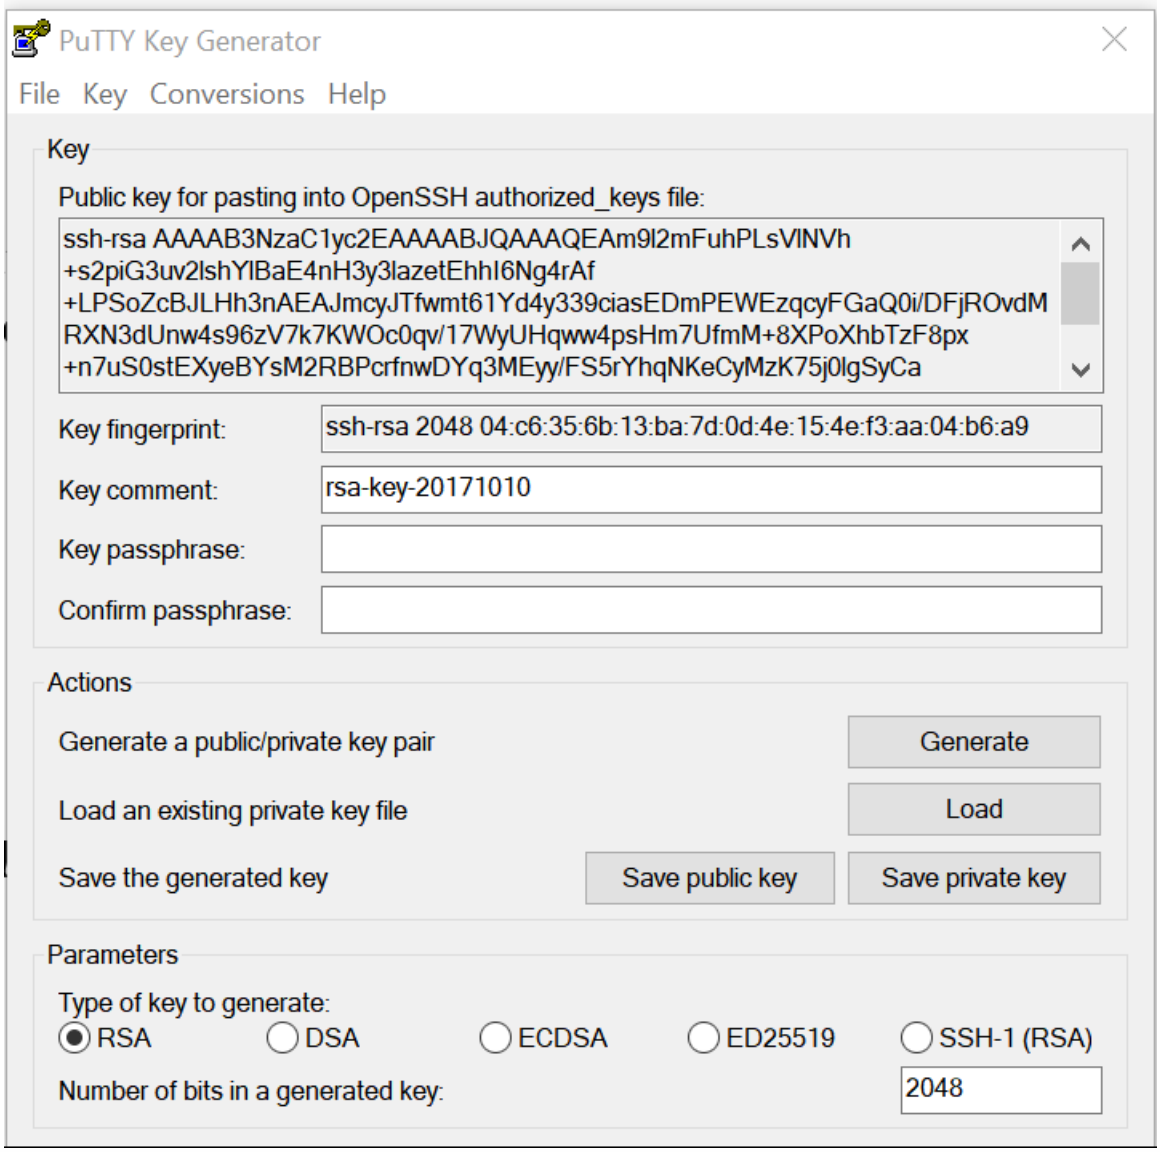

13. Additional steps are required to assign the application a role. Navigate back to More Services, then Subscriptions.

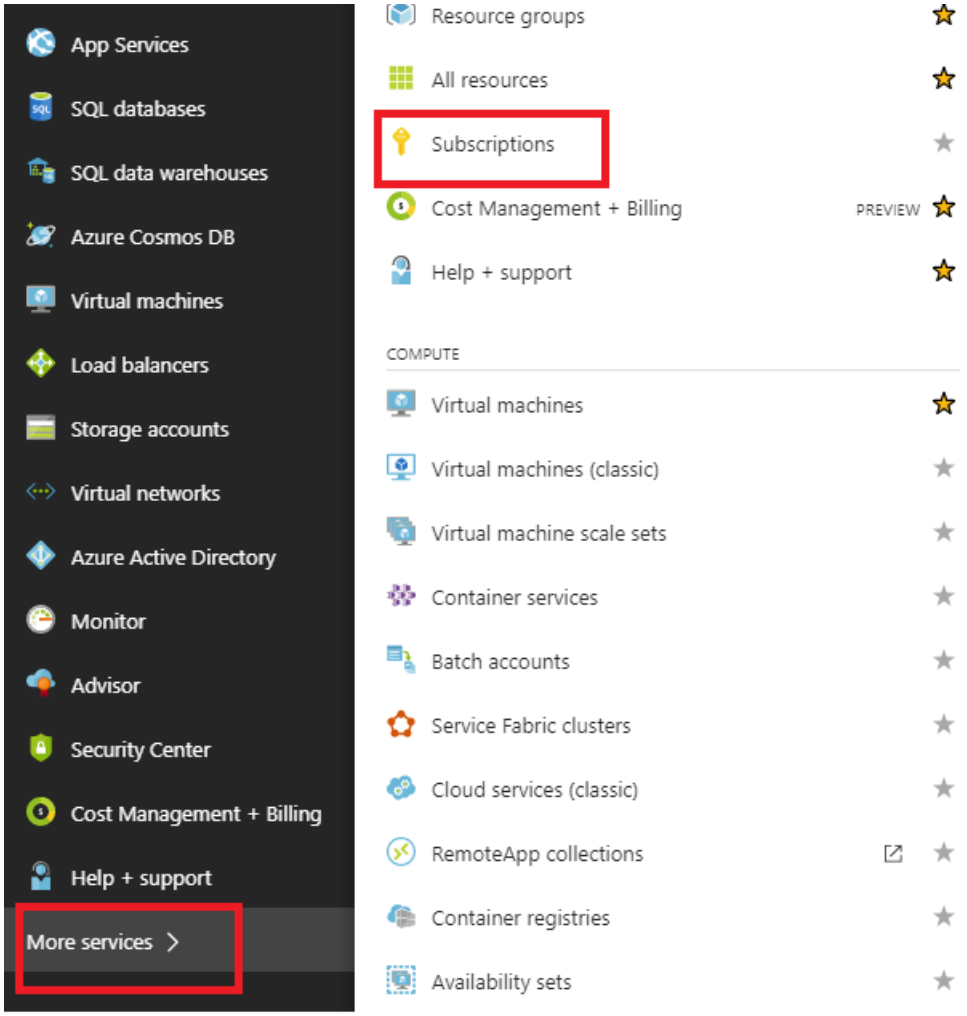

14. Select the active subscription, then **Access control (AIM)**, next click **Add.**

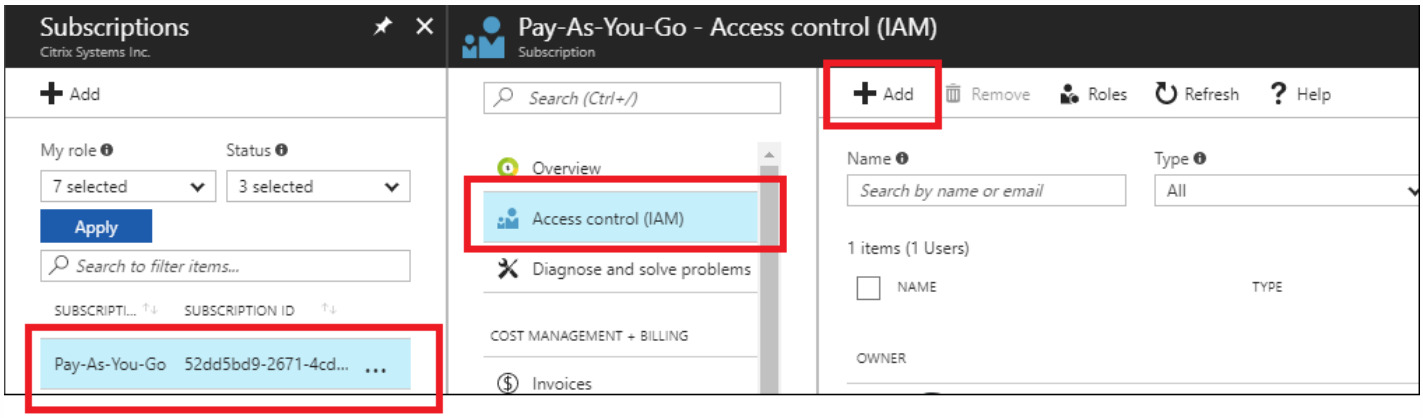

15. In the add permissions pane, select "**Owner**" role, assign access to **"Azure AD user, group, or application"** and search for the registered app in the Select field to allow the Zero Touch Deployment Cloud Service to create and configure the instance on the Azure subscription. Once the app is identified, select it and make sure it populates as a Selected member before clicking **Save**.

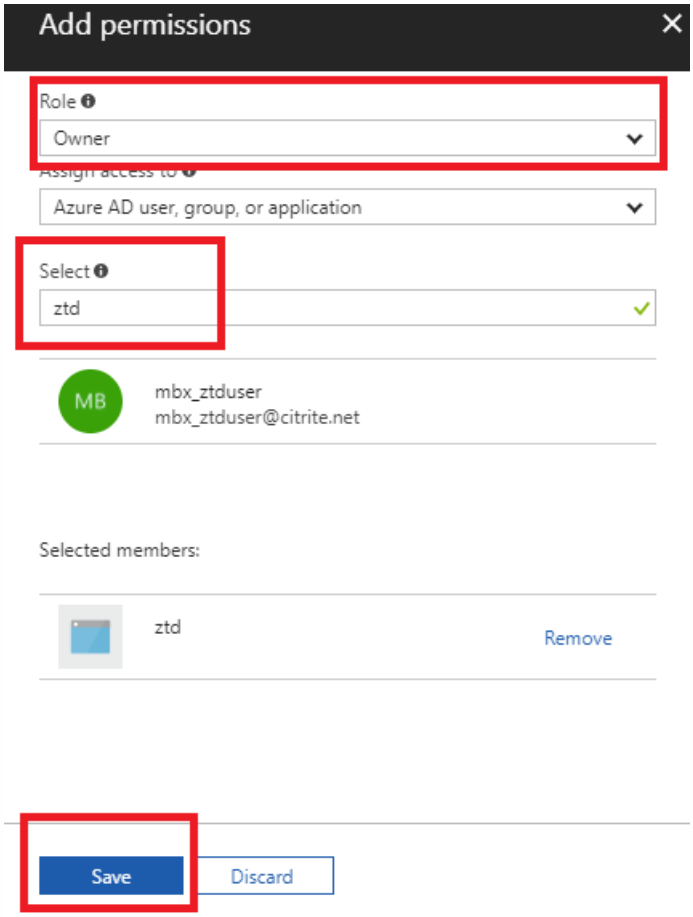

16. After collecting the required inputs and entering them into SD-WAN Center, click **Next**. If the inputs are not correct, you will encounter an authentication failure.

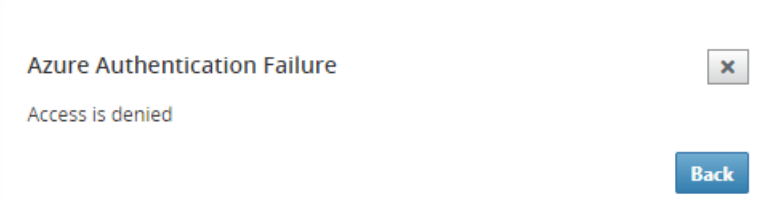

SD-WAN Center Provision and Deploy Azure (Step 2 of 2)

a) Once the Azure authentication is successful, populate the appropriate fields to select the desired Azure Region, and the appropriate Instance Size, then click **Deploy**.

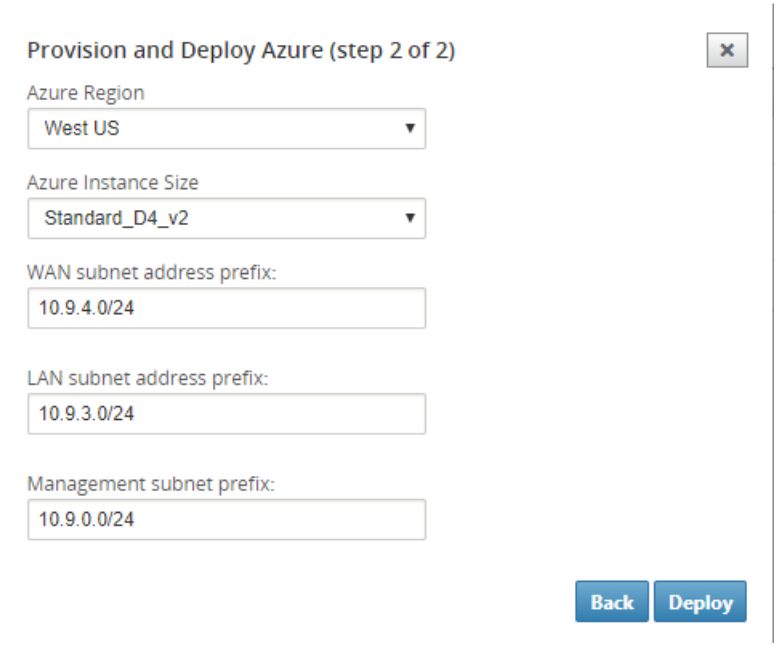

b) Navigating to the **Pending Activation** tab in SD-WAN Center, will help track the current status of the deployment.

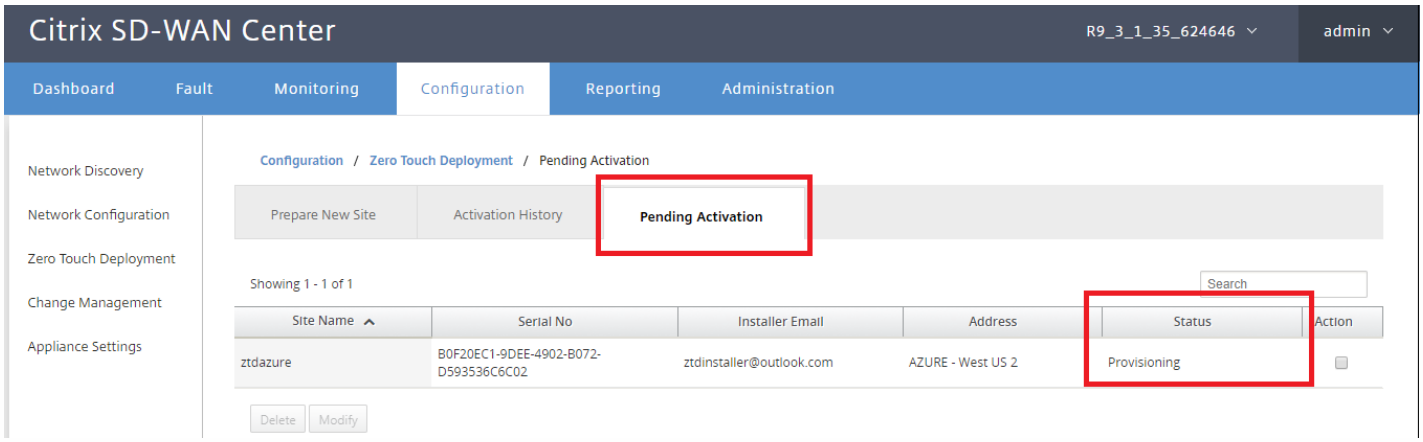

c) An email with an activation code will be delivered to the email address inputted in step 1, obtain the email and open the **activation URL** to trigger the process and check the activation status.

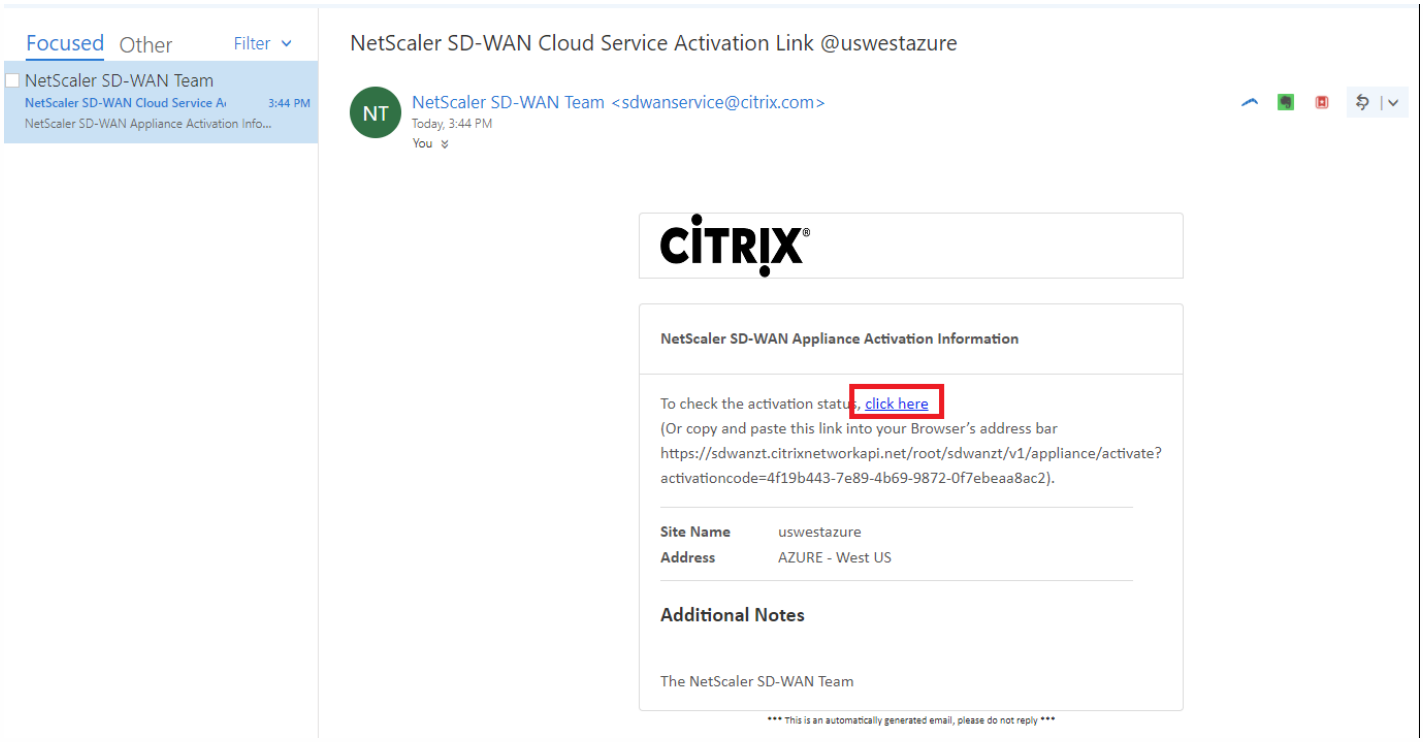

d) An email with an activation URL will be delivered to the email address inputted in step 1. Obtain the email and open the **activation URL** to trigger the process and check the activation status.

e) It will take a few minutes for the instance to be provisioned by the SD-WAN Cloud Service. You can monitor the activity on the Azure portal, under **Activity log** for the **Resource Group** which is automatically created. Any issues or errors with the provisioning will be populated here, as well as replicated to SD-WAN Center in the Activation Status.

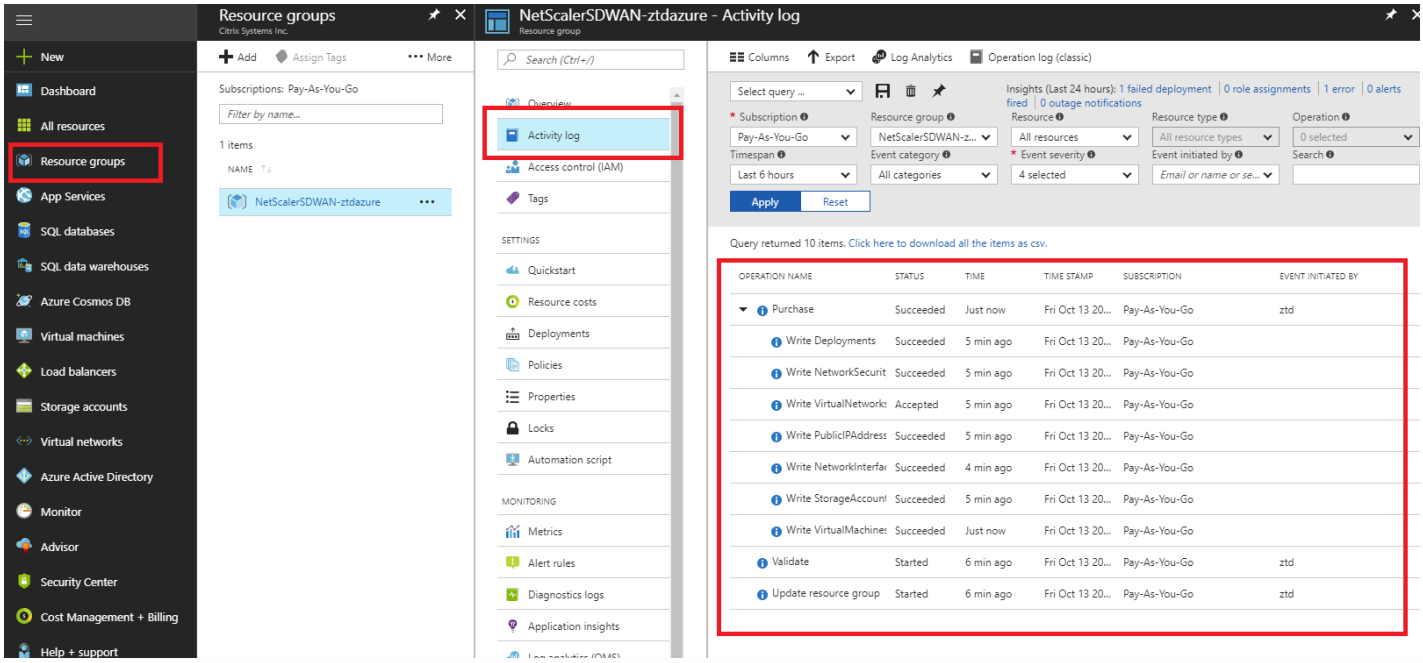

f) In the Azure portal, the successfully launched instance will be available under **Virtual Machines.** To obtain the assigned public IP, navigate to the Overview for the instance.

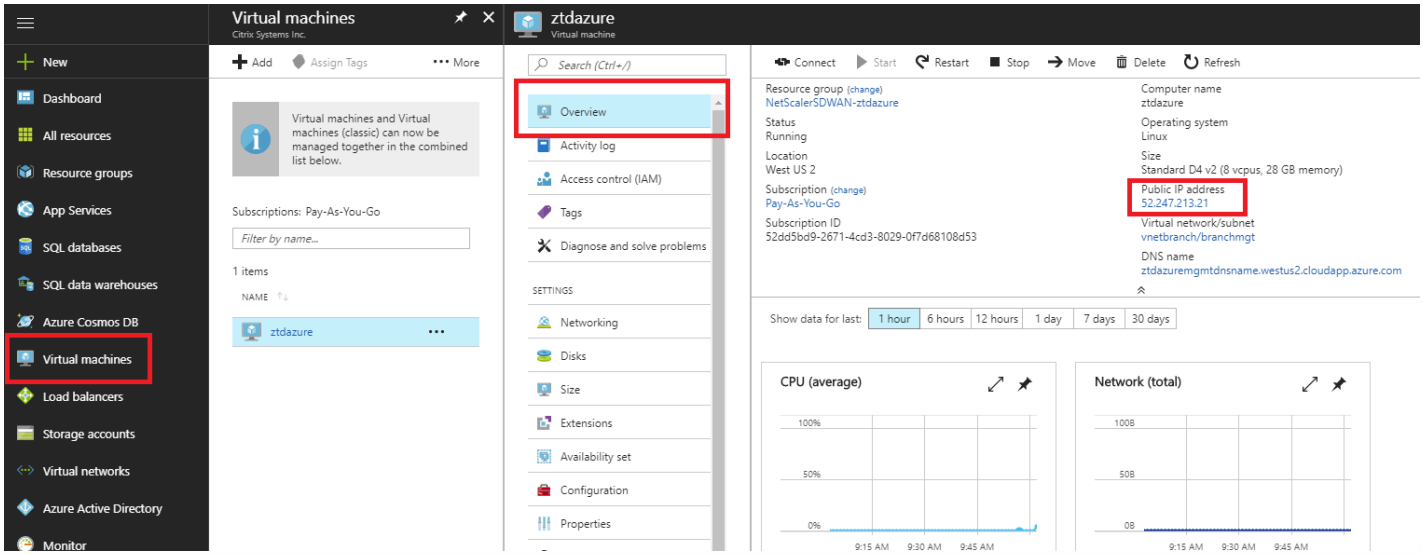

g) After the VM is in a running state, give it a minute before the service will reach out and start the process of downloading the configuration, software and license.

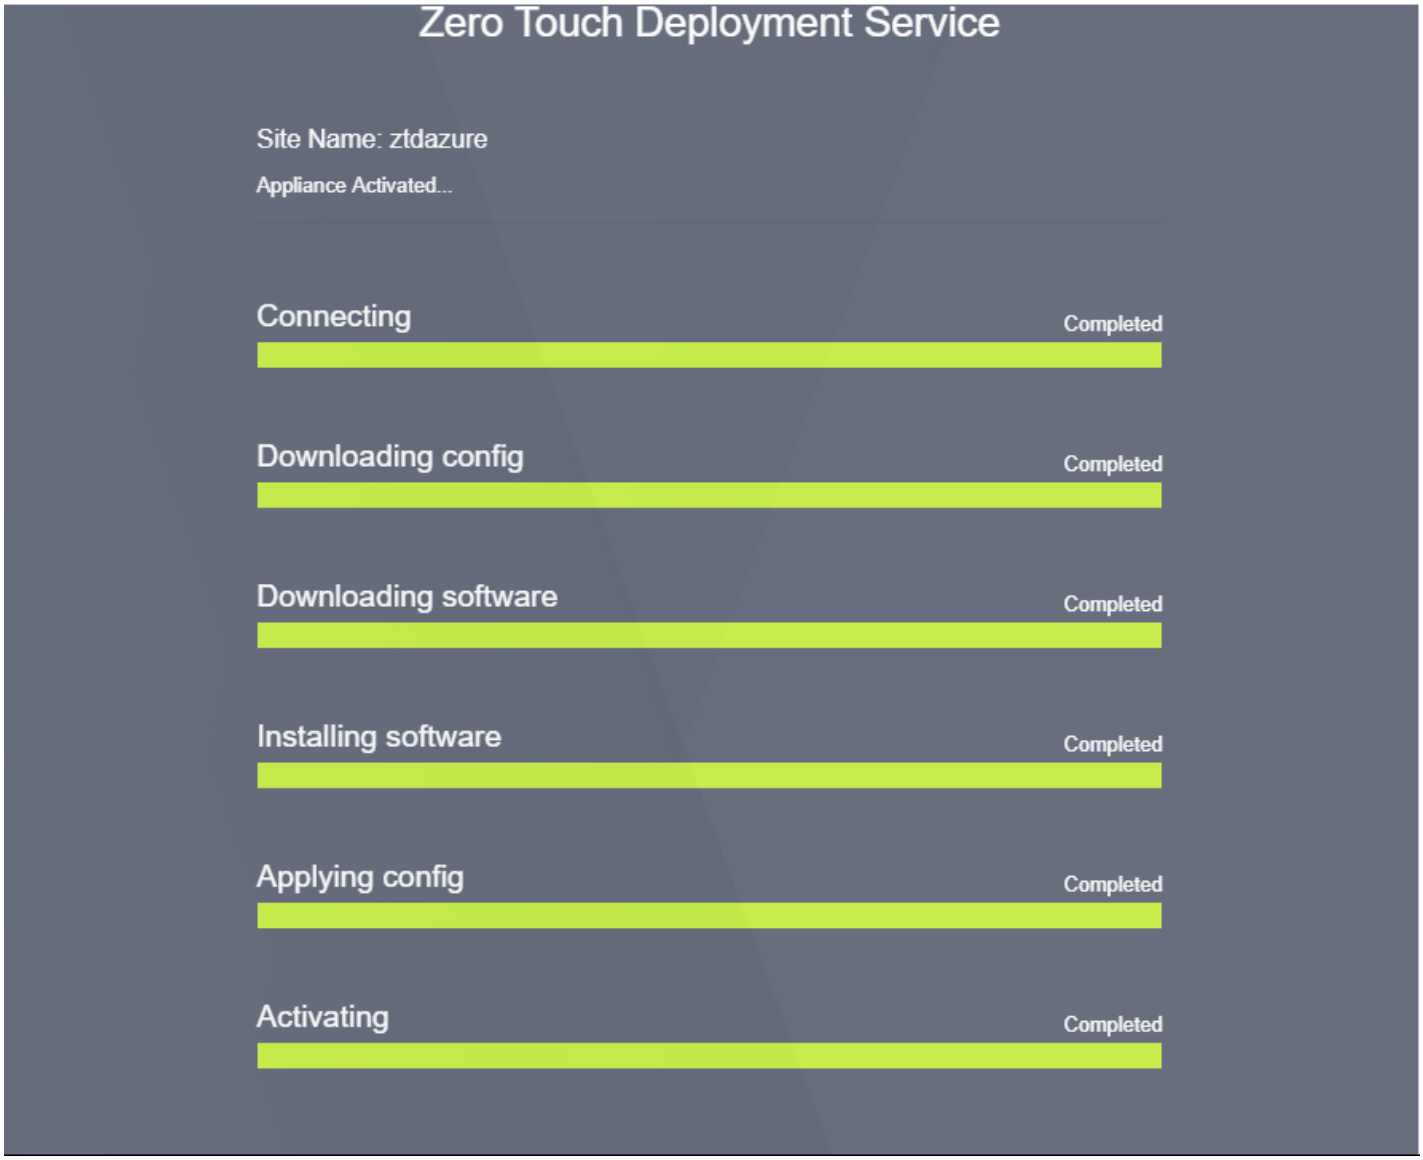

h) After each of the SD-WAN Cloud service steps are automatically complicated, log in to the SD-WAN instances web interface using the public IP obtained from the Azure portal.

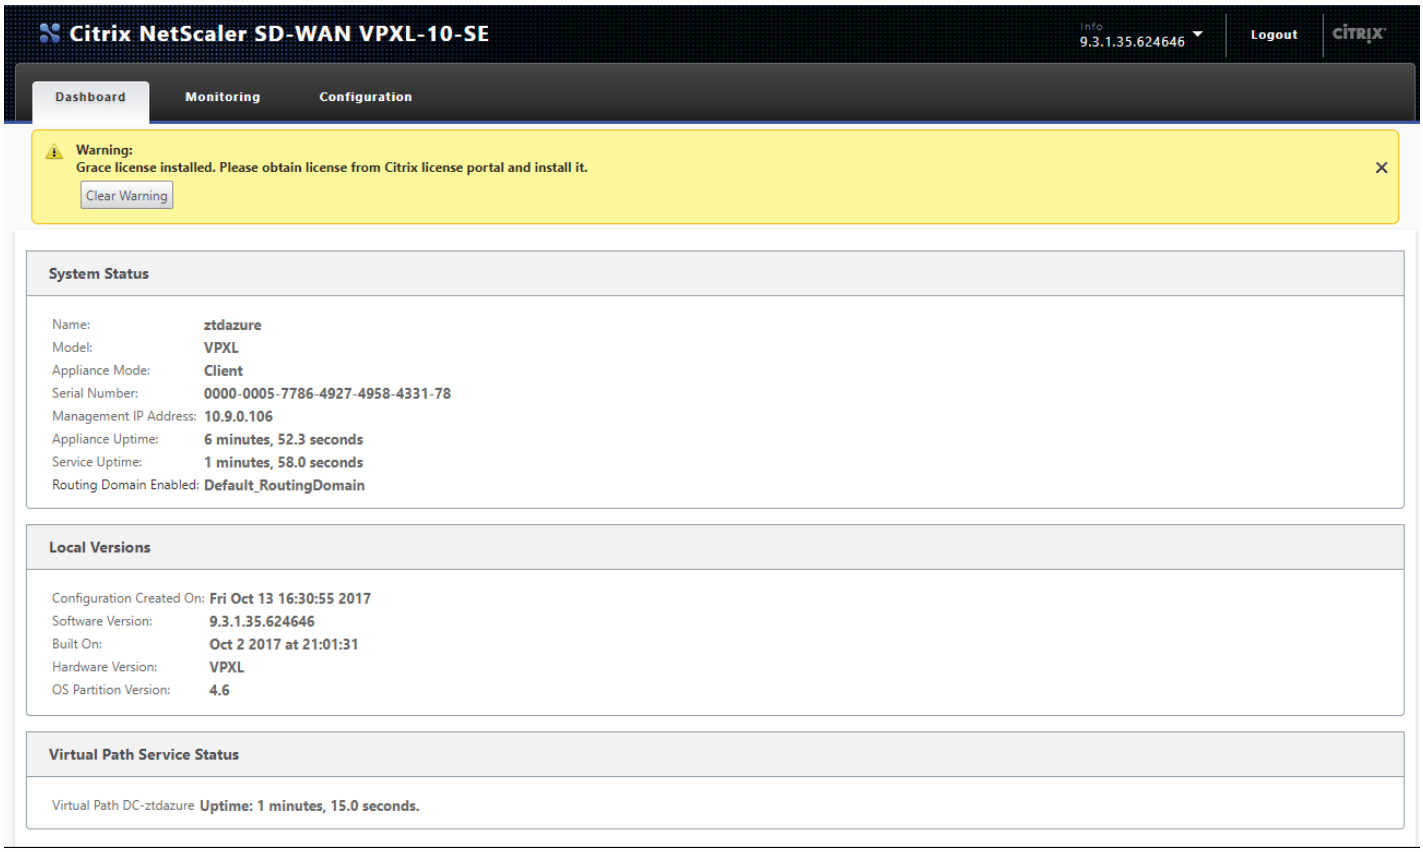

i) The NetScaler SD-WAN Monitoring Statistics page will identify successful connectivity from the MCN to the SD-WAN instance in Azure.

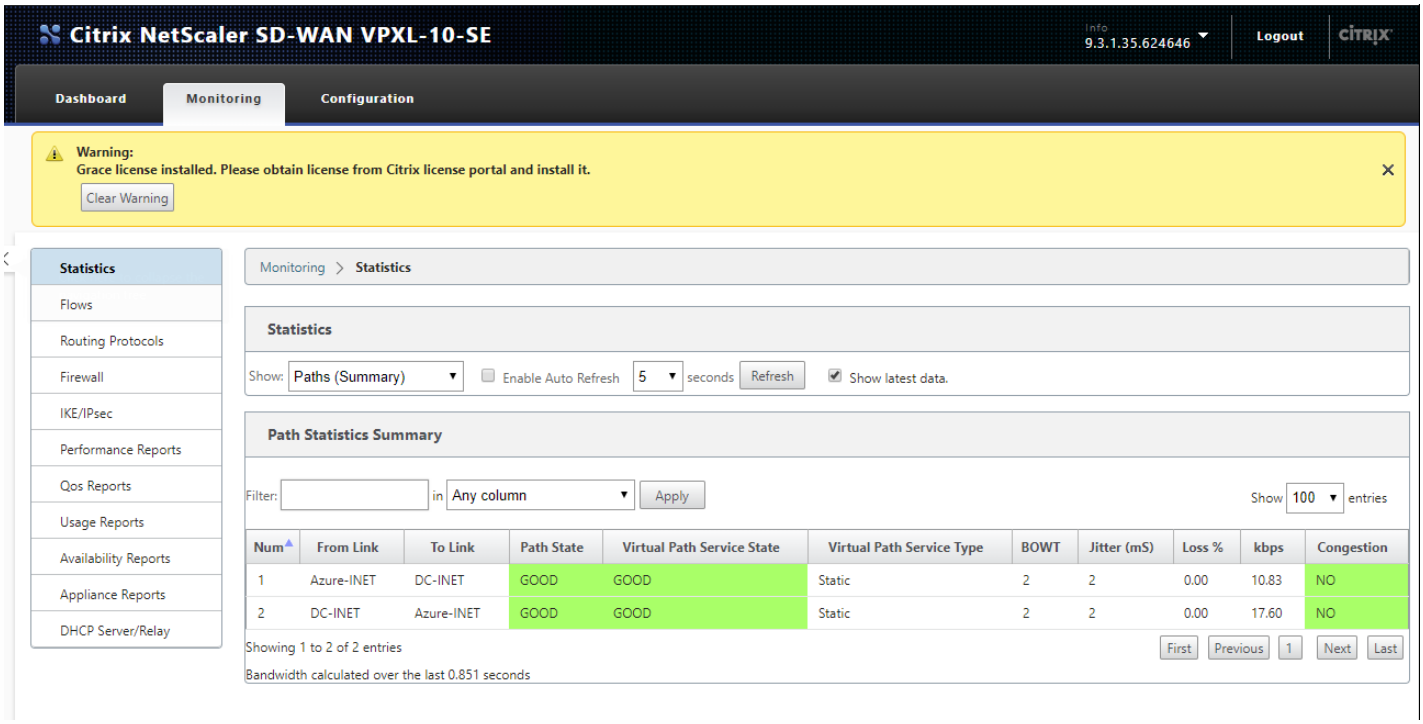

j) Furthermore, the successful (or unsuccessful) provisioning attempt will be logged in the SD-WAN Center's Activation History page.

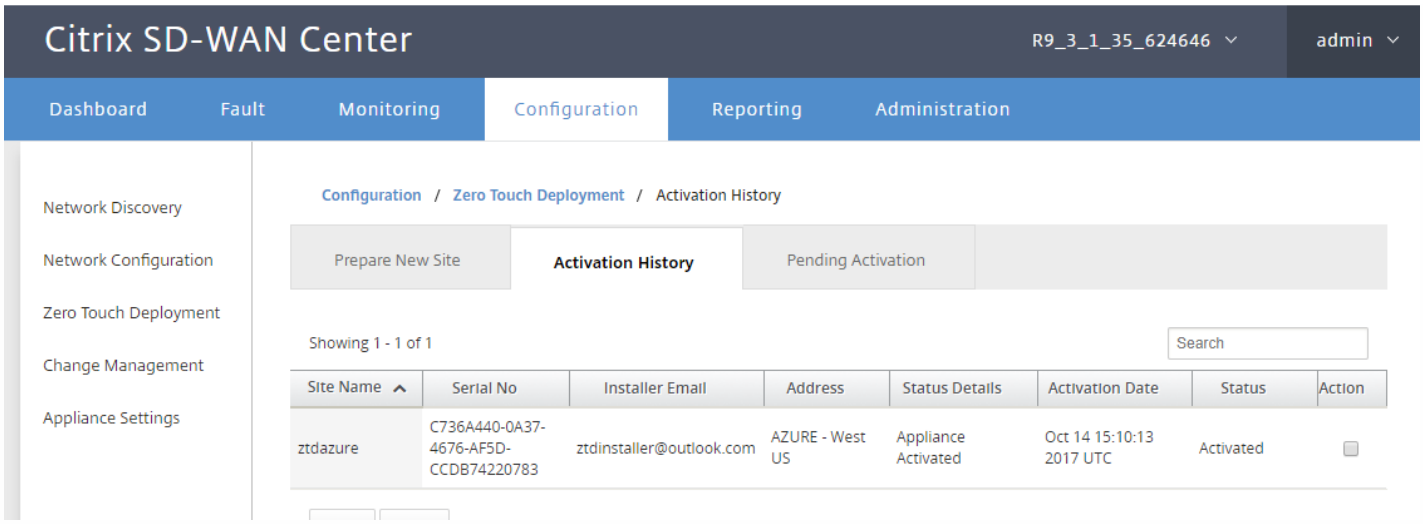

Configure 210-SE LTE

Jun 06, 2018

The following procedure describes the steps and workflow to configure a 210-SE LTE appliance.

Prerequisites

1. Have a standard or 2FF SIM card (15 x 25 mm) from the preferred carrier.
2. Have the Access Point Name (APN) information from the preferred carrier, if different from the default APN settings. The APN information varies from carrier to carrier.

> **Note**
>
> See the 210-SE LTE QSG for more information about installing the appliance.

Add and Configure 210-SE LTE Appliance in the Network

1. Insert the SIM card into the SIM card slot of the appliance. Only a standard or 2FF SIM card (15x25 mm) is supported.
2. In the SD-WAN appliance GUI, navigate to **Configuration** > **Appliance Settings** > **Network Adapters** > **Mobile Broadband**.

   a. If the APN settings are not the default, type the APN settings details.

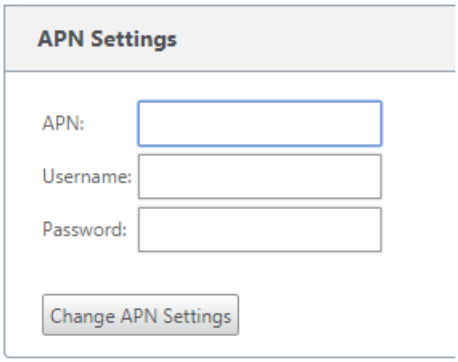

   b. Verify the status info page.

- Status: Enabled means modem tries to establish the data session.
- Card state: If present, it indicates that SIM is properly inserted.
- Signal strength: Excellent, Good, Fair, Poor, or no signal.
- Home network: Carrier of the SIM you inserted.
- APN name: APN used by the LTE modem.
- Session state: Connected indicates that the device has joined the network.

  If the session state is disconnected, check with carrier whether:

  - The account has been activated.

  - The data plan is enabled or not.

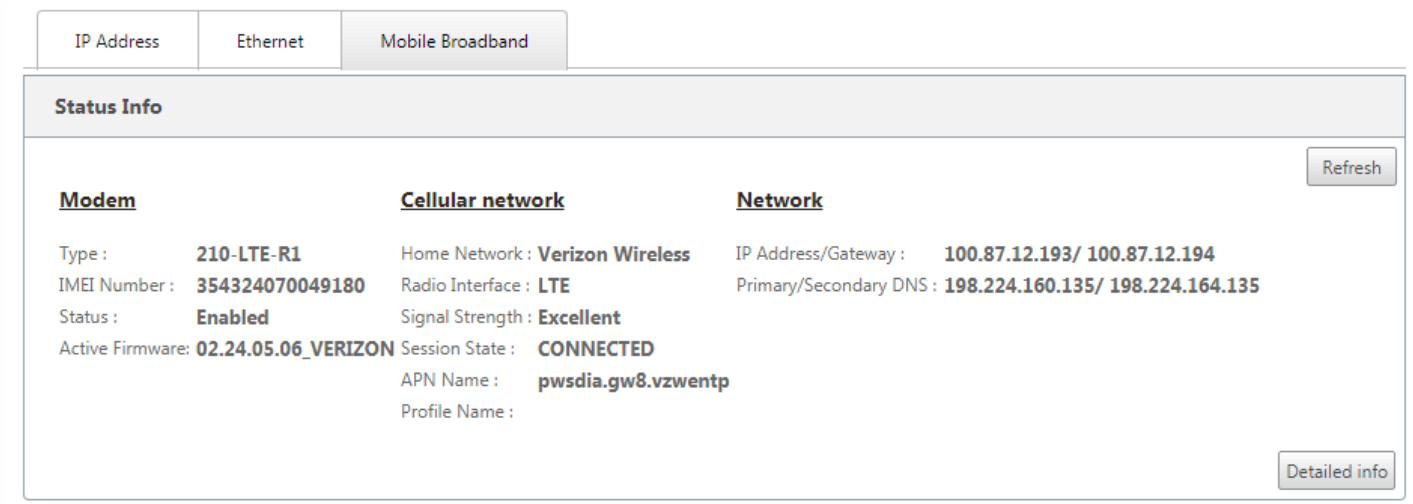

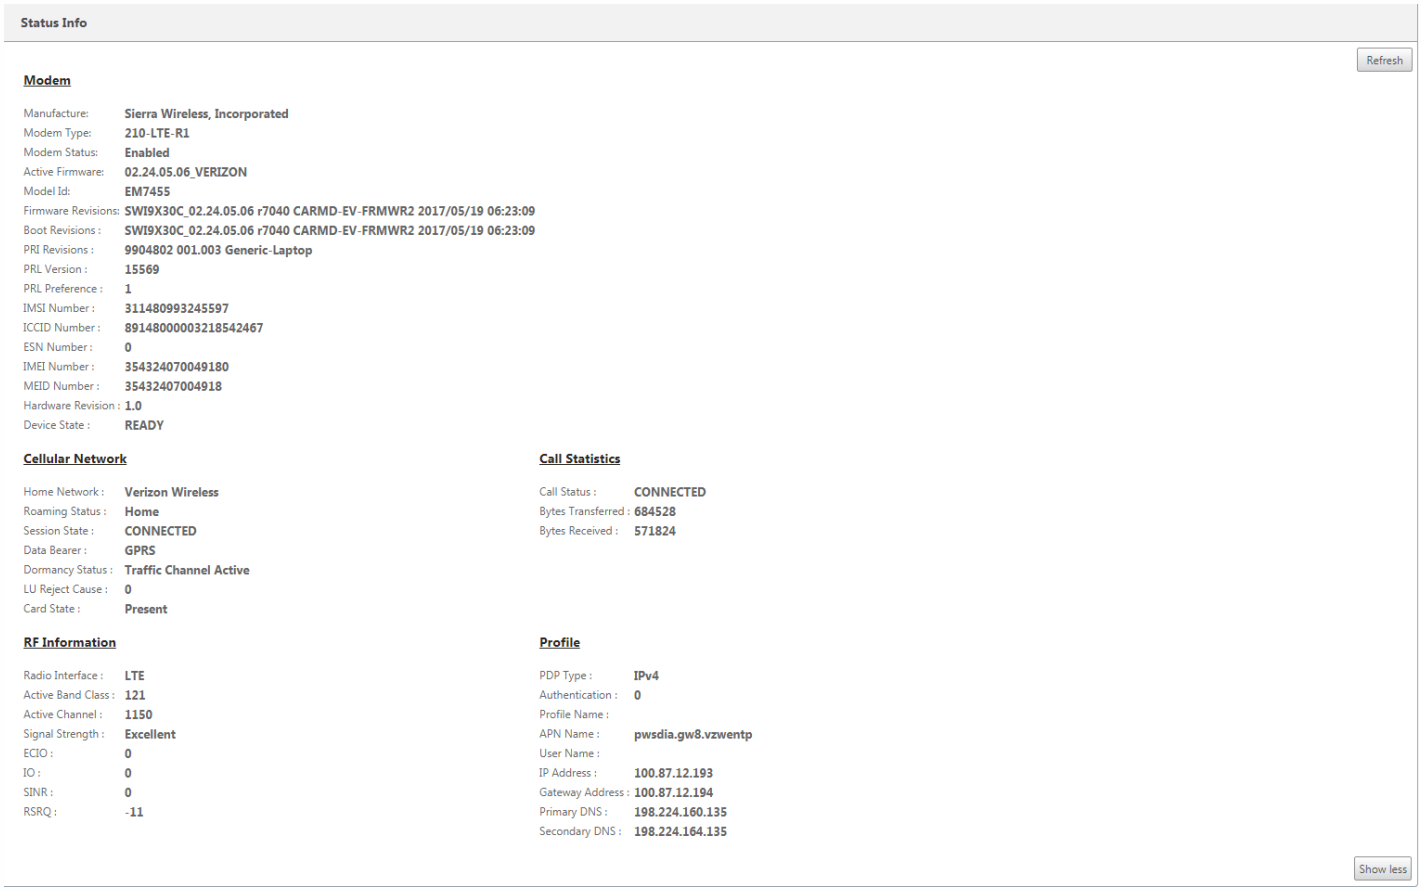

Manage Firmware

The AUTO-SIM option allows the LTE modem to choose the most matching firmware based on the SIM card inserted. Choose **AUTO-SIM,** if you are unsure which firmware to use. You can upload new firmware if available, using the **Upload** option.

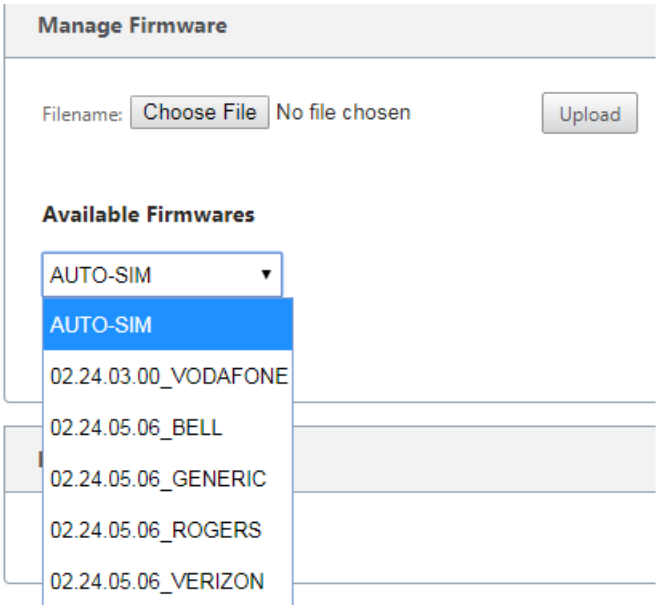

Enable/Disable Modem

Enable/disable modem depending on your intent to use the LTE functionality. By default, the LTE modem is enabled.

Reboot Modem

Reboots the modem. It can take up to 3-5 minutes for the reboot operation to complete.

Refresh SIM

Use this option when you hot swap the SIM card to detect the new SIM card by the 210-SE LTE modem.

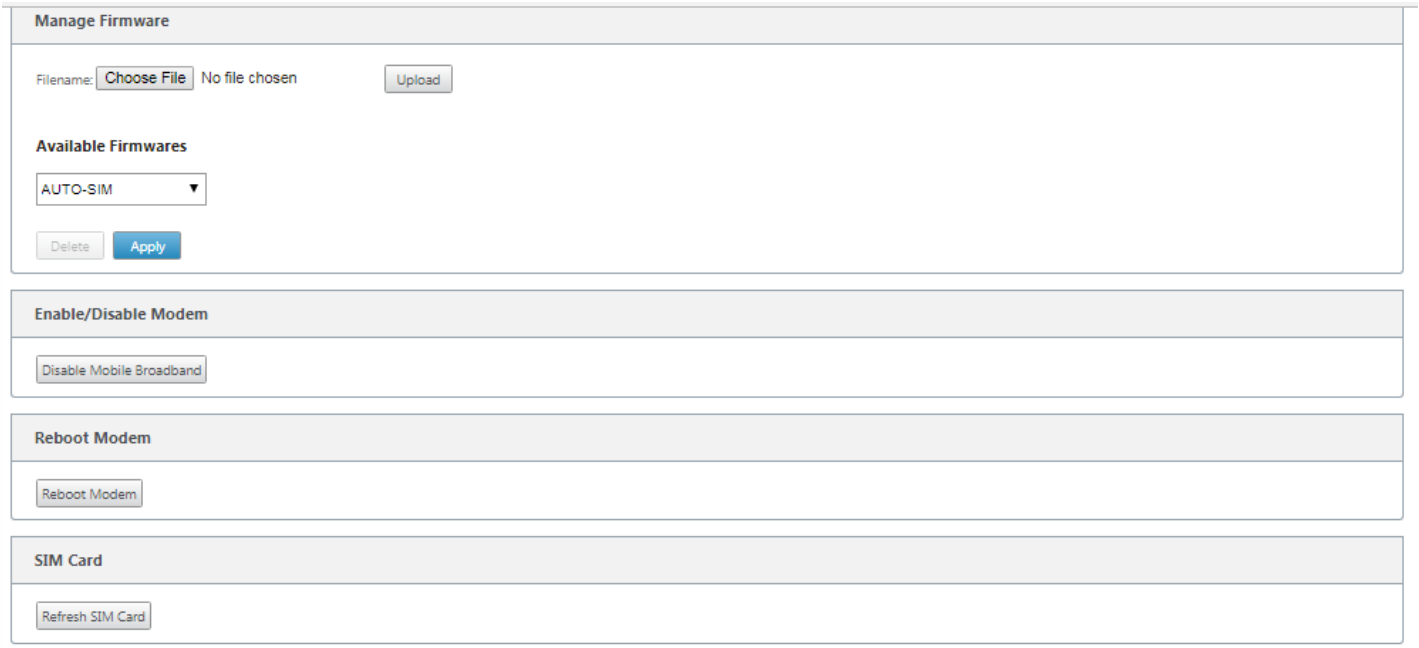

Configure 210-SE LTE Using CLI

To configure 210-SE LTE modem using the CLI:

1. Log in to the SD-WAN appliance console.
2. At the prompt, type the user name and password to gain CLI interface access.
3. At the prompt, type the command lte. Type >help. This displays the list of LTE commands available for configuration.

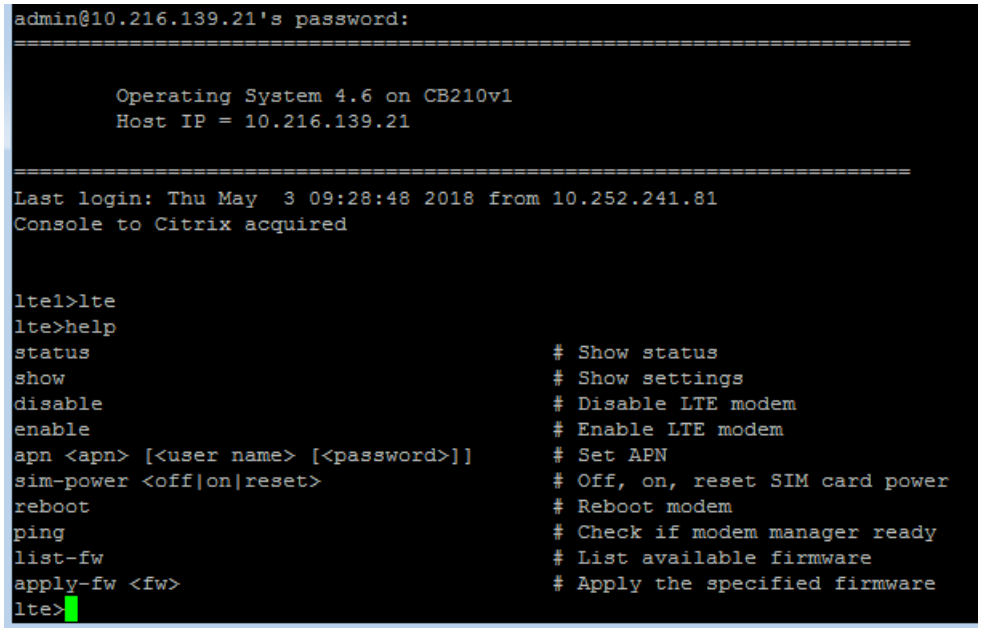

The following table lists the LTE command descriptions.

Command	Description
Help {lte>help}	Lists the available LTE commands and parameters
Status {lte>status}	Displays LTE connectivity status
Show {lte>show}	Displays LTE settings
Disable {lte>disable}	Disables LTE modem
Enable {lte>enable}	Enables LTE modem
Apn {lte>apn}	Configures APN settings information
Sim-power <off | on | reset> {lte>sim-power}	Powers off sim card, Power on sim card, Refresh sim card
Reboot {lte>reboot}	Restarts LTE modem
Ping {lte>ping}	Pings LTE modem
List-fw {lte>list-fw}	Lists firmware available on the R1 or R2 LTE modems
Apply-fw <fw> {lte>apply-fw}	Applies firmware specific to a carrier

Configure MCN for LTE

To configure an MCN:

1. Log in to the SD-WAN appliance GUI. Go to **Configuration Editor**. Complete configuration for the MCN site, see; Configure MCN.

2. Ensure that you provide routable public IP address as part of WAN link configuration.

   You do not have to configure public IP address for client appliances.

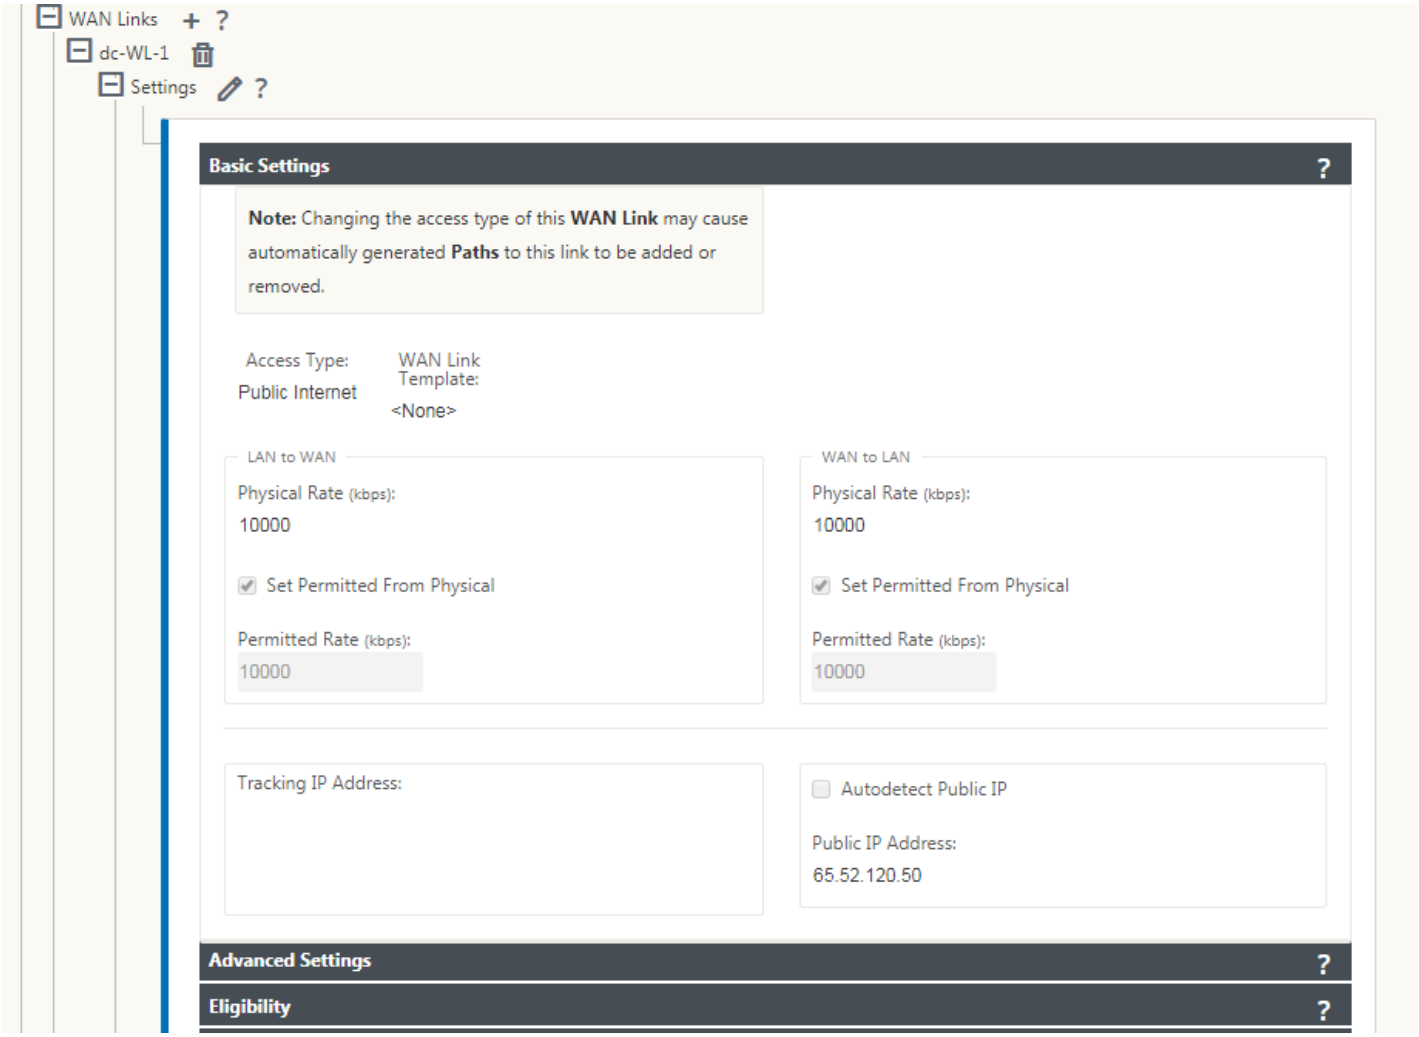

Configure Branch for LTE

To configure the 210-SE LTE appliance as a branch site:

1. In the SD-WAN appliance GUI, go to configuration editor. See; Configure Branch.

   a. Create Interface Groups.

   b. Create up to one Virtual Interface and one Interface Group for the LTE adapter to configure WAN link by selecting the following:

      i. Ethernet Interface – LTE 1

      ii. Security- untrusted (default)

iii. DHCP client - Enabled (default)

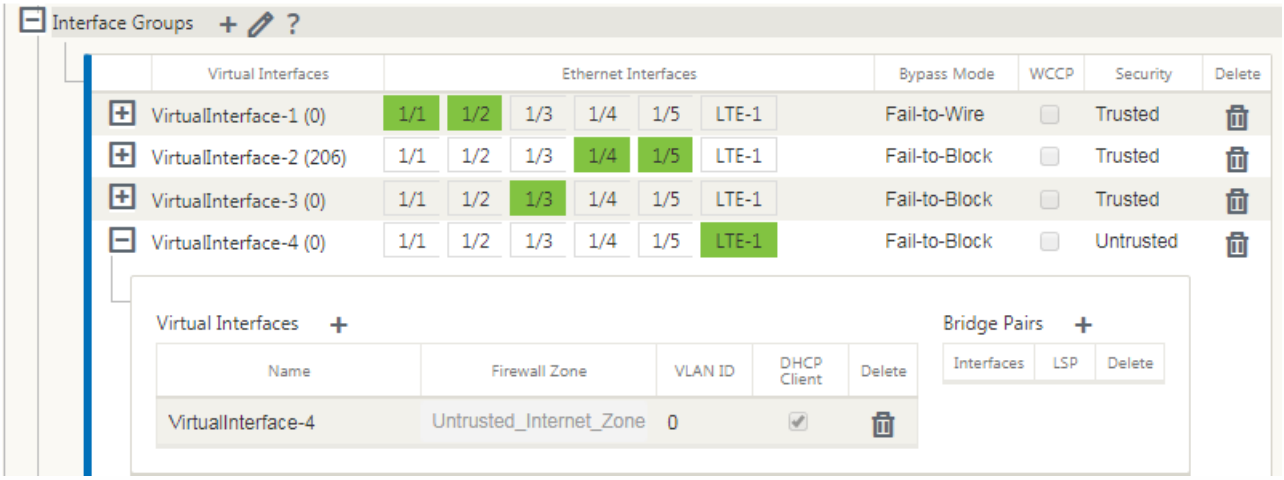

2. Enable **AutoDetect Public IP** for WAN link configuration when configuring WAN link using the virtual interface created for LTE interface.

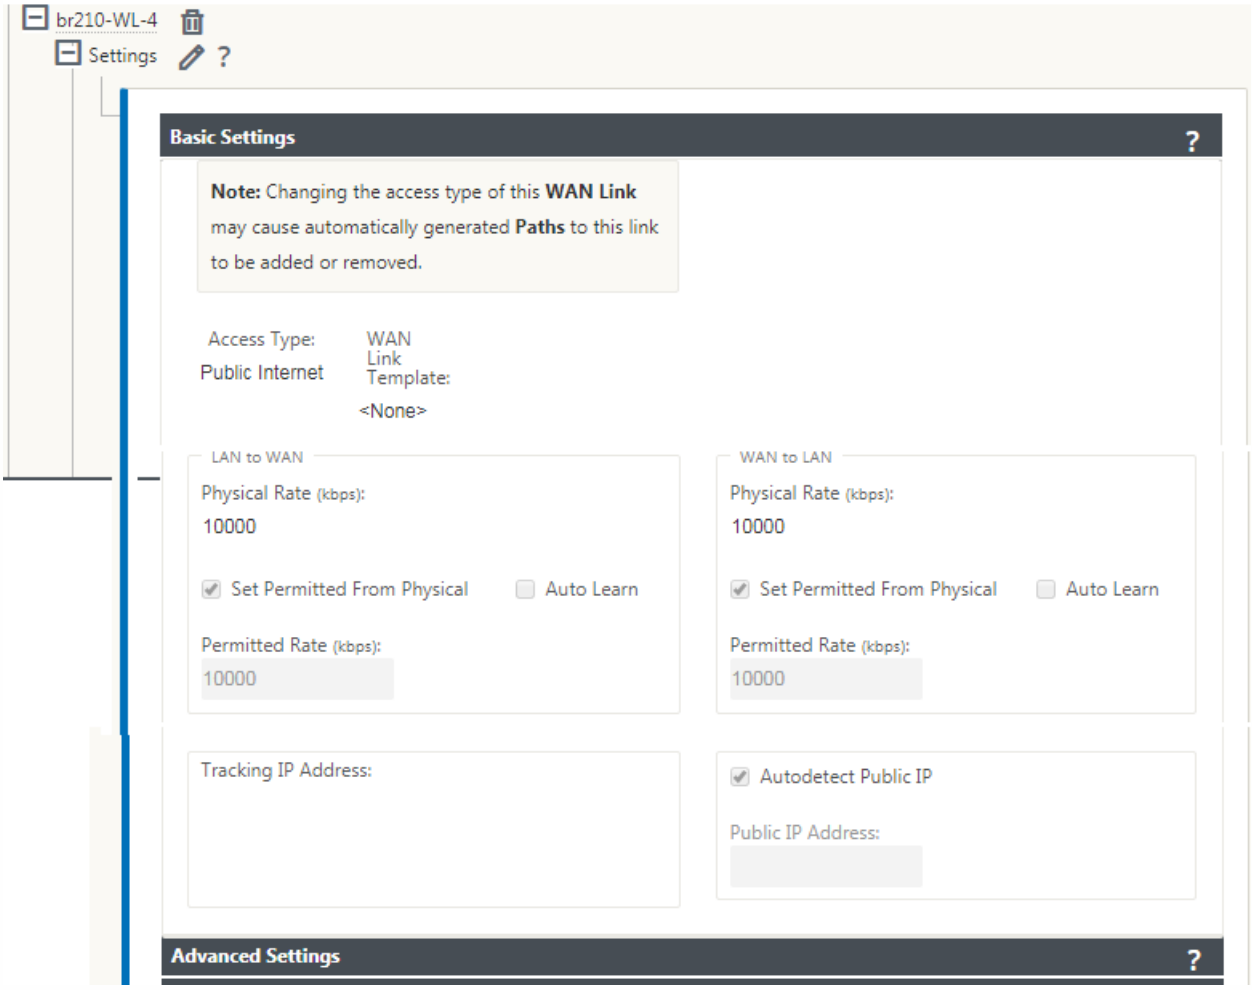

3. By default, when you try to configure WAN link using LTE interface, the WAN link is marked as Metered link and Last Resort Standby mode. You can change these default settings, if necessary.

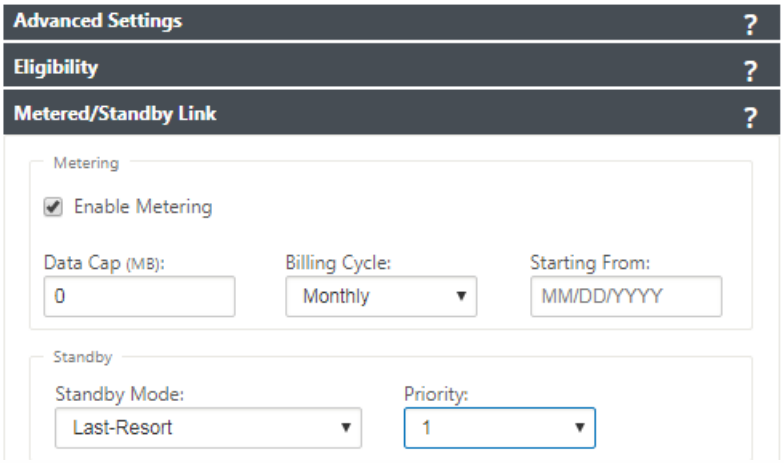

Note

It is recommended to configure the **active MTU detect** option for the WAN links created using the LTE interface. Establishing dynamic virtual paths using LTE interface is not supported.

The IP address and gateway address for the Access Interface of the WAN link need not be configured because it receives that information from the carrier through DHCP.

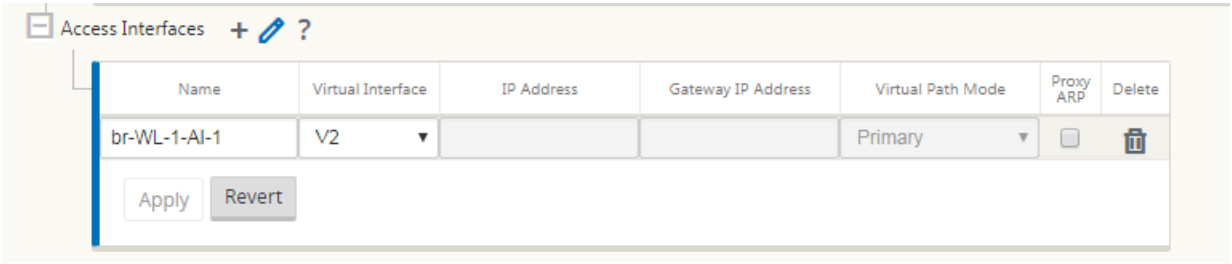

4. Complete rest of the required Branch configuration for the 210-SE LTE appliance. See; configure Branch.
5. Perform Change Management by uploading the SD-WAN software. See the Change Management procedure.
6. Activate configuration through the Local Change Management process. When you perform Change Management, configuration is activated and required configuration is applied.

ZTD over LTE

Pre-requisites for enabling ZTD service over LTE:

1. Install antenna and the SIM card for the 210-SE LTE appliance.
2. Ensure that the SIM card has an activated data plan.
3. Ensure that the management port is not connected.

   - If the management port is connected, disconnect the management port and then restart the appliance.

   - If a static IP address on the Management Interface is configured, you can configure the Management Interface with DHCP, apply the configuration, and then disconnect the Management port, and restart the appliance.

4. Ensure that the 210-SE appliance configuration has internet service defined for LTE interface.

When the appliance is powered on, the ZTD service uses the LTE port to obtain the latest SD-WAN software and SD-WAN configuration only when the management port is not connected.

You can use the SD-WAN Center GUI to deploy and configure 210-SE LTE appliance for the ZTD service.

See the ZTD procedure for more information about deploying and configuring 210-SE LTE appliance using SD-WAN Center.

ZTD Service over Management Interface for 210-SE LTE Appliance

Connect the Management Port and use the standard ZTD procedure that is supported on all other non-LTE platforms.

Manage 210-SE LTE Using SD-WAN Center

The **Mobile Broadband** page under the **Configuration** tab in the SD-WAN Center GUI displays information about the LTE modem configuration. You can perform the LTE modem operations using the same options or buttons displayed in the SD-WAN 210-SE appliance GUI configuration for mobile broadband settings. LTE summary table lists 210-SE LTE appliances and you can select either single or multiple appliances to perform LTE modem operations.

> ### Note
>
> Multi-region site support for uploading firmware is not available.

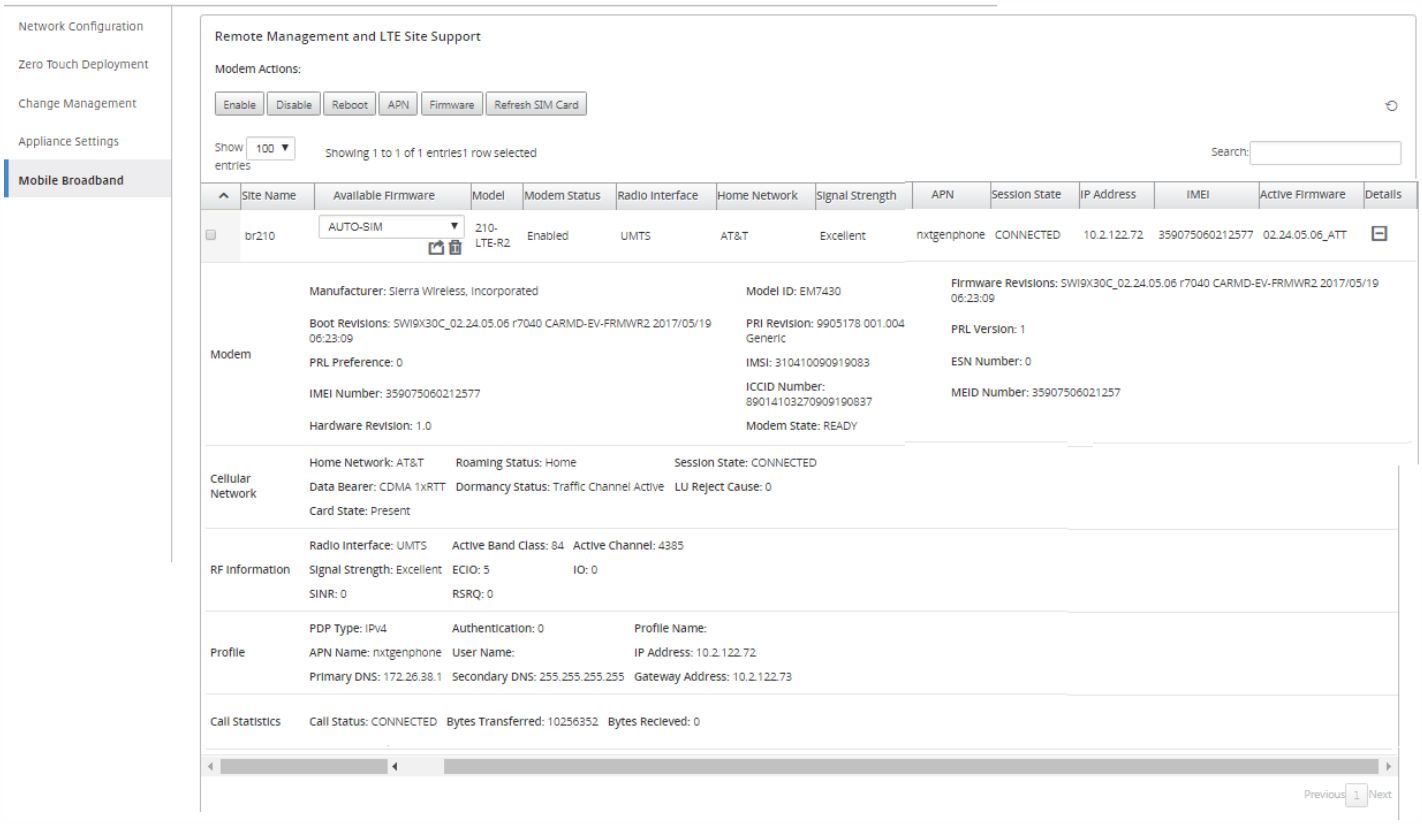

For information about LTE REST API, navigate to the SD-WAN GUI and go to **Configuration** > **Appliance Settings** >**NITRO API**. Click **Download Nitro API Doc**.

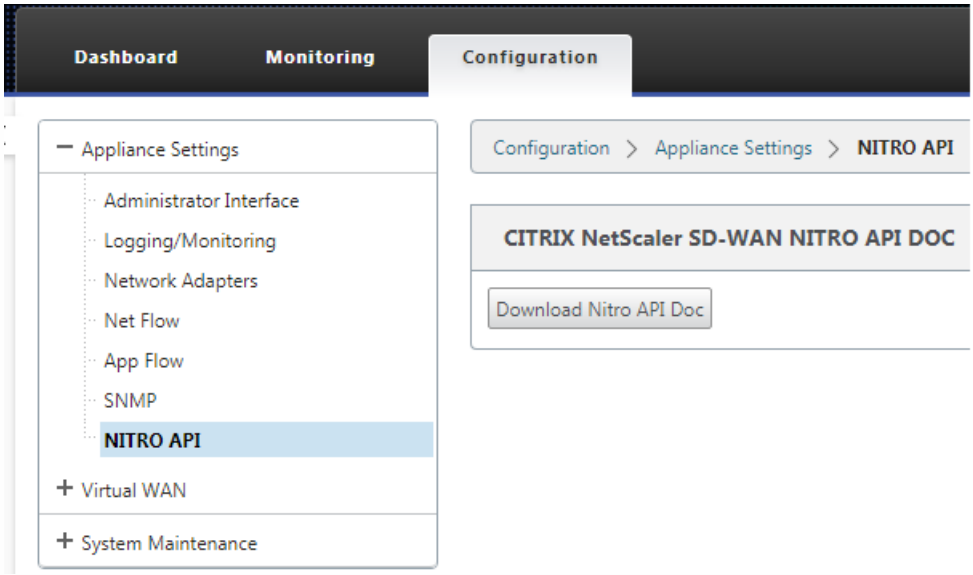

NetScaler SD-WAN WANOP 9.3

Aug 09, 2017

For information about NetScaler SD-WAN WANOP 9.3 installation, deployment, and feature configuration, please refer to the CloudBridge 7.4 documentation. The features and procedures for the NetScaler SD-WAN WANOP 9.3 are similar to the procedures documented in CloudBridge 7.4 release.

Important

Command Center End of Life Notification:

End of Life process for **Citrix Command Center** tool was initiated on 15-May-2017.

It is recommended that you migrate to the new management tool **NetScaler MAS** for your WAN Optimization deployments at the earliest.

Please refer to the articles below for the Command Center EOL calendar and associated details.

- CTX223806 - Notice of Status Change Announcement for Citrix Command Center Software Version 5.2
- CTX223786 - FAQ: Citrix Command Center - End Of Life

Citrix Command Center tool for NetScaler SD-WAN WANOP edition is supported only till the NetScaler SD-WAN 9.2 appliance software release.

Starting with the NetScaler SD-WAN 9.3 software release, NetScaler MAS will be the management tool for SD-WAN WANOP edition appliances.

NetScaler SD-WAN WANOPT Features

Features

Compression

XenApp/XenDesktop Acceleration

HTTP Acceleration

TCP Flow-Control Acceleration

Traffic Shaping

Traffic Classification

Link Definitions

- Secure Traffic Acceleration
- How HTML5 Works
- Video Caching
- Monitoring with AppFlow
- Internet Protocol Version 6 (IPv6) Acceleration
- SCPS Support
- Automatically Configuring CloudBridge Devices
- CloudBridge Connector
- WAN Insight
- Office 365 Acceleration

The topics that are organized under this topic are specific to NetScaler SD-WAN WANOP 9.3.

Configuring SSL Compression

Aug 09, 2017

The Netscaler SD-WAN WO SSL compression feature enables multisession compression of SSL connections (for example, HTTPS traffic), providing a compression ratios of up to 10,000:1. For more information, see SSL Compression.

For SSL compression to work, the SD-WAN WANOP appliance needs certificates from either the server or the client. To support multiple servers, multiple private keys can be installed on the appliance, one per SSL profile. Special SSL rules in the service class definitions match up servers to SSL profiles, and thus SSL profiles to private keys.

SSL compression works in split proxy or transparent proxy mode, you can choose the mode as per your requirement. For more information, see How SSL Compression Works.

> **Note**
>
> Transparent proxy mode is currently not supported.

In NetScaler SD-WAN WANOP 9.3, to enable secure access with SSL tunnel, the latest SSL protocol TLS 1.2 is used in SSL proxy. You can choose to use TLS1.2 protocol only or use TLS1.0, TLS1.1 and TLS1.2 protocols.

> **Note**
>
> SSL protocols SSL v3 and SSL v2 are no longer supported.

To configure SSL compression:

1. Acquire copies of your server's CA certificate and private certificate-key pair and install them on the server-side appliance. These credentials are likely to be application-specific. That is, a server might have different credentials for an Apache Web server than for an Exchange Server running RPC over HTTPS.
2. You can choose to create a split proxy SSL Profile or a Transparent proxy SSL profile.

   For information on configuring split proxy SSL profile, see **Configuring a Split Proxy SSL Profile** section below.

   For information on configuring transparent proxy SSL profile, see **Configuring Transparent Proxy SSL Profile** section below.

   > **Note**
   >
   > Transparent proxy SSL profile is currently not supported.

3. Attach the SSL profile to a service class on the server-side appliance. This can be done by either creating a new service class based on the server IP, or by modifying an existing service class.

For more information see, **Creating or Modifying the Service Class** section below.

4. Set service classes on the client-side appliance. SSL traffic is not compressed unless it falls into a service class, on the client-side appliance, that enables acceleration and compression. This can be an ordinary service-class rule, not an SSL rule (only the server-side appliance needs SSL rules), but it must enable acceleration and compression. The traffic falls into an existing service class, such as “HTTPS” or “Other TCP Traffic.” If this class’s policy enables acceleration and compression, no additional configuration is needed.

5. Verify operation of the rule. Send traffic that should receive SSL acceleration through the appliances. On the server-side appliance, on the Monitoring: Optimization: Connections: Accelerated Connections tab, the Service Class column should match the service class you set up for secure acceleration, and the SSL Proxy column should list True for appropriate connections.

Configuring a Split Proxy SSL Profile

To configure a split proxy SSL profile:

1. In the server-side Netscaler SD-WAN WO appliance, navigate to **Configuration** > **Secure Acceleration** > **SSL Profile** and click **Add Profile**.

> ### Note
>
> You can either manually add an SSL profile or import one that is stored on your local computer.

2. In the **Profile Name** field, enter a name for the SSL profile and select **Profile Enabled**

3. If your SSL server uses more than one virtual host name, In the **Virtual Host Name** field, enter the target virtual host name. This is the host name listed in the server credentials.

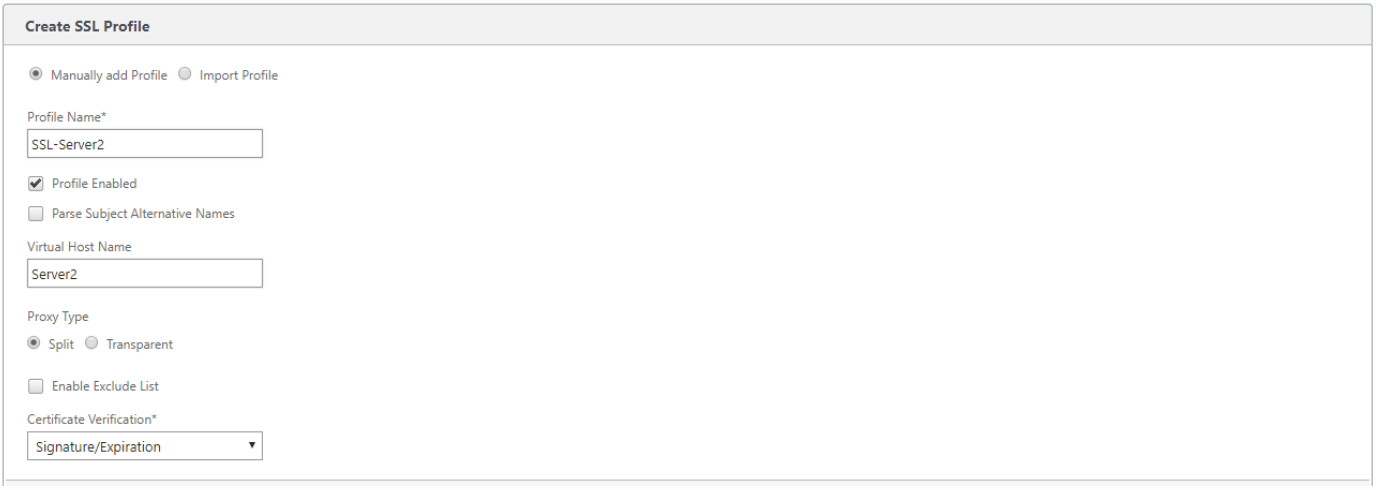

> ### Note
>
> To support multiple virtual hosts, create a separate SSL profile for each host name.

4. Choose **Split** proxy type.

5. In the **Certificate Verification** field, retain the default value (Signature/Expiration) unless your policies dictate otherwise.
6. Perform server-side proxy configuration:
    1. In the **Verification Store** field, select an existing server Certificate Authority (CA), or click **+** to upload a server CA.
    2. Choose **Authentication Required** and in the **Certificate/Private Key** field select a certificate key pair, or click **+** to upload a certificate key pair.
    3. In the **Protocol Version** field, select the protocols your server accepts.

    > **Note**
    >
    > NetScaler SD-WAN WO supports a combination of **TLS1.0, TLS1.1 or TLS1.2**, or **TLS1.2** only. SSL protocols SSLv3 and SSLv2 are not supported.

    4. If necessary, edit the **Cipher Specification** string, using the OpenSSL syntax.
    5. If required, select the type of renegotiation from the **Renegotiation Type** drop-down list to allow client-side SSL session renegotiation.

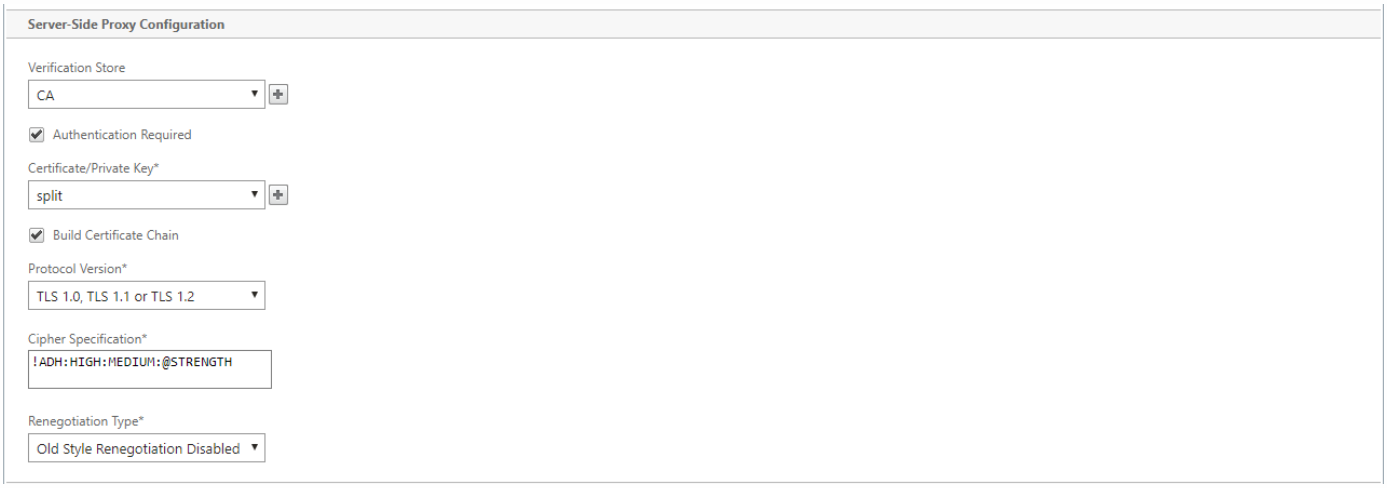

7. Perform client-side proxy configuration:
    1. In the **Certificate/Private Key** field, retain the default value.
    2. Choose **Build Certificate Chain** to allow the server-side appliance to build the SSL certificate chain.
    3. If required, select or upload a CA store to use as the Certificate Chain Store.
    4. In the **Protocol Version** field, select the protocol versions you want to support on the client side.

    > **Note**
    >
    > NetScaler SD-WAN WO supports a combination of **TLS1.0, TLS1.1 or TLS1.2**, or **TLS1.2** only. SSL protocols SSLv3 and SSLv2 are not supported.

5. If necessary, edit the client-side Cipher Specification.
6. If required, select the type of renegotiation from the **Renegotiation Type** drop-down list to allow client-side SSL session renegotiation.

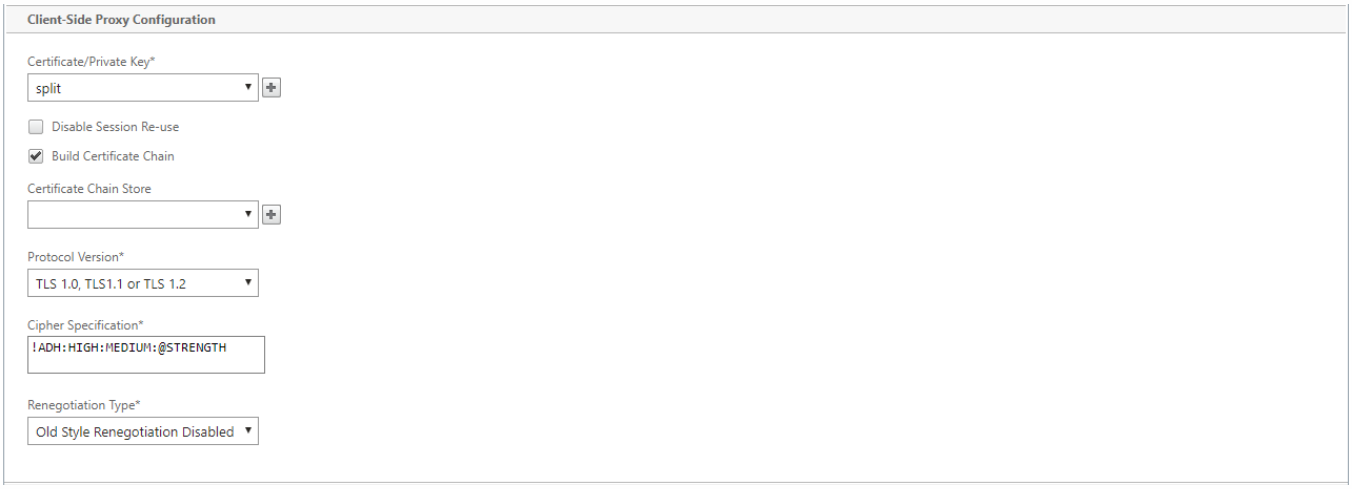

8. Click **Create**.

Configuring Transparent Proxy SSL Profile

To configure a transparent proxy SSL profile:

1. In the server-side Netscaler SD-WAN WO appliance, navigate to **Configuration** > **Secure Acceleration** > **SSL Profile** and click **Add Profile**.

> ### Note
>
> You can either manually add an SSL profile or import one that is stored on your local computer.

2. In the **Profile Name** field, enter a name for the SSL profile and select **Profile Enabled**.
3. If your SSL server uses more than one virtual host name, In the **Virtual Host Name** field, enter the target virtual host name. This is the host name listed in the server credentials.

> ### Note
>
> To support multiple virtual hosts, create a separate SSL profile for each host name.

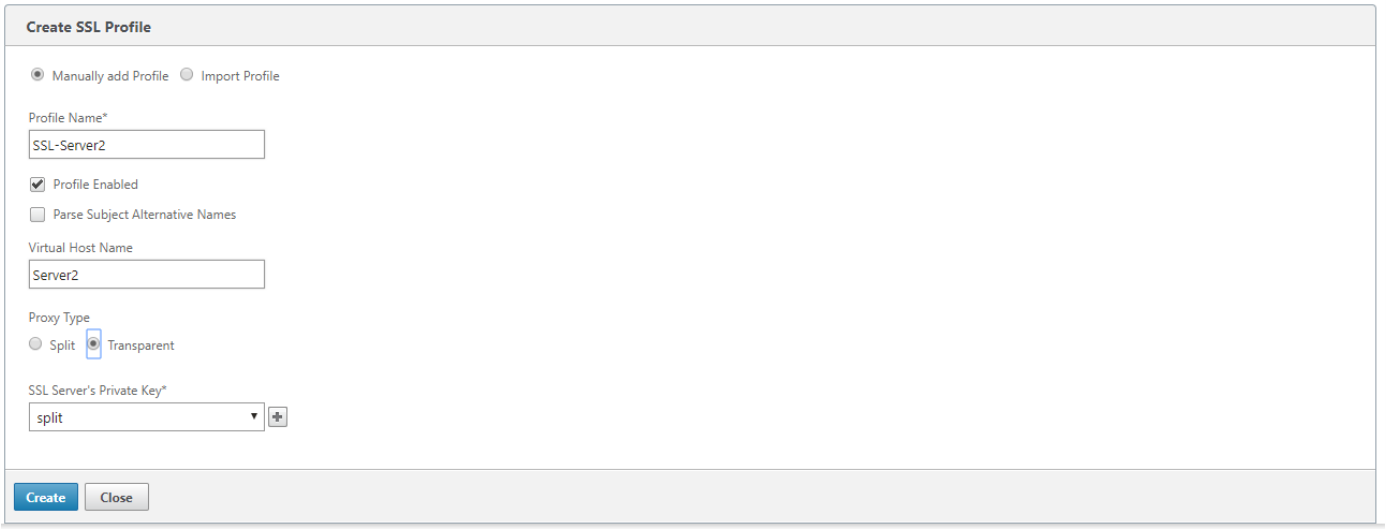

4. Choose **Transparent** proxy type.
5. In the **SSL Server's Private Key** field, select the server's private key from the drop-down menu, or click **+** to upload a new private key.
6. Click **Create**.

Creating or Modifying the Service Class

To create or modify the service class and attach the SSL Profile:

1. In the Netscaler SD-WAN WO appliance web interface, navigate to **Configuration > Optimization Rules > Service Classes** and click **Add**. To edit an existing service class, select the appropriate service class and click **Edit**.
2. In the Name field, enter a name for the new service class (for example, "Accelerated HTTPS").
3. Enable compression by setting the Acceleration Policy to **Disk**, **Memory** or **Flow Control**.
4. In the **Filter Rules** section, click **Add**.
5. In the **Destination IP Address field**, type the server's IP address (for example, 172.16.0.1 or, equivalently, 172.16.0.1/32.
6. In the **Direction** field, set the rule to Unidirectional. SSL profiles are disabled if Bidirectional is specified.
7. In the **SSL Profiles** section, select the SSL profile that you created and move it to the **Configured** section.
8. Click **Create** to create the rule.
9. Click **Create** to create the service class.

Updated CLI Command

NetScaler SD-WAN WO 9.3 supports the latest TLS1.2 SSL protocol. You can choose to use TLS1.2 protocol only or any version of TLS protocols. SSL protocols SSL v3 and SSL v2, and transparent proxy SSL profiles are not supported. The ***add ssl-profile*** and ***set ssl-profile*** CLI commands are updated to reflect these changes.

add ssl-profile

-name "profile-name"

[-state {enable, disable}]

-proxy-type split

[-virtual-hostname "hostname"]

-cert-key “cert-key-pair-name”

[-build-cert-chain {enable, disable}]

[-cert-chain-store {use-all-configured-CA-stores, “store-name”}]

[-cert-verification {none, Signature/Expiration, Signature/Expiration/

Common-Name-White-List, Signature/Expiration/Common-Name-Black-List}]

[-verification-store {use-all-configured-CA-stores, “store-name”}]

[-server-side-protocol { TLS-1.2, TLS-version-any}]

[-server-side-ciphers “ciphers”]

[-server-side-authentication {enable, disable}]

[-server-side-cert-key “cert-key-pair-name”]

[-server-side-build-cert-chain {enable, disable}]

[-server-side-renegotiation {disable-old-style, enable-old-style, new-style,

compatible}]

[-client-side-protocol-version { TLS-1.2, TLS-version-any}]

[-client-side-ciphers “ciphers”]

[-client-side-renegotiation {disable-old-style, enable-old-style, new-style,

compatible}]

set ssl-profile

-name “profile-name” [-state {enable, disable}]

[-proxy-type split]

[-virtual-hostname “hostname”]

[-cert-key “cert-key-pair-name”]

[-build-cert-chain {enable, disable}]

[-cert-chain-store {use-all-configured-CA-stores, “store-name”}]

[-cert-verification {none, Signature/Expiration, Signature/Expiration/

Common-Name-White-List, Signature/Expiration/Common-Name-Black-List}]

[-verification-store {use-all-configured-CA-stores, “store-name”}]

[-server-side-protocol {TLS-1.2, TLS-version-any}]

[-server-side-ciphers “ciphers”]

[-server-side-authentication {enable, disable}]

[-server-side-cert-key “cert-key-pair-name”]

[-server-side-build-cert-chain {enable, disable}]

[-server-side-renegotiation {disable-old-style, enable-old-style, new-style,

compatible}]

[-client-side-protocol-version {TLS-1.2, TLS-version-any}]

[-client-side-ciphers “ciphers”]

[-client-side-renegotiation {disable-old-style, enable-old-style, new-style,

compatible}]

The rest of the SSL Configuration commands remain unchanged. For more information see, SSL Configuration.

XenServer 6.5 Upgrade for SD-WAN WANOP Edition Appliances

Aug 14, 2017

Important

To upgrade to XenServer version 6.5, the appliances must be running NetScaler SD-WAN WANOP software release 9.0.x or later.

Note

Do not attempt upgrading when the appliance is running on software version lower than release 9.0.x to prevent upgrade issues.

How to Upgrade to Citrix XenServer 6.5

To upgrade to XenServer 6.5 on SD-WAN WANOP appliances, ensure that the appliance is running software release version 9.0.x or later. If the appliances are running older software release version, upgrade to the latest software release version first.

1. In NetScaler SD-WAN WANOP GUI, go to C**onfiguration** > **Maintenance** > **Update Software**. Download the *ns-sdw-vw-<Build_No>.upg* file to upgrade the appliance.

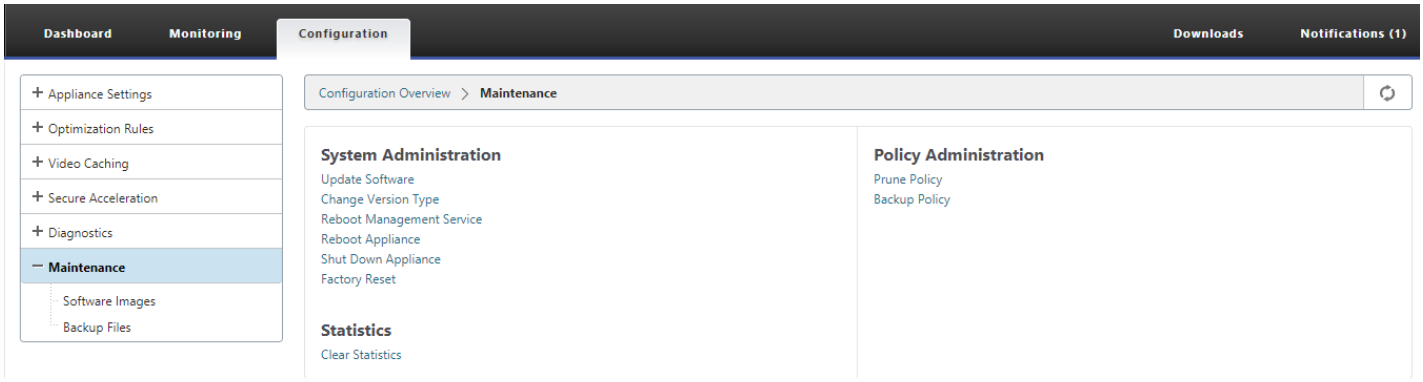

2. After upgrading to the latest software version of WANOP software, navigate to **Configuration > Maintenance > Update Software** in the GUI. Upload *ns-sdw-xen65-pkg_v1.5.upg* file.

3. Wait for approximately 20 mins for the upgrade to complete. The appliance restarts after the upgrade is successfully completed.

The WANOP Client Plug-in

Aug 09, 2017

The WANOP Client Plug-in is a software based network accelerator that runs on Windows laptops and workstations, providing acceleration anywhere, not just at offices with WANOP Client Plug-in appliances. It connects to a Citrix WANOP Client Plug-in appliance at the other end of the link.

The principles of WANOP Client Plug-in operation are generally the same as those of a WANOP Client Plug-in appliance. For topics not included in the plug-in documentation, see the larger documentation set.

The plug-in is distributed as a standard Microsoft installation file (MSI). Plug-in deployment requires some plug-in specific configuration of the WANOP Client Plug-in appliances at the other ends of the links. If you customize the MSI file with the DNS or IP addresses of the WANOP Client Plug-in appliances, and a few other parameters, your users do not have to enter any configuration information when installing the plug-in on their Windows computers.

Figure 1. Typical WANOP Client Plug-in Network Showing the WANOP Client Plug-in

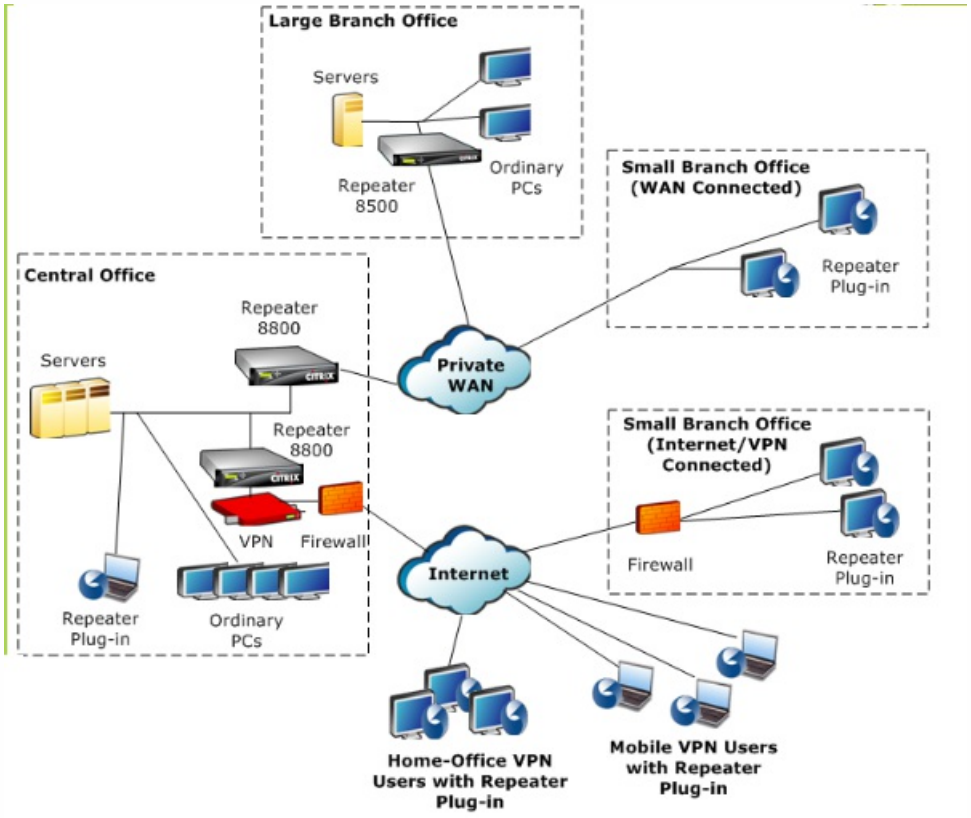

Note

The plug-in is supported by Citrix Receiver 1.2 or later, and can be distributed and managed by Citrix Receiver.

Hardware and Software Requirements

Aug 09, 2017

On the client side of the accelerated link, the WANOP Client Plug-in is supported on Windows desktop and laptop systems, but not on netbooks or thin clients. Citrix recommends the following minimum hardware specifications for the computer running the WANOP Client Plug-in:

- Pentium 4-class CPU
- 2 GB of RAM
- 2 GB of free disk space

WANOP Client Plug-in is supported on Windows 10 platform and needs following system requirements:

- 4GB RAM
- 10GB free disk space

The WANOP Client Plug-in is supported on the following operating systems:

- Windows XP Home
- Windows XP Professional
- Windows Vista (all 32-bit versions of Home Basic, Home Premium, Business, Enterprise, and Ultimate)
- Windows 7 (all 32-bit and 64-bit versions of Home Basic, Home Premium, Professional, Enterprise, and Ultimate)
- Windows 8 (32-bit and 64-bit versions of Enterprise Edition)
- Windows 10 (32-bit and 64-bit versions of Enterprise Edition)

On the server side, the following appliances currently support WANOP Client Plug-in deployments:

- Repeater 8500 Series
- Repeater 8800 Series
- WANOP Client Plug-in VPX
- WANOP Client Plug-in 2000
- WANOP Client Plug-in 3000
- WANOP Client Plug-in 4000
- WANOP Client Plug-in 5000

How the WANOP Plug-in Works

Aug 09, 2017

WANOP Client Plug-in products use your existing WAN/VPN infrastructure. A computer on which the plug-in is installed continues to access the LAN, WAN, and Internet as it did before installation of the plug-in. No changes are required to your routing tables, network settings, client applications, or server applications.

Citrix Access Gateway VPNs require a small amount of WANOP Client Plug-in-specific configuration.

There are two variations on the way connections are handled by the plug-in and appliance: *transparent mode* and *redirector mode.* Redirector is a legacy mode that is not recommended for new deployments.

- **Transparent mode** for plug-in-to-appliance acceleration is very similar to appliance-to-appliance acceleration. The WANOP Client Plug-in appliance must be in the path taken by the packets when traveling between the plug-in and the server. As with appliance-to-appliance acceleration, transparent mode operates as a transparent proxy, preserving the source and destination IP address and port numbers from one end of the connection to the other.
- **Redirector mode** (not recommended) uses an explicit proxy. The plug-in readdresses outgoing packets to the appliance's redirector IP address. The appliance in turn readdresses the packets to the server, while changing the return address to point to itself instead of the plug-in. In this mode, the appliance does not have to be physically inline with the path between the WAN interface and the server (though this is the ideal deployment).

Best Practice: Use transparent mode when you can, and redirector mode when you must.

Transparent Mode

Aug 09, 2017

In transparent mode, the packets for accelerated connections must pass through the target appliance, much as they do in appliance-to-appliance acceleration.

The plug-in is configured with a list of appliances available for acceleration. It attempts to contact each appliance, opening a signaling connection. If the signaling connection is successful, the plug-in downloads the acceleration rules from the appliance, which sends the destination addresses for connections that the appliance can accelerate.

Figure 1. Transparent Mode, Highlighting Three Acceleration Paths

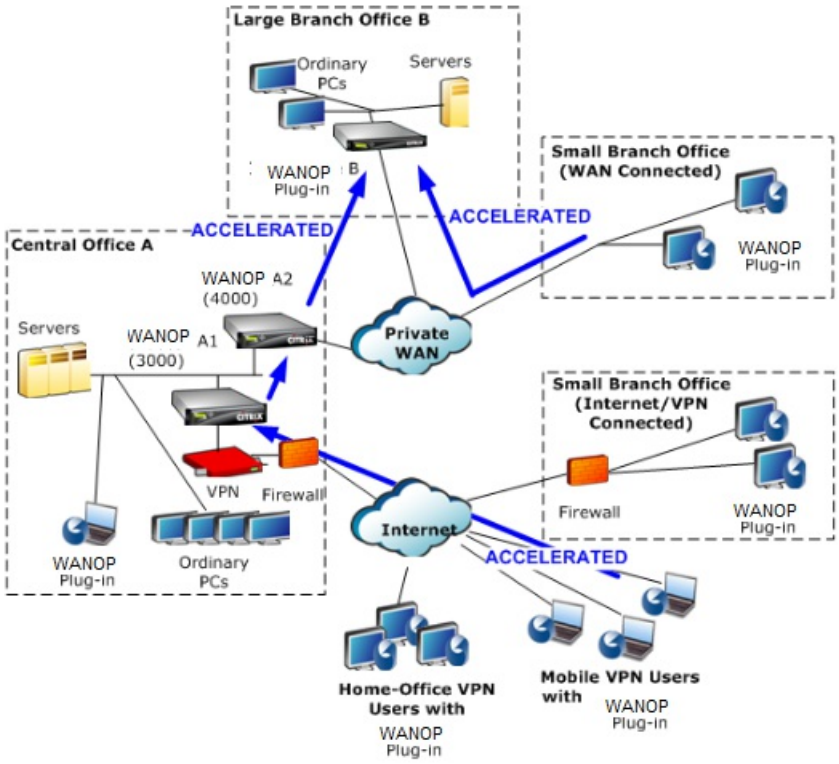

Note

- Traffic flow--Transparent mode accelerates connections between a WANOP Client Plug-in and a plug-in-enabled appliance.
- Licensing--Appliances need a license to support the desired number of plug-ins. In the diagram, Repeater A2 does not need to be licensed for plug-in acceleration, because Repeater A1 provides the plug-in acceleration for site A.
- Daisy-chaining--If the connection passes through multiple appliances on the way to the target appliance, the appliances in the middle must have "daisy-chaining" enabled, or acceleration is blocked. In the diagram, traffic from home-office and mobile VPN users that is destined for Large Branch Office B is accelerated by Repeater B. For this to work, Repeaters A1 and A2 must have daisy-chaining enabled.

Whenever the plug-in opens a new connection, it consults the acceleration rules. If the destination address matches any of

the rules, the plug-in attempts to accelerate the connection by attaching acceleration options to the initial packet in the connection (the SYN packet). If any appliance known to the plug-in attaches acceleration options to the SYN-ACK response packet, an accelerated connection is established with that appliance.

The application and server are unaware that the accelerated connection has been established. Only the plug-in software and the appliance know that acceleration is taking place.

Transparent mode resembles appliance-to-appliance acceleration but is not identical to it. The differences are:

- Client-initiated connections only--Transparent mode accepts connections initiated by the plug-in-equipped system only. If you use a plug-in-equipped system as a server, server connections are not accelerated. Appliance-to-appliance acceleration, on the other hand, works regardless of which side is the client and which is the server. (Active-mode FTP is treated as a special case, because the connection initiating the data transfer requested by the plug-in is opened by the server.)
- Signaling connection--Transparent mode uses a signaling connection between the plug-in and appliance for the transmission of status information. Appliance-to-appliance acceleration does not require a signaling connection, except for secure peer relationships, which are disabled by default. If the plug-in cannot open a signaling connection, it does not attempt to accelerate connections through the appliance.
- Daisy-chaining--For an appliance that is in the path between a plug-in and its selected target appliance, you must enable daisy-chaining on the **Configuration: Tuning** menu.

Transparent mode is often used with VPNs. The WANOP Client Plug-in Plug-in is compatible with most IPSec and PPTP VPNs, and with Citrix Access Gateway VPNs.

The following figure shows packet flow in transparent mode. This packet flow is almost identical to appliance-to-appliance acceleration, except that the decision of whether or not to attempt to accelerate the connection is based on acceleration rules downloaded over the signaling connection.

Figure 2. Packet flow in transparent mode

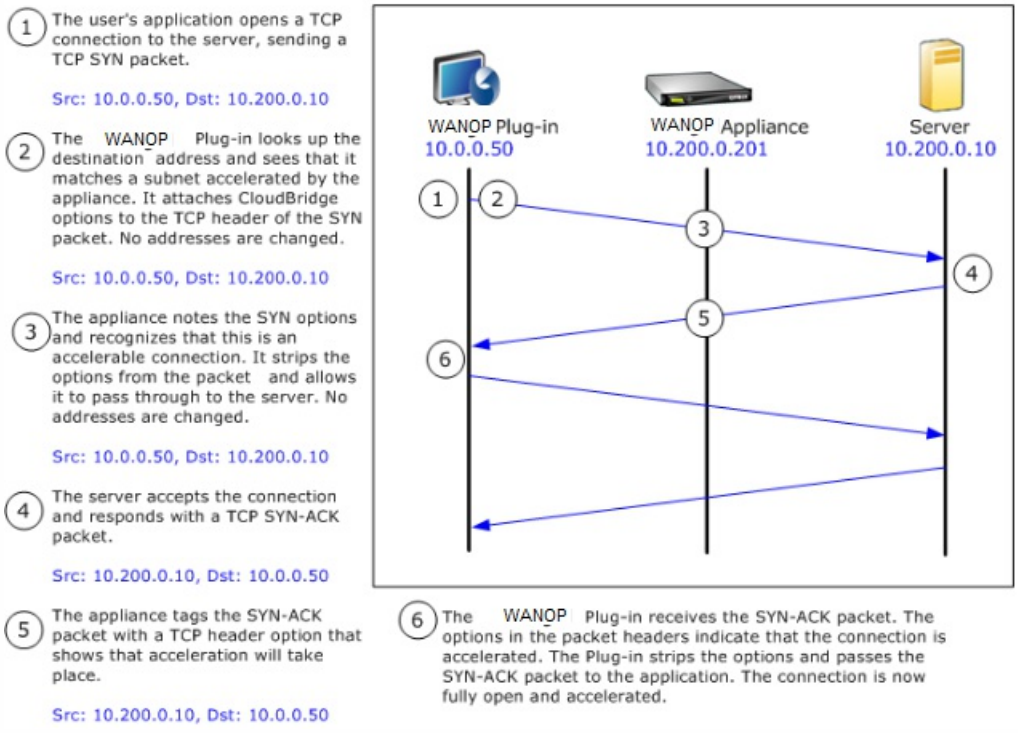

Redirector Mode

Aug 09, 2017

Redirector mode works differently from transparent mode in the following ways:

- The WANOP Client Plug-in Plug-in software redirects the packets by addressing them explicitly to the appliance.
- Therefore, the redirector-mode appliance does not have to intercept all of the WAN-link traffic. Because accelerated connections are addressed to it directly, it can be placed anywhere, as long as it can be reached by both the plug-in and the server.
- The appliance performs its optimizations, then redirects the output packets to the server, replacing the source IP address in the packets with its own address. From the server's point of view, the connection originates at the appliance.
- Return traffic from the server is addressed to the appliance, which performs optimizations in the return direction and forwards the output packets to the plug-in.
- The destination port numbers are not changed, so network monitoring applications can still classify the traffic.

The below figure shows how the Redirector mode works.

Figure 1. Redirector Mode

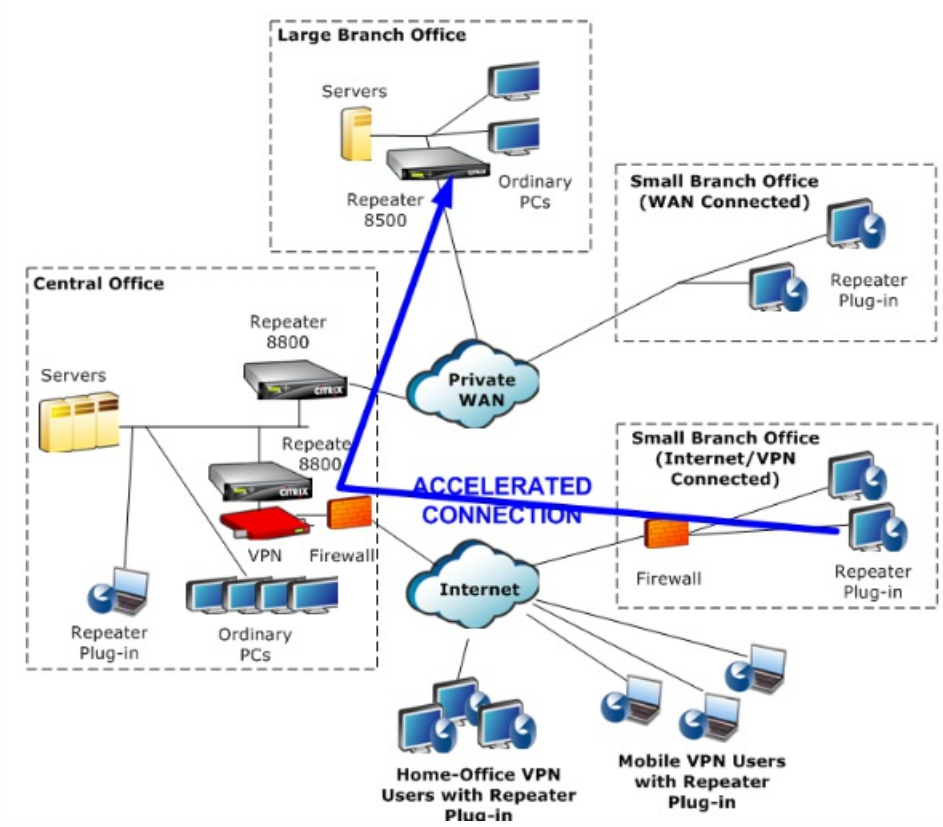

The below figure shows the packet flow and address mapping in *redirector mode*.

Figure 2. Packet Flow in Redirector Mode

1. The user's application opens a TCP connection to the server, sending a TCP SYN packet.

   Src: 10.0.0.50, Dst: 10.200.0.10

2. The Repeater Plug-in looks up the dst address and decides to redirect the connection to the appliance at 10.200.0.201.

   Src: 10.0.0.50, Dst: 10.200.0.201

   (10.200.0.10 is preserved in a TCP option field. Options 24-31 are used for various parameters.)

3. The appliance accepts the connection and forwards the packet to the server (using the dst address from the TCP options field), and giving itself as the src.

   Src: 10.200.0.201, Dst: 10.200.0.10

4. The server accepts the connection and responds with a TCP SYN-ACK packet.

   Src: 10.200.0.10, Dst: 10.200.0.201

5. The appliance rewrites the addresses and forwards the packet to the Plug-in (placing the server address in an option field).

   Src: 10.200.0.201, Dst: 10.0.0.50

6. The connection is now fully open. The client and server send packets back and forth via the appliance.

   While the addresses are altered in Redirector mode, the destination port numbers are not (though the ephemeral port number may be). The data is not encapsulated. Redirector mode is a proxy, not a tunnel.

   There is no 1:1 relationship between packets (though in the end, the data received is always identical to the data sent). Compression may reduce many input packets into a single output packet. CIFS acceleration will perform speculative read-ahead and write-behind operations. Also, if packets are dropped between appliance and the Repeater Plug-in, the retransmission is handled by the appliance, not the server, using advanced recovery algorithms.

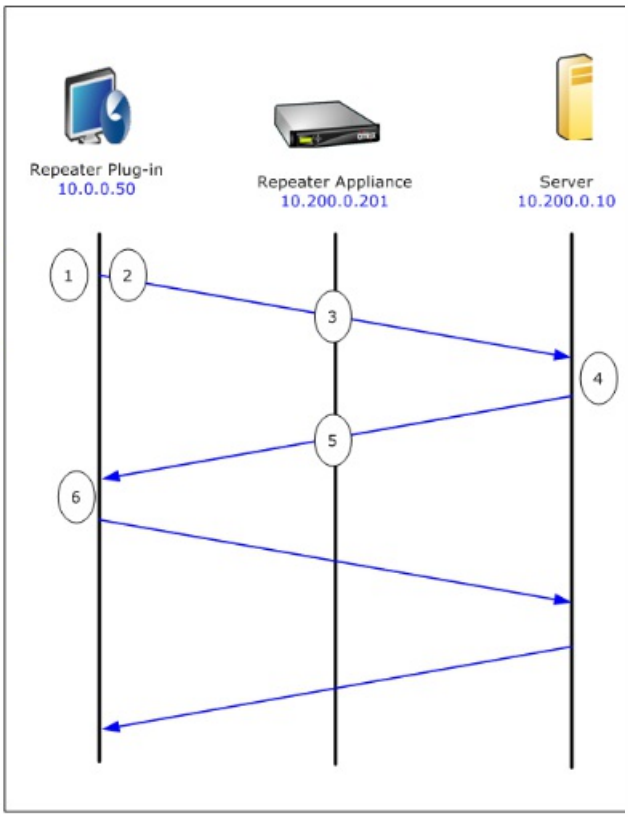

How the Plug-in Selects an Appliance

Aug 09, 2017

Each plug-in is configured with a list of appliances that it can contact to request an accelerated connection.

The appliances each have a list of *acceleration rules*, which is a list of target addresses or ports to which the appliance can establish accelerated connections. The plug-in downloads these rules from the appliances and matches the destination address and port of each connection with each appliance's rule set. If only one appliance offers to accelerate a given connection, selection is easy. If more than one appliance offers to accelerate the connection, the plug-in must choose one of the appliances.

The rules for appliance selection are as follows:

- If all the appliances offering to accelerate the connection are redirector-mode appliances, the leftmost appliance in the plug-in's appliance list is selected. (If the appliances were specified as DNS addresses, and the DNS record has multiple IP addresses, these too are scanned from left to right.)
- If some of the appliances offering to accelerate the connection use redirector mode and some use transparent mode, the transparent-mode appliances are ignored and the selection is made from the redirector-mode appliances.
- If all of the appliances offering to accelerate the connection use transparent mode, the plug-in does not select a specific appliance. It initiates the connection with WANOP Client Plug-in SYN options, and whichever candidate appliance attaches appropriate options to the returning SYN-ACK packet is used. This allows the appliance that is actually in line with the traffic to identify itself to the plug-in. The plug-in must have an open signaling connection with the responding appliance, however, or acceleration does not take place.
- Some configuration information is considered to be global. This configuration information is taken from the leftmost appliance in the list for which a signaling connection can be opened.

Deploying Appliances for Use with Plug-ins

Aug 09, 2017

Client acceleration requires special configuration on the WANOP Client Plug-in appliance. Other considerations include appliance placement. Plug-ins are typically deployed for VPN connections.

Use a Dedicated Appliance When Possible

Attempting to use the same appliance for both plug-in acceleration and link acceleration is often difficult, because the two uses sometimes call for the appliance to be at different points in the data center, and the two uses can call for different service-class rules.

In addition, a single appliance can serve as an endpoint for plug-in acceleration or as an endpoint for site-to-site acceleration, but cannot serve both purposes for the same connection at the same time. Therefore, when you use an appliance for both plug-in acceleration for your VPN and for site-to-site acceleration to a remote data center, plug-in users do not receive site-to-site acceleration. The seriousness of this problem depends on how much of the data used by plug-in users comes from remote sites.

Finally, because a dedicated appliance's resources are not divided between plug-in and site-to-site demands, they provide more resources and thus higher performance to each plug-in user.

Use Inline Mode When Possible

An appliance should be deployed on the same site as the VPN unit that it supports. Typically, the two units are in line with each other. An inline deployment provides the simplest configuration, the most features, and the highest performance. For best results, the appliance should be directly in line with the VPN unit.

However, appliances can use any deployment mode, except group mode or high availability mode. These modes are suitable for both appliance-to-appliance and client-to-appliance acceleration. They can be used alone (*transparent mode*) or in combination with redirector mode.

Place the Appliances in a Secure Part of Your Network

An appliance depends on your existing security infrastructure in the same way that your servers do. It should be placed on the same side of the firewall (and VPN unit, if used) as the servers.

Avoid NAT Problems

Network address translation (NAT) at the plug-in side is handled transparently and is not a concern. At the appliance side, NAT can be troublesome. Apply the following guidelines to ensure a smooth deployment:

- Put the appliance in the same address space as the servers, so that whatever address modifications are used to reach the servers are also applied to the appliance.
- Never access the appliance by using an address that the appliance does not associate with itself.
- The appliance must be able to access the servers by using the same IP addresses at which plug-in users access the same servers.
- In short, do not apply NAT to the addresses of servers or appliances.

Select Softboost Mode

On the Configure Settings: Bandwidth Management page, select Softboost mode. Softboost is the only type of

acceleration supported with the WANOP Client Plug-in Plug-in.

Define Plug-in Acceleration Rules

The appliance maintains a list of acceleration rules that tell the clients which traffic to accelerate. Each rule specifies an address or subnet and a port range that the appliance can accelerate.

What to Accelerate-The choice of what traffic to accelerate depends on the use the appliance is being put to:

- VPN accelerator - If the appliance is being used as a VPN accelerator, with all VPN traffic passing through the appliance, all TCP traffic should be accelerated, regardless of destination.
- Redirector mode - Unlike with transparent mode, an appliance in redirector mode is an explicit proxy, causing the plug-in to forward its traffic to the redirector-mode appliance even when doing so is not desirable. Acceleration can be counterproductive if the client forwards traffic to an appliance that is distant from the server, especially if this "triangle route" introduces a slow or unreliable link. Therefore, Citrix recommends that acceleration rules be configured to allow a given appliance to accelerate its own site only.
- Other uses - When the plug-in is used neither as a VPN accelerator nor in redirector mode, the acceleration rules should include addresses that are remote to the users and local to datacenters.

Defining the Rules- Define acceleration rules on appliance, on the **Configuration: WANOP Client Plug-in: Acceleration Rules** tab.

Rules are evaluated in order, and the action (Accelerate or Exclude) is taken from the first matching rule. For a connection to be accelerated, it must match an Accelerate rule.

The default action is to not accelerate.

Figure 1. Setting Acceleration Rules

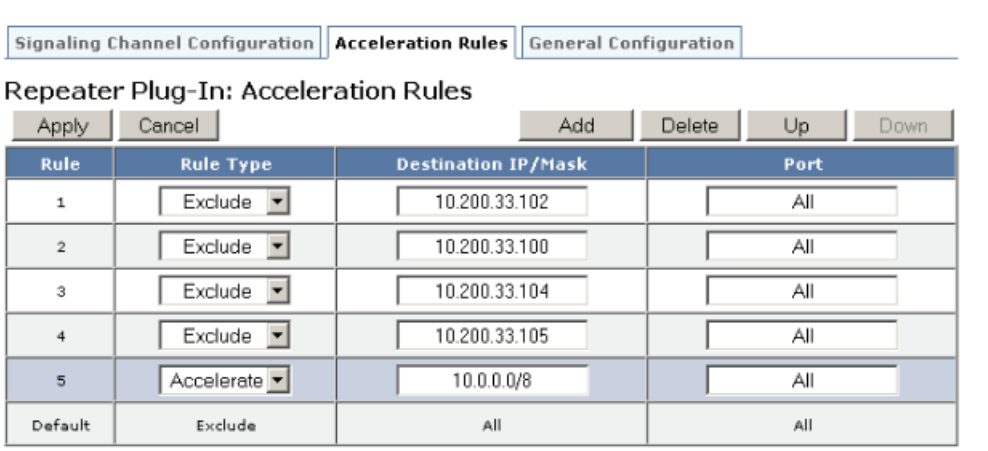

1. On the Configuration: WANOP Plug-in: Acceleration Rules tab:
   - Add an Accelerated rule for each local LAN subnet that can be reached by the appliance. That is, click **Add**, select **Accelerate**, and type the subnet IP address and mask.
   - Repeat for each subnet that is local to the appliance.
2. If you need to exclude some portion of the included range, add an Exclude rule and move it above the more general rule. For example, 10.217.1.99 looks like a local address. If it is really the local endpoint of a VPN unit, create an Exclude rule for it on a line above the Accelerate rule for 10.217.1.0/24.

3. If you want to use acceleration for only a single port (not recommended), such as port 80 for HTTP, replace the wildcard character in the Ports field with the specific port number. You can support additional ports by adding additional rules, one per port.
4. In general, list narrow rules (usually exceptions) before general rules.
5. Click **Apply**. Changes are not saved if you navigate away from this page before applying them.

IP Port Usage

Use the following guidelines for IP port usage:

- **Ports used for communication with WANOP Client Plug-in Plug-in**--The plug-in maintains a dialog with the appliance over a signaling connection, which by default is on port 443 (HTTPS), which is allowed through most firewalls.
- **Ports used for communication with servers**--Communication between the WANOP Client Plug-in Plug-in and the appliance uses the same ports that the client would use for communication with the server if the plug-in and appliance were not present. That is, when a client opens an HTTP connection on port 80, it connects to the appliance on port 80. The appliance in turn contacts the server on port 80.
  In redirector mode, only the well-known port (that is, the destination port on the TCP SYN packet) is preserved. The ephemeral port is not preserved. In transparent mode, both ports are preserved.

  The appliance assumes that it can communicate with the server on any port requested by the client, and the client assumes that it can communicate with the appliance on any desired port. This works well if appliance is subject to the same firewall rules as the servers. When such is the case, any connection that would succeed in a direct connection succeeds in an accelerated connection.

TCP Option Usage and Firewalls

WANOP Client Plug-in parameters are sent in the TCP options. TCP options can occur in any packet and are guaranteed to be present in the SYN and SYN-ACK packets that establish the connection.

Your firewall must not block TCP options in the range of 24-31 (decimal), or acceleration cannot take place. Most firewalls do not block these options. However, a Cisco PIX or ASA firewall with release 7.x firmware might do so by default, and therefore you might have to adjust its configuration.

Customizing the Plug-in MSI File

Aug 09, 2017

You can change parameters in the WANOP Client Plug-in distribution file, which is in the standard Microsoft Installer (MSI) format. Customization requires the use of an MSI editor.

> **Note**
>
> The altered parameters in your edited. MSI file apply only to new installations. When existing plug-in users update to a new release, their existing settings are retained. Therefore, after changing the parameters, you should advise your users to uninstall the old version before installing the new one.

Best Practices

Create a DNS entry that resolves to the nearest plug-in-enabled appliance. For example, define "Repeater.mycompany.com" and have it resolve to your appliance, if you have only one appliance. Or, if you have, say, five appliances, have Repeater.mycompany.com resolve to one of your five appliances, with the appliance selected on the basis of closeness to the client or to the VPN unit. For example, a client using an address associated with a particular VPN should see Repeater.mycompany.com resolve to the IP address of the WANOP Client Plug-in appliance connected to that VPN . Build this address into your plug-in binary with an MSI editor, such as Orca. When you add, move, or remove appliances, changing this single DNS definition on your DNS server updates the appliance list on your plug-ins automatically.

You can also have the DNS entry resolve to multiple appliances, but this is undesirable unless all appliances are configured identically, because the plug-in takes some of it characteristics from the leftmost appliance in the list and applies them globally (including SSL compression characteristics). This can lead to undesirable and confusing results, especially if the DNS server rotates the order of IP addresses for each request.

Installing the Orca MSI Editor

There are many MSI editors, including Orca, which is part of Microsoft's free Platform SDK and can be downloaded from Microsoft.

To install the Orca MSI Editor

1. Download the PSDK-x86.exe version of the SDK and execute it. Follow the installation instructions.
2. Once the SDK is installed, the Orca editor must be installed. It will be under Microsoft Platform SDK\Bin\Orca.Msi. Launch Orca.msi to install the actual Orca editor (orca.exe).
3. **Running Orca**--Microsoft provides its Orca documentation online. The following information describes how to edit the most important WANOP Client Plug-in Plug-in parameters.
4. Launch Orca with **Start > All Programs > Orca.** When a blank Orca window appears, open the WANOP Client Plug-in Plug-in MSI file with **File > Open**.

Figure 1. Using Orca

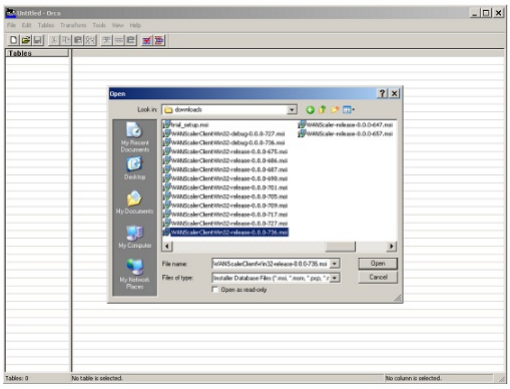

5. On the **Tables** menu, click **Property**. A list of all the editable properties of the .MSI file appears. Edit the parameters shown in the following table. To edit a parameter, double-click on its value, type the new value, and press **Enter**.

Parameter	Description	Default	Comments
WSAPPLIANCES	List of appliances	None	Enter the IP or DNS addresses of your WANOP appliances here, in a comma-separated list in the form of { appliance1, appliance2, appliance3 } . If the port used for signaling connections is different from the default (443), specify the port in the form Appliance1:port_number .
DBCMINSIZE	Minimum amount of disk space to use for compression, in megabytes	250	Changing this to a larger value (for example, 2000) improves compression performance but prevents installation if there is not enough disk space. The plug-in will not install unless there is at least 100 MB of free disk space in addition to the value that you specify for DBCMINSIZE.
EKEYPEM	Private key for the plug-in. Part of the certificate/key pair used with SSL compression	None	Use Orca's Paste Cell command. The normal Paste function does not preserve the key's format. Should be a private key in PEM format (starting with -----BEGIN RSA PRIVATE KEY-----)
X509CERTPEM	Certificate for the plug-in. Part of the certificate/key pair used with SSL compression	None	Use Orca's Paste Cell command. The normal Paste function does not preserve the key's format. Should be a certificate in PEM format (starting with -----BEGIN CERTIFICATE -----)
CACERTPEM	Certification Authority Certificate for	None	Use Orca's Paste Cell command. The normal Paste function does not preserve the key's format. Should be a certificate in

	the plug-in. Used with SSL compression		PEM format (starting with -----BEGIN CERTIFICATE -----)

1.
2. On the Tables menu, click Property. A list of all the editable properties of the .MSI file appears. Edit the parameters shown in the following table. To edit a parameter, double-click on its value, type the new value, and press Enter.

Parameter	Description
WSAPPLIANCES	List of appliances
DBCMINSIZE	Minimum amount of disk space to use for compression, in megabytes
PRIVATEKEYPEM	Private key for the plug-in. Part of the certificate/key pair used with SSL compression
X509CERTPEM	Certificate for the plug-in. Part of the certificate/key pair used with SSL compression
CACERTPEM	Certification Authority Certificate for the plug-in. Used with SSL compression

Figure 2. Editing Parameters in Orca

3. When done, use the **File: Save As** command to save your edited file with a new filename; for example, test.msi.

Figure 2: Editing Parameters in Orca

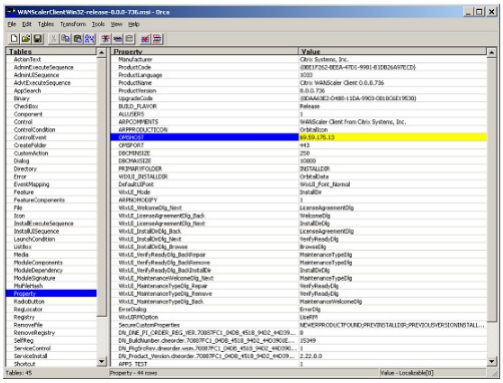

6. When done, use the **File: Save As** command to save your edited file with a new filename; for example, test.msi.

Your plug-in software has now been customized.

Note: Some users have seen a bug in orca that causes it to truncate files to 1 MB. Check the size of the saved file. If it has been truncated, make a copy of the original file and use the Save command to overwrite the original.
Once you have customized the appliance list with Orca and distributed the customized MSI file to your users, the user does not need to type in any configuration information when installing the software.

Deploying Plug-ins On Windows Systems

Aug 09, 2017

The WANOP Client Plug-in is an executable Microsoft installer (MSI) file that you download and install as with any other web-distributed program. Obtain this file from the MyCitrix section of the Citrix.com website.

Note: The WANOP Client Plug-in user interface refers to itself as "Citrix Acceleration Plug-in Manager."

The only user configuration needed by the plug-in is the list of appliance addresses. This list can consist of a comma-separated list of IP or DNS address. The two forms can be mixed. You can customize the distribution file so that the list points to your appliance by default. Once installed, operation is transparent. Traffic to accelerated subnets is sent through an appropriate appliance, and all other traffic is sent directly to the server. The user application is unaware that any of this is happening.

Installation

Aug 09, 2017

To install WANOP Client Plug-in Plug-in accelerator on Windows system:

1. The Repeater*.msi file is an installation file. Close all applications and any windows that might be open, and then launch the installer it in the usual way (double-click on in a file window, or use the run command).

Figure 1. Initial Installation Screen

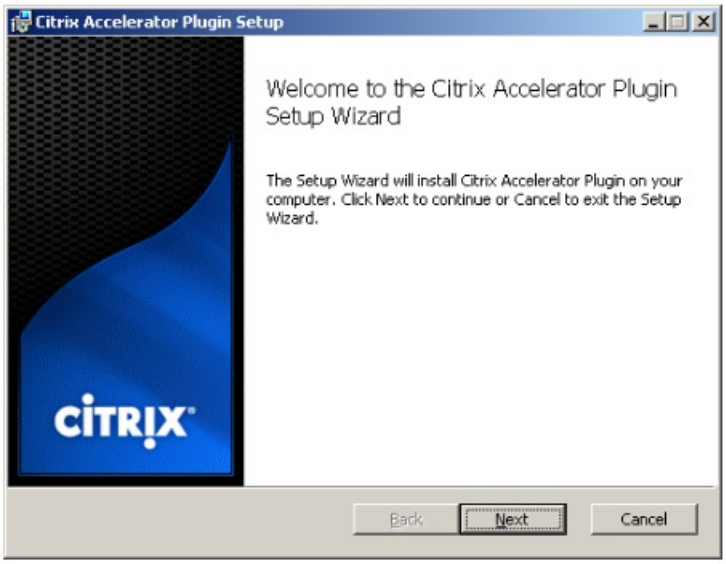

The steps below are for an interactive installation. A silent installation can be performed with the command:

msiexec /i client_msi_file /qn

2. The installation program prompts for the location in which to install the software. The directory that you specify is used for both the client software and the disk-based compression history. Together, they require a minimum of 500 MB of disk space.

3. When the installer finishes, it might ask you to restart the system. After a restart, the WANOP Client Plug-in Plug-in starts automatically.

Figure 2. Final Installation Screen

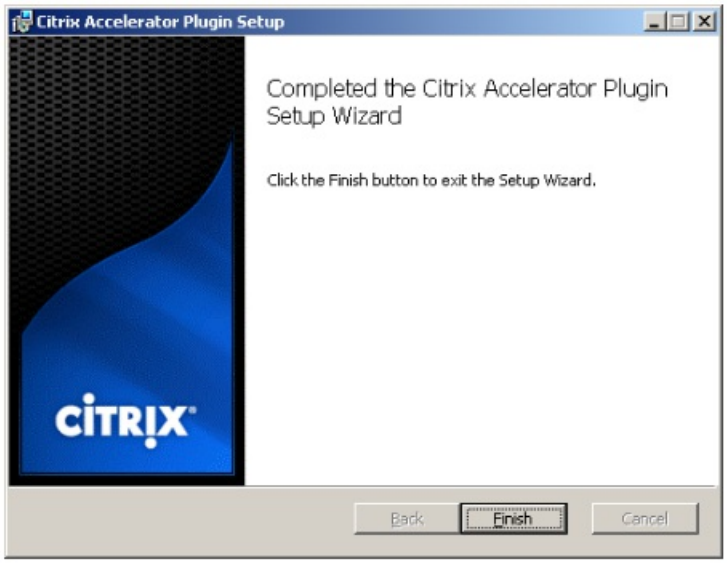

4. Right-click the Accelerator icon in the task bar and select **Manage Acceleration** to launch the Citrix Plug-in Accelerator Manager.

Figure 3. Citrix Accelerator Plug in Manager, Initial (Basic) Display

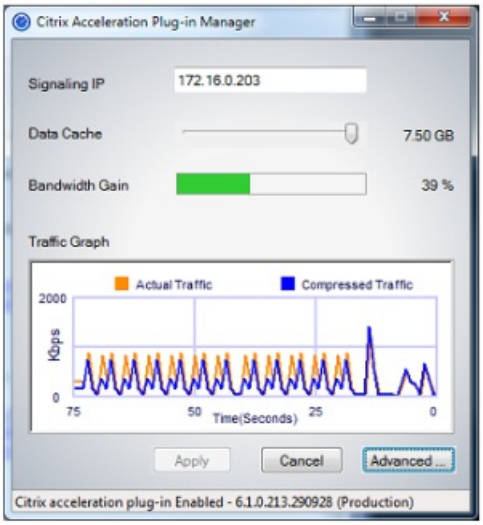

5. If the .MSI file has not been customized for your users, specify the signaling address and the amount of disk space to use for compression:

- In the Appliances: Signaling Addresses field, type the signaling IP address of your appliance. If you have more than one Plug-in-enabled appliance, list them all, separated by commas. Either IP or DNS addresses are acceptable.
- Using the Data Cache slider, select the amount of disk space to use for compression. More is better. 7.5 GB is not too much, if you have that much disk space available.
- Press Apply.

The WANOP Client Plug-in accelerator is now running. All future connections to accelerated subnets will be accelerated

On the plug-in's Advanced Rules tab, the Acceleration Rules list should show each appliance as Connected and each appliance's accelerated subnets as Accelerated. If not, check the Signaling Addresses IP field and your network connectivity in general.

Troubleshooting Plug-ins

Aug 09, 2017

Plug-in installation generally goes smoothly. If not, check for the following issues:

Common problems

- If you do not reboot the system, the WANOP Client Plug-in will not run properly.
- A highly fragmented disk can result in poor compression performance.
- A failure of acceleration (no accelerated connections listed on the **Diagnostics** tab) usually indicates that something is preventing communication with the appliance. Check the **Configuration: Acceleration Rules** listing on the plug-in to make sure that the appliance is being contacted successfully and that the target address is included in one of the acceleration rules. Typical causes of connection failures are:
  - The appliance is not running, or acceleration has been disabled.
  - A firewall is stripping WANOP Client Plug-in TCP options at some point between the plug-in and appliance.
  - The plug-in is using an unsupported VPN.

Deterministic Network Enhancer locking error

On rare occasions, after you install the plug-in and restart your computer, the following error message appears twice:

Deterministic Network Enhancer installation requires a reboot first, to free locked resources. Please run this install again after restarting the computer.

If this occurs, do the following:

1. Go to **Add/Remove Programs** and remove the WANOP Client Plug-in, if present.
2. Go to **Control Panel** > **Network Adapters** > **Local Area Connection** > **Properties**, find the entry for Deterministic Network Enhancer, clear its check box, and click **OK**. (Your network adapter might be called by a name other than "Local Area Connection.")
3. Open a command window and go to c:\windows\inf (or the equivalent directory if you have installed Windows in a non-standard location).
4. Type the following command:
   find "dne2000.cat" oem*.inf
5. Find the highest-numbered oem*.inf file that returned a matching line (the matching line is CatalogFile= dne2000.cat) and edit it. For example:
   notepad oem13.inf
6. Delete everything except the three lines at the top that start with semicolons, and then save the file. This will clear out any inappropriate or obsolete settings and the next installation will use default values.
7. Retry the installation.

Other Installation Problems

Any problem with installing the WANOP Client Plug-in is usually the result of existing networking, firewall, or antivirus software interfering with the installation. Usually, once the installation is complete, there are no further problems.

If the installation fails, try the following steps:

1. Make sure the plug-in installation file has been copied to your local system.
2. Disconnect any active VPN/remote networking clients.

3. Disable any firewall and antivirus software temporarily.
4. If some of this is difficult, do what you can.
5. Reinstall the WANOP Client Plug-in.
6. If this doesn't work, reboot the system and try again.

WANOP Plug-in GUI Commands

Aug 09, 2017

The WANOP Client Plug-in GUI appears when you right-click the **Citrix Accelerator Plug-in** icon and select **Manage Acceleration**. The GUI's Basic display appears first. There is also an Advanced display that can be used if desired.

Basic Display

Aug 09, 2017

On the Basic page, you can set two parameters:

- The Signaling Addresses field specifies the IP address of each appliance that the plug-in can connect to. Citrix recommends listing only one appliance, but you can create a comma-separated list. This is an ordered list, with the leftmost appliances having precedence over the others. Acceleration is attempted with the leftmost appliance for which a signaling connection can be established. You can use both DNS addresses and IP addresses.
  Examples: 10.200.33.200, ws.mycompany.com, ws2.mycompany.com
- The Data Cache slider adjusts the amount of disk space allocated to the plug-in's disk-based compression history. More is better.

In addition, there is a button to move to the Advanced display.

Advanced Display

Aug 09, 2017

The Advanced page contains four tabs: Rules, Connections, Diagnostics, and Certificates.

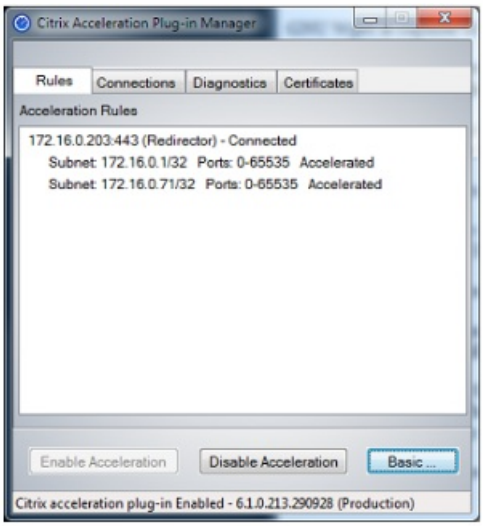

At the bottom of the display are buttons to enable acceleration, disable acceleration, and return to the Basic page.

Rules Tab

The Rules tab displays an abbreviated list of the acceleration rules downloaded from the appliances. Each list item shows the appliance's signaling address and port, acceleration mode (redirector or transparent), and connection state, followed by a summary of the appliance's rules.

Connections Tab

The **Connections** tab lists the number of open connections of different types:

- **Accelerated Connections**--The number of open connections between the WANOP Client Plug-in Plug-in and appliances. This number includes one signaling connection per appliance but does not include accelerated CIFS connections. Clicking More opens a window with a brief summary of each connection. (All of the More buttons allow you to copy the information in the window to the clipboard, should you want to share it with Support.)
- **Accelerated CIFS Connections**--The number of open, accelerated connections with CIFS (Windows file system) servers. This is usually the same as the number of mounted network file systems. Clicking More displays the same information as with accelerated connections, plus a status field that reports Active if the CIFS connection is running with WANOP Client Plug-in's special CIFS optimizations.
- **Accelerated MAPI Connections**--The number of open, accelerated Outlook/Exchange connections.
- **Accelerated ICA connections**--The number of open, accelerated XenApp and XenDesktop connections using the ICA or CGP protocols.
- **Unaccelerated Connections**--Open connections that are not being accelerated. You can click More to display a brief description of why the connection was not accelerated. Typically, the reason is that no appliance accelerates the destination address, which is reported as Service policy rule .
- **Opening/Closing Connections**--Connections that are not fully open, but are in the process of opening or closing (TCP "half-open" or "half-closed" connections). The More button displays some additional information about these connections.

Diagnostics Tab

The Diagnostics page reports the number of connections in different categories, and other useful information.

- **Start Tracing/Stop Tracing**--If you report a problem, your Citrix representative might ask you to perform a connection trace to help pinpoint problems. This button starts and stops the trace. When you stop tracing, a pop-up window shows the trace files. Send them to your Citrix representative by the means he or she recommends.
- **Clear History**--This feature should not be used.
- **Clear Statistics**--Pressing this button clears the statistics on the Performance tab.
- **Console**--A scrollable window with recent status messages, mostly connection-open and connection-close messages, but also error and miscellaneous status messages.

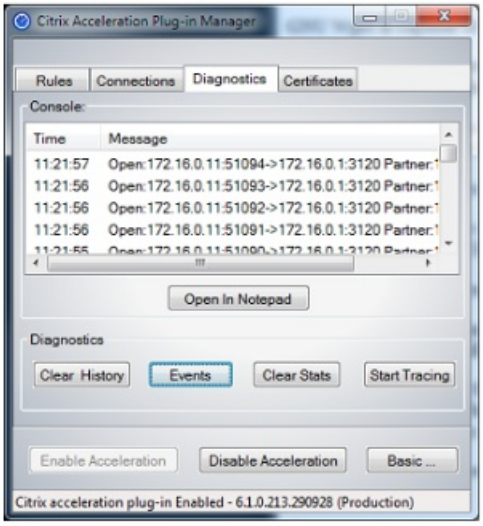

Certificates Tab

On the Certificates tab, you can install security credentials for the optional secure peering feature. The purpose of these security credentials is to enable the appliance to verify whether the plug-in is a trusted client or not.

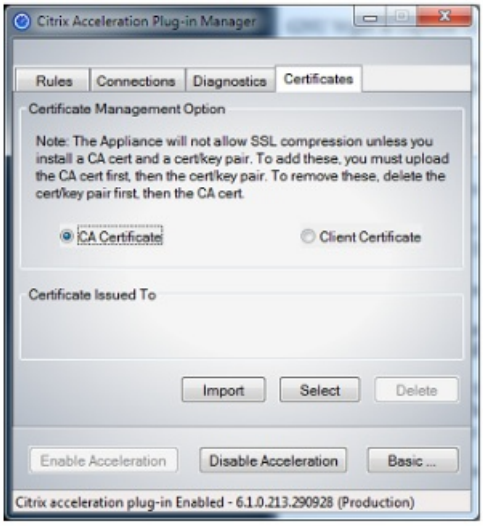

To upload the CA certificate and certificate-key pair:

1. Select CA **Certificate Management**.

2. Click **Import**.
3. Upload a CA certificate. The certificate file must use one of the supported file types (.pem, .crt., .cer, or .spc). A dialog box might appear, asking you to Select the certificate store you want to use and presenting you with a list of keywords. Select the first keyword in the list.
4. Select **Client Certificate Management**.
5. Click **Import**.
6. Select the format of the certificate-key pair (either PKCS12 or PEM/DER).
7. Click **Submit**.

> ## Note
>
> In the case of PEM/DER, there are separate upload boxes for certificate and key. If your certificate-key pair is combined in a single file, specify the file twice, once for each box.

Updating the WANOP Plug-in

Aug 09, 2017

To install a newer version of the WANOP Client Plug-in, follow the same procedure you used when installing the plug-in for the first time.

Uninstalling the WANOP Client Plug-in Plug-in

To uninstall the WANOP Client Plug-in Plug-in To uninstall the WANOP Client Plug-in, use the Windows Add/Remove Programs utility. The WANOP Client Plug-in is listed as **Citrix Acceleration Plug-in** in the list of currently installed programs. Select it and click **Remove**.

You must restart the system to finish uninstalling the client.

Troubleshooting WANOP Plug-in

Aug 09, 2017

- **Issue**: I am facing signaling channel connectivity issues. How can I resolve these issues?
  Resolution: To resolve signaling channel connectivity issues, perform the following troubleshooting steps:

  - Verify that you have correctly configured the signaling IP address. You can do so by pinging the signaling IP address and verifying the response.
  - Verify that the signaling status is enabled on the WANOP appliance.
  - Verify that the firewall installed on the network does not remove the WANOP TCP options.
  - Verify that a valid WANOP plug-in license is installed on the WANOP appliance.
  - Verify that the Signaling Channel Source Filtering configuration does not block the Client Source IP address.
  - If you have enabled LAN Detection, verify that the Round Trip Time between the WANOP plug-in and WANOP appliance is an acceptable value.
- **Issue**: On a WANOP 4000 appliance, I am not able to disable the WANOP plug-in.
  Cause: This is a known issue.

  Resolution: None. You cannot disable the WANOP plug-in on a WANOP 4000 appliance.

- **Issue**: When connecting to the WANOP appliance by using the WANOP plug-in, the following error message entry is logged on the Alerts tab:
  More WANOP Plug-ins than the current limit of <Number> have attempted to connect to this Appliance.

  Cause: The number of connections to the WANOP appliance has exceeded the licensed user limit.

  Resolution: Either wait for a user to disconnect or terminate a connection.

- **Issue**: Incorrect signaling IP address is configured on a WANOP 4000 or 5000 appliance.
  Resolution: To update the signaling IP address on a WANOP 4000 or 5000 appliance, complete the following procedure:

  1. Log on to the NetScaler instance of the WANOP appliance.
  2. Navigate to the Traffic Management > Load Balancing > Virtual Servers > BR_LB_VIP_SIG page.
  3. Update the signaling IP address.
  4. Save the configuration.
- **Issue**: CIFS and ICA traffic is not getting accelerated.
  Resolution: To resolve this issue, perform the following troubleshooting steps:

  - Verify that acceleration rules for IP address and port numbers are correctly defined for the WANOP plug-in.
  - Verify that CIFS or ICA connections are established after signaling connection is successful.
  - Verify the acceleration policy for the service class being used.

Configuring Service Class Association with SSL Profiles

Aug 09, 2017

All SSL related configuration is available through the new configuration editor of the appliance for security and usability. On the SD-WAN Enterprise Edition and two-box deployments, service classes are configured from the configuration editor and hence you cannot attach any SSL profiles. To accommodate the expression of SSL profile mapping to a service class, the work flow for SSL profiles is changed to allow for attaching Service classes in the profile node.

One of the limitations is that the SSL profile will get attached to all rules in a service class. If you need to attach the SSL profile selectively to a particular rule, the service class configuration is split into detailed rules for further selection.

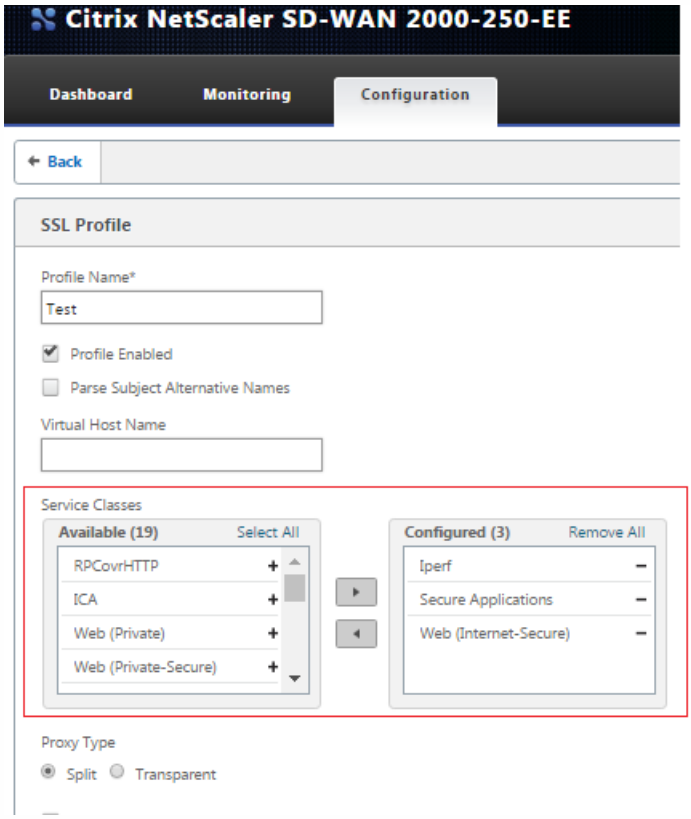

To create SSL profile on new Enterprise Edition appliance at the data center:

1. In the SD-WAN web GUI, go to the **Configuration** > **Secure Acceleration** page. Click **Add Profile**. Create the **SSL Profile**.

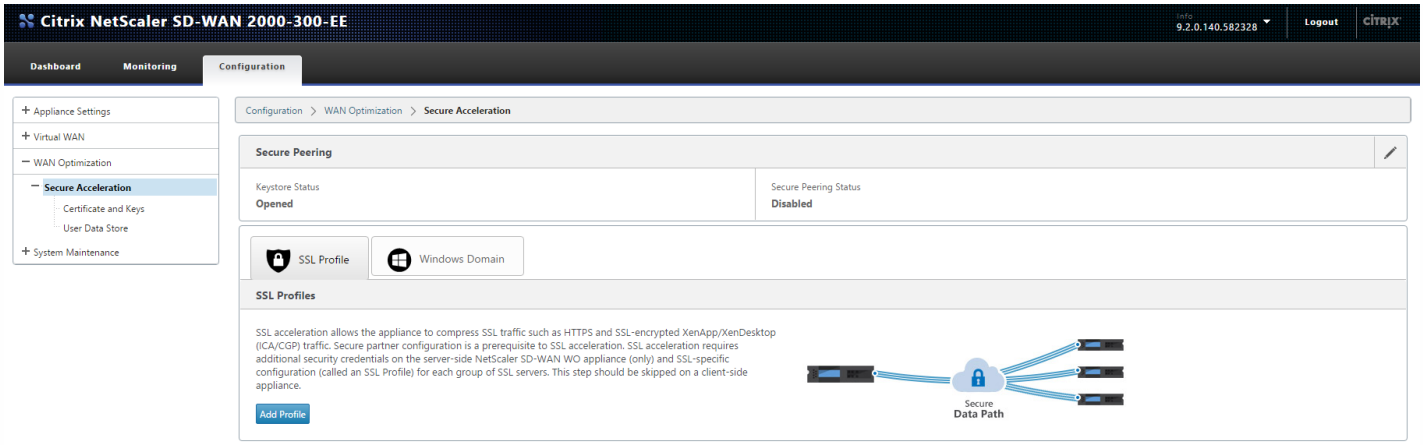

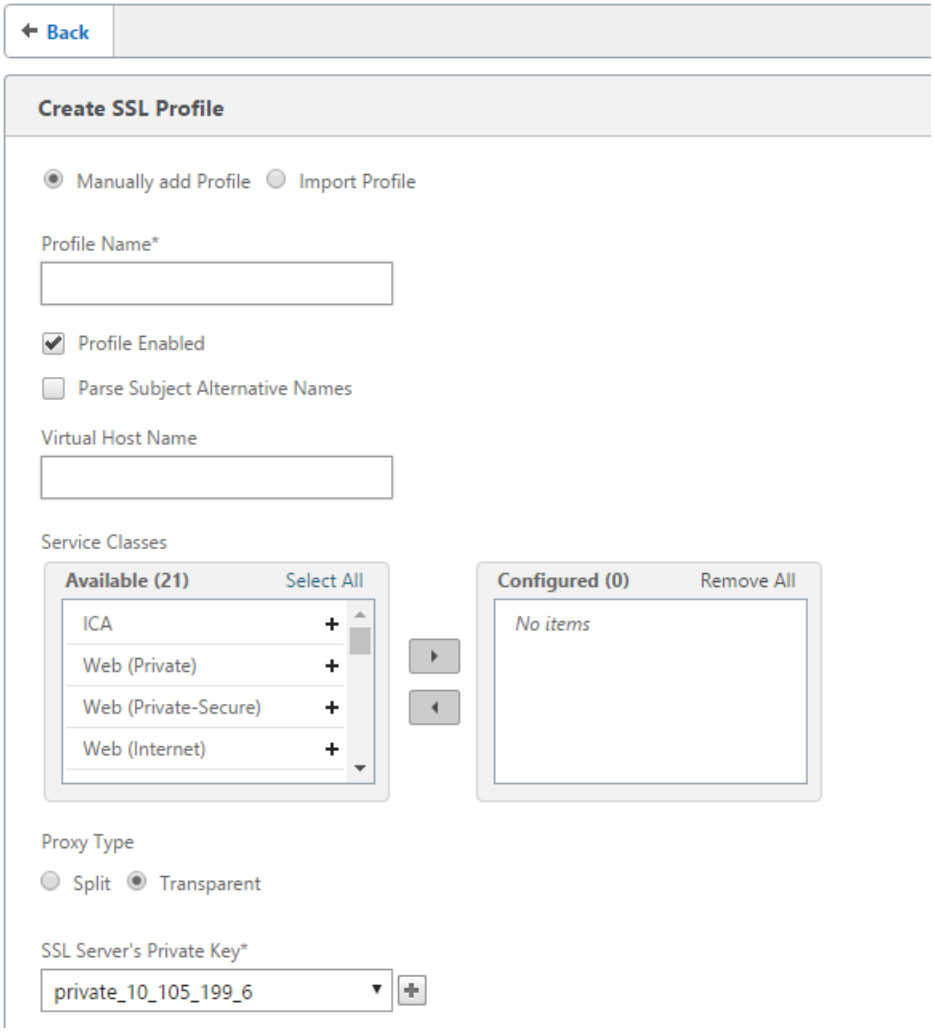

2. In the **Create SSL Profile** page, provide a profile name and select **Service Classes** that will be associated to this profile. Choose **Proxy Type** and provide relevant data and click **Create**.

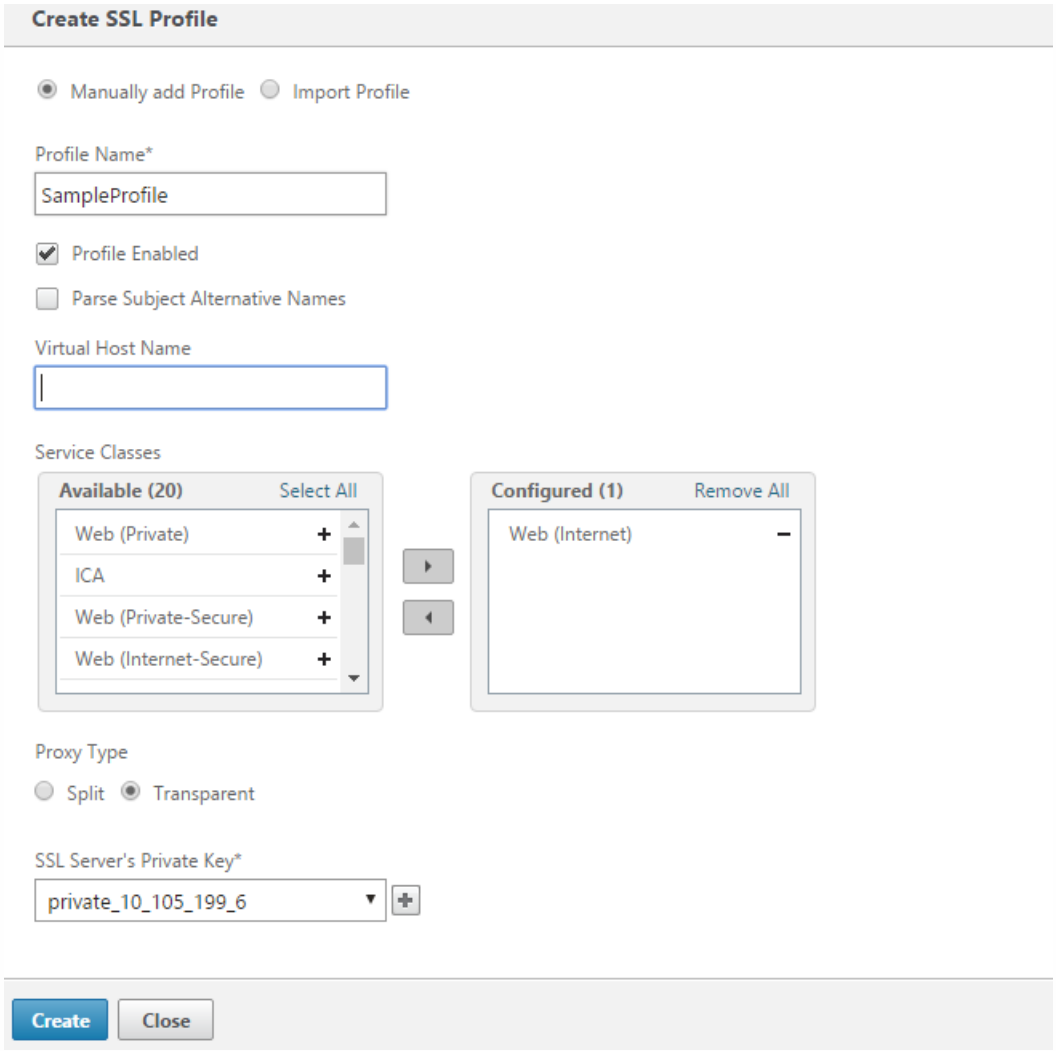

3. After SSL Profile is created successfully and service class is associated, view the SSL profile information as shown below.

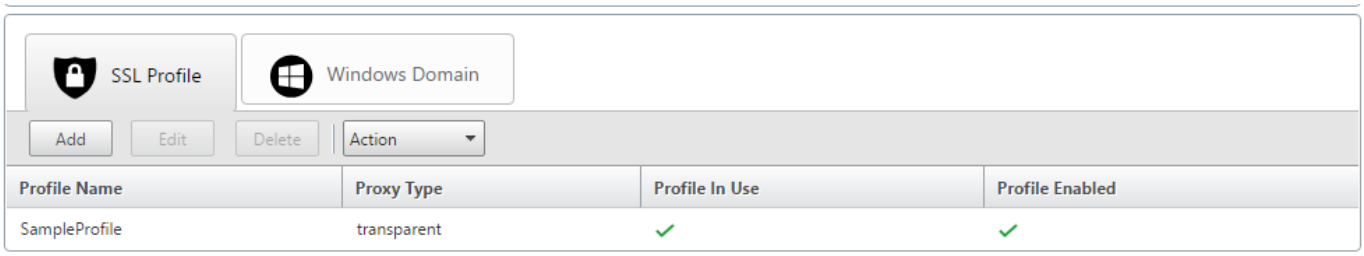

Standard MIB Support

Aug 09, 2017

The following standard MIBs are supported by the SD-WAN Appliances.

MIB	RFC Definition (Link)
DISMAN-EVENT-MIB	https://www.ietf.org/rfc/rfc2981.txt
IF-MIB	https://www.ietf.org/rfc/rfc2863.txt
IP-FORWARD-MIB	https://www.ietf.org/rfc/rfc4292.txt
IP-MIB (Partial)	https://www.ietf.org/rfc/rfc4293.txt
Q-BRIDGE-MIB (Partial)	http://www.ieee802.org/1/files/public/MIBs/IEEE8021-Q-BRIDGE-MIB-201112120000Z.txt
RFC1213-MIB	https://www.ietf.org/rfc/rfc1213.txt
SNMPv2-MIB	https://www.ietf.org/rfc/rfc3418.txt
TCP-MIB	https://www.ietf.org/rfc/rfc4022.txt
P-BRIDGE-MIB.txt	http://www.icir.org/fenner/mibs/extracted/P-BRIDGE-MIB-rfc2674.txt
RMON2-MIB.txt	https://www.ietf.org/rfc/rfc3273.txt
TOKEN-RING-RMON-MIB.txt	http://www.icir.org/fenner/mibs/extracted/TOKEN-RING-RMON-MIB-rmonmib-01.txt

You must download the following SNMP files before you can start monitoring a NetScaler SD-WAN appliance:

- CITRIX-COMMON-MIB.txt
- APPACCELERATION-SMI.txt
- APPACCELERATION-PRODUCTS-MIB.txt
- APPACCELERATION-TC.txt
- APPACCELERATION-STATUS-MIB.txt

- APPCACHE-MIB.txt
- SDX-MIB-smiv2.mib

The MIB files are used by SNMPv3 managers and SNMPv3 trap listeners. The files include the SD-WAN appliance enterprise MIBs, which provides SD-WAN-specific events. To download MIB files, in the SD-WAN web management interface:

1. Navigate to **Configuration** > **Appliance Settings** > **SNMP** > **Download MIB File** page.

2. Select the required **MIB** file.

3. Click **View**.

The MIB file opens in MIB browser.

Additional Notes

- Support for these MIBs is provided by default by the **net-snmp snmpd** daemon process on Linux systems. The MIBs provide the basis for supporting Network Management applications, for example: Nagios or SolarWinds.
- The Ethernet port packet and byte counters are in the **IF-MIB** inside the **ifTable**. System information is in the system object.
- Ethernet ports are included in the **ifTable**, so walking that should be sufficient to ensure that the SNMP subsystem is running.
- Support for the **Q-BRIDGE-MIB** and the **IP-MIB** provides support for the network mapping application in SolarWinds.

References

For additional information about adding SNMP manager, configuring SNMP View/Alarm, and adding SNMP server, see the CloudBridge 7.4 documentation at: http://docs.citrix.com/content/dam/docs/en-us/cloudbridge/7-3/downloads/en.cloudbridge.cb-wrapper-73-con.pdf

Best Practices - Security

Apr 16, 2018

This article outlines security best practices for the NetScaler SD-WAN solution. It provides general security guidance for NetScaler SD-WAN deployments.

NetScaler SD-WAN Deployment Guidelines

To maintain security through the deployment lifecycle, Citrix recommends the following security consideration:

- Physical Security
- Appliance Security
- Network Security
- Administration and Management

Physical Security

Deploy NetScaler SD-WAN Appliances in a Secure Server Room - The appliance or server on which NetScaler SD-WAN is installed, should be placed in a secure server room or restricted data center facility, which protects the appliance from unauthorized access. At the minimum, access should be controlled by an electronic card reader. Access to the appliance should be monitored by CCTV that continuously records all activity for auditing purposes. In the event of a break-in, electronic surveillance system should send an alarm to the security personnel for immediate response.

Protect Front Panel and Console Ports from Unauthorized Access - Secure the appliance in a large cage or rack with physical-key access control.

Protect Power Supply - Make sure that the appliance is protected with an uninterruptable power supply (UPS).

Appliance Security

For appliance security, secure the operating system of any server hosting a NetScaler SD-WAN virtual appliance (VPX), perform remote software updates, and following secure lifecycle management practices:

- Secure the Operating System of Server Hosting a NetScaler SD-WAN VPX Appliance - A NetScaler SD-WAN VPX appliance runs as a virtual appliance on a standard server. Access to the standard server should be protected with role based access control and strong password management. Additionally, Citrix recommends periodic updates to the server with the latest security patches for the operating system, and update-to-date antivirus software on the server.
- Perform Remote Software Updates - Install all security updates to resolve any known issues. Refer to the Security Bulletins web page to sign up and receive upe-to-date security alerts.
- Follow Secure Lifecycle Management Practices - To manage an appliance when redeploying, or initiating RMA, and decommissioning sensitive data, complete the data-reminisce countermeasures by removing the persistent data from the appliance.

Network Security

For network security, do not use the default SSL certificate. Use Transport Layer Security (TLS) when accessing the administrator interface, protect the appliance's non-routable management IP address, configure a high availability setup, and implement Administration and Management safeguards as appropriate for the deployment.

- Do not use the NetScaler Default SSL Certificate - An SSL certificate from a reputable Certificate Authority simplifies

the user experience for Internet-facing Web applications. Unlike the situation with a self-signed certificate or a certificate from the reputable Certificate Authority, web browsers do not require users to install the certificate from the reputable Certificate Authority to initiate secure communication to the Web server.

- Use Transport Layer Security when Accessing Administrator Interface - Make sure that the management IP address is not accessible from the Internet or is at least protected by a secured firewall. Make sure that the LOM IP address is not accessible from the Internet or is at least protected by a secured firewall.
- Secure Administration and Management Accounts – Create an alternative admin account, set strong passwords for admin and viewer accounts. When configure remote account access, consider configuring externally authenticated administrative management of accounts using RADIUS and TACAS. Change the default password for the admin user accounts, configure NTP, use the default session timeout value, use SNMPv3 with SHA Authentication and AES encryption.

NetScaler SD-WAN overlay network protects data traversing the SD-WAN overlay network.

Secure Administrator Interface

For secure web management access, replace default system certificates by uploading and installing certificates from a reputable Certificate Authority.

Configuration > **Appliance Settings** > **Administrator Interface**:

User Accounts:

- Change local user password
- Manage users

HTTPS Certs:

- Certificate
- Key

Miscellaneous:

- Web Console Timeout

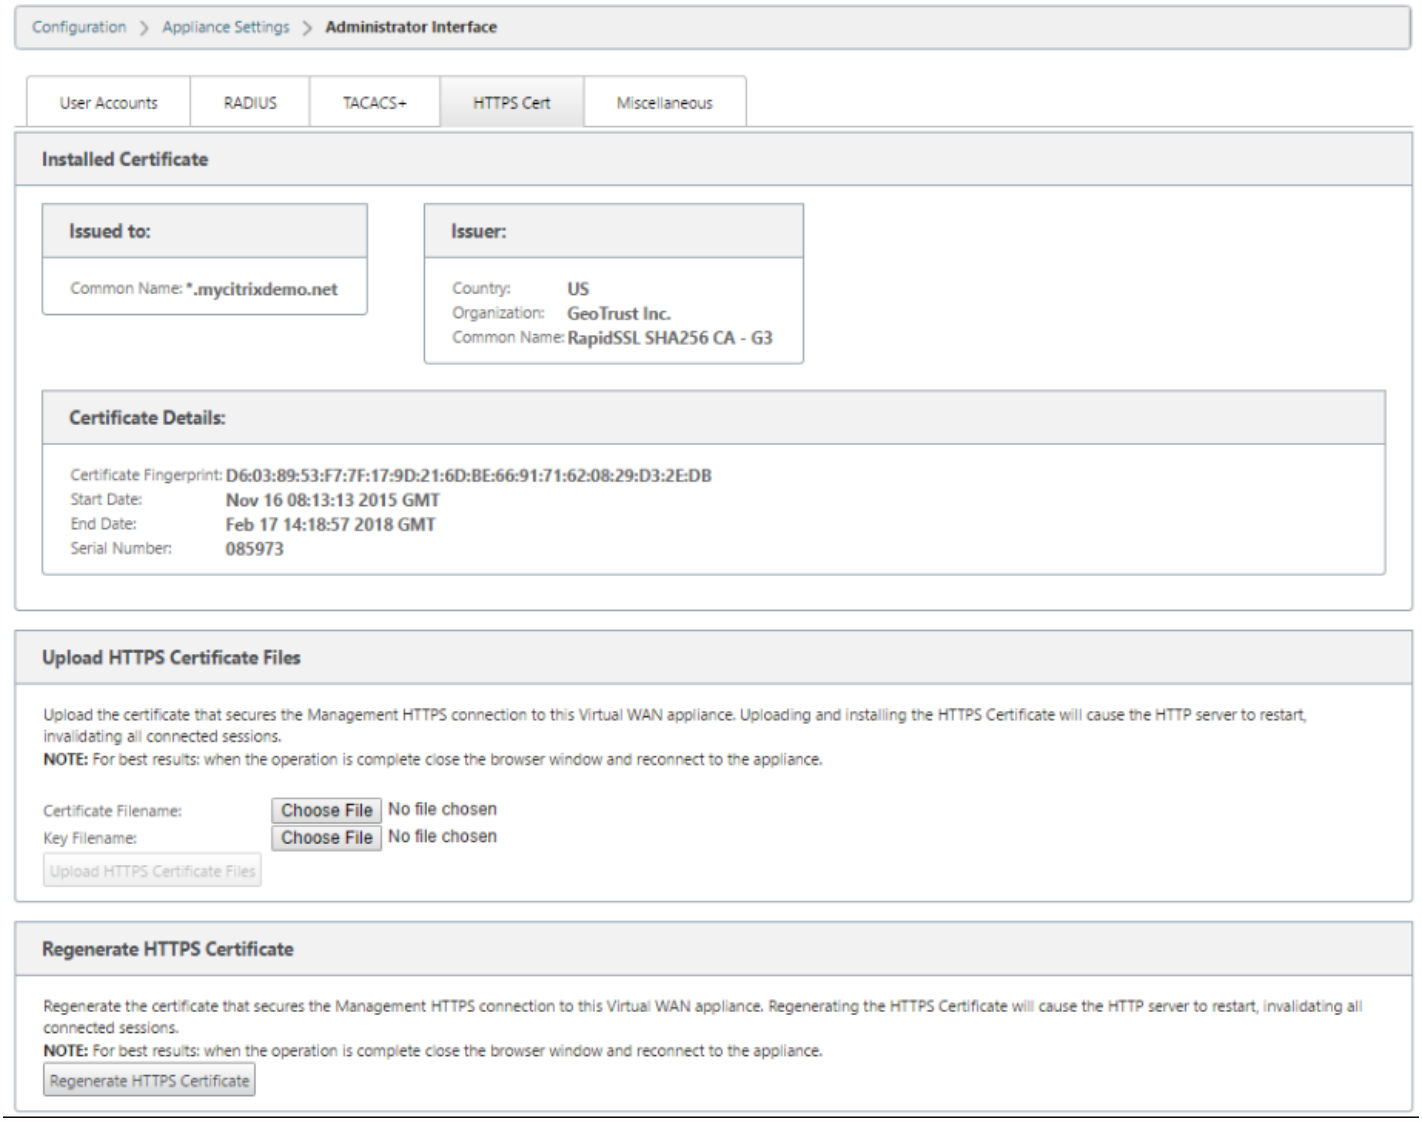

Configuration Editor > **Advanced** > **Global** > **Virtual WAN Network Settings**

Global Firewall Settings

- Global Policy Template
- Default Firewall Actions
- Default Connection State Tracking

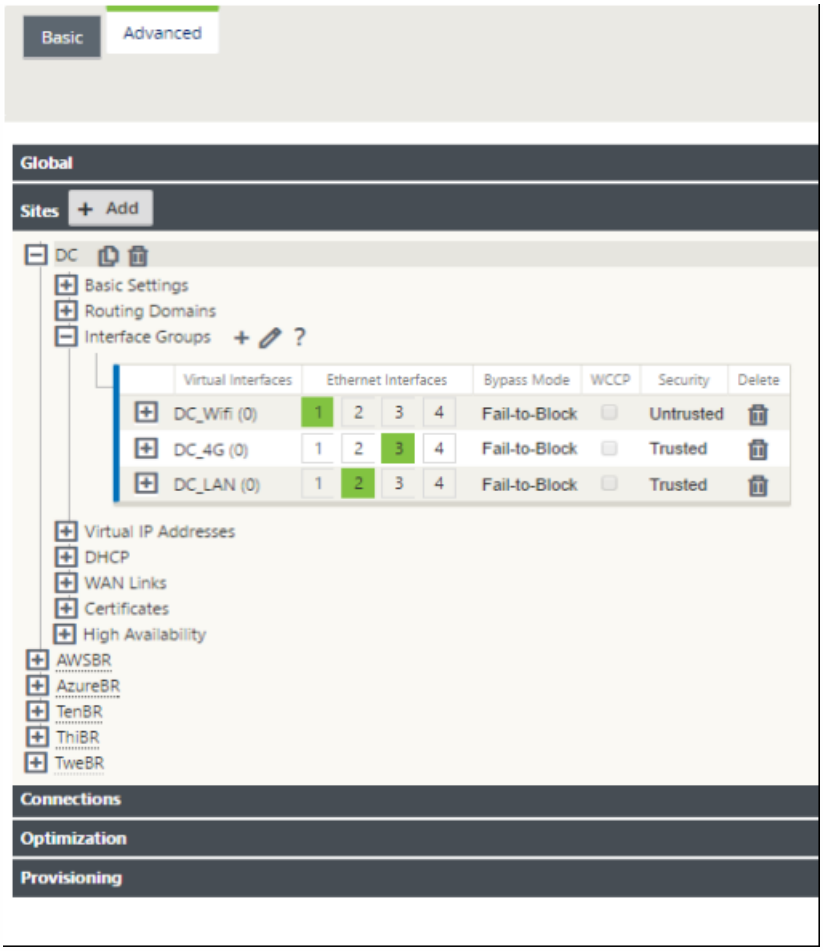

Global Virtual Path Encryption Settings

- AES 128-bit (default)
- Encryption Key Rotation (Default)
- Extended Packet Encryption Header
- Extended Packet Authentication Trailer

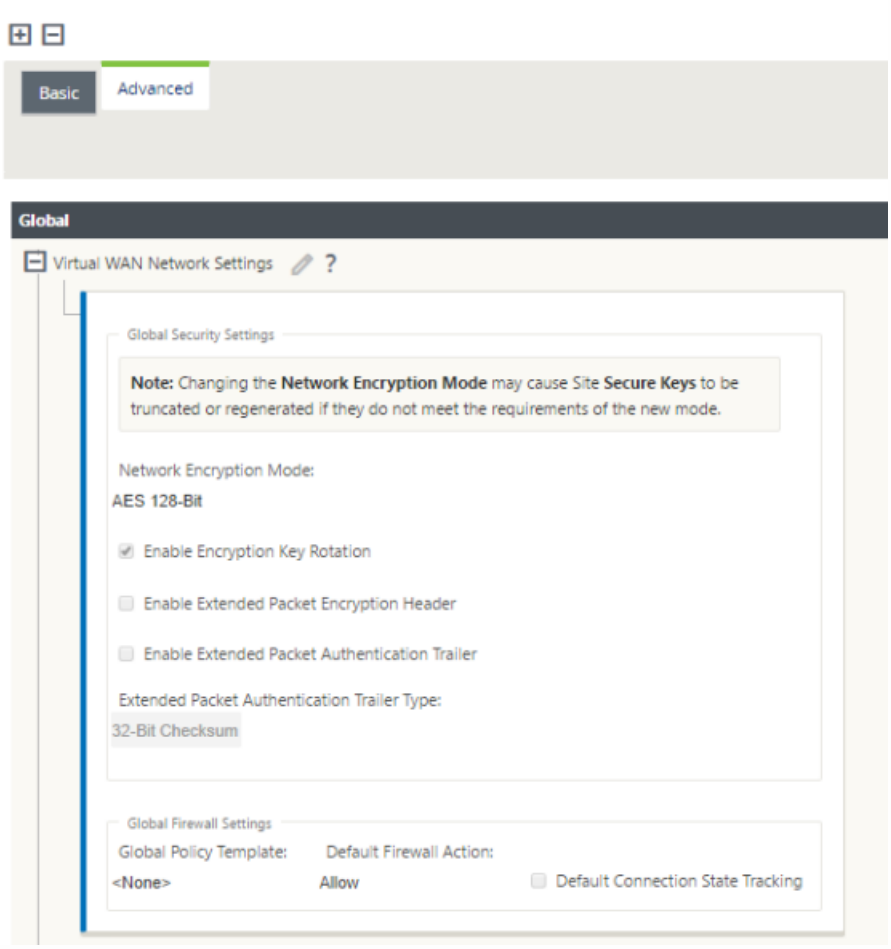

Global virtual Path Encryption Settings

- AES-128 data encryption is enabled by default. It is recommended to use AES-128 or additional protection of AES-256 encryption level for path encryption. Ensure that “enable Encryption Key Rotation” is set to ensure key regeneration for every Virtual Path with encryption enabled using an Elliptic Curve Diffie-Hellman key exchange at intervals of 10-15 minutes.

If the network requires message authentication in addition to confidentiality (i.e. tamper protection), Citrix recommends using IPsec data encryption. If only confidentiality is required, Citrix recommends using the enhanced headers.

- Extended Packet Encryption Header enables a randomly seeded counter to be prepended to the beginning of every encrypted message. When encrypted, this counter will serve as a random initialization vector, deterministic only with the encryption key. This will randomize the output of the encryption, providing strong message indistinguishability. Keep in mind that when enabled this option will increase packet overhead by 16 bytes
- Extended Packet Authentication Trailer appends an authentication code to the end of every encrypted message. This trailer allows for the verification that packets are not modified in transit. Keep in mind this option will increase packet overhead.

Firewall Security

The recommended Firewall configuration is with a default Firewall action as deny all at first, then add exceptions. Prior to adding any rules, document and review the purpose of the firewall rule. Use Stateful inspection and Application level inspection where possible. Simplify rules and eliminate redundant rules. Define and adhere to a change management process that tracks and allows for review of changes to Firewall settings. Set the Firewall for all appliances to track

connections through the appliance using the global settings. Tracking connections verifies that packets are properly formed and are appropriate for the connection state. Create Zones appropriate to the logical hierarchy of the network or functional areas of the organization. Keep in mind that zones are globally significant and can allow geographically disparate networks to be treated as the same security zone. Create the most specific policies possible to reduce the risk of security holes, avoid the use of Any in Allow rules. Configure and maintain a Global Policy Template to create a base level of security for all appliances in the network. Define Policy Templates based on functional roles of appliances in the network and apply them where appropriate. Define Policies at individual sites only when necessary.

Global Firewall Templates - Firewall templates allow for the configuration of global parameters that impact the operation of the firewall on individual appliances operating in the SD-WAN overlay environment.

Default Firewall Actions – Allow enables packets not matching any filter policy are permitted. Deny enables packets not matching any filter policy are dropped.

Default Connection State Tracking – Enables bidirectional connection state tracking for TCP, UDP, and ICMP flows that do not match a filter policy or NAT rule. Asymmetric flows will be blocked when this is enabled even when there are no Firewall policies defined. The settings may be defined at the site level which will override the global setting. If there is a possibility of asymmetric flows at a site, the recommendation is to enable this at a site or policy level and not globally.

Zones - Firewall zones define logical security grouping of networks connected to the NetScaler SD-WAN. Zones can be applied to Virtual Interfaces, Intranet Services, GRE Tunnels, and LAN IPsec Tunnels.

WAN Link Security Zone

Untrusted security zone should be configured on WAN links directly connected to a public (unsecure) network. Untrusted will set the WAN link to its most secure state, allowing only encrypted, authenticated and authorized traffic to be accepted on the interface group. ARP and ICMP to the Virtual IP Address are the only other traffic type allowed. This setting will also ensure that only encrypted traffic will be send out of the interfaces associated with the Interface group.

Routing Domains

Routing Domains are network systems that include a set of routers that are used to segment network traffic. Newly created sires are automatically associated with the default Routing Domain.

Configuration Editor > **Advanced** > **Global**

Routing Domains

- Default_RoutingDomain

IPsec Tunnels

- Default Sets
- Secure Virtual Path User Data with IPsec

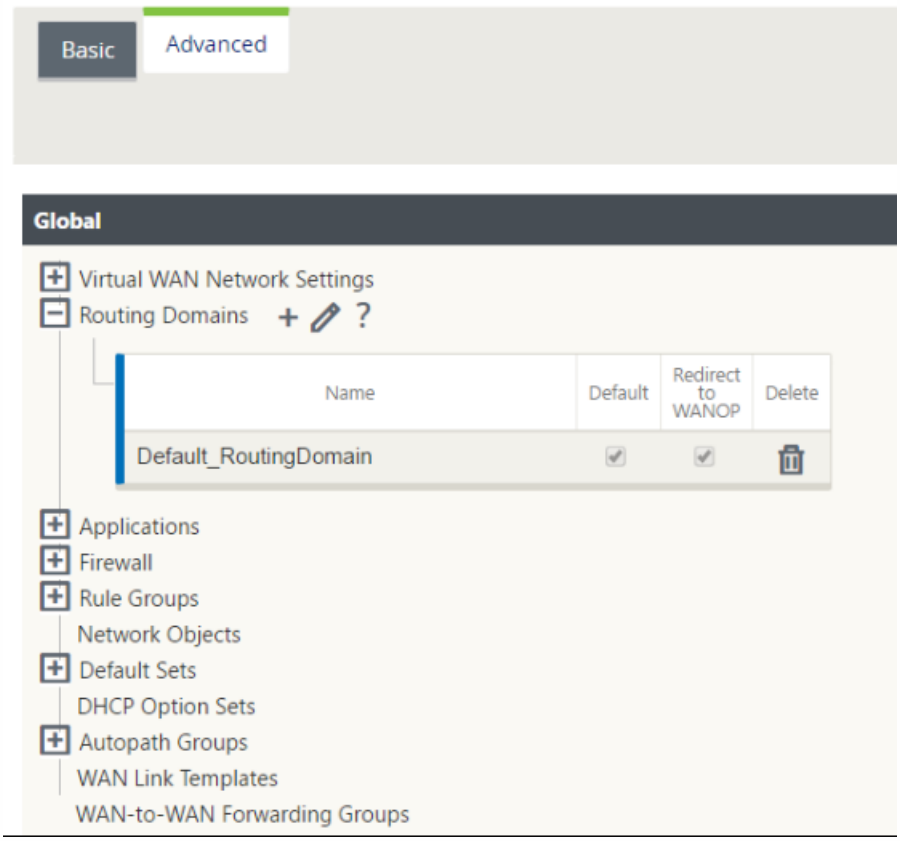

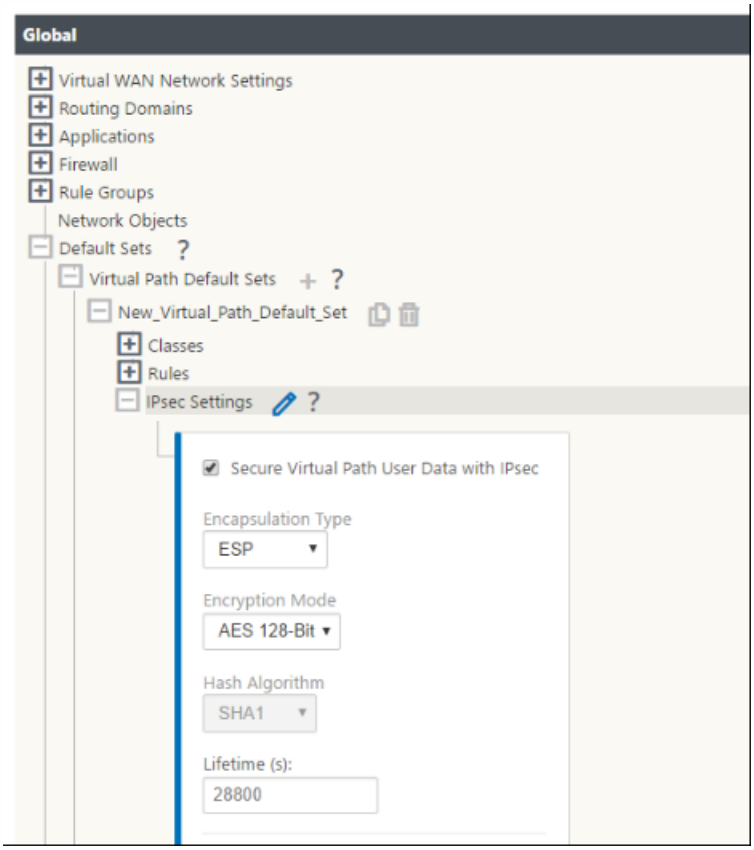

IPSec Tunnels

IPsec Tunnels secure both user data and header information. NetScaler SD-WAN appliances can negotiate fixed IPsec tunnels on the LAN or WAN side with non-SD-WAN peers. For IPsec Tunnels over LAN, a Routing Domain must be

selected. If the IPsec Tunnel uses an Intranet Service, the Routing Domain is pre-determined by the chosen Intranet Service.

IPsec tunnel is established across the Virtual Path before data can flow across the SD-WAN overlay network.

- Tunnel Mode options include ESP - data is encapsulated and encrypted, ESP+Auth – data is encapsulated, encrypted, and validated with an HMAC, AH – data is validated with an HMAC.
- Encryption Mode is the encryption algorithm used when ESP is enabled.
- Hash Algorithm is used to generate an HMAC.
- Lifetime is a preferred duration, in seconds, for an IPsec security association to exist. 0 can be used for unlimited.

IKE Settings

Internet Key Exchange (IKE) is an IPsec protocol used to create a security association (SA). NetScaler SD-WAN appliances support both IKEv1 and IKEv2 protocols.

- Mode can be either Main Mode or Aggressive Mode.
- Identity can be automatic to identify peer, or an IP address can be used to manually specify peer's IP address.
- Authentication enables Pre-Shared Key authentication or certificate as the method of authentication.
- Validate Peer Identity enables validation of the IKE's Peer Identity if the peer's ID type is supported, otherwise do not enable this feature.
- Diffie-Hellman Groups are available for IKE key generation with group 1 at 768-bit, group 2 at 1024-bit, and group 5 at 1536-bit group.
- Hash Algorithm include MD5, SHA1, and SHA-256 has algorithms are available for IKE messages.
- Encryption Modes include AES-128, AES-192, and AES-256 encryption modes are available for IKE messages.
- IKEv2 settings include Peer Authentication and Integrity Algorithm.

Reference Material

Aug 09, 2017

XLSX Application Signature Library

A list of applications that the SD-WAN appliance can identify using Deep Packet Inspection.

Hardware platforms

Sep 07, 2017

The following sections describe the hardware specifications, installation, and initial configuration for all NetScaler SD-WAN hardware platforms:

Hardware Platforms	Describes the NetScaler hardware platforms and provides detailed information about each platform and its components.
Preparing for Installation	Describes how to unpack the NetScaler appliance and prepare the site and rack for installing the appliance. Lists the cautions and warnings that you should review before you install the appliance.
Installing the Hardware	Describes the steps to install the rails, mount the hardware, connect the cables, and turn on the appliance.
Initial Configuration	Describes how to perform initial configuration of your NetScaler appliance and assign management and network IP addresses.

Citrix NetScaler SD-WAN hardware platforms:

- SD-WAN WANOP 400, 800, 1000, 2000, and 3000
- SD-WAN WANOP 1000 WS and 2000 WS
- SD-WAN WANOP 4000 and 5000
- SD-WAN WANOP 4100 and 5100
- SD-WAN Standard Edition 400 and 410
- SD-WAN Standard Edition 1000, 2000, and 2100
- SD-WAN Standard Edition 4000, 4100, and 5100
- SD-WAN Enterprise Edition 1000 and 2000
- SD-WAN VPX Models

The Citrix compliance regulatory models for the appliance editions are:

- SD-WAN 400, 800, 1000: NS 504-2
- SD-WAN 410 Standard Edition: 512-2
- SD-WAN 2000 (all editions): NS 6xCu
- SD-WAN 2100 Standard Edition: 1U1P1A
- SD-WAN 3000 (all editions): NS 6xCu 6xSFP
- SD-WAN 4000 (all editions): 4x10GE SFP+ 8xSFP
- SD-WAN 5000 (all editions): 8x10GE SFP+ 96GB
- SD-WAN 4100 (all editions): 2U1P1B
- SD-WAN 5100 (all editions): 2U1P1D

Common Hardware Components

Aug 09, 2017

Each platform has front panel and back panel hardware components. The various hardware components on the front panel and back panel vary by hardware platform.

On some platforms, the front panel has an LCD display and an RS232 serial console port. The number, type, and location of ports—copper Ethernet, copper and fiber 1G SFP, 10G SFP+, and XFP—vary by hardware platform. The back panel provides access to the fan and the field replaceable units (power supplies, solid-state and hard-disk drives).

This document includes the following details:

LCD Display and LED Status Indicators

The LCD display on the front of every appliance displays messages about the current operating status of the appliance. These messages communicate whether your appliance has started properly and is operating normally. If the appliance is not operating normally, the LCD displays troubleshooting messages.

The LCD displays real-time statistics, diagnostic information, and active alerts. The dimensions of the LCD limit the display to two lines of 16 characters each, causing the displayed information to flow through a sequence of screens. Each screen shows information about a specific function.

The LCD has a neon backlight. Normally, the backlight glows steadily. When there is an active alert, it blinks rapidly. If the alert information exceeds the LCD screen size, the backlight blinks at the beginning of each display screen.When the appliance shuts down, the backlight remains on for one minute and then automatically turns off.

On the appliance's back panel, system status LEDs indicate the overall status of the appliance. The following table describes the indicators of the system status LED.

Note: System status LEDs are available on only the following appliances.

LED Color	LED Indicates
OFF	No power
Green	Appliance is receiving power
Red	Appliance has detected an error

The port LEDs show whether a link is established and traffic is flowing through the port. The following table describes the LED indicators for each port. There are two LED indicators for each port type.

Note: This section applies to the following appliances:

NetScaler SD-WAN WANOP 400, 800, 1000, 2000, and 3000, 4000, 5000, 1000 WS, 2000 WS, VPX 2, 6, 10, 20, 50, 100, 200,

NetScaler SD-WAN Standard Edition 400, 410, 1000, 2000, 4000, 5100, VPX 10, 20, 50, 100, NetScaler SD-WAN Enterprise Edition 1000 and 2000.

Table 1. LED port-status indicators

Port Type	LED Location	LED Function	LED Color	LED Indicates
Ethernet (RJ45)	Left	Speed	Off	No connection, or a traffic rate of 10 megabits per second (Mbps).
			Green	Traffic rate of 100 Mbps.
			Yellow	Traffic rate of 1 gigabit per second.
	Right	Link/ Activity	Off	No link.
			Solid green	Link is established but no traffic is passing through the port.
			Blinking green	Traffic is passing through the port.
Management (RJ45)	Left	Speed	Off	No connection, or a traffic rate of 10 megabits per second (Mbps).
			Green	Traffic rate of 100 Mbps.
			Amber	Traffic rate of 1 gigabit per second.
	Right	Link/ Activity	Off	No link.
			Solid yellow	Link is established but no traffic is passing through the port.
			Blinking yellow	Traffic is passing through the port.

On each power supply, a bicolor LED indicator shows the condition of the power supply. The LEDs of the AC power supplies for each appliance are different from the LEDs of the other appliances.

Table 2. LED Power Supply Indicators

Power Supply Type	LED Color	LED Indicates
AC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing GREEN	Power supply is in standby mode.
	GREEN	Power supply is functional.
	RED	Power supply failure.
DC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing BLUE	Power supply is in standby mode.
	BLUE	Power supply is functional.
	RED	Power supply failure.

Ports

Aug 09, 2017

Note: Some SD-WAN appliances do not require SFP transceivers.

Ports are used to connect the appliance to external devices. Citrix NetScaler SD-WAN appliances support RS232 serial ports, 10/100/1000Base-T copper Ethernet ports, fiber 1G SFP ports and 10-gigabit fiber SFP+ ports. All Citrix NetScaler SD-WAN appliances have a combination of some or all of these ports. For details on the type and number of ports available on your appliance, see the section describing that platform.

RS232 Serial Port

The RS232 serial console port provides a connection between the appliance and a computer, allowing direct access to the appliance for initial configuration or troubleshooting.

All hardware platforms ship with an appropriate serial cable used to connect your computer to the appliance. For instructions on connecting your computer to the appliance, see Installing the Hardware.

Copper Ethernet Ports

The copper Ethernet ports installed on many models of the appliance are standard RJ45 ports.

There are two types of copper Ethernet ports that may be installed on your appliance:

10/100BASE-T portThe 10/100BASE-T port has a maximum transmission speed of 100 megabits per second (Mbps). Most platforms have at least one 10/100BASE-T port.

10/100/1000BASE-T portThe 10/100/1000BASE-T port has a maximum transmission speed of 1 gigabit per second, ten times faster than the other type of copper Ethernet port. Most platforms have at least one 10/100/1000Base-T port.

To connect any of these ports to your network, you plug one end of a standard Ethernet cable into the port and plug the other end into the appropriate network connector.

Management Ports

Management ports are standard copper Ethernet ports (RJ45), which are used for direct access to the appliance for system administration functions.

1G SFP and 10G SFP+ Ports

A 1G SFP port can operate at a speed of 1 Gbps. It accepts either a copper 1G SFP transceiver, for operation as a copper Ethernet port, or a fiber 1G SFP transceiver for operation as a fiber optic port.

The 10G SFP+ ports are high-speed ports that can operate at speeds of up to 10 Gbps. You need a fiber optic cable to connect to a 10G SFP+ port. If the other end of the fiber optic cable is attached to a 1G SFP port, the 10G SFP+ port automatically negotiates to match the speed of the 1G SFP port.

This Netscaler SD-WAN 410-SE appliance can be used as a WAN Optimization device, then the first port pair should have apA labeled as well.

The motherboard port should be labeled MGMT for port Eth0.

Field Replaceable Units

Aug 09, 2017

Citrix NetScaler SD-WAN field replaceable units (FRU) are SD-WAN components that can be quickly and easily removed from the appliance and replaced by the user or a technician at the user's site. The FRUs in a SD-WAN appliance can include an AC power supply and a solid-state drive. The solid-state drive stores your configuration information used to restore from a backup after replacing the unit.

> **Note**
>
> SD-WAN Standard Edition 400 and 410 appliances do not have field replaceable units. The field replaceable SSD and power supplies are not required.

SD-WAN WANOP/SE 4000 and WANOP 5000 field replaceable units (FRU) are components that can be quickly and easily removed from the appliance and replaced by the user or a technician at the user's site. The FRUs in a SD-WAN WANOP/SE 4000 and WANOP 5000 appliance can include DC or AC power supplies, and solid-state and hard-disk drives.

Power Supply

Aug 09, 2017

SD-WAN appliances are configured with a single power supply. For a SD-WAN 3000 WANOP appliance, you can order a second power supply.

SD-WAN 4000/5000 WANOP and 5100 SE appliances are configured with dual power supplies but can operate with only one power supply. The second power supply serves as a backup.

For a SD-WAN Standard Edition 410 appliance, a single chassis power switch is supplied. The device has an external power brick instead of an internal power supply if a desktop form factor is chosen.

The appliances are shipped with a standard power cord that plugs into the appliance's power supply. The other end of the cord has a NEMA 5-15 plug on the other end for connecting to the power outlet on the rack or in the wall.

For power-supply specifications, see Common Hardware Components, which describes the various hardware components, hardware platforms and includes a table summarizing the hardware specifications.

Note: If you suspect that a power-supply fan is not working, see the description of your platform. On some platforms, what appears to be the fan does not turn, and the actual fan turns only when necessary.

Table 1. LED Power Supply Indicators

Power Supply Type	LED Color	LED Indicates
AC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing GREEN	Power supply is in standby mode.
	GREEN	Power supply is functional.
	RED	Power supply failure.
DC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing BLUE	Power supply is in standby mode.
	BLUE	Power supply is functional.
	RED	Power supply failure.

Electrical Safety Precautions for Power Supply Replacement

- Make sure that the appliance has a direct physical connection to earth ground during normal use. When installing or repairing an appliance, always connect the ground circuit first and disconnect it last.

- Always unplug any appliance before performing repairs or upgrades.
- Never touch a power supply when the power cord is plugged in. As long as the power cord is plugged in, line voltages are present in the power supply even if the power switch is turned off.

Replacing an AC Power Supply

A NetScaler SD-WAN 2000 appliance can accommodate only one power supply, which is not field replaceable. A NetScaler 3000 appliance has only one power supply, but you can order and install a second power supply.

Note: Shut down the appliance before replacing the power supply.

To install or replace an AC power supply in a SD-WAN 3000 appliance

1. If replacing an existing power supply, align the semicircular handle, so that it is perpendicular to the power supply, loosen the thumbscrew, press the lever toward the handle and pull out the existing power supply, as shown in the following figure.
   Note: The illustration in the following figures might not represent the actual SD-WAN appliance.
   Figure 1. Removing the Existing AC Power Supply
2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot and press against the semicircular handle until you hear the power supply snap into place.
   Figure 2. Inserting the Replacement AC Power Supply
5. Connect the power supply to a power source.

Replacing AC Power Supply on SD-WAN 4000/5000 appliance

Replace an AC power supply with another AC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace an AC power supply on a SD-WAN 4000/5000 appliance

1. Align the semicircular handle perpendicular to the power supply. Loosen the thumbscrew and press the lever toward the handle and pull out the existing power supply, as shown in the following figure.

   Figure 1. Removing the Existing AC Power Supply
2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.

4. Insert the power supply into the slot and press against the semicircular handle until you hear the power supply snap into place.

   Figure 2. Inserting the Replacement AC Power Supply

5. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Replacing DC Power Support on SD-WAN 4000/5000 appliance

Replace a DC power supply with another DC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace a DC power supply on a SD-WAN 4000/5000 appliance

1. Loosen the thumbscrew and press the lever towards the handle and pull out the existing power supply, as shown in the following figure.

   Figure 3. Removing the Existing DC Power Supply

- Carefully remove the new power supply from its box.
- On the back of the appliance, align the power supply with the power supply slot.
- Insert the power supply into the slot while pressing the lever towards the handle. Apply firm pressure to insert the power supply firmly into the slot.

Figure 4. Inserting the Replacement DC Power Supply

- When the power supply is completely inserted into its slot, release the lever.
- Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power

cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Solid-State Drive

Aug 09, 2017

A solid-state drive (SSD) is a high-performance data storage device that stores data in solid-state flash memory. It stores the SD-WAN software and user data.

For SD-WAN 410 appliance, the on-board SATA disk controller must support at least two devices. Support for SATAv3 (6 Gbps) is available.

Replacing a Solid-State Drive

Replacement solid-state drives (SSDs) contain a pre-installed version of the SD-WAN software.

To replace a Solid-State Drive

1. Shut down the appliance.
2. Locate the SSD on the back panel of the appliance. Push the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Note: The illustration in the following figures might not represent the actual SD-WAN appliance.
   Figure 1. Removing the Existing Solid-State Drive

3. Verify that the replacement SSD is the correct type for the platform.
4. Pick up the new SSD, open the drive handle fully to the left, and insert the drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top if the drive is inserted horizontally or at the right if the drive is inserted vertically.
   Figure 2. Inserting the Replacement Solid-State Drive

5. Turn on the appliance.
6. Perform the initial configuration of the appliance.

Hard Disk Drive

Aug 09, 2017

The NetScaler SD-WAN virtual machines are hosted on the hard-disk drive.

Replacing a Hard Disk Drive

Verify that the replacement hard disk drive is the correct type for the NetScaler SD-WAN WANOP/SE 4000, WANOP 5000, and 5100 SE platforms.

To install a hard disk drive

1. Shut down the appliance.
2. Locate the hard disk drive on the back panel of the appliance.
3. Disengage the hard disk drive by pushing the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Figure 1. Removing the Existing Hard Disk Drive

4. Pick up the new disk drive, open the drive handle fully to the left, and insert the new drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the hard drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top.
   Figure 2. Inserting the Replacement Hard Disk Drive

5. Turn on the appliance.

NetScaler SD-WAN 400, 800, 1000, 2000 and 3000 WANOP Appliances

Aug 09, 2017

The SD-WAN 400, 800, 1000, 2000 and 3000 appliances are 1U accelerators for use in datacenters and larger branch offices.

The SD-WAN 2000 can be thought of as a faster Repeater 8500 appliance with two accelerated bridges, while the SD-WAN WANOP 3000 can be thought of as a faster Repeater 8800 with three accelerated bridges. The configuration process, however, is not the same. Like the high-end Repeater SDX appliance, SD-WAN 2000 and WANOP 3000 appliances use virtual machines for acceleration and management, running under a XenServer hypervisor.

The SD-WAN 400, 800, 1000, 2000, and 3000 Series have the following general characteristics:

- SD-WAN 400 Series. An small, affordable 1U appliance suitable for smaller branch offices, the 400 Series has two accelerated bridges and supports WAN speeds of up to 6 Mbps.
- SD-WAN 800 Series. A small 1U appliance suitable for medium-sized branch offices, the 800 Series has two accelerated bridges and supports WAN speed of up to 10 Mbps.
- SD-WAN 2000 Series. A full-sized 1U appliance suitable for large branch offices and smaller datacenters, the 2000 Series has two accelerated bridges and supports WAN speed of 10-50Mbps.
- Sd-WAN 3000 Series. A full-sized 1U appliance suitable for the largest branch offices and medium-sized datacenters, the 3000 Series has three accelerated bridges and supports WAN speed of 50-155 Mbps.

The Citrix Compliance Regulatory Models are as follows:

- SD-WAN 400 WANOP: CB 504-2
- SD-WAN 800 WANOP: CB 504-2
- SD-WAN 1000 WANOP: CB 504-2
- SD-WAN 2000 WANOP: NS 6xCu
- SD-WAN 3000 WANOP: NS 6xCu 6xSFP

All SD-WAN platforms have similar components and hardware platforms offer a wide range of features, communication ports, and processing capacities. All platforms support the SD-WAN software and have multicore processors.
These appliances have similar architectures, run the same release binaries and are fully supported with release 9.2.

NetScaler SD-WAN 400 and 800

Aug 09, 2017

The Citrix NetScaler SD-WAN 400 and 800 platforms each have a dual-core processor and 8GB of memory. These platforms have a bandwidth of up to 6 Mbps and up to 10 Mbps, respectively.

The following figure shows the front panel of a SD-WAN 400/800 appliance.

Figure 1. Citrix NetScaler SD-WAN 400/800, front panel

The front panel of the NetScaler SD-WAN 400/800 appliance has a power button and five LEDs.

The power button switches main power (the power to the power supply) on or off.

The reset button restarts the appliance.

The LEDs provide critical information about different parts of the appliance.

- Power Fail—Indicates that a power supply unit has failed.
- Information LED—Indicates the following:

Status	Description
Continuously on and red	The appliance is overheated. (This might be a result of cable congestion.)
Blinking red (1Hz)	Fan failure.
Blinking red (0.25Hz)	Power failure.
Solid blue	Local UID has been activated. Use this function to locate the server in a rack mount environment.
Blinking blue (300 m/s)	Remote UID is on. Use this function to identify the server from a remote location.

- NIC1 and NIC2—Indicate network activity on the LAN1 and WAN1 ports.
- HDD—Indicates the status of the hard disk drive.

- Power—When blinking, indicates that the power supply unit is receiving power and operating normally.

The following figure shows the back panel of a SD-WAN 400/800 appliance.

Figure 2. Citrix NetScaler SD-WAN 400/800 appliance, back panel

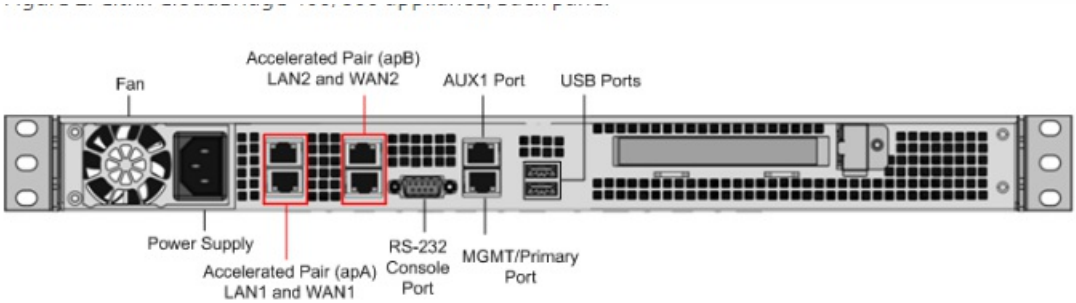

The following components are visible on the back panel of a SD-WAN 400/800 appliance:

- Cooling fan
- Single power supply, rated at 200 watts, 110-240 volts
- Accelerated pairs of Ethernet ports (apA and apB) which function as accelerated bridges. Individual port assignments: LAN1 is apA.1, WAN1 is apA.2, LAN2 is apB.1, LAN2 is apB.2.
- RS-232 serial console port
- One Aux Ethernet port and one management port
- Two USB ports
- One Solid State Drive (SSD)
  - SD-WAN 400 - 160 GB SSD
  - SD-WAN 800 - 240 GB SSD

NetScaler SD-WAN 1000

Mar 19, 2018

The Citrix NetScaler SD-WAN 1000 platform has 3 models: SD-WAN 1000-06, SD-WAN 1000-010, and SD-WAN 1000-020, with bandwidths of 6Mbps, 10Mbps, and 20Mbps, respectively. Each model is a 1U appliance with one quad-core processor and 24 gigabytes (GB) of memory.

The following figure shows the front panel of the SD-WAN 1000 appliance.

The front panel of the NetScaler SD-WAN 1000 appliance has a power button and five LEDs.

- The power button switches main power (the power to the power supply) on or off.
- The reset button restarts the appliance.
- The LEDs provide critical information about different parts of the appliance.

Figure 1. Citrix NetScaler SD-WAN 1000, front panel

The appliance has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) management port, numbered 0/1. The management port is used to connect directly to the appliance for system administration functions.
- Four 10/100/1000Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two accelerated pairs, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), and 1/3 and 1/4 are accelerated pair B (apB).

The following figure shows the back panel of the SD-WAN 1000 appliance.

Figure 2. Citrix SD-WAN 1000 appliance, back panel

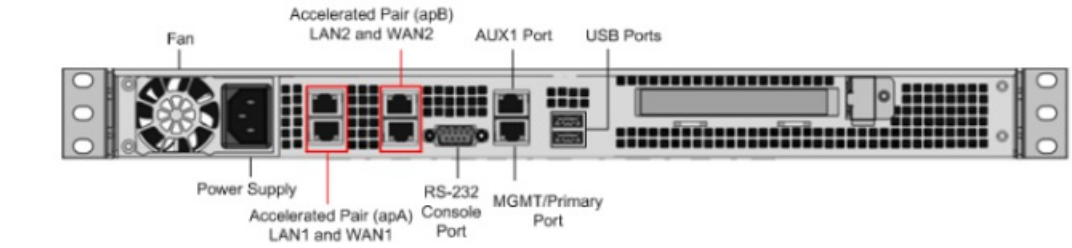

The following components are visible on the back panel of the SD-WAN 1000 appliance:

- 600 GB removable solid-state drive, which stores the appliance's software and user data.
- USB port (reserved for a future release).
- Single power supply, rated at 300 watts, 100-240 volts.

Power on Appliance After a Graceful Shut Down

To power on the appliance after a graceful shut down:

1. Connect a Serial console cable to the rear of the appliance and to the serial port on a management laptop.

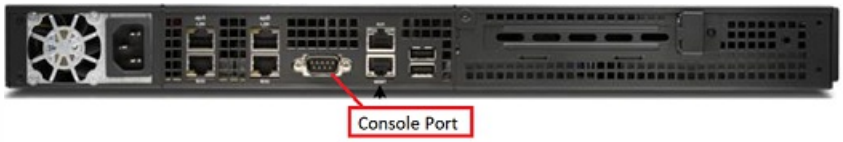

2. On the management laptop, restart a putty session using the following configuration settings:

- Serial line: COM1
- Speed: 9600

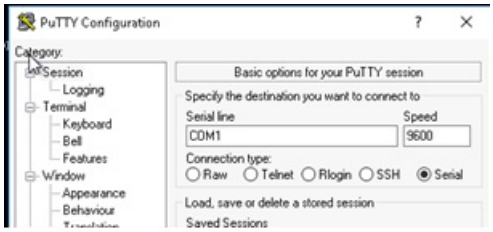

3. Power on the appliance and as it is booting, press the following key in the Putty session to enter the BIOS configuration screen.

   Keypress: **DEL**

4. When in the BIOS, navigate to,

- Advanced Tab > **Select**
- Boot Feature > **Enter**

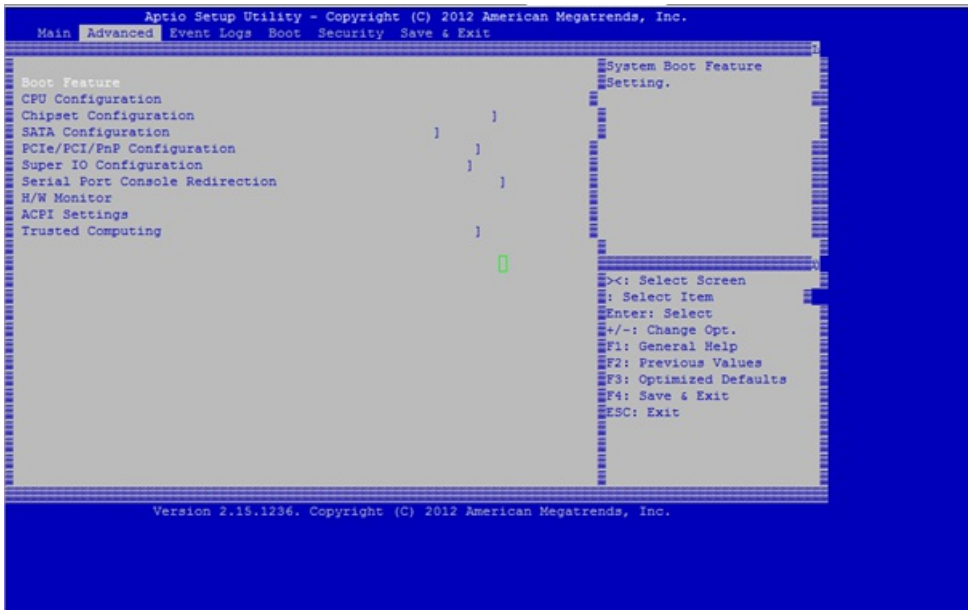

5. When in the Boot Feature screen, change the value of the parameter **Restore on AC Power Loss;** from **Last State** > **Power ON**.

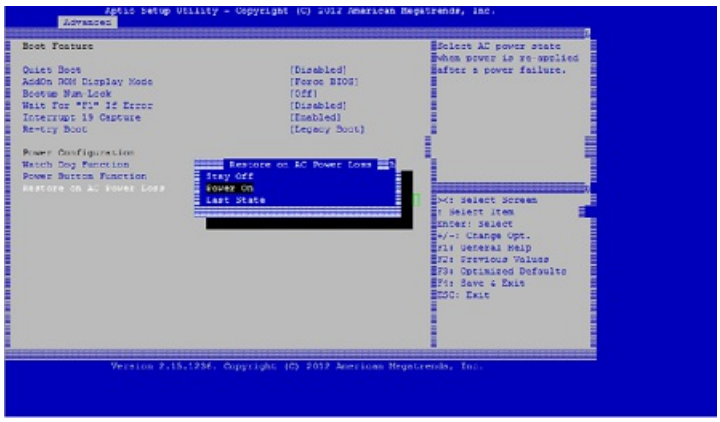

6. Navigate to Save and Exit.

- Select **Save changes** and **Reset**
- Select **Yes**

Allow the system to restart. This takes approximately five minutes.

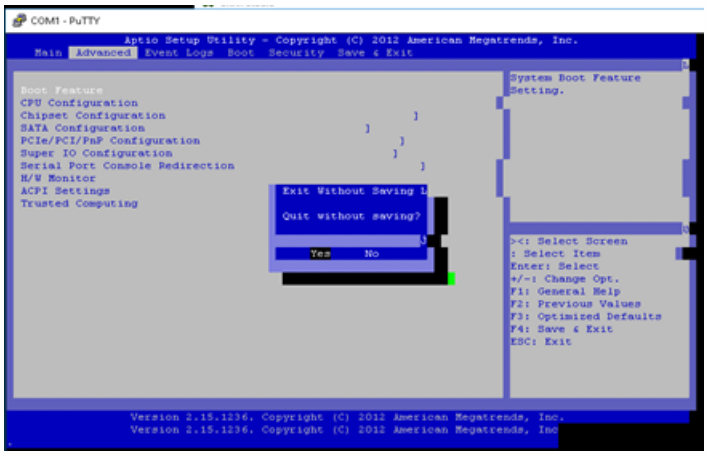

7. After the appliance is powered on, login to the appliance management instance (SVM). The default IP address for

the appliance is: 192.168.100.1, user name is: admin/password.

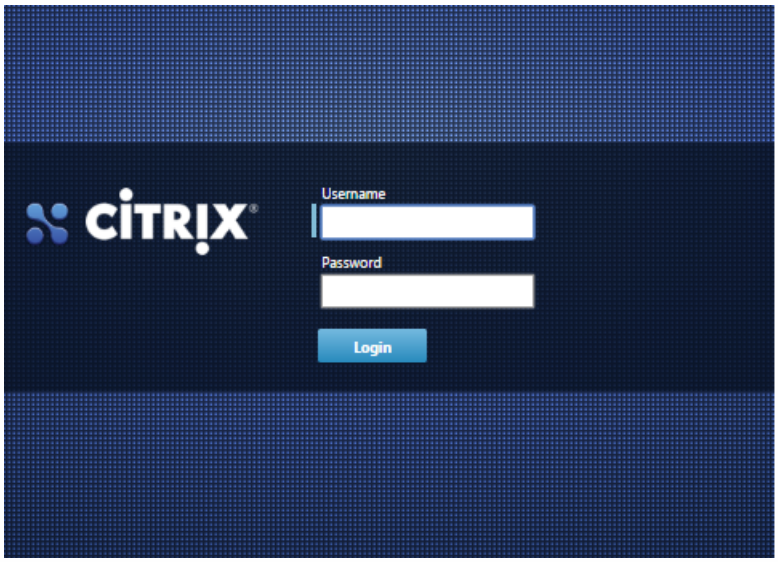

8. In the SD-WAN appliance GUI, navigate to **Configuration** > **Maintenance** > **Reboot Appliance**. Allow the appliance to fully shut down. Ensure that there are no power lights on the appliance when the shut down process has completed.

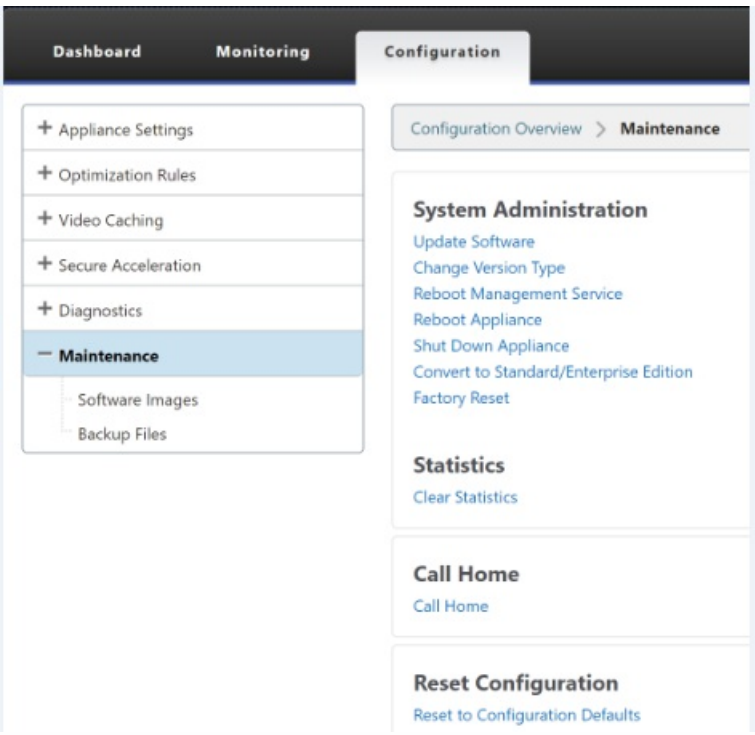

9. Power on the appliance to confirm that the BIOS configuration change has been applied successfully. This can be either done through the APC intelligent PDU Web Management console or by physically pulling the power cable out of the shut down SD-WAN appliance, waiting for 10 seconds and then plugging it back in again. The appliance power ups automatically from all shut down scenarios.

NetScaler SD-WAN 2000

Aug 09, 2017

The Citrix NetScaler SD-WAN 2000 platform has 3 models: SD-WAN 2000-010, SD-WAN 2000-020, and SD-WAN 2000-050, with bandwidths of 10Mbps, 20Mbps, and 50Mbps, respectively. Each model is a 1U appliance with one quad-core processor and 24 gigabytes (GB) of memory.

The following figure shows the front panel of the SD-WAN 2000 appliance.

Figure 1. Citrix NetScaler SD-WAN 2000, front panel

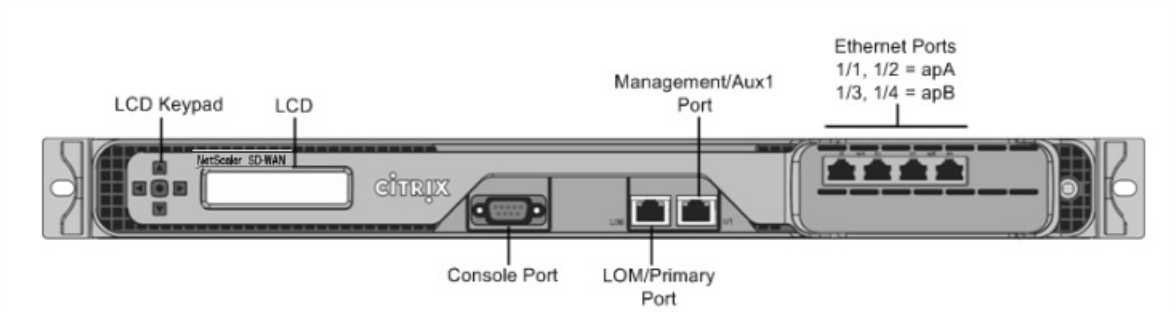

The appliance has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, numbered 0/1. The management port is used to connect directly to the appliance for system administration functions.
  Note: The LOM port also operates as a management port.
- Four 10/100/1000Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two accelerated pairs, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), and 1/3 and 1/4 are accelerated pair B (apB).

The following figure shows the back panel of the SD-WAN 2000 appliance.

Figure 2. Citrix SD-WAN 2000 appliance, back panel

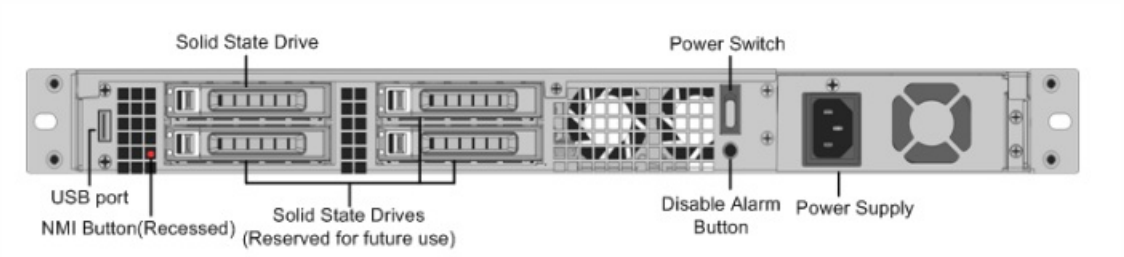

The following components are visible on the back panel of the SD-WAN 2000 appliance:

- 600 GB removable solid-state drive, which stores the appliance's software and user data.
- Power switch, which turns off power to the appliance. Press the switch for five seconds to turn off the power.
- USB port (reserved for a future release).

- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You must use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Single power supply, rated at 300 watts, 100-240 volts.

NetScaler SD-WAN 3000

Aug 09, 2017

The Citrix NetScaler SD-WAN 3000 platform has 3 models: SD-WAN 3000-050, SD-WAN 3000-100, and SD-WAN 3000-155, with bandwidths of 50M bps, 100 Mbps, and 155 Mbps, respectively. Each model is a 1U appliance with one quad-core processor and 32 gigabytes (GB) of memory.

The Citrix NetScaler SD-WAN 3000 appliance is available in two port configurations:

- Six 10/100/1000 Base-T copper Ethernet ports
- Four 1G SX Fiber ports

The following figure shows the front panel of a SD-WAN 3000 with six 10/100/1000 Base-T copper Ethernet ports.

Figure 1. Citrix NetScaler SD-WAN 3000 (6×10/100/1000 Base-T copper Ethernet ports), front panel

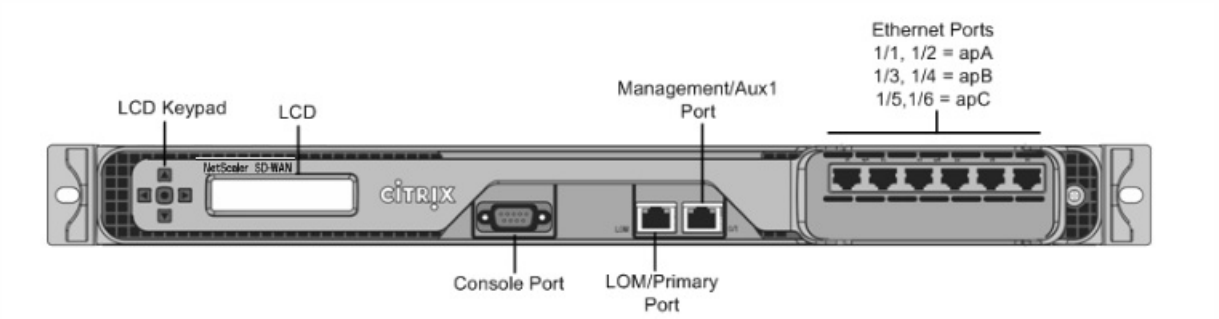

The following figure shows the front panel of a SD-WAN 3000 appliance with four 1G SX fiber ports.

Figure 2. Citrix NetScaler SD-WAN 3000 (4×1G SX Fiber ports), front panel

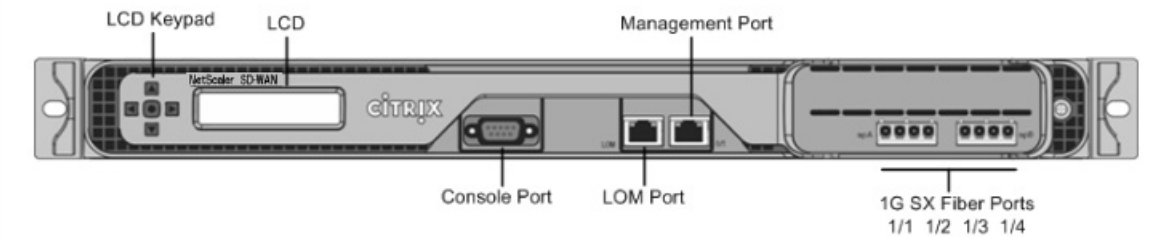

The appliance has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, numbered 0/1. The management port is used to connect directly to the appliance for system administration functions.
  Note: The LOM port also operates as a management port.

- Network Ports, in one of the following configurations:
  - SD-WAN 3000 (6x10/100/1000 Base-T copper Ethernet ports). Six 10/100/1000 Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, 1/4, 1/5, and 1/6 from left to right. The six ports form three accelerated pairs, which function

as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), 1/3 and 1/4 are accelerated pair B (apB), and 1/5 and 1/6 are accelerated pair C (apC).

- SD-WAN 3000 (4x 1G SX Fiber ports). Four 1G SX fiber ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two accelerated pairs, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA) and 1/3 and 1/4 are accelerated pair B (apB).

The following figure shows the back panel of the SD-WAN 3000 appliance.

Figure 3. Citrix NetScaler SD-WAN 3000 appliance, back panel

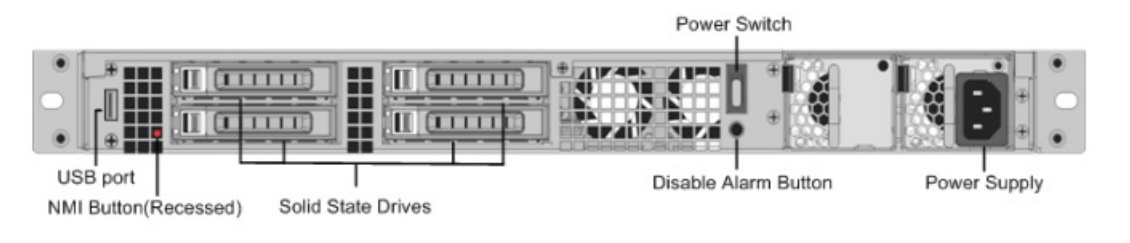

The following components are visible on the back panel of the SD-WAN 3000 appliance:

- Four 600 GB removable solid-state drives. The top left solid-state drive stores both the appliance's software and the user data. The other three store only user data.
- Power switch, which turns power to the appliance on or off. To turn off the power, press the switch for five seconds.
- USB port (reserved for a future release).
- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You must use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Disable alarm button, which is nonfunctional unless you install a second power supply. In that case, it disables the alarm that sounds if the appliance is plugged into only one power outlet or one of the power supplies fails.
- Single power supply, rated at 450 watts, 100-240 volts.

Summary of Hardware Specifications

Aug 09, 2017

The following table summarizes the specifications of theNetScaler SD-WAN 400, 800, 1000, 2000, and 3000 hardware platforms.

Table 1. Citrix NetScaler SD-WAN 400, 800, 2000, 1000, and 3000 Platforms Summary

H/W Specifications	SD-WAN 400	SD-WAN 800	SD-WAN 1000	SD-WAN 2000	SD-WAN 3000
Platform Performance					
Bandwidth	Up to 6 Mbps Model 400-002: 2Mbps Model 400-006: 6 Mbps	Up to 10 Mbps Model 800-002: 2Mbps Model 800-006: 6 Mbps Model 800-010: 10 Mbps	Model 1000-006: 6 Mbps Model 1000-010: 10 Mbps Model 1000-020: 20 Mbps	Model 2000-010: 10 Mbps Model 2000-020: 20 Mbps Model 2000-050: 50 Mbps	Model 3000-050: 50 Mbps Model 3000-100: 100 Mbps Model 3000-155: 155 Mbps
Maximum HDX sessions	Up to 60	Up to 100	200	300	500
Total sessions	500	10,000	10,000	20,000	50,000
Acceleration Plug-in CCUs	NA	NA	NA	750	1,000
Hardware Specifications					
Processor	2 Cores	2 Cores	2 Cores	4 Cores	4 Cores
Total disk space	1 x 160 GB SSD	1 x 240 GB SSD		1 x 600 GB SSD	4 x 600 GB SSD
SSD (dedicated Compression history)	40 GB	80 GB		275 GB	1.5 TB

H/W Specifications	SD-WAN 400	SD-WAN 800	SD-WAN 1000	SD-WAN 2000	SD-WAN 3000
RAM	8 GB	8 GB		24 GB	32 GB
Network Interfaces	2 pair with bypass 10/100/1000	2 pair with bypass 10/100/1000		4 x 10/100/1000 Base-T copper Ethernet	6 x 10/100/1000 Base-T copper Ethernet
Transceiver support	No	No		Yes	Yes
Power supplies	1	1		1	1
Physical Dimensions					
Rack Units	1U	1U		1U	1U
System width	EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks		EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks
System depth	10.5" (26.7 cm)	10.5" (26.7 cm)		25.4" (64.5 cm)	25.4" (64.5 cm)
System weight	8 lbs (3.5 kg)	8 lbs (3.5 kg)		32 lbs (14.5 kg)	32 lbs (14.5 kg)
Shipping dimensions and weight	26L x 18.5W x 6.5" H 14 lbs	26L x 18.5W x 6.5" H 14 lbs		32L x 23.5W x 7.5" H 39 lbs	32L x 23.5W x 7.5" H 39 lbs
Environmental and Regulatory					
Voltage	100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz		100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz
Power consumption (Max.)	200W	200W		300 W	450 W
Operating Temperature (degree Celsius)	10–35	10–35		0–40	0–40

H/W Specifications	SD-WAN 400	SD-WAN 800	SD-WAN 1000	SD-WAN 2000	SD-WAN 3000
Non-operating Temperature (degree Celsius)	-40–70	-40–70	-40–70	-40–70	-40–70
Allowed Relative Humidity	8%–90%	8%–90%		5%–95%	5%–95%
Safety certifications	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)		CSA	TUV
Electromagnetic and susceptibility certifications	FCC (Part 15 Class A), EN 55022 Class A, EN 61000-3-2/-3-3, CISPR 22 Class A	FCC (Part 15 Class A), EN 55022 Class A, EN 61000-3-2/-3-3, CISPR 22 Class A	FCC (Part 15 Class A), EN 55022 Class A, EN 61000-3-2/-3-3, CISPR 22 Class A	FCC (Part 15 Class A), CE, C-Tick, VCCI-A, CCC, KCC, NOM, SASO, SABS, PCT	FCC (Part 15 Class A), CE, C-Tick, VCCI-A, CCC, KCC, NOM, SASO, SABS, PCT
Environmental certifications	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE

Supported Features

Aug 09, 2017

The following table lists various features supported on NetScaler SD-WAN 400, 800, 2000, 1000, and 3000 appliances.

Table 1. Features Table for Citrix NetScaler SD-WAN 400, 800, 2000, 1000, and 3000 Series Appliances

Feature	Citrix NetScaler SD-WAN 400 series	Citrix NetScaler SD-WAN 800 series	Citrix NetScaler SD-WAN 1000 series	Citrix NetScaler SD-WAN 2000 series	Citrix NetScaler SD-WAN 3000 series
AutoConfiguration	Y	Y	Y	Y	Y
SD-WAN Plug-In	N	N	N	Y	Y
Compression	Y	Y	Y	Y	Y
RPC over HTTPS	Y	Y	Y	Y	Y
SSL Compression	Y	Y	Y	Y	Y
TCP Acceleration	Y	Y	Y	Y	Y
Traffic Shaping	Y	Y	Y	Y	Y
Video Caching	N	Y	Y	Y	Y
Windows File System Acceleration	Y	Y	Y	Y	Y
Windows Outlook Acceleration	Y	Y	Y	Y	Y
XenApp/ XenDesktop Acceleration	Y	Y	Y	Y	Y
Group Mode	Y	N	Y	Y	Y
High Availability Mode	Y	Y	Y	Y	Y
Inline Mode	Y	Y	Y	Y	Y
Virtual Inline Mode	Y	Y	Y	Y	Y
WCCP Mode	Y	Y	Y	Y	Y
VLANs	Y	Y	Y	Y	Y

Feature	Citrix NetScaler SD-WAN 400 series	Citrix NetScaler SD-WAN 800 series	Citrix NetScaler SD-WAN 1000 series	Citrix NetScaler SD-WAN 2000 series	Citrix NetScaler SD-WAN 3000 series

Installation

Aug 09, 2017

Within a given series, all models use the same hardware, and the different WAN speed ratings are obtained through different licensing options. For example, both SD-WAN 400 models (the 400-002 and 400-006) use the same hardware, and a given appliance can be licensed either as a 2 Mbps appliance or a 6 Mbps appliance, but a SD-WAN 400 can not be upgraded to a SD-WAN 800, 1000, 2000, or 3000. The same is true of the other series.

The licensed bandwidth applies only to the sending direction, so a SD-WAN 400-002, rated at 2 Mbps in the sending direction, is appropriate for an ADSL link with a 12 Mbps/2 Mbps download/upload bandwidth.

In addition to differences in WAN bandwidth capabilities, the different series vary in CPU power, installed RAM, and installed disk capacity.

All models use solid-state drives instead of conventional hard drives for increased speed and reliability.

Preparing for Installation

Aug 09, 2017

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Once you are satisfied that your appliance has been delivered to your expectations, verify that the location where the appliance will be installed meets temperature and power requirements and that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel should install, maintain, or replace the appliance, and efforts should be taken to ensure that all cautions and warnings are followed.

Unpacking the SD-WAN Appliance

Aug 09, 2017

The hardware accessories for your particular appliance, such as cables, adapters, and rail kit, vary depending on the hardware platform you ordered. Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

If you ordered a SD-WAN 400 or 800 appliance, the box should contain:

- The appliance you ordered
- One DB9 female to DB9 female console port cable
- One 6 ft CAT5 network cable
- One power cable

If you ordered a SD-WAN 1000, 2000 or 3000 appliance, the box should contain:

- The appliance you ordered
- One RJ-45 to DB-9 adapter
- One 6 ft RJ-45/DB-9 cable
- One power cable
- One standard 4-post rail kit
  Note: If the kit that you received does not fit your rack, contact your Citrix sales representative to order the appropriate kit.

In addition to the items included in the box with your new appliance, you will need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you will connect to your network
- One available Ethernet port on your network switch or hub for each Ethernet port you want to connect to your network
- A computer to serve as a management workstation

Preparing the Site and Rack

Aug 09, 2017

A SD-WAN appliance has specific site and rack requirements. You must make sure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and will help ensure a smooth installation.

Site Requirements

The appliance should be installed in a server room or server cabinet with the following features:

Environment control

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment.

Power density

Wiring capable of handling at least 4,000 watts per rack unit in addition to power needs for the CRAC.

Rack Requirements

The rack on which you install your appliance should meet the following criteria:

Rack characteristics

Racks should be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it should be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks should be installed in rows spaced at least 1 meter/3 feet apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

Power connections

At minimum, two standard power outlets per unit.

Network connections

At minimum, four Ethernet connections per rack unit.

Space requirements

One empty rack unit for each SD-WAN 400, 800, 1000, 2000 and 3000 appliances.

Note: You can order the following rail kits separately.

- Compact 4-post rail kit, which fits racks of 23 to 33 inches.
- 2-post rail kit, which fits 2-post racks.

Cautions and Warnings

Aug 09, 2017

Electrical Safety Precautions

Updated: 2014-02-06

Caution: During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.
Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Remove all jewelry and other metal objects that might come into contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.
- Use a regulating, uninterruptible power supply (UPS) to protect the appliance from power surges and voltage spikes, and to keep the appliance operating in case of power failure.
- Never stack the appliance on top of any other server or electronic equipment.
- All appliances are designed to be installed on power systems that use TN earthing. Do not install your device on a power system that uses either TT or IT earthing.
- Make sure that the appliance has a direct physical connection to the earth during normal use. When installing or repairing an appliance, always make sure that the ground circuit is connected first and disconnected last.
- Make sure that a fuse or circuit breaker no larger than 120 VAC, 15 A U.S. (240 VAC, 16 A international) is used on all current-carrying conductors on the power system to which your appliances are connected.
- Do not work alone when working with high voltage components.
- Always disconnect the appliance from power before removing or installing any component. When disconnecting power, first shut down the appliance, and then unplug the power cords of all the power supply units connected to the appliance. As long as the power cord is plugged in, line voltages can be present in the power supply, even when the power switch is OFF.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload. Always unplug any appliance before performing repairs or upgrades.
- Do not overload the wiring in your server cabinet or on your server room rack.
- During thunderstorms, or anticipated thunderstorms, avoid performing any hardware repairs or upgrades until the danger of lightning has passed.
- When you dispose of an old appliance or any components, follow any local and national laws on disposal of electronic waste.
- To prevent possible explosions, replace expired batteries with the same model or a manufacturer-recommended substitute and follow the manufacturer's instructions for battery replacement.
- Never remove a power supply cover or any sealed part that has the following label:

Appliance Precautions

- Determine the placement of each component in the rack before you install the rails.
- Install the heaviest appliance first, at the bottom of the rack, and then work upward. Distribute the load on the rack evenly. An unbalanced rack is hazardous.
- Allow the power supply units and hard drives to cool before touching them.
- Install the equipment near an electrical outlet for easy access.
- Mount equipment in a rack with sufficient airflow for safe operation.
- For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

Rack Precautions

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Oct 12, 2017

A SD-WAN 400, 800, 1000, 2000 or 3000 appliance requires one rack unit. These appliances are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

To mount the appliance, you must first install the rails and then install the appliance in the rack, as follows:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

To remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the locking tabs until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

To attach the inner rails to the appliance

1. Position the right inner rail behind the ear bracket on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws.
4. Repeat steps 1 through 3 to install the left inner rail on the left side of the appliance.

To install the rack rails

1. Position the rack rails at the desired location in the rack, keeping the sliding rail guide facing inward.
2. Snap the rails to the rack. Make sure that both rack rails are at same height and that the rail guides are facing inward.

To install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides, and push the appliance into the rack rails until it locks into place.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.
   Figure 1. Rack Mounting the Appliance

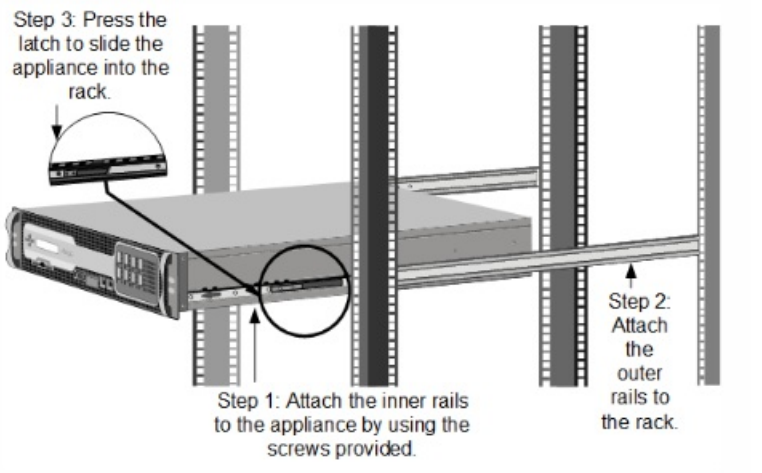

Note

The above illustration might not represent your actual appliance.

Connecting the Cables

Oct 12, 2017

When the appliance is securely mounted on the rack, determine which ports you should use. You are then ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

> ### Warning
>
> Remove all jewelry and other metal objects that might come in contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly, and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ports

A typical installation using a single accelerated-bridge uses three Ethernet ports (the Primary port and apA) and five IP addresses (three on the Primary port's subnet and two on apA's subnet).

The appliance has two motherboard ports and two (SD-WAN 400/800 and SD-WAN 1000/2000) or three (SD-WAN 3000) accelerated bridges. The Primary port is used for initial configuration.

On the SD-WAN 2000 and 3000 appliances, the motherboard ports are labeled "0/1" and "0/2." They are equivalent to the Primary and Aux1 ports, respectively.

The SD-WAN 400/800/1000 and 2000 appliances each have two pairs of accelerated bridge ports. The 3000 appliance has three pairs of accelerated-bridge ports. On the appliance, ports 1/1 and 1/2 are the accelerated pair A (apA) bridge ports, ports 1/3 and 1/4 are the apB ports, and ports 1/5 and 1/6 are the apC bridge ports.

Connecting the Ethernet Cables

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port of the appliance.
2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Figure 1. Inserting an Ethernet cable

Connecting the Console Cable

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port.
   Note: To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.
2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Figure 2. Inserting a console cable

Connecting the Power Cable

The appliances have one power supply, unless you have installed a second. A separate ground cable is not required, because the three-prong plug provides grounding. Provide power to the appliance by installing the power cord.

To connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply.
2. Connect the other end of the power cable to a standard 110V/220V power outlet.

Figure 3. Inserting a power cable

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

To switch on the appliance

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Depending on the appliance, press the ON/OFF toggle power switch or the power button to switch on the appliance.

Caution: Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Initial Configuration

Aug 09, 2017

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

To perform initial configuration:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Perform additional configuration for Windows.
- Assign management IP address through the serial console.
- Troubleshoot initial configuration issues.

By default, the initial configuration deploys the appliance in inline mode.

Prerequisites

Aug 09, 2017

Before you begin configuring the appliance, make sure that the following prerequisites have been met:

- You should have physical access to the appliance.
- You have chosen four IP addresses for management of the appliance.
- In the Worksheet, record all IP addresses and other values you would use to configure the appliance. Preferably, print out the worksheet before you start the configuration process.
- You should already have a SD-WAN license key from Citrix, sent in an email. If you are using remote licensing, you need the IP address of the licensing server.
- WAN Send and Receive Speeds.

Worksheet

Aug 09, 2017

The following table lists the default values of network configuration, SD-WAN configuration, Windows configuration, system settings, password, and setup wizard in your traffic subnet and the management subnet. Record the values applicable to your appliance in the third column.

Deployment Worksheet			
Field	**Default Value**	**Value for your Appliance**	**Description of the Field**
Network Configuration			
XenServer IP Address (Management Subnet)	192.168.100.2		Management IP address of the XenServer.
Management Service IP Address (Management Subnet)	192.168.100.1		Management IP address of the Management Service.
Netmask (Management Subnet)	255.255.0.0		Network mask for the management subnet.
Gateway (Management Subnet)	None		The default gateway IP address of the appliance.
Port Model	2-Port		Select 2-port or 4-port, depending on the model. In 4-port mode, Windows Server does not have access to ports 1/3 and 1/4.
DNS Server	None		IP address of the DNS server. Citrix recommends that you specify a valid DNS server IP address. This is a mandatory parameter.
SD-WAN Configuration			
IP Address (Management Subnet)	192.168.100.20		Primary SD-WAN IP address at which you manage the SD-WAN instance.
Netmask (Management Subnet)	255.255.0.0		Network mask for the management IP address of the appliance. Same as previous netmask.
Gateway (Management Subnet)	None		The default gateway IP address of the appliance. Same as the previous gateway.
System Settings			
NTP Server	(none)		IP address of the NTP server. Citrix recommends that you

Deployment Worksheet			specify a valid NTP server IP address. You can either enter the IP address or the server name.
Time Zone	UTC		Specify the time zone for your location.
Password			
Password	nsroot		New password for access to the appliance.
Confirm Password	nsroot		New password for access to the appliance.
Command Center Configuration			
Command Center IP Address	None		Optional. IP address of the Command Center appliance with which you want to register this appliance. More info.
Command Center Port	8443		Optional. Port number of the Command Center SD-WAN. More info.
Registration Password	None		Password you want to use to register the SD-WAN appliance.
Licensing			
License Server Address	None		IP address of the licensing server. Required only when you select a remote model license type.
Licensing Service Port	27000		Port number of the licensing server. Required only when you select a remote model license type.
Links			
Receive (Download) Speed	None		WAN link download speed.
Send (Upload) Speed	None		WAN link upload speed.

Configuring the Appliance by Connecting a Computer to the Ethernet Port

Aug 09, 2017

For initial configuration of a SD-WAN appliance, perform the following tasks::

- Configure the appliance for use on your site.
- Install the Citrix license.
- Enable acceleration.
- Enable traffic shaping (inline mode only).

With inline deployments, this configuration might be all you need, because most acceleration features are enabled by default and require no additional configuration.

You can configure the appliance connecting the appliance to your computer through either the Ethernet port or the serial console. The following procedure enables you to configure the appliance by connecting it to your computer through the Ethernet port.

If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 25 of the following procedure.

Note: Make sure that you have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.50, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields as blank.

2. Using an Ethernet cable, connect this computer to the port labeled MGMT on a SD-WAN 1000 appliance with Windows Server, or to the port labeled PRI on a SD-WAN 2000 appliance with Windows Server.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance:
   Username: nsroot

   Password: nsroot.

5. Start the configuration wizard by clicking **Get Started**.
6. On the Platform Configuration page, enter the respective values from your worksheet, as shown in the following example:

   Note: If, for configuration, you want to use the same network mask and gateway as those for Network Configuration, select the **Use System Netmask and Gateway** option.
7. Click **Done**. A screen showing the Installation in Progress… message appears. This process takes approximately 2 to 5 minutes, depending on your network speed.
   Note: If you are configuring the appliance by connecting it to your computer through the serial console port, skip step 8

through step 14.

8. A Redirecting to new management IP message appears.
9. Click **OK**.
10. Unplug your computer from the Ethernet port and connect the port to your management network.
11. Reset the IP address of your computer to its previous setting.
12. From a computer on the management network, log on to the appliance by entering the new Management Service IP address, such as https://<Managemnt_IP_Address>, in a web browser.
13. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
14. Log on to the appliance by using the **nsroot** user name and the password from your worksheet.
15. The Configuration wizard starts again. In this wizard, some of the values which you have already provided, appear by default. Specify rest of the values you have recorded in your worksheet.
16. If you want to manage the appliance through Command Center, specify the Command Center IP address, port, and registration password in the **Command Center Configuration** page. Otherwise, skip this step.
17. In System Services section, update the values if necessary.
18. In the **Licensing** section, select the appropriate license type. You can either select a local license or a remote license server to apply a license to the appliance.
    1. If you opt for a local license, you must generate a license by using the host ID of the appliance. To generate a local license for the appliance, see http://support.citrix.com/article/ctx131110. To apply the license, you can navigate to the NetScaler SD-WAN appliance; **Configuration** > **Appliance Settings** > **Licensing** page, after completing the Configuration wizard.
    2. If you opt for a remote licensing server, you must select a remote appliance model and provide the IP address of the licensing server in the **Licensing Server Address** field.
19. In the WAN Link Definition section, specify receive and send speeds for the WAN link in the respective fields. Citrix recommends values 10% lower than the WAN bandwidth, to avoid network congestion.
20. By default, WAN-side adapter settings are configured on the appliance. Accept the default settings.
21. Click **Install**. After the Installation process is complete, the appliance restarts.
22. As soon as the appliance restarts, the Dashboard page appears.

23. To configure the appliance to accelerate the network traffic, open navigate to the **Configuration** tab.
    Note: Make sure that you have already applied the appropriate license to the appliance.
24. On the **Network Adapters** page of the Appliance Settings node, verify and, if necessary, assign IP addresses, subnet masks, and gateways to the accelerated bridges (apA and apB) to be used. Applying these changes restarts the appliance.
    Note: You need to assign IP addresses to apA and apB adapters only if you intended to configure WCCP mode, virtual inline mode, or the Video Caching feature on the appliance.
25. The Initial Configuration is complete. Traffic now flows through the appliance. The Dashboard page shows this traffic.

26. You need additional configuration on the appliance if you intend to use some of the modes and features, such as WCCP mode, virtual inline mode, video caching, secure peering, high availability, encrypted CIFS/MAPI acceleration, AppFlow monitoring, or SNMP monitoring.

Note:

- Inline installations place the appliance between your LAN and WAN routers, using both ports of the accelerated bridge, such as ports LAN1 and WAN1 on a SD-WAN 1000 appliance with Window Server or ports 1/1 and 1/2 on SD-WAN 2000

appliance with Windows Server, for the apA accelerated bridge port.

- WCCP and virtual inline installations connect a single accelerated bridge port to your WAN router.
- Virtual inline installations require that you configure your router to forward WAN traffic to the appliance. See Router Configuration.
- WCCP installations require configuration of your router and the appliance. See WCCP Mode.

Assigning a Management IP Address through the Serial Console

Aug 09, 2017

If you do not want to change the settings of your computer, you can perform initial configuration by connecting the appliance to your computer with a serial null modem cable. Make sure that you have physical access to the appliance.

To configure the appliance through the serial console

1. Connect a serial null modem cable to the appliance's console port.
2. Connect the other end of the cable to the serial COM port of a computer running a terminal emulator, such as Microsoft HyperTerminal, with settings 9600,N,8,1, p.
3. On the HyperTerminal output, press **Enter**. The terminal screen displays the Logon prompt.
   Note: You might have to press **Enter** two or three times, depending on the terminal program you are using.
4. At the logon prompt, log on to the appliance with the following default credentials:
   Username: nsroot

   Password: nsroot.

5. At the **$** prompt, run the following command to switch to the shell prompt of the appliance:
   $ ssh 169.254.0.10

6. Enter **Yes** to continue connecting to the management service.
7. Log on to the shell prompt of the appliance with the following default credentials:
   Password: nsroot.

8. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu:
   # networkconfig

9. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.
10. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the XenServer server.
11. Type **3** and press **Enter** to select option 3, and then specify the network mask for the management service IP address.
12. Type **4** and press **Enter** to select option 4, and then specify the default gateway for the management service IP address.
13. Type **8** and press **Enter** to save the settings and exit.
14. Access the SD-WAN appliance by entering the new management service IP address of the appliance, such as https://<Management_Service_IP_Address>, in a web browser of a computer on the management network.
15. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
16. Run steps 4 through 25 of the Configuring the Appliance by Connecting a Computer to the Ethernet Port procedure to complete the configuration process

Deployment Modes

Aug 09, 2017

A SD-WAN appliance acts as a virtual gateway. It is neither a TCP endpoint nor a router. Like any gateway, its job is to buffer incoming packets and put them onto the outgoing link at the right speed. This packet forwarding can be done in different ways, such as inline mode, virtual inline mode, and WCCP mode. Although these methods are called *modes*, you do not have to disable one forwarding mode to enable another. If your deployment supports more than one mode, the mode that the appliance uses is determined automatically by the Ethernet and IP format of each packet.

Because the appliance supports different forwarding modes and different kinds of non-forwarded connections, it needs a way of distinguishing one kind of traffic from another. It does so by examining the destination IP address and destination Ethernet address (MAC address), as shown in table below. For example, in inline mode, the appliance is acting as a bridge. Unlike other traffic, bridged packets are addressed to a system beyond the appliance, not to the appliance itself. The address fields contain neither the appliance's IP address nor the appliance's Ethernet MAC address.

In addition to pure forwarding modes, the appliance has to account for additional types of connections, including management connections to the GUI and the heartbeat signal that passes between members of a high-availability pair. For completeness, these additional traffic modes are also listed in table below.

Table 1. How Ethernet and IP Addresses Determine the Mode

Destination IP Address	Destination Ethernet Address	Mode
Not appliance	Not appliance	Inline or Pass-through
Not appliance	Appliance	Virtual Inline or L2 WCCP
Appliance	Appliance	Direct (UI access)
Appliance (VIP)	Appliance	High-Availability. Proxy mode
Appliance (WCCP GRE Packet)	Appliance	WCCP GRE Mode
Appliance (Signaling IP)	Appliance	Signaling Connection (SD-WAN plugin Signaling Connection (SD-WAN plugin, Secure Peer) or Redirector Mode Connection (SD-WAN plugin)
All modes can be active simultaneously. The mode used for a given packet is determined by the Ethernet and IP headers.		

The forwarding modes are:

- **Inline mode,** in which the appliance transparently accelerates traffic flowing between its two Ethernet ports. In this mode, the appliance appears (to the rest of the network) to be an Ethernet bridge. Inline mode is recommended, because it requires the least configuration.
- **WCCP mode,** which uses the WCCP v. 2.0 protocol to communicate with the router. This mode is easy to configure on most routers. WCCP has two variants: WCCP-GRE and WCCP-L2. WCCP-GRE encapsulates the WCCP traffic within

generic routing encapsulation (GRE) tunnels. WCCP-L2 uses un-encapsulated network Layer 2 (Ethernet) transport.

- **Virtual inline mode,** in which a router sends WAN traffic to the appliance and the appliance returns it to the router. In this mode, the appliance appears to be a router, but it uses no routing tables. It sends the return traffic to the real router. Virtual inline mode is recommended when inline mode and high-speed WCCP operation are not practical.
- **Group mode,** which allows two appliances to operate together to accelerate a pair of widely separated WAN links.
- **High availability mode,** which allows to appliances to operate as an active/standby high availability pair. If the primary appliance fails, the secondary appliance takes over.

Additional traffic types are listed here for completeness:

- **Pass-through traffic** refers to any traffic that the appliance does not attempt to accelerate. It is a traffic category, not a forwarding mode.
- **Direct access,** where the appliance acts as an ordinary server or client. The GUI and CLI are examples of direct access, using the HTTP, HTTPS, SSH, or SFTP protocols. Direct access traffic can also include the NTP and SNMP protocols.
- **Appliance-to-appliance communication**, which can include signaling connections (used in secure peering and by the SD-WAN plugin), VRRP heartbeats (used in high-availability mode), and encrypted GRE tunnels (used by group mode).
- **Deprecated modes.** Proxy mode and redirector mode are legacy forwarding modes that should not be used in new installations.

Customizing the Ethernet ports

Aug 09, 2017

A typical appliance has four Ethernet ports: two accelerated bridged ports, called *accelerated pair A* (apA.1 and apA.2), with a bypass (fail-to-wire) relay, and two unaccelerated motherboard ports, called Primary and Aux1. The bridged ports provide acceleration, while the motherboard ports are sometimes used for secondary purposes. Most installations use only the bridged ports.

Some SD-WAN units have only the motherboard ports. In this case, the two motherboard ports are bridged.

The appliance's user interface can be accessed by a VLAN or non-VLAN network. You can assign a VLAN to any of the appliance's bridged ports or motherboard ports for management purposes.

Figure 1. Ethernet Ports

Port List

The ports are named as follows:

Table 1. Ethernet Port Names

Motherboard port 1	Primary (or apA.1 if no bypass card is present)
Motherboard port 2	Auxiliary1 or Aux1 (or apA.2 if no bypass card is present)
Bridge #1	Accelerated Pair A (apA, with ports apA.1 and apA.2)
Bridge #2	Accelerated Pair B (apB, with ports apB.1 and apB.2)

Port Parameters

Aug 09, 2017

Each bridge and motherboard port can be:

- Enabled or disabled
- Assigned an IP address and subnet mask
- Assigned a default gateway
- Assigned to a VLAN
- Set to 1000 Mbps, 100 Mbps, or 10 Mbps
- Set to full duplex, half-duplex, or auto (on SD-WAN WANOP 4000/5000 appliances, some ports can be set to 10 Gbps)

All of these parameters except the speed/duplex setting are set on the Configuration: IP Address page. The speed/duplex settings are set on the Configuration: Interface page.

Notes about parameters:

- Disabled ports do not respond to any traffic.
- The browser-based UI can be enabled or disabled independently on all ports.
- To secure the UI on ports with IP addresses, select HTTPS instead of HTTP on the Configuration: Administrator Interface: Web Access page.
- Inline mode works even if a bridge has no IP address. All other modes require that an IP address be assigned to the port.
- Traffic is not routed between interfaces. For example, a connection on bridge apA does not cross over to the Primary or Aux1 ports, but remains on bridge apA. All routing issues are left to your routers.

Accelerated Bridges (apA and apB)

Aug 09, 2017

Every appliance has at least one pair of Ethernet ports that function as an accelerated bridge, called *apA* (for *accelerated pair A*). A bridge can act in inline mode, functioning as a transparent bridge, as if it were an Ethernet switch. Packets flow in one port and out the other. Bridges can also act in one arm mode, in which packets flow in one port and back out the same port.

An appliance that has a bypass card maintains network continuity if a bridge or appliance malfunctions.

Some units have more than one accelerated pair, and these additional accelerated pairs are named apB, apC, and so on.

Bypass Card

If the appliance loses power or fails in some other way, an internal relay closes and the two bridged ports are electrically connected. This connection maintains network continuity but makes the bridge ports inaccessible. Therefore you might want to use one of the motherboard ports for management access.

Caution: Do not enable the Primary port if it is not connected to your network. Otherwise, you cannot access the appliance, as explained in Ethernet Bypass and Link-Down Propagation

Bypass cards are standard on some models and optional on others. Citrix recommends that you purchase appliances with bypass cards for all inline deployments.

The bypass feature is wired as if a cross-over cable connected the two ports, which is the correct behavior in properly wired installations.

Important: Bypass installations must be tested - Improper cabling might work in normal operation but not in bypass mode. The Ethernet ports are tolerant of improper cabling and often silently adjust to it. Bypass mode is hard-wired and has no such adaptability. Test inline installations with the appliance turned off to verify that the cabling is correct for bypass mode.

Using Multiple Bridges

If the appliance is equipped with two accelerated bridges, they can be used to accelerate two different links. These links can either be fully independent or they can be redundant links connecting to the same site. Redundant links can be either load-balanced or used as a main link and a failover link.

Figure 1. Using dual bridges

When it is time for the appliance to send a packet for a given connection, the packet is sent over the same bridge from which the appliance received the most recent input packet for that connection. Thus, the appliance honors whatever link decisions are made by the router, and automatically tracks the prevailing load-balancing or main-link/failover-link algorithm in real time. For non-load-balanced links, the latter algorithm also ensures that packets always use the correct bridge.

WCCP and Virtual Inline Modes

Multiple bridges are supported in both WCCP mode and virtual inline mode. Usage is the same as in the single-bridge case, except that WCCP has the additional limitation that all traffic for a given WCCP service group must arrive on the same bridge.

High Availability with Multiple Bridges

Two units with multiple bridges can be used in a high-availability pair. Simply match up the bridges so that all links pass through both appliances.

Motherboard Ports

Aug 09, 2017

Although the Ethernet ports on a bypass card are inaccessible when the bypass relay is closed, the motherboard ports remain active. You can sometimes access a failed appliance through the motherboard ports if the bridged ports are inaccessible.

The Primary Port

If the Primary port is enabled and has an IP address assigned to it, the appliance uses that IP address to identify itself to other acceleration units. This address is used internally for a variety of purposes, and is most visible to users as the Partner Unit field on the Monitoring: Optimization: Connections page. If no motherboard port is enabled, the appliance uses the IP address of Accelerated Pair A.

The Primary port is used for:

- Administration through the web based UI
- A back channel for group mode
- A back channel for high-availability mode

The Aux1 Port

The Aux1 port is identical to the Primary port. If the Aux1 port is enabled and the Primary port is not, the appliance takes its identity from the Aux1 port's IP address. If both are enabled, the Primary port's IP address is the unit's identity

VLAN Support

Aug 09, 2017

A virtual local area network (VLAN) uses part of the Ethernet header to indicate which virtual network a given Ethernet frame belongs to. SD-WAN appliances support VLAN trunking in all forwarding modes (inline, WCCP, virtual inline, and group mode). Traffic with any combination of VLAN tags is handled and accelerated correctly.

For example, if one traffic stream passing through the accelerated bridge is addressed to 10.0.0.1, VLAN 100, and another is addressed to 10.0.0.1, VLAN 111, the appliance knows that these are two distinct destinations, even though the two VLANs have the same IP address.

You can assign a VLAN to all, some, or none of the appliance's Ethernet ports. If a VLAN is assigned to a port, the management interfaces (GUI and CLI) listen only to traffic on that VLAN. If no VLAN is assigned, the management interfaces listen only to traffic without a VLAN. This selection is made on the Configuration: Appliance Settings: Network Adapters: IP Addresses tab.

Inline Mode

Aug 09, 2017

In inline mode, traffic passes into one of the appliance's Ethernet ports and out of the other. When two sites with inline appliances communicate, every TCP connection passing between them is accelerated. All other traffic is passed through transparently, as if the appliance were not there.

Figure 1. Inline mode, Accelerating All Traffic on a WAN

Note: Any TCP-based traffic passing through both units is accelerated. No address translation, proxying, or per-site setup is required. Inline mode is auto-detecting and auto-configuring.
Configuration is minimized with inline mode, because your WAN router need not be aware of the appliance's existence.

Depending on your configuration, inline mode's link-down propagation can affect management access to the appliance if a link goes down.

Inline mode is most effective when applied to all traffic flowing into and out of a site, but it can be used for only some of the site's traffic.

Ethernet Bypass and Link-Down Propagation

Aug 09, 2017

Note: Link-Down propagation was added to the SD-WAN (formerly SD-WAN) 1000, 2000, 3000, 4000, and 5000 appliances with the 7.2.1 release.

Most appliance models include a "fail-to-wire" (Ethernet bypass) feature for inline mode. If power fails, a relay closes and the input and output ports become electrically connected, allowing the Ethernet signal to pass through from one port to the other as if the appliance were not there. In fail-to-wire mode, the appliance looks like a cross-over cable connecting the two ports.

Any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the appliance's bridge ports are inaccessible.

If carrier is lost on one of the bridge ports, the carrier is dropped on the other bridge port to ensure that the link-down condition is propagated to the device on the other side of the appliance. Units that monitor link state (such as routers) are thus notified of conditions on the other side of the bridge.

Link-down propagation has two operating modes:

- If the Primary port is not enabled, the link-down state on one bridge port is mirrored briefly on the other bridge port, and then the port is re-enabled. This allows the appliance to be reached through the still-connected port for management, HA heartbeat, and other tasks.
- If the Primary port is enabled, the appliance assumes (without checking) that the Primary port is used for management, HA heartbeat, and other tasks. The link-down condition on one bridge port is mirrored persistently on the other port, until carrier is restored or the unit is rebooted. This is true even if the Primary port is enabled in the GUI but not connected to a network, so the Primary port should be disabled (the default) when not in use.

Accelerating an Entire Site

Aug 09, 2017

Inline mode, Accelerating All Traffic on a WAN shows a typical configuration for inline mode. For both sites, the appliances are placed between the LAN and the WAN, so all WAN traffic that can be accelerated is accelerated. This is the simplest method for implementing acceleration, and it should be used when practical.

Because all the link traffic is flowing through the appliances, the benefits of fair queuing and flow control prevent the link from being overrun.

In IP networks, the bottleneck gateway determines the queuing behavior for the entire link. By becoming the bottleneck gateway, the appliance gains control of the link and can manage it intelligently. This is done by setting the bandwidth limit slightly lower than the link speed. When this is done, link performance is ideal, with minimal latency and loss even at full link utilization.

Partial-Site Acceleration

Aug 09, 2017

To reserve the appliance's accelerated bandwidth for a particular group of systems, such as remote backup servers, you can install the appliance on a branch network that includes only those systems. This is shown in the following figure.

Figure 1. Inline Mode, Accelerating Selected Systems Only

SD-WAN traffic shaping relies on controlling the entire link, so traffic shaping is not effective with this topology, because the appliance sees only a portion of link traffic. Latency control is up to the bottleneck gateway, and interactive responsiveness can suffer.

Configuring and Troubleshooting Inline Mode

Aug 09, 2017

Inline mode requires only basic configuration, because it is applied automatically to any packets passing through the accelerated bridge. Troubleshooting is described under .

WCCP Mode

Aug 09, 2017

Web Cache Communication Protocol (WCCP) is a dynamic routing protocol introduced by Cisco. Originally intended only for web caching, WCCP version 2 became a more general-purpose protocol, suitable for use by accelerators such as Citrix NetScaler SD-WAN appliances.

WCCP mode is the simplest way of installing a SD-WAN appliance when inline operation is impractical. It is also useful where asymmetric routing occurs, that is, when packets from the same connection arrive over different WAN links. In WCCP mode, the routers use the WCCP 2.0 protocol to divert traffic through the appliance. Once received by the appliance, the traffic is treated by the acceleration engine and traffic shaper as if it were received in inline mode.

Note:

- For the purposes of this discussion, WCCP version 1 is considered obsolete and only WCCP version 2 is presented.
- The standard WCCP documentation calls WCCP clients "caches." To avoid confusion with actual caches, Citrix generally avoids calling a WCCP client a "cache." Instead, WCCP clients are typically called "appliances."
- This discussion uses the term "router" to indicate WCCP-capable routers and WCCP-capable switches. Though the term "router" is used here, some high-end switches also support WCCP, and can be used with SD-WAN appliances.

The SD-WAN appliances support two WCCP modes:

- WCCP is the original SD-WAN WCCP offering supported since release 3.x. It supports a single appliance service group (no clustering).
- WCCP clustering, introduced in release 7.2, allows your router to load-balance traffic between multiple appliances.

How WCCP Mode Works

The physical mode for WCCP deployment of a SD-WAN appliance is one-arm mode in which the appliance is connected directly to a dedicated port on the WAN router. The WCCP standard includes a protocol negotiation in which the appliance registers itself with the router, and the two negotiate the use of features they support in common. Once this negotiation is successful, traffic is routed between the router and the appliance according to the WCCP router and redirection rules defined on the router.

A WCCP-mode appliance requires only a single Ethernet port. The appliance should either be deployed on a dedicated router port (or WCCP-capable switch port) or isolated from other traffic through a VLAN. Do not mix inline and WCCP modes.

The following figure shows how a router is configured to intercept traffic on selected interfaces and forward it to the WCCP-enabled appliance. Whenever the WCCP-enabled appliance is not available, the traffic is not intercepted, and is forwarded normally.

Figure 1. WCCP Traffic Flow

Traffic Encapsulation

WCCP allows traffic to be forwarded between the router and the appliance in either of the following modes:

- L2 Mode—Requires that the router and appliance be on the same L2 segment (typically an Ethernet segment). The IP packet is unmodified, and only the L2 addressing is altered to forward the packet. In many devices, L2 forwarding is

performed at the hardware layer, giving it the maximum performance. Because of its performance advantage, L2 forwarding is the preferred mode, but not all WCCP-capable devices support it.

- GRE Mode—Generic Routing Encapsulation (GRE) is a routed protocol and the appliance can in theory be placed anywhere, but for performance it should be placed close to the router, on a fast, uncongested path that traverses as few switches and routers as possible. GRE is the original WCCP mode. A GRE header is created and the data packet is appended to it. The receiving device removes the GRE header. With encapsulation, the appliance can be on a subnet that is not directly attached to the router. However, both the encapsulation process and the subsequent routing add CPU overhead to the router, and the addition of the 28-byte GRE header can lead to packet fragmentation, which adds additional overhead.

WCCP mode supports multiple routers and both GRE vs. L2 forwarding. Each router can have multiple WAN links. Each link can have its own WCCP service group.

Traffic shaping is not effective unless the appliance manages UDP traffic as well as TCP traffic. A second service group, with a UDP service group for each WAN link, is recommended if traffic shaping is desired.

Registration and Status Updates

A WCCP client (an appliance) uses UDP port 2048 to register itself with the router and to negotiate which traffic should be sent to it, and also which WCCP features should be used for this traffic. The appliance operates on this traffic and forwards the resulting traffic to the original endpoint. The status of an appliance is tracked through the WCCP registration process and a heartbeat protocol. The appliance first contacts the router over the WCCP control channel (UDP port 2048), and the appliance and router exchange information with packets named “Here_I_Am” and “I_See_You,” respectively. By default, this process is repeated every ten seconds. If the router fails to receive a message from the appliance for three of these intervals, it considers the appliance to have failed and stops forwarding traffic to it until contact is reestablished.

Services and Service Groups

Different appliances using the same router can provide different services. To keep track of which services are assigned to which appliances, the WCCP protocol uses a service group identifier, a one-byte integer. When an appliance registers itself with a router, it includes service group numbers as well.

- A single appliance can support more than one service group.
- A single router can support more than one service group.
- A single appliance can use the same service group with more than one router.
- A single router can use the same service group with more than one appliance. For SD-WAN appliances, multiple appliances are supported in WCCP cluster mode, and a single appliance is supported in WCCP mode.
- Each appliance specifies a “return type” (L2 or GRE) independently for each direction and each service group. SD-WAN 4000/5000 appliances always specify the same return type for both directions. Other SD-WAN appliances allow the return type to be different.

Figure 2. Using different WCCP service groups for different services

Multiple service groups can be used with WCCP on the same appliance. For example, the appliance can receive service-group 51 traffic from one WAN link and service-group 62 traffic from another WAN link. The appliance also supports multiple routers. It is indifferent to whether all the routers use the same service group or different routers use different service groups.

Service Group Tracking. If a packet arrives on one service group, output packets for the same connection are sent on the

same service group. If packets arrive for the same connection on multiple service groups, output packets track the most recently seen service group for that connection.

High Availability Behavior

When WCCP is used with high-availability mode, the primary appliance sends its own apA or apB management IP address, not the virtual address of the HA pair, when it contacts the router. If failover occurs, the new primary appliance contacts the router automatically, reestablishing the WCCP channel. In most cases the WCCP timeout period and the HA failover time overlap. As a result, the network outage is less than the sum of the two delays.

Standard WCCP allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive. WCCP clustering allows multiple appliances per service group.

Deployment Topology

The following figure shows a simple WCCP deployment, suitable for either L2 or GRE. The traffic port (1/1) is connected directly to a dedicated router port (Gig 4/12).
Figure 3. Simple WCCP deployment

In this example, the SD-WAN 4000/5000 is deployed in one-arm mode, with the traffic port (1/1) and the management port (0/1) each connecting to its own dedicated router port.

On the router, WCCP is configured with identical ip wccp redirect in statements on the WAN and LAN ports. Two service groups are used, 71 and 72. Service group 71 is used for TCP traffic and service group 72 is used for UDP traffic. The appliance does not accelerate UDP traffic, but can apply traffic shaping policies to it.

Note: The WCCP specification does not allow protocols other than TCP and UDP to be forwarded, so protocols such as ICMP and GRE always bypass the appliance.

WCCP Clustering

SD-WAN appliances support WCCP clustering, which enables your router to load-balance your traffic between multiple appliances. For more information about deploying SD-WAN appliances as a cluster, see WCCP Clustering.

WCCP Specification

For more information about WCCP, see Web Cache Communication Protocol V2, Revision 1, http://tools.ietf.org/html/draft-mclaggan-wccp-v2rev1-00.

> **Note**
>
> When deploying SD-WAN in WCCP for switch redundancy, we can connect switch 2 to apB. Create a different SG for apB, give it a lower priority than the SG for apA. If apA higher SG is up, that will be used for redirection. If that is down, apB SG will be used. Please note that apA and apB need to be on different subnet.

WCCP Mode (Non-Clustered)

Aug 09, 2017

WCCP mode allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive.

Note: WCCP clustering allows multiple appliances per service group.

Limitations and Best Practices

Following are limitations and best practices for (non-clustered) WCCP mode:

- On appliances with more than one accelerated pair, all the traffic for a given WCCP service group must arrive on the same accelerated pair.
- Do not mix inline and WCCP traffic on the same appliance. The appliance does not enforce this guideline, but violating it can cause difficulties with acceleration. (WCCP and virtual inline modes can be mixed, but only if the WCCP and virtual inline traffic are coming from different routers.)
- For sites with a single WAN router, use WCCP whenever inline mode is not practical.
- Only one appliance is supported per service group. If more than one appliance attempts to connect to the same router with the same service group, the negotiation will succeed only for the first appliance.
- For sites with multiple WAN routers serviced by the same appliance, WCCP can be used to support one, some, or all of your WAN routers. Other routers can use virtual inline mode.

Router Support for WCCP

Aug 09, 2017

Configuring the router for WCCP is very simple. WCCP version 2 support is included in all modern routers, having been added to the Cisco IOS at release 12.0(11)S and 12.1(3)T. The best router-configuration strategy is determined by the characteristics of your router and switches. Traffic shaping requires two service groups.

If your router supports Reverse Path Forwarding, you must disable it on all ports, because it can confuse WCCP traffic with spoofed traffic. This feature is found in newer Cisco routers such as the Cisco 7600.

Router Configuration Strategies

There are two basic approaches to redirecting traffic from the router to the appliance:

- **On the WAN port only**, add a “WCCP redirect in” statement and a “WCCP redirect out” statement.
- **On every port on the router, except the port attached to the appliance**, add a “WCCP redirect in” statement.

The first method redirects only WAN traffic to the appliance, while the second method redirects all router traffic to the appliance, whether it is WAN related or not. On a router with several LAN ports and substantial LAN-to-LAN traffic, sending all traffic to the appliance can overload its LAN segment and burden the appliance with this unnecessary load. If GRE is used, the unnecessary traffic can load down the router as well.

On some routers, the "redirect in" path is faster and puts less of a load on the router's CPU than does the "redirect out" path. If necessary, this can be determined by direct experiment on your router: Try both redirection methods under full network load to see which delivers the highest transfer rates.

Some routers and WCCP-capable switches do not support "WCCP redirect out," so the second method must be used. To avoid overloading the router, the best practice to avoid redirecting large numbers of router ports through the appliance, perhaps by using two routers, one for WAN routing and one for LAN-to-LAN routing.

In general, method 1 is simpler, while method 2 may provide greater performance.

Traffic Shaping and WCCP

A service group can be either TCP or UDP, but not both. For the traffic shaper to be effective, both kinds of WAN traffic must pass through the appliance. Therefore:

- Acceleration requires one service group, for TCP traffic.
- Traffic shaping requires two service groups, one for TCP traffic and one for UDP traffic.
- The difference between the two is configured on the appliance, and the router accepts this configuration.

Configuring the Router

Aug 09, 2017

The appliance negotiates WCCP-GRE or WCCP-L2 automatically. The main choice is between *unicast operation* (in which the appliance is configured with the IP address of each router), or *multicast operation* (in which both the appliance and the routers are configured with the multicast address.)

Normal (Unicast) operation—For normal operation, the procedure is to declare WCCP version 2 and the WCCP group ID for the router as a whole, then enable redirection on each WAN interface. Following is a Cisco IOS example:

```
config term
ip wccp version 2
! We will configure the appliance to use group 51 for TCP and 52 for UDP.
ip wccp 51
ip wccp 52

! Repeat the following three lines for each WAN interface
! you wish to accelerate:
interface your_wan_interface
! If Reverse Path Forwarding is enabled (with an ip verify unicast
! source reachable" statement), delete or comment out the statement:
! ip verify unicast source reachable-via any
! Repeat on all ports.

ip wccp 51 redirect out
ip wccp 51 redirect in
ip wccp 52 redirect out
ip wccp 52 redirect in

! If the appliance is inline with one of the router interfaces
! (NOT SUPPORTED), add the following line for that interface
! to prevent loops:
ip wccp redirect exclude in
^Z
```

If multiple routers are to use the same appliance, each is configured as shown above, using either the same service groups or different ones.

Multicast operation—When giving the appliance and each router a multicast address, the configuration is slightly different than for normal operation. Following is a Cisco IOS example:

```
config term
ip wccp version 2
ip wccp 51 group-address 225.0.0.1

! Repeat the following three lines for each WAN interface
! you wish to accelerate:
interface your_wan_interface
! If Reverse Path Forwarding is enabled (with an ip verify unicast
! source reachable" statement), delete or comment out the statement:
```

```
! ip verify unicast source reachable-via any

ip wccp 51 redirect out
ip wccp 51 redirect in
!
! The following line is needed only on the interface facing the other router,
! if there is another router participating in this service group.
ip wccp 51 group-listen

!If the appliance is inline with one of the router interfaces,
!(which is supported but not recommended), add
!the following line for that interface to prevent loops:
ip wccp redirect exclude in
^Z
```

Basic Configuration Procedure for WCCP Mode on the SD-WAN Appliance

Aug 09, 2017

For most sites, you can use the following procedure to configure the WCCP mode on the appliance. The procedure has you set several parameters to sensible default values. Advanced deployments might require that you set these parameters to other values. For example, if WCCP service group 51 is already used by your router, you need to use a different value for the appliance.

To configure WCCP mode on the appliance

1. On the Configuration: Appliance Settings: WCCP page.
2. If no service groups have been defined, the Select Mode page appears. The options are Single SD-WAN and Cluster (Multiple SD-WAN). Select Single SD-WAN. You are taken to the WCCP page.
   Note: The mode labels are misleading. “Single SD-WAN” mode is also used for high-availability pairs.
3. If WCCP mode is not enabled, click Enable.
4. Click Add Service Group.
5. The default interface (apA), Protocol (TCP), WCCP Priority (0), Router Communication (Unicast), (Password blank) and Time to Live (1) values usually do not have to be changed for the first service group that you create, but if they do, type new values in the fields provided.
6. In the Router Addressing field (if you are using unicast) or the Multicast Address field (if you are using multicast), type the router’s IP address. Use the IP for the router port used for WCCP communication with the appliance.
7. If more than one router is using WCCP to communicate with this appliance, add additional routers now.
8. If your routers have special requirements, set the Router Forwarding (Auto/GRE/Level-2), Router Packet Return (Auto/GRE/Level-2), and Router Assignment (Mask/Hash) fields accordingly. The defaults produce optimal results with most routers.
9. Click Add.
10. Repeat the preceeding steps to create another service group, for UDP traffic (for example, service group Id 52 and Protocol UDP).
11. Go to the Monitoring: Appliance Performance: WCCP page. The Status field should change to Connected within 60 seconds.
12. Send traffic over the link and, on the Connections page, verify that connections are arriving and being accelerated.

WCCP Service Group Configuration Details

Aug 09, 2017

In a service group, a WCCP router and a SD-WAN appliance ("WCCP Cache" in WCCP terminology) negotiate communication attributes (capabilities). The router advertises its capabilities in the "I See You" message. The communication attributes are:

- Forwarding Method: GRE or Level-2
- Packet Return Method (multicast only): GRE or Level-2
- Assignment Method: Hash or Mask
- Password (defaults to none)

The appliance triggers an alert if it detects an incompatibility between its attributes and those of the router. The appliance might be incompatible because of a specific attribute of a service group (such as GRE or Level-2). More rarely, in a multicast service group, an alert can be triggered when the "Auto" selection chooses a particular attribute with a particular router connected, but the attribute is incompatible with a subsequent router.

Following are the basic rules for the communication attributes within a SD-WAN Appliance.

For Router Forwarding:

- When “Auto” is selected, the preference is for Level-2, because it is more efficient for both router and appliance. Level-2 is negotiated if the router supports it and the router is on the same subnet as the appliance.
- Routers in a unicast service group can negotiate different methods if “Auto” is selected.
- Routers in a multicast service group must all use the same method, whether forced with “GRE” or “Level-2”, or, with “Auto,” as determined by the first router in the service group to connect.
- For an incompatibility, an alert announces that the router “has incompatible router forwarding.”

For Router Assignment:

- The default is Hash.
- When “Auto” is selected, the mode is negotiated with the router.
- All routers in a service group must support the same assignment method (Hash or Mask).
- For any service group, if this attribute is configured as “Auto,” the appliance selects “Hash” or “Mask” when the first router is connected. “Hash” is chosen if the router supports it. Otherwise, “Mask” is selected. The problem of subsequent routers being incompatible with the automatically selected method can be minimized by manually selecting a method common to all routers in the service group.
- For an incompatibility, an alert announces that the router “has incompatible router assignment method.”
- With either method, the single appliance in the service group instructs all the routers in the service group to direct all TCP or UDP packets to the appliance. Routers can modify this behavior with access lists or by selecting which interfaces to redirect to the service group.
  For the Mask method, the appliance negotiates the “source IP address” mask. The appliance provides no mechanism to select “destination IP address” or the ports for either source or destination. The “source IP address” mask does not specifically identify any specific IP address or range. The protocol does not provide a means to specify a specific IP address. By default, because there is only a single appliance in the service group, a one-bit mask is used, to conserve router resources. (Release 6.0 used a larger mask.)

For Password:

- If the router requires a password, the password defined on the appliance must match. If the router does not require a password, the password field on the appliance must be blank.

WCCP Testing and Troubleshooting

Aug 09, 2017

When working with WCCP, the appliance provides different ways of monitoring the status of the WCCP interface, and your router should also provide information.

Monitoring: Appliance Performance: WCCP Page—The WCCP page reports the current state of the WCCP link, and reports most problems.

Log Entries—The Monitoring: Appliance Performance: Logging page shows a new entry each time WCCP mode is established or lost.

Figure 1. WCCP Log Entries (format varies somewhat with release)

Router Status—On the router, the "show ip wccp" command shows the status of the WCCP link:

```
Router>enable
Password:
Router#show ip wccp
Global WCCP information:
   Router information:
      Router Identifier:                  172.16.2.4
      Protocol Version:                   2.0

   Service Identifier: 51
      Number of Cache Engines:            0
      Number of routers:                  0
      Total Packets Redirected:           19951
      Redirect access-list:               -none-
      Total Packets Denied Redirect:      0
      Total Packets Unassigned:           0
      Group access-list:                  -none-
      Total Messages Denied to Group:     0
      Total Authentication failures:      0
```

WCCP Clustering

Aug 09, 2017

The WCCP clustering feature enables you to multiply your acceleration capacity by assigning more than one SD-WAN appliance to the same links. You can cluster up to 32 identical appliances, for up to 32 times the capacity. Because it uses the WCCP 2.0 standard, WCCP clustering works on most routers and some smart switches, most likely including those you are already using.

Because it uses a decentralized protocol, WCCP clustering allows SD-WAN appliances to be added or removed at will. If an appliance fails, its traffic is rerouted to the surviving appliances.

Unlike SD-WAN high-availability, an active/passive pair that uses two appliances to provide the performance of a single appliance, the same appliances deployed as a WCCP cluster has twice the performance of a single appliance, delivering both redundancy and improved performance.

In addition to adding more appliances as your site's needs increase, you can use Citrix's "Pay as You Grow" feature to increase your appliances' capabilities through license upgrades.

Citrix Command Center is recommended for managing WCCP clusters. The following figure shows a basic network of a cluster of SD-WAN appliances in WCCP mode, administered by using Citrix Command Center.

Figure 1. SD-WAN Cluster Administered by Using Citrix Command Center

Load-Balanced WCCP Clusters

The WCCP protocol supports up to 32 appliances in a fault-tolerant, load balanced array called a cluster. In the example below, three identical appliances (same model, same software version) are cabled identically and configured identically except for their IP addresses. Appliances using the same service groups with the same router can become a load balanced WCCP cluster. When a new appliance registers itself with the router, it can join the existing pool of appliances and receive its share of traffic. If an appliance leaves the network (as indicated by the absence of heartbeat signals), the cluster is rebalanced so that only the remaining appliances are used.

Figure 2. A load-balanced WCCP cluster with three appliances

One appliance in the cluster is selected as the designated cache, and controls the load-balancing behavior of the appliances in the cluster. The designated cache is the appliance with the lowest IP address. Because the appliances have identical configurations, it doesn't matter which one is the designated cache. If the current designated cache goes offline, a different appliance becomes the designated cache.

The designated cache determines how the load-balanced traffic is allocated and informs the router of these decisions. The router shares information with all members of the cluster, so the cluster can operate even if the designated cache goes offline.

Note: As normally configured, a SD-WAN 4000/5000 appliance appears as two WCCP caches to the router.

Load-Balancing Algorithm

Load balancing in WCCP is static, except when an appliance enters or leaves the cluster, which causes the cluster to be rebalanced among its current members.

The WCCP standard supports load balancing based on a mask or a hash. For example, SD-WAN WCCP clustering uses the mask method only, using a mask of 1-6 bits of the 32-bit IP address. These address bits can be non-consecutive. All addresses yielding the same result when masked are sent to the same appliance. Load balancing effectiveness depends on choosing an appropriate mask value: a poor mask choice can result in poor load-balancing or even none, with all traffic sent to a single appliance.

Deployment Topology

Aug 09, 2017

Depending on your network topology, you can deploy WCCP cluster either with a single router or with multiple routers. Whether connected to a single router or multiple routers, each appliance in the cluster must be connected identically to all routers in use.

Single Router Deployment

In the following diagram, three SD-WAN appliances accelerate the datacenter's 200 Mbps WAN. The site supports 750 XenApp users.

As shown on the Datasheet, a SD-WAN 3000-100 can support 100 Mbps and 400 users, so a pair of these appliances supports 200 Mbps and 800 users, which satisfies the datacenter's requirements of a 200 Mbps link and 750 users.

For fault tolerance, however, the WCCP cluster should continue to operate without becoming overloaded if one appliance fails. That can be accomplished by using three appliances when the calculations call for two. This is called the N+1 rule.

Failure is an unusual event, so usually all three appliances are in operation. In this case, each appliance is supporting only 67 Mbps and 250 users, leaving plenty of headroom, and making good use of the fact that the cluster has three times the CPU power and three times the compression history of a single appliance.

Without WCCP clustering, the same level of capacity and fault-tolerance would require a pair of SD-WAN 4000-500 appliances in high availability mode. Only one of these appliances is active at a time.

Multiple Router Deployment

Using multiple WAN routers is very similar to using a single WAN router. If the previous example is changed to include two 100 Mbps links instead of one 200 Mbps link, the topology changes, but the calculations do not.

Limitations

Aug 09, 2017

Configuring appliances in a WCCP cluster has the following limitations:

- All appliances within a cluster must be the same model and use the same software release.
- Parameter synchronization between appliances within the cluster is not automatic. Use Command Center to manage the appliances as a group.
- SD-WAN traffic shaping is not effective, because it relies on controlling the entire link as a unit, and none of the appliances are in a position to do this. Router QoS can be used instead.
- The WCCP-based load-balancing algorithms do not vary dynamically with load, so achieving a good load balance can require some tuning.
- The hash method of cache assignment is not supported. Mask assignment is the supported method.
- While the WCCP standard allows mask lengths of 1-7 bits, the appliance supports masks of 1-6 bits.
- Multicast service groups are not supported; only unicast service groups are supported.
- All routers using the same service group pair must support the same forwarding method (GRE or L2).
- The forwarding and return method negotiated with the router must match: both must be GRE or both must be L2. Some routers do not support L2 in both directions, resulting in an error of "Router's forward or return or assignment capability mismatch." In this case, the service group must be configured as GRE.
- SD-WAN VPX does not support WCCP clustering.
- The appliance supports (and negotiates) only unweighted (equal) cache assignments. Weighted assignments are not supported.
- Some older appliances, such as the SD-WAN 700, do not support WCCP clustering.
- (SD-WAN 4000/5000 only) Two accelerator instances are required per interface in L2 mode. No more than three interfaces are supported per appliance (and then only on appliances with six or more accelerator instances.)
- (SD-WAN 4000/5000 only) WCCP control packets from the router must match one of the router IP addresses configured on the appliance for the service group. In practice, the router's IP address for the interface that connects it to the appliance should be used. The router's loopback IP should not be used.

Planning Your Deployment

Aug 09, 2017

Deploying appliances in a WCCP cluster requires more planning than does deploying a single appliance. Read the following sections carefully before proceeding.

Selecting Appliances

Aug 09, 2017

The appliances you select for the deployment must all be the same model, running the same software version. Otherwise, management and troubleshooting can become impractical.

Your appliance choice is generally made by comparing your site's WAN bandwidth and number of WAN users to the capacities of the different appliances in the SD-WAN Data Sheet. For fault tolerance, always order one more appliance than is absolutely required according to the data sheet.

The number of appliances you need is found as follows, rounding up all fractions:
appliances = max (appliances_bw, appliances_users),
where
appliances_bw = (WAN_bandwidth / Optimized_WAN_capacity) +1
appliances_users = (WAN_users / Maximum_HDX_sessions) +1
Note that if appliances = 2, you can use just a single appliance instead of WCCP clustering, or an HA pair instead of WCCP clustering, since the equation builds in a spare appliance. In other words, WCCP clustering is not necessary (from a capacity perspective) unless appliances is 3 or more.

Example. Suppose you have 700 users and a 100 Mbps link. Some appliances you might consider are the SD-WAN 2000-050, the SD-WAN 3000-100, and the SD-WAN 4000-310.

Model	Optimized WAN Capacity	Maximum HDX Sessions	Appliances_bw	Appliances_users	Appliances
2000-050	50 Mbps	300	3	4	4
3000-100	100 Mbps	400	2	3	3
4000-310	310 Mbps	750	2	2	2

As you can see from the above table, the higher-performance platforms require fewer appliances to get the job done, as you would expect. The SD-WAN 4000-310 meets the requirements with a single appliance, and evaluates to two appliances only because the equations build in a spare.

You can always add more capacity by adding more appliances, but that is not always necessary. The bandwidth limits of two of the three choices, the SD-WAN 3000-100 and the SD-WAN 4000-310, can be increased through a license upgrade. The SD-WAN 2000-050 however, is already at the high end of the range for SD-WAN 2000 appliances.

Load-Balancing in the WCCP Cluster

Aug 09, 2017

Traffic is distributed among the appliances in the WCCP cluster. If an appliance leaves the cluster (through failure, overload, or being manually disabled), its traffic is rebalanced by distributing it among the surviving members. If an appliance joins the cluster, traffic is rebalanced once more to give the new appliance its fair share.

The Address Mask

Traffic is distributed on the basis of an address mask that is applied to the source and destination addresses of WAN traffic. You must select an appropriate mask field for efficient load-balancing. An inappropriate mask can result in load-balancing that is poor to nonexistent. For example, if the mask matches an address field that is identical at all your remote sites, all your WAN traffic is sent to a single appliance, overloading it. For example, if all of your remote sites have an address in the form of 10.0.x.x, and your mask bits are within the 10.0 portion of the address all traffic is sent to a single appliance.

The address bits extracted by the address mask are used as an index that is used (indirectly) to select one of the WCCP caches (appliances). For example, an address mask with two "one" bits results in four possible values, depending on the address. Each of these values can be thought of as a bucket. With two mask bits, you have four buckets, numbered 0-3. The buckets are assigned to WCCP caches. To be effective, there must be at least as many buckets as caches. If you use a two-bit mask and have five or more caches, some caches are idle, because each bucket is assigned to only one cache, and there are not enough buckets to cover all five caches:

Cache	1	2	3	4	5
Buckets	0	1	2	3	-

If there are more buckets than caches, some caches are assigned multiple buckets. For example, if you set three mask bits, creating eight buckets, and you have four caches, two buckets are assigned to each cache. If you have five caches, three caches are assigned two buckets each, and two caches are assigned just one. If each bucket represents the same number of users, you have a 2:1 load imbalance across caches:

Cache	1	2	3	4	5
Buckets	0-1	2-3	4-5	6	7

Increasing the number of set mask bits reduces this imbalance. With four mask bits (16 index values) and five caches, four caches receive three buckets and one cache receives four buckets, resulting in only a 4:3 imbalance. With six set mask bits (the largest number supported), four caches receive 13 buckets and one receives 12, which is only a 13:12 load imbalance.

Cache	1	2	3	4	5
Buckets	0-12	13-25	26-38	39-51	52-63

Ideally, you would like each remote site to be directed to a single appliance in the WCCP cluster, so that all traffic to and from a given site is stored in the same compression history. With this arrangement, any traffic from one user at the site can be used to compress similar traffic from any other user at that site. In other words, for compressibility, load-balancing works best if it the address mask selects the bits that differentiate one remote site from another. These are often the least-significant bits of the subnet portion of the IP address. Using these bits tends to allocate the same number of remote sites (not users) per local appliance. A mask that aligns with the host portion of the address instead of the subnet results in a

more equal number of remote users (not sites) per appliance, but at the expense of compression effectiveness. (Compression is only effective when connections flow through the same appliances, and splitting traffic from the same remote site between two or more local appliances interferes with this.)

Finally, for good load-balancing, each "one" bit in the address mask must be set to one on 50% of the remote addresses, and set to zero on 50% of the remote addresses. This is not the case on all address bits, since in most WANs, the highest-order network bits never change at all (such as the 10 in 10.x.x.x). Such bits must never be selected by the address mask.

In addition, many subnets are only sparsely populated. For example, if only 50 addresses are used in the subnet 10.1.2.0/24, and they are assigned sequentially, the two higher-order host bits (representing the unused range of 10.1.2.64-10.1.2.255) for this subnet never change, and if these two bits are included in the address mask, three-fourths of the buckets receive no traffic.

Useful compromises between these two extremes can generally be found.

Follow these rules:

- The number of ”one” bits in the address mask must allow at least as many combinations as there are WCCP caches in the cluster. That is, if you have eight appliances, the address mask must contain at least three ”one” bits.
- The ”one” bits in the address mask must each be inside the active address range for most of your remote subnets, or they skew the load-balancing distribution.
- The mask should split the address range of individual remote sites into as few pieces as possible, for best compression performance.
- If a remote appliance is faster than the local members of the WCCP cluster, the mask should be designed to divide its traffic between multiple local appliances. For example, a 100 Mbps remote appliance should have its traffic split between two 50 Mbps local appliances by setting a bit inside the remote appliance’s active address range.
- The “one” bits in the mask are typically contiguous, but this is not required. They can be in any pattern.

Example: Suppose you set an address mask of 0x0000 0f00, which has four “one” bits. This defines a four-bit field that is extracted from the IP address, yielding 16 possible results (16 buckets). These buckets are in turn assigned to the actual WCCP caches in the WCCP cluster.

Address	Masked Address (mask = 0x0000 0f00)	Bucket
10.0.0.5	0.0.0.0	0
10.0.1.128	0.0.1.0	1
155.0.2.55	0.0.2.0	2
253.100.255.2	0.0.15.0	15
10.0.15.1	0.0.15.0	15

Zero bits in the mask are ignored, and the “one” bits are used to define the extracted field. So if the mask is 0x10 10 10 10, these widely separated “one” bits are extracted into a four-bit field, declaring 16 buckets and a bucket numbers in the range of 0-15.

If the mask value is set to zero, a default value of 0x00 00 0f 00 is used.

Assigning Buckets to Appliances

Aug 09, 2017

The mapping of bucket to appliances is subject to several variables:

- Which appliances are available: If an appliance is down, its share of buckets are given to the available appliance. If a new appliance is added to the cluster, it is given its fair share of buckets.
- The mapping algorithm used (deterministic or least-disruptive).
- The order in which appliances come online (least-disruptive mapping only).
- The IP addresses of the appliances. WCCP algorithms can use a sorted list of appliance IP addresses; for example, assigning buckets to appliances in the same order as the appliance IP addresses.

The most important of these factors, from an administrator's point of view, is the mapping algorithm.

Deterministic mapping. The deterministic mapping algorithm is less graceful than the least-disruptive algorithm, but it supports Hot Standby Router Protocol (HSRP) and Global Server Load Balancing (GSLB) routing, and is required when multiple routers using such protocols share the WCCP cluster.

Deterministic mapping is also the preferred method when the cluster has only two appliances.

Assignments are based on the IP addresses of the active appliances. Each appliance gets its fair share of bucket, with the lowest-numbered bucket being assigned to the appliance with the lowest IP address. If there are more appliances than buckets, the leftover appliances (with no bucket assigned to them) are the ones with the highest-numbered IP addresses. This deterministic assignment allows traffic to arrive for a single connection through any of the routers in the service group and be forwarded to the same appliance.

Reassignment can be disruptive to accelerated connections, which are reset if they migrate to a different appliance. With deterministic mapping, the number of buckets that are reassigned to new appliances can be quite high if there are three or more appliances.

Least-disruptive mapping. When a bucket is assigned to a different appliance, any open accelerated connections that used the old appliance is reset. The least-disruptive algorithm keeps the reassignment to a minimum. For example, if you have three appliances, and one appliance fails, the new mapping preserves roughly two-thirds of the assignments and remaps the remaining third (which fails anyway, because their appliance failed). The least-disruptive algorithm does not support HSRP or GSLB routing, because it is not guaranteed to result in identical mappings on all the routers in the service group, and therefore, packets from a single connection might be sent to two different appliances by two different routers, which causes accelerated connections to fail.

Startup and Failover Behavior

Aug 09, 2017

Each appliance registers itself with the routers specified in its service class definitions. The first appliance to register itself, becomes the *designated cache*, and works with the routers to apportion traffic between itself and the other caches (called *subordinate caches*). Because your appliances use identical WCCP algorithms, it does not matter which one becomes the designated cache.

As additional appliances come online, they are added to the WCCP cluster, and the traffic is reapportioned among the active appliances. This happens at ten-second increments. After a cold start of the routers or appliances, all of the appliances might come online within the same ten-second window, or they might arrive over multiple ten-second windows, causing traffic to be reapportioned multiple times before it stabilizes. In the latter case, the appliances that come online first maycan become overloaded until additional appliances come online.

An accelerated connection fails when allocated to a different appliance, making reallocation disruptive. This is not true of non-accelerated connections, which generally experience a delay of thirty seconds or more, and then continue. The least-disruptive mapping option minimizes the amount of reallocation when an appliance fails.

If an appliance fails or otherwise goes offline, its absence is noted, and the designated cache reapportions its traffic to the remaining appliances. If the designated cache itself goes offline, the role of designated cache is also reapportioned. It takes about thirty seconds for the cluster to react to the loss of a cache.

Deployment Worksheet

Aug 09, 2017

On the following worksheet, you can calculate the number of appliances needed for your installation and the recommended mask field size. The recommended mask size is 1-2 bits larger than the minimum mask size for your installation.

Parameter	Value	Notes
Appliance Model Used		—
Supported XenApp and XenDesktop Users Per Appliance	Uspec =	From data sheet
XenApp and XenDesktop Users on WAN Link	Uwan =	—
User overload Factor	Uoverload = Uwan/Uspec =	—
Supported BW Per Appliance	BWspec =	From data sheet
WAN Link BW	BWwan =	—
BW Overload Factor	BWoverload = BWwan/BWspec =	—
Number of appliances required	N = max(Uoverload, BWoverload) +1 =	Includes one spare
		—
Min number of buckets	Bmin = N, rounded up a power of 2 =	—
If SD-WAN 4000 or 5000,	Bmin = 2*N, rounded up to a power of 2 =	—
Recommended value	B = 4 * Bmin if Bmin <= 16, else 2 * Bmin =	—
Number of "one" bits in address mask	M = log2(B)	If B=16, M=4.

Mask value: The mask value is a 32-bit address mask with a number of “one” bits equal to M in the above worksheet. Often

these bits can be the least-significant bits in the WAN subnet mask used by your remote sites. If the masks at your remote sites vary, use the median mask. (Example: With /24 subnets, the least significant bits of the subnet are 0x00 00 nn 00. The number of bits to set to one is log2(mask size): if mask size is 16, set four bits to one. So with a mask size of 16 and a /24 subnet, set the mask value to 0x00 00 0f 00.): _____

The above guidelines work only if the selected subnet field is evenly distributed in your traffic, that is, that each address bit selected by the mask is a one for half the remote hosts, and a zero for the other half. Otherwise, load-balancing is impaired. This even distribution might be true for only a small number of bits in the network field (perhaps only two or three bits). If this is the case with your network, instead of masking bits in the offending area of the subnet field, displace those bits to a portion of the host address field that has the 50/50 property. For example, if only three subnet bits in a /24 subnet have the 50/50 property, and you are using four mask bits, a mask of 0x00 00 07 10 avoids the offending bit at 0x00 00 0800 and displaces it to 0x00 00 00 10, a portion of the address field that is likely to have the 50/50 property if your remote subnets generally use at least 32 IP addresses each.

Parameter	Value	Notes
Final Mask Value		—
Accelerated Bridge		Usually apA
WAN Service Group		A service group not already in use on your router (51-255)
LAN Service Group		Another unused service group
Router IP address		IP address of router interface on port facing the appliance
WCCP Protocol (usually "Auto")		—
DC Algorithm		Use "Deterministic" if you have only two appliances or are using dynamic load balancing like HSRP or GSLB. Otherwise, use "Least Disruptive."

Configuring WCCP Clustering

Aug 09, 2017

After you have finalized the deployment topology, considered all limitations, and filled in the deployment worksheet, you are ready to deploy your appliances in a WCCP cluster. To configure the WCCP cluster, you need to perform the following tasks:

- Configuring the NetScaler Instances
- Configuring the Router
- Configuring the Appliance

Configuring the Router

Aug 09, 2017

WCCP configuration on the router is simple, because most WCCP parameters are set by the appliances.

Unlike legacy SD-WAN WCCP support, WCCP clustering uses two service groups for TCP traffic. One service group is used on the router's WAN interface, and the other is used on the router's LAN interfaces (except for the LAN interface used by the SD-WAN appliances themselves, when deployed in L2-mode WCCP cluster).

As shown in the following figure, you need to configure two service groups because WCCP allows the mask to be applied to either the source IP or the destination IP address, which is not quite what is required. To keep connections between two endpoints together, regardless of which endpoint initiates the connection, the appliance applies the address mask to the source IP address of incoming WAN traffic, and to the destination IP address of incoming LAN traffic. This requires two service groups.

The WAN service group uses WCCP source-ip address masking, while the LAN service group uses dest-ip masking. In some deployments, it may be necessary to reverse the assignments, using the "WAN" service group for your LAN interface and vice versa. This might occur if the number of local IP addresses greatly exceeds the number of remote IP addresses.

Figure 1. SD-WAN WCCP Cluster

To configure WCCP clustering on the router

This procedure assumes Cisco routers, but is similar on other routers. It uses the first of the two methods, discussed above, of redirecting WCCP traffic with an ip wccp redirect in statement on both LAN and WAN ports.

1. Fill in the WCCP clustering Deployment Worksheet.
2. Log on to your router
3. In the global declarations section, declare each service group on the WCCP clustering worksheet, listed as **WAN service group** and **LAN Service group**. For example, ip wccp 61 and ip wccp 62.
   Note: The ip wccp command allows, but does not require, a more elaborate syntax than this, and can specify an ACL name or a password. Both service groups must have the same password, if one is used. The ACLs can be different.
4. Inside the interface declarations for each WAN interface that connects to remote SD-WAN appliances, add an ip wccp x redirect in statement, where x is the WAN service group from the WCCP clustering worksheet.
5. Inside the interface declarations for each LAN interface (except the one connecting to the WCCP cluster, if you are using L2 mode), add an ip wccp y redirect in statement, where y is the LAN service group from the WCCP clustering worksheet.
6. Save your configuration.

Example. The following example uses service group 61 for the WAN service group and service group 62 for the LAN service group. Three router interfaces are used. One is connected to the WAN, one is connected to the LAN, and one is connected to the WCCP cluster.

```
!
! Example is for WCCP clustering using WCCP redirect in statements
! on LAN and WAN interfaces.
! This definition is appropriate for modern Cisco routers.
! Global declarations
ip wccp 61
ip wccp 62
```

```
!
interface GigabitEthernet1/1
description LAN interface. SG 62 is used for LAN
ip address 172.80.1.56 255.255.255.0
ip wccp 62 redirect in
!
interface GigabitEthernet1/2
description LAN interface attaching SD-WAN L2-WCCP appliances
description (No wccp redirect statements are used on this interface)
ip address 172.80.21.56 255.255.255.0
!
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
!
```

Note: If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.

- If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.
- If multiple routers shared the same WCCP cluster, they use the same service groups.

It is also possible to use ip wccp redirect statements on only the WAN interfaces:

```
! Example for WCCP clustering using WCCP redirect in/out statements on
! WAN interface only
! This definition is appropriate for modern Cisco routers.
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN. SG 62 is used for LAN.
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
ip wccp 62 redirect out
!
```

In many routers, the ip wccp redirect out path is not optimized in hardware, but uses the CPU. If the router's capabilities along this path exceeds the WAN speed, this method is practical, and is simpler than using redirect statements on every interface.

Router ACLs can be used to limit redirection. For example, for initial testing, perhaps only a single remote IP address might be allowed to be redirected through WCCP.

Configuring the Appliance

Aug 09, 2017

Repeat the following procedure for each appliance in the cluster:

1. Fill in the WCCP clustering Deployment Worksheet.
2. Navigate to Configuration > Appliance Settings > WCCP page.
3. Click Enable to enable WCCP mode on the appliance.
4. Select **Cluster (Multiple Caches)** option.

5. Fill in parameters in the **Select SD-WAN Cluster** section.

6. Enter T5 from your worksheet as the Cache 1 IP, T6 as the Cache 2 IP, T2 as the Subnet Mask, and T1 as the Gateway. Click **Save**. The **Configure Service Group** section appears.
7. In the Service Group Details section, specify the WAN and LAN service groups (T11 and T12 from your worksheet).
8. In the Priority field, select **100** (in practice this value is somewhat arbitrary).
9. From the Protocol list, select **TCP**.
10. In the DC Algorithm field, select **Deterministic** or **Least Disruptive**. “Deterministic” is always safe to use, and should be used if you are using only two appliances, or are using multiple routers. “Least Disruptive” disrupts fewer user sessions on failover when used with clusters of three or more appliances, but has restrictions on its use.
11. Set **Service Group Pair Status** to On.
12. If your router is configured to require a password, enter the password in the **Service Group Password** field. Otherwise, leave the field blank.
13. In the **Router Communications Details** section, enter the IP address of the router (T8 on your worksheet: often identical to T1 as well). This is the IP address of the appliance-facing router interface. If you use multiple routers to communicate with the appliance, list them all here.
14. From the Router Forwarding list, select Level 2 or GRE, according to the capabilities of your router. Use Level 2 if you can, and GRE if you must.
15. For the Mask Value, enter the value you determined from the WCCP Clustering worksheet. This is a critical value: a poor choice will result in poor load-balancing or none at all.
16. Click Create. This creates the WAN and LAN service groups.
17. On the Configuration > Optimization Rules > Link Definitions page, change the bandwidth limits on each defined WAN to 95% of the aggregate speed of all your WANs. This prevents the link from being under-utilized when load-balancing is imperfect. If ICA (XenApp/XenDesktop) is the dominant use, set each appliance to (95% of WAN bandwidth)/N, where N is the number of appliances (or twice the number of appliances if they are SD-WAN 4000 or 5000 units), to divide the bandwidth equally among the appliances. This latter method is most appropriate for applications with large numbers of active connections that have relatively low bandwidth requirements.

Testing and Troubleshooting

Aug 09, 2017

The **Monitoring** > **Appliance** > **Application Performance** > **WCCP** page shows the current state of not only the local appliance but of all other appliances that have joined the cluster. Select a WCCP cache and click **Get Info**.

The Cache Status tab shows the local appliance's status. When all is well, the status is "25: has assignment." You must refresh the page manually to monitor changes in status. If the appliance does not reach the status of "25: has assignment" within a timeout period, other informative status messages are displayed.

Additional information is displayed when you click on the Service Group or the Routers tabs.

The Cluster Summary tab displays information about the WCCP cluster as a whole. As a side effect of the WCCP protocol, each member of the cluster has information about all the others, so this information can be monitored from any appliance in the cluster.

Your router can also provide status information. See your router documentation.

Virtual Inline Mode

Aug 09, 2017

Note: Use virtual inline mode only when both inline mode and WCCP mode are impractical. Do not mix inline and virtual inline modes within the same appliance. However, you can mix virtual inline and WCCP modes within the same appliance. Citrix does not recommend virtual inline mode with routers that do not support health monitoring.

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router. Almost all of the configuration tasks are performed on the router. The only thing to be configured on the appliance is the forwarding method, and the default method is recommended.

Like WCCP, Virtual inline deployment requires no rewiring and no downtime, and it provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links. Unlike WCCP, it contains no built-in status monitoring or health checking, making troubleshooting difficult. WCCP is thus the recommended mode, and virtual inline is recommended only when inline and WCCP modes are both impractical.

Example

The following figure shows a simple network in which all traffic destined for or received from the remote site is redirected to the appliance. In this example, both the local site and remote site use virtual inline mode.

Figure 1. Virtual Inline Example

Following are some configuration details for the network in this example:

- Endpoint systems have their gateways set to the local router (which is not unique to virtual inline mode).
- Each router is configured to redirect both incoming and outgoing WAN traffic to the local appliance.
- Each appliance processes the traffic received from its local router and forwards it back to the router.
- PBR rules configured on the router prevent routing loops by allowing packets to make only one trip to and from the appliance. The packets that the appliance forwards back to the router are sent to their original (local or remote) destination.
- Each appliance has its default gateway set to the address of the local router, as usual (on the **Configuration: Network Adapters** page). The options for forwarding packets back to the router are Return to Ethernet Sender and Send to Gateway.

Configuring Packet Forwarding on the Appliance

Aug 09, 2017

Virtual inline mode offers two packet-forwarding options:

Return to Ethernet Sender (default)—This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

Send to Gateway (not recommended)—In this mode, virtual inline output packets are forwarded to the default gateway for delivery, even if they are destined for hosts on the local subnet. This option is usually less desirable than the Return to Ethernet Sender option, because it adds an easily forgotten element of complexity to the routing structure.

To specify the packet-forwarding option—On the Configuration: Optimization Rules: Tuning page, next to Virtual Inline, select Return to Ethernet Sender or Send to Gateway.

Router Configuration

Aug 09, 2017

The router has three tasks when supporting virtual inline mode:

1. It must forward both incoming and outgoing WAN traffic to the SD-WAN appliance.
2. It must forward SD-WAN traffic to its destination (WAN or LAN).
3. It must monitor the health of the appliance so that the appliance can be bypassed if it fails.

Policy-Based Rules

In virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the router's Ethernet ports to the appliance and creating routing rules that are based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

The basic routing algorithm is:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward it to the appliance.
- If packet is destined for the WAN, forward to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.
- Traffic shaping is not effective unless all WAN traffic passes through the appliance.

Note: When considering routing options, keep in mind that returning data, not just outgoing data, must flow through the appliance. For example, placing the appliance on the local subnet and designating it as the default router for local systems does not work in a virtual inline deployment. Outgoing data would flow through the appliance, but incoming data would bypass it. To force data through the appliance without router reconfiguration, use inline mode.

Health Monitoring

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does no health monitoring. To enable health monitoring, define a rule to monitor the appliance's availability, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Important: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy-based routing do not support health-checking. The health-monitoring feature is relatively new. It became available in Cisco IOS release 12.3(4)T.

Following is an example of a rule for monitoring the availability of the appliance:

```
!- Use a ping (ICMP echo) to see if appliance is connected track  123 rtr 1 reachabilit y ! rtr 1  type echo protocol Iplcmpecho 192.168.1.200 schedule 1 life forever start-time now
```

This rule pings the appliance at 192.168.1.200 periodically. You can test against 123 to see if the unit is up.

Routing Examples

Aug 09, 2017

The following examples illustrate configuring Cisco routers for the local and remote sites shown in Virtual inline example. To illustrate health monitoring, the configuration for the local site includes health monitoring, but the configuration for the remote site does not.

Note: The configuration for the local site assumes that a ping monitor has already been configured.
The examples conform to the Cisco IOS CLI. They might not be applicable to routers from other vendors.

Local Site, Health-Checking Enabled

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.10.10.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 172.68.1.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended client_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 192.168.1.200 20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 192.168.1.200 10 track 123
```

Remote Site (No Health Checking)

```
! This example does not use health-checking.
! Remember, health-checking is always recommended,
! so this is a configuration of last resort.
!
!
ip cef
!
interface FastEthernet0/0
ip address 20.20.20.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 171.68.2.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.2.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.2.1
!
ip access-list extended client_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
set ip next-hop 192.168.2.200
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop 192.168.2.200
!_
```

Each of the above examples applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP of 10.10.10.0/24 and destination of 20.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended client_side
 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
 ip access-list extended wan_side
 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
```

Note that, for access lists, ordinary masks are not used. Wildcard masks are used instead. Note that when reading a

wildcard mask in binary, "1" is considered a "don't care" bit.

Virtual Inline for Multiple-WAN Environments

Aug 09, 2017

Enterprises with multiple WAN links often have asymmetric routing policies, which seem to require that an inline appliance be in two places at once. Virtual inline mode solves the asymmetric routing problem by using the router configuration to send all WAN traffic through the appliance, regardless of the WAN link used. The below figure shows a simple multiple-WAN link deployment example.

The two local-side routers redirect traffic to the local appliance. The FE 0/0 ports for both routers are in the same broadcast domain as the appliance. The local appliance must use the default virtual inline configuration (Return to Ethernet Sender).

Figure 1. Virtual Inline Mode With Two WAN Routers

Virtual Inline Mode and High-Availability

Sep 10, 2017

Virtual Inline mode can be used in a high availability (HA) configuration. The below figure shows a simple HA deployment. In virtual inline mode, a pair of appliances acts as one virtual appliance. Router configuration is the same for an HA pair as with a single appliance, except that the Virtual IP address of the HA pair, not the IP address of an individual appliance, is used in the router configuration tables. In this example, the local appliances must use default virtual inline configuration (Return to Ethernet Sender).

Figure 1. High-availability Example

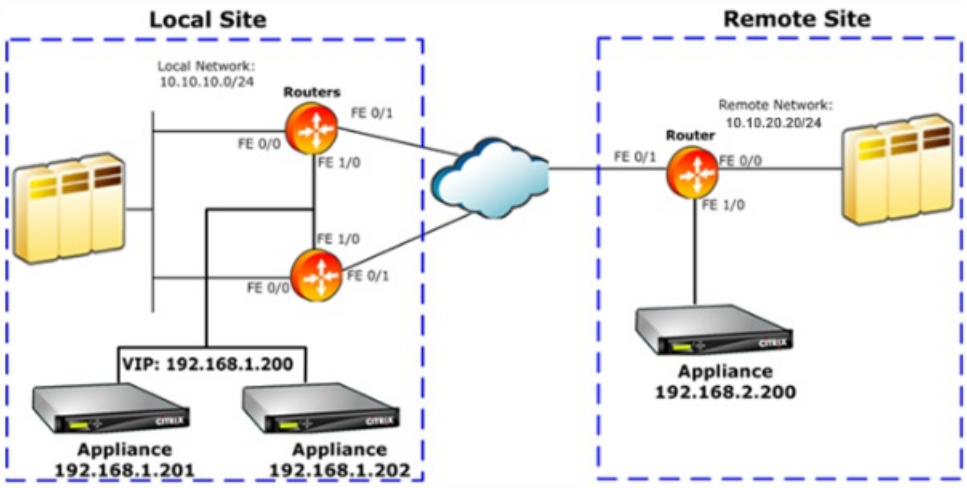

Monitoring and Troubleshooting

Aug 09, 2017

In virtual inline mode, unlike WCCP mode, the appliance provides no virtual inline-specific monitoring. To troubleshoot a virtual inline deployment, log into the appliance and use the Dashboard page to verify that traffic is flowing into and out of the appliance. Traffic forwarding failures are typically caused by errors in router configuration.

If the Monitoring: Usage or Monitoring: Connections pages show that traffic is being forwarded but no acceleration is taking place (assuming that an appliance is already installed on the other end of the WAN link), check to make sure that both incoming WAN traffic and outgoing WAN traffic are being forwarded to the appliance. If only one direction is forwarded, acceleration cannot take place.

To test health-checking, power down the appliance. The router should stop forwarding traffic after the health-checking algorithm times out.

Group Mode

Aug 09, 2017

In group mode, two or more appliances become a single virtual appliance. This mode is one solution to the problem of asymmetric routing, which is defined as any case in which some packets in a given connection pass through a given appliance but others do not. A limitation of the appliance architecture is that acceleration cannot take place unless all packets in a given connection pass through the same two appliances. Group mode overcomes this limitation.

Group mode can be used with multiple or redundant links without reconfiguring your routers.

Note: Group mode is not supported on the SD-WAN 4000 or 5000 appliances.
Group mode applies only to the appliances on one side of the WAN link; the local appliances neither know nor care whether the remote appliances are using group mode.

Group mode uses a heartbeat mechanism to verify that other members of the group are active. Packets are forwarded to active group members only.

Avoiding asymmetric routing is the main reason to use group mode, but group mode is not the only method available for that purpose. If you decide that it is the best method for your environment, you can enable it by setting a few parameters. If the default mechanism for determining which appliance is responsible for a particular connection does not provide optimal acceleration, you can change the forwarding rules.

Figure 1. Group Mode With Redundant Links

Figure 2. Group Mode With Non-Redundant Links with Possible Asymmetric Routing

Figure 3. Group Mode On Nearby Campuses

When to Use Group Mode

Aug 09, 2017

Use group mode in the following set of circumstances:

- You have multiple WAN links.
- There is a chance of asymmetric routing (a packet on a given connection might travel over either link).
- Group mode seems simpler and more practical than alternatives that use a single appliance.

The alternatives are:

- WCCP mode, in which traffic from two or more links is sent to the same appliance by WAN routers, by means of the WCCP protocol.
- Virtual inline mode, in which your routers send traffic from two or more links through the same appliance (or high-availability pair).
- Multiple bridges, where each link passes through a different accelerated bridge in the same appliance.
- LAN-level aggregation, which places an appliance (or high-availability pair) closer to the LAN, before the point where WAN traffic is split into two or more paths.

How Group Mode Works

Aug 09, 2017

In group mode, the appliances that are part of the group each take ownership for a portion of the group's connections. If a given appliance is the owner of a connection, it makes all the acceleration decisions about that connection and is responsible for compression, flow control, packet retransmission, and so on.

If an appliance receives a packet for a connection for which it is not the owner, it forwards the packet to the appliance that is the owner. The owner examines the packet, makes the appropriate acceleration decisions, and forwards any output packets back to the non-owning appliance. This process preserves the link selection made by the router, while allowing all packets in the connection to be managed by the owning appliance. For the routers, the introduction of the appliances has no consequences. The routers do not need to be reconfigured in any way, and the appliances do not need to understand the routing mechanism. They simply accept the routers' forwarding decisions.

Figure 1. Sending-side Traffic in Group Mode

Figure 2. Receiving-side traffic flow in group mode

Group mode has two, user-selectable failure modes, which control how the group members interact with each other if one of them fails. The failure mode also determines whether the failed appliance's bypass card opens (blocking traffic through the appliance) or remains closed (allowing traffic to pass through). The failure modes are:

Continue to accelerate- If a group member fails, its bypass card is opened and no traffic passes through the failed appliance. The result is presumably a fail-over if redundant links are used. Otherwise, the link is simply inaccessible. The other appliances in the group continue to accelerate. The usual hashing algorithm handles the changed conditions. (That is, the old hashing algorithm is used, and if the failed unit is indicated as the owner, a hashing algorithm based on the new, smaller group is applied. This preserves as many older connections as possible.)

Do not accelerate- If a group member fails, its bypass card closes, allowing traffic to pass through without acceleration. Because an unaccelerated path introduces asymmetric routing, the other members of the group also go into pass-through mode when they detect the failure.

Enabling Group Mode

Aug 09, 2017

To enable group mode, create a group of two or more appliances. An appliance can be a member of only one group. Group members are identified by IP address and the SSL common name in the appliance license.

All group mode parameters are on the Settings: Group Mode page, in the Configure Settings: Group Mode table.

Figure 1. Group Mode Page

To enable group mode

1. Select the address to use for group communication. At the top of the Group Mode Configuration table on the Configuration: Advanced Deployments: Group Mode tab, the table cell under Member VIP contains the management address of the port used to communicate with other group members. Use the (unlabeled) drop-down menu to select the correct address (for example, to use the Aux1 port, select the IP address you assigned to the Aux1 port). Then, click Change VIP.
2. Add at least one more group member to the list. (Groups of three or more are supported but are rarely used.) In the next cell of the Member VIP column, type the IP address of the port used by the other appliance for group-mode communication.
3. Type the other group member's SSL common name in the SSL Common Name column. The SSL common name is listed on the other appliance's Configure: Advanced Deployments: High Availability tab. If the other group member is a high-availability pair, the name listed is the SSL common name of the primary appliance.
   Note: If the local appliance is not part of a high-availability pair, the first cell in the HA Secondary SSL Common Name is blank.
   If the other group member is a high-availability pair, specify the SSL Common Name of the HA secondary appliance in the HA Secondary SSL Common Name column.
4. Click Add.
5. Repeat steps 2-4 for any additional appliances or high-availability pairs in the group.
6. The three buttons under the list of group members are toggles, so each is labeled as the opposite of its current setting:
   1. The top button reads either, **Do not accelerate when member failure is detected** or **Continue to accelerate when member failure is detected**. The "Do not accelerate..." setting always works and does not block traffic, but if any member fails, the other group members go into bypass mode, which causes a complete loss of acceleration. With the "Continue to accelerate" option, the failing appliance's bridge becomes an open circuit, and the link fails. This option is appropriate if the WAN router responds by causing a failover. New connections, and open connections belonging to the surviving appliances, are accelerated.
   2. The bottom button should now be labeled Disable Group Mode. If it is not, enable group mode by clicking the button.
7. Refresh the screen. The top of the page should list the group mode partners, but display warnings about their status, because they haven't been configured for group mode yet. For example, it might indicate that the partner cannot be found or is running a different software release.
8. Repeat this procedure with the other members of the group. Within 20 seconds after enabling the last member of the group, the Group Mode Status line should show NORMAL, and the other group mode members should be listed with Status: On-Line and Configuration: OK.

Forwarding Rules

Aug 09, 2017

By default, the *owner* of a group-mode connection is set by a hash of the source and destination IP addresses. Each appliance in the group uses the same algorithm to determine which group member owns a given connection. This method requires no configuration. The owner can optionally be specified through user-settable rules.

Because the group-mode hash is not identical to that used by load balancers, about half of the traffic tends to be forwarded to the owning appliance in a two-Appliance group. In the worst case, forwarding causes the load on the LAN-side interface to be doubled, which halves the appliance's peak forwarding rate for actual WAN traffic.

This speed penalty can be reduced if the Primary or Aux1 Ethernet ports are used for traffic between group members. For example, if you have a group of two appliances, you can use an Ethernet cable to connect the two units' Primary ports, then specify the Primary port on the Group Mode page on each unit. However, maximum performance is achieved if the amount of traffic forwarded between the group-mode members is minimized.

The owner can optionally be set according to specific IP/port-based rules. These rules must be identical on all appliances in the group. Each member of the group verifies that its group-mode configuration is identical to the others. If not all of the configurations are identical, none of the member appliances enter group mode.

If traffic arrives first at the appliance that owns the connection, it is accelerated and forwarded normally. If it arrives first at a different appliance in the group, it is forwarded to its owner over a GRE tunnel, which accelerates it and returns it to the original appliance for forwarding. Thus, group mode leaves the router's link selection unchanged.

Using explicit IP-based forwarding rules can reduce the amount of group-mode forwarding. This is especially useful in primary-link/backup-link scenarios, where each link handles a particular range of IP addresses, but can act as a backup when the other link is down.

Figure 1. IP-Based Owner Selection

Forwarding rules can ensure that group members handle only their "natural" traffic. In many installations, where traffic is usually routed over its normal link and only rarely crosses the other one, these rules can reduce overhead substantially.

Rules are evaluated in order, from top to bottom, and the first matching rule is used. Rules are matched against an optional IP address/mask pair (which is compared against both source and destination addresses), and against an optional port range.

Regardless of the ordering of rules, if the partner appliance is not available, traffic is not forwarded to it, whether a rule matches or not.

For example, in the figure below, member 172.16.1.102 is the owner of all traffic to or from its own subnet (172.16.1.0/24), while member 172.16.0.184 is the owner of all other traffic.

If a packet arrives at unit 172.16.1.102, and it is not addressed to/from net 172.16.1.0/24, it is forwarded to 172.16.0.184.

If unit 172.16.0.184 fails, however, unit 172.16.1.102 no longer forwards packets. It attempts to handle the traffic itself. This behavior can be inhibited by clicking **Do NOT Accelerate When Member Failure Detected** on the Group Mode tab.

In a setup with a primary WAN link and a backup WAN link, write the forwarding rules to send all traffic to the appliance on the primary link. If the primary WAN link fails, but the primary appliance does not, the WAN router fails over and sends traffic over the secondary link. The appliance on the secondary link forwards traffic to the primary-link appliance, and acceleration continues undisturbed. This configuration maintains accelerated connections after the link failover.

Figure 2. Forwarding Rules

Monitoring and Troubleshooting Group Mode

Aug 09, 2017

Two things should be checked in a group-mode installation:

- That the two appliances have entered group mode, which can be determined on either appliance's Configuration: Advanced Deployments: Group Mode page.
- That the behavior of the group-mode pair is as desired when the other member fails, and when one of the links fail, as determined by disabling the other appliance and temporarily disconnecting one of the links, respectively.

High-Availability Mode

Aug 09, 2017

Two identical appliances on the same subnet can be combined as a *high-availability pair.* The appliances each monitor the other's status by using the standard *Virtual Router Redundancy Protocol (VRRP)* heartbeat mechanism. The pair has a common virtual IP address for management, in addition to each appliance's management IP address. If the primary appliance fails, the secondary appliance takes over. Failover takes approximately five seconds.

High availability mode is a standard feature.

How High-Availability Mode Works

Aug 09, 2017

In a high availability (HA) pair, one appliance is primary, and the other is secondary. The primary monitors its own and the secondary's status. If it detects a problem, traffic processing fails over to the secondary appliance. Existing TCP connections are terminated. To ensure successful failover, the two appliances keep their configurations synchronized. In a WCCP mode high availability configuration, the appliance that is processing traffic maintains communication with the upstream router.

Status monitoring—When high availability is enabled, the primary appliance uses the VRRP protocol to send a heartbeat signal to the secondary appliance once per second. In addition, the primary appliance monitors the carrier status of its Ethernet ports. The loss of carrier on a previously active port implies a loss of connectivity.

Failover If the heartbeat signal of the primary appliance should fail, or if the primary appliance loses carrier for five seconds on any previously active Ethernet port, the secondary appliance takes over, becoming the primary. When the failed appliance restarts, it becomes the secondary. The new primary announces itself on the network with an ARP broadcast. MAC spoofing is not used. Ethernet bridging is disabled on the secondary appliance, leaving the primary appliance as the only path for inline traffic. Fail-to-wire is inhibited on both appliances to prevent loops.

Warning

The Ethernet bypass function is disabled in HA mode. If both appliances in an inline HA pair lose power, connectivity is lost. If WAN connectivity is needed during power outages, at least one appliance must be attached to a backup power source.

Note

The secondary appliance in the HA pair has one of its bridge ports, port apA.1, disabled to prevent forwarding loops. If the appliance has dual bridges, apB.1 is also disabled. In a one-arm installation, use port apA.2. Otherwise, the secondary appliance becomes inaccessible when HA is enabled.

Primary/secondary assignment—If both appliances are restarted, the first one to fully initialize itself becomes the primary. That is, the appliances have no assigned roles, and the first one to become available takes over as the primary. The appliance with the highest IP address on the interface used for the VRRP heartbeat is used as a tie-breaker if both become available at the same time.

Connection termination during failover—Both accelerated and unaccelerated TCP connections are terminated as a side effect of failover. Non-TCP sessions are not affected, except for the delay caused by the brief period (several seconds) between the failure of the primary appliance and the failover to the secondary appliance. Users experience the closing of open connections, but they can open new connections.

Configuration synchronization—The two appliances synchronize their settings to ensure that the secondary is ready to take over for the primary. If the configuration of the pair is changed through the browser based interface, the primary appliance updates the secondary appliance immediately.

HA cannot be enabled unless both appliances are running the same software release.

HA in WCCP mode—When WCCP is used with an HA pair, the primary appliance establishes communication with the router. The appliance uses its management IP address on apA or apB, not its virtual IP address, to communicate with the router. Upon failover, the new primary appliance establishes WCCP communication with the router.

Cabling Requirements

Jun 07, 2018

The two appliances in the high availability pair are installed onto the same subnet in either a parallel arrangement or a one-arm arrangement, both of which are shown in the following figure. In a one-arm arrangement, use the apA.2 port (and, optionally, the apB.2 port), not the apA.1 port. Some models require a separate management LAN, whether deployed in inline or one-armed mode. This is depicted only in the middle diagram.

Figure 1. Cabling for High-Availability Pairs

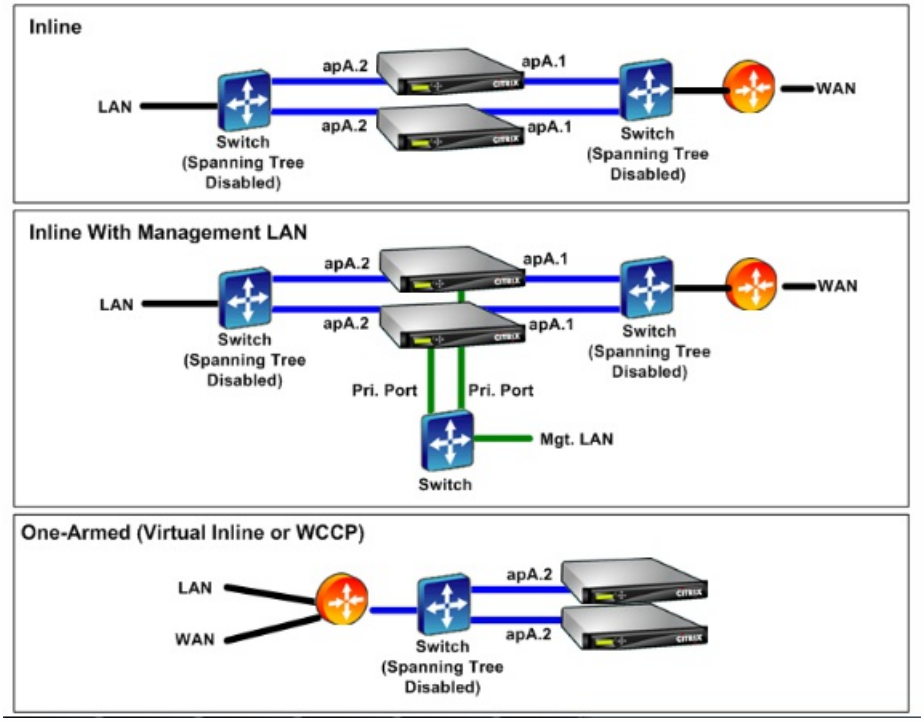

Do not break the above topology with additional switches. Random switch arrangements are not supported. Each of the switches must be either a single, monolithic switch, a single logical switch, or part of the same chassis.

If the spanning-tree protocol (STP) is enabled on the router or switch ports attached to the appliances, failover will work, but the failover time may increase to roughly thirty seconds. Without STP, failover time is roughly five seconds. Thus, to achieve the briefest possible failover interval, disable STP on the ports connecting to the appliances.

Figure 2. Ethernet Port Locations (Older Models)

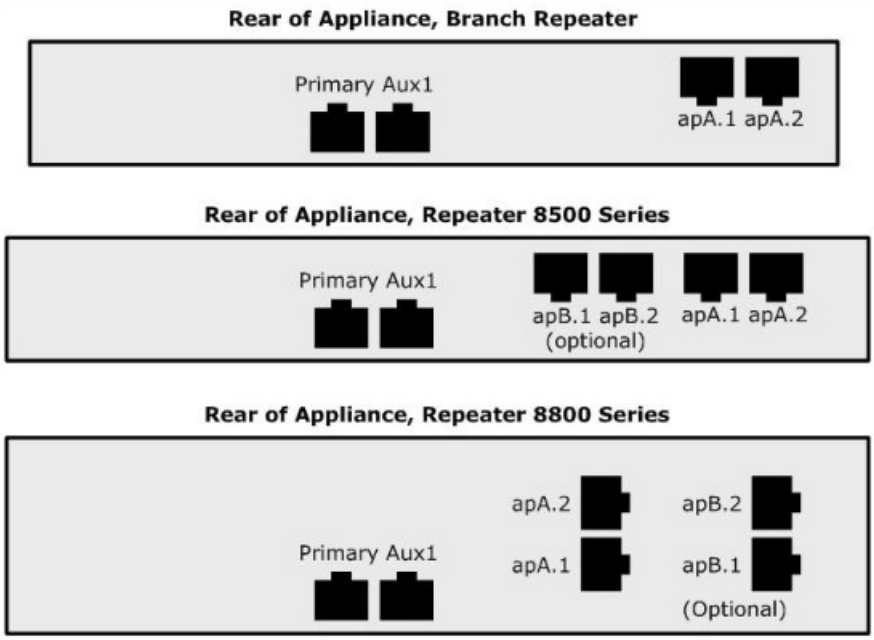

Other Requirements

Aug 09, 2017

Both appliances in an HA pair must meet the following criteria:

- Have identical hardware, as shown by on the System Hardware entry on the Dashboard page.
- Run exactly the same software release.
- Be equipped with Ethernet bypass cards. To determine what is installed in your appliances, see the Dashboard page.

Appliances that do not support HA display a warning on the Configuration: High Availability page.

Management Access to the High-Availability Pair

Aug 09, 2017

When configuring a high-availability (HA) pair, you assign the pair a virtual IP (VIP) address, which enables you to manage the two appliances as if they were a single unit. After you enable high-availability mode, managing the secondary appliance through its IP address is mostly disabled, with most parameters grayed out. A warning message displays the reason on every page. Use the HA VIP for all management tasks. You can, however, disable the secondary appliance's HA state from its management UI.

Configuring the High-Availability Pair

Aug 09, 2017

You can configure two newly installed appliances as a high-availability pair, or you can create an HA pair by adding a second appliance to an existing installation.

Prerequisites: Physical installation and basic configuration procedures

To configure high availability

1. Make sure that no more than one appliance is connected to the traffic networks (on the accelerated bridges). If both are connected, disconnect one bridge cable from the active bridges on the second appliance. This will prevent forwarding loops.
2. On the Features page of the first appliance, disable Traffic Processing. This disables acceleration until the HA pair is configured.
3. Repeat for the second appliance.
4. On the first appliance, go to the Configuration: Advanced Deployments: High Availability tab, show below.
5. Select the Enabled Check box.
6. Click the Configure HA Virtual IP Address link and assign a virtual IP address to the apA interface. This address will be used later to control both appliances as a unit.
7. Return to the High Availability page and, in the VRRP VRID field, assign a VRRP ID to the pair. Although the value defaults to zero, the valid range of VRRP ID numbers is 1 through 255. Within this range, you can specify any value that does not belong to another VRRP device on your network.
8. In the Partner SSL Common Name field, type the other appliance's SSL Common Name, which is displayed on that appliance's Configuration: Advanced Deployments: High Availability tab, in the Partner SSL Common Name field. The SSL credentials used here are factory-installed.
9. Click Update.
10. Repeat steps 3-8 on the second appliance. If you are managing the appliance via an accelerated bridge (such as apA), you may have to reconnect the Ethernet cable that you removed in step 1 to connect to the second appliance. If so, plug this cable in and disconnect the corresponding cable on the first appliance.
11. With your browser, navigate to the virtual IP address of the HA pair. Enable Traffic Processing on the Features page. Any further configuration will be performed from this virtual address.
12. Plug in the cable that was left disconnected.
13. On each appliance, the Configuration: Advanced Deployments: High Availability page should now show that high availability is active and that one appliance is the primary and the other is the secondary. If this is not the case, a warning banner appears at the top of the screen, indicating the nature of the problem.

Figure 1. High-availability configuration page

Updating Software on a High-Availability Pair

Aug 09, 2017

Updating the SD-WAN software on an HA pair causes a failover at one point during the update.

Note: Clicking the Update button terminates all open TCP connections.

To update the software on an HA pair

1. Log on to both appliances.
2. On the secondary appliance, update the software and reboot. After the reboot, the appliance is still the secondary. Verify that the installation succeeded. The primary appliance should show that the secondary appliance exists but that automatic parameter synchronization is not working, due to a version mismatch.
3. On the primary appliance, update the software, and then reboot. The reboot causes a failover, and the secondary appliance becomes the primary. When the reboot is completed, HA should become fully established, because both appliances are running the same software.

Saving/Restoring Parameters of an HA Pair

Aug 09, 2017

The System Maintenance: Backup/Restore function can be used to save and restore parameters of an HA pair as follows:

To back up the parameters

Use the backup feature as usual. That is, log on to the GUI through the HA VIP address (as is normal when managing the HA pair) and, on the System Management: Backup/Restore page, click Download Settings.

To restore the parameters

1. Disable HA on both appliances by clearing the Enabled check box on the Configuration: Advanced Deployments: High Availability (HA) tab.
2. Unplug a network cable from the bridge of one appliance. (Call it “Appliance A.”)
3. Unplug the power cord from Appliance A.
4. Restore the parameters on the other appliance (Appliance B), by uploading a previously saved set of parameters on the System Maintenance: Backup/Restore page and clicking Restore Settings. (Completing this operation requires a restart, which reenables HA).
5. Wait for Appliance B to restart. It becomes the primary.
6. Restart Appliance A.
7. Log on to Appliance A’s GUI and reenable HA on the Configuration: Advanced Deployments: High Availability (HA) tab. The appliance get its parameters from the primary.
8. Plug in the network cable removed in step 2.

Both appliances are now restored and synchronized.

Troubleshooting High Availability Pairs

Aug 09, 2017

If the appliances report any failure to enter high-availability mode, the error message will also note the cause. Some issues that can interfere with high-availability mode are:

- The other appliance is not running.
- The HA parameters on the two appliances are not identical.
- The two appliances are not running the same software release.
- The two appliances do not have the same model number.
- Incorrect or incomplete cabling between the appliances does not allow the HA heartbeat to pass between them.
- The HA/Group Mode SSL Certificates on one or both appliances are damaged or missing.

NetScaler SD-WAN 1000 and 2000 WANOP Appliances with Windows Server

Aug 09, 2017

SD-WAN supports Windows Server on two hardware platforms – 1000WS and 2000WS. Whereas 2000WS has better performance capacity as compared to the 1000WS, the latter has more RAM and hard disk space.

The SD-WAN WANOP 1000 and 2000 appliances with Windows Server combine virtualized instances of the SD-WAN WAN optimization appliance with Windows Server virtual machine installed on the appliance. The Windows Server instance is a fully-licensed version of Microsoft Windows Server 2012 R2 Standard Edition.

The SD-WAN 1000 and 2000 WANOP appliances with Windows Servers are based on the Citrix branch architecture, which supports multiple virtual machines. All branch appliances contain a SD-WAN instance, a management service instance, and a Xen hypervisor. In addition, the SD-WAN 1000 and 2000 appliances with Windows Server include a Windows Server instance, which runs independently of the SD-WAN WANOP instance.

As shown in the below figure, the Windows Server and the WANOP instance are partly isolated from one another, because the accelerated bridges are accessible only to the accelerator. This allows the accelerator and the Windows Server to be placed at different points in your LAN topology.

Figure 1. Architecture of the SD-WAN WANOP 1000 and 2000 Appliances with Windows Server

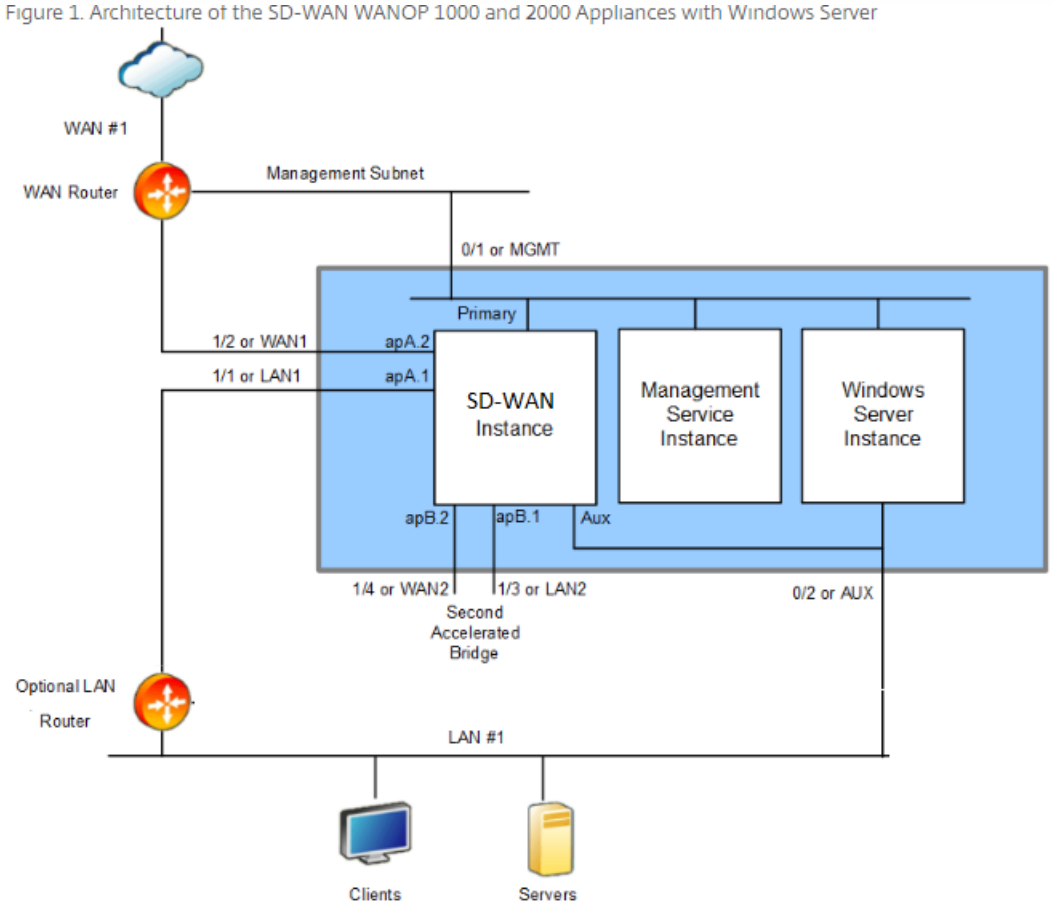

The SD-WAN WANOP instance is typically used in inline mode, with the SD-WAN instance interposed between the WAN

router and the LAN, so WAN traffic flows through the accelerated bridge. The SD-WAN WANOP instance can also be deployed in WCCP or virtual inline modes, using a single accelerated bridge port.

The Windows server is deployed in a one-armed configuration in the same local LAN in which you would deploy any other server.

In addition to the accelerated bridges and the Windows LAN port, a management port connects to all virtual machines (instances) and the hypervisor.

The appliance has two modes, two-port mode and four-port mode, which determine how ports 1/3 and 1/4 are used.

The Citrix Compliance Regulatory Models are:

- SD-WAN 1000WS WANOP: CB 504-2
- SD-WAN 2000WS WANOP: NS 6xCu

SD-WAN 1000 Appliance with Windows Server

Aug 09, 2017

The Citrix SD-WAN 1000 with Windows Server platform has a quad-core processor and 32 GB of memory. This platform has a bandwidth of up to 20 Mbps.

The following figure shows the front panel of a SD-WAN 1000 appliance with Windows Server.

Figure 1. Citrix SD-WAN 1000 with Windows Server, front panel

The front panel of the SD-WAN 1000 with Windows Server appliance has a power button and five LEDs.

The power button is used to switch the appliance on or off.

The reset button restarts the appliance.

The LEDs provide critical information related to different parts of the appliance.

- Power Fail – Indicates the power supply unit has failed.
- Information LED – Indicates the following:

Status	Description
Continuously ON and red	The appliance is overheated. (This might be a result of cable congestion.)
Blinking red (1Hz)	Fan failure, check for an inoperative fan.
Blinking red (0.25Hz)	Power failure, check for the non-operational power supply.
Solid blue	Local UID has been activated. Use this function to locate the server in a rack mount environment.
Blinking blue (300 m/s)	Remote UID is on. Use this function to identify the server from a remote location.

- NIC1 and NIC2 – Indicate network activity on the LAN1 and WAN1 ports.
- HDD – Indicates the status of the hard disk drive.
- Power – Indicates that the power supply units are receiving power and operating normally.

The following figure shows the back panel of a SD-WAN 1000 appliance with Windows Server.

Figure 2. Citrix SD-WAN 1000 appliance with Windows Server , back panel

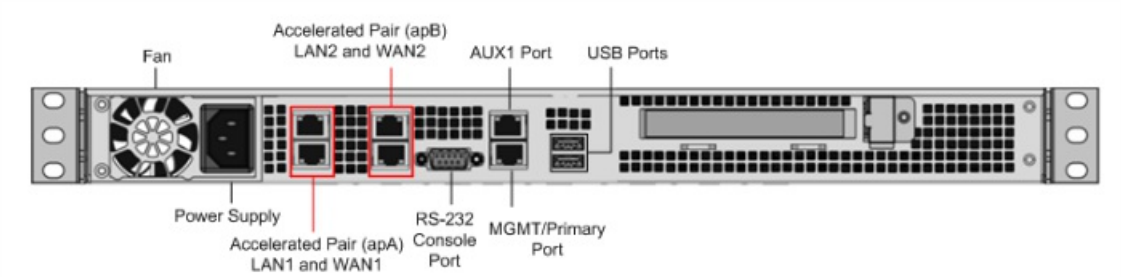

The following components are visible on the back panel of a SD-WAN 1000 appliance with Windows Server:

- Cooling fan
- Single power supply, rated at 200 watts, 110-240 volts
- Accelerated pairs of Ethernet ports (apA and apB) which function as accelerated bridges
- RS-232 serial console port
- One AUX Ethernet port and one management port
- Two USB ports

SD-WAN 2000 Appliance with Windows Server

Aug 09, 2017

The Citrix NetScaler SD-WAN 2000 with Windows Server platform is a 1U appliance with one quad-core processor and 24 gigabytes (GB) of memory.

The following figure shows the front panel of the NetScaler SD-WAN 2000 appliance with Windows Server.

Figure 1. Citrix NetScaler SD-WAN 2000 appliance with Windows Server, front panel

Note: You cannot assign apA ports to Windows Server. However, you can assign AUX port to Windows Server

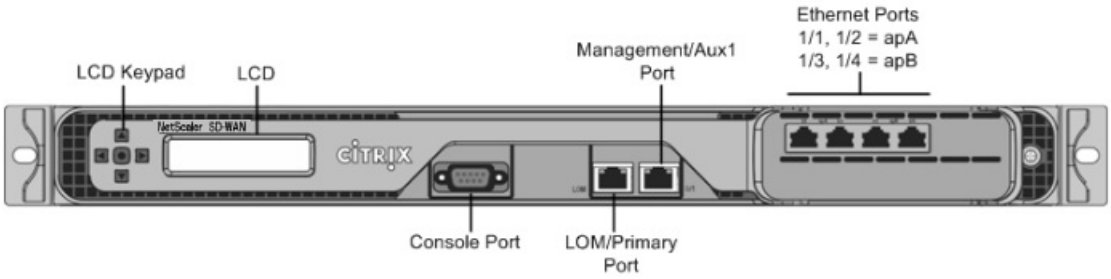

SD-WAN 2000 appliance with Windows Server has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, numbered 0/1, and named PRI (primary). The management port is used to connect directly to the appliance for system administration functions. You can use this port for initial provisioning of WAN optimization and Windows Server.
  Note: The LOM port also operates as a management port.
- Four 10/100/1000Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two *accelerated pairs*, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), and 1/3 and 1/4 are accelerated pair B (apB).

The following figure shows the back panel of the NetScaler SD-WAN 2000 appliance with Windows Server.

Figure 2. Citrix NetScaler SD-WAN 2000 appliance with Windows Server, back panel

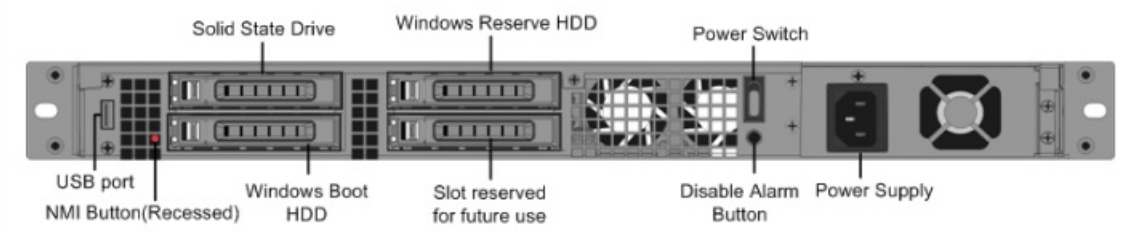

The following components are visible on the back panel of the NetScaler SD-WAN 2000 appliance with Windows Server:

- 600 GB removable solid-state drive, which stores the appliance's software and user data, and 1 TB hard disk drive.
- Power switch, which switches power to the appliance on or off. Press the switch for five seconds to switch off the

power.

- USB port (reserved for a future release).
- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You must use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Single power supply, rated at 300 watts, 100-240 volts.

Ethernet Port Names

Aug 09, 2017

When configuring the appliance, you have to specify IP addresses for various Ethernet ports of the appliance. The Ethernet ports are named differently on the front panel of NetScaler SD-WAN 1000 and 2000 appliances with Windows Server, in the NetScaler SD-WAN instance, and in the Windows Server, as shown in the following table:

Front Panel		SD-WAN Instance	Windows Server
SD-WAN 1000WS	SD-WAN 2000WS		
MGMT (Blue)	0/1 (LOM/PRI)	Primary	Citrix PV Ethernet Adapter #0: 0/1
AUX	0/2 (AUX)	Aux	Citrix PV Ethernet Adapter #1: 0/2
apA LAN1/WCCP (Green)	1/1	apA.1	N/A
apA WAN1	1/2	apA.2	N/A
apB LAN2	1/3	apB.1*	Double-click the Desktop icon nic_mapping.vbs to display the mapping**
apB WAN2	1/4	apB.2*	Double-click the Desktop icon nic_mapping.vbs to display the mapping**

* Available to the SD-WAN instance only in four-port mode.

** Available to the Windows Server only in two-port mode.

Supported Features

Aug 09, 2017

The following table lists various features supported on SD-WAN 1000 and 2000 appliances with Windows Server.

Table 1. Features Table for Citrix NetScaler SD-WAN 1000 and 2000 with Windows Server Series Appliances

	Citrix NetScaler SD-WAN 1000 with Windows Server series	Citrix NetScaler SD-WAN 2000 with Windows Server series
AutoConfiguration	Y	Y
SD-WAN Plug-In	N	Y
Compression	Y	Y
RPC over HTTPS	Y	Y
SSL Compression	Y	Y
TCP Acceleration	Y	Y
Traffic Shaping	Y	Y
Video Caching	Y	Y
Windows File System Acceleration	Y	Y
Windows Outlook Acceleration	Y	Y
XenApp/ XenDesktop Acceleration	Y	Y
Group Mode	Y	Y
High Availability Mode	Y	Y
Inline Mode	Y	Y
Virtual Inline Mode	Y	Y
WCCP Mode	Y	Y
VLANs	Y	Y

Summary of Hardware Specifications

Aug 09, 2017

The following tables summarize the specifications of the SD-WAN 1000 and 2000 with Windows Server hardware platforms.

H/W Specification	SD-WAN 1000 with Windows Server	SD-WAN 2000 with Windows Server
Windows Server version	Windows Server 2012 R2	Windows Server 2012 R2
Platform Performance		
Bandwidth	Up to 20 Mbps	Up to 50 Mbps
Maximum HDX sessions	Up to 100	300
Total sessions	10,000	20,000
Acceleration Plug-in CCUs	N/A	750
Hardware Specifications		
Processor	4 Cores	4 Cores
Total disk space	1x300 GB SSD and 1x1 TB HDD	1 x 600 GB SSD and 1X1 TB HDD
SSD (dedicated Compression history)	123 GB for Disk-Based Compression (DBC) 25 GB for video caching	225 GB for Disk-Based Compression (DBC) 50 GB for video caching
RAM	32 GB	24 GB
Network Interfaces	2 pair with bypass 10/100/1000 2 GigE ports for Management and AUX ports	4 x 10/100/1000 Base-T copper Ethernet 2 GigE ports for Management and AUX ports
Power supplies	1	1
Physical Dimensions		
Rack Units	1U	1U
System width	EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks
System depth	10" (25.4 cm)	25.4" (64.5 cm)
System weight	8.5 lbs (3.9 kg)	32 lbs (14.5 kg)

H/W Specification	SD-WAN 1000 with Windows Server	SD-WAN 2000 with Windows Server
Shipping dimensions and weight	26 L x 18.5 W x 6.5" H 14.5 lbs	32 L x 23.5 W x 7.5" H 39 lbs
Environmental and Regulatory		
Voltage	100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz
Power consumption (Max.)	200 W	300 W
Operating Temperature (degree Celsius)	10–35	0-40
Non-operating Temperature (degree Celsius)	-40 – +70	-40 – +70
Allowed Relative Humidity	8% – 90% non-condensing	5%–95%
Safety certifications	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)
Electromagnetic and susceptibility certifications	FCC (Part 15 Class A), CCC, KCC, NOM, SASO, CITC, EAC, DoC, CE, VCCI, RCM	FCC (Part 15 Class A), CCC, KCC, NOM, SASO, CITC, EAC, DoC, CE, VCCI, RCM
Environmental certifications	RoHS, WEEE	RoHS, WEEE

Installing the Appliance

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. You can also connect the appliance to a computer through Ethernet port for initial configuration. On SD-WAN 1000 appliance with Windows Server, this port is labeled as MGMT (management) port and on SD-WAN 2000 appliance with Windows Server, the port is labeled as PRI (primary) port. To complete the installation, you switch on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 09, 2017

A SD-WAN 1000 or 2000 appliance with Windows Server requires one rack unit. Both are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

Rack Mounting an SD-WAN 1000 Appliance with Windows Server

Aug 09, 2017

SD-WAN 1000 appliance with Windows Server is not shipped with rails. You can mount the appliance to the rack by using the front mounting ports.

Rack Mounting an SD-WAN 2000 Appliance with Windows Server

Oct 12, 2017

A SD-WAN 2000 appliance with Windows Server requires one rack unit. Both are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

To mount a SD-WAN appliance, you must first install the rails and then install the appliance in the rack, as follows:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

To remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the locking tabs until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

To attach the inner rails to the appliance

1. Position the right inner rail behind the ear bracket on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws.
4. Repeat steps 1 through 3 to install the left inner rail on the left side of the appliance.

To install the rack rails

1. Position the rack rails at the desired location in the rack, keeping the sliding rail guide facing inward.
2. Snap the rails to the rack. Ensure that both rack rails are at same height and that the rail guides are facing inward.

To install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides, and push the appliance into the rack rails until it locks into place.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.

Figure 1. Rack Mounting the Appliance

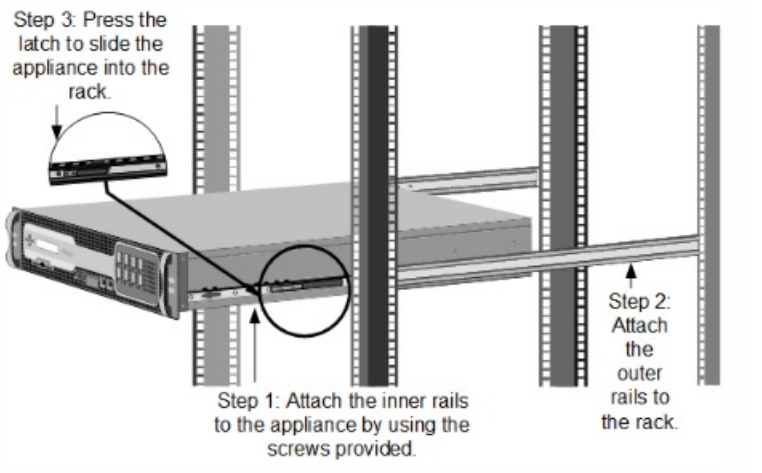

Note

The illustration above might not represent your actual appliance.

Connecting the Cables

Aug 09, 2017

When the appliance is securely mounted on the rack, determine which ports you should use. You are then ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Warning: Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ports

A typical installation using a single accelerated bridge uses four Ethernet ports (the Primary port and apA) and six IP addresses (four on the Primary port's subnet and two on apA's subnet).
The appliance has two motherboard ports and two accelerated bridges.

- The motherboard ports are labeled as MGMT (management) and AUX1 (auxiliary) ports in SD-WAN1000 appliance with Windows Server and PRI (primary) and AUX (auxiliary) in SD-WAN 2000 appliance with Windows Server. You use MGMT port of the SD-WAN 1000 appliance with Windows Server and PRI port of the SD-WAN 2000 appliance with Windows Server for initial configuration.
- Accelerated bridge ports are apA and apB are available on the back panel of SD-WAN 1000 appliance with Windows Server and the front panel of SD-WAN 2000 appliance with Windows Server. On SD-WAN 1000WS appliance with Windows Server, these ports are labeled as LAN1 and WAN1, and LAN2 and WAN2, respectively. However, on SD-WAN 2000WS appliance with Windows Server, these ports are labeled as 1/1 and 1/2, and 1/3 and 1/4, respectively.

Connecting the Ethernet Cables

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port.
   On SD-WAN **1000** appliance with Windows Server, the ports are available on the back panel and labeled as LAN1 and WAN1 for apA bridged port for LAN and WAN links, respectively.

   On SD-WAN **2000** appliance with Windows Server, the ports are available on the front panel. The ports on SD-WAN 2000 with Windows are labeled as 1/1 and 1/2 for the apA bridged port. You can use 1/1 for LAN and 1/2 for WAN link.
2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Connecting the Console Cable

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port.
   On SD-WAN **1000** appliance with Windows Server, the port is located on the back panel.

On SD-WAN 2000 appliance with Windows Server, the port is located on the front panel.

Note: To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.

2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Connecting the Power Cable

A SD-WAN appliance has one power supply. A separate ground cable is not required, because the three-prong plug provides grounding. Provide power to the appliance by installing the power cord. Connect the other end of the power cable to a standard 110V/220V power outlet.

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. After verifying the connections, you are ready to switch on the appliance.

To switch on the appliance

1. Verify that the appliance is connected through a console or Ethernet port, so that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the appliance.
3. On SD-WAN **2000** appliance for Windows Server, verify that the LCD on the front panel is backlit and the start message appears

Caution: Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs, you can quickly remove power from the appliance.

Initial Configuration

Aug 09, 2017

After checking the connections, you are ready to deploy the SD-WAN 1000 and 2000 appliances with Windows Server on the network.

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

To perform initial configuration:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Perform additional configuration for Windows.
- Assign management IP address through the serial console.
- Troubleshoot initial configuration issues.

By default, the initial configuration deploys the appliance in inline mode.

Prerequisites

Aug 09, 2017

Before you begin configuring the appliance, make sure that the following prerequisites have been met:

- You should have physical access to the appliance.
- You have chosen four IP addresses for management of the SD-WAN appliance.
- In the Worksheet, record all IP addresses and other values you would use to configure the appliance. Preferably, print out the worksheet before you start the configuration process.
- You should already have a SD-WAN license key from Citrix, sent in an email. If you are using remote licensing, you need the IP address of the licensing server.
- WAN Send and Receive Speeds.

Worksheet

Aug 09, 2017

The following table lists the default values of network configuration, SD-WAN WANOP configuration, Windows configuration, system settings, password, and setup wizard in your traffic subnet and the management subnet. Record the values applicable to your appliance in the third column.

SD-WAN WANOP 1000 or 2000 Appliance with Windows Server Deployment Worksheet			
Field	**Default Value**	**Value for your Appliance**	**Description of the Field**
Network Configuration			
XenServer IP Address (Management Subnet)	192.168.100.2		Management IP address of the XenServer.
Management Service IP Address (Management Subnet)	192.168.100.1		Management IP address of the Management Service.
Netmask (Management Subnet)	255.255.0.0		Network mask for the management subnet.
Gateway (Management Subnet)	None		The default gateway IP address of the appliance.
Port Model	2-Port		Select 2-port or 4-port, depending on the model. In 4-port mode, Windows Server does not have access to ports 1/3 and 1/4.
DNS Server	None		IP address of the DNS server. Citrix recommends that you specify a valid DNS server IP address. This is a mandatory parameter.
SD-WAN Configuration			
IP Address (Management Subnet)	192.168.100.20		Primary SD-WAN IP address at which you manage the SD-WAN instance.
Netmask (Management Subnet)	255.255.0.0		Network mask for the management IP address of the appliance. Same as previous netmask.
Gateway (Management Subnet)	None		The default gateway IP address of the appliance. Same as the previous gateway.
Windows Configuration			
IP Address	192.168.100.40		IP address to manage Windows Virtual Machine.

(Management Subnet)			
SD-WAN WANOP 1000 or 2000 Appliance with Windows Server Deployment Worksheet			
Netmask (Management Subnet)	255.255.0.0		Network mask for the management IP address of the appliance. Same as the previous netmask.
Gateway (Management Subnet)	None		The default gateway IP address of the appliance. Same as the previous gateway.
System Settings			
NTP Server	(none)		IP address of the NTP server. Citrix recommends that you specify a valid NTP server IP address. You can either enter the IP address or the server name.
Time Zone	UTC		Specify the time zone for your location.
Password			
Password	nsroot		New password for access to the appliance.
Confirm Password	nsroot		New password for access to the appliance.
Command Center Configuration			
Command Center IP Address	None		Optional. IP address of the Command Center appliance with which you want to register this appliance. More info.
Command Center Port	8443		Optional. Port number of the Command Center appliance.
Registration Password	None		Password you want to use to register the SD-WAN appliance.
Licensing			
License Server Address	None		IP address of the licensing server. Required only when you select a remote model license type.
Licensing Service Port	27000		Port number of the licensing server. Required only when you select a remote model license type.
Links			
Receive (Download) Speed	None		WAN link download speed.
Send (Upload) Speed	None		WAN link upload speed.

Configuring the Appliance by Connecting a Computer to the Ethernet Port

Aug 09, 2017

For initial configuration of a SD-WAN appliance, perform the following tasks::

- Configure the appliance for use on your site.
- Install the Citrix license.
- Enable acceleration.
- Enable traffic shaping (inline mode only).

With inline deployments, this configuration might be all you need, because most acceleration features are enabled by default and require no additional configuration.

You can configure the appliance connecting the appliance to your computer through either the Ethernet port or the serial console. The following procedure enables you to configure the appliance by connecting it to your computer through the Ethernet port.

Note: On a SD-WAN 1000 appliance with Windows Server, you use the Ethernet port labeled as MGMT. However, on SD-WAN 2000 appliance with Windows Server, you use the Ethernet port labeled as PRI or LOM.
If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 25 of the following procedure.

Note: Make sure that you have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.50, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields as blank.

2. Using an Ethernet cable, connect this computer to the port labeled MGMT on a SD-WAN WANOP 1000 appliance with Windows Server, or to the port labeled PRI on a SD-WAN WANOP 2000 appliance with Windows Server.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance:
   Username: nsroot

   Password: nsroot.

5. Start the configuration wizard by clicking **Get Started**.
6. On the Platform Configuration page, enter the respective values from your worksheet, as shown in the following example:

   Note: If, for SD-WAN configuration, you want to use the same network mask and gateway as those for Network Configuration, select the **Use System Netmask and Gateway** option.
7. Click **Done**. A screen showing the Installation in Progress… message appears. This process takes approximately 2 to 5

minutes, depending on your network speed.
Note: If you are configuring the appliance by connecting it to your computer through the serial console port, skip step 8 through step 14.

8. A Redirecting to new management IP message appears.
9. Click **OK**.
10. Unplug your computer from the Ethernet port and connect the port to your management network.
11. Reset the IP address of your computer to its previous setting.
12. From a computer on the management network, log on to the appliance by entering the new Management Service IP address, such as https://<Managemnt_IP_Address>, in a web browser.
13. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
14. Log on to the appliance by using the **nsroot** user name and the password from your worksheet.
15. The Configuration wizard starts again. In this wizard, some of the values which you have already provided, appear by default. Specify rest of the values you have recorded in your worksheet.
16. If you want to manage the appliance through Command Center, specify the Command Center IP address, port, and registration password in the **Command Center Configuration** page. Otherwise, skip this step.
17. In System Services section, update the values if necessary.
18. In the **Licensing** section, select the appropriate license type. You can either select a local license or a remote license server to apply a license to the appliance.
    1. If you opt for a local license, you must generate a license by using the host ID of the appliance. To generate a local license for the appliance, see http://support.citrix.com/article/ctx131110. To apply the license, you can navigate to the **SD-WAN** > **Configuration** > **Appliance Settings** > **Licensing page**, after completing the Configuration wizard.
    2. If you opt for a remote licensing server, you must select a remote appliance model and provide the IP address of the licensing server in the **Licensing Server Address** field.
19. In the WAN Link Definition section, specify receive and send speeds for the WAN link in the respective fields. Citrix recommends values 10% lower than the WAN bandwidth, to avoid network congestion.
20. By default, WAN-side adapter settings are configured on the appliance. Accept the default settings.
21. Click **Install**. After the Installation process is complete, the appliance restarts.
22. As soon as the appliance restarts, the Dashboard page appears.

23. To configure the appliance to accelerate the network traffic, open navigate to the **Configuration** tab.
    Note: Make sure that you have already applied the appropriate license to the appliance.
24. On the Network Adapters page of the Appliance Settings node, verify and, if necessary, assign IP addresses, subnet masks, and gateways to the accelerated bridges (apA and apB) to be used. Applying these changes restarts the appliance.
    Note: You need to assign IP addresses to apA and apB adapters only if you intended to configure WCCP mode, virtual inline mode, or the Video Caching feature on the appliance.
25. The Initial Configuration is complete. Traffic now flows through the appliance. The Dashboard page shows this traffic.

26. You need additional configuration on the appliance if you intend to use some of the modes and features, such as WCCP mode, virtual inline mode, video caching, secure peering, high availability, encrypted CIFS/MAPI acceleration, AppFlow monitoring, or SNMP monitoring.

Note:

- Inline installations place the appliance between your LAN and WAN routers, using both ports of the accelerated bridge,

such as ports LAN1 and WAN1 on a SD-WAN 1000 appliance with Window Server or ports 1/1 and 1/2 on SD-WAN 2000 appliance with Windows Server, for the apA accelerated bridge port.

- WCCP and virtual inline installations connect a single accelerated bridge port to your WAN router.
- Virtual inline installations require that you configure your router to forward WAN traffic to the appliance. See Router Configuration.
- WCCP installations require configuration of your router and the appliance. See WCCP Mode.

Additional Configuration for Windows

Aug 09, 2017

You can use the Windows management tools to perform additional configuration, such as configuring ports and other Windows services, on the Windows instance.

To perform additional configurations for Windows

1. Open a remote desktop (RDP) session to the IP address of the Windows instance from your Worksheet.
2. Log on to the Windows instance with the following credentials:
   - **Username**: Administrator
   - **Password**: password
3. Use interface AUX for Windows Server traffic. This port has a Windows Device Description of "Citrix PV Ethernet Adapter #1: 0/2." Set it to use an IP address and network mask in the network that you chose for the Windows adapter.
4. Enable Windows services for access to services, such as domain services, printer definitions, and user rights.
5. Define only a single default gateway. Add non-default routes as appropriate for your installation.

Assigning a Management IP Address through the Serial Console

Aug 09, 2017

If you do not want to change the settings of your computer, you can perform initial configuration by connecting the appliance to your computer with a serial null modem cable. Make sure that you have physical access to the appliance.

To configure the appliance through the serial console

1. Connect a serial null modem cable to the appliance's console port.
2. Connect the other end of the cable to the serial COM port of a computer running a terminal emulator, such as Microsoft HyperTerminal, with settings 9600,N,8,1, p.
3. On the HyperTerminal output, press **Enter**. The terminal screen displays the Logon prompt.
   Note: You might have to press **Enter** two or three times, depending on the terminal program you are using.
4. At the logon prompt, log on to the appliance with the following default credentials:
   Username: nsroot

   Password: nsroot.

5. At the **$** prompt, run the following command to switch to the shell prompt of the appliance:
   $ ssh 169.254.0.10

6. Enter **Yes** to continue connecting to the management service.
7. Log on to the shell prompt of the appliance with the following default credentials:
   Password: nsroot.

8. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu:
   # networkconfig

9. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.
10. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the XenServer server.
11. Type **3** and press **Enter** to select option 3, and then specify the network mask for the management service IP address.
12. Type **4** and press **Enter** to select option 4, and then specify the default gateway for the management service IP address.
13. Type **8** and press **Enter** to save the settings and exit.
14. Access the SD-WAN appliance by entering the new management service IP address of the appliance, such as https://<Management_Service_IP_Address>, in a web browser of a computer on the management network.
15. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
16. Run steps 4 through 25 of the Configuring the Appliance by Connecting a Computer to the Ethernet Port procedure to complete the configuration process

Deployment Modes

Aug 09, 2017

A SD-WAN appliance acts as a virtual gateway. It is neither a TCP endpoint nor a router. Like any gateway, its job is to buffer incoming packets and put them onto the outgoing link at the right speed. This packet forwarding can be done in different ways, such as inline mode, virtual inline mode, and WCCP mode. Although these methods are called *modes*, you do not have to disable one forwarding mode to enable another. If your deployment supports more than one mode, the mode that the appliance uses is determined automatically by the Ethernet and IP format of each packet.

Because the appliance supports different forwarding modes and different kinds of non-forwarded connections, it needs a way of distinguishing one kind of traffic from another. It does so by examining the destination IP address and destination Ethernet address (MAC address), as shown in table below. For example, in inline mode, the appliance is acting as a bridge. Unlike other traffic, bridged packets are addressed to a system beyond the appliance, not to the appliance itself. The address fields contain neither the appliance's IP address nor the appliance's Ethernet MAC address.

In addition to pure forwarding modes, the appliance has to account for additional types of connections, including management connections to the GUI and the heartbeat signal that passes between members of a high-availability pair. For completeness, these additional traffic modes are also listed in table below.

Table 1. How Ethernet and IP Addresses Determine the Mode

Destination IP Address	Destination Ethernet Address	Mode
Not appliance	Not appliance	Inline or Pass-through
Not appliance	Appliance	Virtual Inline or L2 WCCP
Appliance	Appliance	Direct (UI access)
Appliance (VIP)	Appliance	High-Availability. Proxy mode
Appliance (WCCP GRE Packet)	Appliance	WCCP GRE Mode
Appliance (Signaling IP)	Appliance	Signaling Connection (SD-WAN plugin Signaling Connection (SD-WAN plugin, Secure Peer) or Redirector Mode Connection (SD-WAN plugin)
All modes can be active simultaneously. The mode used for a given packet is determined by the Ethernet and IP headers.		

The forwarding modes are:

- **Inline mode,** in which the appliance transparently accelerates traffic flowing between its two Ethernet ports. In this mode, the appliance appears (to the rest of the network) to be an Ethernet bridge. Inline mode is recommended, because it requires the least configuration.
- **WCCP mode,** which uses the WCCP v. 2.0 protocol to communicate with the router. This mode is easy to configure on most routers. WCCP has two variants: WCCP-GRE and WCCP-L2. WCCP-GRE encapsulates the WCCP traffic within

generic routing encapsulation (GRE) tunnels. WCCP-L2 uses un-encapsulated network Layer 2 (Ethernet) transport.

- **Virtual inline mode,** in which a router sends WAN traffic to the appliance and the appliance returns it to the router. In this mode, the appliance appears to be a router, but it uses no routing tables. It sends the return traffic to the real router. Virtual inline mode is recommended when inline mode and high-speed WCCP operation are not practical.
- **Group mode,** which allows two appliances to operate together to accelerate a pair of widely separated WAN links.
- **High availability mode,** which allows to appliances to operate as an active/standby high availability pair. If the primary appliance fails, the secondary appliance takes over.

Additional traffic types are listed here for completeness:

- **Pass-through traffic** refers to any traffic that the appliance does not attempt to accelerate. It is a traffic category, not a forwarding mode.
- **Direct access,** where the appliance acts as an ordinary server or client. The GUI and CLI are examples of direct access, using the HTTP, HTTPS, SSH, or SFTP protocols. Direct access traffic can also include the NTP and SNMP protocols.
- **Appliance-to-appliance communication**, which can include signaling connections (used in secure peering and by the SD-WAN plugin), VRRP heartbeats (used in high-availability mode), and encrypted GRE tunnels (used by group mode).
- **Deprecated modes.** Proxy mode and redirector mode are legacy forwarding modes that should not be used in new installations.

Customizing the Ethernet ports

Aug 09, 2017

A typical appliance has four Ethernet ports: two accelerated bridged ports, called *accelerated pair A* (apA.1 and apA.2), with a bypass (fail-to-wire) relay, and two unaccelerated motherboard ports, called Primary and Aux1. The bridged ports provide acceleration, while the motherboard ports are sometimes used for secondary purposes. Most installations use only the bridged ports.

Some SD-WAN units have only the motherboard ports. In this case, the two motherboard ports are bridged.

The appliance's user interface can be accessed by a VLAN or non-VLAN network. You can assign a VLAN to any of the appliance's bridged ports or motherboard ports for management purposes.

Figure 1. Ethernet Ports

Port List

The ports are named as follows:

Table 1. Ethernet Port Names

Motherboard port 1	Primary (or apA.1 if no bypass card is present)
Motherboard port 2	Auxiliary1 or Aux1 (or apA.2 if no bypass card is present)
Bridge #1	Accelerated Pair A (apA, with ports apA.1 and apA.2)
Bridge #2	Accelerated Pair B (apB, with ports apB.1 and apB.2)

Port Parameters

Aug 09, 2017

Each bridge and motherboard port can be:

- Enabled or disabled
- Assigned an IP address and subnet mask
- Assigned a default gateway
- Assigned to a VLAN
- Set to 1000 Mbps, 100 Mbps, or 10 Mbps
- Set to full duplex, half-duplex, or auto (on SD-WAN 4000/5000 appliances, some ports can be set to 10 Gbps)

All of these parameters except the speed/duplex setting are set on the Configuration: IP Address page. The speed/duplex settings are set on the Configuration: Interface page.

Notes about parameters:

- Disabled ports do not respond to any traffic.
- The browser-based UI can be enabled or disabled independently on all ports.
- To secure the UI on ports with IP addresses, select HTTPS instead of HTTP on the Configuration: Administrator Interface: Web Access page.
- Inline mode works even if a bridge has no IP address. All other modes require that an IP address be assigned to the port.
- Traffic is not routed between interfaces. For example, a connection on bridge apA does not cross over to the Primary or Aux1 ports, but remains on bridge apA. All routing issues are left to your routers.

Accelerated Bridges (apA and apB)

Aug 09, 2017

Every appliance has at least one pair of Ethernet ports that function as an accelerated bridge, called *apA* (for *accelerated pair A*). A bridge can act in inline mode, functioning as a transparent bridge, as if it were an Ethernet switch. Packets flow in one port and out the other. Bridges can also act in one arm mode, in which packets flow in one port and back out the same port.

An appliance that has a bypass card maintains network continuity if a bridge or appliance malfunctions.

Some units have more than one accelerated pair, and these additional accelerated pairs are named apB, apC, and so on.

Bypass Card

If the appliance loses power or fails in some other way, an internal relay closes and the two bridged ports are electrically connected. This connection maintains network continuity but makes the bridge ports inaccessible. Therefore you might want to use one of the motherboard ports for management access.

Caution: Do not enable the Primary port if it is not connected to your network. Otherwise, you cannot access the appliance, as explained in Ethernet Bypass and Link-Down Propagation
Bypass cards are standard on some models and optional on others. Citrix recommends that you purchase appliances with bypass cards for all inline deployments.

The bypass feature is wired as if a cross-over cable connected the two ports, which is the correct behavior in properly wired installations.

Important: Bypass installations must be tested - Improper cabling might work in normal operation but not in bypass mode. The Ethernet ports are tolerant of improper cabling and often silently adjust to it. Bypass mode is hard-wired and has no such adaptability. Test inline installations with the appliance turned off to verify that the cabling is correct for bypass mode.

Using Multiple Bridges

If the appliance is equipped with two accelerated bridges, they can be used to accelerate two different links. These links can either be fully independent or they can be redundant links connecting to the same site. Redundant links can be either load-balanced or used as a main link and a failover link.

Figure 1. Using dual bridges

When it is time for the appliance to send a packet for a given connection, the packet is sent over the same bridge from which the appliance received the most recent input packet for that connection. Thus, the appliance honors whatever link decisions are made by the router, and automatically tracks the prevailing load-balancing or main-link/failover-link algorithm in real time. For non-load-balanced links, the latter algorithm also ensures that packets always use the correct bridge.

WCCP and Virtual Inline Modes

Multiple bridges are supported in both WCCP mode and virtual inline mode. Usage is the same as in the single-bridge case, except that WCCP has the additional limitation that all traffic for a given WCCP service group must arrive on the same bridge.

High Availability with Multiple Bridges

Two units with multiple bridges can be used in a high-availability pair. Simply match up the bridges so that all links pass through both appliances.

Motherboard Ports

Aug 09, 2017

Although the Ethernet ports on a bypass card are inaccessible when the bypass relay is closed, the motherboard ports remain active. You can sometimes access a failed appliance through the motherboard ports if the bridged ports are inaccessible.

The Primary Port

If the Primary port is enabled and has an IP address assigned to it, the appliance uses that IP address to identify itself to other acceleration units. This address is used internally for a variety of purposes, and is most visible to users as the Partner Unit field on the Monitoring: Optimization: Connections page. If no motherboard port is enabled, the appliance uses the IP address of Accelerated Pair A.

The Primary port is used for:

- Administration through the web based UI
- A back channel for group mode
- A back channel for high-availability mode

The Aux1 Port

The Aux1 port is identical to the Primary port. If the Aux1 port is enabled and the Primary port is not, the appliance takes its identity from the Aux1 port's IP address. If both are enabled, the Primary port's IP address is the unit's identity

VLAN Support

Aug 09, 2017

A virtual local area network (VLAN) uses part of the Ethernet header to indicate which virtual network a given Ethernet frame belongs to. SD-WAN appliances support VLAN trunking in all forwarding modes (inline, WCCP, virtual inline, and group mode). Traffic with any combination of VLAN tags is handled and accelerated correctly.

For example, if one traffic stream passing through the accelerated bridge is addressed to 10.0.0.1, VLAN 100, and another is addressed to 10.0.0.1, VLAN 111, the appliance knows that these are two distinct destinations, even though the two VLANs have the same IP address.

You can assign a VLAN to all, some, or none of the appliance's Ethernet ports. If a VLAN is assigned to a port, the management interfaces (GUI and CLI) listen only to traffic on that VLAN. If no VLAN is assigned, the management interfaces listen only to traffic without a VLAN. This selection is made on the Configuration: Appliance Settings: Network Adapters: IP Addresses tab.

Inline Mode

Aug 09, 2017

In inline mode, traffic passes into one of the appliance's Ethernet ports and out of the other. When two sites with inline appliances communicate, every TCP connection passing between them is accelerated. All other traffic is passed through transparently, as if the appliance were not there.

Figure 1. Inline mode, Accelerating All Traffic on a WAN

Note: Any TCP-based traffic passing through both units is accelerated. No address translation, proxying, or per-site setup is required. Inline mode is auto-detecting and auto-configuring.
Configuration is minimized with inline mode, because your WAN router need not be aware of the appliance's existence.

Depending on your configuration, inline mode's link-down propagation can affect management access to the appliance if a link goes down.

Inline mode is most effective when applied to all traffic flowing into and out of a site, but it can be used for only some of the site's traffic.

Ethernet Bypass and Link-Down Propagation

Aug 09, 2017

Note: Link-Down propagation was added to the SD-WAN 2000, 3000, 4000, and 5000 appliances with the 7.2.1 release. Most appliance models include a "fail-to-wire" (Ethernet bypass) feature for inline mode. If power fails, a relay closes and the input and output ports become electrically connected, allowing the Ethernet signal to pass through from one port to the other as if the appliance were not there. In fail-to-wire mode, the appliance looks like a cross-over cable connecting the two ports.

Any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the appliance's bridge ports are inaccessible.

If carrier is lost on one of the bridge ports, the carrier is dropped on the other bridge port to ensure that the link-down condition is propagated to the device on the other side of the appliance. Units that monitor link state (such as routers) are thus notified of conditions on the other side of the bridge.

Link-down propagation has two operating modes:

- If the Primary port is not enabled, the link-down state on one bridge port is mirrored briefly on the other bridge port, and then the port is re-enabled. This allows the appliance to be reached through the still-connected port for management, HA heartbeat, and other tasks.
- If the Primary port is enabled, the appliance assumes (without checking) that the Primary port is used for management, HA heartbeat, and other tasks. The link-down condition on one bridge port is mirrored persistently on the other port, until carrier is restored or the unit is rebooted. This is true even if the Primary port is enabled in the GUI but not connected to a network, so the Primary port should be disabled (the default) when not in use.

Accelerating an Entire Site

Aug 09, 2017

Inline mode, Accelerating All Traffic on a WAN shows a typical configuration for inline mode. For both sites, the appliances are placed between the LAN and the WAN, so all WAN traffic that can be accelerated is accelerated. This is the simplest method for implementing acceleration, and it should be used when practical.

Because all the link traffic is flowing through the appliances, the benefits of fair queuing and flow control prevent the link from being overrun.

In IP networks, the bottleneck gateway determines the queuing behavior for the entire link. By becoming the bottleneck gateway, the appliance gains control of the link and can manage it intelligently. This is done by setting the bandwidth limit slightly lower than the link speed. When this is done, link performance is ideal, with minimal latency and loss even at full link utilization.

Partial-Site Acceleration

Aug 09, 2017

To reserve the appliance's accelerated bandwidth for a particular group of systems, such as remote backup servers, you can install the appliance on a branch network that includes only those systems. This is shown in the following figure.

Figure 1. Inline Mode, Accelerating Selected Systems Only

SD-WAN traffic shaping relies on controlling the entire link, so traffic shaping is not effective with this topology, because the appliance sees only a portion of link traffic. Latency control is up to the bottleneck gateway, and interactive responsiveness can suffer.

Configuring and Troubleshooting Inline Mode

Aug 09, 2017

Inline mode requires only basic configuration, because it is applied automatically to any packets passing through the accelerated bridge. Troubleshooting is described under .

WCCP Mode

Aug 09, 2017

Web Cache Communication Protocol (WCCP) is a dynamic routing protocol introduced by Cisco. Originally intended only for web caching, WCCP version 2 became a more general-purpose protocol, suitable for use by accelerators such as Citrix NetScaler SD-WAN appliances.

WCCP mode is the simplest way of installing a SD-WAN appliance when inline operation is impractical. It is also useful where asymmetric routing occurs, that is, when packets from the same connection arrive over different WAN links. In WCCP mode, the routers use the WCCP 2.0 protocol to divert traffic through the appliance. Once received by the appliance, the traffic is treated by the acceleration engine and traffic shaper as if it were received in inline mode.

Note:

- For the purposes of this discussion, WCCP version 1 is considered obsolete and only WCCP version 2 is presented.
- The standard WCCP documentation calls WCCP clients "caches." To avoid confusion with actual caches, Citrix generally avoids calling a WCCP client a "cache." Instead, WCCP clients are typically called "appliances."
- This discussion uses the term "router" to indicate WCCP-capable routers and WCCP-capable switches. Though the term "router" is used here, some high-end switches also support WCCP, and can be used with SD-WAN WANOP appliances.

The SD-WAN WANOP appliances support two WCCP modes:

- WCCP is the original SD-WAN WANOP WCCP offering supported since release 3.x. It supports a single appliance service group (no clustering).
- WCCP clustering, introduced in release 7.2, allows your router to load-balance traffic between multiple appliances.

How WCCP Mode Works

The physical mode for WCCP deployment of a SD-WAN WANOP appliance is one-arm mode in which the SD-WAN appliance is connected directly to a dedicated port on the WAN router. The WCCP standard includes a protocol negotiation in which the appliance registers itself with the router, and the two negotiate the use of features they support in common. Once this negotiation is successful, traffic is routed between the router and the appliance according to the WCCP router and redirection rules defined on the router.

A WCCP-mode appliance requires only a single Ethernet port. The appliance should either be deployed on a dedicated router port (or WCCP-capable switch port) or isolated from other traffic through a VLAN. Do not mix inline and WCCP modes.

The following figure shows how a router is configured to intercept traffic on selected interfaces and forward it to the WCCP-enabled appliance. Whenever the WCCP-enabled appliance is not available, the traffic is not intercepted, and is forwarded normally.

Figure 1. WCCP Traffic Flow

Traffic Encapsulation

WCCP allows traffic to be forwarded between the router and the appliance in either of the following modes:

- L2 Mode—Requires that the router and appliance be on the same L2 segment (typically an Ethernet segment). The IP packet is unmodified, and only the L2 addressing is altered to forward the packet. In many devices, L2 forwarding is

performed at the hardware layer, giving it the maximum performance. Because of its performance advantage, L2 forwarding is the preferred mode, but not all WCCP-capable devices support it.

- GRE Mode—Generic Routing Encapsulation (GRE) is a routed protocol and the appliance can in theory be placed anywhere, but for performance it should be placed close to the router, on a fast, uncongested path that traverses as few switches and routers as possible. GRE is the original WCCP mode. A GRE header is created and the data packet is appended to it. The receiving device removes the GRE header. With encapsulation, the appliance can be on a subnet that is not directly attached to the router. However, both the encapsulation process and the subsequent routing add CPU overhead to the router, and the addition of the 28-byte GRE header can lead to packet fragmentation, which adds additional overhead.

WCCP mode supports multiple routers and both GRE vs. L2 forwarding. Each router can have multiple WAN links. Each link can have its own WCCP service group.

Traffic shaping is not effective unless the appliance manages UDP traffic as well as TCP traffic. A second service group, with a UDP service group for each WAN link, is recommended if traffic shaping is desired.

Registration and Status Updates

A WCCP client (an appliance) uses UDP port 2048 to register itself with the router and to negotiate which traffic should be sent to it, and also which WCCP features should be used for this traffic. The appliance operates on this traffic and forwards the resulting traffic to the original endpoint. The status of an appliance is tracked through the WCCP registration process and a heartbeat protocol. The appliance first contacts the router over the WCCP control channel (UDP port 2048), and the appliance and router exchange information with packets named “Here_I_Am” and “I_See_You,” respectively. By default, this process is repeated every ten seconds. If the router fails to receive a message from the appliance for three of these intervals, it considers the appliance to have failed and stops forwarding traffic to it until contact is reestablished.

Services and Service Groups

Different appliances using the same router can provide different services. To keep track of which services are assigned to which appliances, the WCCP protocol uses a service group identifier, a one-byte integer. When an appliance registers itself with a router, it includes service group numbers as well.

- A single appliance can support more than one service group.
- A single router can support more than one service group.
- A single appliance can use the same service group with more than one router.
- A single router can use the same service group with more than one appliance. For SD-WAN appliances, multiple appliances are supported in WCCP cluster mode, and a single appliance is supported in WCCP mode.
- Each appliance specifies a “return type” (L2 or GRE) independently for each direction and each service group. SD-WAN WANOP 4000/5000 appliances always specify the same return type for both directions. Other SD-WAN appliances allow the return type to be different.

Figure 2. Using different WCCP service groups for different services

Multiple service groups can be used with WCCP on the same appliance. For example, the appliance can receive service-group 51 traffic from one WAN link and service-group 62 traffic from another WAN link. The appliance also supports multiple routers. It is indifferent to whether all the routers use the same service group or different routers use different service groups.

Service Group Tracking. If a packet arrives on one service group, output packets for the same connection are sent on the

same service group. If packets arrive for the same connection on multiple service groups, output packets track the most recently seen service group for that connection.

High Availability Behavior

When WCCP is used with high-availability mode, the primary appliance sends its own apA or apB management IP address, not the virtual address of the HA pair, when it contacts the router. If failover occurs, the new primary appliance contacts the router automatically, reestablishing the WCCP channel. In most cases the WCCP timeout period and the HA failover time overlap. As a result, the network outage is less than the sum of the two delays.

Standard WCCP allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive. WCCP clustering allows multiple appliances per service group.

Deployment Topology

The following figure shows a simple WCCP deployment, suitable for either L2 or GRE. The traffic port (1/1) is connected directly to a dedicated router port (Gig 4/12).
Figure 3. Simple WCCP deployment

In this example, the SD-WAN WANOP 4000/5000 is deployed in one-arm mode, with the traffic port (1/1) and the management port (0/1) each connecting to its own dedicated router port.

On the router, WCCP is configured with identical ip wccp redirect in statements on the WAN and LAN ports. Two service groups are used, 71 and 72. Service group 71 is used for TCP traffic and service group 72 is used for UDP traffic. The SD-WAN appliance does not accelerate UDP traffic, but can apply traffic shaping policies to it.

Note: The WCCP specification does not allow protocols other than TCP and UDP to be forwarded, so protocols such as ICMP and GRE always bypass the appliance.

WCCP Clustering

SD-WAN release 7.2 or later supports WCCP clustering, which enables your router to load-balance your traffic between multiple appliances. For more information about deploying SD-WAN appliances as a cluster, see WCCP Clustering.

WCCP Specification

For more information about WCCP, see Web Cache Communication Protocol V2, Revision 1, http://tools.ietf.org/html/draft-mclaggan-wccp-v2rev1-00.

WCCP Mode (Non-Clustered)

Aug 09, 2017

WCCP mode allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive.

Note: WCCP clustering allows multiple appliances per service group.

Limitations and Best Practices

Following are limitations and best practices for (non-clustered) WCCP mode:

- On appliances with more than one accelerated pair, all the traffic for a given WCCP service group must arrive on the same accelerated pair.
- Do not mix inline and WCCP traffic on the same appliance. The appliance does not enforce this guideline, but violating it can cause difficulties with acceleration. (WCCP and virtual inline modes can be mixed, but only if the WCCP and virtual inline traffic are coming from different routers.)
- For sites with a single WAN router, use WCCP whenever inline mode is not practical.
- Only one appliance is supported per service group. If more than one appliance attempts to connect to the same router with the same service group, the negotiation will succeed only for the first appliance.
- For sites with multiple WAN routers serviced by the same appliance, WCCP can be used to support one, some, or all of your WAN routers. Other routers can use virtual inline mode.

Router Support for WCCP

Aug 09, 2017

Configuring the router for WCCP is very simple. WCCP version 2 support is included in all modern routers, having been added to the Cisco IOS at release 12.0(11)S and 12.1(3)T. The best router-configuration strategy is determined by the characteristics of your router and switches. Traffic shaping requires two service groups.

If your router supports Reverse Path Forwarding, you must disable it on all ports, because it can confuse WCCP traffic with spoofed traffic. This feature is found in newer Cisco routers such as the Cisco 7600.

Router Configuration Strategies

There are two basic approaches to redirecting traffic from the router to the appliance:

- **On the WAN port only**, add a “WCCP redirect in” statement and a “WCCP redirect out” statement.
- **On every port on the router, except the port attached to the appliance**, add a “WCCP redirect in” statement.

The first method redirects only WAN traffic to the appliance, while the second method redirects all router traffic to the appliance, whether it is WAN related or not. On a router with several LAN ports and substantial LAN-to-LAN traffic, sending all traffic to the appliance can overload its LAN segment and burden the appliance with this unnecessary load. If GRE is used, the unnecessary traffic can load down the router as well.

On some routers, the "redirect in" path is faster and puts less of a load on the router's CPU than does the "redirect out" path. If necessary, this can be determined by direct experiment on your router: Try both redirection methods under full network load to see which delivers the highest transfer rates.

Some routers and WCCP-capable switches do not support "WCCP redirect out," so the second method must be used. To avoid overloading the router, the best practice to avoid redirecting large numbers of router ports through the appliance, perhaps by using two routers, one for WAN routing and one for LAN-to-LAN routing.

In general, method 1 is simpler, while method 2 may provide greater performance.

Traffic Shaping and WCCP

A service group can be either TCP or UDP, but not both. For the traffic shaper to be effective, both kinds of WAN traffic must pass through the appliance. Therefore:

- Acceleration requires one service group, for TCP traffic.
- Traffic shaping requires two service groups, one for TCP traffic and one for UDP traffic.
- The difference between the two is configured on the appliance, and the router accepts this configuration.

Configuring the Router

Aug 09, 2017

The appliance negotiates WCCP-GRE or WCCP-L2 automatically. The main choice is between *unicast operation* (in which the appliance is configured with the IP address of each router), or *multicast operation* (in which both the appliance and the routers are configured with the multicast address.)

Normal (Unicast) operation—For normal operation, the procedure is to declare WCCP version 2 and the WCCP group ID for the router as a whole, then enable redirection on each WAN interface. Following is a Cisco IOS example:

```
config term
ip wccp version 2
! We will configure the appliance to use group 51 for TCP and 52 for UDP.
ip wccp 51
ip wccp 52

! Repeat the following three lines for each WAN interface
! you wish to accelerate:
interface your_wan_interface
! If Reverse Path Forwarding is enabled (with an ip verify unicast
! source reachable" statement), delete or comment out the statement:
! ip verify unicast source reachable-via any
! Repeat on all ports.

ip wccp 51 redirect out
ip wccp 51 redirect in
ip wccp 52 redirect out
ip wccp 52 redirect in

! If the appliance is inline with one of the router interfaces
! (NOT SUPPORTED), add the following line for that interface
! to prevent loops:
ip wccp redirect exclude in
^Z
```

If multiple routers are to use the same appliance, each is configured as shown above, using either the same service groups or different ones.

Multicast operation—When giving the appliance and each router a multicast address, the configuration is slightly different than for normal operation. Following is a Cisco IOS example:

```
config term
ip wccp version 2
ip wccp 51 group-address 225.0.0.1

! Repeat the following three lines for each WAN interface
! you wish to accelerate:
interface your_wan_interface
! If Reverse Path Forwarding is enabled (with an ip verify unicast
! source reachable" statement), delete or comment out the statement:
```

```
! ip verify unicast source reachable-via any

ip wccp 51 redirect out
ip wccp 51 redirect in
!
! The following line is needed only on the interface facing the other router,
! if there is another router participating in this service group.
ip wccp 51 group-listen

!If the appliance is inline with one of the router interfaces,
!(which is supported but not recommended), add
!the following line for that interface to prevent loops:
ip wccp redirect exclude in
^Z
```

Basic Configuration Procedure for WCCP Mode on the SD-WAN Appliance

Aug 09, 2017

For most sites, you can use the following procedure to configure the WCCP mode on the appliance. The procedure has you set several parameters to sensible default values. Advanced deployments might require that you set these parameters to other values. For example, if WCCP service group 51 is already used by your router, you need to use a different value for the appliance.

To configure WCCP mode on the appliance

1. On the Configuration: Appliance Settings: WCCP page.
2. If no service groups have been defined, the Select Mode page appears. The options are Single SD-WAN and Cluster (Multiple SD-WANs). Select Single SD-WAN. You are taken to the WCCP page.
   Note: The mode labels are misleading. “Single SD-WAN” mode is also used for SD-WAN high-availability pairs.
3. If WCCP mode is not enabled, click Enable.
4. Click Add Service Group.
5. The default interface (apA), Protocol (TCP), WCCP Priority (0), Router Communication (Unicast), (Password blank) and Time to Live (1) values usually do not have to be changed for the first service group that you create, but if they do, type new values in the fields provided.
6. In the Router Addressing field (if you are using unicast) or the Multicast Address field (if you are using multicast), type the router’s IP address. Use the IP for the router port used for WCCP communication with the appliance.
7. If more than one router is using WCCP to communicate with this appliance, add additional routers now.
8. If your routers have special requirements, set the Router Forwarding (Auto/GRE/Level-2), Router Packet Return (Auto/GRE/Level-2), and Router Assignment (Mask/Hash) fields accordingly. The defaults produce optimal results with most routers.
9. Click Add.
10. Repeat the preceeding steps to create another service group, for UDP traffic (for example, service group Id 52 and Protocol UDP).
11. Go to the Monitoring: Appliance Performance: WCCP page. The Status field should change to Connected within 60 seconds.
12. Send traffic over the link and, on the Connections page, verify that connections are arriving and being accelerated.

WCCP Service Group Configuration Details

Aug 09, 2017

In a service group, a WCCP router and a SD-WAN appliance ("WCCP Cache" in WCCP terminology) negotiate communication attributes (capabilities). The router advertises its capabilities in the "I See You" message. The communication attributes are:

- Forwarding Method: GRE or Level-2
- Packet Return Method (multicast only): GRE or Level-2
- Assignment Method: Hash or Mask
- Password (defaults to none)

The appliance triggers an alert if it detects an incompatibility between its attributes and those of the router. The appliance might be incompatible because of a specific attribute of a service group (such as GRE or Level-2). More rarely, in a multicast service group, an alert can be triggered when the "Auto" selection chooses a particular attribute with a particular router connected, but the attribute is incompatible with a subsequent router.

Following are the basic rules for the communication attributes within a SD-WAN Appliance.

For Router Forwarding:

- When “Auto” is selected, the preference is for Level-2, because it is more efficient for both router and appliance. Level-2 is negotiated if the router supports it and the router is on the same subnet as the appliance.
- Routers in a unicast service group can negotiate different methods if “Auto” is selected.
- Routers in a multicast service group must all use the same method, whether forced with “GRE” or “Level-2”, or, with “Auto,” as determined by the first router in the service group to connect.
- For an incompatibility, an alert announces that the router “has incompatible router forwarding.”

For Router Assignment:

- The default is Hash.
- When “Auto” is selected, the mode is negotiated with the router.
- All routers in a service group must support the same assignment method (Hash or Mask).
- For any service group, if this attribute is configured as “Auto,” the appliance selects “Hash” or “Mask” when the first router is connected. “Hash” is chosen if the router supports it. Otherwise, “Mask” is selected. The problem of subsequent routers being incompatible with the automatically selected method can be minimized by manually selecting a method common to all routers in the service group.
- For an incompatibility, an alert announces that the router “has incompatible router assignment method.”
- With either method, the single appliance in the service group instructs all the routers in the service group to direct all TCP or UDP packets to the appliance. Routers can modify this behavior with access lists or by selecting which interfaces to redirect to the service group.
  For the Mask method, the appliance negotiates the “source IP address” mask. The appliance provides no mechanism to select “destination IP address” or the ports for either source or destination. The “source IP address” mask does not specifically identify any specific IP address or range. The protocol does not provide a means to specify a specific IP address. By default, because there is only a single appliance in the service group, a one-bit mask is used, to conserve router resources. (Release 6.0 used a larger mask.)

For Password:

- If the router requires a password, the password defined on the appliance must match. If the router does not require a password, the password field on the appliance must be blank.

WCCP Testing and Troubleshooting

Aug 09, 2017

When working with WCCP, the appliance provides different ways of monitoring the status of the WCCP interface, and your router should also provide information.

Monitoring: Appliance Performance: WCCP Page—The WCCP page reports the current state of the WCCP link, and reports most problems.

Log Entries—The Monitoring: Appliance Performance: Logging page shows a new entry each time WCCP mode is established or lost.

Figure 1. WCCP Log Entries (format varies somewhat with release)

Router Status—On the router, the "show ip wccp" command shows the status of the WCCP link:

```
Router>enable
Password:
Router#show ip wccp
Global WCCP information:
   Router information:
      Router Identifier:                  172.16.2.4
      Protocol Version:                   2.0

   Service Identifier: 51
      Number of Cache Engines:            0
      Number of routers:                  0
      Total Packets Redirected:           19951
      Redirect access-list:               -none-
      Total Packets Denied Redirect:      0
      Total Packets Unassigned:           0
      Group access-list:                  -none-
      Total Messages Denied to Group:     0
      Total Authentication failures:      0
```

WCCP Clustering

Aug 09, 2017

The WCCP clustering feature enables you to multiply your acceleration capacity by assigning more than one SD-WAN appliance to the same links. You can cluster up to 32 identical appliances, for up to 32 times the capacity. Because it uses the WCCP 2.0 standard, WCCP clustering works on most routers and some smart switches, most likely including those you are already using.

Because it uses a decentralized protocol, WCCP clustering allows SD-WAN appliances to be added or removed at will. If an appliance fails, its traffic is rerouted to the surviving appliances.

Unlike SD-WAN high-availability, an active/passive pair that uses two appliances to provide the performance of a single appliance, the same appliances deployed as a WCCP cluster has twice the performance of a single appliance, delivering both redundancy and improved performance.

In addition to adding more appliances as your site's needs increase, you can use Citrix's "Pay as You Grow" feature to increase your appliances' capabilities through license upgrades.

Citrix Command Center is recommended for managing WCCP clusters. The following figure shows a basic network of a cluster of SD-WAN appliances in WCCP mode, administered by using Citrix Command Center.

Figure 1. SD-WAN Cluster Administered by Using Citrix Command Center

Load-Balanced WCCP Clusters

The WCCP protocol supports up to 32 appliances in a fault-tolerant, load balanced array called a cluster. In the example below, three identical appliances (same model, same software version) are cabled identically and configured identically except for their IP addresses. Appliances using the same service groups with the same router can become a load balanced WCCP cluster. When a new appliance registers itself with the router, it can join the existing pool of appliances and receive its share of traffic. If an appliance leaves the network (as indicated by the absence of heartbeat signals), the cluster is rebalanced so that only the remaining appliances are used.

Figure 2. A load-balanced WCCP cluster with three appliances

One appliance in the cluster is selected as the designated cache, and controls the load-balancing behavior of the appliances in the cluster. The designated cache is the appliance with the lowest IP address. Because the appliances have identical configurations, it doesn't matter which one is the designated cache. If the current designated cache goes offline, a different appliance becomes the designated cache.

The designated cache determines how the load-balanced traffic is allocated and informs the router of these decisions. The router shares information with all members of the cluster, so the cluster can operate even if the designated cache goes offline.

Note: As normally configured, a SD-WAN WANOP 4000/5000 appliance appears as two WCCP caches to the router.

Load-Balancing Algorithm

Load balancing in WCCP is static, except when an appliance enters or leaves the cluster, which causes the cluster to be rebalanced among its current members.

The WCCP standard supports load balancing based on a mask or a hash. For example, SD-WAN WANOP WCCP clustering uses the mask method only, using a mask of 1-6 bits of the 32-bit IP address. These address bits can be non-consecutive. All addresses yielding the same result when masked are sent to the same appliance. Load balancing effectiveness depends on choosing an appropriate mask value: a poor mask choice can result in poor load-balancing or even none, with all traffic sent to a single appliance.

Deployment Topology

Aug 09, 2017

Depending on your network topology, you can deploy WCCP cluster either with a single router or with multiple routers. Whether connected to a single router or multiple routers, each appliance in the cluster must be connected identically to all routers in use.

Single Router Deployment

In the following diagram, three SD-WAN appliances accelerate the datacenter's 200 Mbps WAN. The site supports 750 XenApp users.

As shown on the NetScaler SD-WAN Datasheet, a SD-WAN 3000-100 can support 100 Mbps and 400 users, so a pair of these appliances supports 200 Mbps and 800 users, which satisfies the datacenter's requirements of a 200 Mbps link and 750 users.

For fault tolerance, however, the WCCP cluster should continue to operate without becoming overloaded if one appliance fails. That can be accomplished by using three appliances when the calculations call for two. This is called the N+1 rule.

Failure is an unusual event, so usually all three appliances are in operation. In this case, each appliance is supporting only 67 Mbps and 250 users, leaving plenty of headroom, and making good use of the fact that the cluster has three times the CPU power and three times the compression history of a single appliance.

Without WCCP clustering, the same level of capacity and fault-tolerance would require a pair of SD-WAN WANOP 4000-500 appliances in high availability mode. Only one of these appliances is active at a time.

Multiple Router Deployment

Using multiple WAN routers is very similar to using a single WAN router. If the previous example is changed to include two 100 Mbps links instead of one 200 Mbps link, the topology changes, but the calculations do not.

Limitations

Aug 09, 2017

Configuring appliances in a WCCP cluster has the following limitations:

- All appliances within a cluster must be the same model and use the same software release.
- Parameter synchronization between appliances within the cluster is not automatic. Use Command Center to manage the appliances as a group.
- SD-WAN traffic shaping is not effective, because it relies on controlling the entire link as a unit, and none of the appliances are in a position to do this. Router QoS can be used instead.
- The WCCP-based load-balancing algorithms do not vary dynamically with load, so achieving a good load balance can require some tuning.
- The hash method of cache assignment is not supported. Mask assignment is the supported method.
- While the WCCP standard allows mask lengths of 1-7 bits, the appliance supports masks of 1-6 bits.
- Multicast service groups are not supported; only unicast service groups are supported.
- All routers using the same service group pair must support the same forwarding method (GRE or L2).
- The forwarding and return method negotiated with the router must match: both must be GRE or both must be L2. Some routers do not support L2 in both directions, resulting in an error of "Router's forward or return or assignment capability mismatch." In this case, the service group must be configured as GRE.
- SD-WAN VPX does not support WCCP clustering.
- The appliance supports (and negotiates) only unweighted (equal) cache assignments. Weighted assignments are not supported.
- Some older appliances, such as the SD-WAN 700, do not support WCCP clustering.
- (SD-WAN WANOP 4000/5000 only) Two accelerator instances are required per interface in L2 mode. No more than three interfaces are supported per appliance (and then only on appliances with six or more accelerator instances.)
- (SD-WAN 4000/5000 only) WCCP control packets from the router must match one of the router IP addresses configured on the appliance for the service group. In practice, the router's IP address for the interface that connects it to the appliance should be used. The router's loopback IP should not be used.

Planning Your Deployment

Aug 09, 2017

Deploying appliances in a WCCP cluster requires more planning than does deploying a single appliance. Read the following sections carefully before proceeding.

Selecting Appliances

Aug 09, 2017

The appliances you select for the deployment must all be the same model, running the same software version. Otherwise, management and troubleshooting can become impractical.

Your appliance choice is generally made by comparing your site's WAN bandwidth and number of WAN users to the capacities of the different appliances in the NetScaler SD-WAN Data Sheet. For fault tolerance, always order one more appliance than is absolutely required according to the data sheet.

The number of appliances you need is found as follows, rounding up all fractions:
appliances = max (appliances_bw, appliances_users),
where
appliances_bw = (WAN_bandwidth / Optimized_WAN_capacity) +1
appliances_users = (WAN_users / Maximum_HDX_sessions) +1
Note that if appliances = 2, you can use just a single appliance instead of WCCP clustering, or an HA pair instead of WCCP clustering, since the equation builds in a spare appliance. In other words, WCCP clustering is not necessary (from a capacity perspective) unless appliances is 3 or more.

Example. Suppose you have 700 users and a 100 Mbps link. Some appliances you might consider are the SD-WAN 2000-050, the SD-WAN 3000-100, and the SD-WAN 4000-310.

Model	Optimized WAN Capacity	Maximum HDX Sessions	Appliances_bw	Appliances_users	Appliances
2000-050	50 Mbps	300	3	4	4
3000-100	100 Mbps	400	2	3	3
4000-310	310 Mbps	750	2	2	2

As you can see from the above table, the higher-performance platforms require fewer appliances to get the job done, as you would expect. The SD-WAN 4000-310 meets the requirements with a single appliance, and evaluates to two appliances only because the equations build in a spare.

You can always add more capacity by adding more appliances, but that is not always necessary. The bandwidth limits of two of the three choices, the SD-WAN 3000-100 and theSD-WAN 4000-310, can be increased through a license upgrade. The SD-WAN 4000-050 however, is already at the high end of the range for NetScaler 2000 appliances.

Load-Balancing in the WCCP Cluster

Aug 09, 2017

Traffic is distributed among the appliances in the WCCP cluster. If an appliance leaves the cluster (through failure, overload, or being manually disabled), its traffic is rebalanced by distributing it among the surviving members. If an appliance joins the cluster, traffic is rebalanced once more to give the new appliance its fair share.

The Address Mask

Traffic is distributed on the basis of an address mask that is applied to the source and destination addresses of WAN traffic. You must select an appropriate mask field for efficient load-balancing. An inappropriate mask can result in load-balancing that is poor to nonexistent. For example, if the mask matches an address field that is identical at all your remote sites, all your WAN traffic is sent to a single appliance, overloading it. For example, if all of your remote sites have an address in the form of 10.0.x.x, and your mask bits are within the 10.0 portion of the address all traffic is sent to a single appliance.

The address bits extracted by the address mask are used as an index that is used (indirectly) to select one of the WCCP caches (appliances). For example, an address mask with two "one" bits results in four possible values, depending on the address. Each of these values can be thought of as a bucket. With two mask bits, you have four buckets, numbered 0-3. The buckets are assigned to WCCP caches. To be effective, there must be at least as many buckets as caches. If you use a two-bit mask and have five or more caches, some caches are idle, because each bucket is assigned to only one cache, and there are not enough buckets to cover all five caches:

Cache	1	2	3	4	5
Buckets	0	1	2	3	-

If there are more buckets than caches, some caches are assigned multiple buckets. For example, if you set three mask bits, creating eight buckets, and you have four caches, two buckets are assigned to each cache. If you have five caches, three caches are assigned two buckets each, and two caches are assigned just one. If each bucket represents the same number of users, you have a 2:1 load imbalance across caches:

Cache	1	2	3	4	5
Buckets	0-1	2-3	4-5	6	7

Increasing the number of set mask bits reduces this imbalance. With four mask bits (16 index values) and five caches, four caches receive three buckets and one cache receives four buckets, resulting in only a 4:3 imbalance. With six set mask bits (the largest number supported), four caches receive 13 buckets and one receives 12, which is only a 13:12 load imbalance.

Cache	1	2	3	4	5
Buckets	0-12	13-25	26-38	39-51	52-63

Ideally, you would like each remote site to be directed to a single appliance in the WCCP cluster, so that all traffic to and from a given site is stored in the same compression history. With this arrangement, any traffic from one user at the site can be used to compress similar traffic from any other user at that site. In other words, for compressibility, load-balancing works best if it the address mask selects the bits that differentiate one remote site from another. These are often the least-significant bits of the subnet portion of the IP address. Using these bits tends to allocate the same number of remote sites (not users) per local appliance. A mask that aligns with the host portion of the address instead of the subnet results in a

more equal number of remote users (not sites) per appliance, but at the expense of compression effectiveness. (Compression is only effective when connections flow through the same appliances, and splitting traffic from the same remote site between two or more local appliances interferes with this.)

Finally, for good load-balancing, each "one" bit in the address mask must be set to one on 50% of the remote addresses, and set to zero on 50% of the remote addresses. This is not the case on all address bits, since in most WANs, the highest-order network bits never change at all (such as the 10 in 10.x.x.x). Such bits must never be selected by the address mask.

In addition, many subnets are only sparsely populated. For example, if only 50 addresses are used in the subnet 10.1.2.0/24, and they are assigned sequentially, the two higher-order host bits (representing the unused range of 10.1.2.64-10.1.2.255) for this subnet never change, and if these two bits are included in the address mask, three-fourths of the buckets receive no traffic.

Useful compromises between these two extremes can generally be found.

Follow these rules:

- The number of ”one” bits in the address mask must allow at least as many combinations as there are WCCP caches in the cluster. That is, if you have eight appliances, the address mask must contain at least three ”one” bits.
- The ”one” bits in the address mask must each be inside the active address range for most of your remote subnets, or they skew the load-balancing distribution.
- The mask should split the address range of individual remote sites into as few pieces as possible, for best compression performance.
- If a remote appliance is faster than the local members of the WCCP cluster, the mask should be designed to divide its traffic between multiple local appliances. For example, a 100 Mbps remote appliance should have its traffic split between two 50 Mbps local appliances by setting a bit inside the remote appliance’s active address range.
- The “one” bits in the mask are typically contiguous, but this is not required. They can be in any pattern.

Example: Suppose you set an address mask of 0x0000 0f00, which has four “one” bits. This defines a four-bit field that is extracted from the IP address, yielding 16 possible results (16 buckets). These buckets are in turn assigned to the actual WCCP caches in the WCCP cluster.

Address	Masked Address (mask = 0x0000 0f00)	Bucket
10.0.0.5	0.0.0.0	0
10.0.1.128	0.0.1.0	1
155.0.2.55	0.0.2.0	2
253.100.255.2	0.0.15.0	15
10.0.15.1	0.0.15.0	15

Zero bits in the mask are ignored, and the “one” bits are used to define the extracted field. So if the mask is 0x10 10 10 10, these widely separated “one” bits are extracted into a four-bit field, declaring 16 buckets and a bucket numbers in the range of 0-15.

If the mask value is set to zero, a default value of 0x00 00 0f 00 is used.

Assigning Buckets to Appliances

Aug 09, 2017

The mapping of bucket to appliances is subject to several variables:

- Which appliances are available: If an appliance is down, its share of buckets are given to the available appliance. If a new appliance is added to the cluster, it is given its fair share of buckets.
- The mapping algorithm used (deterministic or least-disruptive).
- The order in which appliances come online (least-disruptive mapping only).
- The IP addresses of the appliances. WCCP algorithms can use a sorted list of appliance IP addresses; for example, assigning buckets to appliances in the same order as the appliance IP addresses.

The most important of these factors, from an administrator's point of view, is the mapping algorithm.

Deterministic mapping. The deterministic mapping algorithm is less graceful than the least-disruptive algorithm, but it supports Hot Standby Router Protocol (HSRP) and Global Server Load Balancing (GSLB) routing, and is required when multiple routers using such protocols share the WCCP cluster.

Deterministic mapping is also the preferred method when the cluster has only two appliances.

Assignments are based on the IP addresses of the active appliances. Each appliance gets its fair share of bucket, with the lowest-numbered bucket being assigned to the appliance with the lowest IP address. If there are more appliances than buckets, the leftover appliances (with no bucket assigned to them) are the ones with the highest-numbered IP addresses. This deterministic assignment allows traffic to arrive for a single connection through any of the routers in the service group and be forwarded to the same appliance.

Reassignment can be disruptive to accelerated connections, which are reset if they migrate to a different appliance. With deterministic mapping, the number of buckets that are reassigned to new appliances can be quite high if there are three or more appliances.

Least-disruptive mapping. When a bucket is assigned to a different appliance, any open accelerated connections that used the old appliance is reset. The least-disruptive algorithm keeps the reassignment to a minimum. For example, if you have three appliances, and one appliance fails, the new mapping preserves roughly two-thirds of the assignments and remaps the remaining third (which fails anyway, because their appliance failed). The least-disruptive algorithm does not support HSRP or GSLB routing, because it is not guaranteed to result in identical mappings on all the routers in the service group, and therefore, packets from a single connection might be sent to two different appliances by two different routers, which causes accelerated connections to fail.

Startup and Failover Behavior

Aug 09, 2017

Each appliance registers itself with the routers specified in its service class definitions. The first appliance to register itself, becomes the *designated cache*, and works with the routers to apportion traffic between itself and the other caches (called *subordinate caches*). Because your appliances use identical WCCP algorithms, it does not matter which one becomes the designated cache.

As additional appliances come online, they are added to the WCCP cluster, and the traffic is reapportioned among the active appliances. This happens at ten-second increments. After a cold start of the routers or appliances, all of the appliances might come online within the same ten-second window, or they might arrive over multiple ten-second windows, causing traffic to be reapportioned multiple times before it stabilizes. In the latter case, the appliances that come online first maycan become overloaded until additional appliances come online.

An accelerated connection fails when allocated to a different appliance, making reallocation disruptive. This is not true of non-accelerated connections, which generally experience a delay of thirty seconds or more, and then continue. The least-disruptive mapping option minimizes the amount of reallocation when an appliance fails.

If an appliance fails or otherwise goes offline, its absence is noted, and the designated cache reapportions its traffic to the remaining appliances. If the designated cache itself goes offline, the role of designated cache is also reapportioned. It takes about thirty seconds for the cluster to react to the loss of a cache.

Deployment Worksheet

Aug 09, 2017

On the following worksheet, you can calculate the number of appliances needed for your installation and the recommended mask field size. The recommended mask size is 1-2 bits larger than the minimum mask size for your installation.

Parameter	Value	Notes
Appliance Model Used		—
Supported XenApp and XenDesktop Users Per Appliance	Uspec =	From data sheet
XenApp and XenDesktop Users on WAN Link	Uwan =	—
User overload Factor	Uoverload = Uwan/Uspec =	—
Supported BW Per Appliance	BWspec =	From data sheet
WAN Link BW	BWwan =	—
BW Overload Factor	BWoverload = BWwan/BWspec =	—
Number of appliances required	N = max(Uoverload, BWoverload) +1 =	Includes one spare
		—
Min number of buckets	Bmin = N, rounded up a power of 2 =	—
If SD-WAN 4000 or 5000,	Bmin = 2*N, rounded up to a power of 2 =	—
Recommended value	B = 4 * Bmin if Bmin <= 16, else 2 * Bmin =	—
Number of "one" bits in address mask	M = log2(B)	If B=16, M=4.

Mask value: The mask value is a 32-bit address mask with a number of “one” bits equal to M in the above worksheet. Often

these bits can be the least-significant bits in the WAN subnet mask used by your remote sites. If the masks at your remote sites vary, use the median mask. (Example: With /24 subnets, the least significant bits of the subnet are 0x00 00 nn 00. The number of bits to set to one is log2(mask size): if mask size is 16, set four bits to one. So with a mask size of 16 and a /24 subnet, set the mask value to 0x00 00 0f 00.): _____

The above guidelines work only if the selected subnet field is evenly distributed in your traffic, that is, that each address bit selected by the mask is a one for half the remote hosts, and a zero for the other half. Otherwise, load-balancing is impaired. This even distribution might be true for only a small number of bits in the network field (perhaps only two or three bits). If this is the case with your network, instead of masking bits in the offending area of the subnet field, displace those bits to a portion of the host address field that has the 50/50 property. For example, if only three subnet bits in a /24 subnet have the 50/50 property, and you are using four mask bits, a mask of 0x00 00 07 10 avoids the offending bit at 0x00 00 0800 and displaces it to 0x00 00 00 10, a portion of the address field that is likely to have the 50/50 property if your remote subnets generally use at least 32 IP addresses each.

Parameter	Value	Notes
Final Mask Value		—
Accelerated Bridge		Usually apA
WAN Service Group		A service group not already in use on your router (51-255)
LAN Service Group		Another unused service group
Router IP address		IP address of router interface on port facing the appliance
WCCP Protocol (usually "Auto")		—
DC Algorithm		Use "Deterministic" if you have only two appliances or are using dynamic load balancing like HSRP or GSLB. Otherwise, use "Least Disruptive."

Configuring WCCP Clustering

Aug 09, 2017

After you have finalized the deployment topology, considered all limitations, and filled in the deployment worksheet, you are ready to deploy your appliances in a WCCP cluster. To configure the WCCP cluster, you need to perform the following tasks:

- Configuring the NetScaler Instances
- Configuring the Router
- Configuring the Appliance

Configuring the Router

Aug 09, 2017

WCCP configuration on the router is simple, because most WCCP parameters are set by the appliances.

Unlike legacy SD-WAN WCCP support, WCCP clustering uses two service groups for TCP traffic. One service group is used on the router's WAN interface, and the other is used on the router's LAN interfaces (except for the LAN interface used by the SD-WAN appliances themselves, when deployed in L2-mode WCCP cluster).

As shown in the following figure, you need to configure two service groups because WCCP allows the mask to be applied to either the source IP or the destination IP address, which is not quite what is required. To keep connections between two endpoints together, regardless of which endpoint initiates the connection, the appliance applies the address mask to the source IP address of incoming WAN traffic, and to the destination IP address of incoming LAN traffic. This requires two service groups.

The WAN service group uses WCCP source-ip address masking, while the LAN service group uses dest-ip masking. In some deployments, it may be necessary to reverse the assignments, using the "WAN" service group for your LAN interface and vice versa. This might occur if the number of local IP addresses greatly exceeds the number of remote IP addresses.

Figure 1. SD-WAN WCCP Cluster

To configure WCCP clustering on the router

This procedure assumes Cisco routers, but is similar on other routers. It uses the first of the two methods, discussed above, of redirecting WCCP traffic with an ip wccp redirect in statement on both LAN and WAN ports.

1. Fill in the WCCP clustering Deployment Worksheet.
2. Log on to your router
3. In the global declarations section, declare each service group on the WCCP clustering worksheet, listed as **WAN service group** and **LAN Service group**. For example, ip wccp 61 and ip wccp 62.
   Note: The ip wccp command allows, but does not require, a more elaborate syntax than this, and can specify an ACL name or a password. Both service groups must have the same password, if one is used. The ACLs can be different.
4. Inside the interface declarations for each WAN interface that connects to remote SD-WAN appliances, add an ip wccp x redirect in statement, where x is the WAN service group from the WCCP clustering worksheet.
5. Inside the interface declarations for each LAN interface (except the one connecting to the WCCP cluster, if you are using L2 mode), add an ip wccp y redirect in statement, where y is the LAN service group from the WCCP clustering worksheet.
6. Save your configuration.

Example. The following example uses service group 61 for the WAN service group and service group 62 for the LAN service group. Three router interfaces are used. One is connected to the WAN, one is connected to the LAN, and one is connected to the WCCP cluster.

```
!
! Example is for WCCP clustering using WCCP redirect in statements
! on LAN and WAN interfaces.
! This definition is appropriate for modern Cisco routers.
! Global declarations
ip wccp 61
ip wccp 62
```

```
!
interface GigabitEthernet1/1
description LAN interface. SG 62 is used for LAN
ip address 172.80.1.56 255.255.255.0
ip wccp 62 redirect in
!
interface GigabitEthernet1/2
description LAN interface attaching SD-WAN L2-WCCP appliances
description (No wccp redirect statements are used on this interface)
ip address 172.80.21.56 255.255.255.0
!
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
!
```

Note: If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.

- If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.
- If multiple routers shared the same WCCP cluster, they use the same service groups.

It is also possible to use ip wccp redirect statements on only the WAN interfaces:

```
! Example for WCCP clustering using WCCP redirect in/out statements on
! WAN interface only
! This definition is appropriate for modern Cisco routers.
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN. SG 62 is used for LAN.
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
ip wccp 62 redirect out
!
```

In many routers, the ip wccp redirect out path is not optimized in hardware, but uses the CPU. If the router's capabilities along this path exceeds the WAN speed, this method is practical, and is simpler than using redirect statements on every interface.

Router ACLs can be used to limit redirection. For example, for initial testing, perhaps only a single remote IP address might be allowed to be redirected through WCCP.

Configuring the Appliance

Aug 09, 2017

Repeat the following procedure for each appliance in the cluster:

1. Fill in the WCCP clustering Deployment Worksheet.
2. Navigate to Configuration > Appliance Settings > WCCP page.
3. Click Enable to enable WCCP mode on the appliance.
4. Select **Cluster (Multiple Caches)** option.

5. Fill in parameters in the **Select SD-WAN Cluster** section.

6. Enter T5 from your worksheet as the Cache 1 IP, T6 as the Cache 2 IP, T2 as the Subnet Mask, and T1 as the Gateway. Click **Save**. The **Configure Service Group** section appears.
7. In the Service Group Details section, specify the WAN and LAN service groups (T11 and T12 from your worksheet).
8. In the Priority field, select **100** (in practice this value is somewhat arbitrary).
9. From the Protocol list, select **TCP**.
10. In the DC Algorithm field, select **Deterministic** or **Least Disruptive**. “Deterministic” is always safe to use, and should be used if you are using only two appliances, or are using multiple routers. “Least Disruptive” disrupts fewer user sessions on failover when used with clusters of three or more appliances, but has restrictions on its use.
11. Set **Service Group Pair Status** to On.
12. If your router is configured to require a password, enter the password in the **Service Group Password** field. Otherwise, leave the field blank.
13. In the **Router Communications Details** section, enter the IP address of the router (T8 on your worksheet: often identical to T1 as well). This is the IP address of the appliance-facing router interface. If you use multiple routers to communicate with the appliance, list them all here.
14. From the Router Forwarding list, select Level 2 or GRE, according to the capabilities of your router. Use Level 2 if you can, and GRE if you must.
15. For the Mask Value, enter the value you determined from the WCCP Clustering worksheet. This is a critical value: a poor choice will result in poor load-balancing or none at all.
16. Click Create. This creates the WAN and LAN service groups.
17. On the Configuration > Optimization Rules > Link Definitions page, change the bandwidth limits on each defined WAN to 95% of the aggregate speed of all your WANs. This prevents the link from being under-utilized when load-balancing is imperfect. If ICA (XenApp/XenDesktop) is the dominant use, set each appliance to (95% of WAN bandwidth)/N, where N is the number of appliances (or twice the number of appliances if they are SD-WAN 4000 or 5000 units), to divide the bandwidth equally among the appliances. This latter method is most appropriate for applications with large numbers of active connections that have relatively low bandwidth requirements.

Testing and Troubleshooting

Aug 09, 2017

The **Monitoring** > **Appliance** > **Application Performance** > **WCCP** page shows the current state of not only the local appliance but of all other appliances that have joined the cluster. Select a WCCP cache and click **Get Info**.

The Cache Status tab shows the local appliance's status. When all is well, the status is "25: has assignment." You must refresh the page manually to monitor changes in status. If the appliance does not reach the status of "25: has assignment" within a timeout period, other informative status messages are displayed.

Additional information is displayed when you click on the Service Group or the Routers tabs.

The Cluster Summary tab displays information about the WCCP cluster as a whole. As a side effect of the WCCP protocol, each member of the cluster has information about all the others, so this information can be monitored from any appliance in the cluster.

Your router can also provide status information. See your router documentation.

Virtual Inline Mode

Aug 09, 2017

Note: Use virtual inline mode only when both inline mode and WCCP mode are impractical. Do not mix inline and virtual inline modes within the same appliance. However, you can mix virtual inline and WCCP modes within the same appliance. Citrix does not recommend virtual inline mode with routers that do not support health monitoring.

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router. Almost all of the configuration tasks are performed on the router. The only thing to be configured on the appliance is the forwarding method, and the default method is recommended.

Like WCCP, Virtual inline deployment requires no rewiring and no downtime, and it provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links. Unlike WCCP, it contains no built-in status monitoring or health checking, making troubleshooting difficult. WCCP is thus the recommended mode, and virtual inline is recommended only when inline and WCCP modes are both impractical.

Example

The following figure shows a simple network in which all traffic destined for or received from the remote site is redirected to the appliance. In this example, both the local site and remote site use virtual inline mode.

Figure 1. Virtual Inline Example

Following are some configuration details for the network in this example:

- Endpoint systems have their gateways set to the local router (which is not unique to virtual inline mode).
- Each router is configured to redirect both incoming and outgoing WAN traffic to the local appliance.
- Each appliance processes the traffic received from its local router and forwards it back to the router.
- PBR rules configured on the router prevent routing loops by allowing packets to make only one trip to and from the appliance. The packets that the appliance forwards back to the router are sent to their original (local or remote) destination.
- Each appliance has its default gateway set to the address of the local router, as usual (on the **Configuration: Network Adapters** page). The options for forwarding packets back to the router are Return to Ethernet Sender and Send to Gateway.

Configuring Packet Forwarding on the Appliance

Aug 09, 2017

Virtual inline mode offers two packet-forwarding options:

Return to Ethernet Sender (default)—This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

Send to Gateway (not recommended)—In this mode, virtual inline output packets are forwarded to the default gateway for delivery, even if they are destined for hosts on the local subnet. This option is usually less desirable than the Return to Ethernet Sender option, because it adds an easily forgotten element of complexity to the routing structure.

To specify the packet-forwarding option—On the Configuration: Optimization Rules: Tuning page, next to Virtual Inline, select Return to Ethernet Sender or Send to Gateway.

Router Configuration

Aug 09, 2017

The router has three tasks when supporting virtual inline mode:

1. It must forward both incoming and outgoing WAN traffic to the SD-WAN appliance.
2. It must forward traffic to its destination (WAN or LAN).
3. It must monitor the health of the SD-WAN appliance so that the appliance can be bypassed if it fails.

In virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the router's Ethernet ports to the appliance and creating routing rules that are based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

The basic routing algorithm is:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward it to the appliance.
- If packet is destined for the WAN, forward to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.
- Traffic shaping is not effective unless all WAN traffic passes through the appliance.

Note: When considering routing options, keep in mind that returning data, not just outgoing data, must flow through the appliance. For example, placing the appliance on the local subnet and designating it as the default router for local systems does not work in a virtual inline deployment. Outgoing data would flow through the appliance, but incoming data would bypass it. To force data through the appliance without router reconfiguration, use inline mode.

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does no health monitoring. To enable health monitoring, define a rule to monitor the appliance's availability, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Important: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy-based routing do not support health-checking. The health-monitoring feature is relatively new. It became available in Cisco IOS release 12.3(4)T.

Following is an example of a rule for monitoring the availability of the appliance:

```
!- Use a ping (ICMP echo) to see if appliance is connected track  123 rtr 1 reachabilit y ! rtr 1  type echo protocol Iplcmpecho 192.168.1.200 schedule 1 life forever start-time now
```

This rule pings the appliance at 192.168.1.200 periodically. You can test against 123 to see if the unit is up.

Routing Examples

Aug 09, 2017

The following examples illustrate configuring Cisco routers for the local and remote sites shown in Virtual inline example. To illustrate health monitoring, the configuration for the local site includes health monitoring, but the configuration for the remote site does not.

Note: The configuration for the local site assumes that a ping monitor has already been configured.
The examples conform to the Cisco IOS CLI. They might not be applicable to routers from other vendors.

Local Site, Health-Checking Enabled

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.10.10.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 172.68.1.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended client_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 192.168.1.200 20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 192.168.1.200 10 track 123
```

Remote Site (No Health Checking)

```
! This example does not use health-checking.
! Remember, health-checking is always recommended,
! so this is a configuration of last resort.
!
!
ip cef
!
interface FastEthernet0/0
ip address 20.20.20.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 171.68.2.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.2.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.2.1
!
ip access-list extended client_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
set ip next-hop 192.168.2.200
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop 192.168.2.200
!_
```

Each of the above examples applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP of 10.10.10.0/24 and destination of 20.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended client_side
 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
 ip access-list extended wan_side
 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
```

Note that, for access lists, ordinary masks are not used. Wildcard masks are used instead. Note that when reading a

wildcard mask in binary, "1" is considered a "don't care" bit.

Virtual Inline for Multiple-WAN Environments

Aug 09, 2017

Enterprises with multiple WAN links often have asymmetric routing policies, which seem to require that an inline appliance be in two places at once. Virtual inline mode solves the asymmetric routing problem by using the router configuration to send all WAN traffic through the appliance, regardless of the WAN link used. The below figure shows a simple multiple-WAN link deployment example.

The two local-side routers redirect traffic to the local appliance. The FE 0/0 ports for both routers are in the same broadcast domain as the appliance. The local appliance must use the default virtual inline configuration (Return to Ethernet Sender).

Figure 1. Virtual Inline Mode With Two WAN Routers

Virtual Inline Mode and High-Availability

Aug 09, 2017

Virtual Inline mode can be used in a high availability (HA) configuration. The below figure shows a simple HA deployment. In virtual inline mode, a pair of appliances acts as one virtual appliance. Router configuration is the same for an HA pair as with a single appliance, except that the Virtual IP address of the HA pair, not the IP address of an individual appliance, is used in the router configuration tables. In this example, the local appliances must use default virtual inline configuration (Return to Ethernet Sender).

Figure 1. High-availability Example

Monitoring and Troubleshooting

Aug 09, 2017

In virtual inline mode, unlike WCCP mode, the appliance provides no virtual inline-specific monitoring. To troubleshoot a virtual inline deployment, log into the appliance and use the Dashboard page to verify that traffic is flowing into and out of the appliance. Traffic forwarding failures are typically caused by errors in router configuration.

If the Monitoring: Usage or Monitoring: Connections pages show that traffic is being forwarded but no acceleration is taking place (assuming that an appliance is already installed on the other end of the WAN link), check to make sure that both incoming WAN traffic and outgoing WAN traffic are being forwarded to the appliance. If only one direction is forwarded, acceleration cannot take place.

To test health-checking, power down the appliance. The router should stop forwarding traffic after the health-checking algorithm times out.

Group Mode

Aug 09, 2017

In group mode, two or more appliances become a single virtual appliance. This mode is one solution to the problem of asymmetric routing, which is defined as any case in which some packets in a given connection pass through a given appliance but others do not. A limitation of the appliance architecture is that acceleration cannot take place unless all packets in a given connection pass through the same two appliances. Group mode overcomes this limitation.

Group mode can be used with multiple or redundant links without reconfiguring your routers.

Note: Group mode is not supported on the SD-WAN 4000 or 5000 appliances.
Group mode applies only to the appliances on one side of the WAN link; the local appliances neither know nor care whether the remote appliances are using group mode.

Group mode uses a heartbeat mechanism to verify that other members of the group are active. Packets are forwarded to active group members only.

Avoiding asymmetric routing is the main reason to use group mode, but group mode is not the only method available for that purpose. If you decide that it is the best method for your environment, you can enable it by setting a few parameters. If the default mechanism for determining which appliance is responsible for a particular connection does not provide optimal acceleration, you can change the forwarding rules.

Figure 1. Group Mode With Redundant Links

Figure 2. Group Mode With Non-Redundant Links with Possible Asymmetric Routing

Figure 3. Group Mode On Nearby Campuses

When to Use Group Mode

Aug 09, 2017

Use group mode in the following set of circumstances:

- You have multiple WAN links.
- There is a chance of asymmetric routing (a packet on a given connection might travel over either link).
- Group mode seems simpler and more practical than alternatives that use a single appliance.

The alternatives are:

- WCCP mode, in which traffic from two or more links is sent to the same appliance by WAN routers, by means of the WCCP protocol.
- Virtual inline mode, in which your routers send traffic from two or more links through the same appliance (or high-availability pair).
- Multiple bridges, where each link passes through a different accelerated bridge in the same appliance.
- LAN-level aggregation, which places an appliance (or high-availability pair) closer to the LAN, before the point where WAN traffic is split into two or more paths.

How Group Mode Works

Aug 09, 2017

In group mode, the appliances that are part of the group each take ownership for a portion of the group's connections. If a given appliance is the owner of a connection, it makes all the acceleration decisions about that connection and is responsible for compression, flow control, packet retransmission, and so on.

If an appliance receives a packet for a connection for which it is not the owner, it forwards the packet to the appliance that is the owner. The owner examines the packet, makes the appropriate acceleration decisions, and forwards any output packets back to the non-owning appliance. This process preserves the link selection made by the router, while allowing all packets in the connection to be managed by the owning appliance. For the routers, the introduction of the appliances has no consequences. The routers do not need to be reconfigured in any way, and the appliances do not need to understand the routing mechanism. They simply accept the routers' forwarding decisions.

Figure 1. Sending-side Traffic in Group Mode

Figure 2. Receiving-side traffic flow in group mode

Group mode has two, user-selectable failure modes, which control how the group members interact with each other if one of them fails. The failure mode also determines whether the failed appliance's bypass card opens (blocking traffic through the appliance) or remains closed (allowing traffic to pass through). The failure modes are:

Continue to accelerate- If a group member fails, its bypass card is opened and no traffic passes through the failed appliance. The result is presumably a fail-over if redundant links are used. Otherwise, the link is simply inaccessible. The other appliances in the group continue to accelerate. The usual hashing algorithm handles the changed conditions. (That is, the old hashing algorithm is used, and if the failed unit is indicated as the owner, a hashing algorithm based on the new, smaller group is applied. This preserves as many older connections as possible.)

Do not accelerate- If a group member fails, its bypass card closes, allowing traffic to pass through without acceleration. Because an unaccelerated path introduces asymmetric routing, the other members of the group also go into pass-through mode when they detect the failure.

Enabling Group Mode

Aug 09, 2017

To enable group mode, create a group of two or more appliances. An appliance can be a member of only one group. Group members are identified by IP address and the SSL common name in the appliance license.

All group mode parameters are on the Settings: Group Mode page, in the Configure Settings: Group Mode table.

Figure 1. Group Mode Page

To enable group mode

1. Select the address to use for group communication. At the top of the Group Mode Configuration table on the Configuration: Advanced Deployments: Group Mode tab, the table cell under Member VIP contains the management address of the port used to communicate with other group members. Use the (unlabeled) drop-down menu to select the correct address (for example, to use the Aux1 port, select the IP address you assigned to the Aux1 port). Then, click Change VIP.
2. Add at least one more group member to the list. (Groups of three or more are supported but are rarely used.) In the next cell of the Member VIP column, type the IP address of the port used by the other appliance for group-mode communication.
3. Type the other group member's SSL common name in the SSL Common Name column. The SSL common name is listed on the other appliance's Configure: Advanced Deployments: High Availability tab. If the other group member is a high-availability pair, the name listed is the SSL common name of the primary appliance.
   Note: If the local appliance is not part of a high-availability pair, the first cell in the HA Secondary SSL Common Name is blank.
   If the other group member is a high-availability pair, specify the SSL Common Name of the HA secondary appliance in the HA Secondary SSL Common Name column.
4. Click Add.
5. Repeat steps 2-4 for any additional appliances or high-availability pairs in the group.
6. The three buttons under the list of group members are toggles, so each is labeled as the opposite of its current setting:
   1. The top button reads either, **Do not accelerate when member failure is detected** or **Continue to accelerate when member failure is detected**. The "Do not accelerate..." setting always works and does not block traffic, but if any member fails, the other group members go into bypass mode, which causes a complete loss of acceleration. With the "Continue to accelerate" option, the failing appliance's bridge becomes an open circuit, and the link fails. This option is appropriate if the WAN router responds by causing a failover. New connections, and open connections belonging to the surviving appliances, are accelerated.
   2. The bottom button should now be labeled Disable Group Mode. If it is not, enable group mode by clicking the button.
7. Refresh the screen. The top of the page should list the group mode partners, but display warnings about their status, because they haven't been configured for group mode yet. For example, it might indicate that the partner cannot be found or is running a different software release.
8. Repeat this procedure with the other members of the group. Within 20 seconds after enabling the last member of the group, the Group Mode Status line should show NORMAL, and the other group mode members should be listed with Status: On-Line and Configuration: OK.

Forwarding Rules

Aug 09, 2017

By default, the *owner* of a group-mode connection is set by a hash of the source and destination IP addresses. Each appliance in the group uses the same algorithm to determine which group member owns a given connection. This method requires no configuration. The owner can optionally be specified through user-settable rules.

Because the group-mode hash is not identical to that used by load balancers, about half of the traffic tends to be forwarded to the owning appliance in a two-Appliance group. In the worst case, forwarding causes the load on the LAN-side interface to be doubled, which halves the appliance's peak forwarding rate for actual WAN traffic.

This speed penalty can be reduced if the Primary or Aux1 Ethernet ports are used for traffic between group members. For example, if you have a group of two appliances, you can use an Ethernet cable to connect the two units' Primary ports, then specify the Primary port on the Group Mode page on each unit. However, maximum performance is achieved if the amount of traffic forwarded between the group-mode members is minimized.

The owner can optionally be set according to specific IP/port-based rules. These rules must be identical on all appliances in the group. Each member of the group verifies that its group-mode configuration is identical to the others. If not all of the configurations are identical, none of the member appliances enter group mode.

If traffic arrives first at the appliance that owns the connection, it is accelerated and forwarded normally. If it arrives first at a different appliance in the group, it is forwarded to its owner over a GRE tunnel, which accelerates it and returns it to the original appliance for forwarding. Thus, group mode leaves the router's link selection unchanged.

Using explicit IP-based forwarding rules can reduce the amount of group-mode forwarding. This is especially useful in primary-link/backup-link scenarios, where each link handles a particular range of IP addresses, but can act as a backup when the other link is down.

Figure 1. IP-Based Owner Selection

Forwarding rules can ensure that group members handle only their "natural" traffic. In many installations, where traffic is usually routed over its normal link and only rarely crosses the other one, these rules can reduce overhead substantially.

Rules are evaluated in order, from top to bottom, and the first matching rule is used. Rules are matched against an optional IP address/mask pair (which is compared against both source and destination addresses), and against an optional port range.

Regardless of the ordering of rules, if the partner appliance is not available, traffic is not forwarded to it, whether a rule matches or not.

For example, in the figure below, member 172.16.1.102 is the owner of all traffic to or from its own subnet (172.16.1.0/24), while member 172.16.0.184 is the owner of all other traffic.

If a packet arrives at unit 172.16.1.102, and it is not addressed to/from net 172.16.1.0/24, it is forwarded to 172.16.0.184.

If unit 172.16.0.184 fails, however, unit 172.16.1.102 no longer forwards packets. It attempts to handle the traffic itself. This behavior can be inhibited by clicking **Do NOT Accelerate When Member Failure Detected** on the Group Mode tab.

In a setup with a primary WAN link and a backup WAN link, write the forwarding rules to send all traffic to the appliance on the primary link. If the primary WAN link fails, but the primary appliance does not, the WAN router fails over and sends traffic over the secondary link. The appliance on the secondary link forwards traffic to the primary-link appliance, and acceleration continues undisturbed. This configuration maintains accelerated connections after the link failover.

Figure 2. Forwarding Rules

Monitoring and Troubleshooting Group Mode

Aug 09, 2017

Two things should be checked in a group-mode installation:

- That the two appliances have entered group mode, which can be determined on either appliance's Configuration: Advanced Deployments: Group Mode page.
- That the behavior of the group-mode pair is as desired when the other member fails, and when one of the links fail, as determined by disabling the other appliance and temporarily disconnecting one of the links, respectively.

High-Availability Mode

Aug 09, 2017

Two identical appliances on the same subnet can be combined as a *high-availability pair.* The appliances each monitor the other's status by using the standard *Virtual Router Redundancy Protocol (VRRP)* heartbeat mechanism. The pair has a common virtual IP address for management, in addition to each appliance's management IP address. If the primary appliance fails, the secondary appliance takes over. Failover takes approximately five seconds.

High availability mode is a standard feature.

How High-Availability Mode Works

Aug 09, 2017

In a high availability (HA) pair, one appliance is primary, and the other is secondary. The primary monitors its own and the secondary's status. If it detects a problem, traffic processing fails over to the secondary appliance. Existing TCP connections are terminated. To ensure successful failover, the two appliances keep their configurations synchronized. In a WCCP mode high availability configuration, the appliance that is processing traffic maintains communication with the upstream router.

Status monitoring—When high availability is enabled, the primary appliance uses the VRRP protocol to send a heartbeat signal to the secondary appliance once per second. In addition, the primary appliance monitors the carrier status of its Ethernet ports. The loss of carrier on a previously active port implies a loss of connectivity.

Failover If the heartbeat signal of the primary appliance should fail, or if the primary appliance loses carrier for five seconds on any previously active Ethernet port, the secondary appliance takes over, becoming the primary. When the failed appliance restarts, it becomes the secondary. The new primary announces itself on the network with an ARP broadcast. MAC spoofing is not used. Ethernet bridging is disabled on the secondary appliance, leaving the primary appliance as the only path for inline traffic. Fail-to-wire is inhibited on both appliances to prevent loops.

Caution: The Ethernet bypass function is disabled in HA mode. If both appliances in an inline HA pair lose power, connectivity is lost. If WAN connectivity is needed during power outages, at least one appliance must be attached to a backup power source.

Note: The secondary appliance in the HA pair has one of its bridge ports, port apA.1, disabled to prevent forwarding loops. If the appliance has dual bridges, apB.1 is also disabled. In a one-arm installation, use port apA.2. Otherwise, the secondary appliance becomes inaccessible when HA is enabled.

Primary/secondary assignment—If both appliances are restarted, the first one to fully initialize itself becomes the primary. That is, the appliances have no assigned roles, and the first one to become available takes over as the primary. The appliance with the highest IP address on the interface used for the VRRP heartbeat is used as a tie-breaker if both become available at the same time.

Connection termination during failover—Both accelerated and unaccelerated TCP connections are terminated as a side effect of failover. Non-TCP sessions are not affected, except for the delay caused by the brief period (several seconds) between the failure of the primary appliance and the failover to the secondary appliance. Users experience the closing of open connections, but they can open new connections.

Configuration synchronization—The two appliances synchronize their settings to ensure that the secondary is ready to take over for the primary. If the configuration of the pair is changed through the browser based interface, the primary appliance updates the secondary appliance immediately.

HA cannot be enabled unless both appliances are running the same software release.

HA in WCCP mode—When WCCP is used with an HA pair, the primary appliance establishes communication with the router. The appliance uses its management IP address on apA or apB, not its virtual IP address, to communicate with the router. Upon failover, the new primary appliance establishes WCCP communication with the router.

Cabling Requirements

Aug 09, 2017

The two appliances in the high availability pair are installed onto the same subnet in either a parallel arrangement or a one-arm arrangement, both of which are shown in the following figure. In a one-arm arrangement, use the apA.2 port (and, optionally, the apB.2 port), not the apA.1 port. Some models require a separate management LAN, whether deployed in inline or one-armed mode. This is depicted only in the middle diagram.

Figure 1. Cabling for High-Availability Pairs

Do not break the above topology with additional switches. Random switch arrangements are not supported. Each of the switches must be either a single, monolithic switch, a single logical switch, or part of the same chassis.

If the spanning-tree protocol (STP) is enabled on the router or switch ports attached to the appliances, failover will work, but the failover time may increase to roughly thirty seconds. Without STP, failover time is roughly five seconds. Thus, to achieve the briefest possible failover interval, disable STP on the ports connecting to the appliances.

Figure 2. Ethernet Port Locations (Older Models)

Other Requirements

Aug 09, 2017

Both appliances in an HA pair must meet the following criteria:

- Have identical hardware, as shown by on the System Hardware entry on the Dashboard page.
- Run exactly the same software release.
- Be equipped with Ethernet bypass cards. To determine what is installed in your appliances, see the Dashboard page.

Appliances that do not support HA display a warning on the Configuration: High Availability page.

Management Access to the High-Availability Pair

Aug 09, 2017

When configuring a high-availability (HA) pair, you assign the pair a virtual IP (VIP) address, which enables you to manage the two appliances as if they were a single unit. After you enable high-availability mode, managing the secondary appliance through its IP address is mostly disabled, with most parameters grayed out. A warning message displays the reason on every page. Use the HA VIP for all management tasks. You can, however, disable the secondary appliance's HA state from its management UI.

Configuring the High-Availability Pair

Aug 09, 2017

You can configure two newly installed appliances as a high-availability pair, or you can create an HA pair by adding a second appliance to an existing installation.

Prerequisites: Physical installation and basic configuration procedures

To configure high availability

1. Make sure that no more than one appliance is connected to the traffic networks (on the accelerated bridges). If both are connected, disconnect one bridge cable from the active bridges on the second appliance. This will prevent forwarding loops.
2. On the Features page of the first appliance, disable Traffic Processing. This disables acceleration until the HA pair is configured.
3. Repeat for the second appliance.
4. On the first appliance, go to the Configuration: Advanced Deployments: High Availability tab, show below.
5. Select the Enabled Check box.
6. Click the Configure HA Virtual IP Address link and assign a virtual IP address to the apA interface. This address will be used later to control both appliances as a unit.
7. Return to the High Availability page and, in the VRRP VRID field, assign a VRRP ID to the pair. Although the value defaults to zero, the valid range of VRRP ID numbers is 1 through 255. Within this range, you can specify any value that does not belong to another VRRP device on your network.
8. In the Partner SSL Common Name field, type the other appliance's SSL Common Name, which is displayed on that appliance's Configuration: Advanced Deployments: High Availability tab, in the Partner SSL Common Name field. The SSL credentials used here are factory-installed.
9. Click Update.
10. Repeat steps 3-8 on the second appliance. If you are managing the appliance via an accelerated bridge (such as apA), you may have to reconnect the Ethernet cable that you removed in step 1 to connect to the second appliance. If so, plug this cable in and disconnect the corresponding cable on the first appliance.
11. With your browser, navigate to the virtual IP address of the HA pair. Enable Traffic Processing on the Features page. Any further configuration will be performed from this virtual address.
12. Plug in the cable that was left disconnected.
13. On each appliance, the Configuration: Advanced Deployments: High Availability page should now show that high availability is active and that one appliance is the primary and the other is the secondary. If this is not the case, a warning banner appears at the top of the screen, indicating the nature of the problem.

Figure 1. High-availability configuration page

Updating Software on a High-Availability Pair

Aug 09, 2017

Updating the SD-WAN software on an HA pair causes a failover at one point during the update.

Note: Clicking the Update button terminates all open TCP connections.

To update the software on an HA pair

1. Log on to both appliances.
2. On the secondary appliance, update the software and reboot. After the reboot, the appliance is still the secondary. Verify that the installation succeeded. The primary appliance should show that the secondary appliance exists but that automatic parameter synchronization is not working, due to a version mismatch.
3. On the primary appliance, update the software, and then reboot. The reboot causes a failover, and the secondary appliance becomes the primary. When the reboot is completed, HA should become fully established, because both appliances are running the same software.

Saving/Restoring Parameters of an HA Pair

Aug 09, 2017

The System Maintenance: Backup/Restore function can be used to save and restore parameters of an HA pair as follows:

To back up the parameters

Use the backup feature as usual. That is, log on to the GUI through the HA VIP address (as is normal when managing the HA pair) and, on the System Management: Backup/Restore page, click Download Settings.

To restore the parameters

1. Disable HA on both appliances by clearing the Enabled check box on the Configuration: Advanced Deployments: High Availability (HA) tab.
2. Unplug a network cable from the bridge of one appliance. (Call it “Appliance A.”)
3. Unplug the power cord from Appliance A.
4. Restore the parameters on the other appliance (Appliance B), by uploading a previously saved set of parameters on the System Maintenance: Backup/Restore page and clicking Restore Settings. (Completing this operation requires a restart, which reenables HA).
5. Wait for Appliance B to restart. It becomes the primary.
6. Restart Appliance A.
7. Log on to Appliance A’s GUI and reenable HA on the Configuration: Advanced Deployments: High Availability (HA) tab. The appliance get its parameters from the primary.
8. Plug in the network cable removed in step 2.

Both appliances are now restored and synchronized.

Troubleshooting High Availability Pairs

Aug 09, 2017

If the appliances report any failure to enter high-availability mode, the error message will also note the cause. Some issues that can interfere with high-availability mode are:

- The other appliance is not running.
- The HA parameters on the two appliances are not identical.
- The two appliances are not running the same software release.
- The two appliances do not have the same model number.
- Incorrect or incomplete cabling between the appliances does not allow the HA heartbeat to pass between them.
- The HA/Group Mode SSL Certificates on one or both appliances are damaged or missing.

NetScaler SD-WAN 4000 and 5000 WANOP Appliances

Aug 09, 2017

Citrix NetScaler SD-WAN 4000/5000 WANOP appliances are high-performance WAN accelerators for busy datacenters. These appliances combines multiple virtual accelerator instances with a single virtual instance of the NetScaler load-balancer, providing the performance of multiple SD-WAN WANOP appliances in a single package.

SD-WAN 4000/5000 WANOP WAN accelerators are the high end of the Citrix NetScaler SD-WAN product line. They are designed to accelerate sites with WAN links with speeds in excess of 155 Mbps, especially busy datacenters that communicate with a large number of branch and regional sites.

Figure 1. Typical Use Case

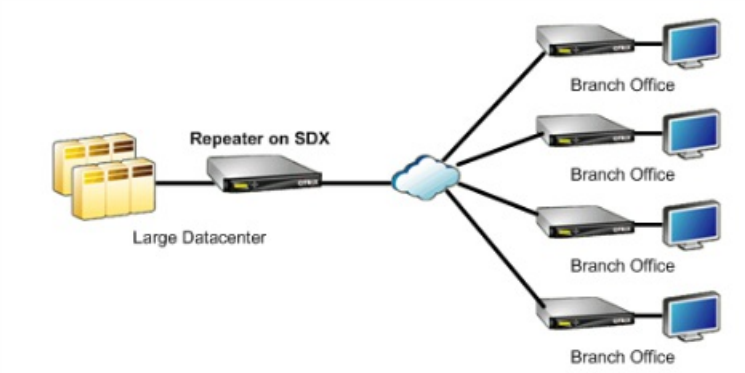

A single SD-WAN WANOP 4000/5000 appliance can support WAN speeds of up to 2 Gbps and up to 5000 XenApp/XenDesktop users.

For datacenters needing even more performance, multiple SD-WAN WANOP 4000/5000 appliances can be deployed as a load-balanced array using the WCCP clustering feature.

Figure 2. Load balancing multiple SD-WAN WANOP 4000/5000 appliances

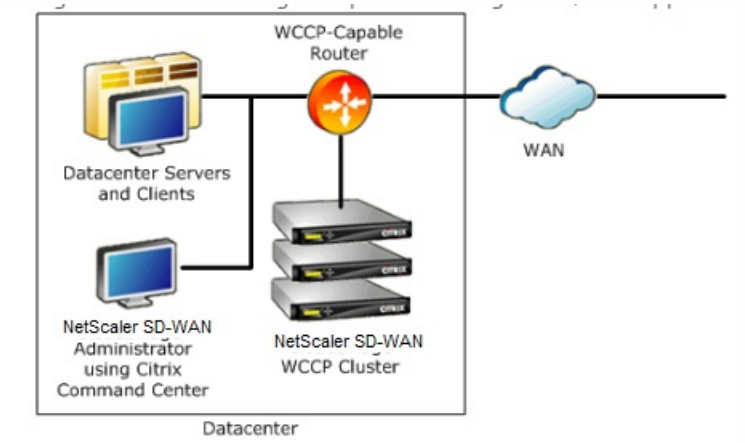

SD-WAN WANOP 4000/5000 is recommended at the hub of a hub-and-spoke deployment, where smaller appliances are used at the spokes, whenever the link speed or the number of XenApp/XenDesktop users is higher than can be supported

by a smaller appliance.

DC to DC Replication

If you require a secondary data center, SD-WAN WANOP 4000/5000 appliances can provide optimization for Data-Center to Data-Center replication. This optimization improves replication time and reduces bandwidth consumption.

For details on how to configure a SD-WAN WANOP appliance for DC-to-DC replication with NetApp SnapMirror, see http://support.citrix.com/article/CTX137181.

Architecture

Aug 09, 2017

Internally, the SD-WAN 4000/5000 appliance contains several virtual machines:

- A Xen hypervisor
- A NetScaler instance
- At least two accelerator instances
- A management server instance that manages the GUI and other tasks
- Internal networking

Figure 2. SD-WAN 4000/5000 virtual machines, internal networks, and external port usage (inline deployment shown)

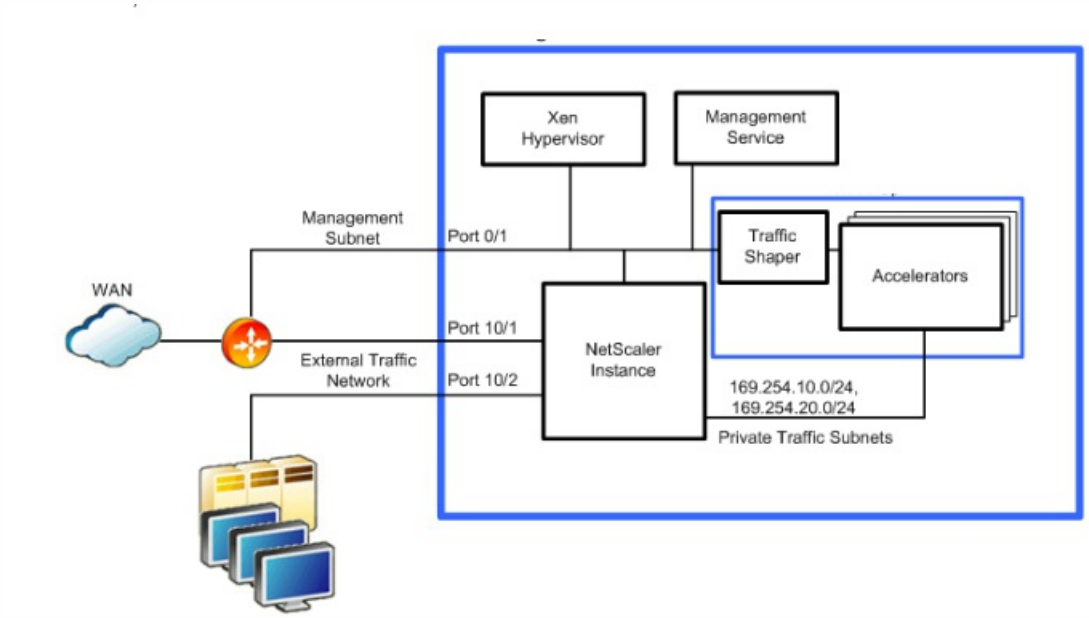

No WAN traffic enters or leaves the accelerators except as configured in the NetScaler instance. When the appliance is first used, the Provisioning Wizard sets up an initial configuration that provides communication and load balancing between the NetScaler instance and the accelerators.

The management service is the management configuration interface for the appliance, and provides access to key operating and monitoring elements of the appliance. The management service displays SD-WAN parameters as if they were from a single accelerator, and all changes made through this interface are applied to all the accelerator instances.

The Xen hypervisor hosts all the virtual machines. The hypervisor is not user-configurable and should not be accessed except at the request of Citrix.

Internal and External Networks

The external network interfaces are divided into two categories: traffic interfaces and management interfaces.

Traffic Interfaces—The traffic interfaces include all the network interfaces except ports 0/1 and 0/2, which are used only for management. Acceleration takes place only on the traffic interfaces.

Note: You must keep the traffic interfaces isolated from the management interface to prevent ARP flapping and other problems. This isolation can be achieved physically or by tagging management interface and traffic interface packets with

different VLANs.

Management subnet—The virtual machines connect directly to the external management subnet, with different IP addresses for the management service, NetScaler instance, and XenServer.

Note: You must keep the traffic interfaces isolated from the management interface to prevent ARP flapping and other problems. This isolation can be achieved physically or by tagging management interface and traffic interface packets with different VLANs.

Private Internal traffic subnet—The accelerators' accelerated ports are connected to the NetScaler instance internally in a one-arm mode, using an internal traffic subnet. There is no direct connection between the instances' accelerated ports and the appliance's external ports. All accelerated traffic to the accelerators is controlled by the NetScaler instance.

Since this internal subnet is not accessible from outside the appliance, it uses non-routable subnets in the 169.254.0.0/16 range. The NetScaler instance provides NAT for features that require routable access to the accelerator Only the following two features of the accelerators require IP addresses that can be reached from the outside world:

- The signaling IP address, used for secure peering and the SD-WAN Plugin.
- IP addresses, used for communication with the router when the WCCP protocol is used.

In both cases, the number of externally visible IP addresses is independent of the number of accelerators the appliance has.

The internal traffic subnet requires two IP addresses per accelerator, plus an address for the NetScaler, plus one or two WCCP VIP addresses if WCCP is used. Since the internal network is private, it has an abundance of address space for these tasks.

Data Flow on the Private Traffic Subnet—The one-arm connection between the NetScaler instance and the accelerators uses the SD-WAN virtual inline mode, in which the NetScaler instance routes packets to the accelerators and the accelerators route them back to the NetScaler instance. Traffic flow over this internal traffic subnet is identical regardless of whether the mode visible to the outside world (on the external interfaces) is inline, virtual inline, or WCCP.

This traffic requires the SD-WAN "Return to Ethernet Sender" option, and the NetScaler MAC Address Forwarding and Use Subnet IP options, which are enabled by the Provisioning Wizard.

Deployment Mode Summary

The differences between WCCP mode, inline mode, and virtual inline mode can be summarized as follows:

- WCCP mode is a one-arm configuration. The accelerators establish WCCP control channels with the router. In WCCP mode, only one or two accelerators manage the WCCP control channel on behalf of all the accelerators. Data traffic is load-balanced across all the accelerators. When GRE encapsulation is used, the NetScaler instance performs GRE encapsulation/decapsulation on the data stream between itself and the router, allowing the data between the NetScaler and the accelerators to use a decapsulated, Level-2 configuration.
- Inline mode operates much the same as WCCP mode internally, but externally the appliance emulates a bridge, and no WCCP control channel is established. A packet that enters the appliance on one bridge port exits through the other bridge port. SD-WAN 4000 and 5000 appliances have multiple bridges to support multiple inline links.
- In virtual inline mode (used when WCCP and inline modes are not feasible), the appliance is deployed in a one-arm configuration, much like WCCP, but without the WCCP control channel. Traffic is sent to the appliance from the router, using policy-based routing (PBR) rules. The appliance processes the traffic and returns it to the router.

Figure 3. WCCP and virtual inine cabling

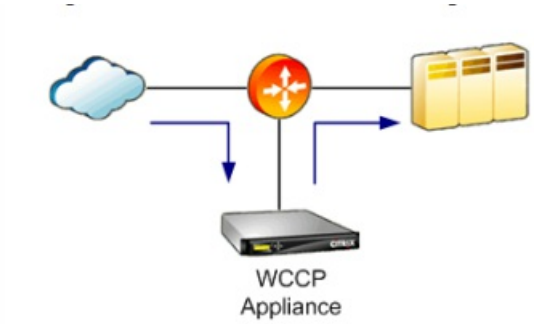

See SD-WAN 4000/5000 virtual machines, internal networks, and external port usage for a diagram of port usage on SD-WAN 4000/5000 appliances. Traffic ports are arranged as a set of accelerated bridges, while the management ports are independent. Typically only one management port is used.

Figure 4. Inline cabling

Accelerated Bridges

SD-WAN 4000/5000 appliances have multiple accelerated bridges. Different models have different numbers and types of bridge ports. The two ports making up such a bridge are called an "accelerated pair." All current models include a built-in network bypass function. (Some older SD-WAN 4000-500 and 4000-1000 units do not include network bypass). The network bypass function (also called "fail to wire") connects pairs of ports together if the appliance fails as a result of either power loss or software failure (as determined by an internal watchdog timer).

Inline deployment. The bypass function allows SD-WAN 4000/5000 to be deployed in line with your WAN, typically between your LAN and your WAN router, without introducing a point of network failure.

The accelerated bridges support either 1 Gbps or 10 Gbps data rates. Ethernet and SFP+ interfaces are supported, depending on model.

One-arm deployment. One-arm deployments are also supported, using WCCP or virtual inline modes. With such deployments, a SD-WAN 4000/5000 traffic port is usually connected directly to a port on the WAN router. The other port on the bridged pair is left unconnected.

Performance considerations. Inline deployments provide higher performance than the one-arm deployments, because the use of two ports instead of one doubles the peak throughput of the interfaces.

Peak throughput is important with SD-WAN 4000/5000 appliances, because the compressor provides acceleration in proportion to the compression ratio. That is, a connection that achieves 100:1 compression transfers data one hundred times faster than an uncompressed connection, provided that the rest of the network path can keep up.

For example, take a datacenter with a 500 Mbps WAN link and a 1 Gbps LAN. The small 2:1 speed ratio between the WAN and LAN allows compression to provide only a 2x speedup on a whole-link basis, because there is no way to get data onto or off of the LAN at speeds above 1 Gbps. A 10 Gbps LAN, which allows a tenfold increase in peak data rates, is recommended for use with SD-WAN 4000/5000 deployments.

When a SD-WAN 4000/5000 appliance is deployed in a one-arm mode, the peak transfer rate is cut in half. A SD-WAN

4000/5000 in one-arm mode, connected to the router with a 1 Gbps LAN interface, saturates this interface when the WAN is running at full speed in both directions. For good performance, a SD-WAN 4000/5000 must have a LAN interface that is much faster than the WAN. When the appliance is connected directly to the router in a one-arm mode, use a 10 Gbps router port.

Note: The 10 Gbps ports support 10 Gbps only. They do not negotiate lower speeds. Use the 1 Gbps ports for 1 Gbps networks.

Other Ports

A SD-WAN 4000/5000 appliance has at least two non-accelerated ports. Port 0/1 is typically used for management, Port 0/2 is present but typically not used. A Light Out Management (LOM) port is also provided. An RS-232 port can be used for management.

SD-WAN 4000

Aug 09, 2017

Citrix SD-WAN 4000 are 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 48 gigabytes (GB) of memory. The Citrix SD-WAN 4000 have a bandwidth of 310Mbps, 500Mbps, and 1Gbps, respectively.

The following figures shows the front panel of the Citrix SD-WAN 4000 appliance.

Figure 1. Citrix SD-WAN 4000, front panel (without FTW cards) and (with FTW cards)

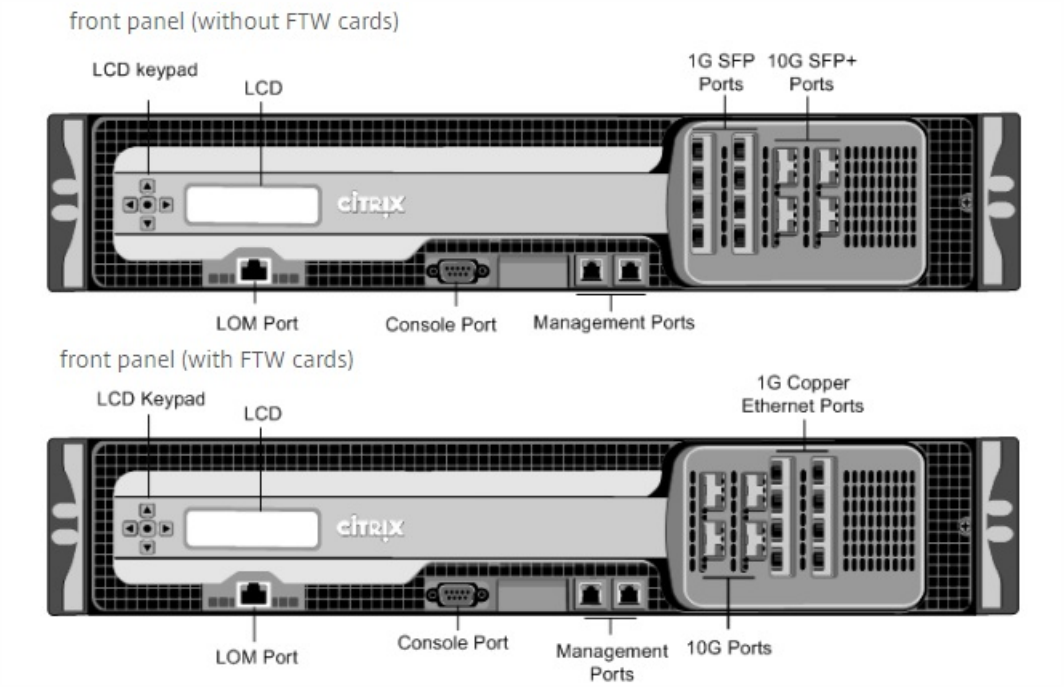

The Citrix SD-WAN 4000 appliances have the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
  Note: The LEDs on the LOM port are not operational by design.
- RS232 serial console port.
- Two 10/100/1000Base-T copper Ethernet management ports (RJ45). These ports are used to connect directly to the appliance for system administration functions.
- Network Ports
  - SD-WAN 4000 (without FTW cards). Eight 1G SFP ports and four 10G SFP+ ports.
  - SD-WAN 4000 (with FTW cards). Eight 1G copper Ethernet ports and four 10G ports.

The following figure shows the back panel of the Citrix SD-WAN 4000 appliance.

Figure 2. Citrix SD-WAN 4000, back panel

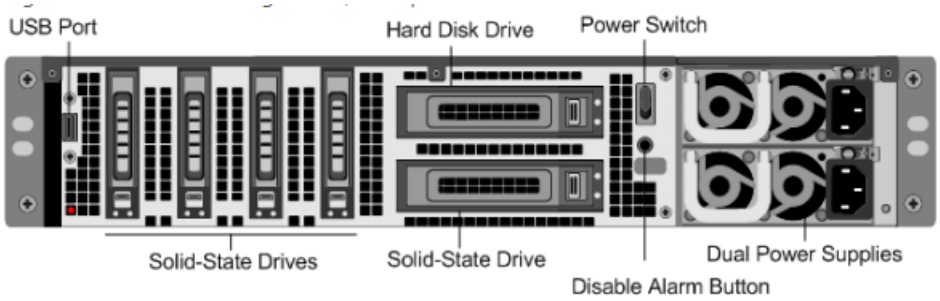

The following components are visible on the back panel of the Citrix SD-WAN 4000 appliance:

- Four 600 GB removable solid-state drives, which store the appliance's compression history. The 256 GB solid-state drive below the hard disk drive stores the appliance's software.
- USB port (reserved for a future release).
- A 1 TB removable hard disk drive.
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.
- Disable alarm button. This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Dual power supplies (either AC or DC), each rated at 850 watts, 100-240 volts.

SD-WAN 5000

Aug 09, 2017

Citrix SD-WAN 5000 are 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 96 gigabytes (GB) of memory. The Citrix SD-WAN 5000 have a bandwidth of 1.5Gbps and 2Gbps respectively.

The following figure shows the front panel of the Citrix SD-WAN 5000 appliance.

Figure 1. Citrix SD-WAN 5000, front panel

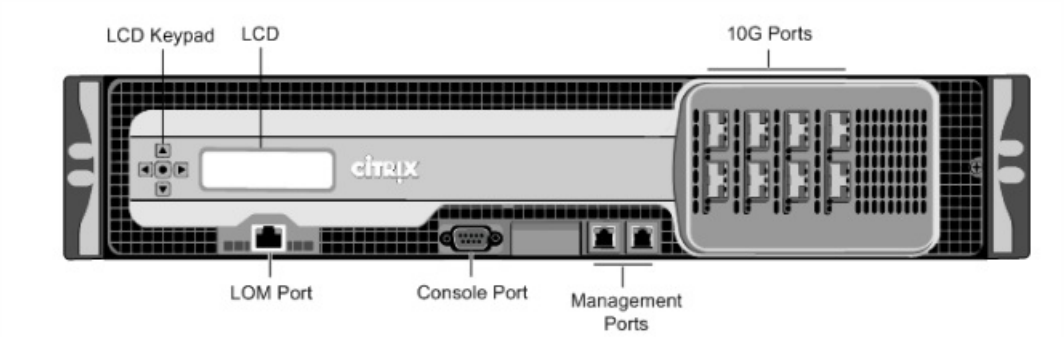

The Citrix SD-WAN 5000 appliance has the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
  Note: The LEDs on the LOM port are not operational by design.
- RS232 serial console port.
- Two 10/100/1000Base-T copper Ethernet management ports (RJ45). These ports are used to connect directly to the appliance for system administration functions.
- Eight 10G ports.

The following figure shows the back panel of the Citrix SD-WAN 5000 appliance.

Figure 2. Citrix SD-WAN 5000, back panel

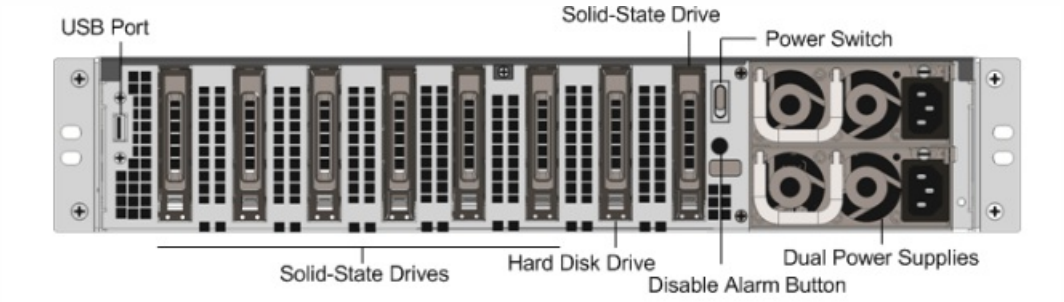

The following components are visible on the back panel of the Citrix SD-WAN 5000 appliance:

- Six 600 GB removable solid-state drives, which store the appliance's compression history. The 256 GB solid-state drive next to the power supplies store the appliance's software.
- USB port (reserved for a future release).

- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.
- A 1 TB removable hard disk drive.
- Disable alarm button. This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Dual power supplies (either AC or DC), each rated at 650 watts, 100-240 volts.

Field Replaceable Units

Aug 09, 2017

Citrix SD-WAN 4000/5000 field replaceable units (FRU) are components that can be quickly and easily removed from the appliance and replaced by the user or a technician at the user's site. The FRUs in a Citrix SD-WAN 4000/5000 appliance can include DC or AC power supplies, and solid-state and hard-disk drives.

Note: By default the appliance ships with AC power supplies. DC power supply is orderable.

Power Supply

Aug 09, 2017

Citrix SD-WAN 4000/5000 appliances are configured with dual power supplies but can operate with only one power supply. The second power supply serves as a backup.

For power-supply specifications, see "Hardware Platforms," which describes the various platforms and includes a table summarizing the hardware specifications.

Table 1. LED Power Supply Indicators

Power Supply Type	LED Color	LED Indicates
AC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing GREEN	Power supply is in standby mode.
	GREEN	Power supply is functional.
	RED	Power supply failure.
DC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing BLUE	Power supply is in standby mode.
	BLUE	Power supply is functional.
	RED	Power supply failure.

Electrical Safety Precautions for Power Supply Replacement

- Make sure that the appliance has a direct physical connection to earth ground during normal use. When installing or repairing an appliance, always connect the ground circuit first and disconnect it last.
- Always unplug any appliance before performing repairs or upgrades.
- Never touch a power supply when the power cord is plugged in. As long as the power cord is plugged in, line voltages are present in the power supply even if the power switch is turned off.

Replacing an AC Power Supply

Replace an AC power supply with another AC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace an AC power supply on a Citrix SD-WAN 4000/5000 appliance

1. Align the semicircular handle perpendicular to the power supply. Loosen the thumbscrew and press the lever toward the handle and pull out the existing power supply, as shown in the following figure.
   Figure 1. Removing the Existing AC Power Supply

2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot and press against the semicircular handle until you hear the power supply snap into place.
   Figure 2. Inserting the Replacement AC Power Supply

5. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Replacing a DC Power Supply

Replace a DC power supply with another DC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace a DC power supply on a Citrix SD-WAN 4000/5000 appliance

1. Loosen the thumbscrew and press the lever towards the handle and pull out the existing power supply, as shown in the following figure.
   Figure 3. Removing the Existing DC Power Supply

2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot while pressing the lever towards the handle. Apply firm pressure to insert the power supply firmly into the slot.
   Figure 4. Inserting the Replacement DC Power Supply

5. When the power supply is completely inserted into its slot, release the lever.
6. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Solid-State Drive

Aug 09, 2017

A solid-state drive (SSD) is a high-performance device that stores data in solid-state flash memory.

Replacing a Solid-State Drive

The SD-WAN 4000/5000 software is stored on the solid-state drive (SSD).

To replace a solid-state drive

1. Shutdown the appliance.
2. Locate the SSD on the back panel of the appliance. Push the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Figure 1. Removing the Existing Solid-State Drive

3. Verify that the replacement SSD is the correct type for the platform.
4. Pick up the new SSD, open the drive handle fully to the left or up, and insert the drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top if the drive is inserted horizontally or at the right if the drive is inserted vertically.
   Figure 2. Inserting the Replacement Solid-State Drive

5. Turn on the appliance.
6. Log on to the default IP address by using a web browser, or connect to the serial console by using a console cable, to perform the initial configuration.

Hard Disk Drive

Aug 09, 2017

The NetScaler and SD-WAN virtual machines are hosted on the hard-disk drive.

Replacing a Hard Disk Drive

Verify that the replacement hard disk drive is the correct type for the SD-WAN 4000/5000 platform.

To install a hard disk drive

1. Shut down the appliance.
2. Locate the hard disk drive on the back panel of the appliance.
3. Disengage the hard disk drive by pushing the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Figure 1. Removing the Existing Hard Disk Drive

4. Pick up the new disk drive, open the drive handle fully to the left, and insert the new drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the hard drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top.
   Figure 2. Inserting the Replacement Hard Disk Drive

5. Turn on the appliance.

Supported Features

Aug 09, 2017

Table 1. Features Table for Citrix SD-WAN 4000 and 5000 Series Appliances

	Citrix SD-WAN 4000 series	Citrix SD-WAN 5000 series
AutoConfiguration	N	N
SD-WAN Connector	Y	Y
SD-WAN Plug-In	Y	Y
Compression	Y	Y
RPC over HTTPS	Y	Y
SSL Compression	Y	Y
TCP Acceleration	Y	Y
Traffic Shaping	Y	Y
Video Caching	N	N
Windows File System Acceleration	Y	Y
Windows Outlook Acceleration	Y	Y
XenApp/ XenDesktop Acceleration	Y	Y
Group Mode Mode	N	N
High Availability Mode	Y	Y
Inline Mode	Y	Y
Virtual Inline Mode	Y	Y
WCCP Mode	Y	Y
VLANs	Y	Y

Summary of Hardware Specifications

Aug 09, 2017

The following tables summarize the specifications of the Citrix NetScaler SD-WAN 4000/5000 WANOP hardware platforms.

Table 1. Citrix NetScaler SD-WAN 4000/5000 WANOP Appliances

	SD-WAN 4000/5000				
Platform Performance					
Bandwidth	310 Mbps	500 Mbps	1.0 Gbps	1.5 Gbps	2.0 Gbps
Maximum HDX sessions	750	1,200	2,500	3,500	5,000
Total sessions	40,000	60,000	120,000	20,000	160,000
Acceleration Plug-in CCUs	1,100	1,800	3,000	3,600	4,800
Hardware Specifications					
Processor	Dual Intel E5645	Dual Intel E5645	Dual Intel E5645	Dual Intel X5680	Dual Intel X5680
Total disk space	3.2 TB	3.2 TB	3.2 TB	4.2 TB	4.2 TB
SSD (dedicated compression history)	2 TB	2 TB	2 TB	3 TB	3 TB
HDD	1 TB	1 TB	1 TB	1 TB	1 TB
RAM	48 GB	48 GB	48 GB	96 GB	96 GB
Network interfaces	4 x 10GigE SX and 8 x 1GigE TX Bypass	4 x 10GigE SX and 8 x 1GigE TX Bypass*	4 x 10GigE SX and 8 x 1GigE TX Bypass*	8 x10GigE SR Bypass	8 x10GigE SR Bypass

*See Citrix NetScaler 4000	**SD-WAN 4000/5000**				
Transceiver support	N/A	N/A	N/A	N/A	N/A
Note: Some non fail-to-wire SD-WAN 4000 appliances require SFP+ transceivers.					
Power supplies	2	2	2	2	2
Physical Dimensions					
Rack units	2	2	2	2	2
System width	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets
System depth	25.4"/64.5 cm	25.4"/64.5 cm	25.4"/64.5 cm	25.4"/64.5 cm	25.4"/64.5 cm
System weight	46 lbs (20.9 kg)	46 lbs (20.9 kg)	46 lbs (20.9 kg)	49 lbs (22.2 kg)	49 lbs (22.2 kg)
Shipping dimensions and weight	37” x 24” by 11” 59 lbs 94 x 61 x 28 cm 26.8 kg	37” x 24” by 11” 59 lbs 94 x 61 x 28 cm 26.8 kg	37” x 24” by 11” 59 lbs 94 x 61 x 28 cm 26.8 kg	37” x 24” by 11” 61 lbs 94 x 61 x 28 cm 27.7 kg	37” x 24” by 11” 61 lbs 94 x 61 x 28 cm 27.7 kg
Environmental and Regulatory					
Input voltage and frequency ranges	100-240 VAC 47-63 Hz	100-240 VAC 47-63 Hz	100-240 VAC 47-63 Hz	100-240 VAC 47-63 Hz	100-240 VAC 47-63 Hz
Power consumption	650 watts 2,200 BTU per hour.	650 watts 2,200 BTU per hour.	650 watts 2,200 BTU per hour.	850 watts 2,900 BTU per hour.	850 watts 2,900 BTU per hour.
Operating temperature	32–104° F	32–104° F	32–104° F	32–104° F	32–104° F

SD-WAN 4000/5000					
	0–40° C	0–40° C	0–40° C	0–40° C	0–40° C
Operating altitude	0–4921' (1,500 m)	0–4921' (1,500 m)	0–4921' (1,500 m)	0–4921' (1,500 m)	0–4921' (1,500 m)
Non-operating temperature	-4–140° F 20–60° C	-4–140° F 20–60° C	-4–140° F 20–60° C	-4–140° F 20–60° C	-4–140° F 20–60° C
Allowed relative humidity	5%-95%, non-condensing	5%-95%, non-condensing	5%-95%, non-condensing	5%-95%, non-condensing	5%-95%, non-condensing
Safety certifications	UL, TUV-C	UL, TUV-C	UL, TUV-C	UL, TUV-C	UL, TUV-C
Electromagnetic emissions certifications and susceptibility standards	FCC (Part 15 Class A), DoC, CE, VCCI, CNS, AN/NES	FCC (Part 15 Class A), DoC, CE, VCCI, CNS, AN/NES	FCC (Part 15 Class A), DoC, CE, VCCI, CNS, AN/NES	FCC (Part 15 Class A), DoC, CE, VCCI, CNS, AN/NES	FCC (Part 15 Class A), DoC, CE, VCCI, CNS, AN/NES
Environmental compliance	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE

Lights Out Management Port of the SD-WAN WAN OP 4000/5000 Appliance

Aug 09, 2017

The SD-WAN 4000/5000 appliances have an Intelligent Platform Management Interface (IPMI), also known as the Lights out Management (LOM), port on the front panel of the appliance. By using the LOM, you can remotely monitor and manage the appliance, independently of the SD-WAN 4000/5000 software. You can remotely change the IP address, perform different power operations, and obtain health monitoring information of the appliance by connecting to the appliance through the LOM port.

By connecting the LOM port over a dedicated channel that is separate from the data channel, you can make sure that connectivity to the appliance is maintained even if the data network is down.

Accessing the LOM Port by using a Web Browser

By using a web browser you can remotely log on to the LOM port to obtain information about the appliance and perform different operations on the appliance.

To access the LOM by using a web browser

1. In a web browser, type the IP address of the LOM port. For initial configuration, type the port's default address: http://192.168.1.3
2. In the User Name box, type **nsroot**.
3. In the Password box, type **nsroot**.

Configuring the LOM Port

You can use the Intelligent Platform Management Interface (IPMI), also known as the Lights Out Management (LOM) port, to remotely monitor and manage the appliance, independently of the NetScaler software. For initial configuration of the lights-out management (LOM) port, connect to the port's default IP address and change it to the address that you want to use for remote monitoring and management. Also specify the administrator credentials and the network settings.

Note: The LEDs on the LOM port are unoperational by design.

To configure the NetScaler LOM Port

1. Connect the LOM port to a management workstation or network.
2. In a web browser, type: http://192.168.1.3.
   Note: The NetScaler LOM port is preconfigured with the IP address 192.168.1.3 and subnet mask 255.255.255.0.
3. In the User Name box, type **nsroot**.
4. In the Password box, type **nsroot**.
5. On the Configuration tab, click Network and type values for the following parameters:
   - IP Address—IP address of the LOM port.
   - Subnet Mask—Subnet mask used to define the subnet of the LOM port.
   - Default Gateway—IP address of the router that connects the LOM port to the network.
6. Click Save.

Power Cycling the Appliance

You can remotely turn off the appliance and turn it back on. The result is similar to pressing the power button on the back

panel of the appliance for less than two seconds.

To power cycle the appliance

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click Remote Control.
4. Under Options, click Power Control, and then click Power Cycle System.
5. Click Perform Action.

Accessing the Appliance by using the Access Console

The LOM port allows you to remotely access and manage the appliance by logging on to a redirected console.

To access the appliance by using the access console

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click Remote Control.
4. Under Options, click Console Redirection.
5. Click Launch Console, and then click Yes.
6. Type the administrator credentials for the appliance.

Obtaining Health Monitoring Information

You can log on to the LOM port to view the health information about the appliance. All system sensor information, such as system temperature, CPU temperature, status of fan and power supplies, appears on the sensor readings page.

To obtain health monitoring information

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click System Health.
4. Under Options, click Sensor Readings.

Power Control Operations using the LOM Port

You can remotely perform different power control operations, such as restarting the appliance, performing a graceful shutdown, and performing a forced shutdown, by using the LOM port.

To perform power control operations

1. In a web browser, log on to the LOM port by using the administrator credentials.
2. In the Menu bar, click Remote Control.
3. Under Options, click Power Control, and then select one of the following options:
   - **Reset System**—Restart the appliance.
   - **Power Off System – Immediate**—Disconnect power to the appliance without shutting down the appliance.
   - **Power On System**—Turn on the appliance.
   - **Power Cycle System**—Turn off the appliance, and then turn it back on.
4. Click Perform Action.

Preparing for Installation

Aug 09, 2017

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Once you are satisfied that your appliance has been delivered to your expectations, verify that the location where the appliance will be installed meets temperature and power requirements and that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel should install, maintain, or replace the appliance, and efforts should be taken to ensure that all cautions and warnings are followed.

Unpacking the Appliance

Aug 09, 2017

Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

Use the following list to verify that you received everything that should have been included in the box.

- The appliance you ordered
- One RJ-45 to DB-9 adapter
- One 6 ft RJ-45/DB-9 cable
- Two power cables
- One fiber patch cable
- One standard 4-post rail kit
  Note: If the kit that you received does not fit your rack, contact your Citrix sales representative to order the appropriate kit.

In addition to the items included in the box with your new appliance, you will need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you will connect to your network
- One available Ethernet port on your network switch or hub for each Ethernet port you want to connect to your network
- A computer to serve as a management workstation

Preparing the Site and Rack

Aug 09, 2017

There are specific site and rack requirements for the SD-WAN 4000/5000 appliance. You must make sure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and help ensure a smooth installation.

Site Requirements

The appliance should be installed in a server room or server cabinet with the following features:

Environment control

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment.

Power density

Wiring capable of handling at least 4,000 watts per rack unit in addition to power needs for the CRAC.

Rack Requirements

The rack on which you install your appliance should meet the following criteria:

Rack characteristics

Racks should be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it should be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks should be installed in rows spaced at least 1 meter/3 feet apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

Power connections

At minimum, two standard power outlets per unit.

Network connections

At minimum, Ethernet connection per rack unit.

Space requirements

Two empty rack units for SD-WAN 4000/5000 appliances.

Note: You can order the following rail kits separately.

- Compact 4-post rail kit, which fits racks of 23 to 33 inches.
- 2-post rail kit, which fits 2-post racks.

Cautions and Warnings

Aug 09, 2017

Electrical Safety Precautions

Updated: 2014-02-06

Caution: During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.
Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Remove all jewelry and other metal objects that might come into contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.
- Use a regulating, uninterruptible power supply (UPS) to protect the appliance from power surges and voltage spikes, and to keep the appliance operating in case of power failure.
- Never stack the appliance on top of any other server or electronic equipment.
- All appliances are designed to be installed on power systems that use TN earthing. Do not install your device on a power system that uses either TT or IT earthing.
- Make sure that the appliance has a direct physical connection to the earth during normal use. When installing or repairing an appliance, always make sure that the ground circuit is connected first and disconnected last.
- Make sure that a fuse or circuit breaker no larger than 120 VAC, 15 A U.S. (240 VAC, 16 A international) is used on all current-carrying conductors on the power system to which your appliances are connected.
- Do not work alone when working with high voltage components.
- Always disconnect the appliance from power before removing or installing any component. When disconnecting power, first shut down the appliance, and then unplug the power cords of all the power supply units connected to the appliance. As long as the power cord is plugged in, line voltages can be present in the power supply, even when the power switch is OFF.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload. Always unplug any appliance before performing repairs or upgrades.
- Do not overload the wiring in your server cabinet or on your server room rack.
- During thunderstorms, or anticipated thunderstorms, avoid performing any hardware repairs or upgrades until the danger of lightning has passed.
- When you dispose of an old appliance or any components, follow any local and national laws on disposal of electronic waste.
- To prevent possible explosions, replace expired batteries with the same model or a manufacturer-recommended substitute and follow the manufacturer's instructions for battery replacement.
- Never remove a power supply cover or any sealed part that has the following label:

Appliance Precautions

- Determine the placement of each component in the rack before you install the rails.
- Install the heaviest appliance first, at the bottom of the rack, and then work upward. Distribute the load on the rack evenly. An unbalanced rack is hazardous.
- Allow the power supply units and hard drives to cool before touching them.
- Install the equipment near an electrical outlet for easy access.
- Mount equipment in a rack with sufficient airflow for safe operation.
- For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

Rack Precautions

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Oct 12, 2017

Most appliances can be installed in standard server racks that conform to EIA-310-D specification. The appliances ship with a set of rails, which you must install before you mount the appliance. The only tools that you need for installing an appliance are a Phillips screwdriver and a flathead screwdriver.

Citrix SD-WAN 4000/5000 appliance requires two rack units.
Each appliance ships with a mounting rail kit that contains two rail assemblies, one for the left side and the other for the right side of the appliance, and screws to attach the rails. An assembly consists of an inner rail and a rack rail. The supplied rail kit is 28 inches long (38 inches extended). Contact your Citrix sales representative to order a 23-inch (33 inches extended) rail kit.

To mount the appliance, you must first install the rails and then install the appliance in the rack.

Perform the following tasks to mount the appliance:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

> ### Warning
>
> If you are installing the appliance as the only unit in the rack, mount it at the bottom. If the rack contains other units, make sure that the heaviest unit is at the bottom. If the rack has stabilizing devices available, install them before mounting the appliance.

> ### Note
>
> The same rail kit is used for both square-hole and round-hole racks. See below for specific instructions for threaded, round-hole racks.

To remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the latch until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

To attach the inner rails to the appliance

1. Position the right inner rail behind the handle on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws: 5 per side, as shown in the following figure.
   Figure 1. Attaching inner rails

4. Repeat steps 1 through 3 to install the left inner rail on the other side of the appliance.

To install the rack rails on the rack

1. If you have a round-hole, threaded rack, skip to step 3.
2. Install square nut retainers into the front post and back post of the rack as shown in the following figures. Before inserting a screw, be sure to align the square nut with the correct hole for your appliance. The three holes are not evenly spaced.
   Figure 2. Installing Retainers into the Front Rack Posts and Figure 3. Installing Retainers into the Rear Rack Posts

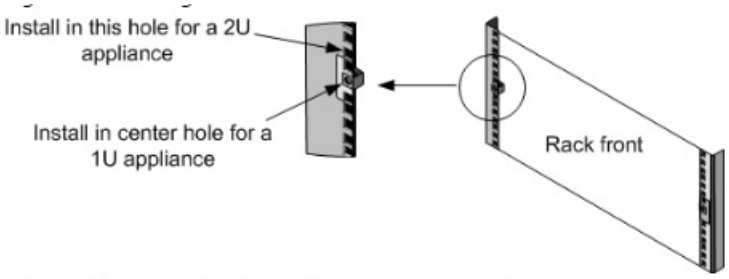

Figure 3. Installing Retainers into the Rear Rack Posts

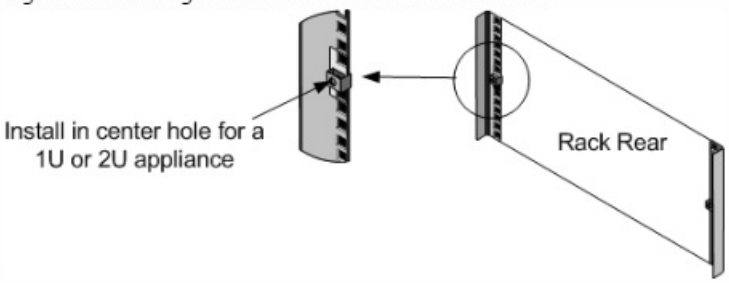

3. Install the adjustable rail assembly into the rack as shown in the following figures. Use a screw to lock the rear rail flange into the rack. With the screw securing the rail in place, you can optionally remove the latching spring.

Figure 4. Installing the Rail Assembly to the Rack

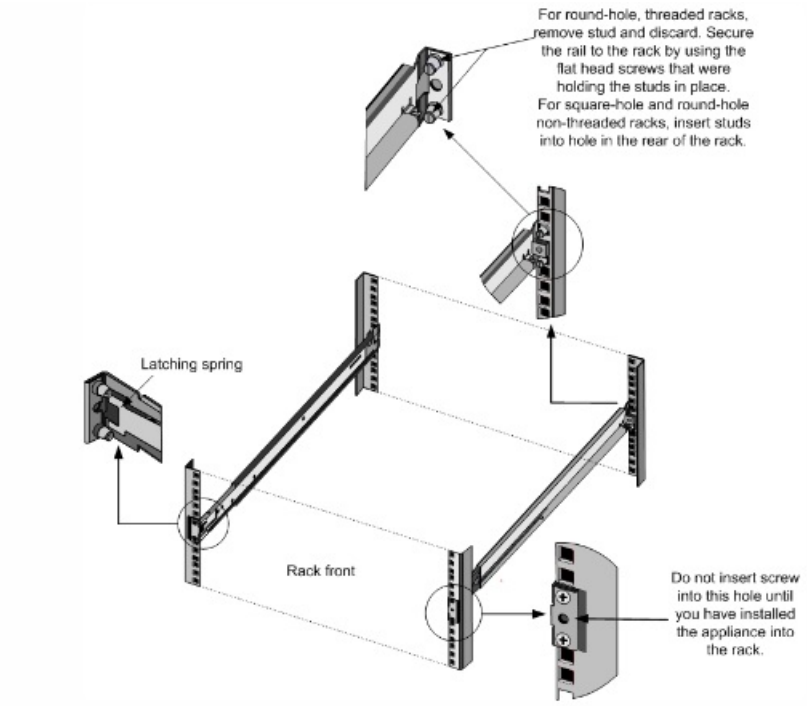

To install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.

Figure 5. Rack Mounting the Appliance

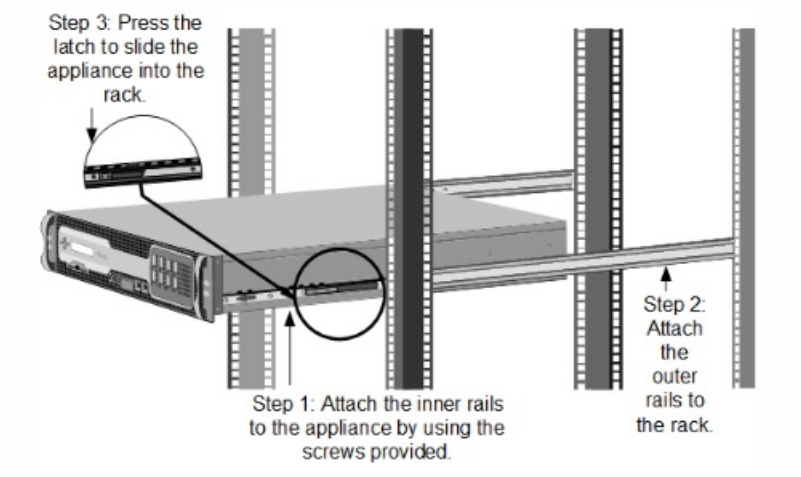

Installing and Removing 1G SFP Transceivers

Oct 12, 2017

A Small Form-Factor Pluggable (SFP) is a compact transceiver that can operate at speeds of up to 1 gigabit per second and is available in both copper and fiber types. Inserting a 1G SFP copper transceiver converts the 1G SFP port to a 1000BASE-T port. Inserting a 1G SFP fiber transceiver converts the 1G SFP port to a 1000BASE-X port. Auto-negotiation is enabled by default on the 1G SFP port into which you insert your 1G SFP transceiver. As soon as a link between the port and the network is established, the speed and mode are matched on both ends of the cable.

Insert 1G SFP transceivers into the 1G SFP ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the 1G SFP transceiver or the appliance.

Note

Some SD-WAN 4000/5000 appliances do not require SFP transceivers.

Important

SD-WAN 4000/5000 appliances do not support 1G SFP transceivers from vendors other than Citrix Systems. Attempting to install third-party 1G SFP transceivers on your SD-WAN 4000/5000 appliance voids the warranty. Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

To install a 1G SFP transceiver

Warning

Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.

Note

The illustration in the following figures might not represent your actual appliance.

1. Remove the 1G SFP transceiver carefully from its box.
2. Align the 1G SFP transceiver to the front of the 1G SFP transceiver port on the front panel of the appliance, as shown in the following figure.
3. Figure 1. Installing a 1G SFP transceiver

3. Hold the 1G SFP transceiver between your thumb and index finger and insert it into the 1G SFP transceiver port, pressing it in until you hear the transceiver snap into place.

4. Lock the transceiver.

5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.

6. If you are using a fiber 1G SFP transceiver, do not remove the dust caps attached to the transceiver and the cable until you are ready to insert the cable.

To remove a 1G SFP transceiver

1. Disconnect the cable from the 1G SFP transceiver. If you are using a fiber optic cable, replace the dust cap on the cable before putting it away.
   Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 1G SFP transceiver.
3. Hold the 1G SFP transceiver between your thumb and index finger and slowly pull it out of the port.
4. If you are removing a fiber 1G SFP transceiver, replace the dust cap before putting it away.
5. Put the 1G SFP transceiver into its original box or another appropriate container.

Installing and Removing 10G SFP+ Transceivers

Aug 09, 2017

Note: Some SD-WAN 4000/5000 appliances do not require SFP+ transceivers.

A 10-Gigabit Small Form-Factor Pluggable (SFP+) is a compact optical transceiver that can operate at speeds of up to 10 gigabits per second. Autonegotiation is enabled by default on the 10G SFP+ ports into which you insert your 10G SFP+ transceiver. As soon as a link between the port and the network is established, the mode is matched on both ends of the cable and for 10G SFP+ transceivers, the speed is also autonegotiated.

Caution: SD-WAN 4000/5000 appliances do not support 10G SFP+ transceivers provided by vendors other than Citrix Systems. Attempting to install third-party 10G SFP+ transceivers on your SD-WAN 4000/5000 appliance voids the warranty. Insert the 10G SFP+ transceivers into the 10G SFP+ ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the transceiver or the appliance.

Caution: Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

To install a 10G SFP+ transceiver

1. Remove the 10G SFP+ transceiver carefully from its box.
   Danger: Do not look directly into fiber optic transceivers and cables. They emit laser beams that can damage your eyes.
2. Align the 10G SFP+ transceiver to the front of the 10G SFP+ transceiver port on the front panel of the appliance.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and insert it into the 10G SFP+ transceiver port, pressing it in until you hear the transceiver snap into place.
4. Move the locking hinge to the DOWN position.
5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. Do not remove the dust caps attached to the transceiver and cable until you are ready to insert the cable.

To remove a 10G SFP+ transceiver

1. Disconnect the cable from the 10G SFP+ transceiver. Replace the dust cap on the cable before putting it away.
   Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 10G SFP+ transceiver by moving the locking hinge to the UP position.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and slowly pull it out of the port.
4. Replace the dust cap on the transceiver before putting it away.
5. Put the 10G SFP+ transceiver into its original box or another appropriate container.

Install Fiber Patch Cable in Ports 10/3 and 10/4

Oct 12, 2017

Through release 7.2.1 and later on an appliance, SD-WAN ports 10/3 and 10/4 must be connected with the provided cable, as shown in the following figure.

Starting with release 7.2.2, the patch cable is no longer required, and can be omitted if:

- The appliance was shipped from the factory with release 7.2.2 or later, or
- The appliance was shipped from the factory with release 7.2.1 or earlier, but you upgrade it to 7.2.2 or later and changed the default loopback in the management service (on **System > Configuration > System > Configure Loopback Settings**).

To install the patch cable

1. Connect the LC-to-LC cable to the ports as shown in the figures above.
2. Insert one end of the cable into port 10/3.
3. Insert the other end of the cable into port 10/4.

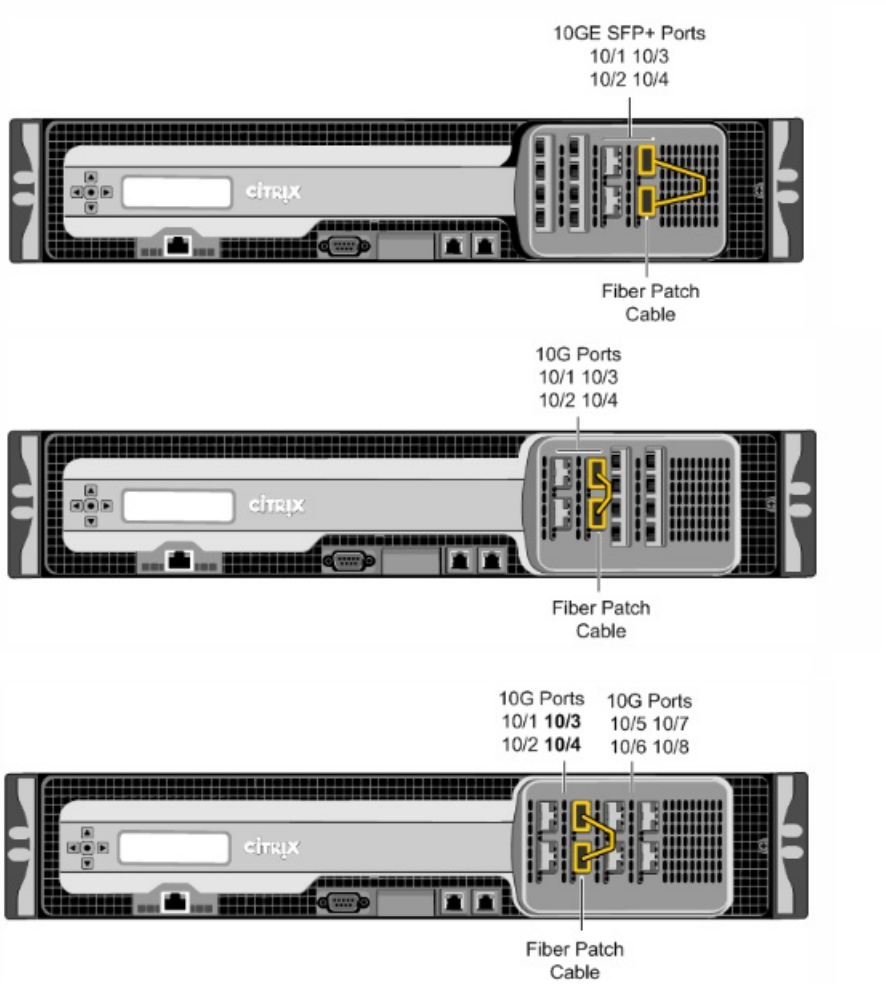

Note

If you decide to eliminate the need to use loopback cable, the ports 10/3 and 10/4 are still reserved. These ports are not available for WAN optimization.

Connecting the Cables

Oct 12, 2017

When the appliance is securely mounted on the rack, you are ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Danger: Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port on the front panel of the appliance, as shown in the following figure.
   Figure 1. Inserting an Ethernet cable

2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.

3. Verify that the LED glows amber when the connection is established.

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port that is located on the front panel of the appliance, as shown in the following figure.
   Figure 2. Inserting a console cable

2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

> **Note**
>
> To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.

The SD-WAN 4000/5000 appliance has two power supplies, with one serving as a backup. A separate ground cable is not required, because the three-prong plug provides grounding. Power up the appliance by installing one or both power cords.

To connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply, as shown in the following figure.
   Figure 3. Inserting a power cable

2. Connect the other end of the power cable to a standard 110V/220V power outlet.

3. Repeat steps 1 and 2 to connect the second power supply.

> **Note**
>
> The appliance emits a high-pitched alert if one power supply fails or if you connect only one power cable to the appliance. To silence the alarm, you can press the small red button located on the back panel of the appliance.

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the back panel of the appliance.

Caution: Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Planning the Deployment

Aug 09, 2017

SD-WAN 4000/5000 deployments require adequate planning, especially for units deployed in large datacenters:

- An appropriate appliance or group of appliances must be selected to support both the current and anticipated load.
- A deployment mode must be selected to match the requirements of your site.
- Other aspects must also be considered.

Sizing Guidelines

Aug 09, 2017

For successful deployment of one or more SD-WAN 4000/5000 appliances in your datacenter, keep the following principles in mind:

- You must provide enough SD-WAN 4000/5000 peak-load capacity, in terms of WAN bandwidth and the number of users. See the current specifications sheet for the capacities of different SD-WAN 4000/5000 models: http://www.citrix.com/content/dam/citrix/en_us/documents/products/SD-WAN-branch-repeater-spec-sheet.pdf (In the spec sheet, the number of users is referred to as "HDX sessions"). Ensure adequate peak-load capacity, both for now and for the time until you expect to upgrade. Acceleration is resource intensive, and performance suffers if the appliance runs short of resources. Never overcommit any SD-WAN appliance, especially in the datacenter. Provision your datacenter to easily accommodate peak loads.
- Provide enough capacity for expected expansion over the life of the deployment. SD-WAN 4000/5000 appliances using the same hardware platform can have their capacity upgraded with a new license as part of the Citrix pay-as-you-grow program. SD-WAN 4000/5000 models 310, 500, and 1000 use one hardware platform, and models 1500 and 2000 use another hardware platform. This means that, for example, a SD-WAN 4000/5000 500 can be converted through a license upgrade to a SD-WAN 4000/5000 1000, but not to a SD-WAN 4000/5000 1500.
- For more capacity than can be provided by a single appliance, multiple SD-WAN 4000/5000 appliances can be cascaded behind a stand-alone NetScaler appliance.
- Different models have differing numbers of traffic ports. If you require multiple bridges, make sure your model has at least as many as you need.

Selecting a Deployment Mode

Aug 09, 2017

The SD-WAN 4000/5000 appliance can be deployed inline or in a one-arm mode. Inline deployments do not require router reconfiguration; one-arm modes do. SD-WAN 4000/5000 offers internal port bypassing (fail-to-wire) to allow traffic to continue flowing in inline mode if the appliance fails.

Note: Only the one-arm WCCP mode (with a single router) is documented at this time. Inline mode is not yet documented. Citrix recommends WCCP mode at this time.
Different SD-WAN 4000/5000 models offer different numbers of accelerated bridges. Models with multiple accelerated bridges can accelerate multiple inline WAN links. See the specifications sheet for more details, http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/SD-WAN-data-sheet.pdf.

A standalone SD-WAN 4000/5000 appliance can be deployed in either of these two recommended modes:

- Inline, bridged (L2 inline). This closely resembles a standard SD-WAN inline deployment. Packets enter one bridge port and exit the other bridge port.
- One-arm, WCCP. This resembles a standard SD-WAN WCCP deployment.

Citrix also supports the following two modes (which are outside the scope of this document):

- Inline, routed. The NetScaler instance uses routing rules instead of bridging rules to determine how to forward packets.
- Virtual inline. This resembles WCCP, but lacks built-in health-checking.

In L2 inline mode, SD-WAN 4000/5000 is placed between your LAN and your WAN router (or other aggregation point at the LAN-WAN boundary). In a one-arm mode, SD-WAN 4000/5000 is generally connected directly to a dedicated port on your WAN router.

In cases where the WAN router ports are not as fast as the LAN (for example, when the WAN router has gigabit Ethernet, but the LAN has10 gigabit Ethernet), inline mode provides better performance, because its LAN-side traffic is not limited to the speed of the router interface. (Compression allows the LAN-side traffic to be much faster than WAN-bound traffic under favorable conditions.)

Considerations:

- The inline modes require no reconfiguration of your routers, but involves a service disruption when bringing the appliance into service.
- One-arm modes require router reconfiguration but do not require a service disruption.
- Inline mode has higher performance than the other modes.
- One-arm modes are limited to half the speed of the router or switch port they are attached to.
- With WCCP mode, configuring the router to send only a fraction of the WAN traffic to SD-WAN 4000/5000 (as little as the traffic from a single remote site or even a single remote IP address) makes it easy to bring up and test the installation gradually. Inline mode requires that all WAN traffic pass through the appliance.
- WCCP mode requires more configuration of the SD-WAN 4000/5000 appliance than do other modes, but is more standardized and provides more status information on the router.

Recommendation:

- The greater control provided by WCCP, and especially the ability to put the deployment into service in stages, makes WCCP the mode of choice for larger, more complex datacenters, especially if there might be a possibility of overloading

the SD-WAN 4000/5000 appliance.

- Inline mode is convenient for smaller WAN networks and simpler datacenters. It is most commonly used with the SD-WAN 4000/5000 310 and 500, and more rarely with the larger appliances.
- Cascaded installations should use WCCP.

Note: Only WCCP mode (with a single router) is currently documented.

Selecting a Load Balancing Method

Aug 09, 2017

By default, the SD-WAN 4000/5000 Provisioning Wizard sets up load balancing to handle different kinds of connections appropriately. This default behavior is adequate for most installations.

Sending all the connections from the same remote accelerator to the same local accelerator maximizes the benefits of SD-WAN compression, and the default load balancing method accomplishes this. If an instance becomes overloaded or unavailable, new connections are reallocated.

By default, the NetScaler instance uses the least-connection method to balance the load across the accelerators. This method applies whether or not the connections are accelerated. Connections are persistent, but persistency is discontinued for an instance that becomes overloaded, and is lost if the local appliance is restarted or when no traffic from a remote appliance is seen for more than 24 hours.

For incoming accelerated connections (that is, connections with SD-WAN options in the header of the SYN packet), all connections from a given remote SD-WAN are sent to the same local accelerator.

The identity of the remote SD-WAN is determined by one of the SD-WAN SYN options: the "AgentID" field, which contains the management IP address of the remote SD-WAN.

This method is used for connections from remote SD-WAN appliances and remote SD-WAN Plug-ins.

Incoming non-accelerated connections and all outgoing connections are also distributed among the accelerators according to the least-connection method, but since they do not contain an AgentID field, they cannot use AgentID persistence. Instead, they use SRCIPDESTIP persistence, meaning that connections with the same IP addresses use the same accelerator.

If an instance is overloaded, the NetScaler instance bypasses it for new connections, sending them through without acceleration. Existing connections continue to be sent to the instance.

This behavior is controlled by the skipPersistency parameter. The default behavior is -skippersistency ReLB. The alternative behavior, -skippersistency bypass, instructs the NetScaler instance to pass the connection through without sending it to an accelerator.

The default load balancing behavior is adequate for most installations, but sometimes customization is needed. This is most commonly true when a few remote sites have much more traffic than the rest. In that case, it can be worthwhile to assign these large sites to accelerators explicitly.

Optional load balancing behavior includes the use of static routing (for hand-crafted load balancing) and variations on the least-connection with AgentID and SRCIPDESTIP persistence methods used in the default configuration. The behavior for

dealing with overloaded instances can be changed from assigning connections to a difference instance to passing them through as unaccelerated.

Gathering Information Needed for Configuration

Aug 09, 2017

Accurate information about both the local and the remote sites is essential to troubleshooting. Before installing the SD-WAN 4000/5000 appliance, make sure that you have done the following:

1. Obtained or drawn an accurate network diagram of your local site (the one in which you are installing SD-WAN 4000/5000). The local network topology and the capabilities of your WAN routers determine which deployment modes are appropriate for the site.
2. Chosen the deployment mode of the local SD-WAN 4000/5000 appliance (for example, WCCP or inline, with or without HA and cascading).
3. Compiled a list of critical applications that must be tested to validate the deployment.
4. Obtained or drawn an accurate network diagram of your WAN, including both the local and the remote WAN links, their bandwidths in both directions, their subnets, and whether they are accelerated. In deployments with many remote sites, an aggregate of the different categories (accelerated and non-accelerated) is probably sufficient, and only the largest remote sites need to be considered individually.
5. Determined whether there are multiple datacenters with datacenter-to-datacenter traffic, and whether any remote datacenters have a SD-WAN 4000/5000 appliance.
6. Decided whether you plan to increase WAN capacity, the number of sites, or the number of users in the next 24 months. If so, the corresponding SD-WAN 4000/5000 capacity should be installed now.
7. If possible, formed an idea of the traffic breakdown over the WAN, including TCP traffic to and from SD-WAN-accelerated sites, other TCP traffic, ICA users, HDX sessions, and real-time traffic such as VoIP. SD-WAN 4000/5000 needs to be provisioned for the peak loads in terms of accelerated TCP connections, ICA users, and total WAN link capacity.
8. Determined the number of WAN links in the local site. Are they independent, or are they load balanced? If so, are they active-active or active-standby?
9. Determined the current, unaccelerated RTT of the remote sites during peak periods.
10. Identified any QoS devices or proxies in the path between the local and remote sites. QoS devices should be on the WAN side of SD-WAN 4000/5000. Proxies should be on the LAN side.

Initial Configuration

Aug 09, 2017

After checking the connections, you are ready to deploy the SD-WAN 4000 and 5000 appliances on the network.

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

Initial configuration consists of the following tasks:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Assign management IP address through the serial console.

By default, the initial configuration deploys the appliance in inline mode.

Prerequisites

Aug 09, 2017

To deploy a Citrix SD-WAN 4000 or 5000 appliance, you must complete the following prerequisite setup before configuring the appliance.

This document covers release of the SD-WAN software. See the release notes for the recommended versions of the NetScaler software corresponding to the desired release of the SD-WAN software. Never use any versions other than those recommended for SD-WAN 4000 and 5000 appliances.

The number of accelerator appliances depend on the hardware platform and the type of license you apply to the appliance. The following list displays the number of accelerators that gets provisioned automatically by the Configuration Wizard:

- Model 310: Two
- Model 500: Three
- Models 1000 and 1500: Six
- Model 2000: Eight

Before you start provisioning the appliance, Citrix recommends that you have the license file with you, as it is required early in the configuration process To download a license file, complete the procedure described in the My Account All Licensing Tools - User Guide.

After you receive the hardware appliance from Citrix, you need to install it in the network. Complete the following procedures to install the hardware.

To install the SD-WAN 4000/5000 appliance hardware, follow the installation procedure at Installing the Hardware.

Plug in the provided loopback cable available with the appliance into ports 10/3 and 10/4 to create a loopback adapter. The loopback adapter is used for the communication between the NetScaler instance and the accelerators.

Note: If you do not want to use the loopback cable, navigate to the System > Configuration > System page and click the **Configure Loopback Settings** link in the **System Settings** section. Select the **Eliminate Loopback Cable** option, and then click **OK**.

Deployment Worksheet

Aug 09, 2017

Note: Use this worksheet only when provisioning a factory-reset appliance with the release 7.1 configuration wizard. If you are simply upgrading a previously configured system to release 7.1, your appliance will retain its previous configuration, which will be different.

The appliance uses at least two ports: the management port (typically 0/1) and the traffic port (such as 10/1). Inline mode uses traffic ports in pairs, such as ports 10/1 and 10/2. Ports must be selected in advance, because the configuration depends on their identity.

The appliance uses three subnets directly: the management subnet, the external traffic subnet, and the internal traffic subnet. Multiple IP addresses are used on each subnet. Each subnet must be specified along with the correct subnet mask.

The following figure is a worksheet for these parameters. It supports inline and WCCP modes, with and without HA. The table below the figure describes what each entry means.

Figure 1. Deployment worksheet

Table 1. Deployment Worksheet Parameters

	Parameter	Example	Your Value	Description
Management Subnet				
M2.	Gateway IP address	10.199.79.254		Default gateway serving the management subnet.
M3.	Subnet Mask	255.255.255.128		Subnet mask for the management subnet.
M4.	Xen Hypervisor IP address	10.199.79.225		IP address of Xen Hypervisor.
M5.	Service VM IP address	10.199.79.226		IP address of Management Service VM, which controls configuration.
M6.	Accelerator UI	10.199.79.227		Accelerator GUI, also called the Broker UI, which manages the instances as a unit.
M7.	NetScaler Management IP address	10.199.79.245		IP address of the NetScaler instance's GUI and CLI interfaces.
External Traffic Subnet				
T1.	Router IP address	172.17.17.1		IP address of router on external traffic subnet.

	Parameter	Example	Your Value	Description
T2.	Subnet Mask	255.255.255.0		Subnet mask of external traffic subnet.
T3.	NetScaler IP address	172.17.17.2		NetScaler IP address on external traffic subnet.
T4.	External Signaling IP address	172.17.17.10		Traffic to this IP address is load-balanced between the signaling IP addresses of the accelerators.
T5.	External WCCP IP address #1	172.17.17.11		Maps through NAT to WCCP VIP on accelerator #1.
T6.	External WCCP IP address #2	172.17.17.12		Maps through NAT to WCCP VIP on accelerator #2.
T7.	Local LAN Subnets	10.200.0.0/16		The local LAN subnet to be accelerated. This is the only subnet that will receive acceleration.
T8.	GRE Router HostID	NA		WCCP-GRE only. Host ID of GRE router.
T9.	Traffic Port	10/1		Port used for accelerated traffic.
T10+.	(Inline) Additional Traffic Port			Other traffic port in pair.
T11, T12	(WCCP) Service Groups: TCP, UDP	71, 72		Service groups used by accelerator #1 for WCCP. First is for TCP traffic, second is for UDP.
T13, T14	(Not used)			
T15, T16	(Inline) Ports used by link #2	10/5, 10/6		If multiple links are used with inline mode, these ports are used for link #2.
T17, T18	(Inline) Ports used by link #3	10/7, 10/8		If multiple links are used with inline mode, these ports are used for link #3.
VLAN1.1, VLAN1.2, VLAN1.3, VLAN1.4	External VLANs for Bridge #1	412		When VLAN trunking is used, these are tagged VLANs crossing bridge #1.
VLAN2.1, VLAN2.2, VLAN2.3, VLAN2.4				When VLAN trunking is used, these are tagged VLANs crossing bridge #2.
VLAN3.1, VLAN3.2, VLAN3.3, VLAN3.4	External VLANs for Bridge #1			When VLAN trunking is used, these are tagged VLANs crossing bridge #3.

Configuring the Appliance

Aug 09, 2017

Before you start configuring the appliance, you must change the IP address of the management service to the one in your management network, so that you can access the appliance over the network. You can change the management IP address by connecting a computer to the appliance through either the Ethernet port or the serial console.

Assigning a Management IP Address through the Ethernet Port

Aug 09, 2017

Use the following procedure for initial configuration of every SD-WAN 1000 or 2000 appliance with Windows Server. The procedure accomplishes the following tasks:

- Configure the appliance for use on your site.
- Install the Citrix license.
- Enable acceleration.
- Enable traffic shaping (inline mode only).

With inline deployments, this configuration might be all you need, because most acceleration features are enabled by default and require no additional configuration.

If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 15 of the following procedure.

Note: You must have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.50, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields blank.

2. Using an Ethernet cable, connect this computer to the port labeled PRI on the SD-WAN appliance.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address, which is http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance:
   Username: nsroot

   Password: nsroot.

5. Start the configuration wizard by clicking **Get Started**.
6. On the **Platform Configuration** page, enter respective values from your worksheet, as shown in the following example:

7. Click **Done**. A screen showing the Installation in Progress... message appears. This process takes approximately 2 to 5 minutes, depending on your network speed.
8. A Redirecting to new management IP message appears.
9. Click **OK**.
10. Unplug your computer from the Ethernet port and connect the port to your management network.
11. Reset the IP address of your computer to its previous setting.
12. From a computer on the management network, log on to the appliance by entering the new management service IP address, such as https://<Managemnt_IP_Address>, in a web browser.

13. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
14. Log on to the appliance by using the **nsroot** user name and the password from your worksheet.
15. To complete the configuration process, see Provisioning the Appliance.

Assigning a Management IP Address through the Serial Port

Aug 09, 2017

If you do not want to change the settings of your computer, you can configure the appliance by connecting it to your computer with a serial null modem cable. You must have physical access to the appliance.

To configure the appliance through the serial console

1. Connect a serial null modem cable to the appliance's console port.
2. Connect the other end of the cable to the serial COM port of a computer running a terminal emulator, such as Microsoft HyperTerminal, with settings 9600,N,8,1, p.
3. In the HyperTerminal output, press **Enter**. The terminal screen displays the logon prompt.
   Note: You might have to press **Enter** two or three times, depending on the terminal program you are using.
4. At the logon prompt, log on to the appliance with the following default credentials:
   Username: nsroot

   Password: nsroot.

5. At the **$** prompt, run the following command to switch to the shell prompt of the appliance:
   $ ssh 169.254.0.10

6. Enter **Yes** to continue connecting to the management service.
7. Log on to the shell prompt of the appliance with the following default credentials:
   Password: nsroot.

8. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu:
   # networkconfig

9. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.
10. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the XenServer.
11. Type **3** and press **Enter** to select option 3, and specify the network mask for the IP addresses.
12. Type **4** and press **Enter** to select option 4, and specify the default gateway for the management service IP address.
13. Type **8** and press **Enter** to save the settings and exit.
14. Access the SD-WAN appliance by entering the new management service IP address of the appliance, such as https://<Management_Service_IP_Address>, in a web browser of a computer on the management network.
15. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
16. To complete the configuration process, see Provisioning the Appliance.

Provisioning the Appliance

Aug 09, 2017

After assigning an IP address to the management service, you are ready to provision the NetScaler and accelerator instances. As soon as you log on to the appliance, the configuration wizard appears.

When using the configuration wizard, keep the following points in mind:

- The following procedure assumes that you have already filled out the configuration worksheet.
- If you change the IP addresses of the Management Network, or change the default gateway to an address not on the Management Network, you lose connectivity to the appliance unless you are on the same Ethernet segment as the management port.
- When using the configuration wizard, check your entries carefully. The wizard has no Back button. If you need to modify the previous screen, use the Back button on your browser. This takes you to the logon page, then to the previous screen.
- The configuration wizard is displayed only when you log on to the appliance for the first time to configure the appliance. After you finish configuring the appliance, this wizard becomes inaccessible, and will reappear only after a factory reset. Check your entries carefully.

This wizard walks you through a fresh configuration of the appliance.

Note: If you receive a #SESS_CORRUPTED error at any time during these procedures, click **Logout**, clear your browser cache, close your browser, and open it again.

To configure the appliance by using the configuration wizard:

1. On the Welcome page, click Get Started.
   Note: All pages after the Get Started page have a heading that says, “Deployment Mode: Inline/L2 Mode,” but this wizard is used for all deployment modes.
2. Follow these steps to configure a fully 7.3-compliant system:
   - Acquire the following release 7.3 software distributions from the release 7.3 downloads page on My Citrix:
     - Management service (as a .tgz file)
     - NetScaler VM (as an .xva file)
     - Accelerator VM (as an .xva file)
     - Upgrade bundle (as a .upg) file
   - Navigate to the System > Configuration > Management Service > Software Images page, and then select **Upload** from the Action list.
   - Upload a release 7.3 Management Service image (distributed as a .tgz file).
   - Navigate to the System > Configuration > NetScaler > Software Images page, and then upload a release 7.3 NetScaler XVA image.
   - Navigate to the System > Configuration > SD-WAN > Software Images page, and then upload the accelerator XVA image.
   - Navigate to the System > Configuration > Management Service page, and then click the **Upgrade Management Service** link.
   - Select the management service image that you recently uploaded and click **OK**.
   - When the lower left-hand corner of the screen displays “Management Service Updated Successfully,” log off and clear your browser cache. Log on after the management service restarts (a few minutes).
   - On the Welcome screen, click **Get Started**.
3. For Management Access Settings, specify values for the various fields according to the network settings. The following

screen shot displays sample values used in this documentation. Enter values as follows:

- **XenServer IP Address**—(Item M4 on your worksheet, or H4 if this is the second appliance in an HA pair.) The management address of the built-in XenServer hypervisor. This must be a valid address on the management network.
- **Management Service IP Address**—(Item M5 on your worksheet, or H5 if this is the second appliance in an HA pair). The address of the Management Service VM that you use to perform most system management tasks. This must be a valid address on the management network.
- **Netmask**—(Item M3 on your worksheet). The subnet mask of the management network.
- **Gateway**— (Item M2 on your worksheet). The default gateway for the management network.
- **DNS Server**—The IP address of the DNS server. This is a mandatory parameter.
- **NTP Server**—IP or FQDN address of your time server. This will be used by all the virtual machines in the appliance. Note that if you use advanced CIFS or MAPI acceleration, the system time of the appliance must be close to that of the Windows domain server, so choose an NTP server that maintains a close relationship to the time on your Windows domain server.
  Note: Unless the NTP server is specified as an IP address, it is not used by the accelerator.
- **Time Zone**—Select your time zone from the pull-down menu.
- **Change Password**—Select this check box and type in a new nsroot password, two times, to change the password. This same password is used on the management service and the NetScaler instance for account nsroot, and on the accelerator for the admin account. If the password is not changed, it remains set to nsroot (the default).

Figure 1. Sample Values for the Fields in Management Access Settings Page of the Configuration

4. Check your settings and click **Continue**.
5. In the Manage Licenses section, see if an appropriate license is already listed in the Name field. If so, select it and skip to step 8.
6. Click **Upload** in the Update Licenses section.
7. Navigate to the folder that contains the license file and open the file.
8. Click **Add License** and upload the license file provided by Citrix. The license is added to the appliance, as shown in the following figure.
   Figure 2. Sample License Added to the Appliance on the Manage License Files Page of the Configuration Wizard

   You can also get a license file from the Citrix.com website by clicking the **here** link and using your My Citrix credentials.

9. Select the license in the Name field and click **Continue**. The SD-WAN Setup page appears. Fill in the fields as follows:
   1. **Network Settings**—This section informs the accelerators of the management network.
      - **SD-WAN Accelerator IP Address**—Enter the value of M6 from your worksheet. This is the IP address of the accelerator
      - **NetScaler IP Address**—Enter the value of M7 from your worksheet. This is the IP address of the NetScaler GUI.
      - **Use System Netmask and Gateway**—Select this option if you want to use the network mask and gateway IP addresses you had specified in the Platform Configuration page.
      - **Netmask**—Enter the value of M3 from your worksheet. This is the subnet mask (netmask) of the management network (note that you have entered this already, on a previous page).
      - **Gateway**—Enter the value of M2 from your worksheet again.
      - **Signaling IP Address**—Enter the value of T4 from your worksheet. This is the external signaling IP address of the accelerator, used by SD-WAN Plug-ins to connect to the appliance.

- **Signaling Netmask**—Enter the value of T2 from your worksheet. This is the subnet mask (netmask) of the external traffic network.

2. **XVA Files**—This section allows you to specify previously uploaded XVA files (Xen virtual machines) for the NetScaler and accelerator instances. Select the XVA images that you uploaded as part of step 2.
   Figure 3. SD-WAN Setup Page

10. Click **Continue**. The wizard starts provisioning the required instances, as shown in the following figure.
    Figure 4. Provisioning Progress Indicator

11. After the instances are provisioned, add one of your local LAN subnets to the Link Configuration section from list T7 in your worksheet, as shown in the following figure. This subnet will be added as a local LAN subnet in the accelerator. If you have more than one LAN subnet, you can add them to the LAN link definition in the Accelerator GUI after the configuration wizard completes. Click Add to add the subnet.
    Figure 5. Link Configuration Is at the Bottom of This Page

12. Log off, and then log back on. If you see a "Version Incompatibility Detected" message, install the upgrade bundle you downloaded in Step 2, using the procedure in Managing the Appliance by using the Management Service.

Basic configuration is complete. Next, perform deployment-mode-specific configuration (such as for WCCP mode).

Note: After the wizard completes, the appliance is configured for the basic setup. To configure the appliance for a specific deployment scenario, see Deployment Modes.

Deployment Modes

Aug 09, 2017

SD-WAN 4000/5000 appliances have two recommended deployment modes: WCCP and inline. These modes are commonly used without high availability (HA), and less commonly with HA.

At this time, Citrix recommends WCCP mode, with a single router and without HA, for most deployments. Use inline mode when WCCP is not available.

Although not all of the following modes are recommended at this time, they are all supported:

- WCCP mode with a single router
- WCCP mode with a single router and high availability
- Cascade of two or more appliances in WCCP mode along with a NetScaler MPX Appliance
- Cascade of two or more appliances in WCCP mode along with a NetScaler MPX Appliance in HA
- Inline mode
- Inline mode in HA
- Virtual inline mode
- Virtual inline mode in HA

Note: While modes other than WCCP and inline are supported, they are incompletely documented and are not recommended for typical installations. Please contact your Citrix representative when considering one of these modes.

Virtual Inline Mode

Aug 09, 2017

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router.

Virtual inline mode provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links.

Note: Citrix recommends that you do not deploy SD-WAN appliances in virtual inline mode with routers that do not support health monitoring.

The tasks for configuring virtual inline mode are performed on the router. On the SD-WAN appliance, just verify that the software version supports virtual inline mode and provision the instances with the necessary IP addresses. Do not change the default forwarding method.

Validated Designs

Aug 09, 2017

The physical mode for virtual inline deployment of a SD-WAN appliance is one-arm mode. In a one-arm topology with multiple subnets, the branches, local Repeater instances, and servers can exist on different subnets. For example, you can deploy the appliance in one-arm mode for managing the Repeater instances, with the NetScaler and local Repeater instances connected by the internal private subnet. A NetScaler-owned subnet IP address (SNIP) is used to communicate with an accelerator instance. You must enable MAC Based Forwarding (MBF) Use Subnet IP address (USNIP), and Return To Ethernet Sender options on the NetScaler instance.

This section contains Citrix validated deployment topologies for virtual inline mode.

Single Router and single link

Aug 09, 2017

A SD-WAN appliance deployed in virtual inline mode with one router and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with one router and one link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router

Two routers and a single link

Aug 09, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router (both routers individually)

Two routers and two links

Aug 09, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers and two links.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure Layer 4 Parameters (only if you expect connection migration between routers)
- Configuring VLANs for Connection Migration
- Configure a Router(both routers individually)

Two routers in High Availability setup

Aug 09, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers in high availability setup and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router (both routers individually)
- Configure Routers in High Availability Setup

Deployment Worksheet

Aug 09, 2017

The appliance uses at least two ports: the management port (typically 0/1) and the traffic port (such as 10/1) . You must select ports in advance, because the configuration depends on the identity of the ports.

The appliance uses three subnets directly: the management subnet, the external traffic subnet, and the internal traffic subnet. Multiple IP addresses are used on each subnet. Each subnet must be specified along with the correct subnet mask.

The following figure is a worksheet for these parameters. It supports inline and virtual inline modes, with and without HA. The table below the figure describes what each entry means.

Deployment Worksheet Parameters

	Parameter	Example	Your Value	Description
Management Subnet				
M2.	Gateway IP address	10.199.79.254		Default gateway serving the management subnet.
M3.	Subnet Mask	255.255.255.128		Subnet mask for the management subnet.
M4.	Xen Hypervisor IP address	10.199.79.225		IP address of Xen Hypervisor.
M5.	Service VM IP address	10.199.79.226		IP address of Management Service VM, which controls configuration.
M6.	Accelerator UI	10.199.79.227		Accelerator GUI, also called the Broker UI, which manages the instances as a unit.
M7.	NetScaler Management IP address	10.199.79.245		IP address of the NetScaler instance's GUI and CLI interfaces.
External Traffic Subnet				
T1.	Router IP address	172.17.17.1		IP address of the first router on external traffic subnet.
T2.	Router IP address	172.17.18.1		IP address of the second router on external traffic subnet.
T3.	Subnet Mask	255.255.255.0		Subnet mask of external traffic subnet.
T4.	Subnet IP address	172.17.17.2		Subnet IP address for NetScaler on external traffic subnet.
T5.	Subnet IP address	172.17.18.2		Subnet IP address for NetScaler on external traffic subnet.

	Parameter	Example	Your Value	Description
T6.	External Signaling IP address	172.17.17.10		Traffic to this IP address is load-balanced between the signaling IP addresses of the accelerators.
T7.	Local LAN Subnets	10.200.0.0/16		The local LAN subnet to be accelerated. This is the only subnet that will receive acceleration.
T8.	Traffic Port	10/1		Port used for accelerated traffic.
T9.	Traffic Port	10/6		Port used for accelerated traffic.

Note: Ports 10/3 and 10/4 are reserved for loopback cable. Do not configure these ports as traffic ports.

Configuring the NetScaler Instance

Aug 09, 2017

By default, the Getting Started wizard configures the appliance in inline mode. To deploy the SD-WAN appliance in virtual inline mode, you must configure the NetScaler instance of the appliance to support this mode.

To configure the NetScaler instance for the virtual inline mode of the appliance:

- Enable L3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configuring the Instance for Connection Migration

Enable Layer 3 Mode

Aug 09, 2017

In the inline mode, the default deployment mode of the appliance, the appliance uses layer 2 bridging. However, to deploy the appliance in virtual inline mode, you must enable layer 3 mode and disable layer 2 mode on the appliance.

To enable layer 3 mode and disable layer 2 mode by using the configuration utility

1. Access the NetScaler instance by clicking the NetScaler instance's IP address on the Configuration > Instances > NetScaler page. You are logged on to the NetScaler instance automatically.
2. Navigate to the System > Settings page.
3. Click the Configure modes link.

4. In the Configure Modes dialog box, select Layer 3 Mode (IP Forwarding) option.
5. Clear the Layer 2 Mode option, as shown in following figure.

6. Click OK.

To enable layer 3 mode and disable layer 2 mode by using the command line interface, run the following commands

```
> enable ns mode L3
> disable ns mode l2
```

Enable the Return to Ethernet Sender Mode

Aug 09, 2017

This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

To enable the Return to Ethernet Sender mode by using the configuration utility

1. Navigate to the System > Network page.
2. Click the Configure Layer 2 Parameter link, as shown in the following figure.

3. In the Configure Layer 2 Parameter dialog box, select the Return to Ethernet Sender option, as shown in the following figure.

4. Click OK.

To enable the Return to Ethernet Sender mode by using the command line interface, run the following command

```
> set L2Param -returnToEthernetSender ENABLED
```

Add a Subnet IP Address

Aug 09, 2017

You must add a Subnet IP (SNIP) address to the NetScaler instance. You assign this IP address to the NetScaler instance on the external traffic subnet. The NetScaler instance uses this address to communicate with the router.

To add a subnet IP address to the instance by using the configuration utility

1. Navigate to the System > Network > IPs page.
2. Click Add, as shown in the following figure.

3. In the Create IP dialog box, specify a subnet IP address in the IP Address field if the address has not already been created by the configuration wizard. (If it has already been created, it is listed on the IPs page). The subnet IP address declares the external traffic network. In the IP Address field, type the NetScaler traffic IP address (entry T4 in your worksheet).
4. In the Netmask field, specify the network mask (entry T3 in your worksheet).
5. From the IP Type group, make sure that the Subnet IP option is selected, as shown in the following figure.

6. Click Create and then Close.
7. If you are configuring the appliance for two links, repeat this procedure to create another subnet IP address (entry T5 of your worksheet).

To add a Subnet IP address to the instance by using the command line interface, run the following command

```
> add ns ip 172.17.17.2 255.255.255.0
> add ns ip 172.17.18.2 255.255.255.0
```

Note: Run the second command only if you are configuring the appliance for two links.

Bind the Subnet IP Address to VLAN of Data Interface

Aug 09, 2017

After adding a subnet IP address to the instance, you must bind the IP address to the VLAN of data interface you are using on the instance. By default, the VLAN 1007 is bound to the 10/1 and 10/2 interfaces.

To bind the subnet IP address to VLAN of data interface from the configuration utility

1. Navigate to the System > Network > VLANs page.
2. Select the VLAN, such as 1007, that was created when provisioning the instances.
3. Click Open, as shown in the following figure.

4. In the Interface Bindings tab of the Configure VLAN dialog box, notice that the VLAN is bound to interfaces 10/1 (entry T9 of your worksheet) and 10/2.

5. In the IP Bindings tab, select the subnet IP address you have created, as shown in the following figure.

6. Click Create and then Close.
7. If you are configuring the appliance for two links, repeat this procedure to bind second subnet IP (entry T5 of your worksheet) to another VLAN, such as 1009.

To bind the subnet IP addresses to VLANs of data interface by using the command line interface, run the following command

```
> bind vlan 1007–IPAddress 172.17.17.2 255.255.255.0
> bind vlan 1009–IPAddress 172.17.18.2 255.255.255.0
```

Note: Run the second command only if you are configuring the appliance for two links.

Configuring the Instance for Connection Migration

Aug 09, 2017

In your network setup, if you expect connection migration between the routers, then you must configure the layer 4 parameters and make changes to the VLAN configurations, appropriately.

Configure Layer 4 Parameters

If you have two routers and expect connection migration between the routers, you must configure layer 4 parameters on the NetScaler instance.

To configure layer 4 parameters from the configuration utility

1. Navigate to the System > Network page.
2. Click the Configure Layer 2 Parameter link, as shown in the following figure.

3. In the Configure Layer 4 Parameter dialog box, select the Vlan from the L2 Connection Method list, as shown in the following figure.

4. Click OK.

To bind the subnet IP address to VLAN of data interface by using the command line interface, run the following command

```
> set L4Param –l2connMethod vlan
```

Configuring VLANs for Connection Migration

If the network has two links and connected to interfaces that are bound to different VLANs, then you must bind the interfaces to the same VLAN to enable connection migration. Additionally, you must bind both subnet IP addresses to the same VLAN. Skip this procedure if you are using interfaces that are bound to the same VLAN.

To configure VLANs for connection migration by using the configuration utility

1. Navigate to the System > Network > VLANs page.
2. Select the VLAN, such as 1009, to which interface and subnet IP address of the second link is bound.
3. Click Open.
4. In the Interface Bindings tab of the Configure VLAN dialog box, clear the interface you are using (entry T9 of your worksheet), as shown in the following screen shot.

5. In the IP Bindings tab, clear the entry for subnet IP address bound to the VLAN.
6. Click OK.
7. Open the VLAN, such as 1007, to which you want to bind the data interface.
8. In the interface Binding tab, select the data interface (entry T9 in your worksheet)
9. In the IP Bindings tab, select the subnet IP address (entry T5 in your worksheet) you just unbound from the other VLAN.
10. Click OK.

To configure VLANs for connection migration by using the command line interface, run the following commands

```
> unbind vlan 1009 –ifnum 10/6
> unbind vlan 1009 –IPAddress 172.17.18.2 255.255.255.0
```

```
> bind vlan 1007 –ifnum 10/6
> bind vlan 1007 –IPAddress 172.17.18.2 255.255.255.0
```

Configuring a Router

Aug 09, 2017

To support virtual inline mode, a router must forward incoming as well as outgoing WAN traffic to the SD-WAN appliance. After the appliance processes the traffic, the router must forward the incoming traffic from the appliance to the LAN and the outgoing traffic from the appliance to the WAN. You have to configure policy based rules to avoid routing loops. In addition, the router must monitor the health of the appliance so that the appliance can be bypassed if it fails .

If the router supports the Reverse Path Forwarding feature, you must disable the feature on the interfaces with policies to redirect traffic to a SD-WAN appliance, including the interface that is connected to the appliance. Otherwise, the router intermittently drops traffic. By default, the Reverse Path Forwarding feature is enabled on the router.

Note: If the network has two routers, configure the following procedures for each router using appropriate IP addresses identified in the worksheet.

In the virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the Ethernet ports of the router to the appliance and creating routing rules based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

Traffic shaping is not effective unless all WAN traffic passes through the appliance. Following is the basic routing algorithm:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward the packet to the appliance.
- If the packet is destined for the WAN, forward the packet to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does not perform health monitoring. To enable health monitoring, define a rule to monitor the availability of the appliance, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Note: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy based routing do not support health checking. The health-monitoring feature is relatively new. It was first available in Cisco IOS release 12.3(4)T.
Following is an example of a rule for monitoring the availability of the appliance by using Cisco Router model 7600 with IOS Software version:

```
!- Use a ping (ICMP echo) to see if appliance is in the connected track
123 rtr 1 reachability
!
rtr 1
type echo protocol IpIcmpecho 17.17.17.2    schedule 1 life forever start-time now
```

This rule pings the appliance at 17.17.17.2 periodically. You can test against 123 to see if the unit is up.

Following is an example of configuring a Cisco router for virtual inline mode:

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.200.51.0   255.255.255.0
ip policy route-map server_side_map
!
interface FastEthernet0/1
ip address 17.17.17.1   255.255.255.0!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
ip policy route-map wan_side_map
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended server_side
permit ip 10.100.51.0 0.0.0.255 10.20.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.20.20.0 0.0.0.255 10.100.51.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 17.17.17.1   20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 17.17.17.1   10 track 123
```

This example applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP address of 10.100.51.0/24 and destination IP address of 10.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended server_side
 permit tcp 10.200.51.0   0.0.0.255 10.20.20.0 0.0.0.255
```



```
ip access-list extended wan_side
 permit tcp 10.20.20.0 0.0.0.255 10.200.51.0 0.0.0.255
!
```

To configure high availability between routers, see the router-specific high availably configuration manual.

Verifying the Virtual Inline Mode

Aug 09, 2017

By default, the HTTP and CIFS acceleration is enabled on the appliance. After you have successfully configured the appliance in the virtual inline deployment mode, the appliance starts accelerating these connections.

Note: To configure application-specific connections, you must configure service classes for the respective applications. To configure a service class, see http://support.citrix.com/proddocs/topic/SD-WAN-72/cb-traffic-classification-con.html.

To verify that you have successfully configure the appliance in the virtual inline mode

1. Send network traffic through the appliance.
2. Navigate to the SD-WAN > Monitoring > Optimization > Connections page.
3. Verify that the Accelerated Connections tab displays entries for the accelerated connections, as shown in the following screen shot. This tab displays an entry each for all accelerated connections.

Figure 1. The Accelerated Connections tab

WCCP Mode

Aug 09, 2017

Web Cache Communication Protocol (WCCP) is a dynamic routing protocol introduced by Cisco. Originally intended only for web caching, WCCP version 2 became a more general-purpose protocol, suitable for use by accelerators such as Citrix SD-WAN appliances.

WCCP mode is the simplest way of installing a SD-WAN appliance when inline operation is impractical. It is also useful where asymmetric routing occurs, that is, when packets from the same connection arrive over different WAN links. In WCCP mode, the routers use the WCCP 2.0 protocol to divert traffic through the appliance. Once received by the appliance, the traffic is treated by the acceleration engine and traffic shaper as if it were received in inline mode.

Note:

- For the purposes of this discussion, WCCP version 1 is considered obsolete and only WCCP version 2 is presented.
- The standard WCCP documentation calls WCCP clients "caches." To avoid confusion with actual caches, Citrix generally avoids calling a WCCP client a "cache." Instead, WCCP clients are typically called "appliances."
- This discussion uses the term "router" to indicate WCCP-capable routers and WCCP-capable switches. Though the term "router" is used here, some high-end switches also support WCCP, and can be used with SD-WAN appliances.

The SD-WAN appliances support two WCCP modes:

- WCCP is the original SD-WAN WCCP offering supported since release 3.x. It supports a single appliance service group (no clustering).
- WCCP clustering, introduced in release 7.2, allows your router to load-balance traffic between multiple appliances.

The physical mode for WCCP deployment of a SD-WAN appliance is one-arm mode in which the SD-WAN appliance is connected directly to a dedicated port on the WAN router. The WCCP standard includes a protocol negotiation in which the appliance registers itself with the router, and the two negotiate the use of features they support in common. Once this negotiation is successful, traffic is routed between the router and the appliance according to the WCCP router and redirection rules defined on the router.

A WCCP-mode appliance requires only a single Ethernet port. The appliance should either be deployed on a dedicated router port (or WCCP-capable switch port) or isolated from other traffic through a VLAN. Do not mix inline and WCCP modes.

The following figure shows how a router is configured to intercept traffic on selected interfaces and forward it to the WCCP-enabled appliance. Whenever the WCCP-enabled appliance is not available, the traffic is not intercepted, and is forwarded normally.

Figure 1. WCCP Traffic Flow

WCCP allows traffic to be forwarded between the router and the appliance in either of the following modes:

- L2 Mode—Requires that the router and appliance be on the same L2 segment (typically an Ethernet segment). The IP packet is unmodified, and only the L2 addressing is altered to forward the packet. In many devices, L2 forwarding is

performed at the hardware layer, giving it the maximum performance. Because of its performance advantage, L2 forwarding is the preferred mode, but not all WCCP-capable devices support it.

- GRE Mode—Generic Routing Encapsulation (GRE) is a routed protocol and the appliance can in theory be placed anywhere, but for performance it should be placed close to the router, on a fast, uncongested path that traverses as few switches and routers as possible. GRE is the original WCCP mode. A GRE header is created and the data packet is appended to it. The receiving device removes the GRE header. With encapsulation, the appliance can be on a subnet that is not directly attached to the router. However, both the encapsulation process and the subsequent routing add CPU overhead to the router, and the addition of the 28-byte GRE header can lead to packet fragmentation, which adds additional overhead.

WCCP mode supports multiple routers and both GRE vs. L2 forwarding. Each router can have multiple WAN links. Each link can have its own WCCP service group.

Traffic shaping is not effective unless the appliance manages UDP traffic as well as TCP traffic. A second service group, with a UDP service group for each WAN link, is recommended if traffic shaping is desired.

A WCCP client (an appliance) uses UDP port 2048 to register itself with the router and to negotiate which traffic should be sent to it, and also which WCCP features should be used for this traffic. The appliance operates on this traffic and forwards the resulting traffic to the original endpoint. The status of an appliance is tracked through the WCCP registration process and a heartbeat protocol. The appliance first contacts the router over the WCCP control channel (UDP port 2048), and the appliance and router exchange information with packets named “Here_I_Am” and “I_See_You,” respectively. By default, this process is repeated every ten seconds. If the router fails to receive a message from the appliance for three of these intervals, it considers the appliance to have failed and stops forwarding traffic to it until contact is reestablished.

Different appliances using the same router can provide different services. To keep track of which services are assigned to which appliances, the WCCP protocol uses a service group identifier, a one-byte integer. When an appliance registers itself with a router, it includes service group numbers as well.

- A single appliance can support more than one service group.
- A single router can support more than one service group.
- A single appliance can use the same service group with more than one router.
- A single router can use the same service group with more than one appliance. For SD-WAN appliances, multiple appliances are supported in WCCP cluster mode, and a single appliance is supported in WCCP mode.
- Each appliance specifies a “return type” (L2 or GRE) independently for each direction and each service group. SD-WAN 4000/5000 appliances always specify the same return type for both directions. Other SD-WAN appliances allow the return type to be different.

Figure 2. Using different WCCP service groups for different services

Multiple service groups can be used with WCCP on the same appliance. For example, the appliance can receive service-group 51 traffic from one WAN link and service-group 62 traffic from another WAN link. The appliance also supports multiple routers. It is indifferent to whether all the routers use the same service group or different routers use different service groups.

Service Group Tracking. If a packet arrives on one service group, output packets for the same connection are sent on

the same service group. If packets arrive for the same connection on multiple service groups, output packets track the most recently seen service group for that connection.

When WCCP is used with high-availability mode, the primary appliance sends its own apA or apB management IP address, not the virtual address of the HA pair, when it contacts the router. If failover occurs, the new primary appliance contacts the router automatically, reestablishing the WCCP channel. In most cases the WCCP timeout period and the HA failover time overlap. As a result, the network outage is less than the sum of the two delays.

Standard WCCP allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive. WCCP clustering allows multiple appliances per service group.

The following figure shows a simple WCCP deployment, suitable for either L2 or GRE. The traffic port (1/1) is connected directly to a dedicated router port (Gig 4/12).
Figure 3. Simple WCCP deployment

In this example, the SD-WAN 4000/5000 is deployed in one-arm mode, with the traffic port (1/1) and the management port (0/1) each connecting to its own dedicated router port.

On the router, WCCP is configured with identical ip wccp redirect in statements on the WAN and LAN ports. Two service groups are used, 71 and 72. Service group 71 is used for TCP traffic and service group 72 is used for UDP traffic. The SD-WAN appliance does not accelerate UDP traffic, but can apply traffic shaping policies to it.

Note: The WCCP specification does not allow protocols other than TCP and UDP to be forwarded, so protocols such as ICMP and GRE always bypass the appliance.

SD-WAN release 7.2 or later supports WCCP clustering, which enables your router to load-balance your traffic between multiple appliances. For more information about deploying SD-WAN appliances as a cluster, see WCCP Clustering.

For more information about WCCP, see Web Cache Communication Protocol V2, Revision 1, http://tools.ietf.org/html/draft-mclaggan-wccp-v2rev1-00.

WCCP Mode (Non-Clustered)

Aug 09, 2017

WCCP mode allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive.

Note: WCCP clustering allows multiple appliances per service group.

Following are limitations and best practices for (non-clustered) WCCP mode:

- On appliances with more than one accelerated pair, all the traffic for a given WCCP service group must arrive on the same accelerated pair.
- Do not mix inline and WCCP traffic on the same appliance. The appliance does not enforce this guideline, but violating it can cause difficulties with acceleration. (WCCP and virtual inline modes can be mixed, but only if the WCCP and virtual inline traffic are coming from different routers.)
- For sites with a single WAN router, use WCCP whenever inline mode is not practical.
- Only one appliance is supported per service group. If more than one appliance attempts to connect to the same router with the same service group, the negotiation will succeed only for the first appliance.
- For sites with multiple WAN routers serviced by the same appliance, WCCP can be used to support one, some, or all of your WAN routers. Other routers can use virtual inline mode.

Configuring WCCP

Aug 09, 2017

A WCCP deployment follows the same initial steps as an inline deployment, but has additional steps beyond the basic inline procedure.

Perform the following tasks if you have not done so already:

- Install the SD-WAN hardware. See Installing the Hardware.
- Fill out the Deployment Worksheet and perform the initial configuration. See .
- If you are using high-availability mode, see the Configuring the High Availability Setup on the Appliances section before proceeding.

The following high-level procedure summarizes the WCCP installation process, which works for both GRE and L2 forwarding and for any number of routers and links.

To configure WCCP mode

Note: You must follow the hyperlinks and follow the detailed instructions for each step.

1. Configuring the Router.
    1. Enable WCCP globally.
    2. Disable reverse path forwarding if your router supports it.
    3. Configure WCCP service groups.
2. Configuring Accelerators for WCCP Negotiation.
    1. Enable WCCP.
    2. For each service group:
        1. Create a service group definition on the SD-WAN appliance.
        2. Verify that this service group establishes WCCP communication with its associated routers.
3. Verifying WCCP Mode.
4. If using high-availability mode, configure and test the second appliance, then complete the Configuring the High Availability Setup on the Appliances procedure.

Note: This information is for WCCP mode. For WCCP Clustering, see the SD-WAN WCCP Clustering topics, especially the Configuring the Router subtopic.

For each WCCP router:

1. Enable WCCP in global router configuration.
2. For each WCCP service group on the worksheet for this router, declare ip wccp <sg> in global router configuration
3. Referring to the configuration examples below, for each WAN interface on this router:
    1. You can either use Method A or Method B to configure the router:
        - Method A: On the WAN interface only, declare ip wccp <sg> redirect in and ip wccp <sg> redirect out for the two service groups associated with the WAN interface.
        - Method B: If there is only one WAN interface, you can alternatively declare ip wccp <sg> redirect in on the WAN interface and on every LAN interface except the appliance's traffic interface.
    2. If your router supports reverse path forwarding, disable it on this interface by changing any ip verify unicast reverse-

path commands to no ip verify unicast reverse-path commands on each interface that has an ip wccp redirect command.

4. Save the configuration.

Router Configuration Examples

For normal operation, you must declare WCCP version 2 and the WCCP group ID for the router as a whole, and then enable redirection on each WAN interface.

Either one or two WCCP service groups can be used, but two are recommended, so that both TCP and UDP can be redirected, allowing more accurate traffic shaping. The WCCP standard requires that TCP and UDP traffic use different service groups.

Method A is required if the router has multiple WAN interfaces.

Example

Following is an example of configuring a Cisco IOS router:

```
! This example is for WCCP mode, not WCCP clustering
! (which is covered elsewhere)
config term
ip wccp version 2
! The two service groups are T11 and T12 on the
! configuration worksheet
! We will use group 72 for TCP and 73 for UDP.
ip wccp 72
ip wccp 73


! Repeat the following lines for each WAN interface
! you wish to accelerate:
interface <WAN_Interface>

! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect out
ip wccp 72 redirect in
ip wccp 73 redirect out
ip wccp 73 redirect in


^Z
```

Method B is preferred in circumstances when the routers do not support the **wccp redirect out** statement.

Example

Following is an example of configuring a Cisco IOS router:

```
! This example is for WCCP mode, not WCCP clustering
```



```
! (which is covered elsewhere)
config term
ip wccp version 2
!
! The two service groups are T11 and T12 on the
! configuration worksheet
! We will use group 72 for TCP and 73 for UDP.
ip wccp 72
ip wccp 73


! Repeat the following lines for the WAN interface
! you wish to accelerate:
interface <WAN_Interface>
!
! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect in
ip wccp 73 redirect in

! Repeat the following lines for all the LAN interfaces
! EXCEPT those connected to the SD-WAN appliance:
interface <LAN_Interface>
!
! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect in
ip wccp 73 redirect in

^Z
```

Remember to save your router configuration when you are satisfied that it is correct.

One accelerator instance manages WCCP control traffic on behalf of all the instances. The WCCP control traffic is negligible. The actual data traffic is divided among all the accelerators.

Note: The GUI calls WCCP mode “single cache.”

Summary: To configure the accelerators for WCCP mode, first enable WCCP mode. Then, configure service groups and create a WAN link definition, on the SD-WAN appliance, for each WAN link on each WCCP router. (Each link has two service groups, one for TCP and one for UDP.) If a service group is already defined for a given link, add the current router’s IP address to the definition. Test the service group’s WCCP status before creating the WAN link definition, and verify link traffic and acceleration status before configuring the next WAN link.

To configure the accelerators for WCCP mode

1. On the SD-WAN appliance, Navigate to Configuration > Appliance Settings > Advance Deployments > WCCP page.
2. If the **Enable** button is displayed, click it to enable WCCP mode on the appliance. (If the Disable button is displayed, WCCP mode is already enabled.)
   Note: We will actually be configuring two caches.
3. In the Select Mode area, select **Single Cache**.
4. Starting with accelerator instance #1 (labeled "WCCP Cache 1" on the page), configure the SD-WAN IP Details by entering the external VIP you defined for accelerator instance 1 (T5 on your worksheet for instance #1, T6 for instance #2). Set the subnet mask for the external traffic network (T2 on your worksheet). Set the gateway IP for the external traffic network (T1 on your worksheet). Click Save. The Configure Service Group controls appear.
5. In the Configure Service Group section, click Add. An Add Service Group popup appears.
6. In the Service Group Details area, specify a WCCP service group ID in the ID field. This ID must match one of the service groups that you have defined on your router. Start with the lowest-numbered service group in your list (T11 for the TCP service class, T12 for the UDP service class.).
7. In the WCCP Priority field, set the WCCP priority to 100 for instance #1, or to 80 for instance #2. (Other values work. Use a priority for instance #1 that is greater than the priority for instance #2, and use a priority for instance #2 that is greater than zero.)
8. From the Protocol list, select a protocol. You will perform this step for both TCP and UDP. Start with TCP.
9. In the Service Group Password field, enter a password if your router is configured to require one. Otherwise, leave the field blank.
10. In the Router Communications Details area, in the Router IP Address field, enter the IP address of the router. This is the router's address for its appliance-facing interface (T8 on your worksheet). If you use multiple routers to communicate with the appliance, list them all here.
11. From the Router Assignment list, select a router assignment (Hash, Mask, or Auto). Auto is recommended. If Auto is selected, Hash is negotiated if the router supports it. Otherwise Mask is used.
12. From the Router Forwarding list, select **Level 2** or **GRE**. The same method must be used for both outbound and inbound packets. L2 is recommended whenever possible, as GRE adds overhead to both the router and the appliance. L2 requires that your router support Level 2, and that the router's IP and the VIP addresses be in the same subnet. Otherwise, use GRE.
13. Click Create.
14. Repeat steps 6-13 with the next service group in sequence, but selecting UDP instead of TCP.
15. Repeat steps 4-14 for instance #2 (called "WCCP Cache 2" in the GUI), except that the Cache IP is T6 from your worksheet (instead of T5), and the WCCP priority value is 80 (instead of 100).
16. If desired, click Advanced Settings on the WCCP page and select a quicker timeout (Responsive or Tolerant, rather than Default). This is a WCCP 2.1 feature and is not supported by all routers. If the appliance has trouble connecting to the router, set this parameter back to Default.

Note: You must consider the following points when configuring a Citrix SD-WAN 4000/5000 appliance:

- Traffic is load balanced across the accelerators on the basis of NetScaler load balancing policies.
- The WCCP service group ID that you assign to the accelerator must match a service group defined on your router, or the WCCP negotiation fails.

Verifying WCCP Mode

Aug 09, 2017

You can monitor the WCCP configuration from the SD-WAN GUI.

To monitor the WCCP configuration

1. Navigate to the Monitoring > Appliance Performance > WCCP page.
2. Select a cache and click Get Info. A Cache Status page displays the WCCP configuration, as shown in the following figure.

3. Start traffic that should be forwarded through the SD-WAN appliance and monitor the connection on the Monitoring > Optimization > Connections page.
   - If the connections are shown on the Accelerated Connections tab, that is an indicator that everything is working.
   - If the connections are on the Unaccelerated Connections tab, look at the Details column. A routing asymmetry detected message implies that one of the ip wccp redirect lines on the router is missing or has an error, or that different paths are taken by client-server and server-client traffic.
   - If no connections are shown, but the appliance reports that it is connected to the router, and the WCCP monitoring page shows no errors, the issue is probably with the router configuration.

WCCP Clustering

Aug 09, 2017

The WCCP clustering feature enables you to multiply your acceleration capacity by assigning more than one SD-WAN appliance to the same links. You can cluster up to 32 identical appliances, for up to 32 times the capacity. Because it uses the WCCP 2.0 standard, WCCP clustering works on most routers and some smart switches, most likely including those you are already using.

Because it uses a decentralized protocol, WCCP clustering allows SD-WAN appliances to be added or removed at will. If an appliance fails, its traffic is rerouted to the surviving appliances.

Unlike SD-WAN high-availability, an active/passive pair that uses two appliances to provide the performance of a single appliance, the same appliances deployed as a WCCP cluster has twice the performance of a single appliance, delivering both redundancy and improved performance.

In addition to adding more appliances as your site's needs increase, you can use Citrix's "Pay as You Grow" feature to increase your appliances' capabilities through license upgrades.

Citrix Command Center is recommended for managing WCCP clusters. The following figure shows a basic network of a cluster of SD-WAN appliances in WCCP mode, administered by using Citrix Command Center.

Figure 1. SD-WAN Cluster Administered by Using Citrix Command Center

The WCCP protocol supports up to 32 appliances in a fault-tolerant, load balanced array called a cluster. In the example below, three identical appliances (same model, same software version) are cabled identically and configured identically except for their IP addresses. Appliances using the same service groups with the same router can become a load balanced WCCP cluster. When a new appliance registers itself with the router, it can join the existing pool of appliances and receive its share of traffic. If an appliance leaves the network (as indicated by the absence of heartbeat signals), the cluster is rebalanced so that only the remaining appliances are used.

Figure 2. A load-balanced WCCP cluster with three appliances

One appliance in the cluster is selected as the designated cache, and controls the load-balancing behavior of the appliances in the cluster. The designated cache is the appliance with the lowest IP address. Because the appliances have identical configurations, it doesn't matter which one is the designated cache. If the current designated cache goes offline, a different appliance becomes the designated cache.

The designated cache determines how the load-balanced traffic is allocated and informs the router of these decisions. The router shares information with all members of the cluster, so the cluster can operate even if the designated cache goes offline.

Note: As normally configured, a SD-WAN 4000/5000 appliance appears as two WCCP caches to the router.

Load-Balancing Algorithm

Load balancing in WCCP is static, except when an appliance enters or leaves the cluster, which causes the cluster to be rebalanced among its current members.

The WCCP standard supports load balancing based on a mask or a hash. For example, SD-WAN WCCP clustering uses the mask method only, using a mask of 1-6 bits of the 32-bit IP address. These address bits can be non-consecutive. All addresses yielding the same result when masked are sent to the same appliance. Load balancing effectiveness depends on choosing an appropriate mask value: a poor mask choice can result in poor load-balancing or even none, with all traffic sent to a single appliance.

Deployment Topology

Aug 09, 2017

Depending on your network topology, you can deploy WCCP cluster either with a single router or with multiple routers. Whether connected to a single router or multiple routers, each appliance in the cluster must be connected identically to all routers in use.

In the following diagram, three SD-WAN appliances accelerate the datacenter's 200 Mbps WAN. The site supports 750 XenApp users.

As shown on the SD-WAN Datasheet, a SD-WAN 3000-100 can support 100 Mbps and 400 users, so a pair of these appliances supports 200 Mbps and 800 users, which satisfies the datacenter's requirements of a 200 Mbps link and 750 users.

For fault tolerance, however, the WCCP cluster should continue to operate without becoming overloaded if one appliance fails. That can be accomplished by using three appliances when the calculations call for two. This is called the N+1 rule.

Failure is an unusual event, so usually all three appliances are in operation. In this case, each appliance is supporting only 67 Mbps and 250 users, leaving plenty of headroom, and making good use of the fact that the cluster has three times the CPU power and three times the compression history of a single appliance.

Without WCCP clustering, the same level of capacity and fault-tolerance would require a pair of SD-WAN 4000-500 appliances in high availability mode. Only one of these appliances is active at a time.

Using multiple WAN routers is very similar to using a single WAN router. If the previous example is changed to include two 100 Mbps links instead of one 200 Mbps link, the topology changes, but the calculations do not.

Limitations

Aug 09, 2017

Configuring appliances in a WCCP cluster has the following limitations:

- All appliances within a cluster must be the same model and use the same software release.
- Parameter synchronization between appliances within the cluster is not automatic. Use Command Center to manage the appliances as a group.
- SD-WAN traffic shaping is not effective, because it relies on controlling the entire link as a unit, and none of the appliances are in a position to do this. Router QoS can be used instead.
- The WCCP-based load-balancing algorithms do not vary dynamically with load, so achieving a good load balance can require some tuning.
- The hash method of cache assignment is not supported. Mask assignment is the supported method.
- While the WCCP standard allows mask lengths of 1-7 bits, the appliance supports masks of 1-6 bits.
- Multicast service groups are not supported; only unicast service groups are supported.
- All routers using the same service group pair must support the same forwarding method (GRE or L2).
- The forwarding and return method negotiated with the router must match: both must be GRE or both must be L2. Some routers do not support L2 in both directions, resulting in an error of "Router's forward or return or assignment capability mismatch." In this case, the service group must be configured as GRE.
- SD-WAN VPX does not support WCCP clustering.
- The appliance supports (and negotiates) only unweighted (equal) cache assignments. Weighted assignments are not supported.
- Some older appliances, such as the SD-WAN 700, do not support WCCP clustering.
- (SD-WAN 4000/5000 only) Two accelerator instances are required per interface in L2 mode. No more than three interfaces are supported per appliance (and then only on appliances with six or more accelerator instances.)
- (SD-WAN 4000/5000 only) WCCP control packets from the router must match one of the router IP addresses configured on the appliance for the service group. In practice, the router's IP address for the interface that connects it to the appliance should be used. The router's loopback IP should not be used.

Planning Your Deployment

Aug 09, 2017

Deploying appliances in a WCCP cluster requires more planning than does deploying a single appliance. Read the following sections carefully before proceeding.

Selecting Appliances

Aug 09, 2017

The appliances you select for the deployment must all be the same model, running the same software version. Otherwise, management and troubleshooting can become impractical.

Your appliance choice is generally made by comparing your site's WAN bandwidth and number of WAN users to the capacities of the different appliances in the SD-WAN Data Sheet. For fault tolerance, always order one more appliance than is absolutely required according to the data sheet.

The number of appliances you need is found as follows, rounding up all fractions:
appliances = max (appliances_bw, appliances_users),
where
appliances_bw = (WAN_bandwidth / Optimized_WAN_capacity) +1
appliances_users = (WAN_users / Maximum_HDX_sessions) +1
Note that if appliances = 2, you can use just a single appliance instead of WCCP clustering, or an HA pair instead of WCCP clustering, since the equation builds in a spare appliance. In other words, WCCP clustering is not necessary (from a capacity perspective) unless appliances is 3 or more.

Example. Suppose you have 700 users and a 100 Mbps link. Some appliances you might consider are the SD-WAN 2000-050, the SD-WAN 3000-100, and the SD-WAN 4000-310.

Model	Optimized WAN Capacity	Maximum HDX Sessions	Appliances_bw	Appliances_users	Appliances
2000-050	50 Mbps	300	3	4	4
3000-100	100 Mbps	400	2	3	3
4000-310	310 Mbps	750	2	2	2

As you can see from the above table, the higher-performance platforms require fewer appliances to get the job done, as you would expect. The SD-WAN 4000-310 meets the requirements with a single appliance, and evaluates to two appliances only because the equations build in a spare.

You can always add more capacity by adding more appliances, but that is not always necessary. The bandwidth limits of two of the three choices, the SD-WAN 3000-100 and the SD-WAN 4000-310, can be increased through a license upgrade. The SD-WAN 2000-050 however, is already at the high end of the range for SD-WAN 2000 appliances.

Load-Balancing in the WCCP Cluster

Aug 09, 2017

Traffic is distributed among the appliances in the WCCP cluster. If an appliance leaves the cluster (through failure, overload, or being manually disabled), its traffic is rebalanced by distributing it among the surviving members. If an appliance joins the cluster, traffic is rebalanced once more to give the new appliance its fair share.

The Address Mask

Traffic is distributed on the basis of an address mask that is applied to the source and destination addresses of WAN traffic. You must select an appropriate mask field for efficient load-balancing. An inappropriate mask can result in load-balancing that is poor to nonexistent. For example, if the mask matches an address field that is identical at all your remote sites, all your WAN traffic is sent to a single appliance, overloading it. For example, if all of your remote sites have an address in the form of 10.0.x.x, and your mask bits are within the 10.0 portion of the address all traffic is sent to a single appliance.

The address bits extracted by the address mask are used as an index that is used (indirectly) to select one of the WCCP caches (appliances). For example, an address mask with two "one" bits results in four possible values, depending on the address. Each of these values can be thought of as a bucket. With two mask bits, you have four buckets, numbered 0-3. The buckets are assigned to WCCP caches. To be effective, there must be at least as many buckets as caches. If you use a two-bit mask and have five or more caches, some caches are idle, because each bucket is assigned to only one cache, and there are not enough buckets to cover all five caches:

Cache	1	2	3	4	5
Buckets	0	1	2	3	-

If there are more buckets than caches, some caches are assigned multiple buckets. For example, if you set three mask bits, creating eight buckets, and you have four caches, two buckets are assigned to each cache. If you have five caches, three caches are assigned two buckets each, and two caches are assigned just one. If each bucket represents the same number of users, you have a 2:1 load imbalance across caches:

Cache	1	2	3	4	5
Buckets	0-1	2-3	4-5	6	7

Increasing the number of set mask bits reduces this imbalance. With four mask bits (16 index values) and five caches, four caches receive three buckets and one cache receives four buckets, resulting in only a 4:3 imbalance. With six set mask bits (the largest number supported), four caches receive 13 buckets and one receives 12, which is only a 13:12 load imbalance.

Cache	1	2	3	4	5
Buckets	0-12	13-25	26-38	39-51	52-63

Ideally, you would like each remote site to be directed to a single appliance in the WCCP cluster, so that all traffic to and from a given site is stored in the same compression history. With this arrangement, any traffic from one user at the site can be used to compress similar traffic from any other user at that site. In other words, for compressibility, load-balancing works best if it the address mask selects the bits that differentiate one remote site from another. These are often the least-significant bits of the subnet portion of the IP address. Using these bits tends to allocate the same number of remote sites (not users) per local appliance. A mask that aligns with the host portion of the address instead of the subnet results in a

more equal number of remote users (not sites) per appliance, but at the expense of compression effectiveness. (Compression is only effective when connections flow through the same appliances, and splitting traffic from the same remote site between two or more local appliances interferes with this.)

Finally, for good load-balancing, each "one" bit in the address mask must be set to one on 50% of the remote addresses, and set to zero on 50% of the remote addresses. This is not the case on all address bits, since in most WANs, the highest-order network bits never change at all (such as the 10 in 10.x.x.x). Such bits must never be selected by the address mask.

In addition, many subnets are only sparsely populated. For example, if only 50 addresses are used in the subnet 10.1.2.0/24, and they are assigned sequentially, the two higher-order host bits (representing the unused range of 10.1.2.64-10.1.2.255) for this subnet never change, and if these two bits are included in the address mask, three-fourths of the buckets receive no traffic.

Useful compromises between these two extremes can generally be found.

Follow these rules:

- The number of ”one” bits in the address mask must allow at least as many combinations as there are WCCP caches in the cluster. That is, if you have eight appliances, the address mask must contain at least three ”one” bits.
- The ”one” bits in the address mask must each be inside the active address range for most of your remote subnets, or they skew the load-balancing distribution.
- The mask should split the address range of individual remote sites into as few pieces as possible, for best compression performance.
- If a remote appliance is faster than the local members of the WCCP cluster, the mask should be designed to divide its traffic between multiple local appliances. For example, a 100 Mbps remote appliance should have its traffic split between two 50 Mbps local appliances by setting a bit inside the remote appliance’s active address range.
- The “one” bits in the mask are typically contiguous, but this is not required. They can be in any pattern.

Example: Suppose you set an address mask of 0x0000 0f00, which has four “one” bits. This defines a four-bit field that is extracted from the IP address, yielding 16 possible results (16 buckets). These buckets are in turn assigned to the actual WCCP caches in the WCCP cluster.

Address	Masked Address (mask = 0x0000 0f00)	Bucket
10.0.0.5	0.0.0.0	0
10.0.1.128	0.0.1.0	1
155.0.2.55	0.0.2.0	2
253.100.255.2	0.0.15.0	15
10.0.15.1	0.0.15.0	15

Zero bits in the mask are ignored, and the “one” bits are used to define the extracted field. So if the mask is 0x10 10 10 10, these widely separated “one” bits are extracted into a four-bit field, declaring 16 buckets and a bucket numbers in the range of 0-15.

If the mask value is set to zero, a default value of 0x00 00 0f 00 is used.

Assigning Buckets to Appliances

Aug 09, 2017

The mapping of bucket to appliances is subject to several variables:

- Which appliances are available: If an appliance is down, its share of buckets are given to the available appliance. If a new appliance is added to the cluster, it is given its fair share of buckets.
- The mapping algorithm used (deterministic or least-disruptive).
- The order in which appliances come online (least-disruptive mapping only).
- The IP addresses of the appliances. WCCP algorithms can use a sorted list of appliance IP addresses; for example, assigning buckets to appliances in the same order as the appliance IP addresses.

The most important of these factors, from an administrator's point of view, is the mapping algorithm.

Deterministic mapping. The deterministic mapping algorithm is less graceful than the least-disruptive algorithm, but it supports Hot Standby Router Protocol (HSRP) and Global Server Load Balancing (GSLB) routing, and is required when multiple routers using such protocols share the WCCP cluster.

Deterministic mapping is also the preferred method when the cluster has only two appliances.

Assignments are based on the IP addresses of the active appliances. Each appliance gets its fair share of bucket, with the lowest-numbered bucket being assigned to the appliance with the lowest IP address. If there are more appliances than buckets, the leftover appliances (with no bucket assigned to them) are the ones with the highest-numbered IP addresses. This deterministic assignment allows traffic to arrive for a single connection through any of the routers in the service group and be forwarded to the same appliance.

Reassignment can be disruptive to accelerated connections, which are reset if they migrate to a different appliance. With deterministic mapping, the number of buckets that are reassigned to new appliances can be quite high if there are three or more appliances.

Least-disruptive mapping. When a bucket is assigned to a different appliance, any open accelerated connections that used the old appliance is reset. The least-disruptive algorithm keeps the reassignment to a minimum. For example, if you have three appliances, and one appliance fails, the new mapping preserves roughly two-thirds of the assignments and remaps the remaining third (which fails anyway, because their appliance failed). The least-disruptive algorithm does not support HSRP or GSLB routing, because it is not guaranteed to result in identical mappings on all the routers in the service group, and therefore, packets from a single connection might be sent to two different appliances by two different routers, which causes accelerated connections to fail.

Startup and Failover Behavior

Aug 09, 2017

Each appliance registers itself with the routers specified in its service class definitions. The first appliance to register itself, becomes the *designated cache*, and works with the routers to apportion traffic between itself and the other caches (called *subordinate caches*). Because your appliances use identical WCCP algorithms, it does not matter which one becomes the designated cache.

As additional appliances come online, they are added to the WCCP cluster, and the traffic is reapportioned among the active appliances. This happens at ten-second increments. After a cold start of the routers or appliances, all of the appliances might come online within the same ten-second window, or they might arrive over multiple ten-second windows, causing traffic to be reapportioned multiple times before it stabilizes. In the latter case, the appliances that come online first maycan become overloaded until additional appliances come online.

An accelerated connection fails when allocated to a different appliance, making reallocation disruptive. This is not true of non-accelerated connections, which generally experience a delay of thirty seconds or more, and then continue. The least-disruptive mapping option minimizes the amount of reallocation when an appliance fails.

If an appliance fails or otherwise goes offline, its absence is noted, and the designated cache reapportions its traffic to the remaining appliances. If the designated cache itself goes offline, the role of designated cache is also reapportioned. It takes about thirty seconds for the cluster to react to the loss of a cache.

Deployment Worksheet

Aug 09, 2017

On the following worksheet, you can calculate the number of appliances needed for your installation and the recommended mask field size. The recommended mask size is 1-2 bits larger than the minimum mask size for your installation.

Parameter	Value	Notes
Appliance Model Used		—
Supported XenApp and XenDesktop Users Per Appliance	Uspec =	From data sheet
XenApp and XenDesktop Users on WAN Link	Uwan =	—
User overload Factor	Uoverload = Uwan/Uspec =	—
Supported BW Per Appliance	BWspec =	From data sheet
WAN Link BW	BWwan =	—
BW Overload Factor	BWoverload = BWwan/BWspec =	—
Number of appliances required	N = max(Uoverload, BWoverload) +1 =	Includes one spare
		—
Min number of buckets	Bmin = N, rounded up a power of 2 =	—
If SD-WAN 4000 or 5000,	Bmin = 2*N, rounded up to a power of 2 =	—
Recommended value	B = 4 * Bmin if Bmin <= 16, else 2 * Bmin =	—
Number of "one" bits in address mask	M = log2(B)	If B=16, M=4.

Mask value: The mask value is a 32-bit address mask with a number of “one” bits equal to M in the above worksheet. Often

these bits can be the least-significant bits in the WAN subnet mask used by your remote sites. If the masks at your remote sites vary, use the median mask. (Example: With /24 subnets, the least significant bits of the subnet are 0x00 00 nn 00. The number of bits to set to one is log2(mask size): if mask size is 16, set four bits to one. So with a mask size of 16 and a /24 subnet, set the mask value to 0x00 00 0f 00.): _____

The above guidelines work only if the selected subnet field is evenly distributed in your traffic, that is, that each address bit selected by the mask is a one for half the remote hosts, and a zero for the other half. Otherwise, load-balancing is impaired. This even distribution might be true for only a small number of bits in the network field (perhaps only two or three bits). If this is the case with your network, instead of masking bits in the offending area of the subnet field, displace those bits to a portion of the host address field that has the 50/50 property. For example, if only three subnet bits in a /24 subnet have the 50/50 property, and you are using four mask bits, a mask of 0x00 00 07 10 avoids the offending bit at 0x00 00 0800 and displaces it to 0x00 00 00 10, a portion of the address field that is likely to have the 50/50 property if your remote subnets generally use at least 32 IP addresses each.

Parameter	Value	Notes
Final Mask Value		—
Accelerated Bridge		Usually apA
WAN Service Group		A service group not already in use on your router (51-255)
LAN Service Group		Another unused service group
Router IP address		IP address of router interface on port facing the appliance
WCCP Protocol (usually "Auto")		—
DC Algorithm		Use "Deterministic" if you have only two appliances or are using dynamic load balancing like HSRP or GSLB. Otherwise, use "Least Disruptive."

Configuring WCCP Clustering

Aug 09, 2017

After you have finalized the deployment topology, considered all limitations, and filled in the deployment worksheet, you are ready to deploy your appliances in a WCCP cluster. To configure the WCCP cluster, you need to perform the following tasks:

- Configuring the NetScaler Instances
- Configuring the Router
- Configuring the Appliance

Configuring the Router

Aug 09, 2017

WCCP configuration on the router is simple, because most WCCP parameters are set by the appliances.

Unlike legacy SD-WAN WCCP support, WCCP clustering uses two service groups for TCP traffic. One service group is used on the router's WAN interface, and the other is used on the router's LAN interfaces (except for the LAN interface used by the SD-WAN appliances themselves, when deployed in L2-mode WCCP cluster).

As shown in the following figure, you need to configure two service groups because WCCP allows the mask to be applied to either the source IP or the destination IP address, which is not quite what is required. To keep connections between two endpoints together, regardless of which endpoint initiates the connection, the appliance applies the address mask to the source IP address of incoming WAN traffic, and to the destination IP address of incoming LAN traffic. This requires two service groups.

The WAN service group uses WCCP source-ip address masking, while the LAN service group uses dest-ip masking. In some deployments, it may be necessary to reverse the assignments, using the "WAN" service group for your LAN interface and vice versa. This might occur if the number of local IP addresses greatly exceeds the number of remote IP addresses.

Figure 1. SD-WAN WCCP Cluster

To configure WCCP clustering on the router

This procedure assumes Cisco routers, but is similar on other routers. It uses the first of the two methods, discussed above, of redirecting WCCP traffic with an ip wccp redirect in statement on both LAN and WAN ports.

1. Fill in the WCCP clustering Deployment Worksheet.
2. Log on to your router
3. In the global declarations section, declare each service group on the WCCP clustering worksheet, listed as **WAN service group** and **LAN Service group**. For example, ip wccp 61 and ip wccp 62.
   Note: The ip wccp command allows, but does not require, a more elaborate syntax than this, and can specify an ACL name or a password. Both service groups must have the same password, if one is used. The ACLs can be different.
4. Inside the interface declarations for each WAN interface that connects to remote SD-WAN appliances, add an ip wccp x redirect in statement, where x is the WAN service group from the WCCP clustering worksheet.
5. Inside the interface declarations for each LAN interface (except the one connecting to the WCCP cluster, if you are using L2 mode), add an ip wccp y redirect in statement, where y is the LAN service group from the WCCP clustering worksheet.
6. Save your configuration.

Example. The following example uses service group 61 for the WAN service group and service group 62 for the LAN service group. Three router interfaces are used. One is connected to the WAN, one is connected to the LAN, and one is connected to the WCCP cluster.

```
!
! Example is for WCCP clustering using WCCP redirect in statements
! on LAN and WAN interfaces.
! This definition is appropriate for modern Cisco routers.
! Global declarations
ip wccp 61
ip wccp 62
```

```
!
interface GigabitEthernet1/1
description LAN interface. SG 62 is used for LAN
ip address 172.80.1.56 255.255.255.0
ip wccp 62 redirect in
!
interface GigabitEthernet1/2
description LAN interface attaching SD-WAN L2-WCCP appliances
description (No wccp redirect statements are used on this interface)
ip address 172.80.21.56 255.255.255.0
!
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
!
```

Note: If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.

- If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.
- If multiple routers shared the same WCCP cluster, they use the same service groups.

It is also possible to use ip wccp redirect statements on only the WAN interfaces:

```
! Example for WCCP clustering using WCCP redirect in/out statements on
! WAN interface only
! This definition is appropriate for modern Cisco routers.
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN. SG 62 is used for LAN.
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
ip wccp 62 redirect out
!
```

In many routers, the ip wccp redirect out path is not optimized in hardware, but uses the CPU. If the router's capabilities along this path exceeds the WAN speed, this method is practical, and is simpler than using redirect statements on every interface.

Router ACLs can be used to limit redirection. For example, for initial testing, perhaps only a single remote IP address might be allowed to be redirected through WCCP.

Configuring the Appliance

Aug 09, 2017

Repeat the following procedure for each appliance in the cluster:

1. Fill in the WCCP clustering Deployment Worksheet.
2. Navigate to Configuration > Appliance Settings > WCCP page.
3. Click Enable to enable WCCP mode on the appliance.
4. Select **Cluster (Multiple Caches)** option.

5. Fill in parameters in the **Select SD-WAN Cluster** section.

6. Enter T5 from your worksheet as the Cache 1 IP, T6 as the Cache 2 IP, T2 as the Subnet Mask, and T1 as the Gateway. Click **Save**. The **Configure Service Group** section appears.
7. In the Service Group Details section, specify the WAN and LAN service groups (T11 and T12 from your worksheet).
8. In the Priority field, select **100** (in practice this value is somewhat arbitrary).
9. From the Protocol list, select **TCP**.
10. In the DC Algorithm field, select **Deterministic** or **Least Disruptive**. “Deterministic” is always safe to use, and should be used if you are using only two appliances, or are using multiple routers. “Least Disruptive” disrupts fewer user sessions on failover when used with clusters of three or more appliances, but has restrictions on its use.
11. Set **Service Group Pair Status** to On.
12. If your router is configured to require a password, enter the password in the **Service Group Password** field. Otherwise, leave the field blank.
13. In the **Router Communications Details** section, enter the IP address of the router (T8 on your worksheet: often identical to T1 as well). This is the IP address of the appliance-facing router interface. If you use multiple routers to communicate with the appliance, list them all here.
14. From the Router Forwarding list, select Level 2 or GRE, according to the capabilities of your router. Use Level 2 if you can, and GRE if you must.
15. For the Mask Value, enter the value you determined from the WCCP Clustering worksheet. This is a critical value: a poor choice will result in poor load-balancing or none at all.
16. Click Create. This creates the WAN and LAN service groups.
17. On the Configuration > Optimization Rules > Link Definitions page, change the bandwidth limits on each defined WAN to 95% of the aggregate speed of all your WANs. This prevents the link from being under-utilized when load-balancing is imperfect. If ICA (XenApp/XenDesktop) is the dominant use, set each appliance to (95% of WAN bandwidth)/N, where N is the number of appliances (or twice the number of appliances if they are SD-WAN 4000 or 5000 units), to divide the bandwidth equally among the appliances. This latter method is most appropriate for applications with large numbers of active connections that have relatively low bandwidth requirements.

Testing and Troubleshooting

Aug 09, 2017

The **Monitoring** > **Appliance** > **Application Performance** > **WCCP** page shows the current state of not only the local appliance but of all other appliances that have joined the cluster. Select a WCCP cache and click **Get Info**.

The Cache Status tab shows the local appliance's status. When all is well, the status is "25: has assignment." You must refresh the page manually to monitor changes in status. If the appliance does not reach the status of "25: has assignment" within a timeout period, other informative status messages are displayed.

Additional information is displayed when you click on the Service Group or the Routers tabs.

The Cluster Summary tab displays information about the WCCP cluster as a whole. As a side effect of the WCCP protocol, each member of the cluster has information about all the others, so this information can be monitored from any appliance in the cluster.

Your router can also provide status information. See your router documentation.

Inline Mode

Aug 09, 2017

Note: This documentation is valid only for appliances that were provisioned from a factory-reset state with the release 7.1 Configuration Wizard. For appliances that have been updated to release 7.1, but not reprovisioned from a factory-reset state, see the release 7.0 documentation.

When you deploy a SD-WAN 4000/5000 appliance in inline mode, pairs of Ethernet ports on the appliance function as accelerated bridges. Traffic flows into one bridge port and out the other. When two sites with appliances communicate, TCP connections between the sites can be accelerated. Traffic that cannot be accelerated is passed through transparently, as if the appliance were not there.

For maximum reliability, the bridge pairs are equipped with a bypass feature that causes the two ports to be connected to each other should the appliance fail or lose power, allowing traffic to continue flowing even during such an outage.

Starting in release 7.1, inline mode depends on the NetScaler “add interfacePair” command to isolate bridge traffic, ensuring that traffic from one bridge pair stays on that pair. In previous releases, this was done with VLAN definitions. Appliances that were provisioned with release 7.1 have this feature enabled by default. Appliances that were upgraded to release 7.1 without reprovisioning retain their old configuration.

Inline mode is currently recommended only for sites where WCCP is not practical, and which have a single WAN link, or have fully independent WAN links that do not use dynamic routing, load-balancing, or fail-over.

Deployment Topology

Aug 09, 2017

The following figure shows a SD-WAN 4000/5000 appliance in inline mode.

Figure 1. Basic cabling for inline mode

As shown in the above figure, inline mode is a two-arm mode. For inline deployments, the NetScaler instance is configured in L2 (bridged) mode, but the accelerators are connected internally to the NetScaler instance in a one-arm configuration.

Inline mode is the easiest mode to configure. You connect one port of an accelerated pair to the WAN router and the other to the LAN network. The appliance transparently accelerates traffic flowing between the two ports, which to the rest of the network appear to be an Ethernet bridge.

You can also deploy the appliance to accelerate traffic from certain resources only, such as back-end servers, and not the traffic of the entire network. Such an arrangement reserves the appliance's resources for the selected traffic. In this case, you install the appliance on the branch network that includes the resources for with you want to accelerate traffic.

The following figure shows partial site acceleration:

Figure 2. Partial site acceleration

Port Affinity

Aug 09, 2017

Port traffic on a given bridge must be isolated to that bridge. In release 7.1, this is done as part of the provisioning process. It can also be done manually with the “add interfacePair” command in the NetScaler command line interface.

The following examples show how this command is used to create port affinity on all bridged pairs in the appliance:

SD-WAN 4000

add interfacePair 1 -ifnum 1/1 1/2
add interfacePair 2 -ifnum 1/3 1/4
add interfacePair 3 -ifnum 1/5 1/6
add interfacePair 4 -ifnum 1/7 1/8
add interfacePair 5 -ifnum 10/1 10/2

SD-WAN 5000

add interfacePair 1 -ifnum 10/1 10/2
add interfacePair 2 -ifnum 10/4 10/5
add interfacePair 3 -ifnum 10/6 10/7

VLAN Trunking

Aug 09, 2017

VLAN trunking is also known as tagged VLAN and 802.1Q tagging. The 802.1Q tagging enables a networking device to add information to a frame at Layer 2 to identify the VLAN membership of the frame. Tagging also enables network environments to have VLANs that span multiple devices. A device that receives the packet reads the tag and recognizes the VLAN to which the frame belongs.

When you configure tagging on bridged interfaces, the VLAN configuration must be identical on both ports of the bridge.

Tagged VLANs are not supported on the management interfaces (ports 0/1 and 0/2).

For example, if your WAN link uses VLAN 412, you declare VLAN 412 as a tagged VLAN in the NetScaler instance, and bind it to both ports of the bridge (such as ports 10/1 and 10/2), as shown in the example below.

Figure 1. Tagged VLANs for VLAN trunking. VLAN 412 is tagged

VLANs can be declared in either of two ways:

1. From the System > Settings > Configure NSVLAN Settings dialog box. This method declares a VLAN whose broadcast traffic is isolated from other VLANS. This method is recommended for the management subnet. It requires a restart to take effect.
   Note: This VLAN configuration method is neither synchronized nor propagated in high availability mode. Therefore, you must perform the configuration independently on each appliance of a high availability setup.
2. From the Create VLANs dialog box (reached from Network > VLANs > Add..). This method does not create an isolated broadcast domain, from traffic originating in the NetScaler instance until we bind the NetScaler IP addresses to the VLAN. Adding such a VLAN does not require a restart. This method is recommended for all VLANs except the management subnet.

Ethernet Bypass

Aug 09, 2017

The appliance includes a bypass feature for inline mode. In a power failure, a relay closes and the input and output ports become electrically connected. This feature allows the Ethernet signal to pass through from one port to the other, as if the appliance were not there. The appliance functions like a cross-over cable connecting the two ports.

Besides a power failure, any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the bridge ports of the appliance are inaccessible.

- The bypass feature is disabled when the NetScaler instance is set to L3 mode. Because L3 mode is the factory default, inline mode should be configured before the appliance is placed in line with data traffic.
- The bypass feature is disabled when the appliance is in HA mode.
- A bypass event causes all bypass-enabled port pairs (except the loopback ports) to enter the bypass mode.
- The loopback ports never enter bypass mode.
- A bypass event occurs if the NetScaler instance or the bypass daemon in Dom-0 becomes unresponsive.
- A bypass event is not triggered by accelerators becoming unresponsive.
- The 1-Gigabit bypass ports are copper, and 10-Gigabit bypass ports are fiber ports.

Configuring Inline Mode

Aug 09, 2017

Note: These instructions are valid only for appliances that were provisioned from a factory-reset state with the release 7.1 Configuration Wizard. For appliances that have been updated to release 7.1, but not reprovisioned from a factory-reset state, see the release 7.0 documentation.

Basic configurationis performed by the release 7.1 Configuration Wizard. Additional configuration is required only if you use VLAN trunking on data passing through the bridges.

1. Do not attach the bridges to your traffic networks yet. Begin by provisioning the appliance with the configuration wizard.
2. If your traffic does not contain tagged VLAN traffic, skip to the last step of this procedure.
3. Navigate to the NetScaler instance at Configuration > NetScaler > Instances and click on the IP address of the NetScaler instance.
4. If the Citrix SD-WAN Connector Get Started page appears, ignore it.
5. Click Configuration > Network > VLANs > Add.
6. In the **Create VLAN** dialog box, configure the tagged VLANs to use bridge #1. In the **VLAN Id** field, enter the VLAN IDof the first tagged VLAN (VLAN1.1 on your worksheet). Under **Interfaces**, clear all the check boxes. Then, select **Active** and **Tagged** for the first port of the bridge (T9 on your worksheet) and the second port of the bridge (T10). Click **Create**.
7. Repeat the previous step for any remaining VLANs using bridge #1 (VLAN1.2, VLAN1.3, and so on).
8. Repeat for any additional bridge pairs.
9. Click Close.
10. Click Save to save your configuration.
11. Connect the bridges to your traffic networks. Configuration is complete.

Configuring the High Availability Setup on the Appliances

Aug 09, 2017

High availability(HA) works directly between the NetScaler instances of two SD-WAN 4000 or 5000 appliances. As shown in the configuration , the two appliances are configured almost identically, except for management network IP addresses.

Note: The accelerator instances on the two appliances are not synchronized, and must be kept consistent manually. Take this into account when deciding whether to use HA.
Note: For a smooth installation, install and test one appliance before adding the second one, noting all configuration changes, especially to the accelerator.
Note: You must use the same aPA and signaling addresses on both appliances. However, all management subnet IP addresses must be unique on each appliances.
If the active appliance becomes unavailable, the passive appliance transparently takes over the function of the primary appliance. This is called “failover.” As a result, disruption of services over the network is minimal. After a failover, all clients must reestablish their connections to the managed servers, but the session persistence rules are maintained as they were before the failover.

HA is supported in all deployment modes, and the HA configuration procedure is the same for all modes. The two appliances should be running identical hardware, licensing, and software releases, and must be deployed identically, using the same deployment modes on the same subnets.

When you enable HA, the configuration of the primary appliance’s NetScaler instance is copied to the secondary appliance as part of the NetScaler HA synchronization process.

To configure a high availability setup of NetScaler instances

1. Complete the configuration for your chosen deployment mode (inline or WCCP). Note that all parameters for the external traffic subnet and the private traffic subnet are identical for both appliances, but some management subnet values are different on the two appliances.
2. Fully configure appliance #1 and test it thoroughly. If inline mode is used, do not connect the traffic network to the bridge ports on appliance #2.
3. Fully configure appliance #2. Note that some parameters are different on the two appliances: H1, H4, H5, H6, and H17 are used on the second appliance in place of M1, M4, M5, M6, and M17. Make sure that the accelerators are configured identically to the ones on appliance #1, and that both appliances have identical VLAN definitions in the NetScaler instances.
4. Access the NetScaler instance on appliance #1, by specifying its IP address (M17) in a web browser.
5. Log on to the NetScaler instance.
6. In the Navigation pane, expand the System node.
7. Select the High Availability node.
8. Click Add, as shown in the following figure.
   Figure 1. Configuring a high availability setup of the NetScaler instances

9. In the remote Node IP Address field of the High Availability Setup dialog box, specify the NSIP address of the NetScaler instance of the other appliance #2 (H17 on your worksheet), as shown in the following figure.

Figure 2. Configuring a high availability setup of the NetScaler instances

10. Click OK. The appliances are now configured as a high availability pair, as shown in the following figure.
    Figure 3. Configuring high availability on the NetScaler instance

Note: To learn more about setting up high availability on a NetScaler instance, see the High Availability node of the Citrix eDocs website.

Evaluating the Configuration

Aug 09, 2017

Putting your appliance online in a production network requires special attention to prevent disruption or confusion, especially in a complex environment

When deploying SD-WAN 4000/5000, the basic rollout decision is whether to activate the entire deployment at once or to roll it out in stages. In a large or complex environment, a phased approach avoids trouble, and the deployment can be extended at will. This type of approach calls for the use of WCCP. The following example illustrates one approach for such a site:

1. Configure the system as described in the installation procedure, except for the router. There, instead of setting up WCCP redirection for all incoming and outgoing WAN traffic, set it up for traffic to and from either a single remote site or a single IP address at that site. The remote site must already contain an enabled SD-WAN appliance.
2. The accelerator page. If not, check your WCCP configuration on the router and on the accelerators, and check your NAT definitions on the NetScaler instance by using **Monitoring: WCCP** page. If not, check your WCCP configuration on the router and on the accelerators, and check your NAT definitions on the NetScaler instance by using **nstrace**. If **nstrace** reveals an issue, and your definitions look correct, rebooting the appliance may resolve the issue.
3. Test acceleration between the new site and the remote site, with the remote site as the client side and the SD-WAN 4000/5000 equipped site as the server side, as described in General Monitoring.
4. If traffic does not appear, the router is not sending traffic to the SD-WAN 4000/5000 properly. The error could be in the Router configuration, the NetScaler configuration, or the SD-WAN WCCP configuration. Double-check these settings.
5. If traffic appears but is not accelerated, you might have a problem with asymmetrical routing, with not having a SD-WAN license installed, or with having acceleration disabled either globally or on the service classes associated with the traffic.
6. When all is working properly, test reverse connections, where a site on the SD-WAN 4000/5000 side is the client and the remote site is the server, if applicable.
7. If using NetScaler HA, save the configuration of the individual WCCP-enabled instances from the individual instances' GUIs, and save the configuration of the accelerator, do basic configuration manually, then restore the saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance. Be careful to avoid IP address conflicts.SD-WAN 4000/5000, do basic configuration manually, then restore the saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance. Be careful to avoid IP address conflicts.
8. If using NetScaler HA, save the configuration of the individual WCCP-enabled instances from the individual instances' GUIs, and save the configuration of the accelerator, restore these saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance.SD-WAN 4000/5000, restore these saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance.
9. Expand the scope of acceleration to include more remote sites, and repeat the above testing. When doing so, also examine the **Monitoring: System Load** page, especially during peak periods, to verify that the SD-WAN 4000/5000 is not heavily loaded.
10. Continue this process until the entire WAN is being accelerated.

Use the SD-WAN 4000/5000Use the SD-WAN 4000/5000 GUI to monitor traffic after you configure a LAN link and a WAN link. SD-WAN 4000/5000 allows a very simple link definition.

Configuring the Links

To enable monitoring, you must first configure one LAN link and one WAN link. To do so, edit the default links on the Configure: Links page as follows:

1. Edit one link so its name is "LAN," its type is "LAN," and its speed is 10 Gbps in both directions. Delete its existing filter rule, then click Add Rule, and then click Save to save a link definition that matches all traffic.
2. Edit the other link so that its name is "WAN," its type is "WAN," its speed is 95% of the aggregate speed of your site's WAN links in each direction. Delete its existing filter rule, then click Add Rule, and then click Save to save a link definition that matches all traffic.

To verify that link configuration is working correctly, traffic must be flowing. If the network does not have enough traffic to fill the WAN link to capacity, run test traffic to fill the network to capacity. Then look at the link reports on the Reports: Link Usage tab. The following figure shows these reports.

General Monitoring

1. If WCCP is configured, verify that the service groups are in operation and the routers are redirecting traffic. (Note that the SD-WAN WCCP page packet counts are not present in SD-WAN 4000/5000. Check traffic by other means, such as on the Monitoring: Active Connections page, and on the router.)
2. On the **remote** SD-WANs, verify that outgoing connections are being accelerated, and that all accelerated connections to the datacenter report the same Partner Unit on the remote appliance's Monitoring: Connections page. When load-balancing is working properly, all outgoing accelerated connections show the same Partner Unit. (However, incoming accelerated connections might show different units.)
3. Double-check remote SD-WANs for correctly set bandwidth limits, to prevent remote issues from being misidentified as datacenter issues.
4. Generally monitor the SD-WAN 4000/5000 unit for alerts.
5. In the broker UI, use the Dashboard, the Monitoring: Remote Partners, and perhaps the Monitoring: Appliance Load pages to monitor the overall activity and load of the system.

Troubleshooting Tips

Aug 09, 2017

While most installations complete smoothly, some installations require knowledge of the appliance's internal structure or the use of little-known features before you can perform additional monitoring and troubleshooting. These troubleshooting tips provide information and techniques that allow a more in-depth analysis of the appliance.

Understanding Internal Addresses

Some reports show addresses on the private subnets within the SD-WAN 4000/5000, so it's good to know what these addresses mean. These subnets connect the virtual machines together, without connecting to external ports.

All these addresses are on the local link subnet 169.254.0.0/16, described in RFC3927. This address space is segmented into three partly overlapping subnets: system management, private traffic, and accelerator management subnets.

System Management Subnet, 169.254.0.0/16

Function	Address
Management Service	169.254.0.10/16
NetScaler Instance	169.254.0.11/16
XenServer	169.254.0.1/16

Private Traffic Subnet, 169.254.10.0/24

Function	Address
apA IP, accelerators 1-8	169.254.10.21/24 - 169.254.10.28/24
apA Signaling IP, accelerators 1-8	169.254.10.121/24 - 169.254.10.128/24
NetScaler Instance	169.254.10.11/24

Accelerator Management Subnet, 169.254.0.21-28/24

Function	Address
Accelerator unified management IP (controls all accelerators)	169.254.0.20/24
Primary Port IP, accelerators 1-8	169.254.0.21/24 - 169.254.0.28/24

Figure 1. Virtual machines in the SD-WAN 4000 and 5000. The system management subnet is not shown in this diagram. The traffic shaper manages traffic from all accelerators and is controlled via the accelerator GUI.

Checking and Correcting Accelerator Instance Status

Sometimes an error message may indicate an issue with one of the virtual machines in the appliance. To check their status, go to the System Configuration page and select an Instance view of either the SD-WAN or NetScaler subsystems. For example, the SD-WAN page is shown below:

- A fully active instance will show a green circle for VM State, Instance State, and Licensed.
- Your appliance may have more instances present than are licensed; ignore the unlicensed instances.
- If the VM State or Instance State of the remaining instances are not green, use the “Rediscover” action to attempt to bring these instances back into operation.

You can also get detailed information for each instance, as shown below:

- Every instance should have a Status of “Inventory from SD-WAN Instance completed.
- Every instance should be running the same version of the software.
- Every instance should have the netmask (255.255.255.0) and gateway (169.254.0.20).
- Instances that show an uptime shorter than other instances have rebooted since the last whole-system reboot.

Logging Into the NetScaler Instance

Sometimes it is useful to log into the NetScaler instance to check its status or do configuration. You can log into the NetScaler instance from the NetScaler Instances page of the System Configuration view, as shown below. Click the IP Address link.

You can also log into the NetScaler instance directly from your browser if you know its IP address on the management port (port 0/1).

Once logged in, you will see the NetScaler GUI, which identifies itself as NetScaler VPX at the top of the page.

This is the standard NetScaler user interface. Using monitoring features is safe. Configuration changes should be made with caution, as the SD-WAN 4000/5000 makes undocumented assumptions about how the NetScaler instance is configured.

Using Ping and Traceroute

The ping and traceroute utilities are not available on the accelerator instances, as they are on other SD-WAN products. Instead, you can use the equivalent features on the NetScaler instance, using the Diagnostics page as shown below.

These features will work over your external network and on the appliance’s internal subnets.

Using the System Dashboard

Unlike the SD-WAN Dashboard, the System Dashboard page is devoted largely to hardware monitoring.

- The System Health tables show a status summary, with a Details link for expanded information in graphical form.
- The Events tables show a status summary, with a Show Events link to see the related log entries.

The system below shows a conditions worth investigating:

- The power supply was struggling, as can be seen in both the Hardware Sensors Summary and the System Health Events log. (The count in the System Health Events heading shows that there was only one event, the Date field shows that it

happened long before the screen was displayed, and the Message indicates that the incident was Non-Critical, so replacing the power supply is likely not called for.)

- Several ports are marked as Down, which is only an error if a cable is supposed to be present. Most appliances have several unused ports.
- Fail To Wire lists FTW Disabled for all ports. This means that the network bypassing feature is not enabled on this appliance. Examination of the FTW Events showed that there were no actual events, indicating that the feature is probably disabled.

For each warning or error, additional details are available through the Details links or Show Events buttons.

Logging In To Different Instances via SSH

You can log into some of the virtual machines from the management port (port 0/1) using an ssh utility (such as PuTTY on Windows), logging in either as root or nsroot and using the administrative password. This will give you a shell prompt.

The most common use for logging on via SSH is to restore the IP address of an instance, typically the management service, that has become unreachable due to misconfigured network parameters. Otherwise, SSH is not recommended, as configuration changes can render the appliance unstable or unusable.

If neither of the two instances below are accessible over the network, you can log into the XenServer instance using the RS-232 port, which will give you a shell prompt.

Instance	Login	Password	Actual username
Management Service	nsroot	Admin password	root
Management Service	root	Admin password	root
XenServer	nsroot	Admin password	nsroot
XenServer	root	Admin password	root

Once logged into one of these virtual machines, you can use SSH from the shell prompt to reach the NetScaler instance or the accelerator at the appropriate 169.254.x.x address.

The usual UNIX/Linux commands are available, including the vi text editor.

Monitoring Individual Accelerator Instances

Logging into the accelerator GUI IP allows you to manage all the accelerator instances as a unit. Changes are automatically propagated to all the accelerator instances.

On rare occasions, you may wish to troubleshoot individual accelerator instances. To do this, use the following URL's:

Accelerator Instance	URL
1	https://<accelerator_ip>:4001
2	https://<accelerator_ip>:4002
...	

8	
Accelerator Instance	https://<accelerator_ip>:4008 URL

The login for the instances is admin. The password is the same admin password as is used on the other instances.

This is recommended for monitoring, not for making permanent changes, since any parameter you set in an instance may be overwritten later by the synchronization process.

Using Individual Elements of the Update Bundle

The update bundles distributed by Citrix are in a simple .tgz format (a tar archive compressed with gzip). It is sometimes useful to extract individual components from the archive, rather than going back to the the Citrix Web site and downloading them individually. This is most commonly useful with the management service (build-svm*.tgz) or the accelerator release (orbital*.bin).

The update bundle can be managed by tar/gzip or by archiving utilities like 7-zip.

NetScaler SD-WAN 4100 and 5100 WANOP Appliances

Aug 31, 2017

Citrix NetScaler SD-WAN 4100/5100 WANOP appliances are high-performance WAN accelerators for busy datacenters. These appliances combines multiple virtual accelerator instances with a single virtual instance of the NetScaler load-balancer, providing the performance of multiple SD-WAN WANOP appliances in a single package.

SD-WAN 4100/5100 WANOP WAN accelerators are the high end of the Citrix NetScaler SD-WAN product line. They are designed to accelerate sites with WAN links with speeds in excess of 1 Gbps, especially busy datacenters that communicate with a large number of branch and regional sites.

A single SD-WAN WANOP 4100/5100 appliance can support WAN speeds of up to 2 Gbps and up to 5000 XenApp/XenDesktop users.

For datacenters needing even more performance, multiple SD-WAN WANOP 4100/5100 appliances can be deployed as a load-balanced array using the WCCP clustering feature.

SD-WAN WANOP 4100/5100 is recommended at the hub of a hub-and-spoke deployment, where smaller appliances are used at the spokes, whenever the link speed or the number of XenApp/XenDesktop users is higher than can be supported by a smaller appliance.

If you require a secondary data center, SD-WAN WANOP 4000/5000 appliances can provide optimization for Data-Center to Data-Center replication. This optimization improves replication time and reduces bandwidth consumption.

For details on how to configure a SD-WAN WANOP appliance for DC-to-DC replication with NetApp SnapMirror, see http://support.citrix.com/article/CTX137181.

Architecture

Aug 31, 2017

Internally, the SD-WAN 4000/5000 appliance contains several virtual machines:

- A Xen hypervisor
- A NetScaler instance
- At least two accelerator instances
- A management server instance that manages the GUI and other tasks
- Internal networking

Figure 2. SD-WAN 4100/5100 virtual machines, internal networks, and external port usage (inline deployment shown)

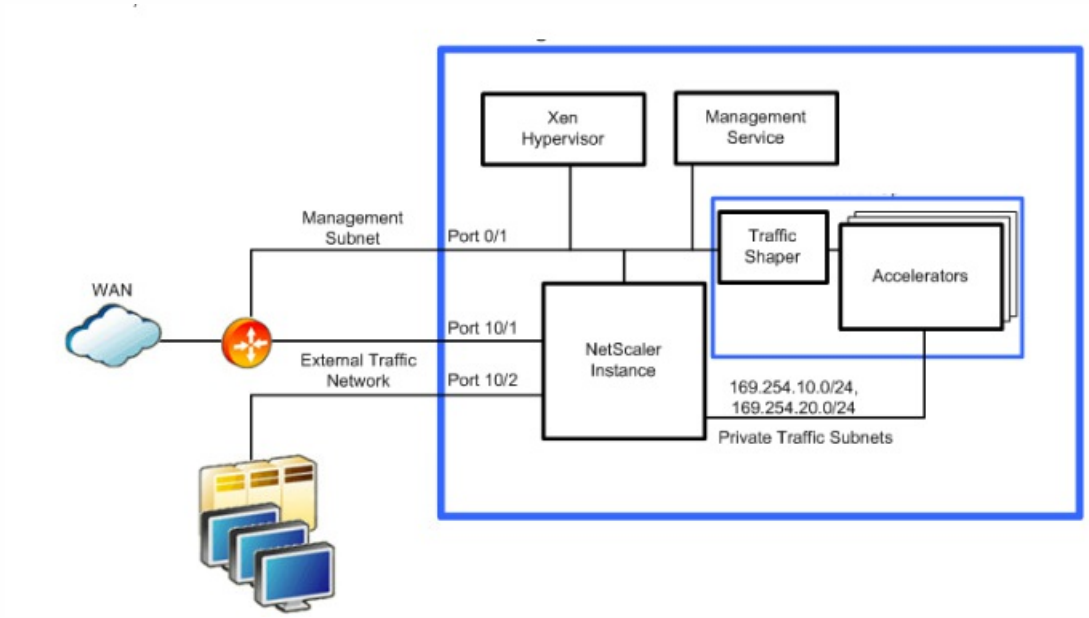

No WAN traffic enters or leaves the accelerators except as configured in the NetScaler instance. When the appliance is first used, the Provisioning Wizard sets up an initial configuration that provides communication and load balancing between the NetScaler instance and the accelerators.

The management service is the management configuration interface for the appliance, and provides access to key operating and monitoring elements of the appliance. The management service displays SD-WAN parameters as if they were from a single accelerator, and all changes made through this interface are applied to all the accelerator instances.

The Xen hypervisor hosts all the virtual machines. The hypervisor is not user-configurable and should not be accessed except at the request of Citrix.

The external network interfaces are divided into two categories: traffic interfaces and management interfaces.

Traffic Interfaces—The traffic interfaces include all the network interfaces except ports 0/1 and 0/2, which are used only for management. Acceleration takes place only on the traffic interfaces.

Note: You must keep the traffic interfaces isolated from the management interface to prevent ARP flapping and other problems. This isolation can be achieved physically or by tagging management interface and traffic interface packets with

different VLANs.

Management subnet—The virtual machines connect directly to the external management subnet, with different IP addresses for the management service, NetScaler instance, and XenServer.

Note: You must keep the traffic interfaces isolated from the management interface to prevent ARP flapping and other problems. This isolation can be achieved physically or by tagging management interface and traffic interface packets with different VLANs.

Private Internal traffic subnet—The accelerators' accelerated ports are connected to the NetScaler instance internally in a one-arm mode, using an internal traffic subnet. There is no direct connection between the instances' accelerated ports and the appliance's external ports. All accelerated traffic to the accelerators is controlled by the NetScaler instance.

Since this internal subnet is not accessible from outside the appliance, it uses non-routable subnets in the 169.254.0.0/16 range. The NetScaler instance provides NAT for features that require routable access to the accelerator Only the following two features of the accelerators require IP addresses that can be reached from the outside world:

- The signaling IP address, used for secure peering and the SD-WAN Plugin.
- IP addresses, used for communication with the router when the WCCP protocol is used.

In both cases, the number of externally visible IP addresses is independent of the number of accelerators the appliance has.

The internal traffic subnet requires two IP addresses per accelerator, plus an address for the NetScaler, plus one or two WCCP VIP addresses if WCCP is used. Since the internal network is private, it has an abundance of address space for these tasks.

Data Flow on the Private Traffic Subnet—The one-arm connection between the NetScaler instance and the accelerators uses the SD-WAN virtual inline mode, in which the NetScaler instance routes packets to the accelerators and the accelerators route them back to the NetScaler instance. Traffic flow over this internal traffic subnet is identical regardless of whether the mode visible to the outside world (on the external interfaces) is inline, virtual inline, or WCCP.

This traffic requires the SD-WAN "Return to Ethernet Sender" option, and the NetScaler MAC Address Forwarding and Use Subnet IP options, which are enabled by the Provisioning Wizard.

Deployment Mode Summary

The differences between WCCP mode, inline mode, and virtual inline mode can be summarized as follows:

- WCCP mode is a one-arm configuration. The accelerators establish WCCP control channels with the router. In WCCP mode, only one or two accelerators manage the WCCP control channel on behalf of all the accelerators. Data traffic is load-balanced across all the accelerators. When GRE encapsulation is used, the NetScaler instance performs GRE encapsulation/decapsulation on the data stream between itself and the router, allowing the data between the NetScaler and the accelerators to use a decapsulated, Level-2 configuration.
- Inline mode operates much the same as WCCP mode internally, but externally the appliance emulates a bridge, and no WCCP control channel is established. A packet that enters the appliance on one bridge port exits through the other bridge port. SD-WAN 4000 and 5000 appliances have multiple bridges to support multiple inline links.
- In virtual inline mode (used when WCCP and inline modes are not feasible), the appliance is deployed in a one-arm configuration, much like WCCP, but without the WCCP control channel. Traffic is sent to the appliance from the router, using policy-based routing (PBR) rules. The appliance processes the traffic and returns it to the router.

Figure 3. WCCP and virtual inine cabling

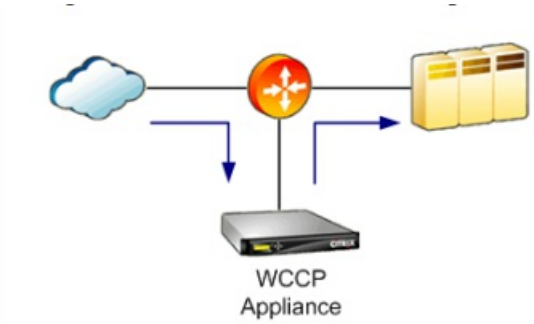

See SD-WAN 4100/5100 virtual machines, internal networks, and external port usage for a diagram of port usage on SD-WAN 4100/5100 appliances. Traffic ports are arranged as a set of accelerated bridges, while the management ports are independent. Typically only one management port is used.

Figure 4. Inline cabling

SD-WAN 4100/5100 appliances have multiple accelerated bridges. Different models have different numbers and types of bridge ports. The two ports making up such a bridge are called an "accelerated pair." All current models include a built-in network bypass function. (Some older SD-WAN 4100-500 and 4100-1000 units do not include network bypass). The network bypass function (also called "fail to wire") connects pairs of ports together if the appliance fails as a result of either power loss or software failure (as determined by an internal watchdog timer).

Inline deployment. The bypass function allows SD-WAN 4100/5100 to be deployed in line with your WAN, typically between your LAN and your WAN router, without introducing a point of network failure.

The accelerated bridges support either 1 Gbps or 10 Gbps data rates. Ethernet and SFP+ interfaces are supported, depending on model.

One-arm deployment. One-arm deployments are also supported, using WCCP or virtual inline modes. With such deployments, a SD-WAN 4000/5000 traffic port is usually connected directly to a port on the WAN router. The other port on the bridged pair is left unconnected.

Performance considerations. Inline deployments provide higher performance than the one-arm deployments, because the use of two ports instead of one doubles the peak throughput of the interfaces.

Peak throughput is important with SD-WAN 4100/5100 appliances, because the compressor provides acceleration in proportion to the compression ratio. That is, a connection that achieves 100:1 compression transfers data one hundred times faster than an uncompressed connection, provided that the rest of the network path can keep up.

For example, take a datacenter with a 500 Mbps WAN link and a 1 Gbps LAN. The small 2:1 speed ratio between the WAN and LAN allows compression to provide only a 2x speedup on a whole-link basis, because there is no way to get data onto or off of the LAN at speeds above 1 Gbps. A 10 Gbps LAN, which allows a tenfold increase in peak data rates, is recommended for use with SD-WAN 4100/5100 deployments.

When a SD-WAN 4100/5100 appliance is deployed in a one-arm mode, the peak transfer rate is cut in half. A SD-WAN

4100/5100 in one-arm mode, connected to the router with a 1 Gbps LAN interface, saturates this interface when the WAN is running at full speed in both directions. For good performance, SD-WAN 4100/5100 must have a LAN interface that is much faster than the WAN. When the appliance is connected directly to the router in a one-arm mode, use a 10 Gbps router port.

> **Note**
>
> The 10 Gbps ports support 10 Gbps only. They do not negotiate lower speeds. Use the 1 Gbps ports for 1 Gbps networks.

A SD-WAN 4100/5100 appliance has at least two non-accelerated ports. Port 0/1 is typically used for management, Port 0/2 is present but typically not used. A Light Out Management (LOM) port is also provided. An RS-232 port can be used for management.

SD-WAN 4100 WANOP

Aug 31, 2017

Citrix NetScaler SD-WAN 4100 WANOP are 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 48 gigabytes (GB) of memory. The Citrix SD-WAN 4100 WANOP has a bandwidth of 310Mbps, 500Mbps, and 1Gbps, respectively.

The following figure shows the front panel of the Citrix SD-WAN 4100 appliance.

Figure 1. Citrix SD-WAN 4100, front panel

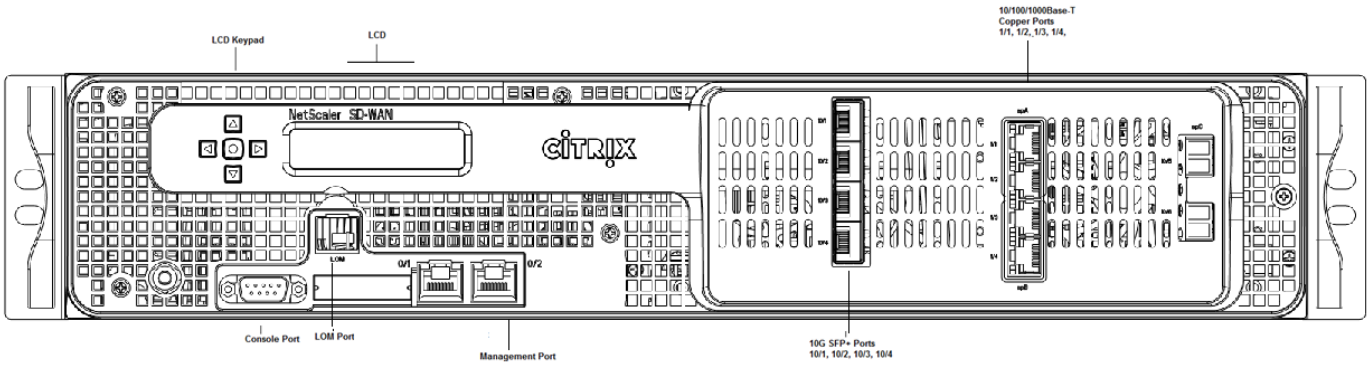

The Citrix SD-WAN 4100 WANOP appliances have the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- RS232 serial console port.
- Network Ports
  - 1x 2-port 10/1G Bypass
  - 1x 4-port 1G Bypass
  - 1x 4-port 10G/1G
  - 1x 2-port 10G (Hidden)

The following figure shows the back panel of the Citrix SD-WAN 4100 WANOP appliance.

Figure 2. Citrix SD-WAN 4100 WANOP back panel

The following components are visible on the back panel of the Citrix SD-WAN 4100 WANOP appliance:

- Four 800 GB removable solid-state drives, which store the appliance's compression history.
- Two 1 TB removable hard disk drives.
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.
- Disable alarm button. This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Two power supplies (either AC or DC). providing full hot swap redundancy.

SD-WAN 5100 WANOP

Aug 31, 2017

Citrix SD-WAN 5100 WANOP is a 2U appliance. Each model has 10-core processor with 2.80GHz and 128 gigabytes (GB) of memory. The Citrix SD-WAN 5100 WANOP appliance has a bandwidth of 2Gbps.

The following figure shows the front panel of the Citrix SD-WAN 5100 WANOP appliance.

Figure 1. Citrix SD-WAN 5100 WANOP, front panel

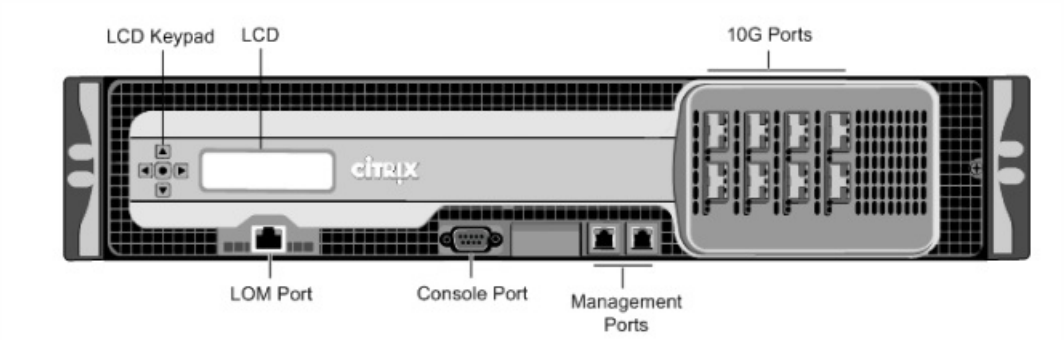

The Citrix SD-WAN 5100 WANOP appliance has the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- RS232 serial console port.
- Ethernet ports:
  - Two 2-port 10/1G bypass, One 4-port 10G/1G, One 2-port 10G.

These ports are used to connect directly to the appliance for system administration functions.

The following figure shows the back panel of the Citrix SD-WAN 5100 WANOP appliance.

Figure 2. Citrix SD-WAN 5100 WANOP, back panel

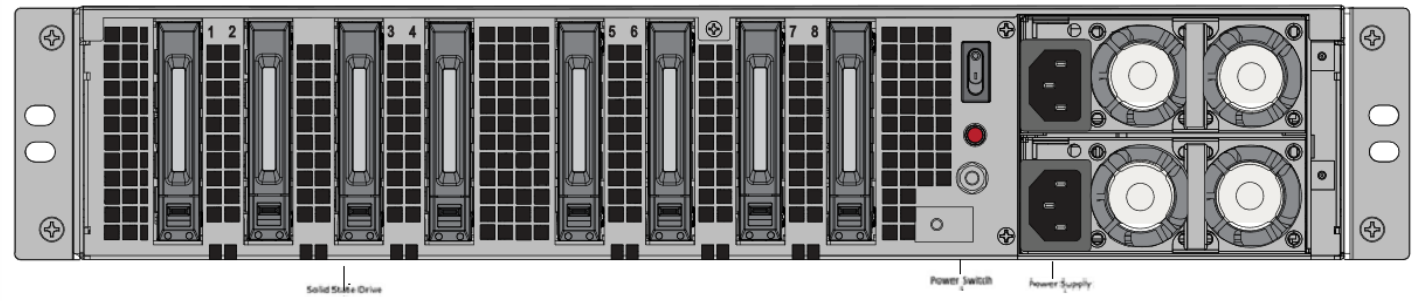

The following components are visible on the back panel of the Citrix SD-WAN 5100 WANOP appliance:

- Six 800 GB removable solid-state drives, which store the appliance's compression history.
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.

- Two 1 TB removable hard disk drive.
- Disable alarm button. This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Two power supplies (either AC or DC). providing full hot swap redundancy. Each power supply has an LED indicating its status.

Field Replaceable Units

Aug 31, 2017

Citrix SD-WAN 4100/5100 WANOP field replaceable units (FRU) are components that can be quickly and easily removed from the appliance and replaced by the user or a technician at the user's site. The FRUs in a Citrix SD-WAN 4100/5100 WANOP appliance can include DC or AC power supplies, and solid-state and hard-disk drives.

> ## Note
>
> By default the appliance ships with AC power supplies. DC power supply is orderable.

Power Supply

Aug 31, 2017

Citrix SD-WAN 4000/5000 appliances are configured with dual power supplies but can operate with only one power supply. The second power supply serves as a backup.

For power-supply specifications, see "Hardware Platforms," which describes the various platforms and includes a table summarizing the hardware specifications.

Table 1. LED Power Supply Indicators

Power Supply Type	LED Color	LED Indicates
AC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing GREEN	Power supply is in standby mode.
	GREEN	Power supply is functional.
	RED	Power supply failure.
DC	OFF	No power to any power supply.
	Flashing RED	No power to this power supply.
	Flashing BLUE	Power supply is in standby mode.
	BLUE	Power supply is functional.
	RED	Power supply failure.

- Make sure that the appliance has a direct physical connection to earth ground during normal use. When installing or repairing an appliance, always connect the ground circuit first and disconnect it last.
- Always unplug any appliance before performing repairs or upgrades.
- Never touch a power supply when the power cord is plugged in. As long as the power cord is plugged in, line voltages are present in the power supply even if the power switch is turned off.

Replace an AC power supply with another AC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace an AC power supply on a Citrix SD-WAN 4000/5000 appliance

1. Align the semicircular handle perpendicular to the power supply. Loosen the thumbscrew and press the lever toward the handle and pull out the existing power supply, as shown in the following figure.
   Figure 1. Removing the Existing AC Power Supply

2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot and press against the semicircular handle until you hear the power supply snap into place.
   Figure 2. Inserting the Replacement AC Power Supply

5. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Replace a DC power supply with another DC power supply. All power supplies must be of the same type (AC or DC).

Note: You can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace a DC power supply on a Citrix SD-WAN 4000/5000 appliance

1. Loosen the thumbscrew and press the lever towards the handle and pull out the existing power supply, as shown in the following figure.
   Figure 3. Removing the Existing DC Power Supply

2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot while pressing the lever towards the handle. Apply firm pressure to insert the power supply firmly into the slot.
   Figure 4. Inserting the Replacement DC Power Supply

5. When the power supply is completely inserted into its slot, release the lever.
6. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

Note: SD-WAN 4000/5000 appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to an appliance in which two power supplies are installed. To silence the alarm, press the small red button on the back panel of the appliance. The disable alarm button is functional only when the appliance has two power supplies.

Solid-State Drive

Aug 31, 2017

A solid-state drive (SSD) is a high-performance device that stores data in solid-state flash memory.

The SD-WAN 4000/5000 software is stored on the solid-state drive (SSD).

To replace a solid-state drive

1. Shutdown the appliance.
2. Locate the SSD on the back panel of the appliance. Push the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Figure 1. Removing the Existing Solid-State Drive

3. Verify that the replacement SSD is the correct type for the platform.
4. Pick up the new SSD, open the drive handle fully to the left or up, and insert the drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top if the drive is inserted horizontally or at the right if the drive is inserted vertically.
   Figure 2. Inserting the Replacement Solid-State Drive

5. Turn on the appliance.
6. Log on to the default IP address by using a web browser, or connect to the serial console by using a console cable, to perform the initial configuration.

Hard Disk Drive

Aug 31, 2017

The NetScaler and SD-WAN virtual machines are hosted on the hard-disk drive.

Verify that the replacement hard disk drive is the correct type for the SD-WAN 4000/5000 platform.

To install a hard disk drive

1. Shut down the appliance.
2. Locate the hard disk drive on the back panel of the appliance.
3. Disengage the hard disk drive by pushing the safety latch of the drive cover to the right or down, depending on the platform, while pulling out on the drive handle to disengage. Pull out the faulty drive.
   Figure 1. Removing the Existing Hard Disk Drive

4. Pick up the new disk drive, open the drive handle fully to the left, and insert the new drive into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the hard drive locks securely into the slot.
   Important: When you insert the drive, make sure that the Citrix product label is at the top.
   Figure 2. Inserting the Replacement Hard Disk Drive

5. Turn on the appliance.

Supported Features

Apr 06, 2018

Features	Citrix SD-WAN 4100	Citrix SD-WAN 5100
AutoConfiguration	N	N
SD-WAN Connector	Y	Y
SD-WAN Plug in	Y	Y
Compression	Y	Y
RPC Over HTTPS	Y	Y
SSL Compression	Y	Y
TCP Acceleration	Y	Y
Traffic Shaping	Y	Y
Video Caching	N	N
Windows File System Acceleration	Y	Y
Windows Outlook Acceleration	Y	Y
XenApp/XenDesktop Acceleration	N	N
Group Mode	N	N
High Availability Mode	Y	Y
Inline Mode	Y	Y
Virtual Inline Mode	Y	Y
WCCP Mode	Y	Y
VLANs	Y	Y

Summary of Hardware Specifications

Aug 31, 2017

The following tables summarize the specifications of the Citrix NetScaler SD-WAN 4100/5100 WANOP hardware platforms.

Specifications	4100 WANOP	5100 WANOP
Bandwidth	Up to 1 Gbps	Up to 2 Gbps
Regulatory model number	2U1P1B	2U1P1D
Processors	2 X 6 Core 2.60GHz	2 X 10 Core 2.80GHz
HDD	2x 1TB HDD boot drives in RAID 1 (mirroring) mode	2x 1TB HDD boot drives in RAID 1 (mirroring) mode
SSD	4x 800GB	6x 800GB
Memory	96 GB	128 GB
Number of power supplies	2 power supplies providing full hot swap redundancy	2 power supplies providing full hot swap redundancy
AC power supply	100 to 240Vac, 50 to 60Hz, 2 x 9.0 to 4.5A	100 to 240Vac, 50 to 60Hz, 2 x 9.0 to 4.5A
DC power supply	-36Vdc to -72Vdc, 2 x 25.5 to 13.0A	-36Vdc to -72Vdc, 2 x 25.5 to 13.0A
Maximum AC power consumption	633W	822W
Maximum DC power consumption	712W	895W
Airflow (front to rear)	65 CFM, typical	65 CFM, typical
Heat Dissipation	137 W/FT 2/FT, typical	144 W/FT 2/FT, typical
Package weight (lbs.) Shipping dimensions and weight	62 lbs (28.1 kgs)	64 lbs (29.10 kgs)
Dimensions	36.5" x 24.5" by 11" (93 cm x 63 cm x 28cm)	36.5" x 24.5" by 11" (93 cm x 63 cm x 28cm)

Specifications	4100 WANOP	5100 WANOP
System weight (lbs.)	45 lbs (20.4 kg)	47 lbs (27 kg)
Rack Units	2U	2U
Width	EIA 310-D, IEC 60297, DIN 41494 SC48D rack 17.25" (44 cm)	EIA 310-D, IEC 60297, DIN 41494 SC48D rack 17.25" (44 cm)
Depth	28" (71.1 cm)	28" (71.1 cm)
Operating temperature	32 – 104 F (0 – 40 C)	32 – 104 F (0 – 40 C)
Non-operating temperature	14F to 140F (-10C to 60C)	14F to 140F (-10C to 60C)
Humidity range (non-condensing)	5%-95% non-condensing	5%-95% non-condensing
Safety certifications	IEC 60950-1, 2nd Edition CSA 60950-1, 2nd Edition UL 60950-1, 2nd Edition AS/NZS 6050-1	IEC 60950-1, 2nd Edition CSA 60950-1, 2nd Edition UL 60950-1, 2nd Edition AS/NZS 6050-1
EMC & susceptibility	US (FCC (Part 15 Class A)) Europe (CE (EN55022/55024)) Australia (RCM), Japan (VCCI), Korea (KCC), Taiwan (BSMI), China (CCC), India (BIS), Russia (EAC), Saudi Arabia (CITC), Brazil (Anatel), South Africa (ICASA), Mexico (NOM), Egypt (NTRA), Israel (MoC)	US (FCC (Part 15 Class A)) Europe (CE (EN55022/55024)) Australia (RCM), Japan (VCCI), Korea (KCC), Taiwan (BSMI), China (CCC), India (BIS), Russia (EAC), Saudi Arabia (CITC), Brazil (Anatel), South Africa (ICASA), Mexico (NOM), Egypt (NTRA), Israel (MoC)
Environmental compliance	RoHS, REACH, WEEE	RoHS, REACH, WEEE

Lights Out Management Port of the SD-WAN WANOP 4100/5100 Appliance

Aug 31, 2017

The SD-WAN 4100/5100 appliances have an Intelligent Platform Management Interface (IPMI), also known as the Lights out Management (LOM), port on the front panel of the appliance. By using the LOM, you can remotely monitor and manage the appliance, independently of the SD-WAN 4100/5100 software. You can remotely change the IP address, perform different power operations, and obtain health monitoring information of the appliance by connecting to the appliance through the LOM port.

By connecting the LOM port over a dedicated channel that is separate from the data channel, you can make sure that connectivity to the appliance is maintained even if the data network is down.

By using a web browser you can remotely log on to the LOM port to obtain information about the appliance and perform different operations on the appliance.

To access the LOM by using a web browser

1. In a web browser, type the IP address of the LOM port. For initial configuration, type the port's default address: http://192.168.1.3
2. In the User Name box, type **nsroot**.
3. In the Password box, type **nsroot**.

You can use the Intelligent Platform Management Interface (IPMI), also known as the Lights Out Management (LOM) port, to remotely monitor and manage the appliance, independently of the NetScaler software. For initial configuration of the lights-out management (LOM) port, connect to the port's default IP address and change it to the address that you want to use for remote monitoring and management. Also specify the administrator credentials and the network settings.

Note: The LEDs on the LOM port are unoperational by design.

To configure the NetScaler LOM Port

1. Connect the LOM port to a management workstation or network.
2. In a web browser, type: http://192.168.1.3.
   Note: The NetScaler LOM port is preconfigured with the IP address 192.168.1.3 and subnet mask 255.255.255.0.
3. In the User Name box, type **nsroot**.
4. In the Password box, type **nsroot**.
5. On the Configuration tab, click Network and type values for the following parameters:
   - IP Address—IP address of the LOM port.
   - Subnet Mask—Subnet mask used to define the subnet of the LOM port.
   - Default Gateway—IP address of the router that connects the LOM port to the network.
6. Click Save.

You can remotely turn off the appliance and turn it back on. The result is similar to pressing the power button on the back

panel of the appliance for less than two seconds.

To power cycle the appliance

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click Remote Control.
4. Under Options, click Power Control, and then click Power Cycle System.
5. Click Perform Action.

The LOM port allows you to remotely access and manage the appliance by logging on to a redirected console.

To access the appliance by using the access console

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click Remote Control.
4. Under Options, click Console Redirection.
5. Click Launch Console, and then click Yes.
6. Type the administrator credentials for the appliance.

You can log on to the LOM port to view the health information about the appliance. All system sensor information, such as system temperature, CPU temperature, status of fan and power supplies, appears on the sensor readings page.

To obtain health monitoring information

1. In a web browser, type the IP address of the LOM port.
2. In the User Name and Password boxes, type the administrator credentials.
3. In the Menu bar, click System Health.
4. Under Options, click Sensor Readings.

You can remotely perform different power control operations, such as restarting the appliance, performing a graceful shutdown, and performing a forced shutdown, by using the LOM port.

To perform power control operations

1. In a web browser, log on to the LOM port by using the administrator credentials.
2. In the Menu bar, click Remote Control.
3. Under Options, click Power Control, and then select one of the following options:
   - **Reset System**—Restart the appliance.
   - **Power Off System – Immediate**—Disconnect power to the appliance without shutting down the appliance.
   - **Power On System**—Turn on the appliance.
   - **Power Cycle System**—Turn off the appliance, and then turn it back on.
4. Click Perform Action.

Preparing for Installation

Aug 31, 2017

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Once you are satisfied that your appliance has been delivered to your expectations, verify that the location where the appliance will be installed meets temperature and power requirements and that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel should install, maintain, or replace the appliance, and efforts should be taken to ensure that all cautions and warnings are followed.

Unpacking the Appliance

Aug 31, 2017

Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

Use the following list to verify that you received everything that should have been included in the box.

- The appliance you ordered
- One RJ-45 to DB-9 adapter
- One 6 ft RJ-45/DB-9 cable
- Two power cables
- One fiber patch cable
- One standard 4-post rail kit

In addition to the items included in the box with your new appliance, you will need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you will connect to your network
- One available Ethernet port on your network switch or hub for each Ethernet port you want to connect to your network
- A computer to serve as a management workstation

> **Note**
>
> If the kit that you received does not fit your rack, contact your Citrix sales representative to order the appropriate kit.

Preparing the Site and Rack

Aug 31, 2017

There are specific site and rack requirements for the SD-WAN 4100/5100 appliance. You must make sure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and help ensure a smooth installation.

The appliance should be installed in a server room or server cabinet with the following features:

Environment control

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment.

Power density

Wiring capable of handling at least 4,000 watts per rack unit in addition to power needs for the CRAC.

The rack on which you install your appliance should meet the following criteria:

Rack characteristics

Racks should be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it should be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks should be installed in rows spaced at least 1 meter/3 feet apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

Power connections

At minimum, two standard power outlets per unit.

Network connections

At minimum, Ethernet connection per rack unit.

Space requirements

Two empty rack units for SD-WAN 4100/5100 appliances.

You can order the following rail kits separately.

- Compact 4-post rail kit, which fits racks of 23 to 33 inches.
- 2-post rail kit, which fits 2-post racks.

Cautions and Warnings

Aug 31, 2017

> ## Warning
>
> During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.

Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Remove all jewelry and other metal objects that might come into contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.
- Use a regulating, uninterruptible power supply (UPS) to protect the appliance from power surges and voltage spikes, and to keep the appliance operating in case of power failure.
- Never stack the appliance on top of any other server or electronic equipment.
- All appliances are designed to be installed on power systems that use TN earthing. Do not install your device on a power system that uses either TT or IT earthing.
- Make sure that the appliance has a direct physical connection to the earth during normal use. When installing or repairing an appliance, always make sure that the ground circuit is connected first and disconnected last.
- Make sure that a fuse or circuit breaker no larger than 120 VAC, 15 A U.S. (240 VAC, 16 A international) is used on all current-carrying conductors on the power system to which your appliances are connected.
- Do not work alone when working with high voltage components.
- Always disconnect the appliance from power before removing or installing any component. When disconnecting power, first shut down the appliance, and then unplug the power cords of all the power supply units connected to the appliance. As long as the power cord is plugged in, line voltages can be present in the power supply, even when the power switch is OFF.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload. Always unplug any appliance before performing repairs or upgrades.
- Do not overload the wiring in your server cabinet or on your server room rack.
- During thunderstorms, or anticipated thunderstorms, avoid performing any hardware repairs or upgrades until the danger of lightning has passed.
- When you dispose of an old appliance or any components, follow any local and national laws on disposal of electronic waste.
- To prevent possible explosions, replace expired batteries with the same model or a manufacturer-recommended substitute and follow the manufacturer's instructions for battery replacement.
- Never remove a power supply cover or any sealed part that has the following label:

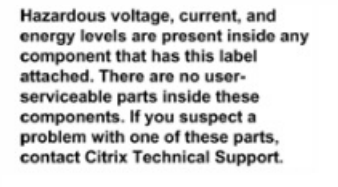

- Determine the placement of each component in the rack before you install the rails.
- Install the heaviest appliance first, at the bottom of the rack, and then work upward. Distribute the load on the rack evenly. An unbalanced rack is hazardous.
- Allow the power supply units and hard drives to cool before touching them.
- Install the equipment near an electrical outlet for easy access.
- Mount equipment in a rack with sufficient airflow for safe operation.
- For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

Aug 31, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 31, 2017

Most appliances can be installed in standard server racks that conform to EIA-310-D specification. The appliances ship with a set of rails, which you must install before you mount the appliance. The only tools that you need for installing an appliance are a Phillips screwdriver and a flathead screwdriver.

Caution: If you are installing the appliance as the only unit in the rack, mount it at the bottom. If the rack contains other units, make sure that the heaviest unit is at the bottom. If the rack has stabilizing devices available, install them before mounting the appliance.
Citrix SD-WAN 4000/5000 appliance requires two rack units.

Each appliance ships with a mounting rail kit that contains two rail assemblies, one for the left side and the other for the right side of the appliance, and screws to attach the rails. An assembly consists of an inner rail and a rack rail. The supplied rail kit is 28 inches long (38 inches extended). Contact your Citrix sales representative to order a 23-inch (33 inches extended) rail kit.
Note: The same rail kit is used for both square-hole and round-hole racks. See figure 4 for specific instructions for threaded, round-hole racks.
To mount the appliance, you must first install the rails and then install the appliance in the rack.

Perform the following tasks to mount the appliance:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the latch until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

1. Position the right inner rail behind the handle on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws: 5 per side, as shown in the following figure.
4. Repeat steps 1 through 3 to install the left inner rail on the other side of the appliance.
   Figure 1. Attaching inner rails

1. If you have a round-hole, threaded rack, skip to step 3.
2. Install square nut retainers into the front post and back post of the rack as shown in the following figures. Before inserting a screw, be sure to align the square nut with the correct hole for your appliance. The three holes are not evenly spaced.
   Figure 2. Installing Retainers into the Front Rack Posts

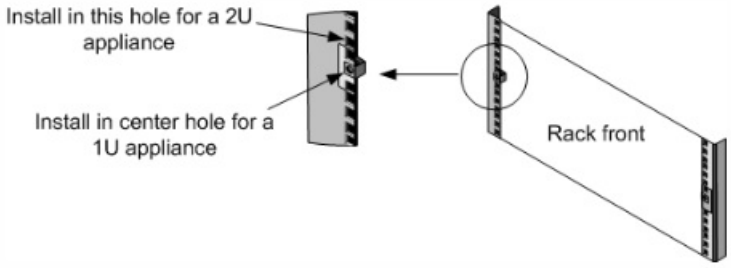

Figure 3. Installing Retainers into the Rear Rack Posts

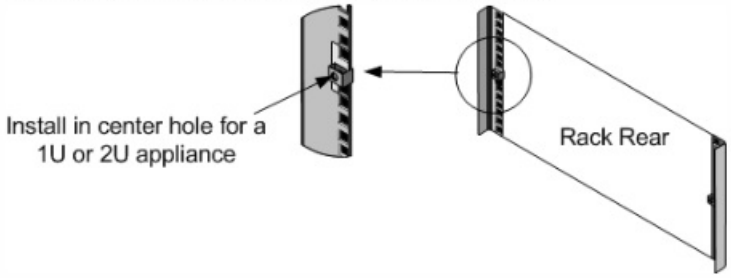

3. Install the adjustable rail assembly into the rack as shown in the following figures. Use a screw to lock the rear rail flange into the rack. With the screw securing the rail in place, you can optionally remove the latching spring.

Figure 4: Assembling the Rack

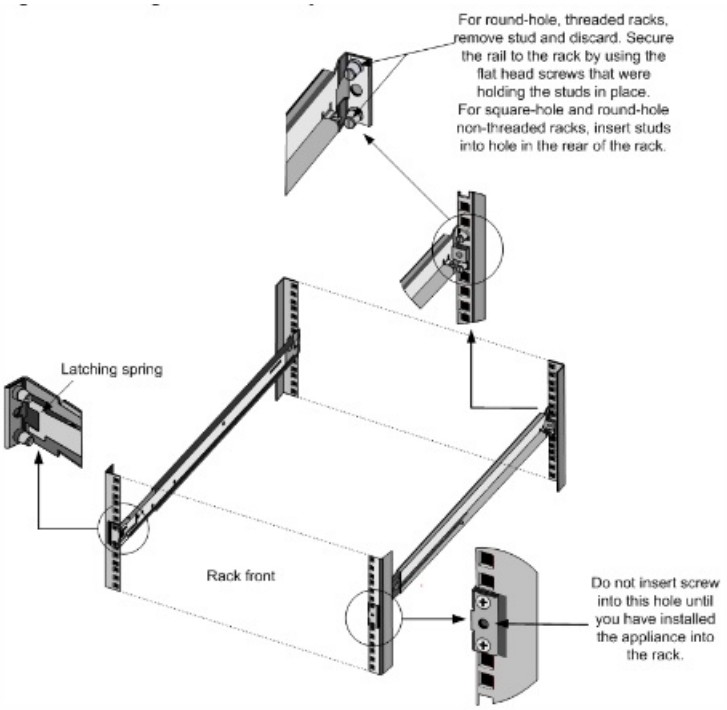

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.

Figure 5. Rack Mounting the Appliance

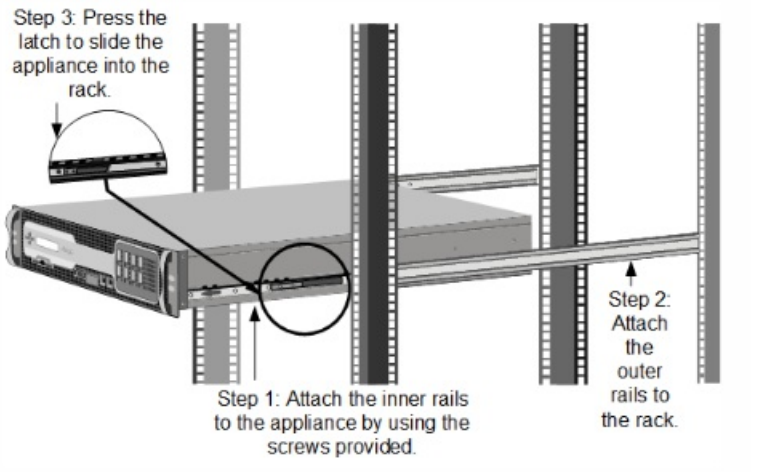

Installing and Removing 1G SFP Transceivers

Aug 31, 2017

Note: Some SD-WAN 4000/5000 appliances do not require SFP transceivers.

A Small Form-Factor Pluggable (SFP) is a compact transceiver that can operate at speeds of up to 1 gigabit per second and is available in both copper and fiber types. Inserting a 1G SFP copper transceiver converts the 1G SFP port to a 1000BASE-T port. Inserting a 1G SFP fiber transceiver converts the 1G SFP port to a 1000BASE-X port. Auto-negotiation is enabled by default on the 1G SFP port into which you insert your 1G SFP transceiver. As soon as a link between the port and the network is established, the speed and mode are matched on both ends of the cable.

Caution: SD-WAN 4000/5000 appliances do not support 1G SFP transceivers from vendors other than Citrix Systems. Attempting to install third-party 1G SFP transceivers on your SD-WAN 4000/5000 appliance voids the warranty.

Insert 1G SFP transceivers into the 1G SFP ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the 1G SFP transceiver or the appliance.

Caution: Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

1. Remove the 1G SFP transceiver carefully from its box.
   Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Align the 1G SFP transceiver to the front of the 1G SFP transceiver port on the front panel of the appliance, as shown in the following figure.
3. Hold the 1G SFP transceiver between your thumb and index finger and insert it into the 1G SFP transceiver port, pressing it in until you hear the transceiver snap into place.
4. Lock the transceiver.
5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. If you are using a fiber 1G SFP transceiver, do not remove the dust caps attached to the transceiver and the cable until you are ready to insert the cable.

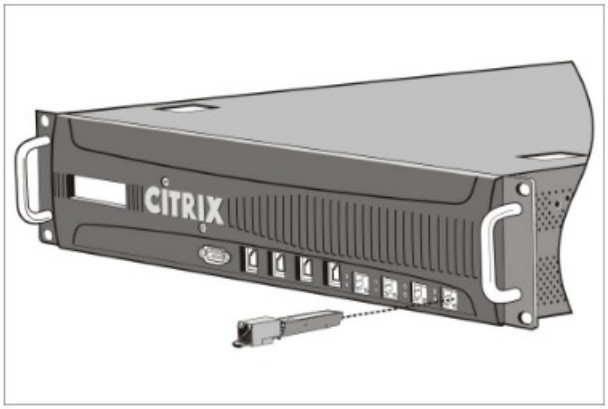

Note

The illustration in the following figures might not represent your actual appliance.

1. Disconnect the cable from the 1G SFP transceiver. If you are using a fiber optic cable, replace the dust cap on the cable before putting it away.
   Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 1G SFP transceiver.
3. Hold the 1G SFP transceiver between your thumb and index finger and slowly pull it out of the port.
4. If you are removing a fiber 1G SFP transceiver, replace the dust cap before putting it away.
5. Put the 1G SFP transceiver into its original box or another appropriate container.

Installing and Removing 10G SFP+ Transceivers

Aug 31, 2017

Warning

Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

Note

Some SD-WAN 4100/5100 appliances do not require SFP+ transceivers.

A 10-Gigabit Small Form-Factor Pluggable (SFP+) is a compact optical transceiver that can operate at speeds of up to 10 gigabits per second. Autonegotiation is enabled by default on the 10G SFP+ ports into which you insert your 10G SFP+ transceiver. As soon as a link between the port and the network is established, the mode is matched on both ends of the cable and for 10G SFP+ transceivers, the speed is also autonegotiated.
Insert the 10G SFP+ transceivers into the 10G SFP+ ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the transceiver or the appliance.

Important

SD-WAN 4100/5100 appliances do not support 10G SFP+ transceivers provided by vendors other than Citrix Systems. Attempting to install third-party 10G SFP+ transceivers on your SD-WAN 4100/5100 appliance voids the warranty.

1. Remove the 10G SFP+ transceiver carefully from its box.
2. Align the 10G SFP+ transceiver to the front of the 10G SFP+ transceiver port on the front panel of the appliance.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and insert it into the 10G SFP+ transceiver port, pressing it in until you hear the transceiver snap into place.
4. Move the locking hinge to the DOWN position.
5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. Do not remove the dust caps attached to the transceiver and cable until you are ready to insert the cable.

Warning

Do not look directly into fiber optic transceivers and cables. They emit laser beams that can damage your eyes.

1. Disconnect the cable from the 10G SFP+ transceiver. Replace the dust cap on the cable before putting it away.
2. Unlock the 10G SFP+ transceiver by moving the locking hinge to the UP position.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and slowly pull it out of the port.
4. Replace the dust cap on the transceiver before putting it away.
5. Put the 10G SFP+ transceiver into its original box or another appropriate container.

Install Fiber Patch Cable in Ports 10/3 and 10/4

Aug 31, 2017

Through release 9.3, on an appliance, SD-WAN ports 10/3 and 10/4 must be connected with the provided cable, as shown in the following figure.

Starting with release 9.3, the patch cable is no longer required, and can be omitted if:

- The appliance was shipped from the factory with release 9.3 or later, or
- The appliance was shipped from the factory with release 9.3 or earlier, but you upgrade it to later version and change the default loopback in the management service (on **System > Configuration > System > Configure Loopback Settings**).

If you decide to eliminate the need to use loopback cable, the ports 10/3 and 10/4 are still reserved. These ports are not available for WAN optimization.

To install the patch cable

1. Connect the LC-to-LC cable to the ports as shown in the figures above.
2. Insert one end of the cable into port 10/3.
3. Insert the other end of the cable into port 10/4.

Connecting the Cables

Aug 31, 2017

When the appliance is securely mounted on the rack, you are ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Danger: Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port on the front panel of the appliance, as shown in the following figure.
2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Figure 1. Inserting an Ethernet cable

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port that is located on the front panel of the appliance, as shown in the following figure.
2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Figure 2. Inserting a console cable

> **Note**
>
> To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.

The SD-WAN 4000/5000 appliance has two power supplies, with one serving as a backup. A separate ground cable is not required, because the three-prong plug provides grounding. Power up the appliance by installing one or both power cords.

To connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply, as shown in the following figure.
2. Connect the other end of the power cable to a standard 110V/220V power outlet.
3. Repeat steps 1 and 2 to connect the second power supply.

Figure 3. Inserting a power cable

> **Note**
>
> The appliance emits a high-pitched alert if one power supply fails or if you connect only one power cable to the appliance. To silence the alarm, you can press the small red button located on the back panel of the appliance.

Switching on the Appliance

Aug 31, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the back panel of the appliance.

> ## Warning
>
> Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Planning the Deployment

Aug 31, 2017

SD-WAN 4100/5100 deployments require adequate planning, especially for units deployed in large datacenters:

- An appropriate appliance or group of appliances must be selected to support both the current and anticipated load.
- A deployment mode must be selected to match the requirements of your site.
- Other aspects must also be considered.

Sizing Guidelines

Mar 20, 2018

For successful deployment of one or more SD-WAN 4100/5100 appliances in your datacenter, keep the following principles in mind:

- You must provide enough SD-WAN 4100/5100 peak-load capacity, in terms of WAN bandwidth and the number of users. See the current specifications sheet for the capacities of different SD-WAN 4100/5100 models. Ensure adequate peak-load capacity, both for now and for the time until you expect to upgrade. Acceleration is resource intensive, and performance suffers if the appliance runs short of resources. Never overcommit any SD-WAN appliance, especially in the datacenter. Provision your datacenter to easily accommodate peak loads.
- Provide enough capacity for expected expansion over the life of the deployment. SD-WAN 4100/5100 appliances using the same hardware platform can have their capacity upgraded with a new license as part of the Citrix pay-as-you-grow program. SD-WAN 4100/5100 models 310, 500, and 1000 use one hardware platform, and models 1500 and 2000 use another hardware platform. This means that, for example, a SD-WAN 4100/5100 500 can be converted through a license upgrade to a SD-WAN 4100/5100 1000, but not to a SD-WAN 4100/5010 1500.
- For more capacity than can be provided by a single appliance, multiple SD-WAN 4100/5100 appliances can be cascaded behind a stand-alone NetScaler appliance.
- Different models have differing numbers of traffic ports. If you require multiple bridges, make sure your model has at least as many as you need.

Selecting a Deployment Mode

Aug 31, 2017

The SD-WAN 4100/5100 appliance can be deployed inline or in a one-arm mode. Inline deployments do not require router reconfiguration; one-arm modes do. SD-WAN 4100/5100 offers internal port bypassing (fail-to-wire) to allow traffic to continue flowing in inline mode if the appliance fails.

> **Note**
>
> Only the one-arm WCCP mode (with a single router) is documented at this time. Inline mode is not yet documented. Citrix recommends WCCP mode at this time.

Different SD-WAN 4100/5100 models offer different numbers of accelerated bridges. Models with multiple accelerated bridges can accelerate multiple inline WAN links. See the specifications sheet for more details, http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/SD-WAN-data-sheet.pdf.

A standalone SD-WAN 4100/5100 appliance can be deployed in either of these two recommended modes:

- Inline, bridged (L2 inline). This closely resembles a standard SD-WAN inline deployment. Packets enter one bridge port and exit the other bridge port.
- One-arm, WCCP. This resembles a standard SD-WAN WCCP deployment.

Citrix also supports the following two modes (which are outside the scope of this document):

- Inline, routed. The NetScaler instance uses routing rules instead of bridging rules to determine how to forward packets.
- Virtual inline. This resembles WCCP, but lacks built-in health-checking.

In L2 inline mode, SD-WAN 4100/5100 is placed between your LAN and your WAN router (or other aggregation point at the LAN-WAN boundary). In a one-arm mode, SD-WAN 4100/5100 is generally connected directly to a dedicated port on your WAN router.

In cases where the WAN router ports are not as fast as the LAN (for example, when the WAN router has gigabit Ethernet, but the LAN has10 gigabit Ethernet), inline mode provides better performance, because its LAN-side traffic is not limited to the speed of the router interface. (Compression allows the LAN-side traffic to be much faster than WAN-bound traffic under favorable conditions.)

Considerations:

- The inline modes require no reconfiguration of your routers, but involves a service disruption when bringing the appliance into service.
- One-arm modes require router reconfiguration but do not require a service disruption.
- Inline mode has higher performance than the other modes.
- One-arm modes are limited to half the speed of the router or switch port they are attached to.
- With WCCP mode, configuring the router to send only a fraction of the WAN traffic to SD-WAN 4100/5100 (as little as the traffic from a single remote site or even a single remote IP address) makes it easy to bring up and test the installation gradually. Inline mode requires that all WAN traffic pass through the appliance.
- WCCP mode requires more configuration of the SD-WAN 4100/5100 appliance than do other modes, but is more

standardized and provides more status information on the router.

Recommendation:

- The greater control provided by WCCP, and especially the ability to put the deployment into service in stages, makes WCCP the mode of choice for larger, more complex datacenters, especially if there might be a possibility of overloading the SD-WAN 4100/5100 appliance.
- Inline mode is convenient for smaller WAN networks and simpler datacenters. It is most commonly used with the SD-WAN 4100/5100 310 and 500, and more rarely with the larger appliances.
- Cascaded installations should use WCCP.

Selecting a Load Balancing Method

Aug 31, 2017

By default, the SD-WAN 4100/5100 Provisioning Wizard sets up load balancing to handle different kinds of connections appropriately. This default behavior is adequate for most installations.

Sending all the connections from the same remote accelerator to the same local accelerator maximizes the benefits of SD-WAN compression, and the default load balancing method accomplishes this. If an instance becomes overloaded or unavailable, new connections are reallocated.

By default, the NetScaler instance uses the least-connection method to balance the load across the accelerators. This method applies whether or not the connections are accelerated. Connections are persistent, but persistency is discontinued for an instance that becomes overloaded, and is lost if the local appliance is restarted or when no traffic from a remote appliance is seen for more than 24 hours.

For incoming accelerated connections (that is, connections with SD-WAN options in the header of the SYN packet), all connections from a given remote SD-WAN are sent to the same local accelerator.

The identity of the remote SD-WAN is determined by one of the SD-WAN SYN options: the "AgentID" field, which contains the management IP address of the remote SD-WAN.

This method is used for connections from remote SD-WAN appliances and remote SD-WAN Plug-ins.

Incoming non-accelerated connections and all outgoing connections are also distributed among the accelerators according to the least-connection method, but since they do not contain an AgentID field, they cannot use AgentID persistence. Instead, they use SRCIPDESTIP persistence, meaning that connections with the same IP addresses use the same accelerator.

If an instance is overloaded, the NetScaler instance bypasses it for new connections, sending them through without acceleration. Existing connections continue to be sent to the instance.

This behavior is controlled by the skipPersistency parameter. The default behavior is -skippersistency ReLB. The alternative behavior, -skippersistency bypass, instructs the NetScaler instance to pass the connection through without sending it to an accelerator.

The default load balancing behavior is adequate for most installations, but sometimes customization is needed. This is most commonly true when a few remote sites have much more traffic than the rest. In that case, it can be worthwhile to assign these large sites to accelerators explicitly.

Optional load balancing behavior includes the use of static routing (for hand-crafted load balancing) and variations on the least-connection with AgentID and SRCIPDESTIP persistence methods used in the default configuration. The behavior for

dealing with overloaded instances can be changed from assigning connections to a difference instance to passing them through as unaccelerated.

Gathering Information Needed for Configuration

Aug 31, 2017

Accurate information about both the local and the remote sites is essential to troubleshooting. Before installing the SD-WAN 4100/5100 appliance, make sure that you have done the following:

1. Obtained or drawn an accurate network diagram of your local site (the one in which you are installing SD-WAN 4100/5100). The local network topology and the capabilities of your WAN routers determine which deployment modes are appropriate for the site.
2. Chosen the deployment mode of the local SD-WAN 4100/5100 appliance (for example, WCCP or inline, with or without HA and cascading).
3. Compiled a list of critical applications that must be tested to validate the deployment.
4. Obtained or drawn an accurate network diagram of your WAN, including both the local and the remote WAN links, their bandwidths in both directions, their subnets, and whether they are accelerated. In deployments with many remote sites, an aggregate of the different categories (accelerated and non-accelerated) is probably sufficient, and only the largest remote sites need to be considered individually.
5. Determined whether there are multiple datacenters with datacenter-to-datacenter traffic, and whether any remote datacenters have a SD-WAN 4100/5100 appliance.
6. Decided whether you plan to increase WAN capacity, the number of sites, or the number of users in the next 24 months. If so, the corresponding SD-WAN 4100/5100 capacity should be installed now.
7. If possible, formed an idea of the traffic breakdown over the WAN, including TCP traffic to and from SD-WAN-accelerated sites, other TCP traffic, ICA users, HDX sessions, and real-time traffic such as VoIP. SD-WAN 4100/5100 needs to be provisioned for the peak loads in terms of accelerated TCP connections, ICA users, and total WAN link capacity.
8. Determined the number of WAN links in the local site. Are they independent, or are they load balanced? If so, are they active-active or active-standby?
9. Determined the current, unaccelerated RTT of the remote sites during peak periods.
10. Identified any QoS devices or proxies in the path between the local and remote sites. QoS devices should be on the WAN side of SD-WAN 4100/5100. Proxies should be on the LAN side.

Initial Configuration

Aug 31, 2017

After checking the connections, you are ready to deploy the SD-WAN 4100 and 5100 appliances on the network.

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

Initial configuration consists of the following tasks:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Assign management IP address through the serial console.

By default, the initial configuration deploys the appliance in inline mode.

Prerequisites

Mar 16, 2018

To deploy a Citrix SD-WAN 4100 or 5100 appliance, you must complete the following prerequisite setup before configuring the appliance.

This document covers release of the SD-WAN software. See the release notes for the recommended versions of the NetScaler software corresponding to the desired release of the SD-WAN software. Never use any versions other than those recommended for SD-WAN 4100 and 5100 appliances.

The number of accelerator appliances depend on the hardware platform and the type of license you apply to the appliance. The following list displays the number of accelerators that gets provisioned automatically by the Configuration Wizard:

- Model 310: Two
- Model 500: Three
- Models 1000 and 1500: Six
- Model 2000: Eight

Before you start provisioning the appliance, Citrix recommends that you have the license file with you, as it is required early in the configuration process To download a license file, complete the procedure described in the *My Account All Licensing Tools* - User Guide.

After you receive the hardware appliance from Citrix, you need to install it in the network. To install the SD-WAN 4100/5100 appliance hardware, follow the installation procedure at Installing the Hardware.

Deployment Worksheet

Aug 31, 2017

> **Note**
>
> Use this worksheet only when provisioning a factory-reset appliance with the release 9.3 configuration wizard. If you are simply upgrading a previously configured system to release 9.3, your appliance will retain its previous configuration, which will be different.

The appliance uses at least two ports: the management port (typically 0/1) and the traffic port (such as 10/1). Inline mode uses traffic ports in pairs, such as ports 10/1 and 10/2. Ports must be selected in advance, because the configuration depends on their identity.

The appliance uses three subnets directly: the management subnet, the external traffic subnet, and the internal traffic subnet. Multiple IP addresses are used on each subnet. Each subnet must be specified along with the correct subnet mask.

The following figure is a worksheet for these parameters. It supports inline and WCCP modes, with and without HA. The table below the figure describes what each entry means.

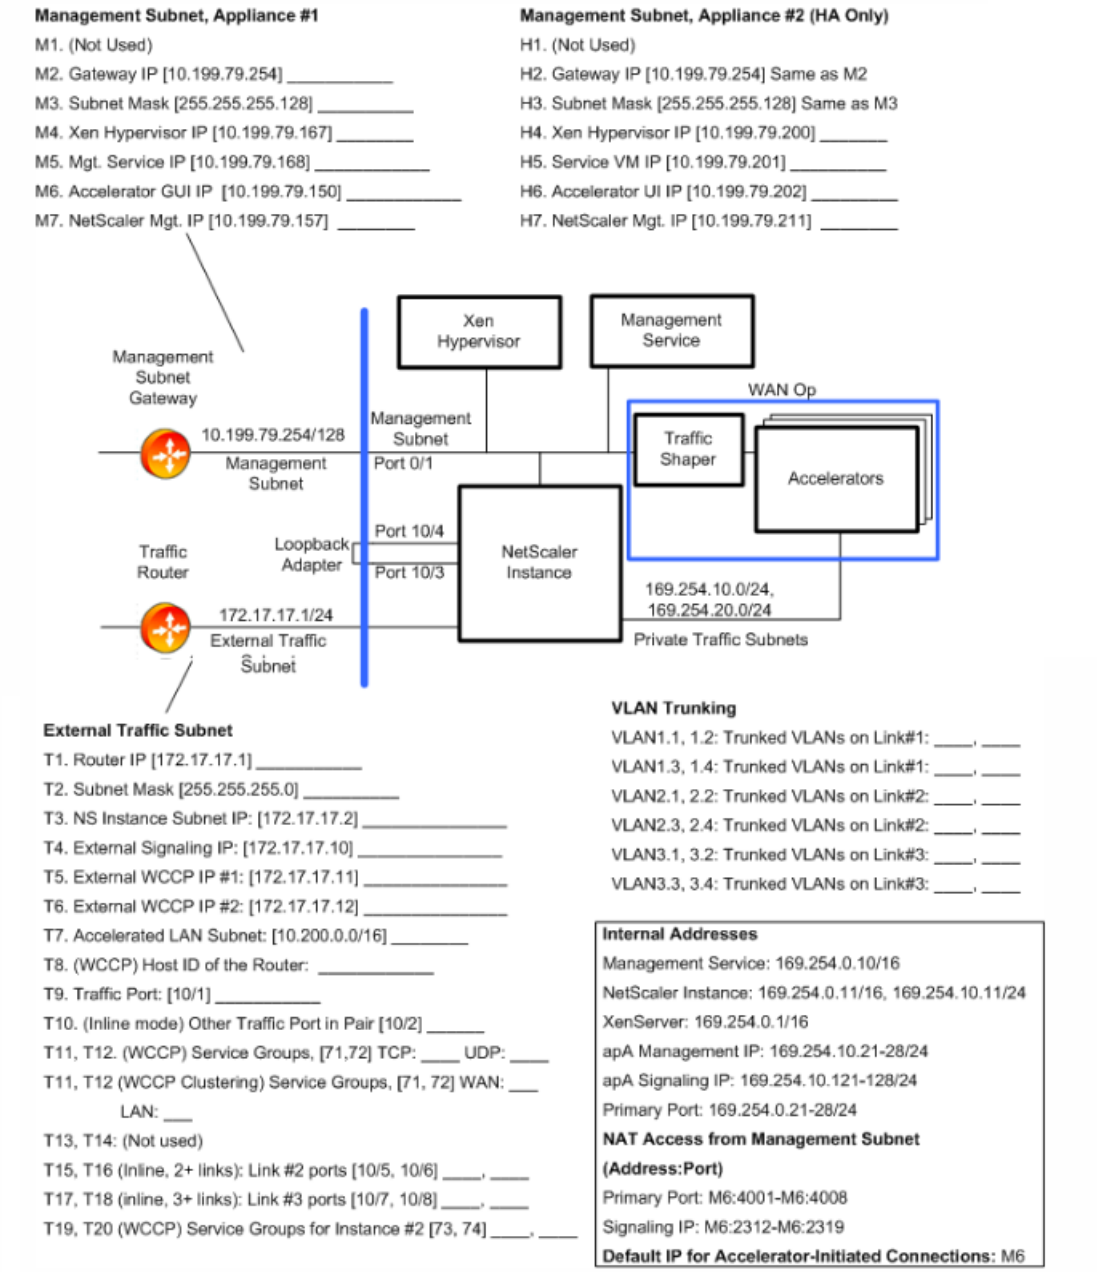

Table 1. Deployment Worksheet Parameters

	Parameter	Example	Your Value	Description
Management Subnet				
M2.	Gateway IP address	10.199.79.254		Default gateway serving the management subnet.
M3.	Subnet Mask	255.255.255.128		Subnet mask for the management subnet.
M4.	Xen Hypervisor IP	10.199.79.225		IP address of Xen Hypervisor.

	Parameter	Example	Your Value	Description
	address			
M5.	Service VM IP address	10.199.79.226		IP address of Management Service VM, which controls configuration.
M6.	Accelerator UI	10.199.79.227		Accelerator GUI, also called the Broker UI, which manages the instances as a unit.
M7.	NetScaler Management IP address	10.199.79.245		IP address of the NetScaler instance's GUI and CLI interfaces.
External Traffic Subnet				
T1.	Router IP address	172.17.17.1		IP address of router on external traffic subnet.
T2.	Subnet Mask	255.255.255.0		Subnet mask of external traffic subnet.
T3.	NetScaler IP address	172.17.17.2		NetScaler IP address on external traffic subnet.
T4.	External Signaling IP address	172.17.17.10		Traffic to this IP address is load-balanced between the signaling IP addresses of the accelerators.
T5.	External WCCP IP address #1	172.17.17.11		Maps through NAT to WCCP VIP on accelerator #1.
T6.	External WCCP IP address #2	172.17.17.12		Maps through NAT to WCCP VIP on accelerator #2.
T7.	Local LAN Subnets	10.200.0.0/16		The local LAN subnet to be accelerated. This is the only subnet that will receive acceleration.
T8.	GRE Router HostID	NA		WCCP-GRE only. Host ID of GRE router.
T9.	Traffic Port	10/1		Port used for accelerated traffic.
T10+.	(Inline) Additional Traffic Port			Other traffic port in pair.
T11, T12	(WCCP) Service Groups: TCP, UDP	71, 72		Service groups used by accelerator #1 for WCCP. First is for TCP traffic, second is for UDP.
T13, T14	(Not used)			
T15, T16	(Inline) Ports used	10/5, 10/6		If multiple links are used with inline mode,

Parameter	Example	Your Value	Description
T17, T18	10/7, 10/8		If multiple links are used with inline mode, these ports are used for link #3.
VLAN1.1, VLAN1.2, VLAN1.3, VLAN1.4	External VLANs for Bridge #1	412	When VLAN trunking is used, these are tagged VLANs crossing bridge #1.
VLAN2.1, VLAN2.2, VLAN2.3, VLAN2.4			When VLAN trunking is used, these are tagged VLANs crossing bridge #2.
VLAN3.1, VLAN3.2, VLAN3.3, VLAN3.4	External VLANs for Bridge #1		When VLAN trunking is used, these are tagged VLANs crossing bridge #3.

Configuring the Appliance

Aug 31, 2017

Before you start configuring the appliance, you must change the IP address of the management service to the one in your management network, so that you can access the appliance over the network. You can change the management IP address by connecting a computer to the appliance through either the Ethernet port or the serial console.

Assigning a Management IP Address through the Ethernet Port

Aug 31, 2017

Use the following procedure for initial configuration of every SD-WAN 1000 or 2000 appliance with Windows Server. The procedure accomplishes the following tasks:

- Configure the appliance for use on your site.
- Install the Citrix license.
- Enable acceleration.
- Enable traffic shaping (inline mode only).

With inline deployments, this configuration might be all you need, because most acceleration features are enabled by default and require no additional configuration.

If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 15 of the following procedure.

Note: You must have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.50, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields blank.

2. Using an Ethernet cable, connect this computer to the port labeled PRI on the SD-WAN appliance.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address, which is http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance:
   Username: nsroot

   Password: nsroot.

5. Start the configuration wizard by clicking **Get Started**.
6. On the **Platform Configuration** page, enter respective values from your worksheet, as shown in the following example:

7. Click **Done**. A screen showing the Installation in Progress... message appears. This process takes approximately 2 to 5 minutes, depending on your network speed.
8. A Redirecting to new management IP message appears.
9. Click **OK**.
10. Unplug your computer from the Ethernet port and connect the port to your management network.
11. Reset the IP address of your computer to its previous setting.
12. From a computer on the management network, log on to the appliance by entering the new management service IP address, such as https://<Managemnt_IP_Address>, in a web browser.

13. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
14. Log on to the appliance by using the **nsroot** user name and the password from your worksheet.
15. To complete the configuration process, see Provisioning the Appliance.

Assigning a Management IP Address through the Serial Port

Aug 31, 2017

If you do not want to change the settings of your computer, you can configure the appliance by connecting it to your computer with a serial null modem cable. You must have physical access to the appliance.

To configure the appliance through the serial console

1. Connect a serial null modem cable to the appliance's console port.
2. Connect the other end of the cable to the serial COM port of a computer running a terminal emulator, such as Microsoft HyperTerminal, with settings 9600,N,8,1, p.
3. In the HyperTerminal output, press **Enter**. The terminal screen displays the logon prompt.
   Note: You might have to press **Enter** two or three times, depending on the terminal program you are using.
4. At the logon prompt, log on to the appliance with the following default credentials:
   Username: nsroot

   Password: nsroot.

5. At the **$** prompt, run the following command to switch to the shell prompt of the appliance:
   $ ssh 169.254.0.10

6. Enter **Yes** to continue connecting to the management service.
7. Log on to the shell prompt of the appliance with the following default credentials:
   Password: nsroot.

8. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu:
   # networkconfig

9. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.
10. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the XenServer.
11. Type **3** and press **Enter** to select option 3, and specify the network mask for the IP addresses.
12. Type **4** and press **Enter** to select option 4, and specify the default gateway for the management service IP address.
13. Type **8** and press **Enter** to save the settings and exit.
14. Access the SD-WAN appliance by entering the new management service IP address of the appliance, such as https://<Management_Service_IP_Address>, in a web browser of a computer on the management network.
15. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
16. To complete the configuration process, see Provisioning the Appliance.

Provisioning the Appliance

Aug 31, 2017

After assigning an IP address to the management service, you are ready to provision the NetScaler and accelerator instances. As soon as you log on to the appliance, the configuration wizard appears.

When using the configuration wizard, keep the following points in mind:

- The following procedure assumes that you have already filled out the configuration worksheet.
- If you change the IP addresses of the Management Network, or change the default gateway to an address not on the Management Network, you lose connectivity to the appliance unless you are on the same Ethernet segment as the management port.
- When using the configuration wizard, check your entries carefully. The wizard has no Back button. If you need to modify the previous screen, use the Back button on your browser. This takes you to the logon page, then to the previous screen.
- The configuration wizard is displayed only when you log on to the appliance for the first time to configure the appliance. After you finish configuring the appliance, this wizard becomes inaccessible, and will reappear only after a factory reset. Check your entries carefully.

This wizard walks you through a fresh configuration of the appliance.

Note: If you receive a #SESS_CORRUPTED error at any time during these procedures, click **Logout**, clear your browser cache, close your browser, and open it again.

To configure the appliance by using the configuration wizard:

1. On the Welcome page, click Get Started.
   Note: All pages after the Get Started page have a heading that says, “Deployment Mode: Inline/L2 Mode,” but this wizard is used for all deployment modes.
2. Follow these steps to configure a fully 7.3-compliant system:
   - Acquire the following release 7.3 software distributions from the release 7.3 downloads page on My Citrix:
     - Management service (as a .tgz file)
     - NetScaler VM (as an .xva file)
     - Accelerator VM (as an .xva file)
     - Upgrade bundle (as a .upg) file
   - Navigate to the System > Configuration > Management Service > Software Images page, and then select **Upload** from the Action list.
   - Upload a release 7.3 Management Service image (distributed as a .tgz file).
   - Navigate to the System > Configuration > NetScaler > Software Images page, and then upload a release 7.3 NetScaler XVA image.
   - Navigate to the System > Configuration > SD-WAN > Software Images page, and then upload the accelerator XVA image.
   - Navigate to the System > Configuration > Management Service page, and then click the **Upgrade Management Service** link.
   - Select the management service image that you recently uploaded and click **OK**.
   - When the lower left-hand corner of the screen displays “Management Service Updated Successfully,” log off and clear your browser cache. Log on after the management service restarts (a few minutes).
   - On the Welcome screen, click **Get Started**.
3. For Management Access Settings, specify values for the various fields according to the network settings. The following

screen shot displays sample values used in this documentation. Enter values as follows:

- **XenServer IP Address**—(Item M4 on your worksheet, or H4 if this is the second appliance in an HA pair.) The management address of the built-in XenServer hypervisor. This must be a valid address on the management network.
- **Management Service IP Address**—(Item M5 on your worksheet, or H5 if this is the second appliance in an HA pair). The address of the Management Service VM that you use to perform most system management tasks. This must be a valid address on the management network.
- **Netmask**—(Item M3 on your worksheet). The subnet mask of the management network.
- **Gateway**— (Item M2 on your worksheet). The default gateway for the management network.
- **DNS Server**—The IP address of the DNS server. This is a mandatory parameter.
- **NTP Server**—IP or FQDN address of your time server. This will be used by all the virtual machines in the appliance. Note that if you use advanced CIFS or MAPI acceleration, the system time of the appliance must be close to that of the Windows domain server, so choose an NTP server that maintains a close relationship to the time on your Windows domain server.
  Note: Unless the NTP server is specified as an IP address, it is not used by the accelerator.
- **Time Zone**—Select your time zone from the pull-down menu.
- **Change Password**—Select this check box and type in a new nsroot password, two times, to change the password. This same password is used on the management service and the NetScaler instance for account nsroot, and on the accelerator for the admin account. If the password is not changed, it remains set to nsroot (the default).

Figure 1. Sample Values for the Fields in Management Access Settings Page of the Configuration

4. Check your settings and click **Continue**.
5. In the Manage Licenses section, see if an appropriate license is already listed in the Name field. If so, select it and skip to step 8.
6. Click **Upload** in the Update Licenses section.
7. Navigate to the folder that contains the license file and open the file.
8. Click **Add License** and upload the license file provided by Citrix. The license is added to the appliance, as shown in the following figure.
   Figure 2. Sample License Added to the Appliance on the Manage License Files Page of the Configuration Wizard

   You can also get a license file from the Citrix.com website by clicking the **here** link and using your My Citrix credentials.

9. Select the license in the Name field and click **Continue**. The SD-WAN Setup page appears. Fill in the fields as follows:
   1. **Network Settings**—This section informs the accelerators of the management network.
      - **SD-WAN Accelerator IP Address**—Enter the value of M6 from your worksheet. This is the IP address of the accelerator
      - **NetScaler IP Address**—Enter the value of M7 from your worksheet. This is the IP address of the NetScaler GUI.
      - **Use System Netmask and Gateway**—Select this option if you want to use the network mask and gateway IP addresses you had specified in the Platform Configuration page.
      - **Netmask**—Enter the value of M3 from your worksheet. This is the subnet mask (netmask) of the management network (note that you have entered this already, on a previous page).
      - **Gateway**—Enter the value of M2 from your worksheet again.
      - **Signaling IP Address**—Enter the value of T4 from your worksheet. This is the external signaling IP address of the accelerator, used by SD-WAN Plug-ins to connect to the appliance.

- **Signaling Netmask**—Enter the value of T2 from your worksheet. This is the subnet mask (netmask) of the external traffic network.

2. **XVA Files**—This section allows you to specify previously uploaded XVA files (Xen virtual machines) for the NetScaler and accelerator instances. Select the XVA images that you uploaded as part of step 2.
   Figure 3. SD-WAN Setup Page

10. Click **Continue**. The wizard starts provisioning the required instances, as shown in the following figure.
    Figure 4. Provisioning Progress Indicator

11. After the instances are provisioned, add one of your local LAN subnets to the Link Configuration section from list T7 in your worksheet, as shown in the following figure. This subnet will be added as a local LAN subnet in the accelerator. If you have more than one LAN subnet, you can add them to the LAN link definition in the Accelerator GUI after the configuration wizard completes. Click Add to add the subnet.
    Figure 5. Link Configuration Is at the Bottom of This Page

12. Log off, and then log back on. If you see a "Version Incompatibility Detected" message, install the upgrade bundle you downloaded in Step 2, using the procedure in Managing the Appliance by using the Management Service.

Basic configuration is complete. Next, perform deployment-mode-specific configuration (such as for WCCP mode).

Note: After the wizard completes, the appliance is configured for the basic setup. To configure the appliance for a specific deployment scenario, see Deployment Modes.

Deployment Modes

Aug 31, 2017

SD-WAN 4100/5100 appliances have two recommended deployment modes: WCCP and inline. These modes are commonly used without high availability (HA), and less commonly with HA.

At this time, Citrix recommends WCCP mode, with a single router and without HA, for most deployments. Use inline mode when WCCP is not available.

Although not all of the following modes are recommended at this time, they are all supported:

- WCCP mode with a single router
- WCCP mode with a single router and high availability
- Cascade of two or more appliances in WCCP mode along with a NetScaler MPX Appliance
- Cascade of two or more appliances in WCCP mode along with a NetScaler MPX Appliance in HA
- Inline mode
- Inline mode in HA
- Virtual inline mode
- Virtual inline mode in HA

> ## Note
>
> While modes other than WCCP and inline are supported, they are incompletely documented and are not recommended for typical installations. Please contact your Citrix representative when considering one of these modes.

Virtual Inline Mode

Aug 31, 2017

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router.

Virtual inline mode provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links.

The tasks for configuring virtual inline mode are performed on the router. On the SD-WAN appliance, just verify that the software version supports virtual inline mode and provision the instances with the necessary IP addresses. Do not change the default forwarding method.

> ## Note
>
> Citrix recommends that you do not deploy SD-WAN appliances in virtual inline mode with routers that do not support health monitoring.

Validated Designs

Aug 31, 2017

The physical mode for virtual inline deployment of a SD-WAN appliance is one-arm mode. In a one-arm topology with multiple subnets, the branches, local Repeater instances, and servers can exist on different subnets. For example, you can deploy the appliance in one-arm mode for managing the Repeater instances, with the NetScaler and local Repeater instances connected by the internal private subnet. A NetScaler-owned subnet IP address (SNIP) is used to communicate with an accelerator instance. You must enable MAC Based Forwarding (MBF) Use Subnet IP address (USNIP), and Return To Ethernet Sender options on the NetScaler instance.

This section contains Citrix validated deployment topologies for virtual inline mode.

Single Router and single link

Aug 31, 2017

A SD-WAN appliance deployed in virtual inline mode with one router and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with one router and one link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router

Two routers and a single link

Aug 31, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router (both routers individually)

Two routers and two links

Aug 31, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers and two links.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure Layer 4 Parameters (only if you expect connection migration between routers)
- Configuring VLANs for Connection Migration
- Configure a Router(both routers individually)

Two routers in High Availability setup

Aug 31, 2017

A SD-WAN appliance deployed in virtual inline mode with two routers in high availability setup and one link.

If you are deploying a SD-WAN appliance in virtual inline mode with two routers and a single link, complete the following procedures:

- Enable Layer 3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configure a Router (both routers individually)
- Configure Routers in High Availability Setup

Deployment Worksheet

Aug 31, 2017

The appliance uses at least two ports: the management port (typically 0/1) and the traffic port (such as 10/1) . You must select ports in advance, because the configuration depends on the identity of the ports.

The appliance uses three subnets directly: the management subnet, the external traffic subnet, and the internal traffic subnet. Multiple IP addresses are used on each subnet. Each subnet must be specified along with the correct subnet mask.

The following figure is a worksheet for these parameters. It supports inline and virtual inline modes, with and without HA. The table below the figure describes what each entry means.

Deployment Worksheet Parameters

	Parameter	Example	Your Value	Description
Management Subnet				
M2.	Gateway IP address	10.199.79.254		Default gateway serving the management subnet.
M3.	Subnet Mask	255.255.255.128		Subnet mask for the management subnet.
M4.	Xen Hypervisor IP address	10.199.79.225		IP address of Xen Hypervisor.
M5.	Service VM IP address	10.199.79.226		IP address of Management Service VM, which controls configuration.
M6.	Accelerator UI	10.199.79.227		Accelerator GUI, also called the Broker UI, which manages the instances as a unit.
M7.	NetScaler Management IP address	10.199.79.245		IP address of the NetScaler instance's GUI and CLI interfaces.
External Traffic Subnet				
T1.	Router IP address	172.17.17.1		IP address of the first router on external traffic subnet.
T2.	Router IP address	172.17.18.1		IP address of the second router on external traffic subnet.
T3.	Subnet Mask	255.255.255.0		Subnet mask of external traffic subnet.
T4.	Subnet IP address	172.17.17.2		Subnet IP address for NetScaler on external traffic subnet.
T5.	Subnet IP address	172.17.18.2		Subnet IP address for NetScaler on external traffic subnet.

	Parameter	Example	Your Value	Description
T6.	External Signaling IP address	172.17.17.10		Traffic to this IP address is load-balanced between the signaling IP addresses of the accelerators.
T7.	Local LAN Subnets	10.200.0.0/16		The local LAN subnet to be accelerated. This is the only subnet that will receive acceleration.
T8.	Traffic Port	10/1		Port used for accelerated traffic.
T9.	Traffic Port	10/6		Port used for accelerated traffic.

Note: Ports 10/3 and 10/4 are reserved for loopback cable. Do not configure these ports as traffic ports.

Configuring the NetScaler Instance

Aug 31, 2017

By default, the Getting Started wizard configures the appliance in inline mode. To deploy the SD-WAN appliance in virtual inline mode, you must configure the NetScaler instance of the appliance to support this mode.

To configure the NetScaler instance for the virtual inline mode of the appliance:

- Enable L3 Mode
- Enable the Return to Ethernet Sender Mode
- Add a Subnet IP Address
- Bind the Subnet IP Address to VLAN of Data Interface
- Configuring the Instance for Connection Migration

Enable Layer 3 Mode

Aug 31, 2017

In the inline mode, the default deployment mode of the appliance, the appliance uses layer 2 bridging. However, to deploy the appliance in virtual inline mode, you must enable layer 3 mode and disable layer 2 mode on the appliance.

To enable layer 3 mode and disable layer 2 mode by using the configuration utility

1. Access the NetScaler instance by clicking the NetScaler instance's IP address on the Configuration > Instances > NetScaler page. You are logged on to the NetScaler instance automatically.
2. Navigate to the System > Settings page.
3. Click the Configure modes link.

4. In the Configure Modes dialog box, select Layer 3 Mode (IP Forwarding) option.
5. Clear the Layer 2 Mode option, as shown in following figure.

6. Click OK.

To enable layer 3 mode and disable layer 2 mode by using the command line interface, run the following commands

```
> enable ns mode L3
> disable ns mode l2
```

Enable the Return to Ethernet Sender Mode

Aug 31, 2017

This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

To enable the Return to Ethernet Sender mode by using the configuration utility

1. Navigate to the System > Network page.
2. Click the Configure Layer 2 Parameter link, as shown in the following figure.

3. In the Configure Layer 2 Parameter dialog box, select the Return to Ethernet Sender option, as shown in the following figure.

4. Click OK.

To enable the Return to Ethernet Sender mode by using the command line interface, run the following command

```
> set L2Param -returnToEthernetSender ENABLED
```

Add a Subnet IP Address

Aug 31, 2017

You must add a Subnet IP (SNIP) address to the NetScaler instance. You assign this IP address to the NetScaler instance on the external traffic subnet. The NetScaler instance uses this address to communicate with the router.

To add a subnet IP address to the instance by using the configuration utility

1. Navigate to the System > Network > IPs page.
2. Click Add, as shown in the following figure.

3. In the Create IP dialog box, specify a subnet IP address in the IP Address field if the address has not already been created by the configuration wizard. (If it has already been created, it is listed on the IPs page). The subnet IP address declares the external traffic network. In the IP Address field, type the NetScaler traffic IP address (entry T4 in your worksheet).
4. In the Netmask field, specify the network mask (entry T3 in your worksheet).
5. From the IP Type group, make sure that the Subnet IP option is selected, as shown in the following figure.

6. Click Create and then Close.
7. If you are configuring the appliance for two links, repeat this procedure to create another subnet IP address (entry T5 of your worksheet).

To add a Subnet IP address to the instance by using the command line interface, run the following command

```
> add ns ip 172.17.17.2 255.255.255.0
> add ns ip 172.17.18.2 255.255.255.0
```

Note: Run the second command only if you are configuring the appliance for two links.

Bind the Subnet IP Address to VLAN of Data Interface

Aug 31, 2017

After adding a subnet IP address to the instance, you must bind the IP address to the VLAN of data interface you are using on the instance. By default, the VLAN 1007 is bound to the 10/1 and 10/2 interfaces.

To bind the subnet IP address to VLAN of data interface from the configuration utility

1. Navigate to the System > Network > VLANs page.
2. Select the VLAN, such as 1007, that was created when provisioning the instances.
3. Click Open, as shown in the following figure.

4. In the Interface Bindings tab of the Configure VLAN dialog box, notice that the VLAN is bound to interfaces 10/1 (entry T9 of your worksheet) and 10/2.

5. In the IP Bindings tab, select the subnet IP address you have created, as shown in the following figure.

6. Click Create and then Close.
7. If you are configuring the appliance for two links, repeat this procedure to bind second subnet IP (entry T5 of your worksheet) to another VLAN, such as 1009.

To bind the subnet IP addresses to VLANs of data interface by using the command line interface, run the following command

```
> bind vlan 1007–IPAddress 172.17.17.2 255.255.255.0
> bind vlan 1009–IPAddress 172.17.18.2 255.255.255.0
```

Note: Run the second command only if you are configuring the appliance for two links.

Configuring the Instance for Connection Migration

Aug 31, 2017

In your network setup, if you expect connection migration between the routers, then you must configure the layer 4 parameters and make changes to the VLAN configurations, appropriately.

Configure Layer 4 Parameters

If you have two routers and expect connection migration between the routers, you must configure layer 4 parameters on the NetScaler instance.

To configure layer 4 parameters from the configuration utility

1. Navigate to the System > Network page.
2. Click the Configure Layer 2 Parameter link, as shown in the following figure.

3. In the Configure Layer 4 Parameter dialog box, select the Vlan from the L2 Connection Method list, as shown in the following figure.

4. Click OK.

To bind the subnet IP address to VLAN of data interface by using the command line interface, run the following command

```
> set L4Param –l2connMethod vlan
```

Configuring VLANs for Connection Migration

If the network has two links and connected to interfaces that are bound to different VLANs, then you must bind the interfaces to the same VLAN to enable connection migration. Additionally, you must bind both subnet IP addresses to the same VLAN. Skip this procedure if you are using interfaces that are bound to the same VLAN.

To configure VLANs for connection migration by using the configuration utility

1. Navigate to the System > Network > VLANs page.
2. Select the VLAN, such as 1009, to which interface and subnet IP address of the second link is bound.
3. Click Open.
4. In the Interface Bindings tab of the Configure VLAN dialog box, clear the interface you are using (entry T9 of your worksheet), as shown in the following screen shot.

5. In the IP Bindings tab, clear the entry for subnet IP address bound to the VLAN.
6. Click OK.
7. Open the VLAN, such as 1007, to which you want to bind the data interface.
8. In the interface Binding tab, select the data interface (entry T9 in your worksheet)
9. In the IP Bindings tab, select the subnet IP address (entry T5 in your worksheet) you just unbound from the other VLAN.
10. Click OK.

To configure VLANs for connection migration by using the command line interface, run the following commands

```
> unbind vlan 1009 –ifnum 10/6
> unbind vlan 1009 –IPAddress 172.17.18.2 255.255.255.0
> bind vlan 1007 –ifnum 10/6
```

```
> bind vlan 1007 –IPAddress 172.17.18.2 255.255.255.0
```

Configuring a Router

Aug 31, 2017

To support virtual inline mode, a router must forward incoming as well as outgoing WAN traffic to the SD-WAN appliance. After the appliance processes the traffic, the router must forward the incoming traffic from the appliance to the LAN and the outgoing traffic from the appliance to the WAN. You have to configure policy based rules to avoid routing loops. In addition, the router must monitor the health of the appliance so that the appliance can be bypassed if it fails .

If the router supports the Reverse Path Forwarding feature, you must disable the feature on the interfaces with policies to redirect traffic to a SD-WAN appliance, including the interface that is connected to the appliance. Otherwise, the router intermittently drops traffic. By default, the Reverse Path Forwarding feature is enabled on the router.

Note: If the network has two routers, configure the following procedures for each router using appropriate IP addresses identified in the worksheet.

Policy-based Rules

In the virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the Ethernet ports of the router to the appliance and creating routing rules based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

Traffic shaping is not effective unless all WAN traffic passes through the appliance. Following is the basic routing algorithm:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward the packet to the appliance.
- If the packet is destined for the WAN, forward the packet to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.

Health Monitoring

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does not perform health monitoring. To enable health monitoring, define a rule to monitor the availability of the appliance, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Note: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy based routing do not support health checking. The health-monitoring feature is relatively new. It was first available in Cisco IOS release 12.3(4)T.

Following is an example of a rule for monitoring the availability of the appliance by using Cisco Router model 7600 with IOS Software version:

```
!- Use a ping (ICMP echo) to see if appliance is in the connected track
123 rtr 1 reachability
!
rtr 1
type echo protocol IpIcmpecho 17.17.17.2    schedule 1 life forever start-time now
```

This rule pings the appliance at 17.17.17.2 periodically. You can test against 123 to see if the unit is up.

Example of Router Configuration

Following is an example of configuring a Cisco router for virtual inline mode:

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.200.51.0   255.255.255.0
ip policy route-map server_side_map
!
interface FastEthernet0/1
ip address 17.17.17.1   255.255.255.0!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
ip policy route-map wan_side_map
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended server_side
permit ip 10.100.51.0 0.0.0.255 10.20.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.20.20.0 0.0.0.255 10.100.51.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 17.17.17.1   20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 17.17.17.1   10 track 123
```

This example applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP address of 10.100.51.0/24 and destination IP address of 10.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended server_side
 permit tcp 10.200.51.0   0.0.0.255 10.20.20.0 0.0.0.255
```



```
ip access-list extended wan_side
 permit tcp 10.20.20.0 0.0.0.255 10.200.51.0 0.0.0.255
!
```

Configuring Routers in a High Availability Setup

To configure high availability between routers, see the router-specific high availably configuration manual.

Verifying the Virtual Inline Mode

Aug 31, 2017

By default, the HTTP and CIFS acceleration is enabled on the appliance. After you have successfully configured the appliance in the virtual inline deployment mode, the appliance starts accelerating these connections.

Note: To configure application-specific connections, you must configure service classes for the respective applications. To configure a service class, see http://support.citrix.com/proddocs/topic/SD-WAN-72/cb-traffic-classification-con.html.

To verify that you have successfully configure the appliance in the virtual inline mode

1. Send network traffic through the appliance.
2. Navigate to the SD-WAN > Monitoring > Optimization > Connections page.
3. Verify that the Accelerated Connections tab displays entries for the accelerated connections, as shown in the following screen shot. This tab displays an entry each for all accelerated connections.

Figure 1. The Accelerated Connections tab

WCCP Mode

Aug 31, 2017

Web Cache Communication Protocol (WCCP) is a dynamic routing protocol introduced by Cisco. Originally intended only for web caching, WCCP version 2 became a more general-purpose protocol, suitable for use by accelerators such as Citrix SD-WAN appliances.

WCCP mode is the simplest way of installing a SD-WAN appliance when inline operation is impractical. It is also useful where asymmetric routing occurs, that is, when packets from the same connection arrive over different WAN links. In WCCP mode, the routers use the WCCP 2.0 protocol to divert traffic through the appliance. Once received by the appliance, the traffic is treated by the acceleration engine and traffic shaper as if it were received in inline mode.

> ### Note
>
> - For the purposes of this discussion, WCCP version 1 is considered obsolete and only WCCP version 2 is presented.
> - The standard WCCP documentation calls WCCP clients "caches." To avoid confusion with actual caches, Citrix generally avoids calling a WCCP client a "cache." Instead, WCCP clients are typically called "appliances."
> - This discussion uses the term "router" to indicate WCCP-capable routers and WCCP-capable switches. Though the term "router" is used here, some high-end switches also support WCCP, and can be used with SD-WAN appliances.

The SD-WAN appliances support two WCCP modes:

- WCCP is the original SD-WAN WCCP offering supported since release 3.x. It supports a single appliance service group (no clustering).
- WCCP clustering, introduced in release 7.2, allows your router to load-balance traffic between multiple appliances.

How WCCP Mode Works

The physical mode for WCCP deployment of a SD-WAN appliance is one-arm mode in which the SD-WAN appliance is connected directly to a dedicated port on the WAN router. The WCCP standard includes a protocol negotiation in which the appliance registers itself with the router, and the two negotiate the use of features they support in common. Once this negotiation is successful, traffic is routed between the router and the appliance according to the WCCP router and redirection rules defined on the router.

A WCCP-mode appliance requires only a single Ethernet port. The appliance should either be deployed on a dedicated router port (or WCCP-capable switch port) or isolated from other traffic through a VLAN. Do not mix inline and WCCP modes.

The following figure shows how a router is configured to intercept traffic on selected interfaces and forward it to the WCCP-enabled appliance. Whenever the WCCP-enabled appliance is not available, the traffic is not intercepted, and is forwarded normally.

Figure 1. WCCP Traffic Flow

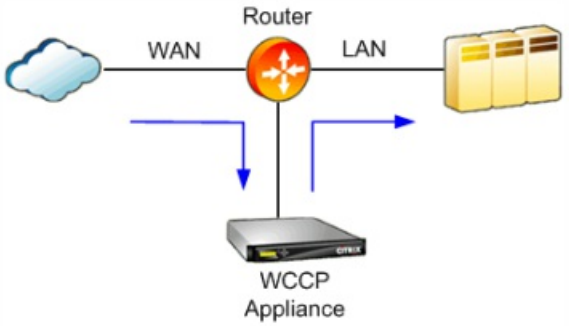

Traffic Encapsulation

WCCP allows traffic to be forwarded between the router and the appliance in either of the following modes:

- L2 Mode—Requires that the router and appliance be on the same L2 segment (typically an Ethernet segment). The IP packet is unmodified, and only the L2 addressing is altered to forward the packet. In many devices, L2 forwarding is performed at the hardware layer, giving it the maximum performance. Because of its performance advantage, L2 forwarding is the preferred mode, but not all WCCP-capable devices support it.
- GRE Mode—Generic Routing Encapsulation (GRE) is a routed protocol and the appliance can in theory be placed anywhere, but for performance it should be placed close to the router, on a fast, uncongested path that traverses as few switches and routers as possible. GRE is the original WCCP mode. A GRE header is created and the data packet is appended to it. The receiving device removes the GRE header. With encapsulation, the appliance can be on a subnet that is not directly attached to the router. However, both the encapsulation process and the subsequent routing add CPU overhead to the router, and the addition of the 28-byte GRE header can lead to packet fragmentation, which adds additional overhead.

WCCP mode supports multiple routers and both GRE vs. L2 forwarding. Each router can have multiple WAN links. Each link can have its own WCCP service group.

Traffic shaping is not effective unless the appliance manages UDP traffic as well as TCP traffic. A second service group, with a UDP service group for each WAN link, is recommended if traffic shaping is desired.

Registration and Status Updates

A WCCP client (an appliance) uses UDP port 2048 to register itself with the router and to negotiate which traffic should be sent to it, and also which WCCP features should be used for this traffic. The appliance operates on this traffic and forwards the resulting traffic to the original endpoint. The status of an appliance is tracked through the WCCP registration process and a heartbeat protocol. The appliance first contacts the router over the WCCP control channel (UDP port 2048), and the appliance and router exchange information with packets named “Here_I_Am” and “I_See_You,” respectively. By default, this process is repeated every ten seconds. If the router fails to receive a message from the appliance for three of these intervals, it considers the appliance to have failed and stops forwarding traffic to it until contact is reestablished.

Services and Service Groups

Different appliances using the same router can provide different services. To keep track of which services are assigned to which appliances, the WCCP protocol uses a service group identifier, a one-byte integer. When an appliance registers itself with a router, it includes service group numbers as well.

- A single appliance can support more than one service group.
- A single router can support more than one service group.
- A single appliance can use the same service group with more than one router.
- A single router can use the same service group with more than one appliance. For SD-WAN appliances, multiple

appliances are supported in WCCP cluster mode, and a single appliance is supported in WCCP mode.

- Each appliance specifies a “return type” (L2 or GRE) independently for each direction and each service group. SD-WAN 4000/5000 appliances always specify the same return type for both directions. Other SD-WAN appliances allow the return type to be different.

Multiple service groups can be used with WCCP on the same appliance. For example, the appliance can receive service-group 51 traffic from one WAN link and service-group 62 traffic from another WAN link. The appliance also supports multiple routers. It is indifferent to whether all the routers use the same service group or different routers use different service groups.

Service Group Tracking. If a packet arrives on one service group, output packets for the same connection are sent on the same service group. If packets arrive for the same connection on multiple service groups, output packets track the most recently seen service group for that connection.

Figure 2. Using different WCCP service groups for different services

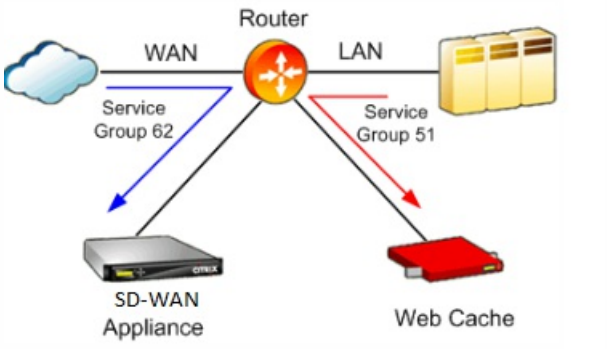

High Availability Behavior

When WCCP is used with high-availability mode, the primary appliance sends its own apA or apB management IP address, not the virtual address of the HA pair, when it contacts the router. If failover occurs, the new primary appliance contacts the router automatically, reestablishing the WCCP channel. In most cases the WCCP timeout period and the HA failover time overlap. As a result, the network outage is less than the sum of the two delays.

Standard WCCP allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive. WCCP clustering allows multiple appliances per service group.

Deployment Topology

The following figure shows a simple WCCP deployment, suitable for either L2 or GRE. The traffic port (1/1) is connected directly to a dedicated router port (Gig 4/12).
In this example, the SD-WAN 4100/5100 is deployed in one-arm mode, with the traffic port (1/1) and the management port (0/1) each connecting to its own dedicated router port.
On the router, WCCP is configured with identical `ip wccp redirect in` statements on the WAN and LAN ports. Two service groups are used, 71 and 72. Service group 71 is used for TCP traffic and service group 72 is used for UDP traffic. The SD-WAN appliance does not accelerate UDP traffic, but can apply traffic shaping policies to it.

Figure 3. Simple WCCP deployment

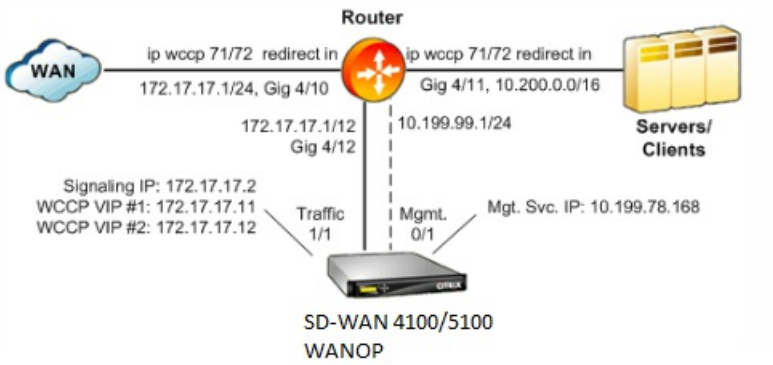

Note

The WCCP specification does not allow protocols other than TCP and UDP to be forwarded, so protocols such as ICMP and GRE always bypass the appliance.

WCCP Clustering

SD-WAN release 7.2 or later supports WCCP clustering, which enables your router to load-balance your traffic between multiple appliances. For more information about deploying SD-WAN appliances as a cluster, see WCCP Clustering.

WCCP Specification

For more information about WCCP, see Web Cache Communication Protocol V2, Revision 1, http://tools.ietf.org/html/draft-mclaggan-wccp-v2rev1-00.

WCCP Mode (Non-Clustered)

Aug 31, 2017

WCCP mode allows only a single appliance in a WCCP service group. If a new appliance attempts to contact the router, it discovers that the other appliance is handling the service group, and the new appliance sets an Alert. It periodically checks to determine whether the service group is still active with the other appliance, and the new appliance handles the service group when the other appliance becomes inactive.

Note: WCCP clustering allows multiple appliances per service group.

Limitations and Best Practices

Following are limitations and best practices for (non-clustered) WCCP mode:

- On appliances with more than one accelerated pair, all the traffic for a given WCCP service group must arrive on the same accelerated pair.
- Do not mix inline and WCCP traffic on the same appliance. The appliance does not enforce this guideline, but violating it can cause difficulties with acceleration. (WCCP and virtual inline modes can be mixed, but only if the WCCP and virtual inline traffic are coming from different routers.)
- For sites with a single WAN router, use WCCP whenever inline mode is not practical.
- Only one appliance is supported per service group. If more than one appliance attempts to connect to the same router with the same service group, the negotiation will succeed only for the first appliance.
- For sites with multiple WAN routers serviced by the same appliance, WCCP can be used to support one, some, or all of your WAN routers. Other routers can use virtual inline mode.

Configuring WCCP

Aug 31, 2017

A WCCP deployment follows the same initial steps as an inline deployment, but has additional steps beyond the basic inline procedure.

Prerequisites

Perform the following tasks if you have not done so already:

- Install the SD-WAN hardware. See Installing the Hardware.
- Fill out the Deployment Worksheet and perform the initial configuration. See .
- If you are using high-availability mode, see the Configuring the High Availability Setup on the Appliances section before proceeding.

Configuration Summary

The following high-level procedure summarizes the WCCP installation process, which works for both GRE and L2 forwarding and for any number of routers and links.

To configure WCCP mode

Note: You must follow the hyperlinks and follow the detailed instructions for each step.

1. Configuring the Router.
    1. Enable WCCP globally.
    2. Disable reverse path forwarding if your router supports it.
    3. Configure WCCP service groups.
2. Configuring Accelerators for WCCP Negotiation.
    1. Enable WCCP.
    2. For each service group:
        1. Create a service group definition on the SD-WAN appliance.
        2. Verify that this service group establishes WCCP communication with its associated routers.
3. Verifying WCCP Mode.
4. If using high-availability mode, configure and test the second appliance, then complete the Configuring the High Availability Setup on the Appliances procedure.

Configuring the Router

Note: This information is for WCCP mode. For WCCP Clustering, see the SD-WAN WCCP Clustering topics, especially the Configuring the Router subtopic.

For each WCCP router:

1. Enable WCCP in global router configuration.
2. For each WCCP service group on the worksheet for this router, declare ip wccp <sg> in global router configuration
3. Referring to the configuration examples below, for each WAN interface on this router:
    1. You can either use Method A or Method B to configure the router:
        - Method A: On the WAN interface only, declare ip wccp <sg> redirect in and ip wccp <sg> redirect out for the two service groups associated with the WAN interface.
        - Method B: If there is only one WAN interface, you can alternatively declare ip wccp <sg> redirect in on the WAN interface and on every LAN interface except the appliance's traffic interface.
    2. If your router supports reverse path forwarding, disable it on this interface by changing any ip verify unicast reverse-

path commands to no ip verify unicast reverse-path commands on each interface that has an ip wccp redirect command.

4. Save the configuration.

Router Configuration Examples

For normal operation, you must declare WCCP version 2 and the WCCP group ID for the router as a whole, and then enable redirection on each WAN interface.

Either one or two WCCP service groups can be used, but two are recommended, so that both TCP and UDP can be redirected, allowing more accurate traffic shaping. The WCCP standard requires that TCP and UDP traffic use different service groups.

Method A is required if the router has multiple WAN interfaces.

Example

Following is an example of configuring a Cisco IOS router:

```
! This example is for WCCP mode, not WCCP clustering
! (which is covered elsewhere)
config term
ip wccp version 2
! The two service groups are T11 and T12 on the
! configuration worksheet
! We will use group 72 for TCP and 73 for UDP.
ip wccp 72
ip wccp 73


! Repeat the following lines for each WAN interface
! you wish to accelerate:
interface <WAN_Interface>

! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect out
ip wccp 72 redirect in
ip wccp 73 redirect out
ip wccp 73 redirect in


^Z
```

Method B is preferred in circumstances when the routers do not support the wccp redirect out statement.

Example

Following is an example of configuring a Cisco IOS router:

```
! This example is for WCCP mode, not WCCP clustering
```



```
! (which is covered elsewhere)
config term
ip wccp version 2
!
! The two service groups are T11 and T12 on the
! configuration worksheet
! We will use group 72 for TCP and 73 for UDP.
ip wccp 72
ip wccp 73


! Repeat the following lines for the WAN interface
! you wish to accelerate:
interface <WAN_Interface>
!
! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect in
ip wccp 73 redirect in

! Repeat the following lines for all the LAN interfaces
! EXCEPT those connected to the SD-WAN appliance:
interface <LAN_Interface>
!
! If reverse-path forwarding is enabled, change any
!  ip verify unicast reverse-path commands to
! no ip verify unicates reverse-path commands:
no ip verify unicast reverse-path
ip wccp 72 redirect in
ip wccp 73 redirect in

^Z
```

Remember to save your router configuration when you are satisfied that it is correct.

Configuring Accelerators for WCCP Negotiation

One accelerator instance manages WCCP control traffic on behalf of all the instances. The WCCP control traffic is negligible. The actual data traffic is divided among all the accelerators.

Note: The GUI calls WCCP mode “single cache.”

Summary: To configure the accelerators for WCCP mode, first enable WCCP mode. Then, configure service groups and create a WAN link definition, on the SD-WAN appliance, for each WAN link on each WCCP router. (Each link has two service groups, one for TCP and one for UDP.) If a service group is already defined for a given link, add the current router’s IP address to the definition. Test the service group’s WCCP status before creating the WAN link definition, and verify link traffic and acceleration status before configuring the next WAN link.

To configure the accelerators for WCCP mode

1. On the SD-WAN appliance, Navigate to Configuration > Appliance Settings > Advance Deployments > WCCP page.
2. If the **Enable** button is displayed, click it to enable WCCP mode on the appliance. (If the Disable button is displayed, WCCP mode is already enabled.)
   Note: We will actually be configuring two caches.
3. In the Select Mode area, select **Single Cache**.
4. Starting with accelerator instance #1 (labeled "WCCP Cache 1" on the page), configure the SD-WAN IP Details by entering the external VIP you defined for accelerator instance 1 (T5 on your worksheet for instance #1, T6 for instance #2). Set the subnet mask for the external traffic network (T2 on your worksheet). Set the gateway IP for the external traffic network (T1 on your worksheet). Click Save. The Configure Service Group controls appear.
5. In the Configure Service Group section, click Add. An Add Service Group popup appears.
6. In the Service Group Details area, specify a WCCP service group ID in the ID field. This ID must match one of the service groups that you have defined on your router. Start with the lowest-numbered service group in your list (T11 for the TCP service class, T12 for the UDP service class.).
7. In the WCCP Priority field, set the WCCP priority to 100 for instance #1, or to 80 for instance #2. (Other values work. Use a priority for instance #1 that is greater than the priority for instance #2, and use a priority for instance #2 that is greater than zero.)
8. From the Protocol list, select a protocol. You will perform this step for both TCP and UDP. Start with TCP.
9. In the Service Group Password field, enter a password if your router is configured to require one. Otherwise, leave the field blank.
10. In the Router Communications Details area, in the Router IP Address field, enter the IP address of the router. This is the router's address for its appliance-facing interface (T8 on your worksheet). If you use multiple routers to communicate with the appliance, list them all here.
11. From the Router Assignment list, select a router assignment (Hash, Mask, or Auto). Auto is recommended. If Auto is selected, Hash is negotiated if the router supports it. Otherwise Mask is used.
12. From the Router Forwarding list, select **Level 2** or **GRE**. The same method must be used for both outbound and inbound packets. L2 is recommended whenever possible, as GRE adds overhead to both the router and the appliance. L2 requires that your router support Level 2, and that the router's IP and the VIP addresses be in the same subnet. Otherwise, use GRE.
13. Click Create.
14. Repeat steps 6-13 with the next service group in sequence, but selecting UDP instead of TCP.
15. Repeat steps 4-14 for instance #2 (called "WCCP Cache 2" in the GUI), except that the Cache IP is T6 from your worksheet (instead of T5), and the WCCP priority value is 80 (instead of 100).
16. If desired, click Advanced Settings on the WCCP page and select a quicker timeout (Responsive or Tolerant, rather than Default). This is a WCCP 2.1 feature and is not supported by all routers. If the appliance has trouble connecting to the router, set this parameter back to Default.

Note: You must consider the following points when configuring a Citrix SD-WAN 4000/5000 appliance:

- Traffic is load balanced across the accelerators on the basis of NetScaler load balancing policies.
- The WCCP service group ID that you assign to the accelerator must match a service group defined on your router, or the WCCP negotiation fails.

Verifying WCCP Mode

Aug 31, 2017

You can monitor the WCCP configuration from the SD-WAN GUI.

To monitor the WCCP configuration

1. Navigate to the Monitoring > Appliance Performance > WCCP page.
2. Select a cache and click Get Info. A Cache Status page displays the WCCP configuration, as shown in the following figure.

3. Start traffic that should be forwarded through the SD-WAN appliance and monitor the connection on the Monitoring > Optimization > Connections page.
   - If the connections are shown on the Accelerated Connections tab, that is an indicator that everything is working.
   - If the connections are on the Unaccelerated Connections tab, look at the Details column. A routing asymmetry detected message implies that one of the ip wccp redirect lines on the router is missing or has an error, or that different paths are taken by client-server and server-client traffic.
   - If no connections are shown, but the appliance reports that it is connected to the router, and the WCCP monitoring page shows no errors, the issue is probably with the router configuration.

WCCP Clustering

Aug 31, 2017

The WCCP clustering feature enables you to multiply your acceleration capacity by assigning more than one SD-WAN appliance to the same links. You can cluster up to 32 identical appliances, for up to 32 times the capacity. Because it uses the WCCP 2.0 standard, WCCP clustering works on most routers and some smart switches, most likely including those you are already using.

Because it uses a decentralized protocol, WCCP clustering allows SD-WAN appliances to be added or removed at will. If an appliance fails, its traffic is rerouted to the surviving appliances.

Unlike SD-WAN high-availability, an active/passive pair that uses two appliances to provide the performance of a single appliance, the same appliances deployed as a WCCP cluster has twice the performance of a single appliance, delivering both redundancy and improved performance.

In addition to adding more appliances as your site's needs increase, you can use Citrix's "Pay as You Grow" feature to increase your appliances' capabilities through license upgrades.

Citrix Command Center is recommended for managing WCCP clusters. The following figure shows a basic network of a cluster of SD-WAN appliances in WCCP mode, administered by using Citrix Command Center.

Figure 1. SD-WAN Cluster Administered by Using Citrix Command Center

Load-Balanced WCCP Clusters

The WCCP protocol supports up to 32 appliances in a fault-tolerant, load balanced array called a cluster. In the example below, three identical appliances (same model, same software version) are cabled identically and configured identically except for their IP addresses. Appliances using the same service groups with the same router can become a load balanced WCCP cluster. When a new appliance registers itself with the router, it can join the existing pool of appliances and receive its share of traffic. If an appliance leaves the network (as indicated by the absence of heartbeat signals), the cluster is rebalanced so that only the remaining appliances are used.

Figure 2. A load-balanced WCCP cluster with three appliances

One appliance in the cluster is selected as the designated cache, and controls the load-balancing behavior of the appliances in the cluster. The designated cache is the appliance with the lowest IP address. Because the appliances have identical configurations, it doesn't matter which one is the designated cache. If the current designated cache goes offline, a different appliance becomes the designated cache.

The designated cache determines how the load-balanced traffic is allocated and informs the router of these decisions. The router shares information with all members of the cluster, so the cluster can operate even if the designated cache goes offline.

Note: As normally configured, a SD-WAN 4000/5000 appliance appears as two WCCP caches to the router.

Load-Balancing Algorithm

Load balancing in WCCP is static, except when an appliance enters or leaves the cluster, which causes the cluster to be rebalanced among its current members.

The WCCP standard supports load balancing based on a mask or a hash. For example, SD-WAN WCCP clustering uses the mask method only, using a mask of 1-6 bits of the 32-bit IP address. These address bits can be non-consecutive. All addresses yielding the same result when masked are sent to the same appliance. Load balancing effectiveness depends on choosing an appropriate mask value: a poor mask choice can result in poor load-balancing or even none, with all traffic sent to a single appliance.

Deployment Topology

Aug 31, 2017

Depending on your network topology, you can deploy WCCP cluster either with a single router or with multiple routers. Whether connected to a single router or multiple routers, each appliance in the cluster must be connected identically to all routers in use.

Single Router Deployment

In the following diagram, three SD-WAN appliances accelerate the datacenter's 200 Mbps WAN. The site supports 750 XenApp users.

As shown on the SD-WAN Datasheet, a SD-WAN 3000-100 can support 100 Mbps and 400 users, so a pair of these appliances supports 200 Mbps and 800 users, which satisfies the datacenter's requirements of a 200 Mbps link and 750 users.

For fault tolerance, however, the WCCP cluster should continue to operate without becoming overloaded if one appliance fails. That can be accomplished by using three appliances when the calculations call for two. This is called the N+1 rule.

Failure is an unusual event, so usually all three appliances are in operation. In this case, each appliance is supporting only 67 Mbps and 250 users, leaving plenty of headroom, and making good use of the fact that the cluster has three times the CPU power and three times the compression history of a single appliance.

Without WCCP clustering, the same level of capacity and fault-tolerance would require a pair of SD-WAN 4000-500 appliances in high availability mode. Only one of these appliances is active at a time.

Multiple Router Deployment

Using multiple WAN routers is very similar to using a single WAN router. If the previous example is changed to include two 100 Mbps links instead of one 200 Mbps link, the topology changes, but the calculations do not.

Limitations

Aug 31, 2017

Configuring appliances in a WCCP cluster has the following limitations:

- All appliances within a cluster must be the same model and use the same software release.
- Parameter synchronization between appliances within the cluster is not automatic. Use Command Center to manage the appliances as a group.
- SD-WAN traffic shaping is not effective, because it relies on controlling the entire link as a unit, and none of the appliances are in a position to do this. Router QoS can be used instead.
- The WCCP-based load-balancing algorithms do not vary dynamically with load, so achieving a good load balance can require some tuning.
- The hash method of cache assignment is not supported. Mask assignment is the supported method.
- While the WCCP standard allows mask lengths of 1-7 bits, the appliance supports masks of 1-6 bits.
- Multicast service groups are not supported; only unicast service groups are supported.
- All routers using the same service group pair must support the same forwarding method (GRE or L2).
- The forwarding and return method negotiated with the router must match: both must be GRE or both must be L2. Some routers do not support L2 in both directions, resulting in an error of "Router's forward or return or assignment capability mismatch." In this case, the service group must be configured as GRE.
- SD-WAN VPX does not support WCCP clustering.
- The appliance supports (and negotiates) only unweighted (equal) cache assignments. Weighted assignments are not supported.
- Some older appliances, such as the SD-WAN 700, do not support WCCP clustering.
- (SD-WAN 4000/5000 only) Two accelerator instances are required per interface in L2 mode. No more than three interfaces are supported per appliance (and then only on appliances with six or more accelerator instances.)
- (SD-WAN 4000/5000 only) WCCP control packets from the router must match one of the router IP addresses configured on the appliance for the service group. In practice, the router's IP address for the interface that connects it to the appliance should be used. The router's loopback IP should not be used.

Planning Your Deployment

Aug 31, 2017

Deploying appliances in a WCCP cluster requires more planning than does deploying a single appliance. Read the following sections carefully before proceeding.

Selecting Appliances

Aug 31, 2017

The appliances you select for the deployment must all be the same model, running the same software version. Otherwise, management and troubleshooting can become impractical.

Your appliance choice is generally made by comparing your site's WAN bandwidth and number of WAN users to the capacities of the different appliances in the SD-WAN Data Sheet. For fault tolerance, always order one more appliance than is absolutely required according to the data sheet.

The number of appliances you need is found as follows, rounding up all fractions:
appliances = max (appliances_bw, appliances_users),
where
appliances_bw = (WAN_bandwidth / Optimized_WAN_capacity) +1
appliances_users = (WAN_users / Maximum_HDX_sessions) +1
Note that if appliances = 2, you can use just a single appliance instead of WCCP clustering, or an HA pair instead of WCCP clustering, since the equation builds in a spare appliance. In other words, WCCP clustering is not necessary (from a capacity perspective) unless appliances is 3 or more.

Example. Suppose you have 700 users and a 100 Mbps link. Some appliances you might consider are the SD-WAN 2000-050, the SD-WAN 3000-100, and the SD-WAN 4000-310.

Model	Optimized WAN Capacity	Maximum HDX Sessions	Appliances_bw	Appliances_users	Appliances
2000-050	50 Mbps	300	3	4	4
3000-100	100 Mbps	400	2	3	3
4000-310	310 Mbps	750	2	2	2

As you can see from the above table, the higher-performance platforms require fewer appliances to get the job done, as you would expect. The SD-WAN 4000-310 meets the requirements with a single appliance, and evaluates to two appliances only because the equations build in a spare.

You can always add more capacity by adding more appliances, but that is not always necessary. The bandwidth limits of two of the three choices, the SD-WAN 3000-100 and the SD-WAN 4000-310, can be increased through a license upgrade. The SD-WAN 2000-050 however, is already at the high end of the range for SD-WAN 2000 appliances.

Load-Balancing in the WCCP Cluster

Aug 31, 2017

Traffic is distributed among the appliances in the WCCP cluster. If an appliance leaves the cluster (through failure, overload, or being manually disabled), its traffic is rebalanced by distributing it among the surviving members. If an appliance joins the cluster, traffic is rebalanced once more to give the new appliance its fair share.

The Address Mask

Traffic is distributed on the basis of an address mask that is applied to the source and destination addresses of WAN traffic. You must select an appropriate mask field for efficient load-balancing. An inappropriate mask can result in load-balancing that is poor to nonexistent. For example, if the mask matches an address field that is identical at all your remote sites, all your WAN traffic is sent to a single appliance, overloading it. For example, if all of your remote sites have an address in the form of 10.0.x.x, and your mask bits are within the 10.0 portion of the address all traffic is sent to a single appliance.

The address bits extracted by the address mask are used as an index that is used (indirectly) to select one of the WCCP caches (appliances). For example, an address mask with two "one" bits results in four possible values, depending on the address. Each of these values can be thought of as a bucket. With two mask bits, you have four buckets, numbered 0-3. The buckets are assigned to WCCP caches. To be effective, there must be at least as many buckets as caches. If you use a two-bit mask and have five or more caches, some caches are idle, because each bucket is assigned to only one cache, and there are not enough buckets to cover all five caches:

Cache	1	2	3	4	5
Buckets	0	1	2	3	-

If there are more buckets than caches, some caches are assigned multiple buckets. For example, if you set three mask bits, creating eight buckets, and you have four caches, two buckets are assigned to each cache. If you have five caches, three caches are assigned two buckets each, and two caches are assigned just one. If each bucket represents the same number of users, you have a 2:1 load imbalance across caches:

Cache	1	2	3	4	5
Buckets	0-1	2-3	4-5	6	7

Increasing the number of set mask bits reduces this imbalance. With four mask bits (16 index values) and five caches, four caches receive three buckets and one cache receives four buckets, resulting in only a 4:3 imbalance. With six set mask bits (the largest number supported), four caches receive 13 buckets and one receives 12, which is only a 13:12 load imbalance.

Cache	1	2	3	4	5
Buckets	0-12	13-25	26-38	39-51	52-63

Ideally, you would like each remote site to be directed to a single appliance in the WCCP cluster, so that all traffic to and from a given site is stored in the same compression history. With this arrangement, any traffic from one user at the site can be used to compress similar traffic from any other user at that site. In other words, for compressibility, load-balancing works best if it the address mask selects the bits that differentiate one remote site from another. These are often the least-significant bits of the subnet portion of the IP address. Using these bits tends to allocate the same number of remote sites (not users) per local appliance. A mask that aligns with the host portion of the address instead of the subnet results in a

more equal number of remote users (not sites) per appliance, but at the expense of compression effectiveness. (Compression is only effective when connections flow through the same appliances, and splitting traffic from the same remote site between two or more local appliances interferes with this.)

Finally, for good load-balancing, each "one" bit in the address mask must be set to one on 50% of the remote addresses, and set to zero on 50% of the remote addresses. This is not the case on all address bits, since in most WANs, the highest-order network bits never change at all (such as the 10 in 10.x.x.x). Such bits must never be selected by the address mask.

In addition, many subnets are only sparsely populated. For example, if only 50 addresses are used in the subnet 10.1.2.0/24, and they are assigned sequentially, the two higher-order host bits (representing the unused range of 10.1.2.64-10.1.2.255) for this subnet never change, and if these two bits are included in the address mask, three-fourths of the buckets receive no traffic.

Useful compromises between these two extremes can generally be found.

Follow these rules:

- The number of ”one” bits in the address mask must allow at least as many combinations as there are WCCP caches in the cluster. That is, if you have eight appliances, the address mask must contain at least three ”one” bits.
- The ”one” bits in the address mask must each be inside the active address range for most of your remote subnets, or they skew the load-balancing distribution.
- The mask should split the address range of individual remote sites into as few pieces as possible, for best compression performance.
- If a remote appliance is faster than the local members of the WCCP cluster, the mask should be designed to divide its traffic between multiple local appliances. For example, a 100 Mbps remote appliance should have its traffic split between two 50 Mbps local appliances by setting a bit inside the remote appliance’s active address range.
- The “one” bits in the mask are typically contiguous, but this is not required. They can be in any pattern.

Example: Suppose you set an address mask of 0x0000 0f00, which has four “one” bits. This defines a four-bit field that is extracted from the IP address, yielding 16 possible results (16 buckets). These buckets are in turn assigned to the actual WCCP caches in the WCCP cluster.

Address	Masked Address (mask = 0x0000 0f00)	Bucket
10.0.0.5	0.0.0.0	0
10.0.1.128	0.0.1.0	1
155.0.2.55	0.0.2.0	2
253.100.255.2	0.0.15.0	15
10.0.15.1	0.0.15.0	15

Zero bits in the mask are ignored, and the “one” bits are used to define the extracted field. So if the mask is 0x10 10 10 10, these widely separated “one” bits are extracted into a four-bit field, declaring 16 buckets and a bucket numbers in the range of 0-15.

If the mask value is set to zero, a default value of 0x00 00 0f 00 is used.

Assigning Buckets to Appliances

Aug 31, 2017

The mapping of bucket to appliances is subject to several variables:

- Which appliances are available: If an appliance is down, its share of buckets are given to the available appliance. If a new appliance is added to the cluster, it is given its fair share of buckets.
- The mapping algorithm used (deterministic or least-disruptive).
- The order in which appliances come online (least-disruptive mapping only).
- The IP addresses of the appliances. WCCP algorithms can use a sorted list of appliance IP addresses; for example, assigning buckets to appliances in the same order as the appliance IP addresses.

The most important of these factors, from an administrator's point of view, is the mapping algorithm.

Deterministic mapping. The deterministic mapping algorithm is less graceful than the least-disruptive algorithm, but it supports Hot Standby Router Protocol (HSRP) and Global Server Load Balancing (GSLB) routing, and is required when multiple routers using such protocols share the WCCP cluster.

Deterministic mapping is also the preferred method when the cluster has only two appliances.

Assignments are based on the IP addresses of the active appliances. Each appliance gets its fair share of bucket, with the lowest-numbered bucket being assigned to the appliance with the lowest IP address. If there are more appliances than buckets, the leftover appliances (with no bucket assigned to them) are the ones with the highest-numbered IP addresses. This deterministic assignment allows traffic to arrive for a single connection through any of the routers in the service group and be forwarded to the same appliance.

Reassignment can be disruptive to accelerated connections, which are reset if they migrate to a different appliance. With deterministic mapping, the number of buckets that are reassigned to new appliances can be quite high if there are three or more appliances.

Least-disruptive mapping. When a bucket is assigned to a different appliance, any open accelerated connections that used the old appliance is reset. The least-disruptive algorithm keeps the reassignment to a minimum. For example, if you have three appliances, and one appliance fails, the new mapping preserves roughly two-thirds of the assignments and remaps the remaining third (which fails anyway, because their appliance failed). The least-disruptive algorithm does not support HSRP or GSLB routing, because it is not guaranteed to result in identical mappings on all the routers in the service group, and therefore, packets from a single connection might be sent to two different appliances by two different routers, which causes accelerated connections to fail.

Startup and Failover Behavior

Aug 31, 2017

Each appliance registers itself with the routers specified in its service class definitions. The first appliance to register itself, becomes the *designated cache*, and works with the routers to apportion traffic between itself and the other caches (called *subordinate caches*). Because your appliances use identical WCCP algorithms, it does not matter which one becomes the designated cache.

As additional appliances come online, they are added to the WCCP cluster, and the traffic is reapportioned among the active appliances. This happens at ten-second increments. After a cold start of the routers or appliances, all of the appliances might come online within the same ten-second window, or they might arrive over multiple ten-second windows, causing traffic to be reapportioned multiple times before it stabilizes. In the latter case, the appliances that come online first maycan become overloaded until additional appliances come online.

An accelerated connection fails when allocated to a different appliance, making reallocation disruptive. This is not true of non-accelerated connections, which generally experience a delay of thirty seconds or more, and then continue. The least-disruptive mapping option minimizes the amount of reallocation when an appliance fails.

If an appliance fails or otherwise goes offline, its absence is noted, and the designated cache reapportions its traffic to the remaining appliances. If the designated cache itself goes offline, the role of designated cache is also reapportioned. It takes about thirty seconds for the cluster to react to the loss of a cache.

Deployment Worksheet

Aug 31, 2017

On the following worksheet, you can calculate the number of appliances needed for your installation and the recommended mask field size. The recommended mask size is 1-2 bits larger than the minimum mask size for your installation.

Parameter	Value	Notes
Appliance Model Used		—
Supported XenApp and XenDesktop Users Per Appliance	Uspec =	From data sheet
XenApp and XenDesktop Users on WAN Link	Uwan =	—
User overload Factor	Uoverload = Uwan/Uspec =	—
Supported BW Per Appliance	BWspec =	From data sheet
WAN Link BW	BWwan =	—
BW Overload Factor	BWoverload = BWwan/BWspec =	—
Number of appliances required	N = max(Uoverload, BWoverload) +1 =	Includes one spare
		—
Min number of buckets	Bmin = N, rounded up a power of 2 =	—
If SD-WAN 4000 or 5000,	Bmin = 2*N, rounded up to a power of 2 =	—
Recommended value	B = 4 * Bmin if Bmin <= 16, else 2 * Bmin =	—
Number of "one" bits in address mask	M = log2(B)	If B=16, M=4.

Mask value: The mask value is a 32-bit address mask with a number of “one” bits equal to M in the above worksheet. Often

these bits can be the least-significant bits in the WAN subnet mask used by your remote sites. If the masks at your remote sites vary, use the median mask. (Example: With /24 subnets, the least significant bits of the subnet are 0x00 00 nn 00. The number of bits to set to one is log2(mask size): if mask size is 16, set four bits to one. So with a mask size of 16 and a /24 subnet, set the mask value to 0x00 00 0f 00.): _____

The above guidelines work only if the selected subnet field is evenly distributed in your traffic, that is, that each address bit selected by the mask is a one for half the remote hosts, and a zero for the other half. Otherwise, load-balancing is impaired. This even distribution might be true for only a small number of bits in the network field (perhaps only two or three bits). If this is the case with your network, instead of masking bits in the offending area of the subnet field, displace those bits to a portion of the host address field that has the 50/50 property. For example, if only three subnet bits in a /24 subnet have the 50/50 property, and you are using four mask bits, a mask of 0x00 00 07 10 avoids the offending bit at 0x00 00 0800 and displaces it to 0x00 00 00 10, a portion of the address field that is likely to have the 50/50 property if your remote subnets generally use at least 32 IP addresses each.

Parameter	Value	Notes
Final Mask Value		—
Accelerated Bridge		Usually apA
WAN Service Group		A service group not already in use on your router (51-255)
LAN Service Group		Another unused service group
Router IP address		IP address of router interface on port facing the appliance
WCCP Protocol (usually "Auto")		—
DC Algorithm		Use "Deterministic" if you have only two appliances or are using dynamic load balancing like HSRP or GSLB. Otherwise, use "Least Disruptive."

Configuring WCCP Clustering

Aug 31, 2017

After you have finalized the deployment topology, considered all limitations, and filled in the deployment worksheet, you are ready to deploy your appliances in a WCCP cluster. To configure the WCCP cluster, you need to perform the following tasks:

- Configuring the NetScaler Instances
- Configuring the Router
- Configuring the Appliance

Configuring the Router

Aug 31, 2017

WCCP configuration on the router is simple, because most WCCP parameters are set by the appliances.

Unlike legacy SD-WAN WCCP support, WCCP clustering uses two service groups for TCP traffic. One service group is used on the router's WAN interface, and the other is used on the router's LAN interfaces (except for the LAN interface used by the SD-WAN appliances themselves, when deployed in L2-mode WCCP cluster).

As shown in the following figure, you need to configure two service groups because WCCP allows the mask to be applied to either the source IP or the destination IP address, which is not quite what is required. To keep connections between two endpoints together, regardless of which endpoint initiates the connection, the appliance applies the address mask to the source IP address of incoming WAN traffic, and to the destination IP address of incoming LAN traffic. This requires two service groups.

The WAN service group uses WCCP source-ip address masking, while the LAN service group uses dest-ip masking. In some deployments, it may be necessary to reverse the assignments, using the "WAN" service group for your LAN interface and vice versa. This might occur if the number of local IP addresses greatly exceeds the number of remote IP addresses.

Figure 1. SD-WAN WCCP Cluster

To configure WCCP clustering on the router

This procedure assumes Cisco routers, but is similar on other routers. It uses the first of the two methods, discussed above, of redirecting WCCP traffic with an ip wccp redirect in statement on both LAN and WAN ports.

1. Fill in the WCCP clustering Deployment Worksheet.
2. Log on to your router
3. In the global declarations section, declare each service group on the WCCP clustering worksheet, listed as **WAN service group** and **LAN Service group**. For example, ip wccp 61 and ip wccp 62.
   Note: The ip wccp command allows, but does not require, a more elaborate syntax than this, and can specify an ACL name or a password. Both service groups must have the same password, if one is used. The ACLs can be different.
4. Inside the interface declarations for each WAN interface that connects to remote SD-WAN appliances, add an ip wccp x redirect in statement, where x is the WAN service group from the WCCP clustering worksheet.
5. Inside the interface declarations for each LAN interface (except the one connecting to the WCCP cluster, if you are using L2 mode), add an ip wccp y redirect in statement, where y is the LAN service group from the WCCP clustering worksheet.
6. Save your configuration.

Example. The following example uses service group 61 for the WAN service group and service group 62 for the LAN service group. Three router interfaces are used. One is connected to the WAN, one is connected to the LAN, and one is connected to the WCCP cluster.

```
!
! Example is for WCCP clustering using WCCP redirect in statements
! on LAN and WAN interfaces.
! This definition is appropriate for modern Cisco routers.
! Global declarations
ip wccp 61
ip wccp 62
```

```
!
interface GigabitEthernet1/1
description LAN interface. SG 62 is used for LAN
ip address 172.80.1.56 255.255.255.0
ip wccp 62 redirect in
!
interface GigabitEthernet1/2
description LAN interface attaching SD-WAN L2-WCCP appliances
description (No wccp redirect statements are used on this interface)
ip address 172.80.21.56 255.255.255.0
!
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
!
```

Note: If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.

- If the router used multiple ports for LAN traffic, each port is configured with an ip wccp 62 redirect in statement. Similarly, if the router used multiple ports for WAN traffic, each port is configured with an ip wccp 61 redirect in statement.
- If multiple routers shared the same WCCP cluster, they use the same service groups.

It is also possible to use ip wccp redirect statements on only the WAN interfaces:

```
! Example for WCCP clustering using WCCP redirect in/out statements on
! WAN interface only
! This definition is appropriate for modern Cisco routers.
interface GigabitEthernet1/3
description WAN interface. SG 61 is used for WAN. SG 62 is used for LAN.
ip address 172.80.22.56 255.255.255.0
ip wccp 61 redirect in
ip wccp 62 redirect out
!
```

In many routers, the ip wccp redirect out path is not optimized in hardware, but uses the CPU. If the router's capabilities along this path exceeds the WAN speed, this method is practical, and is simpler than using redirect statements on every interface.

Router ACLs can be used to limit redirection. For example, for initial testing, perhaps only a single remote IP address might be allowed to be redirected through WCCP.

Configuring the Appliance

Aug 31, 2017

Repeat the following procedure for each appliance in the cluster:

1. Fill in the WCCP clustering Deployment Worksheet.
2. Navigate to Configuration > Appliance Settings > WCCP page.
3. Click Enable to enable WCCP mode on the appliance.
4. Select **Cluster (Multiple Caches)** option.

5. Fill in parameters in the **Select SD-WAN Cluster** section.

6. Enter T5 from your worksheet as the Cache 1 IP, T6 as the Cache 2 IP, T2 as the Subnet Mask, and T1 as the Gateway. Click **Save**. The **Configure Service Group** section appears.
7. In the Service Group Details section, specify the WAN and LAN service groups (T11 and T12 from your worksheet).
8. In the Priority field, select **100** (in practice this value is somewhat arbitrary).
9. From the Protocol list, select **TCP**.
10. In the DC Algorithm field, select **Deterministic** or **Least Disruptive**. “Deterministic” is always safe to use, and should be used if you are using only two appliances, or are using multiple routers. “Least Disruptive” disrupts fewer user sessions on failover when used with clusters of three or more appliances, but has restrictions on its use.
11. Set **Service Group Pair Status** to On.
12. If your router is configured to require a password, enter the password in the **Service Group Password** field. Otherwise, leave the field blank.
13. In the **Router Communications Details** section, enter the IP address of the router (T8 on your worksheet: often identical to T1 as well). This is the IP address of the appliance-facing router interface. If you use multiple routers to communicate with the appliance, list them all here.
14. From the Router Forwarding list, select Level 2 or GRE, according to the capabilities of your router. Use Level 2 if you can, and GRE if you must.
15. For the Mask Value, enter the value you determined from the WCCP Clustering worksheet. This is a critical value: a poor choice will result in poor load-balancing or none at all.
16. Click Create. This creates the WAN and LAN service groups.
17. On the Configuration > Optimization Rules > Link Definitions page, change the bandwidth limits on each defined WAN to 95% of the aggregate speed of all your WANs. This prevents the link from being under-utilized when load-balancing is imperfect. If ICA (XenApp/XenDesktop) is the dominant use, set each appliance to (95% of WAN bandwidth)/N, where N is the number of appliances (or twice the number of appliances if they are SD-WAN 4000 or 5000 units), to divide the bandwidth equally among the appliances. This latter method is most appropriate for applications with large numbers of active connections that have relatively low bandwidth requirements.

Testing and Troubleshooting

Aug 31, 2017

The **Monitoring** > **Appliance** > **Application Performance** > **WCCP** page shows the current state of not only the local appliance but of all other appliances that have joined the cluster. Select a WCCP cache and click **Get Info**.

The Cache Status tab shows the local appliance's status. When all is well, the status is "25: has assignment." You must refresh the page manually to monitor changes in status. If the appliance does not reach the status of "25: has assignment" within a timeout period, other informative status messages are displayed.

Additional information is displayed when you click on the Service Group or the Routers tabs.

The Cluster Summary tab displays information about the WCCP cluster as a whole. As a side effect of the WCCP protocol, each member of the cluster has information about all the others, so this information can be monitored from any appliance in the cluster.

Your router can also provide status information. See your router documentation.

Inline Mode

Aug 31, 2017

When you deploy a SD-WAN 4100/5100 appliance in inline mode, pairs of Ethernet ports on the appliance function as accelerated bridges. Traffic flows into one bridge port and out the other. When two sites with appliances communicate, TCP connections between the sites can be accelerated. Traffic that cannot be accelerated is passed through transparently, as if the appliance were not there.

For maximum reliability, the bridge pairs are equipped with a bypass feature that causes the two ports to be connected to each other should the appliance fail or lose power, allowing traffic to continue flowing even during such an outage.

Inline mode is currently recommended only for sites where WCCP is not practical, and which have a single WAN link, or have fully independent WAN links that do not use dynamic routing, load-balancing, or fail-over.

> ## Note
>
> For CloudBridge 4000 and 5000 WANOP appliances in release 7.1, the inline mode depends on the NetScaler “add interfacePair” command to isolate bridge traffic, ensuring that traffic from one bridge pair stays on that pair. In previous releases, this was done with VLAN definitions. Appliances that were provisioned with release 7.1 have this feature enabled by default. Appliances that were upgraded to release 7.1 without reprovisioning retain their old configuration. For more information about 4000 and 5000 WANOP appliances, see the CloudBridge 7.4 documentation.

Deployment Topology

Aug 31, 2017

The following figure shows a SD-WAN 4000/5000 appliance in inline mode.

Figure 1. Basic cabling for inline mode

As shown in the above figure, inline mode is a two-arm mode. For inline deployments, the NetScaler instance is configured in L2 (bridged) mode, but the accelerators are connected internally to the NetScaler instance in a one-arm configuration.

Inline mode is the easiest mode to configure. You connect one port of an accelerated pair to the WAN router and the other to the LAN network. The appliance transparently accelerates traffic flowing between the two ports, which to the rest of the network appear to be an Ethernet bridge.

You can also deploy the appliance to accelerate traffic from certain resources only, such as back-end servers, and not the traffic of the entire network. Such an arrangement reserves the appliance's resources for the selected traffic. In this case, you install the appliance on the branch network that includes the resources for with you want to accelerate traffic.

The following figure shows partial site acceleration:

Figure 2. Partial site acceleration

Port Affinity

Aug 31, 2017

Port traffic on a given bridge must be isolated to that bridge. In release 7.1, this is done as part of the provisioning process. It can also be done manually with the “add interfacePair” command in the NetScaler command line interface.

The following examples show how this command is used to create port affinity on all bridged pairs in the appliance:

SD-WAN 4000

add interfacePair 1 -ifnum 1/1 1/2
add interfacePair 2 -ifnum 1/3 1/4
add interfacePair 3 -ifnum 1/5 1/6
add interfacePair 4 -ifnum 1/7 1/8
add interfacePair 5 -ifnum 10/1 10/2

SD-WAN 5000

add interfacePair 1 -ifnum 10/1 10/2
add interfacePair 2 -ifnum 10/4 10/5
add interfacePair 3 -ifnum 10/6 10/7

VLAN Trunking

Aug 31, 2017

VLAN trunking is also known as tagged VLAN and 802.1Q tagging. The 802.1Q tagging enables a networking device to add information to a frame at Layer 2 to identify the VLAN membership of the frame. Tagging also enables network environments to have VLANs that span multiple devices. A device that receives the packet reads the tag and recognizes the VLAN to which the frame belongs.

When you configure tagging on bridged interfaces, the VLAN configuration must be identical on both ports of the bridge.

Tagged VLANs are not supported on the management interfaces (ports 0/1 and 0/2).

For example, if your WAN link uses VLAN 412, you declare VLAN 412 as a tagged VLAN in the NetScaler instance, and bind it to both ports of the bridge (such as ports 10/1 and 10/2), as shown in the example below.

Figure 1. Tagged VLANs for VLAN trunking. VLAN 412 is tagged

VLANs can be declared in either of two ways:

1. From the System > Settings > Configure NSVLAN Settings dialog box. This method declares a VLAN whose broadcast traffic is isolated from other VLANS. This method is recommended for the management subnet. It requires a restart to take effect.
   Note: This VLAN configuration method is neither synchronized nor propagated in high availability mode. Therefore, you must perform the configuration independently on each appliance of a high availability setup.
2. From the Create VLANs dialog box (reached from Network > VLANs > Add..). This method does not create an isolated broadcast domain, from traffic originating in the NetScaler instance until we bind the NetScaler IP addresses to the VLAN. Adding such a VLAN does not require a restart. This method is recommended for all VLANs except the management subnet.

Ethernet Bypass

Aug 31, 2017

The appliance includes a bypass feature for inline mode. In a power failure, a relay closes and the input and output ports become electrically connected. This feature allows the Ethernet signal to pass through from one port to the other, as if the appliance were not there. The appliance functions like a cross-over cable connecting the two ports.

Besides a power failure, any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the bridge ports of the appliance are inaccessible.

Bypass Considerations

- The bypass feature is disabled when the NetScaler instance is set to L3 mode. Because L3 mode is the factory default, inline mode should be configured before the appliance is placed in line with data traffic.
- The bypass feature is disabled when the appliance is in HA mode.
- A bypass event causes all bypass-enabled port pairs (except the loopback ports) to enter the bypass mode.
- The loopback ports never enter bypass mode.
- A bypass event occurs if the NetScaler instance or the bypass daemon in Dom-0 becomes unresponsive.
- A bypass event is not triggered by accelerators becoming unresponsive.
- The 1-Gigabit bypass ports are copper, and 10-Gigabit bypass ports are fiber ports.

Configuring Inline Mode

Aug 31, 2017

Note: These instructions are valid only for appliances that were provisioned from a factory-reset state with the release 7.1 Configuration Wizard. For appliances that have been updated to release 7.1, but not reprovisioned from a factory-reset state, see the release 7.0 documentation.

Basic configurationis performed by the release 7.1 Configuration Wizard. Additional configuration is required only if you use VLAN trunking on data passing through the bridges.

To configure inline mode

1. Do not attach the bridges to your traffic networks yet. Begin by provisioning the appliance with the configuration wizard.
2. If your traffic does not contain tagged VLAN traffic, skip to the last step of this procedure.
3. Navigate to the NetScaler instance at Configuration > NetScaler > Instances and click on the IP address of the NetScaler instance.
4. If the Citrix SD-WAN Connector Get Started page appears, ignore it.
5. Click Configuration > Network > VLANs > Add.
6. In the **Create VLAN**dialog box, configure the tagged VLANs to use bridge #1. In the **VLAN Id** field, enter the VLAN IDof the first tagged VLAN (VLAN1.1 on your worksheet). Under **Interfaces**, clear all the check boxes. Then, select **Active**and **Tagged**for the first port of the bridge (T9 on your worksheet) and the second port of the bridge (T10). Click **Create**.
7. Repeat the previous step for any remaining VLANs using bridge #1 (VLAN1.2, VLAN1.3, and so on).
8. Repeat for any additional bridge pairs.
9. Click Close.
10. Click Save to save your configuration.
11. Connect the bridges to your traffic networks. Configuration is complete.

Configuring the High Availability Setup on the Appliances

Aug 31, 2017

High availability(HA) works directly between the NetScaler instances of two SD-WAN 4100 or 5100 appliances. As shown in the configuration , the two appliances are configured almost identically, except for management network IP addresses.

> ## Note
>
> 1. The accelerator instances on the two appliances are not synchronized, and must be kept consistent manually. Take this into account when deciding whether to use HA.
> 2. For a smooth installation, install and test one appliance before adding the second one, noting all configuration changes, especially to the accelerator.
> 3. You must use the same aPA and signaling addresses on both appliances. However, all management subnet IP addresses must be unique on each appliances.

If the active appliance becomes unavailable, the passive appliance transparently takes over the function of the primary appliance. This is called “failover.” As a result, disruption of services over the network is minimal. After a failover, all clients must reestablish their connections to the managed servers, but the session persistence rules are maintained as they were before the failover.

HA is supported in all deployment modes, and the HA configuration procedure is the same for all modes. The two appliances should be running identical hardware, licensing, and software releases, and must be deployed identically, using the same deployment modes on the same subnets.

When you enable HA, the configuration of the primary appliance’s NetScaler instance is copied to the secondary appliance as part of the NetScaler HA synchronization process.

To configure a high availability setup of NetScaler instances

1. Complete the configuration for your chosen deployment mode (inline or WCCP). Note that all parameters for the external traffic subnet and the private traffic subnet are identical for both appliances, but some management subnet values are different on the two appliances.
2. Fully configure appliance #1 and test it thoroughly. If inline mode is used, do not connect the traffic network to the bridge ports on appliance #2.
3. Fully configure appliance #2. Note that some parameters are different on the two appliances: H1, H4, H5, H6, and H17 are used on the second appliance in place of M1, M4, M5, M6, and M17. Make sure that the accelerators are configured identically to the ones on appliance #1, and that both appliances have identical VLAN definitions in the NetScaler instances.
4. Access the NetScaler instance on appliance #1, by specifying its IP address (M17) in a web browser.
5. Log on to the NetScaler instance.
6. In the Navigation pane, expand the System node.
7. Select the High Availability node.
8. Click Add to configure a high availability setup of the NetScaler instances.
9. In the remote Node IP Address field of the High Availability Setup dialog box, specify the NSIP address of the NetScaler

instance of the other appliance #2 (H17 on your worksheet).

10. Click OK. The appliances are now configured as a high availability pair.

> **Note**
>
> To learn more about setting up high availability on a NetScaler instance, see the High Availability node of the Citrix eDocs website.

Evaluating the Configuration

Aug 31, 2017

Putting your appliance online in a production network requires special attention to prevent disruption or confusion, especially in a complex environment

Rollout Example

When deploying SD-WAN 4100/5100, the basic rollout decision is whether to activate the entire deployment at once or to roll it out in stages. In a large or complex environment, a phased approach avoids trouble, and the deployment can be extended at will. This type of approach calls for the use of WCCP. The following example illustrates one approach for such a site:

1. Configure the system as described in the installation procedure, except for the router. There, instead of setting up WCCP redirection for all incoming and outgoing WAN traffic, set it up for traffic to and from either a single remote site or a single IP address at that site. The remote site must already contain an enabled SD-WAN appliance.
2. The accelerator page. If not, check your WCCP configuration on the router and on the accelerators, and check your NAT definitions on the NetScaler instance by using **Monitoring: WCCP** page. If not, check your WCCP configuration on the router and on the accelerators, and check your NAT definitions on the NetScaler instance by using **nstrace**. If **nstrace** reveals an issue, and your definitions look correct, rebooting the appliance may resolve the issue.
3. Test acceleration between the new site and the remote site, with the remote site as the client side and the SD-WAN 4100/5100 equipped site as the server side, as described in General Monitoring.
4. If traffic does not appear, the router is not sending traffic to the SD-WAN 4100/5100 properly. The error could be in the Router configuration, the NetScaler configuration, or the SD-WAN WCCP configuration. Double-check these settings.
5. If traffic appears but is not accelerated, you might have a problem with asymmetrical routing, with not having a SD-WAN license installed, or with having acceleration disabled either globally or on the service classes associated with the traffic.
6. When all is working properly, test reverse connections, where a site on the SD-WAN 4100/5100 side is the client and the remote site is the server, if applicable.
7. If using NetScaler HA, save the configuration of the individual WCCP-enabled instances from the individual instances' GUIs, and save the configuration of the accelerator, do basic configuration manually, then restore the saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance. Be careful to avoid IP address conflicts.SD-WAN 4100/5100, do basic configuration manually, then restore the saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance. Be careful to avoid IP address conflicts.
8. If using NetScaler HA, save the configuration of the individual WCCP-enabled instances from the individual instances' GUIs, and save the configuration of the accelerator, restore these saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance - SD-WAN 4100/5100, restore these saved configurations, first of the accelerators as a whole, and then restore the two WCCP-enabled instances. Once this is done (and NetScaler HA is enabled), test failover by powering down the primary appliance.
9. Expand the scope of acceleration to include more remote sites, and repeat the above testing. When doing so, also examine the **Monitoring: System Load** page, especially during peak periods, to verify that the SD-WAN 4100/5100 is not heavily loaded.
10. Continue this process until the entire WAN is being accelerated.

Monitoring

Use the SD-WAN 4100/5100 GUI to monitor traffic after you configure a LAN link and a WAN link. SD-WAN 4100/5100 allows a very simple link definition.

Configuring the Links

To enable monitoring, you must first configure one LAN link and one WAN link. To do so, edit the default links on the Configure: Links page as follows:

1. Edit one link so its name is "LAN," its type is "LAN," and its speed is 10 Gbps in both directions. Delete its existing filter rule, then click Add Rule, and then click Save to save a link definition that matches all traffic.
2. Edit the other link so that its name is "WAN," its type is "WAN," its speed is 95% of the aggregate speed of your site's WAN links in each direction. Delete its existing filter rule, then click Add Rule, and then click Save to save a link definition that matches all traffic.

To verify that link configuration is working correctly, traffic must be flowing. If the network does not have enough traffic to fill the WAN link to capacity, run test traffic to fill the network to capacity. Then look at the link reports on the Reports: Link Usage tab.

General Monitoring

1. If WCCP is configured, verify that the service groups are in operation and the routers are redirecting traffic. (Note that the SD-WAN WCCP page packet counts are not present in SD-WAN 4100/5100. Check traffic by other means, such as on the Monitoring: Active Connections page, and on the router.)
2. On the **remote** SD-WANs, verify that outgoing connections are being accelerated, and that all accelerated connections to the datacenter report the same Partner Unit on the remote appliance's Monitoring: Connections page. When load-balancing is working properly, all outgoing accelerated connections show the same Partner Unit. (However, incoming accelerated connections might show different units.)
3. Double-check remote SD-WANs for correctly set bandwidth limits, to prevent remote issues from being misidentified as datacenter issues.
4. Generally monitor the SD-WAN 4100/5100 unit for alerts.
5. In the broker UI, use the Dashboard, the Monitoring: Remote Partners, and perhaps the Monitoring: Appliance Load pages to monitor the overall activity and load of the system.

Troubleshooting Tips

Aug 31, 2017

While most installations complete smoothly, some installations require knowledge of the appliance's internal structure or the use of little-known features before you can perform additional monitoring and troubleshooting. These troubleshooting tips provide information and techniques that allow a more in-depth analysis of the appliance.

Understanding Internal Addresses

Some reports show addresses on the private subnets within the SD-WAN 4100/5100, so it's good to know what these addresses mean. These subnets connect the virtual machines together, without connecting to external ports.

All these addresses are on the local link subnet 169.254.0.0/16, described in RFC3927. This address space is segmented into three partly overlapping subnets: system management, private traffic, and accelerator management subnets.

Virtual machines in the SD-WAN 4100 and 5100. The system management subnet is not shown in this diagram. The traffic shaper manages traffic from all accelerators and is controlled via the accelerator GUI.

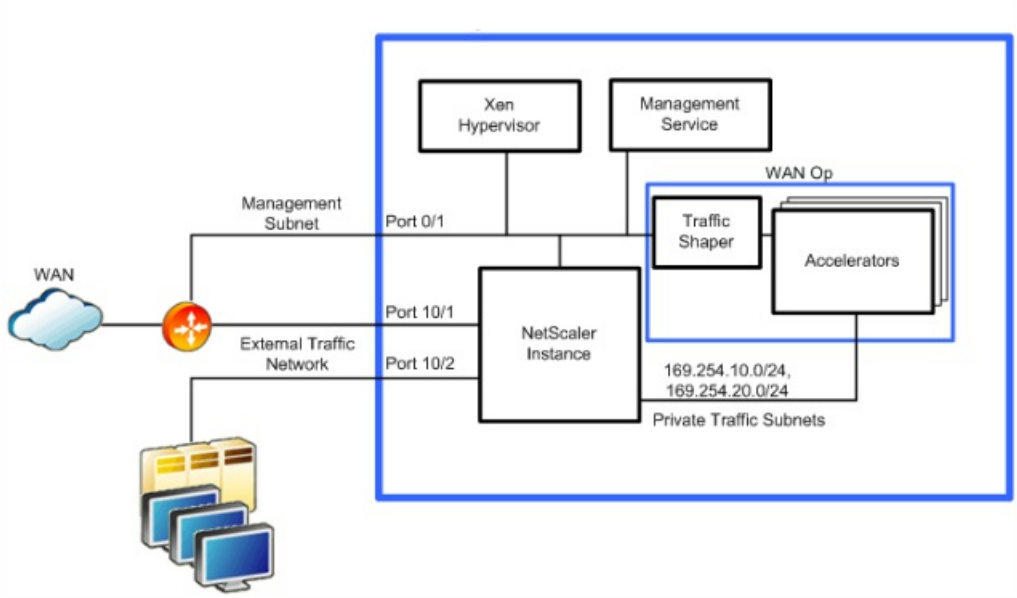

System Management Subnet

Function	Address
Management Service	169.254.0.10/16
NetScaler Instance	169.254.0.11/16
XenServer	169.254.0.1/16

Private Traffic Subnet

Function	Address
apA IP, accelerators 1-8	169.254.10.21/24 - 169.254.10.28/24
apA Signaling IP, accelerators 1-8	169.254.10.121/24 - 169.254.10.128/24
NetScaler Instance	169.254.10.11/24

Accelerator Management Subnet

Function	Address
Accelerator unified management IP (controls all accelerators)	169.254.0.20/24
Primary Port IP, accelerators 1-8	169.254.0.21/24 - 169.254.0.28/24

Checking and Correcting Accelerator Instance Status

Sometimes an error message may indicate an issue with one of the virtual machines in the appliance. To check their status, go to the System Configuration page and select an Instance view of either the SD-WAN or NetScaler subsystems. For example, the SD-WAN page.

- A fully active instance will show a green circle for VM State, Instance State, and Licensed.
- Your appliance may have more instances present than are licensed; ignore the unlicensed instances.
- If the VM State or Instance State of the remaining instances are not green, use the “Rediscover” action to attempt to bring these instances back into operation.

You can also get detailed information for each instance:

- Every instance should have a Status of “Inventory from SD-WAN Instance completed.
- Every instance should be running the same version of the software.
- Every instance should have the netmask (255.255.255.0) and gateway (169.254.0.20).
- Instances that show an uptime shorter than other instances have rebooted since the last whole-system reboot.

Logging Into the NetScaler Instance

Sometimes it is useful to log into the NetScaler instance to check its status or do configuration. You can log into the NetScaler instance from the NetScaler Instances page of the bview, as shown below. Click the **IP Address** link.

You can also log into the NetScaler instance directly from your browser if you know its IP address on the management port (port 0/1).
Once logged in, you will see the NetScaler GUI, which identifies itself as NetScaler VPX at the top of the page.

This is the standard NetScaler user interface. Using monitoring features is safe. Configuration changes should be made with caution, as the SD-WAN 4100/5100 makes undocumented assumptions about how the NetScaler instance is configured.

Using Ping and Traceroute

The ping and traceroute utilities are not available on the accelerator instances, as they are on other SD-WAN products. Instead, you can use the equivalent features on the NetScaler instance, using the Diagnostics page as shown below.

These features will work over your external network and on the appliance's internal subnets.

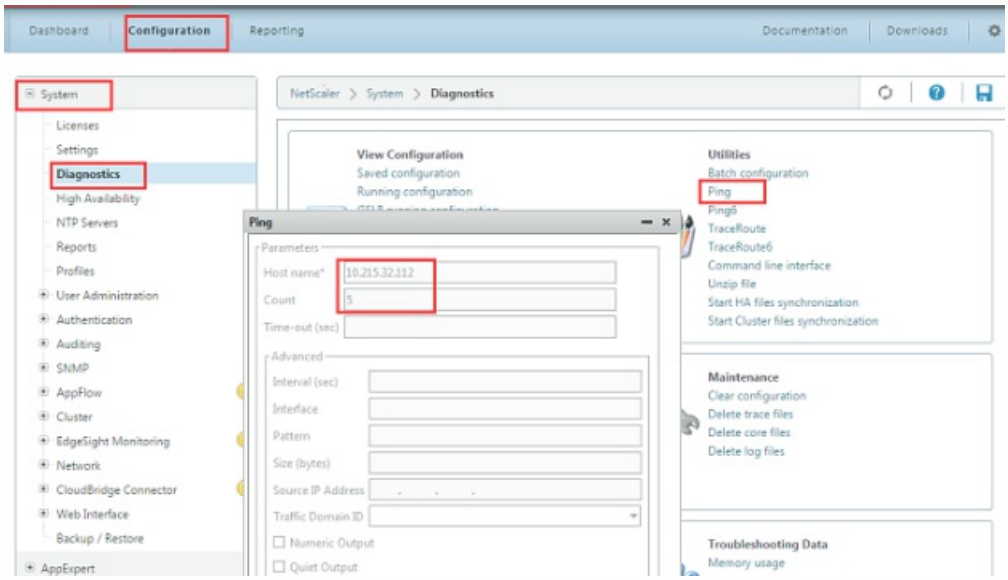

Using the System Dashboard

Unlike the SD-WAN Dashboard, the System Dashboard page is devoted largely to hardware monitoring.

- The System Health tables show a status summary, with a Details link for expanded information in graphical form.
- The Events tables show a status summary, with a Show Events link to see the related log entries.
- If several ports are marked as Down, which is only an error if a cable is supposed to be present. Most appliances have several unused ports.
- Fail To Wire lists FTW Disabled for all ports. This means that the network bypassing feature is not enabled on this appliance. Examination of the FTW Events showed that there were no actual events, indicating that the feature is probably disabled.

For each warning or error, additional details are available through the Details links or Show Events buttons.

Logging In To Different Instances via SSH

You can log into some of the virtual machines from the management port (port 0/1) using an ssh utility (such as PuTTY on Windows), logging in either as root or nsroot and using the administrative password. This will give you a shell prompt.

The most common use for logging on via SSH is to restore the IP address of an instance, typically the management service, that has become unreachable due to misconfigured network parameters. Otherwise, SSH is not recommended, as configuration changes can render the appliance unstable or unusable.

If neither of the two instances below are accessible over the network, you can log into the XenServer instance using the RS-232 port, which will give you a shell prompt.

Once logged into one of these virtual machines, you can use SSH from the shell prompt to reach the NetScaler instance or the accelerator at the appropriate 169.254.x.x address.
The usual UNIX/Linux commands are available, including the vi text editor.

Instance	Login	Password	Actual Username
Management Service	nsroot	Admin password	root
Management Service	root	Admin password	root
XenServer	nsroot	Admin password	nsroot
XenServer	root	Admin password	root

Monitoring Individual Accelerator Instances

Logging into the accelerator GUI IP allows you to manage all the accelerator instances as a unit. Changes are automatically propagated to all the accelerator instances.

On rare occasions, you may wish to troubleshoot individual accelerator instances.

The login for the instances is admin. The password is the same admin password as is used on the other instances.
This is recommended for monitoring, not for making permanent changes, since any parameter you set in an instance may be overwritten later by the synchronization process. To do this, use the following URL's:

Accelerator Instance	URL
1	https://<accelerator_ip>:4001
2	https://<accelerator_ip>:4002
8	https://<accelerator_ip>:4008

Using Individual Elements of the Update Bundle

The update bundles distributed by Citrix are in a simple .tgz format (a tar archive compressed with gzip). It is sometimes useful to extract individual components from the archive, rather than going back to the the Citrix Web site and downloading them individually. This is most commonly useful with the management service (build-svm*.tgz) or the accelerator release (orbital*.bin).

The update bundle can be managed by tar/gzip or by archiving utilities like 7-zip.

NetScaler SD-WAN 1000, 2000, and 2100 Standard Edition Appliances

Aug 09, 2017

The SD-WAN Standard Edition 1000, 2000, and 2100 appliances combine virtualized instances of the SD-WAN appliance.

The SD-WAN Standard Edition 1000, 2000, and 2100 appliances are based on the Citrix branch architecture, which supports multiple virtual machines. All branch appliances contain a SD-WAN Standard Edition instance and management service instance.

The SD-WAN instance is typically used in inline mode, with the SD-WAN instance interposed between the WAN router and the LAN. The SD-WAN instance can also be deployed in virtual inline mode.

The appliance has two modes; two-port mode and four-port mode, which determine how ports 1/3 and 1/4 are used.

Netscaler SD-WAN 1000 SE

Mar 23, 2018

The SD-WAN 1000-SE with platform has a quad-core processor and 32 GB of memory. This platform has a bandwidth of up to 100 Mbps.

The following figure shows the front panel of a SD-WAN 1000-SE appliance.

Figure 1. Citrix NetScaler SD-WAN 1000-SE front panel

The front panel of the SD-WAN 1000-SE appliance has a power button and five LEDs.

The power button is used to switch the appliance on or off.

The reset button restarts the appliance.

The LEDs provide critical information related to different parts of the appliance.

- Power Fail – Indicates the power supply unit has failed.
- Information LED – Indicates the following:

Status	Description
Continuously ON and red	The appliance is overheated. (This might be a result of cable congestion.)
Blinking red (1Hz)	Fan failure, check for an inoperative fan.
Blinking red (0.25Hz)	Power failure, check for the non-operational power supply.
Solid blue	Local UID has been activated. Use this function to locate the server in a rack mount environment.
Blinking blue (300 m/s)	Remote UID is on. Use this function to identify the server from a remote location.

- NIC1 and NIC2 – Indicate network activity on the LAN1 and WAN1 ports.
- HDD – Indicates the status of the hard disk drive.

- Power – Indicates that the power supply units are receiving power and operating normally.

The following figure shows the back panel of a SD-WAN 1000-SE appliance.

Figure 2. Citrix NetScaler SD-WAN 1000-SE appliance , back panel

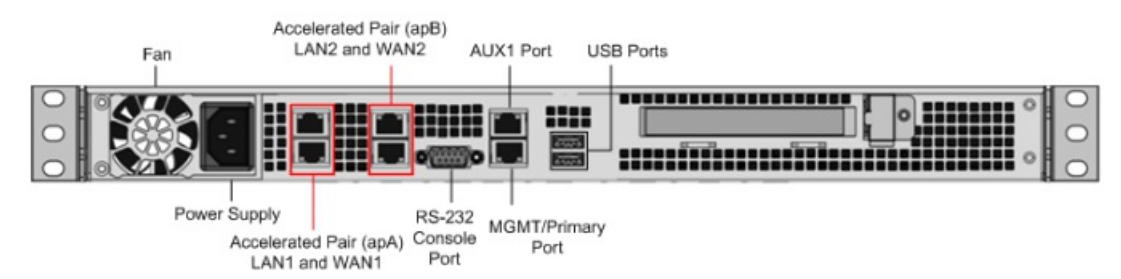

The following components are visible on the back panel of an SD-WAN 1000-SE appliance:

- Cooling fan.
- Single power supply, rated at 200 watts, 110-240 volts.
- Accelerated pairs of Ethernet ports (apA and apB).
- RS-232 serial console port.
- One AUX Ethernet port and one management port.
- Two USB ports.

Power on Appliance After a Graceful Shut Down

To power on the appliance after a graceful shut down:

1. Connect a Serial console cable to the rear of the appliance and to the serial port on a management laptop.

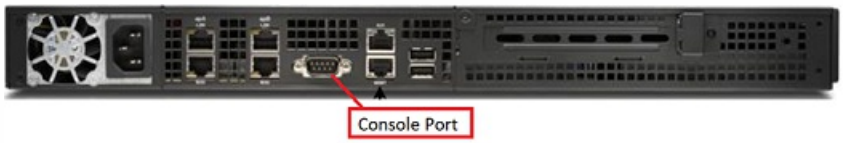

2. On the management laptop, restart a putty session using the following configuration settings:

- Serial line: COM1
- Speed: 9600

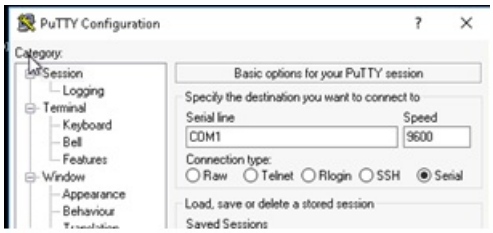

3. Power on the appliance and as it is booting, press the following key in the Putty session to enter the BIOS configuration screen.

   Keypress: **DEL**

4. When in the BIOS, navigate to,

- Advanced Tab > **Select**

- Boot Feature > **Enter**

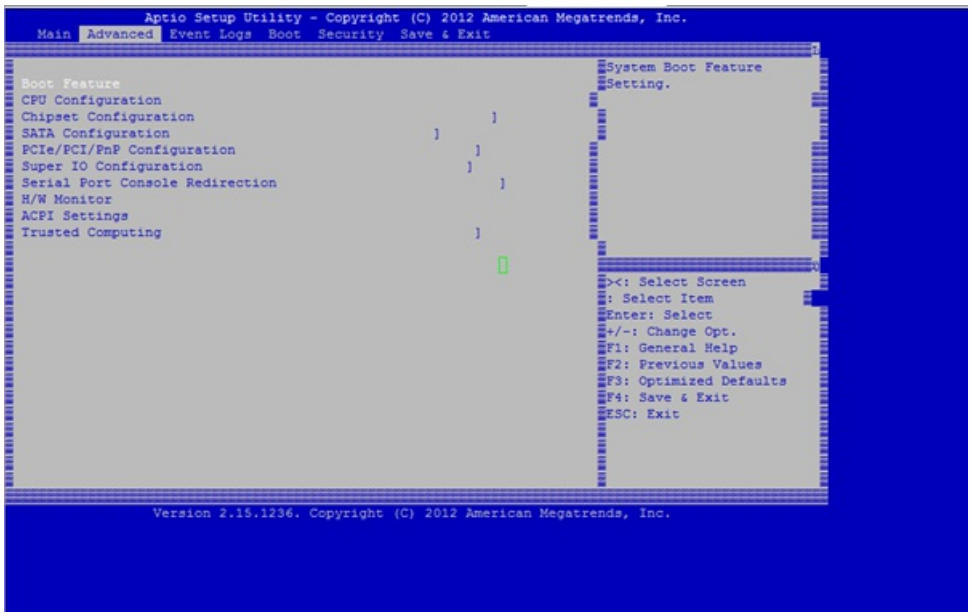

5. When in the Boot Feature screen, change the value of the parameter **Restore on AC Power Loss;** from **Last State** > **Power ON**.

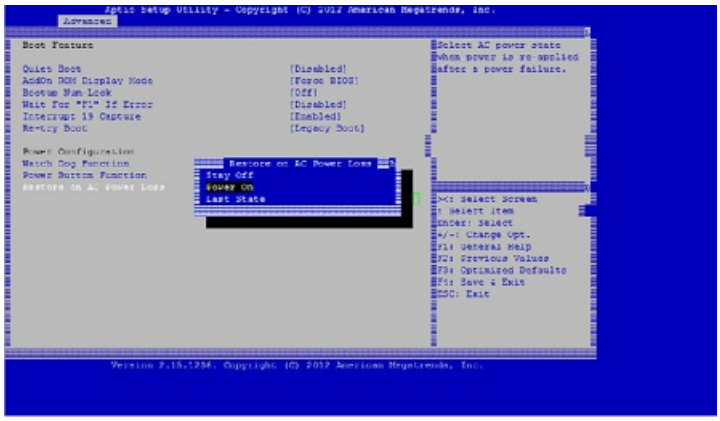

6. Navigate to Save and Exit.

- Select **Save changes** and **Reset**
- Select **Yes**

Allow the system to restart. This takes approximately five minutes.

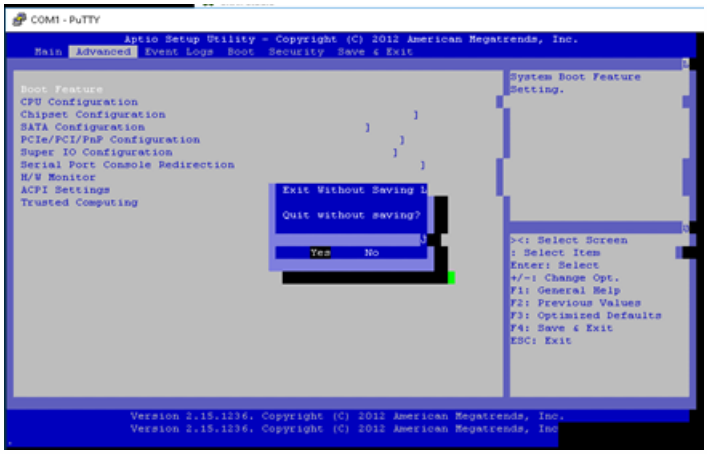

7. After the appliance is powered on, login to the appliance management instance (SVM). The default IP address for the appliance is: 192.168.100.1, user name is: admin/password.

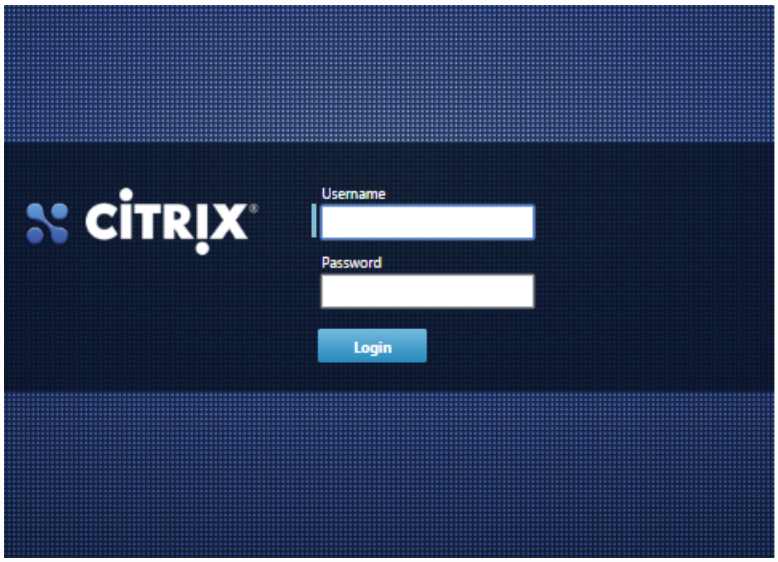

8. In the SD-WAN appliance GUI, navigate to **Configuration** > **Maintenance** > **Reboot Appliance**. Allow the appliance to fully shut down. Ensure that there are no power lights on the appliance when the shut down process has completed.

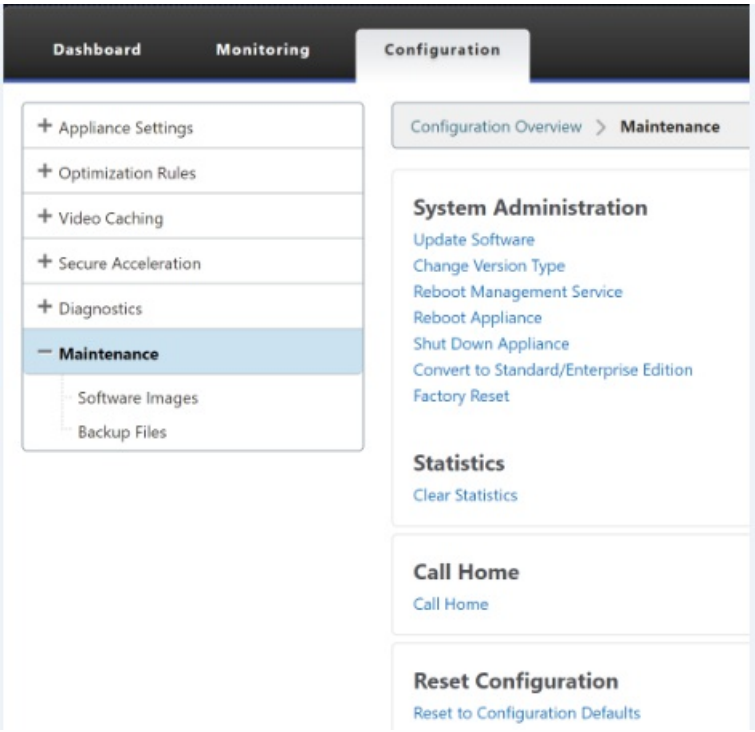

9. Power on the appliance to confirm that the BIOS configuration change has been applied successfully. This can be either done through the APC intelligent PDU Web Management console or by physically pulling the power cable out of the shut down SD-WAN appliance, waiting for 10 seconds and then plugging it back in again. The appliance power ups automatically from all shut down scenarios.

NetScaler SD-WAN 2000 SE

Aug 09, 2017

The Citrix NetScaler SD-WAN 2000-SE platform is a 1U appliance with one quad-core processor and 24 gigabytes (GB) of memory.

The following figure shows the front panel of the SD-WAN 2000-SE appliance.

Figure 1. Citrix NetScaler SD-WAN 2000-SE appliance, front panel

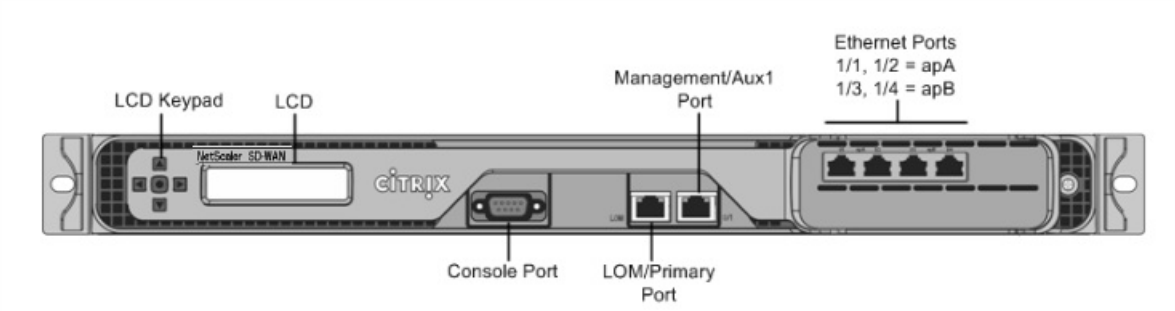

SD-WAN 2000-SE appliance has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, numbered 0/1, and named PRI (primary). The management port is used to connect directly to the appliance for system administration functions. You can use this port for initial provisioning of Virtual WAN. The LOM port also operates as a management port.
- Four 10/100/1000Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two *accelerated pairs*, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), and 1/3 and 1/4 are accelerated pair B (apB).

Figure 2. Citrix NetScaler SD-WAN 2000-SE appliance, back panel

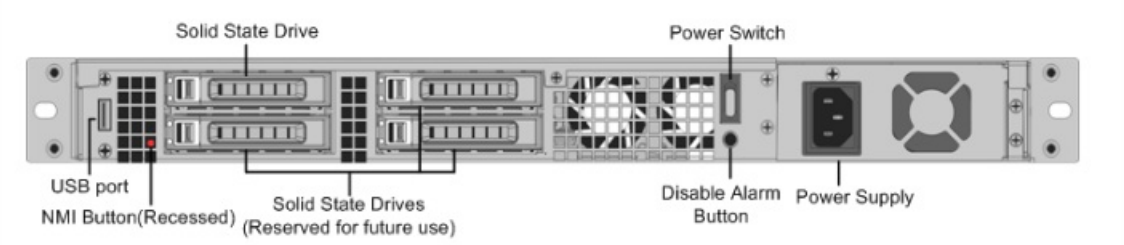

The following components are visible on the back panel of the SD-WAN 2000-SE appliance:

- 600 GB removable solid-state drive, which stores the appliance's software and user data, and 1 TB hard disk drive.
- Power switch. Press the switch for five seconds to switch off the power.
- USB port (reserved for a future release).

- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You must use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Single power supply, rated at 300 watts, 100-240 volts.

NetScaler SD-WAN 2100 SE

Jun 12, 2018

The Citrix NetScaler SD-WAN 2100-SE platform is a 1U appliance with 8 core processor and 32 GB (GB) of memory.

The following figure shows the front panel of the SD-WAN 2100-SE appliance.

Figure 1. Citrix NetScaler SD-WAN 2100-SE appliance, front panel

SD-WAN 2100-SE appliance has the following ports:

- An RS232 serial console port.
- Copper Ethernet (RJ45) Port called the Lights out Management (LOM) port labeled LOM, and management port labeled 0/1. You can use these ports to remotely monitor and manage the appliance independently of the appliance's software.
- USB ports.
- Four 1000Base-TX copper Ethernet ports.
- Four 1GE SFP ports.

Port Labels - old 2100-SE Front Bezel	Description
LOM	Lights out management
0/1	Management

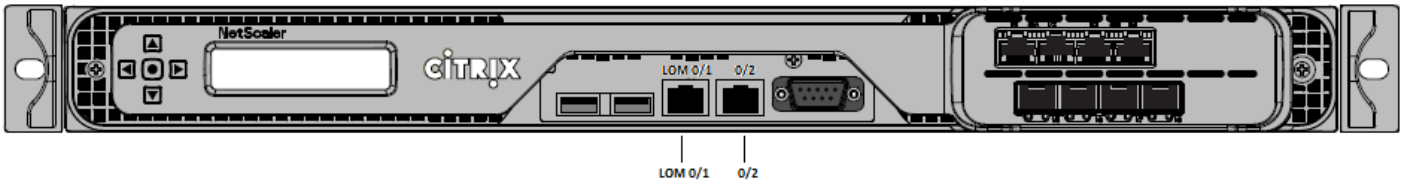

Port Labels - new 2100-SE Front Bezel	Description
LOM 0/1	Lights out management
0/2	Reserved for future use (Management)

- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port labeled lights out management and 0/1. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, labeled 0/2. This management port cannot be used for system

administration functions. This port is reserved for future use.

Figure 2. Citrix NetScaler SD-WAN 2100-SE appliance, back panel

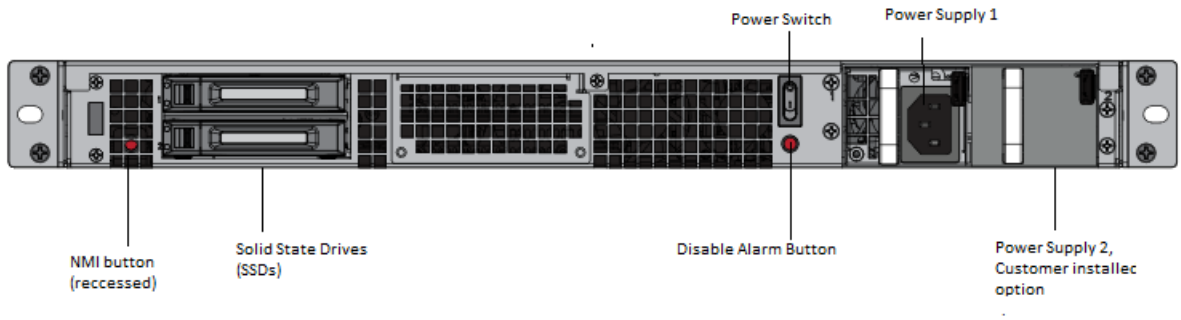

The following components are visible on the back panel of the SD-WAN 2100-SE appliance:

- 240 GB removable solid-state drive, which stores the appliance's software and user data, and 1 TB hard disk drive.
- Power switch, which switches power to the appliance on or off. Press the switch for five seconds to switch off the power.
- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You can use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Single power supply, rated at 450 watts, 100–240 volts.

Upgrade 2100 SE Appliance to 2100 EE Appliance

> ### Important
>
> To use EE functionality, you need SD-WAN release 10.0 on the 2100 EE appliances and install EE licenses. The SD-WAN release 9.3 and higher software release versions support 2100 EE.

- 2100 SE ships with only one SSD (240GB) and one blank carrier.
- If you want to upgrade to an EE appliance, you can order the kit.
  - Kit includes an additional SSD (480GB) and appropriate license for EE.
- Upon receiving the kit, install the new 480GB drive in the empty slot (leaving original SSD as is), upgrade to SD-WAN release 10.0, and apply new EE license.

Insert Solid State Drive (SSD)

1. Insert the required SSD in the standard edition appliance. For instructions about how to insert SSD, see Solid State Drive (Field Replaceable Unit).

- 2100 SE appliance requires 480 GB or more SSD.

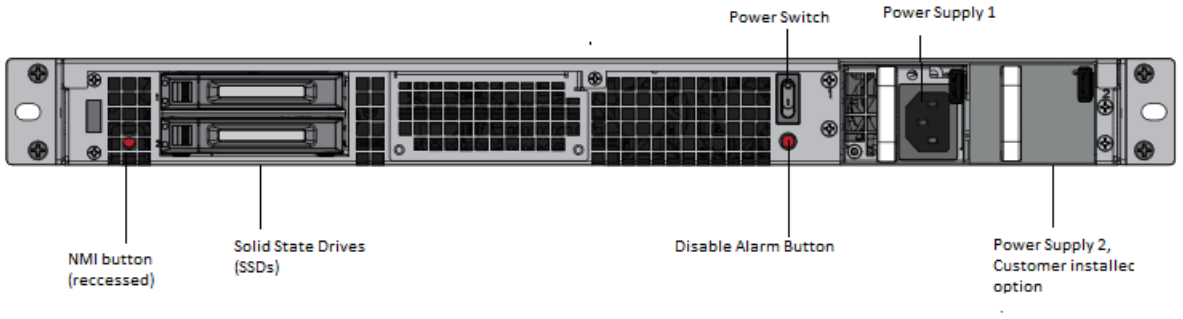

2. Restart the appliance through the SD-WAN web management interface.

3. Ensure that the software release version installed on the appliance is SD-WAN release version 9.3.3.

Appliances shipped with software release version 9.3.3

For appliances shipped with 9.3.3 or new manufacture image. Follow the steps provided here; upgrade new appliance.

1. Upgrade the network using single step upgrade to software release version 10.0 or later.
2. Install the Enterprise Edition platform license. For license information, see the Citrix SD-WAN product downloads site.

Configure Management IP Address Using Serial Console

1. Access serial console of the appliance.
2. Log in using the **root/nsroot** credentials.
3. Type the **ssh admin@169.254.0.60 -l admin** command.
4. Type password: **password**.
5. Type the **management_ip** command.
6. Type the **set interface 192.168.100.1 255.255.255.0 192.168.100.254** command.
7. Type the **apply** command.

Summary of Hardware Specifications

Aug 09, 2017

The following table summarizes the specifications of the SD-WAN 1000-SE, 2000-SE, and 2100-SE hardware platforms.

Specifications	SD-WAN 1000-SE	SD-WAN 2000-SE	SD-WAN 2100-SE
Bandwidth	Upto 100 Mbps	Upto 300 Mbps	Upto 1.5 Gbps
Total sessions, Max Virtual Paths (Static/Dynamic)	10,000	20,000	128/32
Processor	4 Cores	4 Cores	8-Core 2.1GHz
Total Disk Space	1X480 GB SSD	1X800 GB SSD	1X240 GB SSD
RAM	16 GB	24 GB	32 GB
Network Interfaces	2 pair with bypass 10/100/1000 2 GigE ports for Management and AUX ports	2 pair with bypass 10/100/1000	2 pair with bypass of 1G 4 x 1GE SFP 2 GigE ports for Management and AUX ports
Power Supplies	1	1	1 module and 1 optional FRU
Rack Units	1U	1U	1U
System Width	EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks
System Depth	10" (25.4 cm)	28" (72 cm)	28" (72 cm)
System Weight	8.5 lbs (3.9 kg)	32 lbs (14.5 kg)	32 lbs (14.5 kg)
Shipping dimensions and weight	26 L x 18.5 W x 6.5" H 14.5 lbs	32 L x 23.5 W x 7.5" H 39 lbs	33" L x 24" W x 8" H 40 lbs
Voltage	100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz
Power consumption (Max.)	200 W	300 W	450 W
Operating Temperature (degree Celsius)	10–35	0-40	0-40
Non-operating			

(degree Celsius)			
Specifications	**SD-WAN 1000-SE**	**SD-WAN 2000-SE**	**SD-WAN 2100-SE**
Allowed Relative Humidity	8% – 90% non-condensing	5%–95% non condensing	20%-80% non-condensing
Safety certifications	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)
Electromagnetic and susceptibility certifications	FCC Class A, EN 55022 Class A, EN 61000-3-2/-3-3, CISPR 22 Class A	FCC (Part 15 Class A), CE, C-Tick, VCCI-A, CCC, KCC, NOM, SASO, SABS, PCT	FCC (Part 15 Class A), CCC, KCC, NOM, CITC, EAC, DoC, CE, VCCI, RCM
Environmental certifications	RoHS, WEEE	RoHS, WEEE	RoHS, WEEE

Installing the Appliance

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. You can also connect the appliance to a computer through Ethernet port for initial configuration. On SD-WAN 1000-SE appliance, this port is labeled as MGMT (management) port and on SD-WAN 2000-SE and SD-WAN 2100-SE appliances, the port is labeled as PRI (primary) port. To complete the installation, you switch on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 09, 2017

A SD-WAN 1000-SE or 2000-SE appliance requires one rack unit. Both are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

Rack Mounting an SD-WAN 1000-SE Appliance

Aug 09, 2017

SD-WAN 1000-SE appliance is not shipped with rails. You can mount the appliance to the rack by using the front mounting ports.

-
-
-
-

-
-

-
-
-
-
-
-

-
-
-
-

-
-
-
-
-
-
-

-
-
-
-
-

!- Use a ping (ICMP echo) to see if appliance is connected track 123 rtr 1 reachabilit y ! rtr 1 type echo protocol Iplcmpecho 192.168.1.200 schedule 1 life forever start-time now

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.10.10.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 172.68.1.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended client_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 192.168.1.200 20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 192.168.1.200 10 track 123
```

```
! This example does not use health-checking.
! Remember, health-checking is always recommended,
! so this is a configuration of last resort.
!
!
ip cef
!
interface FastEthernet0/0
ip address 20.20.20.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 171.68.2.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.2.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.2.1
!
ip access-list extended client_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
set ip next-hop 192.168.2.200
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop 192.168.2.200
!_
```

```
!
 ip access-list extended client_side
 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
 ip access-list extended wan_side
 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
```

-
-

-
-

-
-

-
-
-

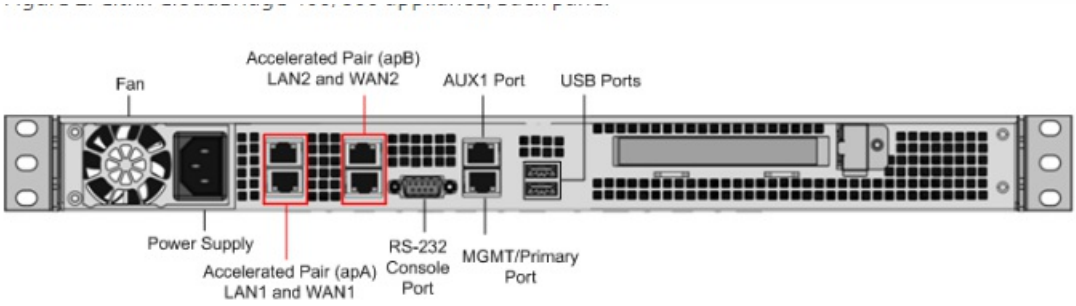

-
-
-

-
-
-
-
  -

-
-

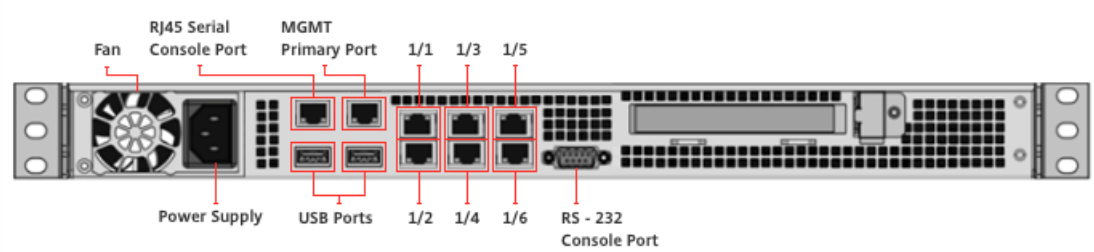

-
-
-
  -
  -
  -
-
-
-

-
-
-
-
-

-
-
-

-
-

Cautions and Warnings

Aug 09, 2017

Updated: 2014-02-06

Caution: During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.
Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Remove all jewelry and other metal objects that might come into contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.
- Use a regulating, uninterruptible power supply (UPS) to protect the appliance from power surges and voltage spikes, and to keep the appliance operating in case of power failure.
- Never stack the appliance on top of any other server or electronic equipment.
- All appliances are designed to be installed on power systems that use TN earthing. Do not install your device on a power system that uses either TT or IT earthing.
- Make sure that the appliance has a direct physical connection to the earth during normal use. When installing or repairing an appliance, always make sure that the ground circuit is connected first and disconnected last.
- Make sure that a fuse or circuit breaker no larger than 120 VAC, 15 A U.S. (240 VAC, 16 A international) is used on all current-carrying conductors on the power system to which your appliances are connected.
- Do not work alone when working with high voltage components.
- Always disconnect the appliance from power before removing or installing any component. When disconnecting power, first shut down the appliance, and then unplug the power cords of all the power supply units connected to the appliance. As long as the power cord is plugged in, line voltages can be present in the power supply, even when the power switch is OFF.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload. Always unplug any appliance before performing repairs or upgrades.
- Do not overload the wiring in your server cabinet or on your server room rack.
- During thunderstorms, or anticipated thunderstorms, avoid performing any hardware repairs or upgrades until the danger of lightning has passed.
- When you dispose of an old appliance or any components, follow any local and national laws on disposal of electronic waste.
- To prevent possible explosions, replace expired batteries with the same model or a manufacturer-recommended substitute and follow the manufacturer's instructions for battery replacement.
- Never remove a power supply cover or any sealed part that has the following label:

- Determine the placement of each component in the rack before you install the rails.
- Install the heaviest appliance first, at the bottom of the rack, and then work upward. Distribute the load on the rack evenly. An unbalanced rack is hazardous.
- Allow the power supply units and hard drives to cool before touching them.
- Install the equipment near an electrical outlet for easy access.
- Mount equipment in a rack with sufficient airflow for safe operation.
- For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 09, 2017

SD-WAN 400 and 410 SE appliances require one rack unit. These appliances are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

To mount a SD-WAN appliance, you must first install the rails and then install the appliance in the rack, as follows:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

To remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the locking tabs until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

To attach the inner rails to the appliance

1. Position the right inner rail behind the ear bracket on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws.
4. Repeat steps 1 through 3 to install the left inner rail on the left side of the appliance.

To install the rack rails

1. Position the rack rails at the desired location in the rack, keeping the sliding rail guide facing inward.
2. Snap the rails to the rack.
   Note: Make sure that both rack rails are at same height and that the rail guides are facing inward.

To install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides, and push the appliance into the rack rails until it locks into place.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.
   Note: The illustration in the following figure might not represent your actual appliance.
   Figure 1. Rack Mounting the Appliance

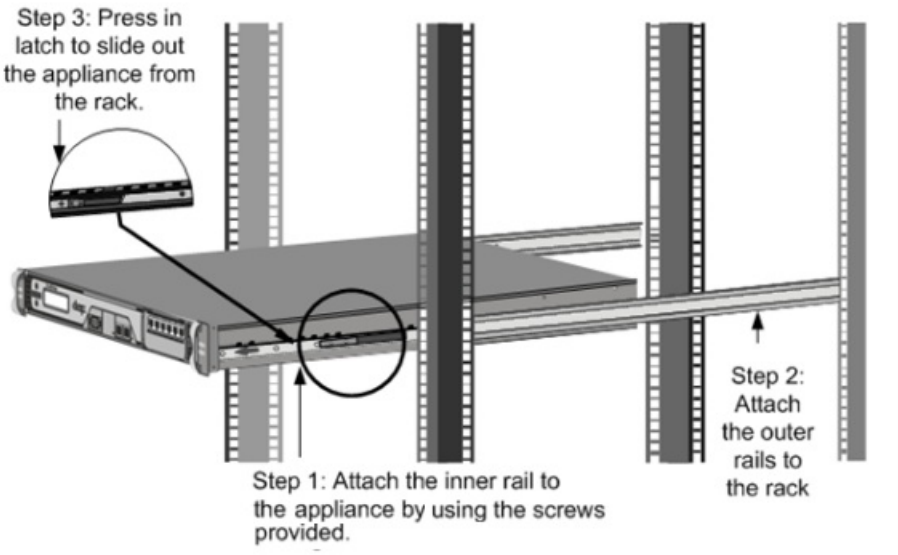

Connecting the Cables

Aug 09, 2017

When the appliance is securely mounted on the rack, determine which ports you should use. You are then ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Danger: Remove all jewelry and other metal objects that might come in contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly, and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

A typical installation using a single accelerated-bridge uses three Ethernet ports (the Primary port and apA).

The appliance has two motherboard ports. The Primary port is used for initial configuration.

On a SD-WAN 410-SE appliance, ports 1/1 and 1/2 are the accelerated pair A (apA) bridge ports, ports 1/3 and 1/4 are the apB ports, and ports 1/5 and 1/6 are the apC bridge ports.

Updated: 2014-01-20

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port of the appliance.
2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Updated: 2014-01-20

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port.
   Note: To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.
2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Updated: 2014-01-20

A SD-WAN appliance has one power supply, unless you have installed a second. A separate ground cable is not required, because the three-prong plug provides grounding. Provide power to the appliance by installing the power cord.

To connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply.
2. Connect the other end of the power cable to a standard 110V/220V power outlet.

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Depending on the appliance, press the ON/OFF toggle power switch or the power button to switch on the appliance.

Caution: Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Initial Configuration

Aug 09, 2017

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

You can connect the ethernet cable to the appliance management port, and login to the web management interface using the default IP address.

To perform initial configuration:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Perform additional configuration for Windows.
- Assign management IP address through the serial console.
- Troubleshoot initial configuration issues.

By default, the initial configuration deploys the appliance in inline mode.

> ## Note
>
> On the SD-WAN 410-SE appliance, the default static IP address is 192.168.100.1. It also has DHCP enabled (by default) for Management access. When you enable the SD-WAN service, or install SD-WAN license using default static IP (192.168.100.1), this IP is lost.
>
> You should assign a static IP address or be notified of the DHCP IP address that is assigned to this appliance. In order to see the DHCP IP address, you can log into the SD-WAN command line interface using admin/password, to view the management IP address. The LCD console on the 410-SE appliance will not display the IP address.

Prerequisites

Aug 09, 2017

Note

The default static IP address is 192.168.100.1. The static IP address has DHCP enabled (by default) for Management access.

When you enable the SD-WAN service, or install SD-WAN license using the default static IP (192.168.100.1), this IP is disabled, and you can obtain DHCP address within the network.

Always assign a Static IP address or be aware of the DHCP address that is assigned to the appliance. To view the DHCP IP address, you can login to SD-WAN CLI (using admin/password) that displays the management IP or go to **Configuration** -> **Appliance Settings** > **Network Adapters** > **Ethernet**.

1. Ensure that you have permanent DHCP address assigned to SD-WAN appliances.
2. The DHCP address should be associated to the management NIC address.
3. Connect the management NIC address to the DHCP enabled LAN or reboot the applaince when ready.

Before you begin configuring the appliance, make sure that the following prerequisites have been met:

- You should have physical access to the appliance.
- In the Worksheet, record all IP addresses and other values you would use to configure the appliance. Preferably, print out the worksheet before you start the configuration process.
- You should already have a SD-WAN license key from Citrix, sent in an email. If you are using remote licensing, you need the IP address of the licensing server.
- WAN Send and Receive Speeds.

Setting up the SD-WAN Appliance

Aug 09, 2017

To set up your NetScaler SD-WAN Appliance hardware, see the instructions documented in the Setting up the Appliance Hardware section.

Deployment Modes

Aug 09, 2017

A SD-WAN appliance acts as a virtual gateway. It is neither a TCP endpoint nor a router. Like any gateway, its job is to buffer incoming packets and put them onto the outgoing link at the right speed. This packet forwarding can be done in different ways, such as inline mode, virtual inline mode, and WCCP mode (WANOP appliance only). Although these methods are called *modes*, you do not have to disable one forwarding mode to enable another. If your deployment supports more than one mode, the mode that the appliance uses is determined automatically by the Ethernet and IP format of each packet.

Because the appliance supports different forwarding modes and different kinds of non-forwarded connections, it needs a way of distinguishing one kind of traffic from another. It does so by examining the destination IP address and destination Ethernet address (MAC address), as shown in table below. For example, in inline mode, the appliance is acting as a bridge. Unlike other traffic, bridged packets are addressed to a system beyond the appliance, not to the appliance itself. The address fields contain neither the appliance's IP address nor the appliance's Ethernet MAC address.

In addition to pure forwarding modes, the appliance has to account for additional types of connections, including management connections to the GUI and the heartbeat signal that passes between members of a high-availability pair. For completeness, these additional traffic modes are also listed in table below.

Table 1. How Ethernet and IP Addresses Determine the Mode

Destination IP Address	Destination Ethernet Address	Mode
Not appliance	Not appliance	Inline or Pass-through
Not appliance	Appliance	Virtual Inline or L2 WCCP
Appliance	Appliance	Direct (UI access)
Appliance (VIP)	Appliance	High-Availability. Proxy mode
Appliance (Signaling IP)	Appliance	Signaling Connection (SD-WAN plugin Signaling Connection (SD-WAN plugin, Secure Peer) or Redirector Mode Connection (SD-WAN plugin)
All modes can be active simultaneously. The mode used for a given packet is determined by the Ethernet and IP headers.		

The forwarding modes are:

- **Inline mode,** in which the appliance transparently accelerates traffic flowing between its two Ethernet ports. In this mode, the appliance appears (to the rest of the network) to be an Ethernet bridge. Inline mode is recommended, because it requires the least configuration.
- **Virtual inline mode,** in which a router sends WAN traffic to the appliance and the appliance returns it to the router. In this mode, the appliance appears to be a router, but it uses no routing tables. It sends the return traffic to the real router. Virtual inline mode is recommended when inline mode and high-speed WCCP operation are not practical.
- **High availability mode,** which allows to appliances to operate as an active/standby high availability pair. If the primary

appliance fails, the secondary appliance takes over.

Additional traffic types are listed here for completeness:

- **Pass-through traffic** refers to any traffic that the appliance does not attempt to accelerate. It is a traffic category, not a forwarding mode.
- **Direct access,** where the appliance acts as an ordinary server or client. The GUI and CLI are examples of direct access, using the HTTP, HTTPS, SSH, or SFTP protocols. Direct access traffic can also include the NTP and SNMP protocols.
- **Appliance-to-appliance communication**, which can include signaling connections (used in secure peering and by the SD-WAN plugin), VRRP heartbeats (used in high-availability mode), and encrypted GRE tunnels (used by group mode).
- **Deprecated modes.** Proxy mode and redirector mode are legacy forwarding modes that should not be used in new installations.

Customizing the Ethernet ports

Aug 09, 2017

A typical appliance has four Ethernet ports: two accelerated bridged ports, called *accelerated pair A* (apA.1 and apA.2), with a bypass (fail-to-wire) relay, and two unaccelerated motherboard ports, called Primary and Aux1. The bridged ports provide acceleration, while the motherboard ports are sometimes used for secondary purposes. Most installations use only the bridged ports.

Some SD-WAN applainces have only the motherboard ports. In this case, the two motherboard ports are bridged.

The appliance's user interface can be accessed by a VLAN or non-VLAN network. You can assign a VLAN to any of the appliance's bridged ports or motherboard ports for management purposes.

Figure 1. Ethernet Ports

The ports are named as follows:

Table 1. Ethernet Port Names

Motherboard port 1	Primary (or apA.1 if no bypass card is present)
Motherboard port 2	Auxiliary1 or Aux1 (or apA.2 if no bypass card is present)
Bridge #1	Accelerated Pair A (apA, with ports apA.1 and apA.2)
Bridge #2 Bridge #3	Accelerated Pair B (apB, with ports apB.1 and apB.2) Accelerated Pair C (apC, with ports apC.1 and apC.2)

Port Parameters

Aug 09, 2017

Each bridge and motherboard port can be:

- Enabled or disabled
- Assigned an IP address and subnet mask
- Assigned a default gateway
- Assigned to a VLAN
- Set to 1000 Mbps, 100 Mbps, or 10 Mbps
- Set to full duplex, half-duplex, or auto (on SD-WAN 4000/5000 WANOP and SD-WAN 4000 SE/5100 SE appliances, some ports can be set to 10 Gbps)

All of these parameters except the speed/duplex setting are set on the Configuration: IP Address page. The speed/duplex settings are set on the Configuration: Interface page.

Notes about parameters:

- Disabled ports do not respond to any traffic.
- The browser-based UI can be enabled or disabled independently on all ports.
- To secure the UI on ports with IP addresses, select HTTPS instead of HTTP on the Configuration: Administrator Interface: Web Access page.
- Inline mode works even if a bridge has no IP address. All other modes require that an IP address be assigned to the port.
- Traffic is not routed between interfaces. For example, a connection on bridge apA does not cross over to the Primary or Aux1 ports, but remains on bridge apA. All routing issues are left to your routers.

Accelerated Bridges (apA, apB, and apC)

Aug 09, 2017

Every appliance has at least one pair of Ethernet ports that function as an accelerated bridge, called *apA* (for *accelerated pair A*). SD-WAN 410-SE appliance has three pairs of ethernet ports (apA, apB, and apC). A bridge can act in inline mode, functioning as a transparent bridge, as if it were an Ethernet switch. Packets flow in one port and out the other. Bridges can also act in one arm mode, in which packets flow in one port and back out the same port.

An appliance that has a bypass card maintains network continuity if a bridge or appliance malfunctions.

Some units have more than one accelerated pair, and these additional accelerated pairs are named apB, apC, and so on.

If the appliance loses power or fails in some other way, an internal relay closes and the two bridged ports are electrically connected. This connection maintains network continuity but makes the bridge ports inaccessible. Therefore you might want to use one of the motherboard ports for management access.

Caution: Do not enable the Primary port if it is not connected to your network. Otherwise, you cannot access the appliance, as explained in Ethernet Bypass and Link-Down Propagation
Bypass cards are standard on some models and optional on others. Citrix recommends that you purchase appliances with bypass cards for all inline deployments.

The bypass feature is wired as if a cross-over cable connected the two ports, which is the correct behavior in properly wired installations.

Important: Bypass installations must be tested - Improper cabling might work in normal operation but not in bypass mode. The Ethernet ports are tolerant of improper cabling and often silently adjust to it. Bypass mode is hard-wired and has no such adaptability. Test inline installations with the appliance turned off to verify that the cabling is correct for bypass mode.

If the appliance is equipped with two accelerated bridges, they can be used to accelerate two different links. These links can either be fully independent or they can be redundant links connecting to the same site. Redundant links can be either load-balanced or used as a main link and a failover link.

Figure 1. Using dual bridges

When it is time for the appliance to send a packet for a given connection, the packet is sent over the same bridge from which the appliance received the most recent input packet for that connection. Thus, the appliance honors whatever link decisions are made by the router, and automatically tracks the prevailing load-balancing or main-link/failover-link algorithm in real time. For non-load-balanced links, the latter algorithm also ensures that packets always use the correct bridge.

Multiple bridges are supported in both WCCP mode (WANOP) and virtual inline mode. Usage is the same as in the single-bridge case, except that WCCP (WANOP) has the additional limitation that all traffic for a given WCCP (WANOP) service group must arrive on the same bridge.

Two units with multiple bridges can be used in a high-availability pair. Simply match up the bridges so that all links pass through both appliances.

Motherboard Ports

Aug 09, 2017

Although the Ethernet ports on a bypass card are inaccessible when the bypass relay is closed, the motherboard ports remain active. You can sometimes access a failed appliance through the motherboard ports if the bridged ports are inaccessible.

The Primary Port

If the Primary port is enabled and has an IP address assigned to it, the appliance uses that IP address to identify itself to other acceleration units. This address is used internally for a variety of purposes, and is most visible to users as the Partner Unit field on the Monitoring: Optimization: Connections page. If no motherboard port is enabled, the appliance uses the IP address of Accelerated Pair A.

The Primary port is used for:

- Administration through the web based UI
- A back channel for group mode
- A back channel for high-availability mode

VLAN Support

Aug 09, 2017

A virtual local area network (VLAN) uses part of the Ethernet header to indicate which virtual network a given Ethernet frame belongs to. SD-WAN appliances support VLAN trunking in all forwarding modes (inline, WCCP (WANOP), virtual inline, and group mode). Traffic with any combination of VLAN tags is handled and accelerated correctly.

For example, if one traffic stream passing through the accelerated bridge is addressed to 10.0.0.1, VLAN 100, and another is addressed to 10.0.0.1, VLAN 111, the appliance knows that these are two distinct destinations, even though the two VLANs have the same IP address.

You can assign a VLAN to all, some, or none of the appliance's Ethernet ports. If a VLAN is assigned to a port, the management interfaces (GUI and CLI) listen only to traffic on that VLAN. If no VLAN is assigned, the management interfaces listen only to traffic without a VLAN. This selection is made on the Configuration: Appliance Settings: Network Adapters: IP Addresses tab.

Inline Mode

Jun 12, 2018

In inline mode, traffic passes into one of the appliance's Ethernet ports and out of the other. When two sites with inline appliances communicate, every TCP connection passing between them is accelerated. All other traffic is passed through transparently, as if the appliance were not there.

For example; when you have an SD-WAN appliance deployed in Inline-Mode with only 4 VLANs on interface group. If there are 6 VLANs going through the traffic, only trafic that matches the 4 VLANs configured is recognized by the appliance.

Figure 1. Inline mode, Accelerating All Traffic on a WAN

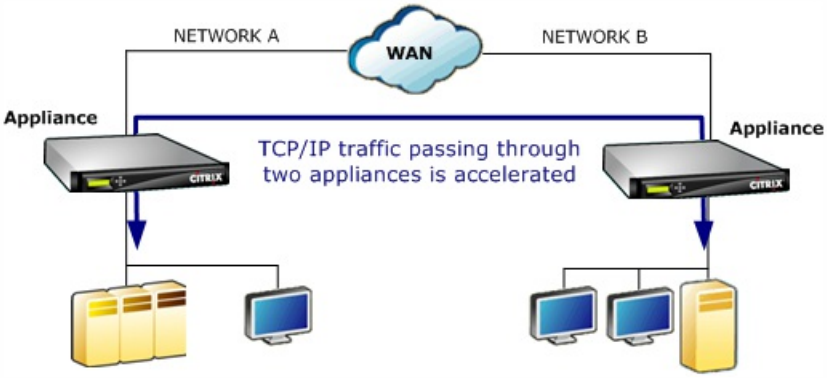

Note

Any TCP-based traffic passing through both units is accelerated. No address translation, proxying, or per-site setup is required. Inline mode is auto-detecting and auto-configuring.

Configuration is minimized with inline mode, because your WAN router need not be aware of the appliance's existence.

Depending on your configuration, inline mode's link-down propagation can affect management access to the appliance if a link goes down.

Inline mode is most effective when applied to all traffic flowing into and out of a site, but it can be used for only some of the site's traffic.

Ethernet Bypass and Link-Down Propagation

Aug 09, 2017

Most appliance models include a "fail-to-wire" (Ethernet bypass) feature for inline mode. If power fails, a relay closes and the input and output ports become electrically connected, allowing the Ethernet signal to pass through from one port to the other as if the appliance were not there. In fail-to-wire mode, the appliance looks like a cross-over cable connecting the two ports.

Any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the appliance's bridge ports are inaccessible.

If carrier is lost on one of the bridge ports, the carrier is dropped on the other bridge port to ensure that the link-down condition is propagated to the device on the other side of the appliance. Units that monitor link state (such as routers) are thus notified of conditions on the other side of the bridge.

Link-down propagation has two operating modes:

- If the Primary port is not enabled, the link-down state on one bridge port is mirrored briefly on the other bridge port, and then the port is re-enabled. This allows the appliance to be reached through the still-connected port for management, HA heartbeat, and other tasks.
- If the Primary port is enabled, the appliance assumes (without checking) that the Primary port is used for management, HA heartbeat, and other tasks. The link-down condition on one bridge port is mirrored persistently on the other port, until carrier is restored or the unit is rebooted. This is true even if the Primary port is enabled in the GUI but not connected to a network, so the Primary port should be disabled (the default) when not in use.

Accelerating an Entire Site

Aug 09, 2017

Inline mode, Accelerating All Traffic on a WAN shows a typical configuration for inline mode. For both sites, the appliances are placed between the LAN and the WAN, so all WAN traffic that can be accelerated is accelerated. This is the simplest method for implementing acceleration, and it should be used when practical.

Because all the link traffic is flowing through the appliances, the benefits of fair queuing and flow control prevent the link from being overrun.

In IP networks, the bottleneck gateway determines the queuing behavior for the entire link. By becoming the bottleneck gateway, the appliance gains control of the link and can manage it intelligently. This is done by setting the bandwidth limit slightly lower than the link speed. When this is done, link performance is ideal, with minimal latency and loss even at full link utilization.

Partial-Site Acceleration

Aug 09, 2017

To reserve the appliance's accelerated bandwidth for a particular group of systems, such as remote backup servers, you can install the appliance on a branch network that includes only those systems. This is shown in the following figure.

Figure 1. Inline Mode, Accelerating Selected Systems Only

SD-WAN traffic shaping relies on controlling the entire link, so traffic shaping is not effective with this topology, because the appliance sees only a portion of link traffic. Latency control is up to the bottleneck gateway, and interactive responsiveness can suffer.

Configuring and Troubleshooting Inline Mode

Aug 09, 2017

Inline mode requires only basic configuration, because it is applied automatically to any packets passing through the accelerated bridge.

Virtual Inline Mode

Aug 09, 2017

Note: Use virtual inline mode only when both inline mode and WCCP mode are impractical. Do not mix inline and virtual inline modes within the same appliance. However, you can mix virtual inline and WCCP modes within the same appliance. Citrix does not recommend virtual inline mode with routers that do not support health monitoring.

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router. Almost all of the configuration tasks are performed on the router. The only thing to be configured on the appliance is the forwarding method, and the default method is recommended.

Like WCCP, Virtual inline deployment requires no rewiring and no downtime, and it provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links. Unlike WCCP, it contains no built-in status monitoring or health checking, making troubleshooting difficult. WCCP is thus the recommended mode, and virtual inline is recommended only when inline and WCCP modes are both impractical.

The following figure shows a simple network in which all traffic destined for or received from the remote site is redirected to the appliance. In this example, both the local site and remote site use virtual inline mode.

Figure 1. Virtual Inline Example

Following are some configuration details for the network in this example:

- Endpoint systems have their gateways set to the local router (which is not unique to virtual inline mode).
- Each router is configured to redirect both incoming and outgoing WAN traffic to the local appliance.
- Each appliance processes the traffic received from its local router and forwards it back to the router.
- PBR rules configured on the router prevent routing loops by allowing packets to make only one trip to and from the appliance. The packets that the appliance forwards back to the router are sent to their original (local or remote) destination.
- Each appliance has its default gateway set to the address of the local router, as usual (on the **Configuration: Network Adapters** page). The options for forwarding packets back to the router are Return to Ethernet Sender and Send to Gateway.

Configuring Packet Forwarding on the Appliance

Aug 09, 2017

Virtual inline mode offers two packet-forwarding options:

Return to Ethernet Sender (default)—This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

Send to Gateway (not recommended)—In this mode, virtual inline output packets are forwarded to the default gateway for delivery, even if they are destined for hosts on the local subnet. This option is usually less desirable than the Return to Ethernet Sender option, because it adds an easily forgotten element of complexity to the routing structure.

To specify the packet-forwarding option—On the Configuration: Optimization Rules: Tuning page, next to Virtual Inline, select Return to Ethernet Sender or Send to Gateway.

Router Configuration

Aug 09, 2017

The router has three tasks when supporting virtual inline mode:

1. It must forward both incoming and outgoing WAN traffic to the SD-WAN appliance.
2. It must forward SD-WAN traffic to its destination (WAN or LAN).
3. It must monitor the health of the SD-WAN appliance so that the appliance can be bypassed if it fails.

In virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the router's Ethernet ports to the appliance and creating routing rules that are based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

The basic routing algorithm is:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward it to the appliance.
- If packet is destined for the WAN, forward to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.
- Traffic shaping is not effective unless all WAN traffic passes through the appliance.

Note: When considering routing options, keep in mind that returning data, not just outgoing data, must flow through the appliance. For example, placing the appliance on the local subnet and designating it as the default router for local systems does not work in a virtual inline deployment. Outgoing data would flow through the appliance, but incoming data would bypass it. To force data through the appliance without router reconfiguration, use inline mode.

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does no health monitoring. To enable health monitoring, define a rule to monitor the appliance's availability, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Important: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy-based routing do not support health-checking. The health-monitoring feature is relatively new. It became available in Cisco IOS release 12.3(4)T.

Following is an example of a rule for monitoring the availability of the appliance:

```
!- Use a ping (ICMP echo) to see if appliance is connected track  123 rtr 1 reachabilit y ! rtr 1  type echo protocol Iplcmpecho 192.168.1.200 schedule 1 life forever start-time now
```

This rule pings the appliance at 192.168.1.200 periodically. You can test against 123 to see if the unit is up.

Routing Examples

Aug 09, 2017

The following examples illustrate configuring Cisco routers for the local and remote sites shown in Virtual inline example. To illustrate health monitoring, the configuration for the local site includes health monitoring, but the configuration for the remote site does not.

Note: The configuration for the local site assumes that a ping monitor has already been configured.
The examples conform to the Cisco IOS CLI. They might not be applicable to routers from other vendors.

Local Site, Health-Checking Enabled

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.10.10.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 172.68.1.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended client_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 192.168.1.200 20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 192.168.1.200 10 track 123
```

Remote Site (No Health Checking)

```
! This example does not use health-checking.
! Remember, health-checking is always recommended,
! so this is a configuration of last resort.
!
!
ip cef
!
interface FastEthernet0/0
ip address 20.20.20.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 171.68.2.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.2.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.2.1
!
ip access-list extended client_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
set ip next-hop 192.168.2.200
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop 192.168.2.200
!_
```

Each of the above examples applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP of 10.10.10.0/24 and destination of 20.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended client_side
 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
 ip access-list extended wan_side
 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
```

Note that, for access lists, ordinary masks are not used. Wildcard masks are used instead. Note that when reading a

wildcard mask in binary, "1" is considered a "don't care" bit.

Virtual Inline for Multiple-WAN Environments

Aug 09, 2017

Enterprises with multiple WAN links often have asymmetric routing policies, which seem to require that an inline appliance be in two places at once. Virtual inline mode solves the asymmetric routing problem by using the router configuration to send all WAN traffic through the appliance, regardless of the WAN link used. The below figure shows a simple multiple-WAN link deployment example.

The two local-side routers redirect traffic to the local appliance. The FE 0/0 ports for both routers are in the same broadcast domain as the appliance. The local appliance must use the default virtual inline configuration (Return to Ethernet Sender).

Figure 1. Virtual Inline Mode With Two WAN Routers

Virtual Inline Mode and High-Availability

Aug 09, 2017

Virtual Inline mode can be used in a high availability (HA) configuration. The below figure shows a simple HA deployment. In virtual inline mode, a pair of appliances acts as one virtual appliance. Router configuration is the same for an HA pair as with a single appliance, except that the Virtual IP address of the HA pair, not the IP address of an individual appliance, is used in the router configuration tables. In this example, the local appliances must use default virtual inline configuration (Return to Ethernet Sender).

Figure 1. High-availability Example

Monitoring and Troubleshooting

Aug 09, 2017

In virtual inline mode, unlike WCCP mode (WANOP), the appliance provides no virtual inline-specific monitoring. To troubleshoot a virtual inline deployment, log into the appliance and use the Dashboard page to verify that traffic is flowing into and out of the appliance. Traffic forwarding failures are typically caused by errors in router configuration.

If the Monitoring: Usage or Monitoring: Connections pages show that traffic is being forwarded but no acceleration is taking place (assuming that an appliance is already installed on the other end of the WAN link), check to make sure that both incoming WAN traffic and outgoing WAN traffic are being forwarded to the appliance. If only one direction is forwarded, acceleration cannot take place.

To test health-checking, power down the appliance. The router should stop forwarding traffic after the health-checking algorithm times out.

High-Availability Mode

Aug 09, 2017

Two identical appliances on the same subnet can be combined as a *high-availability pair.* The appliances each monitor the other's status by using the standard *Virtual Router Redundancy Protocol (VRRP)* heartbeat mechanism. The pair has a common virtual IP address for management, in addition to each appliance's management IP address. If the primary appliance fails, the secondary appliance takes over. Failover takes approximately five seconds.

High availability mode is a standard feature.

How High-Availability Mode Works

Aug 09, 2017

In a high availability (HA) pair, one appliance is primary, and the other is secondary. The primary monitors its own and the secondary's status. If it detects a problem, traffic processing fails over to the secondary appliance. Existing TCP connections are terminated. To ensure successful failover, the two appliances keep their configurations synchronized. In a WCCP mode (WANOP) high availability configuration, the appliance that is processing traffic maintains communication with the upstream router.

Status monitoring—When high availability is enabled, the primary appliance uses the VRRP protocol to send a heartbeat signal to the secondary appliance once per second. In addition, the primary appliance monitors the carrier status of its Ethernet ports. The loss of carrier on a previously active port implies a loss of connectivity.

Failover If the heartbeat signal of the primary appliance should fail, or if the primary appliance loses carrier for five seconds on any previously active Ethernet port, the secondary appliance takes over, becoming the primary. When the failed appliance restarts, it becomes the secondary. The new primary announces itself on the network with an ARP broadcast. MAC spoofing is not used. Ethernet bridging is disabled on the secondary appliance, leaving the primary appliance as the only path for inline traffic. Fail-to-wire is inhibited on both appliances to prevent loops.

Caution: The Ethernet bypass function is disabled in HA mode. If both appliances in an inline HA pair lose power, connectivity is lost. If WAN connectivity is needed during power outages, at least one appliance must be attached to a backup power source.

Note: The secondary appliance in the HA pair has one of its bridge ports, port apA.1, disabled to prevent forwarding loops. If the appliance has dual bridges, apB.1 is also disabled. In a one-arm installation, use port apA.2. Otherwise, the secondary appliance becomes inaccessible when HA is enabled.

Primary/secondary assignment—If both appliances are restarted, the first one to fully initialize itself becomes the primary. That is, the appliances have no assigned roles, and the first one to become available takes over as the primary. The appliance with the highest IP address on the interface used for the VRRP heartbeat is used as a tie-breaker if both become available at the same time.

Connection termination during failover—Both accelerated and unaccelerated TCP connections are terminated as a side effect of failover. Non-TCP sessions are not affected, except for the delay caused by the brief period (several seconds) between the failure of the primary appliance and the failover to the secondary appliance. Users experience the closing of open connections, but they can open new connections.

Configuration synchronization—The two appliances synchronize their settings to ensure that the secondary is ready to take over for the primary. If the configuration of the pair is changed through the browser based interface, the primary appliance updates the secondary appliance immediately.

HA cannot be enabled unless both appliances are running the same software release.

Cabling Requirements

Aug 09, 2017

The two appliances in the high availability pair are installed onto the same subnet in either a parallel arrangement or a one-arm arrangement, both of which are shown in the following figure. In a one-arm arrangement, use the apA.2 port (and, optionally, the apB.2 port), not the apA.1 port. Some models require a separate management LAN, whether deployed in inline or one-armed mode. This is depicted only in the middle diagram.

Figure 1. Cabling for High-Availability Pairs

Do not break the above topology with additional switches. Random switch arrangements are not supported. Each of the switches must be either a single, monolithic switch, a single logical switch, or part of the same chassis.

If the spanning-tree protocol (STP) is enabled on the router or switch ports attached to the appliances, failover will work, but the failover time may increase to roughly thirty seconds. Without STP, failover time is roughly five seconds. Thus, to achieve the briefest possible failover interval, disable STP on the ports connecting to the appliances.

Figure 2. Ethernet Port Locations (Older Models)

Other Requirements

Aug 09, 2017

Both appliances in an HA pair must meet the following criteria:

- Have identical hardware, as shown by on the System Hardware entry on the Dashboard page.
- Run exactly the same software release.
- Be equipped with Ethernet bypass cards. To determine what is installed in your appliances, see the Dashboard page.

Appliances that do not support HA display a warning on the Configuration: High Availability page.

Management Access to the High-Availability Pair

Aug 09, 2017

When configuring a high-availability (HA) pair, you assign the pair a virtual IP (VIP) address, which enables you to manage the two appliances as if they were a single unit. After you enable high-availability mode, managing the secondary appliance through its IP address is mostly disabled, with most parameters grayed out. A warning message displays the reason on every page. Use the HA VIP for all management tasks. You can, however, disable the secondary appliance's HA state from its management UI.

Configuring the High-Availability Pair

Aug 09, 2017

You can configure two newly installed appliances as a high-availability pair, or you can create an HA pair by adding a second appliance to an existing installation.

Prerequisites: Physical installation and basic configuration procedures

To configure high availability

1. Make sure that no more than one appliance is connected to the traffic networks (on the accelerated bridges). If both are connected, disconnect one bridge cable from the active bridges on the second appliance. This will prevent forwarding loops.
2. On the Features page of the first appliance, disable Traffic Processing. This disables acceleration until the HA pair is configured.
3. Repeat for the second appliance.
4. On the first appliance, go to the Configuration: Advanced Deployments: High Availability tab, show below.
5. Select the Enabled Check box.
6. Click the Configure HA Virtual IP Address link and assign a virtual IP address to the apA interface. This address will be used later to control both appliances as a unit.
7. Return to the High Availability page and, in the VRRP VRID field, assign a VRRP ID to the pair. Although the value defaults to zero, the valid range of VRRP ID numbers is 1 through 255. Within this range, you can specify any value that does not belong to another VRRP device on your network.
8. In the Partner SSL Common Name field, type the other appliance's SSL Common Name, which is displayed on that appliance's Configuration: Advanced Deployments: High Availability tab, in the Partner SSL Common Name field. The SSL credentials used here are factory-installed.
9. Click Update.
10. Repeat steps 3-8 on the second appliance. If you are managing the appliance via an accelerated bridge (such as apA), you may have to reconnect the Ethernet cable that you removed in step 1 to connect to the second appliance. If so, plug this cable in and disconnect the corresponding cable on the first appliance.
11. With your browser, navigate to the virtual IP address of the HA pair. Enable Traffic Processing on the Features page. Any further configuration will be performed from this virtual address.
12. Plug in the cable that was left disconnected.
13. On each appliance, the Configuration: Advanced Deployments: High Availability page should now show that high availability is active and that one appliance is the primary and the other is the secondary. If this is not the case, a warning banner appears at the top of the screen, indicating the nature of the problem.

Figure 1. High-availability configuration page

Updating Software on a High-Availability Pair

Aug 09, 2017

Updating the SD-WAN software on an HA pair causes a failover at one point during the update.

Note: Clicking the Update button terminates all open TCP connections.

To update the software on an HA pair

1. Log on to both appliances.
2. On the secondary appliance, update the software and reboot. After the reboot, the appliance is still the secondary. Verify that the installation succeeded. The primary appliance should show that the secondary appliance exists but that automatic parameter synchronization is not working, due to a version mismatch.
3. On the primary appliance, update the software, and then reboot. The reboot causes a failover, and the secondary appliance becomes the primary. When the reboot is completed, HA should become fully established, because both appliances are running the same software.

Saving/Restoring Parameters of an HA Pair

Aug 09, 2017

The System Maintenance: Backup/Restore function can be used to save and restore parameters of an HA pair as follows:

To back up the parameters

Use the backup feature as usual. That is, log on to the GUI through the HA VIP address (as is normal when managing the HA pair) and, on the System Management: Backup/Restore page, click Download Settings.

To restore the parameters

1. Disable HA on both appliances by clearing the Enabled check box on the Configuration: Advanced Deployments: High Availability (HA) tab.
2. Unplug a network cable from the bridge of one appliance. (Call it “Appliance A.”)
3. Unplug the power cord from Appliance A.
4. Restore the parameters on the other appliance (Appliance B), by uploading a previously saved set of parameters on the System Maintenance: Backup/Restore page and clicking Restore Settings. (Completing this operation requires a restart, which reenables HA).
5. Wait for Appliance B to restart. It becomes the primary.
6. Restart Appliance A.
7. Log on to Appliance A’s GUI and reenable HA on the Configuration: Advanced Deployments: High Availability (HA) tab. The appliance get its parameters from the primary.
8. Plug in the network cable removed in step 2.

Both appliances are now restored and synchronized.

Troubleshooting High Availability Pairs

Aug 09, 2017

If the appliances report any failure to enter high-availability mode, the error message will also note the cause. Some issues that can interfere with high-availability mode are:

- The other appliance is not running.
- The HA parameters on the two appliances are not identical.
- The two appliances are not running the same software release.
- The two appliances do not have the same model number.
- Incorrect or incomplete cabling between the appliances does not allow the HA heartbeat to pass between them.
- The HA SSL Certificates on one or both appliances are damaged or missing.

Factory Reset on 410-SE

Dec 20, 2017

The factory reset option is applicable only to the SD-WAN 410-Standard Edition applaince.

To perform factory reset on 410-SE:

1. Once the power LED is flashing, power OFF the appliance with using the power button to the right of the NMI reset button as shown in the image above. Hold the power button for few seconds (10-20 seconds) until the green LED and other LED lights are turned off.

> ## Note
>
> Press and hold the power button for a few seconds to turn off the appliance.

2. Press the NMI reset button once and then power ON the appliance using the power button.

The green LED starts blinking on and off for the next 20-25 minutes until the eUSB recovery process is finished.

3. Wait for few minutes, apporximately 5 minutes initially as no activity will happen on the CLI. Non-activity in CLI does not mean nothing is happening. The system is initializing the process.

> ## Tip
>
> - Pressing the reset button even number of times cancels the reset action and results in normal appliance reboot.
>
> - Pressing the reset button odd number of times performs a factory reset.
>
> - Power LED flash indicates that the appliance is being reset.

4. After 5 mins the appliance restarts and the CLI is displayed. There will be couple of reboots (apporximately, 4-5) for extracting the software image from the eUSB (sdb) and copying, programming, and re-flashing to the SATADOM (sda).

```
>searching /misc/rescue.img on /dev/sda /dev/sda1 Waiting for 2 seconds and retry ...

Searching /misc/rescue.img on /dev/sda /dev/sda1 /dev/sdb /dev/sdb1

Found /misc/rescue.img on /dev/sdb1.  Extracting ... Found /dev/sdb1.


>task=mfg TARGET_DRIVE= eth= IP_MODE= ip= mask= gw=

Running Manufacturing script: /usr/local/bin/mfg.sh ...


>Copying files. Please be patient ...Updated extlinux config and installed extlinux on /tmp/tmp.hTg61R2uPh/primary/boot

Updated extlinux config and installed extlinux on /tmp/tmp.hTg61R2uPh/secondary/boot

Success        Rebooting for baremetal OS installation ...

Rebooting in 5 seconds ...: 1     seconds ...

Will now restart

[  145.238780] reboot: Restarting system
```

The appliance restarts 4 to 5 times as it extracts, copies, and initializes the boot process.

5. At the login prompt, you can start configuring the appliance using CLI or the web management interface.

NetScaler SD-WAN 4000, 4100, and 5100 Standard Edition Appliances

Aug 29, 2017

Citrix NetScaler SD-WAN Standard Edition 4000/4100/5100 appliances are high-performance appliances for busy datacenters.

SD-WAN 4000/4100/5100 Standard Edition appliances are designed to with Virtual WAN links with speeds in excess of 1 Gbps, especially for busy datacenters that communicate with a large number of branch and regional sites.

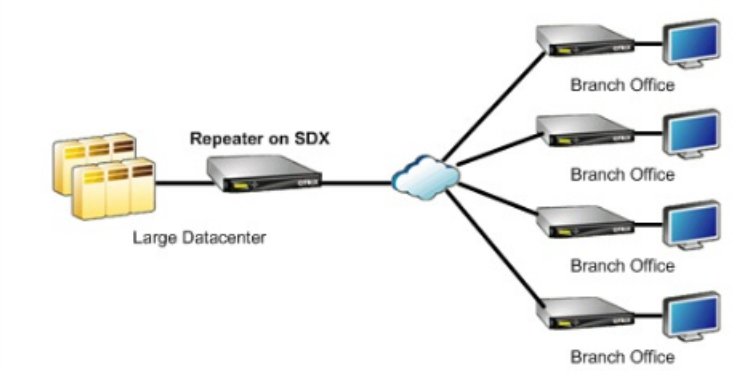

SD-WAN 4000/4100/5100 SE is recommended at the hub of a hub-and-spoke deployment, where smaller appliances are used at the spokes, whenever the link speed or the number of XenApp/XenDesktop users is higher than that can be supported by a smaller appliance.

Citrix NetScaler SD-WAN 4000 SE

Aug 09, 2017

Citrix NetScaler SD-WAN 4000 is a 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 48 gigabytes (GB) of memory. The Citrix NetScaler SD-WAN 4000 SE has a bandwidth of 300Mbps, 500Mbps, 1Gbps, and 2Gbps respectively.

The following figures shows the front panel of the Citrix NetScaler SD-WAN 4000 SE appliance.

Figure 1. Citrix NetScaler SD-WAN 4000 SE, front panel

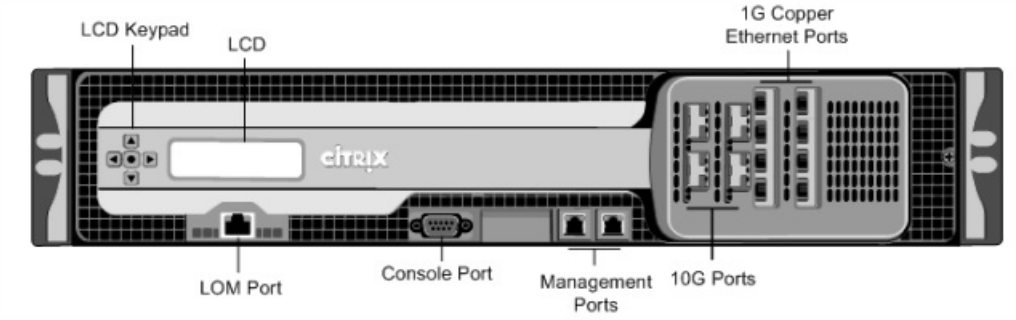

The Citrix NetScaler SD-WAN 4000 SE appliances have the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
  Note: The LEDs on the LOM port are not operational by design.
- RS232 serial console port.
- Two 10/100/1000Base-T copper Ethernet management ports (RJ45). These ports are used to connect directly to the appliance for system administration functions.
- Network Ports
  - NetScaler SD-WAN 4000 SE (without FTW cards). Eight 1G SFP ports and four 10G SFP+ ports.
  - NetScaler SD-WAN 4000 SE (with FTW cards). Eight 1G copper Ethernet ports and four 10G ports.

The following figure shows the back panel of the Citrix NetScaler SD-WAN 4000 SE appliance.

Figure 2. Citrix NetScaler SD-WAN 4000 SE, back panel

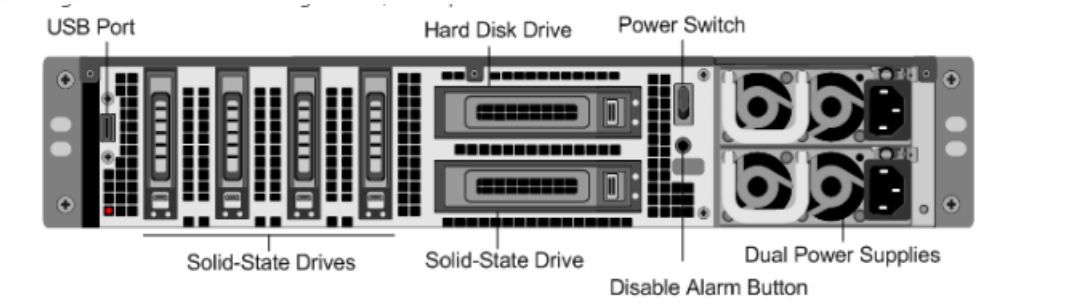

The following components are visible on the back panel of the Citrix NetScsler SD-WAN 4000 SE appliance:

- Four 600 GB removable solid-state drives. The 256 GB solid-state drive below the hard disk drive stores the appliance's software.. Newer editions of 4000-SE have 800 GB removable SSD and 240 GB SSD.
- USB port (reserved for a future release).

- A 1 TB removable hard disk drive.
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.
- Disable alarm button. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Dual power supplies (either AC or DC), each rated at 850 watts, 100-240 volts.

Citrix NetScaler SD-WAN 4100 SE

Aug 09, 2017

Citrix NetScaler SD-WAN 4100 is a 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 96 gigabytes (GB) of memory. The Citrix NetScaler SD-WAN 4100 SE has a virtual WAN bandwidth of 1Gbps and 2Gbps.

The following figures shows the front panel of the Citrix NetScaler SD-WAN 4100 SE appliance.

Figure 1. Citrix NetScaler SD-WAN 4100 SE, front panel

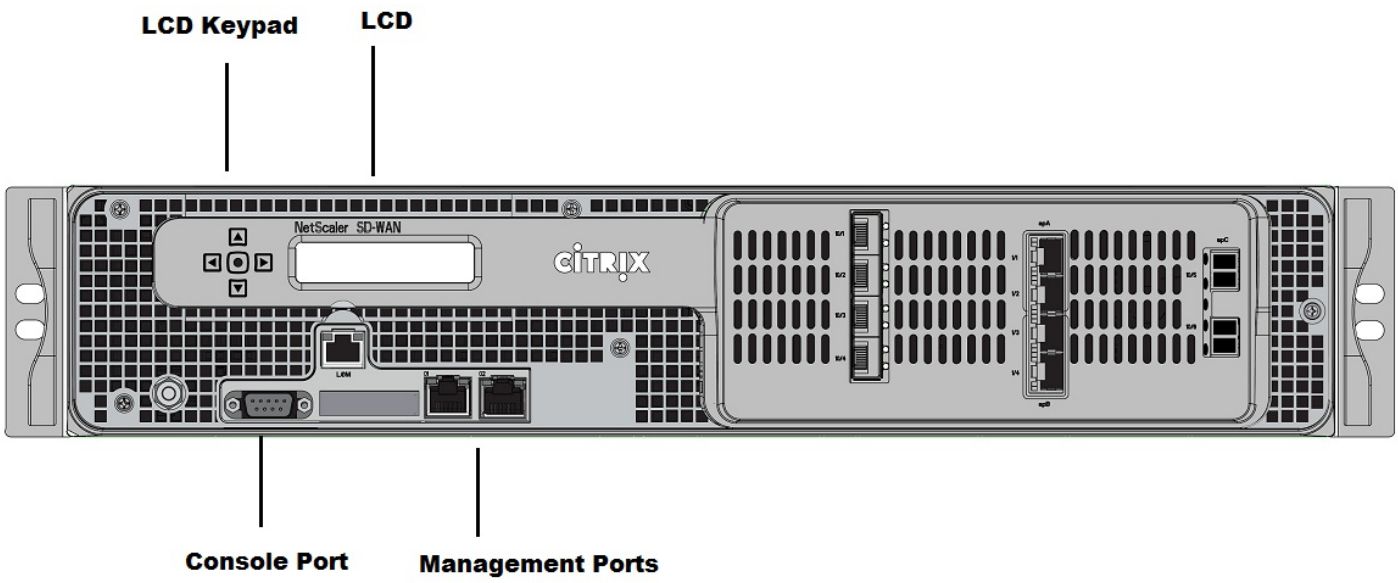

The Citrix NetScaler SD-WAN 4100 SE appliances have the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software. The LEDs on the LOM port are not operational by design.
- RS232 serial console port.
- Two 10/1000Base-T copper Ethernet management ports (RJ45). These ports are used to connect directly to the appliance for system administration functions.
- 2 port 10G FTW
- 4 port 10G/1G SFP+
- 4 port 10/100/1000 FTW RJ 45

The following figure shows the back panel of the Citrix NetScaler SD-WAN 4100 SE appliance.

Figure 2. Citrix NetScaler SD-WAN 4100 SE, back panel

The following components are visible on the back panel of the Citrix NetScsler SD-WAN 4100 SE appliance:

- 2 X 1 TB HDD in RAID 1.
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to turn off the power.
- Disable alarm button. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Dual power supplies (either AC or DC), each with max power of 850 watts, 100-240 volts.

Citrix NetScaler SD-WAN 5100 SE

Apr 16, 2018

Citrix NetScaler SD-WAN 5100 SE is a 2U appliance. Each model has two 10-core processors for a total of 20 physical cores (40 cores with hyper-threading), and 128 gigabytes (GB) of memory. For latest performance and bandwidth capacity details, please see also the latest data sheet that gets updated more regularly at: https://www.citrix.com/products/netscaler-sd-wan/netscaler-data-sheet.html.

The Citrix NetScaler SD-WAN 5100 SE appliance front panel has the following ports:

- 10/100Base-T copper Ethernet Port (RJ45), also called LOM port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software. The LEDs on the lights out management port are not operational by design.
- RS232 serial console port.
- Two 10/100/1000Base-T copper Ethernet management ports (RJ45). These ports are used to connect directly to the appliance for system administration functions.
- Eight 10G ports.

The below figure illustrates NetScaler SD-WAN 5100-SE Appliance model.

The following components are visible on the back panel of the Citrix NetScaler SD-WAN 5100 SE appliance:

- 2 X 1 TB removable hard disk drive.
- USB port (reserved for a future release).
- Power switch, which turns off power to the appliance, just as if you were to unplug the power supply. Press the switch for five seconds to clear the power.
- Disable alarm button. This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when you have plugged the appliance into only one power outlet or when one power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Dual power supplies (either AC or DC), each rated at 850 watts, 100–240 volts.

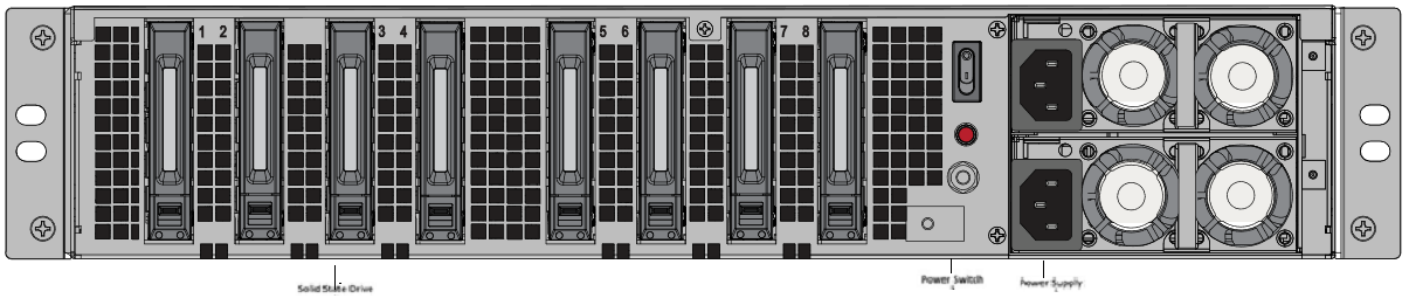

Insert Solid State Drive (SSD)

1. Insert the required SSD in the standard edition appliance. For instructions about how to insert SSD, see Solid State Drive (Field Replaceable Unit).

- 5100 SE appliance requires 800 GB more SSD. Insert the SSD into the third bay.

2. Restart the appliance through the SD-WAN web management interface.

3. Ensure that the software release version installed on the appliance is SD-WAN release version 9.3.3.

Appliances shipped with software release version 9.3.3

For appliances shipped with 9.3.3 or new manufacture image. Follow the steps provided here; upgrade new appliance.

1. Install the Enterprise Edition platform license. For license information, see the Citrix SD-WAN product downloads site.
2. Upgrade the network using single step upgrade to software release version 10.0 or later.

Configure Management IP Address Using Serial Console

1. Access serial console of the appliance.
2. Log in using the **root/nsroot** credentials.
3. Type the **ssh admin@169.254.0.60 -l** administrator command.
4. Type password: **password**.
5. Type the **management_ip** command.
6. Type the **set interface 192.168.100.1 255.255.255.0 192.168.100.254** command.
7. Type the **apply** command.

Summary of Hardware Specifications

Aug 09, 2017

The following table summarizes the specifications of Citrix NetScaler SD-WAN 4000, 4100, and 5100 SE hardware platforms.

Specifications	SD-WAN 4000 SE	SD-WAN 4100 SE	SD-WAN 5100 SE
Regulatory Model Number	4x10GE SFP+ 8xSFP	2U1P1B	2U1P1D
Processors	Two 6-core	Two 6-core	Two 10-core
Memory	96 GB	96 GB	128 GB
Number of power supplies	1 (optional second power supply for redundancy)	1 (optional second power supply for redundancy)	1 (optional second power supply for redundancy)
AC power supply, input voltage, frequency and current	100-240VAC, 47-63 hz	100-240VAC, 47-63 hz 7.0-3.5A	100-240VAC, 47-63 hz 9.0-4.5A
Maximum AC power consumption	650 W	850 W	850 W
Package weight	69 lbs	69 lbs	69 lbs
Shipping dimensions	36.5' L X 24.5' W X 11' H	36.5' L X 24.5' W X 11' H	36.5' L X 24.5' W X 11' H
System weight	60 lbs	60 lbs	60 lbs
Rack unit	2RU	2RU	2RU
Rack options - Width	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting brackets	EIA 310-D, IEC 60297, DIN 41494 SC48D rack width with mounting bracket
Depth	28" (72 cm)	28" (72 cm)	28" (72 cm)
Operating temperature	32 – 104 F (0 – 40 C)	32 – 104 F (0 – 40 C)	32 – 104 F (0 – 40 C)
Humidity (non-condensing)	20% - 80%	20% - 80%	20% - 80%

Specifications	SD-WAN 4000 SE	SD-WAN 4100 SE	SD-WAN 5100 SE
Safety certifications	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)
EMC and susceptibility	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)	USA (FCC), Europe (CE), Japan (VCCI), Australia (RCM), China (CCC), Korea (KCC), India (BIS), Mexico (NOM), Saudi Arabia (CITC), South Africa (ICASA), Russia (EAC), Taiwan (BSMI), Brazil (Anatel), Israel (MoC)
Environmental compliance	ROHS, WEEE	ROHS, WEEE	ROHS, WEEE

Preparing for Installation

Aug 09, 2017

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Once you are satisfied that your appliance has been delivered to your expectations, verify that the location where the appliance will be installed meets temperature and power requirements and that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel should install, maintain, or replace the appliance, and efforts should be taken to ensure that all cautions and warnings are followed.

Unpacking the Appliance

Aug 09, 2017

Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

Use the following list to verify that you received everything that should have been included in the box.

- The appliance you ordered
- One RJ-45 to DB-9 adapter
- One 6 ft RJ-45/DB-9 cable
- One or Two power cables depending on the platform edition
- One standard 4-post rail kit

In addition to the items included in the box with your new appliance, you will need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you will connect to your network.
- One available Ethernet port on your network switch or hub for each Ethernet port you want to connect to your network.
- A computer to serve as a management workstation.

> ## Note
>
> If the kit that you received does not fit your rack, contact your Citrix sales representative to order the appropriate kit.

Preparing the Site and Rack

Aug 09, 2017

There are specific site and rack requirements for the SD-WAN 4000/4100/5100 SE appliances. You must make sure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and help ensure a smooth installation.

The appliance should be installed in a server room or server cabinet with the following features:

Environment control

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment.

Power density

Wiring capable of handling at least 4,000 watts per rack unit in addition to power needs for the CRAC.

The rack on which you install your appliance should meet the following criteria:

Rack characteristics

Racks should be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it should be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks should be installed in rows spaced at least 1 meter/3 feet apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

Power connections

At minimum, two standard power outlets per unit.

Network connections

At minimum, Ethernet connection per rack unit.

Space requirements

Two empty rack units for SD-WAN 4000/4100/5100 SE appliances.

> **Note**
>
> You can order the following rail kits separately.
>
> - Compact 4-post rail kit, which fits racks of 23 to 33 inches.
> - 2-post rail kit, which fits 2-post racks.

Cautions and Warnings

Aug 09, 2017

> ## Warning
>
> During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.

Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Remove all jewelry and other metal objects that might come into contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.
- Use a regulating, uninterruptible power supply (UPS) to protect the appliance from power surges and voltage spikes, and to keep the appliance operating in case of power failure.
- Never stack the appliance on top of any other server or electronic equipment.
- All appliances are designed to be installed on power systems that use TN earthing. Do not install your device on a power system that uses either TT or IT earthing.
- Make sure that the appliance has a direct physical connection to the earth during normal use. When installing or repairing an appliance, always make sure that the ground circuit is connected first and disconnected last.
- Make sure that a fuse or circuit breaker no larger than 120 VAC, 15 A U.S. (240 VAC, 16 A international) is used on all current-carrying conductors on the power system to which your appliances are connected.
- Do not work alone when working with high voltage components.
- Always disconnect the appliance from power before removing or installing any component. When disconnecting power, first shut down the appliance, and then unplug the power cords of all the power supply units connected to the appliance. As long as the power cord is plugged in, line voltages can be present in the power supply, even when the power switch is OFF.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload. Always unplug any appliance before performing repairs or upgrades.
- Do not overload the wiring in your server cabinet or on your server room rack.
- During thunderstorms, or anticipated thunderstorms, avoid performing any hardware repairs or upgrades until the danger of lightning has passed.
- When you dispose of an old appliance or any components, follow any local and national laws on disposal of electronic waste.
- To prevent possible explosions, replace expired batteries with the same model or a manufacturer-recommended substitute and follow the manufacturer's instructions for battery replacement.
- Never remove a power supply cover or any sealed part.

- Determine the placement of each component in the rack before you install the rails.
- Install the heaviest appliance first, at the bottom of the rack, and then work upward. Distribute the load on the rack evenly. An unbalanced rack is hazardous.
- Allow the power supply units and hard drives to cool before touching them.
- Install the equipment near an electrical outlet for easy access.
- Mount equipment in a rack with sufficient airflow for safe operation.
- For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 09, 2017

Most appliances can be installed in standard server racks that conform to EIA-310-D specification. The appliances ship with a set of rails, which you must install before you mount the appliance. The only tools that you need for installing an appliance are a Phillips screwdriver and a flathead screwdriver.

Citrix SD-WAN 4000/4100/5100 appliances requires two rack units.

Each appliance ships with a mounting rail kit that contains two rail assemblies, one for the left side and the other for the right side of the appliance, and screws to attach the rails. An assembly consists of an inner rail and a rack rail. The supplied rail kit is 28 inches long (38 inches extended). Contact your Citrix sales representative to order a 23-inch (33 inches extended) rail kit.

To mount the appliance, you must first install the rails and then install the appliance in the rack.

Perform the following tasks to mount the appliance:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

> ## Note
>
> The same rail kit is used for both square-hole and round-hole racks. See figure 4 for specific instructions for threaded, round-hole racks.

> ## Warning
>
> If you are installing the appliance as the only unit in the rack, mount it at the bottom. If the rack contains other units, make sure that the heaviest unit is at the bottom. If the rack has stabilizing devices available, install them before mounting the appliance.

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the latch until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

1. Position the right inner rail behind the handle on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws: 5 per side, as shown in the following figure.

Figure 1. Attaching inner rails

1. If you have a round-hole, threaded rack, skip to step 3.
2. Install square nut retainers into the front post and back post of the rack as shown in the following figures. Before inserting a screw, be sure to align the square nut with the correct hole for your appliance. The three holes are not evenly spaced.
   Figure 2. Installing Retainers into the Front Rack Posts

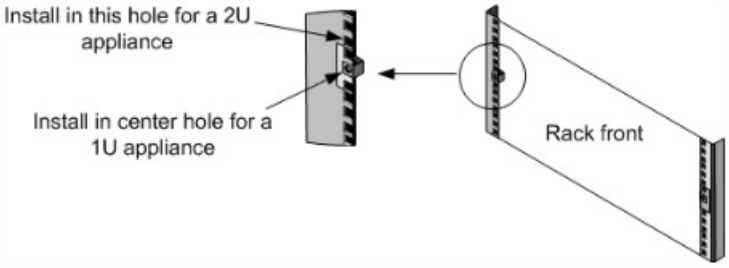

Figure 3. Installing Retainers into the Rear Rack Posts

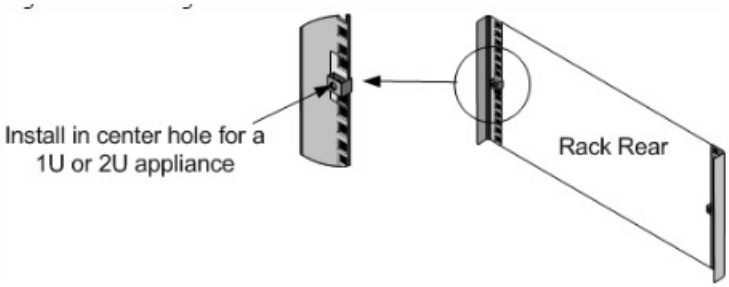

3. Install the adjustable rail assembly into the rack as shown in the following figures. Use a screw to lock the rear rail flange into the rack. With the screw securing the rail in place, you can optionally remove the latching spring.

Figure 4. Installing the Rail Assembly to the Rack

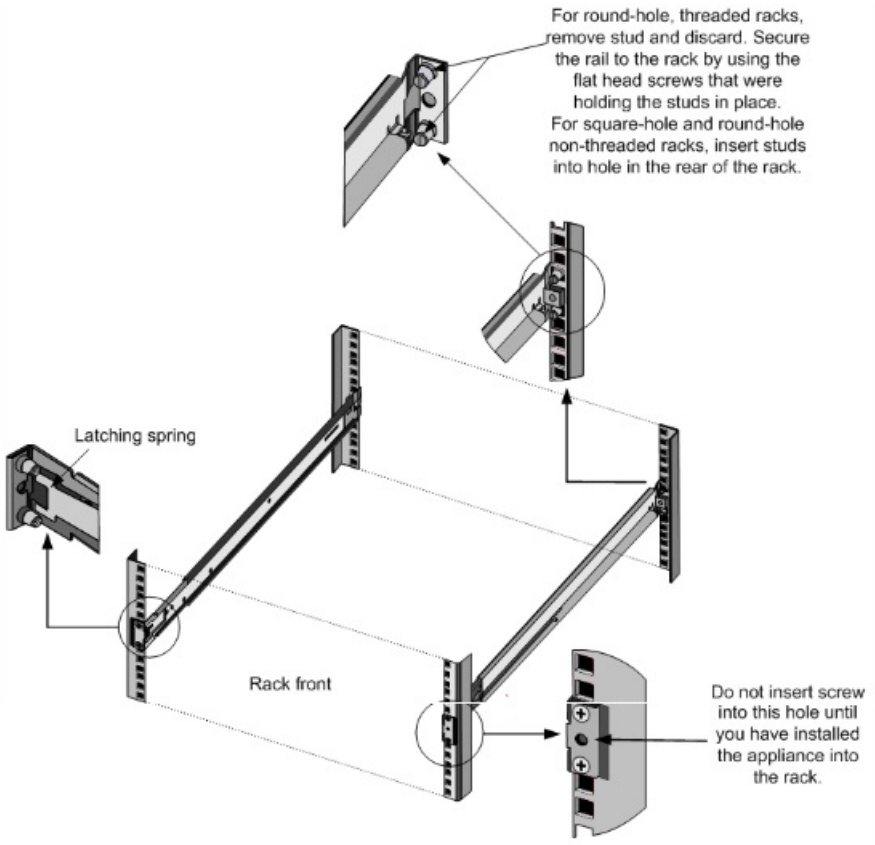

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.

Figure 5. Rack Mounting the Appliance

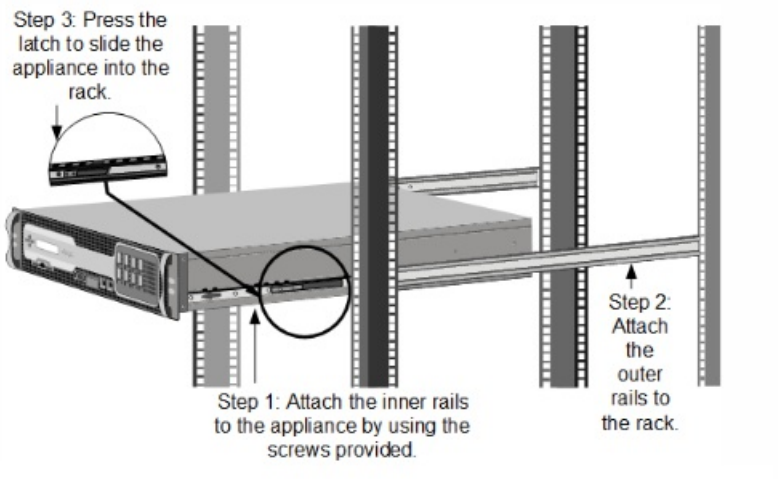

Installing and Removing 1G SFP Transceivers

Aug 09, 2017

A Small Form-Factor Pluggable (SFP) is a compact transceiver that can operate at speeds of up to 1 gigabit per second and is available in both copper and fiber types. Inserting a 1G SFP copper transceiver converts the 1G SFP port to a 1000BASE-T port. Inserting a 1G SFP fiber transceiver converts the 1G SFP port to a 1000BASE-X port. Auto-negotiation is enabled by default on the 1G SFP port into which you insert your 1G SFP transceiver. As soon as a link between the port and the network is established, the speed and mode are matched on both ends of the cable.

Note

Some SD-WAN Standard Edition appliances do not require SFP transceivers.

Important

SFP transceivers must be ordered separately. With exception to fiber FTW ports that have integrated transceivers, which are not removable.

1. Remove the 1G SFP transceiver carefully from its box.
   Align the 1G SFP transceiver to the front of the 1G SFP transceiver port on the front panel of the appliance, as shown in the following figure.

Warning

Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.

2. Hold the 1G SFP transceiver between your thumb and index finger and insert it into the 1G SFP transceiver port, pressing it in until you hear the transceiver snap into place.

3. Lock the transceiver.

4. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.

5. If you are using a fiber 1G SFP transceiver, do not remove the dust caps attached to the transceiver and the cable until you are ready to insert the cable.

1. Disconnect the cable from the 1G SFP transceiver. If you are using a fiber optic cable, replace the dust cap on the cable before putting it away.
   Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 1G SFP transceiver.
3. Hold the 1G SFP transceiver between your thumb and index finger and slowly pull it out of the port.
4. If you are removing a fiber 1G SFP transceiver, replace the dust cap before putting it away.
5. Put the 1G SFP transceiver into its original box or another appropriate container.

Warning

Insert 1G SFP transceivers into the 1G SFP ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the 1G SFP transceiver or the appliance.

Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

Installing and Removing 10G SFP+ Transceivers

Aug 09, 2017

A 10-Gigabit Small Form-Factor Pluggable (SFP+) is a compact optical transceiver that can operate at speeds of up to 10 gigabits per second. Autonegotiation is enabled by default on the 10G SFP+ ports into which you insert your 10G SFP+ transceiver. As soon as a link between the port and the network is established, the mode is matched on both ends of the cable and for 10G SFP+ transceivers, the speed is also autonegotiated.

Important

SFP transceivers must be ordered separately. With exception to fiber FTW ports that have integrated transceivers, which are not removable.

1. Remove the 10G SFP+ transceiver carefully from its box. Align the 10G SFP+ transceiver to the front of the 10G SFP+ transceiver port on the front panel of the appliance.
2. Hold the 10G SFP+ transceiver between your thumb and index finger and insert it into the 10G SFP+ transceiver port, pressing it in until you hear the transceiver snap into place.
3. Move the locking hinge to the DOWN position.
4. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
5. Do not remove the dust caps attached to the transceiver and cable until you are ready to insert the cable.

Warning

Do not look directly into fiber optic transceivers and cables. They emit laser beams that can damage your eyes.

1. Disconnect the cable from the 10G SFP+ transceiver. Replace the dust cap on the cable before putting it away. Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 10G SFP+ transceiver by moving the locking hinge to the UP position.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and slowly pull it out of the port.
4. Replace the dust cap on the transceiver before putting it away.
5. Put the 10G SFP+ transceiver into its original box or another appropriate container.

Warning

Insert the 10G SFP+ transceivers into the 10G SFP+ ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the transceiver or the appliance.

Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

Connecting the Cables

Aug 09, 2017

When the appliance is securely mounted on the rack, you are ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

> ### Warning
>
> Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port on the front panel of the appliance, as shown in the following figure.
   Figure 1. Inserting an Ethernet cable

2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port that is located on the front panel of the

appliance, as shown in the following figure.
Figure 2. Inserting a console cable

2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

> **Note**
>
> To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.

The SD-WAN 4100/5100 SE appliances have two power supplies, with one serving as a backup. A separate ground cable is not required, because the three-prong plug provides grounding. Power up the appliance by installing one or both power cords.

To connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply, as shown in the following figure.
   Figure 3. Inserting a power cable

3. Connect the other end of the power cable to a standard 110V/220V power outlet.

4. Repeat steps 1 and 2 to connect the second power supply.

> **Note**
>
> The appliance emits a high-pitched alert if one power supply fails or if you connect only one power cable to the appliance. To silence the alarm, you can press the small red button located on the back panel of the appliance.

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the back panel of the appliance.

> ## Warning
>
> Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Planning the Deployment

Aug 09, 2017

SD-WAN 4000/4100/5100 deployments require adequate planning, especially for units deployed in large datacenters:

- An appropriate appliance or group of appliances must be selected to support both the current and anticipated load.
- A deployment mode must be selected to match the requirements of your site.
- Other aspects must also be considered.

Sizing Guidelines

Aug 09, 2017

For successful deployment of one or more SD-WAN 4000/4100/5100 appliances in your datacenter, keep the following principles in mind:

- You must provide enough SD-WAN 4000/4100/5100 peak-load capacity, in terms of WAN bandwidth and the number of users. See the current specifications sheet for the capacities of different SD-WAN 4000/4100/5100 models: https://www.citrix.com/content/dam/citrix/en_us/documents/data-sheet/netscaler-sd-wan-datasheet.pdf
- Provide enough capacity for expected expansion over the life of the deployment. SD-WAN 4000/4100/5000 appliances using the same hardware platform can have their capacity upgraded with a new license as part of the Citrix pay-as-you-grow program.
- For more capacity than can be provided by a single appliance, multiple SD-WAN 4000/4100/5100 appliances can be cascaded behind a stand-alone NetScaler appliance.
- Different models have differing numbers of traffic ports. If you require multiple bridges, make sure your model has at least as many as you need.

Selecting a Deployment Mode

Aug 09, 2017

The SD-WAN 4000/4100/5100 appliance can be deployed inline or in a one-arm mode. Inline deployments do not require router reconfiguration; one-arm modes do. SD-WAN 4000/4100/5100 offers internal port bypassing (fail-to-wire) to allow traffic to continue flowing in inline mode if the appliance fails.

A standalone SD-WAN 4000/4100/5100 appliance can be deployed in either of these recommended modes:

- Inline, bridged (L2 inline). This closely resembles a standard SD-WAN inline deployment. Packets enter one bridge port and exit the other bridge port.
- Inline, routed. The NetScaler instance uses routing rules instead of bridging rules to determine how to forward packets.
- Virtual inline. This resembles WCCP, but lacks built-in health-checking.

In L2 inline mode, SD-WAN 4000/4100/5100 is placed between your LAN and your WAN router (or other aggregation point at the LAN-WAN boundary). In a one-arm mode, SD-WAN 4000/4100/5100 is generally connected directly to a dedicated port on your WAN router.

In cases where the WAN router ports are not as fast as the LAN (for example, when the WAN router has gigabit Ethernet, but the LAN has 10 gigabit Ethernet), inline mode provides better performance, because its LAN-side traffic is not limited to the speed of the router interface. (Compression allows the LAN-side traffic to be much faster than WAN-bound traffic under favorable conditions.)

Considerations:

- The inline modes require no reconfiguration of your routers, but involves a service disruption when bringing the appliance into service.
- One-arm modes require router reconfiguration but do not require a service disruption.
- Inline mode has higher performance than the other modes.
- One-arm modes are limited to half the speed of the router or switch port they are attached to.

Recommendation:

- Inline mode is convenient for smaller WAN networks and simpler datacenters.

Gathering Information Needed for Configuration

Aug 09, 2017

Accurate information about both the local and the remote sites is essential to troubleshooting. Before installing the SD-WAN 4000/4100/5100 appliance, make sure that you have done the following:

1. Obtained or drawn an accurate network diagram of your local site (the one in which you are installing SD-WAN 4000/4100/5100). The local network topology and the capabilities of your WAN routers determine which deployment modes are appropriate for the site.
2. Chosen the deployment mode of the local SD-WAN 4000/4100/5100 appliance (inline, with or without HA and cascading).
3. Compiled a list of critical applications that must be tested to validate the deployment.
4. Obtained or drawn an accurate network diagram of your WAN, including both the local and the remote WAN links, their bandwidths in both directions, their subnets. In deployments with many remote sites, an aggregate of the different categories (accelerated and non-accelerated) is probably sufficient, and only the largest remote sites need to be considered individually.
5. Determined whether there are multiple datacenters with datacenter-to-datacenter traffic, and whether any remote datacenters have a SD-WAN 4000/4100/5100 appliance.
6. Decided whether you plan to increase WAN capacity, the number of sites, or the number of users in the next 24 months. If so, the corresponding SD-WAN 4000/4100/5100 capacity should be installed now.
7. If possible, formed an idea of the traffic breakdown over the WAN, including TCP traffic to and from SD-WAN sites, other TCP traffic, ICA users, and real-time traffic such as VoIP. SD-WAN 4000/4100/5100 needs to be provisioned for the peak loads in terms of accelerated TCP connections, ICA users, and total WAN link capacity.
8. Determined the number of WAN links in the local site. Are they independent, or are they load balanced? If so, are they active-active or active-standby?
9. Determined the current, unaccelerated RTT of the remote sites during peak periods.
10. Identified any QoS devices or proxies in the path between the local and remote sites. QoS devices should be on the WAN side of SD-WAN 4000/4100/5100. Proxies should be on the LAN side.

Initial Configuration

Aug 09, 2017

After checking the connections, you are ready to deploy the SD-WAN 4000, 4100, and 5100 appliances on the network.

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

Initial configuration consists of the following tasks:

- Initial configuration consists of the following tasks:
- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Assign management IP address through the serial console.

By default, the fail-to-wire network adapters are in bypass state.

Prerequisites

Aug 09, 2017

To deploy a Citrix NetScaler SD-WAN 4000, 4100, or 5100 appliance, you must complete the following prerequisite setup before configuring the appliance.

This document covers release of the SD-WAN software. See the release notes for the recommended versions of the NetScaler software corresponding to the desired release of the SD-WAN software.

Before you start provisioning the appliance, Citrix recommends that you have the license file with you, as it is required early in the configuration process. To download a license file, complete the procedure described in the Licensing section.

After you receive the hardware appliance from Citrix, you need to install it in the network. Complete the following procedures to install the hardware.

To install the SD-WAN 4000/4100/5100 appliance hardware, follow the installation procedure at Installing the Hardware.

Configuring the Appliance

Aug 09, 2017

Before you start configuring the appliance, you must change the IP address of the management service to the one in your management network, so that you can access the appliance over the network. You can change the management IP address by connecting a computer to the appliance through either the Ethernet port or the serial console.

Assigning a Management IP Address through the Ethernet Port

Aug 09, 2017

Use the following procedure for initial configuration of every SD-WAN 4000, 4100, or 5100 appliance. The procedure accomplishes the following tasks:

- Configure the appliance for use on your site.
- Install the Citrix license.

For detailed steps, see Assigning a Management IP Address.

If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 11 of the following procedure.

Note: You must have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.1, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields blank.
2. Using an Ethernet cable, connect this computer to the port labeled PRI on the SD-WAN appliance.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address, which is http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance, user: *admin* and password: *password.*
5. A redirection to new management IP message appears.
6. Click **OK**.
7. Unplug your computer from the Ethernet port and connect the port to your management network.
8. Reset the IP address of your computer to its previous setting.
9. From a computer on the management network, log on to the appliance by entering the new management service IP address, such as https://<Managemnt_IP_Address>, in a web browser.
10. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
11. Log on to the appliance.

Assigning a Management IP Address through the Serial Port

Aug 09, 2017

If you do not want to change the settings of your computer, you can configure the appliance by connecting it to your computer with a serial null modem cable. You must have physical access to the appliance.

For more information, see the Assigning a Management IP address through the serial port procedure.

Setting up the SD-WAN Appliance

Aug 09, 2017

To set up your NetScaler SD-WAN Appliance hardware, see the instructions documented in the Setting up the Appliance Hardware section.

Deployment Modes

Aug 09, 2017

SD-WAN 4000/4100/5100 SE appliances have one recommended deployment mode: Inline. This mode is commonly used without high availability (HA), and less commonly with HA.

Although not all of the following modes are recommended at this time, they are all supported:

- Inline mode
- Inline mode in HA
- Virtual inline mode
- Virtual inline mode in HA

Virtual Inline Mode

Aug 09, 2017

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance and the appliance forwards the processed packets back to the router.

The tasks for configuring virtual inline mode are performed on the router.

See the Virtual Inline Deployment mode instructions for more information.

Configuring the High Availability Setup on the Appliances

Aug 09, 2017

High Availaiblity (HA) is supported in all deployment modes, and the HA configuration procedure is the same for all modes. The two appliances should be running identical hardware, licensing, and software releases, and must be deployed identically, using the same deployment modes on the same subnets.

To learn more about setting up high availability on SD-WAN appliances, see the High Availability section.

NetScaler SD-WAN 210 Standard Edition Appliances

May 09, 2018

The 210-SE appliance is a 1U appliance for use in small branch offices. This appliance has 2-core processor with 4 GB memory and 64 gigabytes (GB) of storage.

The Citrix Compliance Regulatory models are:

- SD-WAN 210-SE (non-LTE) - NS-SDW-210
- SD-WAN 210-SE LTE - NS-SDW-210-LTE-R1, NS-SDW-210-LTE-R2

For more information, see the NetScaler product platform datasheet.

NetScaler SD-WAN 210 SE

May 09, 2018

This appliance has 2-core processors with 4 GB memory and 64 gigabytes (GB) of storage.

The following figure shows the front panel of the 210 SE appliance.

> **Note**
>
> The 210-SE appliance cannot be configured as an MCN appliance.

Figure 1. SD-WAN 210-SE front panel

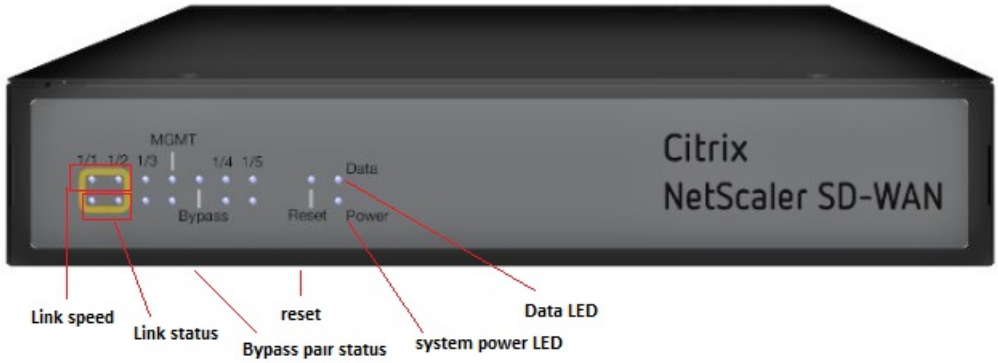

LED	Description
Ethernet Copper Ports LED	Active/Link: **Green** Speed 1000: **Orange** Speed 100: **Green** Speed 10 : **off**
Bypass LEDs	Normal Mode: **Green** Bypass Mode: **Orange**
Ethernet Fiber Ports	Active/Link: **Green** Speed 1000: **Orange**
System and Data LEDs	System Power on: **Green** System Power off: **off** Data access storage: **Blue**

Figure 2. SD-WAN 210 SE back panel

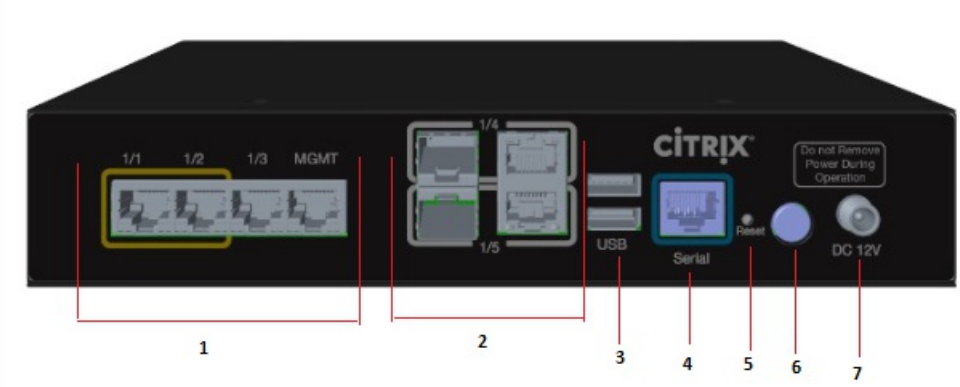

The following components are visible on the back panel of the 210 SE appliance:

Interface	Port Labels	Type	Description
1	1/1 and 1/2	Bypass/FTW	Fail-to-wire
	1/3	Traffic	Network traffic
	Management	RJ-45	A copper Ethernet (RJ45) management port. The management port is used to connect directly to the appliance for system administration functions. You can use this port for initial provisioning of Virtual WAN.
2	1/4 and 1/5	SFP and Ethernet (combination ports)	Used as a combination of SFP and Ethernet one each on the top and bottom.
3	USB	2	USB ports
4	Console	RS-232 serial	A RS232 serial console port
5	Reset	Reset button	Consult Citrix technical support for more information.
6	Power	Power button	Power button to power on or off the appliance. Press the switch for five seconds to switch off the power.
7	Power Supply	DC Power Supply	Single power adapter. Power rating: 40 W, voltage: 12 V, and current: 3.33 A.

The Citrix NetScaler SD-WAN 210 SE LTE appliance is a 1U appliance. This appliance has 2-core processor with 4 GB memory and 64 gigabytes (GB) of storage.

The following figure shows the front panel of the 210 SE LTE appliance.

Figure 3. Citrix NetScaler SD-WAN 210 SE (LTE) front panel with Antenna

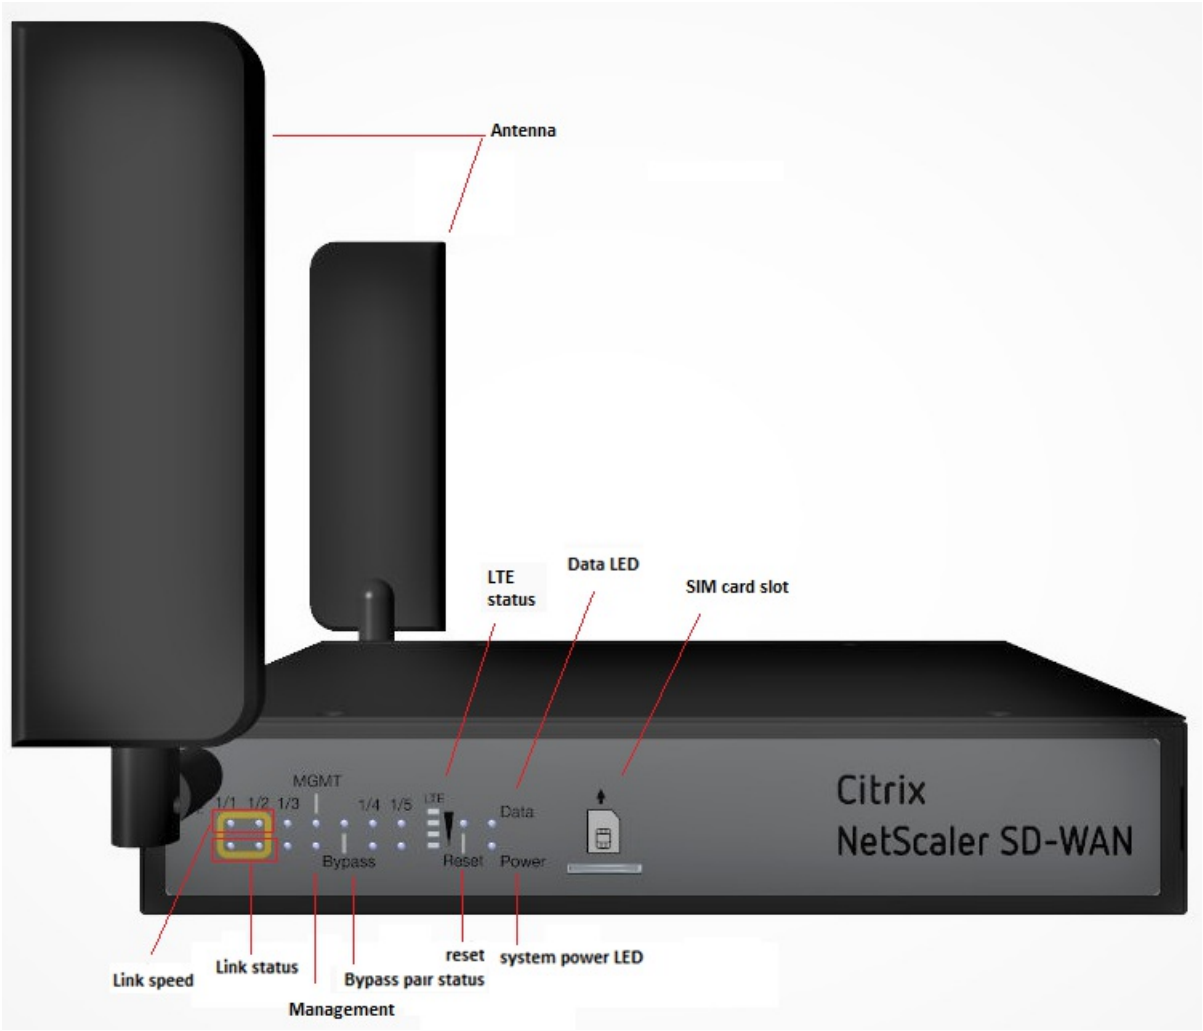

LED	Description
Ethernet Copper Ports LED	Active/Link: **Green** Speed 1000: **Orange** Speed 100: **Green** Speed 10: off
Bypass LEDs	Normal Mode: **Green** Bypass Mode: **Orange**
Ethernet Fiber Ports	Active/Link: **Green** Speed 1000: **Orange**
System and Data LEDs	System Power on: **Green** System Power off: off Data access storage: Blue

Figure 4. Citrix NetScaler SD-WAN 210 SE (LTE) back panel with Antenna

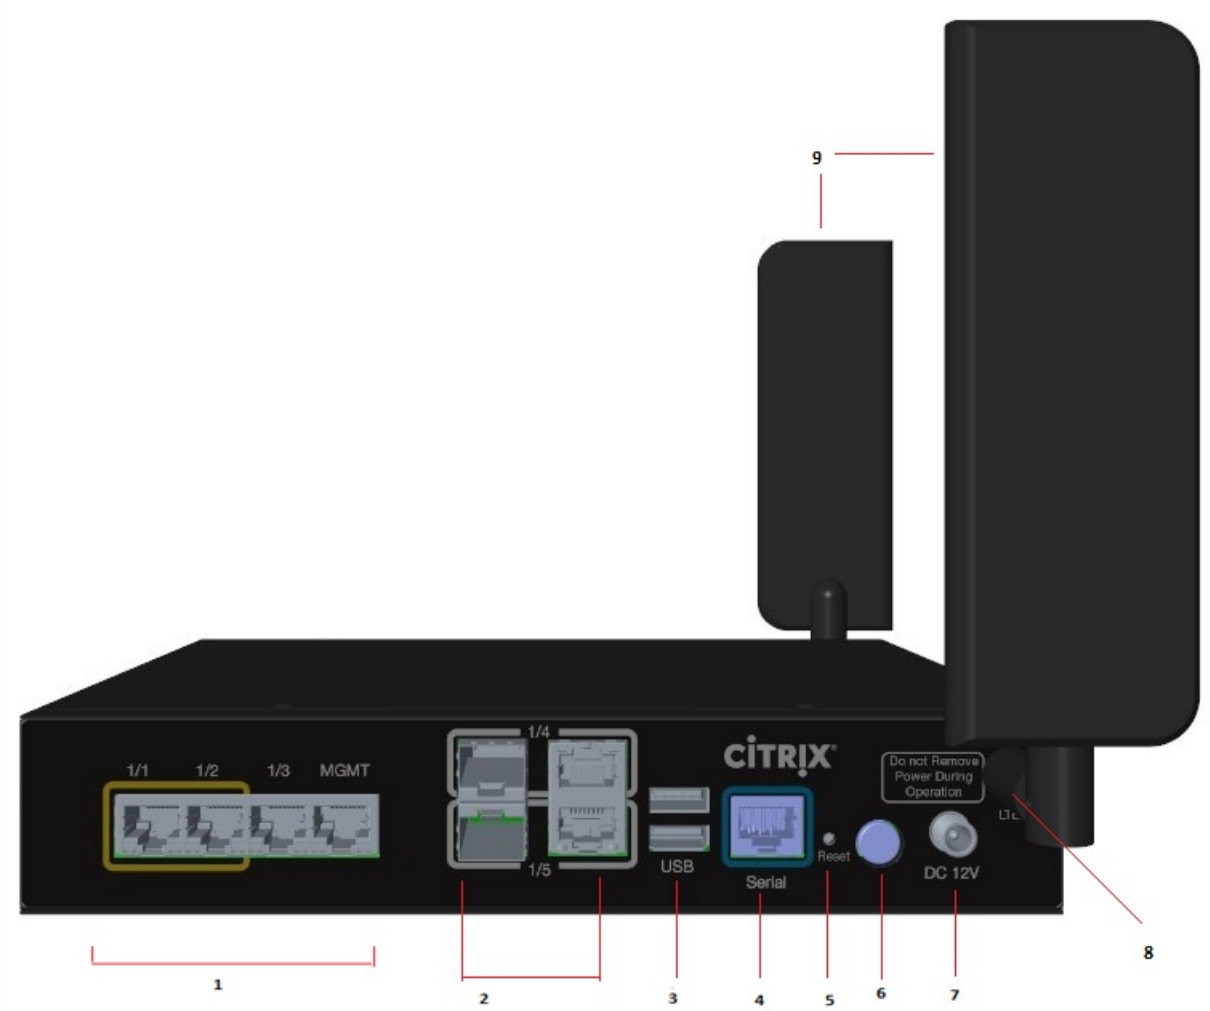

The following components are visible on the back panel of the 210 SE LTE appliance:

Interface	Port Labels	Type	Description
1	1/1 and 1/2	Bypass/FTW	FTW ports
	1/3	Traffic	Traffic port
	Management	RJ-45	A copper Ethernet (RJ45) management port. The management port is used to connect directly to the appliance for system administration functions. You can use this port for initial provisioning of Virtual WAN.
2	1/4 and 1/5	SFP and Ethernet (combination ports)	Used as a combination of SFP and Ethernet one each on the top and bottom.
3	USB	2	USB ports
4	Console	RS-232 serial	A RS232 serial console port
5	Reset	Reset button	Consult Citrix technical support for more information.
6	Power	Power button	Power button to power on or off the appliance. Press the switch for five seconds to switch off the power.
7	Power supply	DC power supply	Single power adapter. Power rating: 40 W, voltage: 12 V, and current: 3.33 A.
8	Antenna Connectors	Male connectors	Connectors for antenna
9	Two antenna	LTE antennas	Antennas shipped with the appliance.

Summary of Hardware Specifications

May 09, 2018

The following table summarizes the specifications of the SD-WAN 210-SE hardware platform.

Specifications	210-SE and 210-SE LTE
Regulatory Model Number	NS-SDW-210 NS-SDW-210-LTE-R1 NS-SDW-210-LTE-R2
Processors	Intel C3338
Memory	4GB DDR4 PC4-2400 SODIMM
Number of power adapters	1
AC power supply (adapter) voltage, frequency and current	115-230V 50-60Hz 1.5A
Maximum AC power consumption	18.9W
Maximum DC power consumption	13.2W
Airflow (front to rear)	Fan-less
Heat dissipation	45.0384 BTU
Package weight (lbs.)	5 lbs
System weight (lbs.)	2.9 lb
Height	4.2 mm
Width	231.6mm
Depth	174.3mm
Operating temperature	0 to 40° C
Humidity range (non-condensing)	5% to 90% RH
Safety certifications	UL
Regulatory Certifications	FCC, CE
Environmental compliance	RoHS/REACH/PFOS/CoM /WEEE
Safety certifications	c UL us listed
[illegible]	Sierra Wireless™ EM7455

Specifications	210-SE and 210-SE LTE
LTE Bands	1, 2, 3, 4, 5, 7, 8, 12, 13, 20, 25, 26, 29, 30, 41 1, 3, 5, 7, 8, 18, 19, 21, 28, 38, 39, 40, 41
Regulatory Compliance	FCC, IC, CE, ICASA, ENACOM, IFETEL ANATEL, RCM, BIS
Industry Certifications	PTCRB, GCF

Preparing for Installation

May 09, 2018

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Once you are satisfied that your appliance has been delivered to your expectations, verify that the location where the appliance will be installed meets temperature and power requirements and that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel should install, maintain, or replace the appliance, and efforts should be taken to ensure that all cautions and warnings are followed.

Unpacking the SD-WAN Appliance

May 09, 2018

The hardware accessories for your particular appliance, such as cables, adapters, and rail kit, vary depending on the hardware platform you ordered. Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

If you ordered a SD-WAN 210 appliance, the box should contain:

- The appliance you ordered
- One RJ45 console cable
- One power cable
- One power adapter
- Rubber feet -4

Accessory FRUs: Optional

- Rackmount Kit
- Wall mount Kit

If you ordered a 210-SE LTE appliance, verify that the following components and accessories are included:

- The NetScaler SD-WAN 210 SE appliance

accessory contents:

- Power adapter - 1
- Power cable - 1
- Rubber feet – 4
- Two LTE antenna – self assembly

The mounting LTE male connectors are assembled to the shipped appliance.

In addition to the items included in the box with your new appliance, you need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you will connect to your network.
- One available Ethernet port on your network switch or hub for each Ethernet port you want to connect to your network.
- A computer to serve as a management workstation.

Preparing the Site and Rack

May 09, 2018

SD-WAN appliances have specific site and rack requirements. You must make sure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and will help ensure a smooth installation.

The appliance should be installed in a server room or server cabinet with the following features:

Environment control

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment.

Power density

Wiring capable of handling at least 4,000 watts per rack unit in addition to power needs for the CRAC.

The rack on which you install your appliance should meet the following criteria:

Rack characteristics

Racks should be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it should be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks should be installed in rows spaced at least 1 meter/3 feet apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

Power connections

At minimum, two standard power outlets per unit.

Network connections

At minimum, four Ethernet connections per rack unit.

Space requirements

One empty rack unit for each SD-WAN 210 appliances.

Appliance Installation Precautions

1. Determine the placement of each component in the rack before you install the rack.
2. Install the equipment near an electrical outlet for easy access.
3. Mount equipment in a rack with sufficient airflow for safe operation.
4. For a closed or multiple-unit rack assembly, the ambient operating temperature of the rack environment might be greater than the ambient temperature of the room. Therefore, consider the lowest and highest operating temperatures of the equipment when making a decision about where to install the appliance in the rack.

The NetScaler SD-WAN 210-SE appliance can be installed in the following three installation modes:

- Desktop mount
- Rackmount

- Wall mount

Cautions and Warnings

May 09, 2018

During installation or maintenance procedures, wear a grounding wrist strap to avoid ESD damage to the electronics of the appliance. Use a conductive wrist strap attached to a good earth ground or to the appliance. You can attach it to the connector beside the ESD symbol on the back.

Follow basic electrical safety precautions to protect yourself from harm and the appliance from damage.

- Be aware of the location of the emergency power off (EPO) switch, so that you can quickly remove power to the appliance if an electrical accident occurs.
- Do not use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as electrical insulators.
- Make sure that the power supply cords include grounding plugs and are plugged into grounded electrical outlets.
- Make sure that the power source can handle the appliance's maximum power consumption rating with no danger of an overload.
- A reliable ground must be maintained at all times. Therefore, the rack should be grounded. Pay particular attention to power supply connections other than the direct connection to the branch circuit (for example, connections to power strips).

Warning

- Electrostatic discharge (ESD) can damage your equipment.
- Do not place any objects on the appliance.
- Do not cover vent holes on the side of the appliance.
- Metal surface of the appliance can get heated up.
- Use caution when touching the metal surface of the appliance.

- Make sure that the leveling jacks on the bottom of the rack are fully extended to the floor, with the full weight of the rack resting on them.
- For a single-rack installation, attach a stabilizer to the rack.
- For a multiple-rack installation, couple (attach) the racks together.
- Always make sure that the rack is stable before extending a component from the rack.
- Extend only one component at a time. Extending two or more simultaneously might cause the rack to become unstable.
- The handles on the left and right of the front panel of the appliance should be used only for extending the appliance out of the rack. Do not use these handles for mounting the appliance on the rack. Use the rack-rail hardware, described later, instead.

Installing the Hardware

May 09, 2018

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

The NetScaler SD-WAN 210-SE appliance can be installed in the following three installation modes:

- Desktop mount
- Rackmount
- Wall mount

Mounting the Appliance

May 09, 2018

The rackmount chassis of SD-WAN 210-SE fits a standard rack and takes 1U of racking height.

The NetScaler SD-WAN 210-SE appliance can be placed on any flat surface, or mounted in any standard rack unit with the provided rack-mount brackets and screws.

To install the appliance into a rack:

1. Ensure that the SD-WAN 210-SE appliance is placed on a stable surface prior to rack-mount installation.

2. Attach the provided rack-mount brackets to the sides of the appliance using the provided bracket screws.

    a. If you are installing the appliance into a four-post rack, attach the rack-mount brackets with the handles aligned with the front of the SD-WAN 210-SE appliance.

    b. If you are installing the appliance into a two-post rack, attach the rack-mount brackets with the handles aligned with the middle of the SD-WAN 210-SE appliance.

3. Position the SD-WAN 210-SE appliance in the rack. Ensure there is enough room around the device to allow for sufficient air flow.

4. Line up the rack-mount bracket holes to the holes on the rack and ensure that the SD-WAN 210-SE appliance is level.

5. Finger tighten four rack-mount screws to attach the appliance to the rack.

6. Tighten the rack-mount screws with an appropriate screwdriver.

7. Plug the provided power cable.

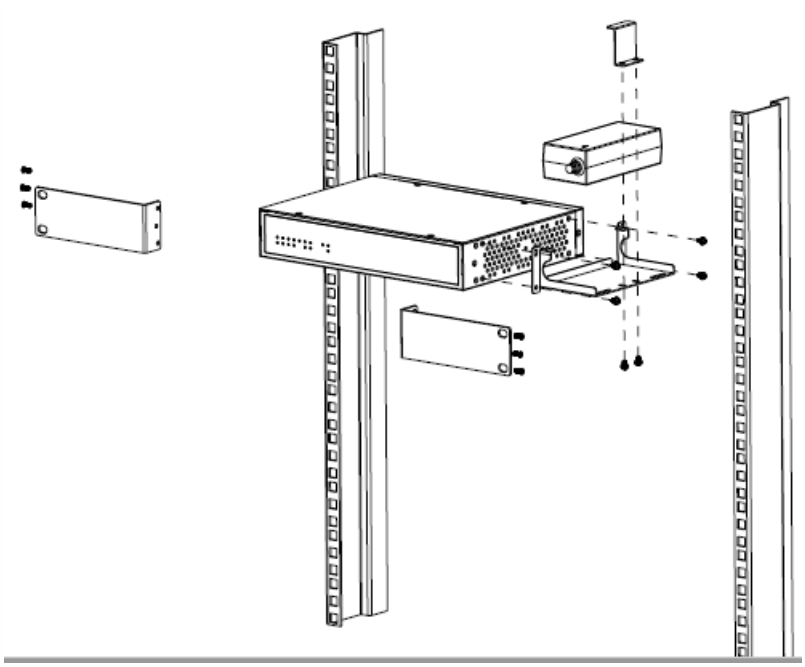

Rack Mount the Appliance with Antenna

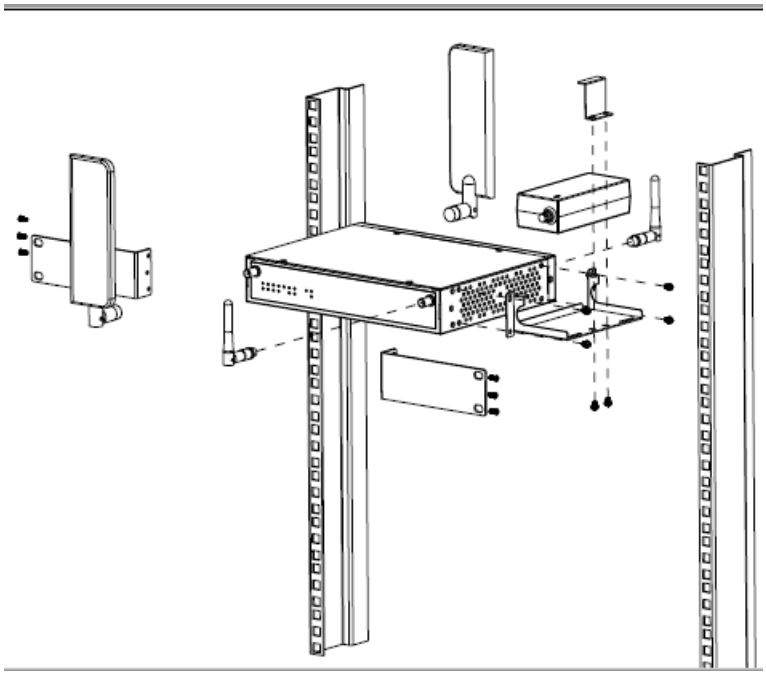

To install the appliance into the wall:

1. Using the supplied drywall screws and anchors, drill 2 holes into your wall lined up with the two matching screw holes on the rear of the mount housing of the appliance. Hang the appliance with the 4 (2 on each side) screws.

a) Align the chassis ears brackets against the front and back of the mounting holes provided and properly attach ears bracket to the SD-WAN chassis by tightening the 4 screws (each side 2) provided.

b) Mount the SD-WAN 210-SE chassis using the 4 ear brackets (each side 2) as shown below on both sides to the wall.

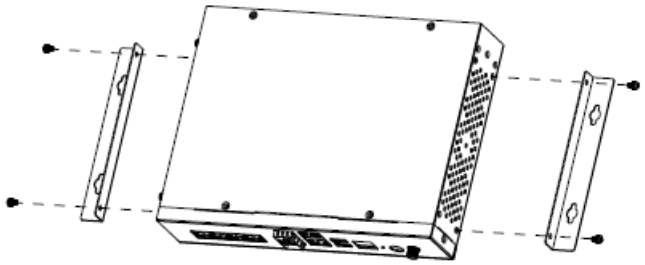

Wall Mount the Appliance with Antenna

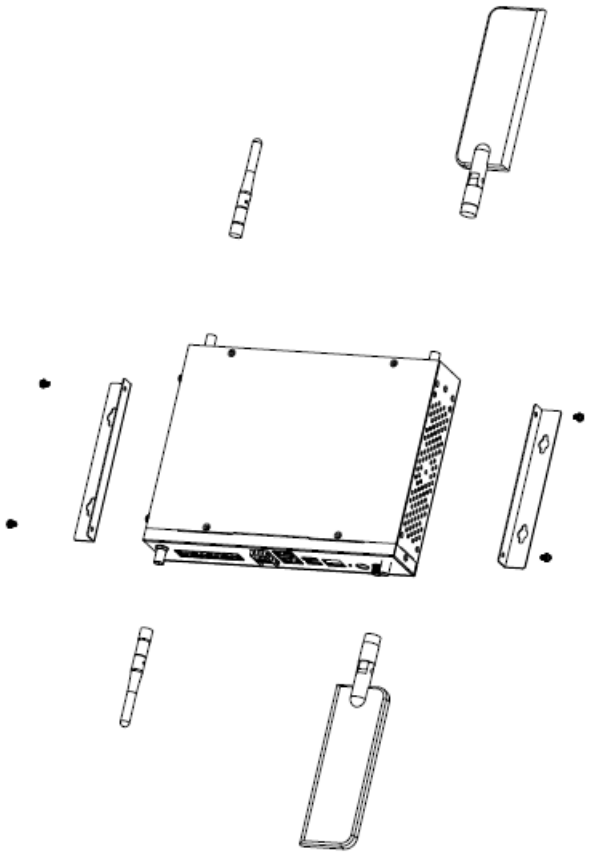

SD-WAN 210-SE appliance can be desktop mounted using the rubber feet shipped in the appliance package.

Connecting the Cables

May 09, 2018

When the appliance is securely mounted on the rack, determine which ports you should use. You are then ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

> ### Warning
>
> Remove all jewelry and other metal objects that might come in contact with power sources or wires before installing or repairing the appliance. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly, and may cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port.
   Note: To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.
2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

> ### Warning
>
> Connecting multiple network ports to the same switch or VLAN can result in a network loop.

A SD-WAN appliance has one power supply, unless you have installed a second. A separate ground cable is not required, because the three-prong plug provides grounding. Provide power to the appliance by installing the power cord.

To connect the appliance to the power source

Connect the power cable to one of the inlet receptacles on the back of the appliance, and connect the other end of the power cable to a power outlet on the back panel of the appliance, next to the power supply.

The Citrix logo and the LCD on the front of the NetScaler SD-WAN illuminate after the appliance starts, and the LCD indicates the operational status of the appliance.

Switching on the Appliance

May 09, 2018

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

1. Verify that the appliance is connected through a console or Ethernet port. This will ensure that you can configure the appliance after it is switched on.
2. Depending on the appliance, press the ON/OFF toggle power switch or the power button to switch on the appliance.

Connecting the Appliance to the Network

Connect the interfaces on the appliance to the network ports on the appropriate switches by using Ethernet/fiber optic cables.

> **Warning**
>
> Connecting multiple network ports to the same switch or VLAN can result in a network loop.

Initial Configuration

May 09, 2018

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

You can connect the ethernet cable to the appliance management port, and login to the web management interface using the default IP address.

To perform initial configuration:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Perform additional configuration for Windows.
- Assign management IP address through the serial console.
- Troubleshoot initial configuration issues.

By default, the initial configuration deploys the appliance in inline mode.

> ## Note
>
> On the SD-WAN 210-SE appliance, the default static IP address is 192.168.100.1. It also has DHCP enabled (by default) for Management access. When you enable the SD-WAN service, or install SD-WAN license using default static IP (192.168.100.1), this IP is lost.
>
> You should assign a static IP address or be notified of the DHCP IP address that is assigned to this appliance. In order to see the DHCP IP address, you can log into the SD-WAN command line interface using admin/password, to view the management IP address.

Prerequisites

May 09, 2018

> ## Note
>
> The default static IP address is 192.168.100.1. The static IP address has DHCP enabled (by default) for Management access.
>
> When you enable the SD-WAN service, or install SD-WAN license using the default static IP (192.168.100.1), this IP is disabled, and you can obtain DHCP address within the network.
>
> Always assign a Static IP address or be aware of the DHCP address that is assigned to the appliance. To view the DHCP IP address, you can login to SD-WAN CLI (using admin/password) that displays the management IP or go to **Configuration** -> **Appliance Settings** > **Network Adapters** > **Ethernet**.
>
> 1. Ensure that you have permanent DHCP address assigned to SD-WAN appliances.
> 2. The DHCP address should be associated to the management NIC address.
> 3. Connect the management NIC address to the DHCP enabled LAN or reboot the applaince when ready.

Before you begin configuring the appliance, make sure that the following prerequisites have been met:

- You should have physical access to the appliance.
- In the Worksheet, record all IP addresses and other values you would use to configure the appliance. Preferably, print out the worksheet before you start the configuration process.
- You should already have a SD-WAN license key from Citrix, sent in an email. If you are using remote licensing, you need the IP address of the licensing server.
- WAN Send and Receive Speeds.

Setting up the SD-WAN Appliance

May 09, 2018

1. If you are configuring the appliance using Zero Touch Deployment (ZTD), refer to the following link on the docs.citrix.com site; https://docs.citrix.com/en-us/netscaler-sd-wan/9-3/zero-touch-deployment-service.html.

2. If you are configuring a hardware SD-WAN appliance, physically connect the appliance to a PC. Refer to the following link on the docs.citrix.com site; https://docs.citrix.com/en-us/netscaler-sd-wan/9-3/getting-started/setting-up-virtual-wan-appliances/appliance-hardware.html

> **Note**
>
> When you connect AC power to SD-WAN 210-SE, it will boot immediately without pressing the power button.

NetScaler SD-WAN 1000 and 2000 Enterprise Edition Appliances

Aug 09, 2017

The SD-WAN Enterprise Edition 1000 and 2000 appliances combine virtualized instances of WAN optimization and Virtual WAN functionality installed on the appliance.

It offers a combination of Virtual WAN and WAN Optimization capabilities.

The SD-WAN 1000 EE and 2000 EE appliances are based on the Citrix branch architecture, which supports multiple virtual machines. All branch appliances contain a SD-WAN instance, a management service instance, and a Xen hypervisor.

The SD-WAN instance is typically used in inline mode, with the SD-WAN instance interposed between the WAN router and the LAN, so WAN traffic flows through the accelerated bridge. The SD-WAN instance can also be deployed in virtual inline mode, using a single accelerated bridge port.

In addition to the accelerated bridges and the Windows LAN port, a management port connects to all virtual machines (instances) and the hypervisor.

The appliance has two modes, two-port mode and four-port mode, which determine how ports 1/3 and 1/4 are used.

SD-WAN 2000 EE Appliance

Aug 09, 2017

The Citrix NetScaler SD-WAN 2000 EE platform is a 1U appliance with one quad-core processor and 24 gigabytes (GB) of memory.

The following figure shows the front panel of the SD-WAN 2000 EE appliance.

Figure 1. Citrix NetScaler SD-WAN 2000 EE appliance, front panel

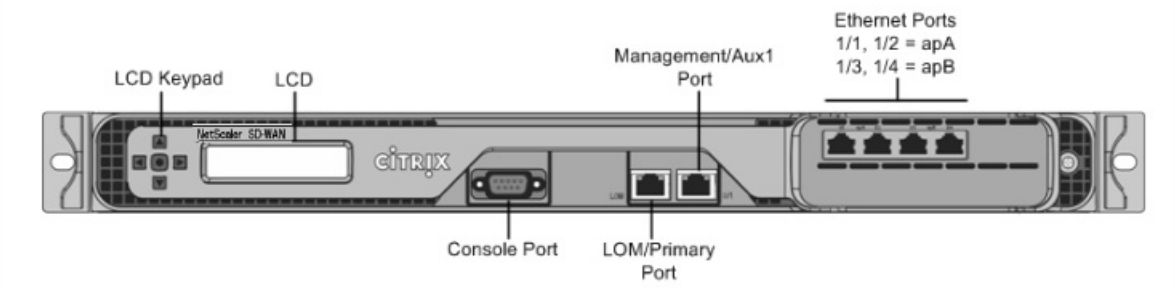

SD-WAN 2000 EE appliance has the following ports:

- An RS232 serial console port.
- A copper Ethernet (RJ45) Port called the Lights out Management (LOM) port. You can use this port to remotely monitor and manage the appliance independently of the appliance's software.
- A copper Ethernet (RJ45) management port, numbered 0/1, and named PRI (primary). The management port is used to connect directly to the appliance for system administration functions. You can use this port for initial provisioning of WAN optimization and Windows Server.
  Note: The LOM port also operates as a management port.
- Four 10/100/1000Base-T copper Ethernet ports numbered 1/1, 1/2, 1/3, and 1/4 from left to right. The four ports form two *accelerated pairs*, which function as accelerated bridges. Ports 1/1 and 1/2 are accelerated pair A (apA), and 1/3 and 1/4 are accelerated pair B (apB).

The following figure shows the back panel of the SD-WAN 2000 EE appliance.

Figure 2. Citrix SD-WAN 2000 EE appliance, back panel

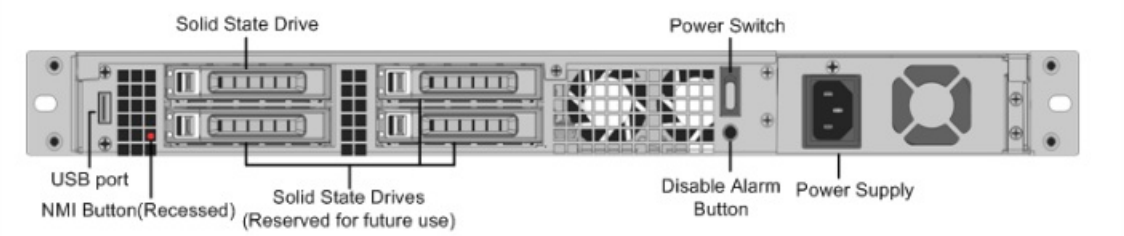

The following components are visible on the back panel of the SD-WAN 2000 EE appliance:

- 600 GB removable solid-state drive, which stores the appliance's software and user data, and 1 TB hard disk drive.
- Power switch, which switches power to the appliance on or off. Press the switch for five seconds to switch off the power.

- USB port (reserved for a future release).
- Non-maskable interrupt (NMI) button, for use at the request of Technical Support to produce a core dump. You must use a pen, pencil, or other pointed object to press this red button, which is recessed to prevent unintentional activation.
- Single power supply, rated at 300 watts, 100-240 volts.

SD-WAN 1000 EE Appliance

Mar 23, 2018

The Citrix NetScaler SD-WAN 1000 EE platform has a quad-core processor and 32 GB of memory. This platform has a bandwidth of up to 100 Mbps.

The following figure shows the front panel of an SD-WAN 1000 EE appliance.

Figure 1. Citrix NetScaler SD-WAN 1000 EE, front panel

The front panel of the SD-WAN 1000 EE appliance has a power button and five LEDs. The power button is used to switch the appliance on or off. The reset button restarts the appliance.

The LEDs provide critical information related to different parts of the appliance.

- Power Fail – Indicates that the power supply unit has failed.

Information LED – Indicates the following:

Status	Description
Continuously ON and red	The appliance is overheated. (This might be a result of cable congestion.)
Blinking red (1 Hz)	Fan failure, check for an inoperative fan.
Blinking red (0.25 Hz)	Power failure, check for the non-operational power supply.
Solid blue	Local UID has been activated. Use this function to locate the server in a rack mount environment.
Blinking blue (300 m/s)	Remote UID is on. Use this function to identify the server from a remote location.

- NIC1 and NIC2 – Indicate network activity on the LAN1 and WAN1 ports.
- HDD – Indicates the status of the hard disk drive.
- Power – Indicates that the power supply units are receiving power and operating normally.

The following figure shows the back panel of an SD-WAN 1000 EE appliance.

Figure 2. Citrix NetScaler SD-WAN 1000 EE appliance, back panel

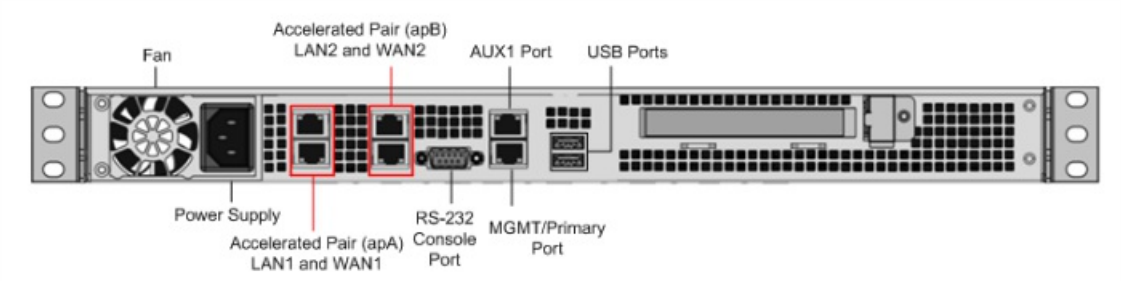

The following components are visible on the back panel of an SD-WAN 1000 EE appliance:

- Cooling fan.
- Single power supply, rated at 200 watts, 110–240 volts.
- Accelerated pairs of Ethernet ports (apA and apB) which function as accelerated bridges.
- RS-232 serial console port.
- One AUX Ethernet port and one management port.
- Two USB ports.

To power on the appliance after a graceful shut down:

1. Connect a Serial console cable to the rear of the appliance and to the serial port on a management laptop.

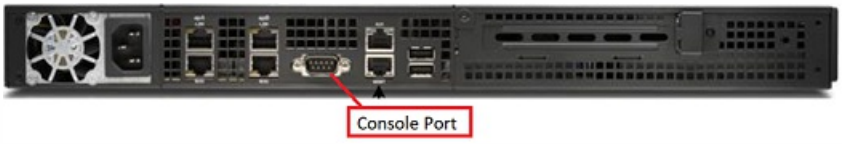

2. On the management laptop, restart a putty session using the following configuration settings:

- Serial line: COM1
- Speed: 9600

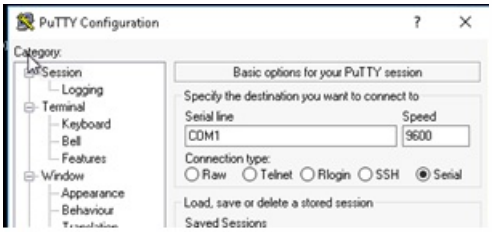

3. Power on the appliance and as it is booting, press the following key in the Putty session to enter the BIOS configuration screen.

Keypress: **DEL**

4. When in the BIOS, navigate to,

- Advanced Tab > **Select**
- Boot Feature > **Enter**

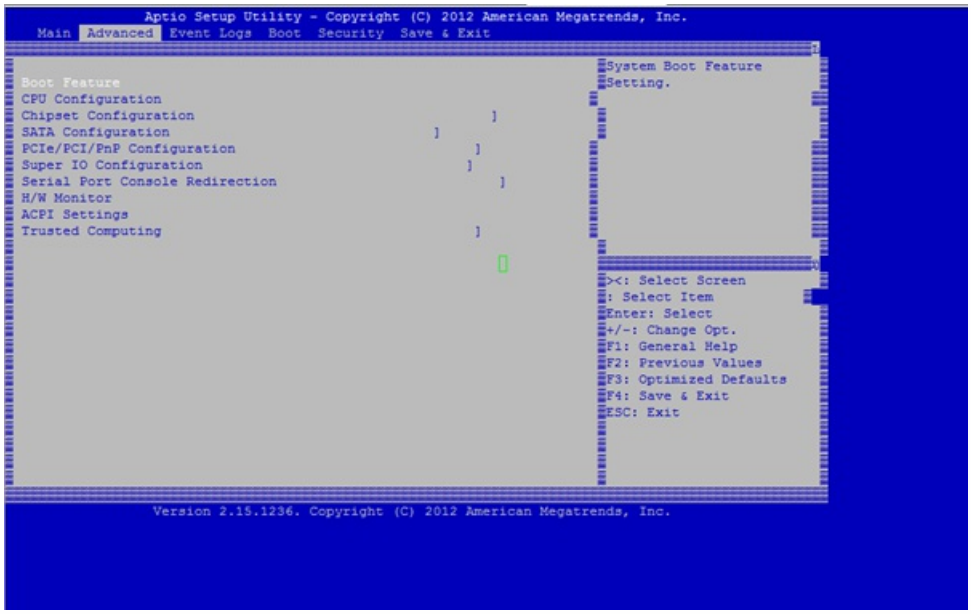

5. When in the Boot Feature screen, change the value of the parameter **Restore on AC Power Loss**; from **Last State** > **Power ON**.

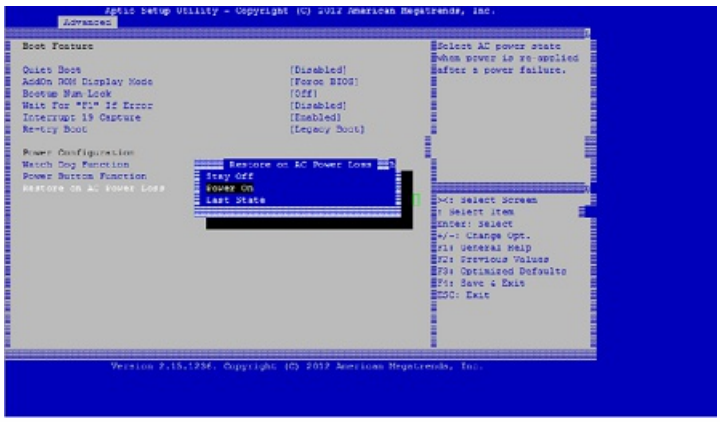

6. Navigate to Save and Exit.

- Select **Save changes** and **Reset**
- Select **Yes**

Allow the system to restart. This takes approximately five minutes.

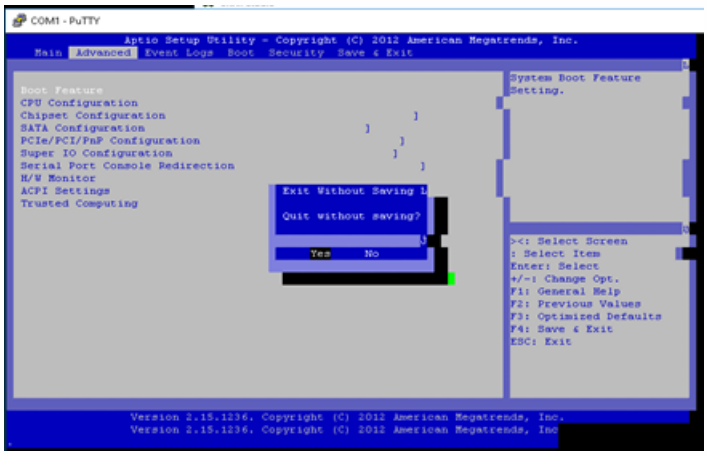

7. After the appliance is powered on, login to the appliance management instance (SVM). The default IP address for

the appliance is: 192.168.100.1, user name is: admin/password.

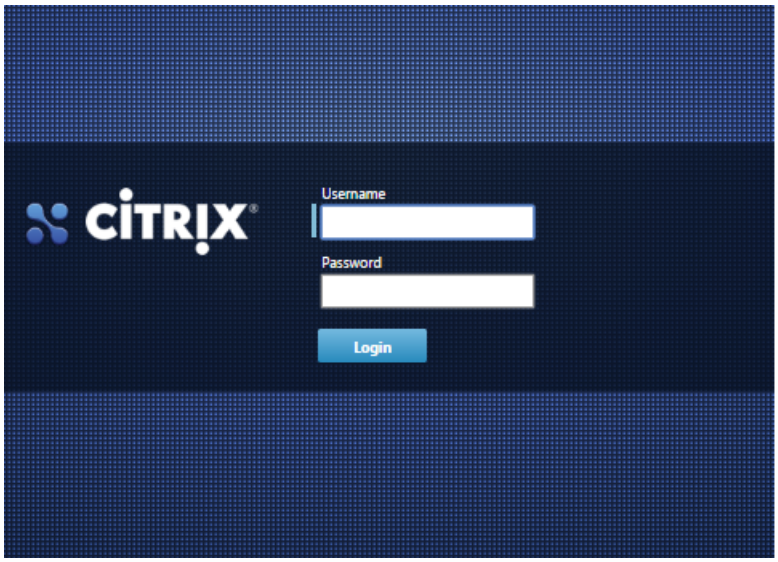

8. In the SD-WAN appliance GUI, navigate to **Configuration** > **Maintenance** > **Reboot Appliance**. Allow the appliance to fully shut down. Ensure that there are no power lights on the appliance when the shut down process has completed.

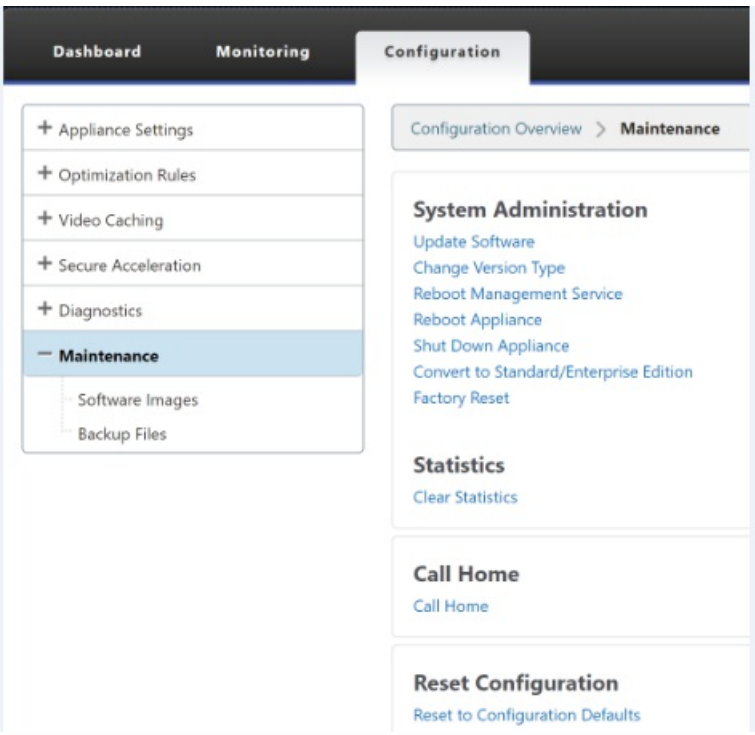

9. Power on the appliance to confirm that the BIOS configuration change has been applied successfully. This can be either done through the APC intelligent PDU Web Management console or by physically pulling the power cable out of the shut down SD-WAN appliance, waiting for 10 seconds and then plugging it back in again. The appliance power ups automatically from all shut down scenarios.

Summary of Hardware Specifications

Aug 09, 2017

The following tables summarize the specifications of the SD-WAN 1000 EE and 2000 EE platforms.

	SD-WAN 1000 EE	SD-WAN 2000 EE
Platform Performance		
Bandwidth	Up to 100 Mbps	Up to 250 Mbps
Maximum HDX sessions	Up to 100	300
Total sessions	10,000	20,000
Acceleration Plug-in CCUs	N/A	750
Hardware Specifications		
Processor	4 Cores	4 Cores
Total disk space	1x300 GB SSD and 1x1 TB HDD	1 x 600 GB SSD and 1X1 TB HDD
SSD (dedicated Compression history)	123 GB for Disk-Based Compression (DBC) 25 GB for video caching	225 GB for Disk-Based Compression (DBC) 50 GB for video caching
RAM	32 GB	24 GB
Network Interfaces	2 pair with bypass 10/100/1000 2 GigE ports for Management and AUX ports	4 x 10/100/1000 Base-T copper Ethernet 2 GigE ports for Management and AUX ports
Power supplies	1	1
Physical Dimensions		
Rack Units	1U	1U
System width	EIA 310-D for 19-inch racks	EIA 310-D for 19-inch racks
System depth	10" (25.4 cm)	25.4" (64.5 cm)
System weight	8.5 lbs (3.9 kg)	32 lbs (14.5 kg)
Shipping dimensions and weight	26 L x 18.5 W x 6.5" H 14.5 lbs	32 L x 23.5 W x 7.5" H 39 lbs

Environmental and Regulatory	SD-WAN 1000 EE	SD-WAN 2000 EE
Voltage	100/240 VAC, 50-60 Hz	100/240 VAC, 50-60 Hz
Power consumption (Max.)	200 W	300 W
Operating Temperature (degree Celsius)	10–35	0-40
Non-operating Temperature (degree Celsius)	-40 – +70	-40 – +70
Allowed Relative Humidity	8% – 90% non-condensing	5%–95%
Safety certifications	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)	CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)
Electromagnetic and susceptibility certifications	FCC (Part 15 Class A), CCC, KCC, NOM, SASO, CITC, EAC, DoC, CE, VCCI, RCM	FCC (Part 15 Class A), CCC, KCC, NOM, SASO, CITC, EAC, DoC, CE, VCCI, RCM
Environmental certifications	RoHS, WEEE	RoHS, WEEE

Ethernet Port Names

Aug 09, 2017

When configuring the appliance, you have to specify IP addresses for various Ethernet ports of the appliance. The Ethernet ports are named differently on the front panel of SD-WAN 1000 EE and 2000 EE appliances in the NetScaler SD-WAN instance, as shown in the following table:

Front Panel		SD-WAN Instance
SD-WAN 1000 EE	SD-WAN 2000 EE	
MGMT (Blue)	0/1 (LOM/PRI)	Primary
AUX	0/2 (AUX)	Aux
apA LAN1/WCCP (Green)	1/1	apA.1
apA WAN1	1/2	apA.2
apB LAN2	1/3	apB.1*
apB WAN2	1/4	apB.2*

* Available to the SD-WAN instance only in four-port mode.

Installing the Appliance

Aug 09, 2017

After you have determined that the location where you will install your appliance meets the environmental standards and the server rack is in place according to the instructions, you are ready to install the hardware. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. You can also connect the appliance to a computer through Ethernet port for initial configuration. On SD-WAN 1000 EE appliance, this port is labeled as MGMT (management) port and on SD-WAN 2000 EE, the port is labeled as PRI (primary) port. To complete the installation, you switch on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Rack Mounting the Appliance

Aug 09, 2017

A SD-WAN 1000 EE or 2000 EE appliance requires one rack unit. Both are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

Rack Mounting a SD-WAN 1000 EE Appliance

Aug 09, 2017

SD-WAN 1000 EE appliance is not shipped with rails. You can mount the appliance to the rack by using the front mounting ports.

Rack Mounting a SD-WAN 2000 EE Appliance

Aug 09, 2017

A SD-WAN 2000 appliance requires one rack unit. Both are rack-mount devices that can be installed into two-post relay racks or four-post EIA-310 server racks. Verify that the rack is compatible with your appliance.

To mount a SD-WAN appliance, you must first install the rails and then install the appliance in the rack, as follows:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.

To remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Depress the locking tabs until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

To attach the inner rails to the appliance

1. Position the right inner rail behind the ear bracket on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws.
4. Repeat steps 1 through 3 to install the left inner rail on the left side of the appliance.

To install the rack rails

1. Position the rack rails at the desired location in the rack, keeping the sliding rail guide facing inward.
2. Snap the rails to the rack.
   Note: Make sure that both rack rails are at same height and that the rail guides are facing inward.

To install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides, and push the appliance into the rack rails until it locks into place.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack.
   Note: The illustration in the following figure might not represent your actual appliance.
   Figure 1. Rack Mounting the Appliance

Connecting the Cables

Aug 09, 2017

When the appliance is securely mounted on the rack, determine which ports you should use. You are then ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Warning: Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.

Ports

A typical installation using a single accelerated bridge uses four Ethernet ports (the Primary port and apA) and six IP addresses (four on the Primary port's subnet and two on apA's subnet).
The appliance has two motherboard ports and two accelerated bridges.

- The motherboard ports are labeled as MGMT (management) and AUX1 (auxiliary) ports in SD-WAN 1000 appliance and PRI (primary) and AUX (auxiliary) in SD-WAN 2000 appliance. You use MGMT port of the SD-WAN 1000 appliance and PRI port of the SD-WAN 2000 appliance for initial configuration.
- Accelerated bridge ports are apA and apB are available on the back panel of SD-WAN 1000 appliance and the front panel of SD-WAN 2000 appliance. On SD-WAN 1000 appliance, these ports are labeled as LAN1 and WAN1, and LAN2 and WAN2, respectively. However, on SD-WAN 2000 appliance, these ports are labeled as 1/1 and 1/2, and 1/3 and 1/4, respectively.

Connecting the Ethernet Cables

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port.

To connect an Ethernet cable to a 10/100/1000BASE-T port

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port.
   On SD-WAN **1000** appliance, the ports are available on the back panel and labeled as LAN1 and WAN1 for apA bridged port for LAN and WAN links, respectively.

   On SD-WAN **2000** appliance, the ports are available on the front panel. The ports on SD-WAN 2000 are labeled as 1/1 and 1/2 for the apA bridged port. You can use 1/1 for LAN and 1/2 for WAN link.

2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Connecting the Console Cable

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port.
   On SD-WAN **1000** appliance, the port is located on the back panel.

   On SD-WAN **2000** appliance, the port is located on the front panel.

Note: To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.

2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Connecting the Power Cable

A SD-WAN appliance has one power supply. A separate ground cable is not required, because the three-prong plug provides grounding. Provide power to the appliance by installing the power cord. Connect the other end of the power cable to a standard 110V/220V power outlet.

Switching on the Appliance

Aug 09, 2017

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. After verifying the connections, you are ready to switch on the appliance.

To switch on the appliance

1. Verify that the appliance is connected through a console or Ethernet port, so that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the appliance.
3. On SD-WAN **2000** appliance, verify that the LCD on the front panel is backlit and the start message appears

Caution: Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs, you can quickly remove power from the appliance.

Initial Configuration

Aug 09, 2017

After checking the connections, you are ready to deploy the SD-WAN 1000 and 2000 appliances on the network.

The appliance shipped from Citrix has default IP addresses configured on it. To deploy the appliance on the network, you must configure the appropriate IP addresses on the appliance to accelerate the network traffic.

To perform initial configuration:

- Identify the prerequisites for the initial configuration.
- Record various values required in the initial configuration procedure.
- Configure the appliance by connecting it to the Ethernet port.
- Perform additional configuration for Windows.
- Assign management IP address through the serial console.
- Troubleshoot initial configuration issues.

By default, the initial configuration deploys the appliance in inline mode.

Prerequisites

Aug 09, 2017

Before you begin configuring the appliance, make sure that the following prerequisites have been met:

- You should have physical access to the appliance.
- In the Worksheet, record all IP addresses and other values you would use to configure the appliance. Preferably, print out the worksheet before you start the configuration process.
- You should already have a SD-WAN license key from Citrix, sent in an email. If you are using remote licensing, you need the IP address of the licensing server.
- WAN Send and Receive Speeds.

Configuring the Appliance by Connecting a Computer to the Ethernet Port

Aug 09, 2017

For initial configuration of a SD-WAN appliance, perform the following tasks::

- Configure the appliance for use on your site.
- Install the Citrix license.
- Enable acceleration.
- Enable traffic shaping (inline mode only).

With inline deployments, this configuration might be all you need, because most acceleration features are enabled by default and require no additional configuration.

You can configure the appliance connecting the appliance to your computer through either the Ethernet port or the serial console. The following procedure enables you to configure the appliance by connecting it to your computer through the Ethernet port.

Note: On a SD-WAN 1000 appliance, you use the Ethernet port labeled as MGMT. However, on SD-WAN 2000 appliance, you use the Ethernet port labeled as PRI or LOM.
If you want to configure the appliance by connecting it to the computer through the serial console, assign the management service IP address from your Worksheet by completing the Assigning a Management IP Address through the Serial Console procedure, and then run steps 4 through 25 of the following procedure.

Note: Make sure that you have physical access to the appliance.

To configure the appliance by connecting a computer to the SD-WAN appliance's Ethernet port 0/1

1. Set the Ethernet port address of a computer (or other browser-equipped device with an Ethernet port), to 192.168.100.1, with a network mask of 255.255.0.0. On a Windows device, this is done by changing the Internet Protocol Version 4 properties of the LAN connection, as shown below. You can leave the gateway and DNS server fields as blank.

2. Using an Ethernet cable, connect this computer to the port labeled MGMT on a SD-WAN 1000 appliance, or to the port labeled PRI on a SD-WAN 2000 appliance.
3. Switch on the appliance. Using the web browser on the computer, access the appliance by using the default management service IP address http://192.168.100.1.
4. On the login page, use the following default credentials to log on to the appliance.
5. Start the configuration wizard by clicking **Get Started**.
6. On the Platform Configuration page, enter the respective values from your worksheet, as shown in the following example:

   Note: If, for SD-WAN configuration, you want to use the same network mask and gateway as those for Network Configuration, select the **Use System Netmask and Gateway** option.
7. Click **Done**. A screen showing the Installation in Progress... message appears. This process takes approximately 2 to 5 minutes, depending on your network speed.
   Note: If you are configuring the appliance by connecting it to your computer through the serial console port, skip step 8 through step 14.
8. A Redirecting to new management IP message appears.

9. Click **OK**.
10. Unplug your computer from the Ethernet port and connect the port to your management network.
11. Reset the IP address of your computer to its previous setting.
12. From a computer on the management network, log on to the appliance by entering the new Management Service IP address, such as https://<Managemnt_IP_Address>, in a web browser.
13. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
14. Log on to the appliance.
15. The Configuration wizard starts again. In this wizard, some of the values which you have already provided, appear by default. Specify rest of the values you have recorded in your worksheet.
16. In System Services section, update the values if necessary.
17. In the **Licensing** section, select the appropriate license type. You can either select a local license or a remote license server to apply a license to the appliance.
    1. If you opt for a local license, you must generate a license by using the host ID of the appliance. To generate a local license for the appliance, see http://support.citrix.com/article/ctx131110. To apply the license, you can navigate to the SD-WAN > Configuration > Appliance Settings > Licensing page, after completing the Configuration wizard.
    2. If you opt for a remote licensing server, you must select a remote appliance model and provide the IP address of the licensing server in the **Licensing Server Address** field.
18. In the WAN Link Definition section, specify receive and send speeds for the WAN link in the respective fields. Citrix recommends values 10% lower than the WAN bandwidth, to avoid network congestion.
19. By default, WAN-side adapter settings are configured on the appliance. Accept the default settings.
20. Click **Install**. After the Installation process is complete, the appliance restarts.
21. As soon as the appliance restarts, the Dashboard page appears.

22. To configure the appliance to accelerate the network traffic, open navigate to the **Configuration** tab.
    Note: Make sure that you have already applied the appropriate license to the appliance.
23. On the Network Adapters page of the Appliance Settings node, verify and, if necessary, assign IP addresses, subnet masks, and gateways to the accelerated bridges (apA and apB) to be used. Applying these changes restarts the appliance.
    Note: You need to assign IP addresses to apA and apB adapters only if you intended to configure WCCP mode, virtual inline mode, or the Video Caching feature on the appliance.
24. The Initial Configuration is complete. Traffic now flows through the appliance. The Dashboard page shows this traffic.

25. You need additional configuration on the appliance if you intend to use some of the modes and features, such as, virtual inline mode, video caching, secure peering, high availability, encrypted CIFS/MAPI acceleration, AppFlow monitoring, or SNMP monitoring.

Note:

- Inline installations place the appliance between your LAN and WAN routers, using both ports of the accelerated bridge, such as ports LAN1 and WAN1 on a SD-WAN 1000 appliance with Window Server or ports 1/1 and 1/2 on SD-WAN 2000 appliance with Windows Server, for the apA accelerated bridge port.
- WCCP and virtual inline installations connect a single accelerated bridge port to your WAN router.
- Virtual inline installations require that you configure your router to forward WAN traffic to the appliance. See Router Configuration.
- WCCP installations require configuration of your router and the appliance. See WCCP Mode.

Assigning a Management IP Address through the Serial Console

Aug 09, 2017

If you do not want to change the settings of your computer, you can perform initial configuration by connecting the appliance to your computer with a serial null modem cable. Make sure that you have physical access to the appliance.

To configure the appliance through the serial console

1. Connect a serial null modem cable to the appliance's console port.
2. Connect the other end of the cable to the serial COM port of a computer running a terminal emulator, such as Microsoft HyperTerminal, with settings 9600,N,8,1, p.
3. On the HyperTerminal output, press **Enter**. The terminal screen displays the Logon prompt.
   Note: You might have to press **Enter** two or three times, depending on the terminal program you are using.
4. At the logon prompt, log on to the appliance with the following default credentials:
   Username: nsroot

   Password: nsroot.

5. At the **$** prompt, run the following command to switch to the shell prompt of the appliance:
   $ ssh 169.254.0.10

6. Enter **Yes** to continue connecting to the management service.
7. Log on to the shell prompt of the appliance with the following default credentials:
   Password: nsroot.

8. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu:
   # networkconfig

9. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.
10. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the XenServer server.
11. Type **3** and press **Enter** to select option 3, and then specify the network mask for the management service IP address.
12. Type **4** and press **Enter** to select option 4, and then specify the default gateway for the management service IP address.
13. Type **8** and press **Enter** to save the settings and exit.
14. Access the SD-WAN appliance by entering the new management service IP address of the appliance, such as https://<Management_Service_IP_Address>, in a web browser of a computer on the management network.
15. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
16. Run steps 4 through 25 of the Configuring the Appliance by Connecting a Computer to the Ethernet Port procedure to complete the configuration process

Setting up the SD-WAN Appliance

Aug 09, 2017

To set up your NetScaler SD-WAN Appliance hardware, see the instructions documented in the Setting up the Appliance Hardware section.

Deployment Modes

Aug 09, 2017

A SD-WAN appliance acts as a virtual gateway. It is neither a TCP endpoint nor a router. Like any gateway, its job is to buffer incoming packets and put them onto the outgoing link at the right speed. This packet forwarding can be done in different ways, such as inline mode, virtual inline mode. Although these methods are called *modes*, you do not have to disable one forwarding mode to enable another. If your deployment supports more than one mode, the mode that the appliance uses is determined automatically by the Ethernet and IP format of each packet.

Because the appliance supports different forwarding modes and different kinds of non-forwarded connections, it needs a way of distinguishing one kind of traffic from another. It does so by examining the destination IP address and destination Ethernet address (MAC address), as shown in table below. For example, in inline mode, the appliance is acting as a bridge. Unlike other traffic, bridged packets are addressed to a system beyond the appliance, not to the appliance itself. The address fields contain neither the appliance's IP address nor the appliance's Ethernet MAC address.

In addition to pure forwarding modes, the appliance has to account for additional types of connections, including management connections to the GUI and the heartbeat signal that passes between members of a high-availability pair. For completeness, these additional traffic modes are also listed in table below.

Table 1. How Ethernet and IP Addresses Determine the Mode

Destination IP Address	Destination Ethernet Address	Mode
Not appliance	Not appliance	Inline or Pass-through
Not appliance	Appliance	Virtual Inline or L2 WCCP (WAN OP)
Appliance	Appliance	Direct (UI access)
Appliance (VIP)	Appliance	High-Availability. Proxy mode
Appliance (WCCP GRE Packet)	Appliance	WCCP GRE Mode
Appliance (Signaling IP)	Appliance	Signaling Connection (SD-WAN plugin Signaling Connection (SD-WAN plugin, Secure Peer) or Redirector Mode Connection (SD-WAN plugin)
All modes can be active simultaneously. The mode used for a given packet is determined by the Ethernet and IP headers.		

The forwarding modes are:

- **Inline mode,** in which the appliance transparently accelerates traffic flowing between its two Ethernet ports. In this mode, the appliance appears (to the rest of the network) to be an Ethernet bridge. Inline mode is recommended, because it requires the least configuration.
- **WCCP mode (WAN OP),** which uses the WCCP v. 2.0 protocol to communicate with the router. This mode is easy to configure on most routers. WCCP has two variants: WCCP-GRE and WCCP-L2. WCCP-GRE encapsulates the WCCP

traffic within generic routing encapsulation (GRE) tunnels. WCCP-L2 uses un-encapsulated network Layer 2 (Ethernet) transport.

- **Virtual inline mode,** in which a router sends WAN traffic to the appliance and the appliance returns it to the router. In this mode, the appliance appears to be a router, but it uses no routing tables. It sends the return traffic to the real router. Virtual inline mode is recommended when inline mode and high-speed WCCP operation are not practical.
- **Group mode,** which allows two appliances to operate together to accelerate a pair of widely separated WAN links.
- **High availability mode,** which allows to appliances to operate as an active/standby high availability pair. If the primary appliance fails, the secondary appliance takes over.

Additional traffic types are listed here for completeness:

- **Pass-through traffic** refers to any traffic that the appliance does not attempt to accelerate. It is a traffic category, not a forwarding mode.
- **Direct access,** where the appliance acts as an ordinary server or client. The GUI and CLI are examples of direct access, using the HTTP, HTTPS, SSH, or SFTP protocols. Direct access traffic can also include the NTP and SNMP protocols.
- **Appliance-to-appliance communication**, which can include signaling connections (used in secure peering and by the SD-WAN plugin), VRRP heartbeats (used in high-availability mode), and encrypted GRE tunnels (used by group mode).
- **Deprecated modes.** Proxy mode and redirector mode are legacy forwarding modes that should not be used in new installations.

Customizing the Ethernet ports

Aug 09, 2017

A typical appliance has four Ethernet ports: two accelerated bridged ports, called *accelerated pair A* (apA.1 and apA.2), with a bypass (fail-to-wire) relay, and two unaccelerated motherboard ports, called Primary and Aux1. The bridged ports provide acceleration, while the motherboard ports are sometimes used for secondary purposes. Most installations use only the bridged ports.

Some SD-WAN units have only the motherboard ports. In this case, the two motherboard ports are bridged.

The appliance's user interface can be accessed by a VLAN or non-VLAN network. You can assign a VLAN to any of the appliance's bridged ports or motherboard ports for management purposes.

Figure 1. Ethernet Ports

The ports are named as follows:

Table 1. Ethernet Port Names

Motherboard port 1	Primary (or apA.1 if no bypass card is present)
Motherboard port 2	Auxiliary1 or Aux1 (or apA.2 if no bypass card is present)
Bridge #1	Accelerated Pair A (apA, with ports apA.1 and apA.2)
Bridge #2	Accelerated Pair B (apB, with ports apB.1 and apB.2)

Port Parameters

Aug 09, 2017

Each bridge and motherboard port can be:

- Enabled or disabled
- Assigned an IP address and subnet mask
- Assigned a default gateway
- Assigned to a VLAN
- Set to 1000 Mbps, 100 Mbps, or 10 Mbps
- Set to full duplex, half-duplex, or auto (on SD-WAN 4000 WAN OP/SE/ 5000 WAN OP appliances, some ports can be set to 10 Gbps)

All of these parameters except the speed/duplex setting are set on the Configuration: IP Address page. The speed/duplex settings are set on the Configuration: Interface page.

Notes about parameters:

- Disabled ports do not respond to any traffic.
- The browser-based UI can be enabled or disabled independently on all ports.
- To secure the UI on ports with IP addresses, select HTTPS instead of HTTP on the Configuration: Administrator Interface: Web Access page.
- Inline mode works even if a bridge has no IP address. All other modes require that an IP address be assigned to the port.
- Traffic is not routed between interfaces. For example, a connection on bridge apA does not cross over to the Primary or Aux1 ports, but remains on bridge apA. All routing issues are left to your routers.

Accelerated Bridges (apA and apB)

Aug 09, 2017

Every appliance has at least one pair of Ethernet ports that function as an accelerated bridge, called *apA* (for *accelerated pair A*). A bridge can act in inline mode, functioning as a transparent bridge, as if it were an Ethernet switch. Packets flow in one port and out the other. Bridges can also act in one arm mode, in which packets flow in one port and back out the same port.

An appliance that has a bypass card maintains network continuity if a bridge or appliance malfunctions.

Some units have more than one accelerated pair, and these additional accelerated pairs are named apB, apC, and so on.

If the appliance loses power or fails in some other way, an internal relay closes and the two bridged ports are electrically connected. This connection maintains network continuity but makes the bridge ports inaccessible. Therefore you might want to use one of the motherboard ports for management access.

Caution: Do not enable the Primary port if it is not connected to your network. Otherwise, you cannot access the appliance, as explained in Ethernet Bypass and Link-Down Propagation
Bypass cards are standard on some models and optional on others. Citrix recommends that you purchase appliances with bypass cards for all inline deployments.

The bypass feature is wired as if a cross-over cable connected the two ports, which is the correct behavior in properly wired installations.

Important: Bypass installations must be tested - Improper cabling might work in normal operation but not in bypass mode. The Ethernet ports are tolerant of improper cabling and often silently adjust to it. Bypass mode is hard-wired and has no such adaptability. Test inline installations with the appliance turned off to verify that the cabling is correct for bypass mode.

If the appliance is equipped with two accelerated bridges, they can be used to accelerate two different links. These links can either be fully independent or they can be redundant links connecting to the same site. Redundant links can be either load-balanced or used as a main link and a failover link.

Figure 1. Using dual bridges

When it is time for the appliance to send a packet for a given connection, the packet is sent over the same bridge from which the appliance received the most recent input packet for that connection. Thus, the appliance honors whatever link decisions are made by the router, and automatically tracks the prevailing load-balancing or main-link/failover-link algorithm in real time. For non-load-balanced links, the latter algorithm also ensures that packets always use the correct bridge.

Multiple bridges are supported in virtual inline mode.

Two units with multiple bridges can be used in a high-availability pair. Simply match up the bridges so that all links pass through both appliances.

Motherboard Ports

Aug 09, 2017

Although the Ethernet ports on a bypass card are inaccessible when the bypass relay is closed, the motherboard ports remain active. You can sometimes access a failed appliance through the motherboard ports if the bridged ports are inaccessible.

The Primary Port

If the Primary port is enabled and has an IP address assigned to it, the appliance uses that IP address to identify itself to other acceleration units. This address is used internally for a variety of purposes, and is most visible to users as the Partner Unit field on the Monitoring: Optimization: Connections page. If no motherboard port is enabled, the appliance uses the IP address of Accelerated Pair A.

The Primary port is used for:

- Administration through the web based UI
- A back channel for group mode
- A back channel for high-availability mode

The Aux1 Port

The Aux1 port is identical to the Primary port. If the Aux1 port is enabled and the Primary port is not, the appliance takes its identity from the Aux1 port's IP address. If both are enabled, the Primary port's IP address is the unit's identity

VLAN Support

Aug 09, 2017

A virtual local area network (VLAN) uses part of the Ethernet header to indicate which virtual network a given Ethernet frame belongs to. SD-WAN appliances support VLAN trunking in all forwarding modes (inline, virtual inline, and group mode). Traffic with any combination of VLAN tags is handled and accelerated correctly.

For example, if one traffic stream passing through the accelerated bridge is addressed to 10.0.0.1, VLAN 100, and another is addressed to 10.0.0.1, VLAN 111, the appliance knows that these are two distinct destinations, even though the two VLANs have the same IP address.

You can assign a VLAN to all, some, or none of the appliance's Ethernet ports. If a VLAN is assigned to a port, the management interfaces (GUI and CLI) listen only to traffic on that VLAN. If no VLAN is assigned, the management interfaces listen only to traffic without a VLAN. This selection is made on the Configuration: Appliance Settings: Network Adapters: IP Addresses tab.

Inline Mode

Aug 09, 2017

In inline mode, traffic passes into one of the appliance's Ethernet ports and out of the other. When two sites with inline appliances communicate, every TCP connection passing between them is accelerated. All other traffic is passed through transparently, as if the appliance were not there.

Figure 1. Inline mode, Accelerating All Traffic on a WAN

Note: Any TCP-based traffic passing through both units is accelerated. No address translation, proxying, or per-site setup is required. Inline mode is auto-detecting and auto-configuring.
Configuration is minimized with inline mode, because your WAN router need not be aware of the appliance's existence.

Depending on your configuration, inline mode's link-down propagation can affect management access to the appliance if a link goes down.

Inline mode is most effective when applied to all traffic flowing into and out of a site, but it can be used for only some of the site's traffic.

Ethernet Bypass and Link-Down Propagation

Aug 09, 2017

Note: Link-Down propagation was added to the SD-WAN 2000, 3000, 4000, and 5000 appliances with the 7.2.1 release. Most appliance models include a "fail-to-wire" (Ethernet bypass) feature for inline mode. If power fails, a relay closes and the input and output ports become electrically connected, allowing the Ethernet signal to pass through from one port to the other as if the appliance were not there. In fail-to-wire mode, the appliance looks like a cross-over cable connecting the two ports.

Any failure of the appliance hardware or software also closes the relay. When the appliance is restarted, the bypass relay remains closed until the appliance is fully initialized, maintaining network continuity at all times. This feature is automatic and requires no user configuration.

When the bypass relay is closed, the appliance's bridge ports are inaccessible.

If carrier is lost on one of the bridge ports, the carrier is dropped on the other bridge port to ensure that the link-down condition is propagated to the device on the other side of the appliance. Units that monitor link state (such as routers) are thus notified of conditions on the other side of the bridge.

Link-down propagation has two operating modes:

- If the Primary port is not enabled, the link-down state on one bridge port is mirrored briefly on the other bridge port, and then the port is re-enabled. This allows the appliance to be reached through the still-connected port for management, HA heartbeat, and other tasks.
- If the Primary port is enabled, the appliance assumes (without checking) that the Primary port is used for management, HA heartbeat, and other tasks. The link-down condition on one bridge port is mirrored persistently on the other port, until carrier is restored or the unit is rebooted. This is true even if the Primary port is enabled in the GUI but not connected to a network, so the Primary port should be disabled (the default) when not in use.

Accelerating an Entire Site

Aug 09, 2017

Inline mode, Accelerating All Traffic on a WAN shows a typical configuration for inline mode. For both sites, the appliances are placed between the LAN and the WAN, so all WAN traffic that can be accelerated is accelerated. This is the simplest method for implementing acceleration, and it should be used when practical.

Because all the link traffic is flowing through the appliances, the benefits of fair queuing and flow control prevent the link from being overrun.

In IP networks, the bottleneck gateway determines the queuing behavior for the entire link. By becoming the bottleneck gateway, the appliance gains control of the link and can manage it intelligently. This is done by setting the bandwidth limit slightly lower than the link speed. When this is done, link performance is ideal, with minimal latency and loss even at full link utilization.

Partial-Site Acceleration

Aug 09, 2017

To reserve the appliance's accelerated bandwidth for a particular group of systems, such as remote backup servers, you can install the appliance on a branch network that includes only those systems. This is shown in the following figure.

Figure 1. Inline Mode, Accelerating Selected Systems Only

SD-WAN traffic shaping relies on controlling the entire link, so traffic shaping is not effective with this topology, because the appliance sees only a portion of link traffic. Latency control is up to the bottleneck gateway, and interactive responsiveness can suffer.

Configuring and Troubleshooting Inline Mode

Aug 09, 2017

Inline mode requires only basic configuration, because it is applied automatically to any packets passing through the accelerated bridge. Troubleshooting is described under .

Virtual Inline Mode

Aug 09, 2017

Note: Use virtual inline mode only when both inline mode and WCCP mode are impractical. Do not mix inline and virtual inline modes within the same appliance. However, you can mix virtual inline and WCCP modes within the same appliance. Citrix does not recommend virtual inline mode with routers that do not support health monitoring.

In virtual inline mode, the router uses policy based routing (PBR) rules to redirect incoming and outgoing WAN traffic to the appliance for acceleration, and the appliance forwards the processed packets back to the router. Almost all of the configuration tasks are performed on the router. The only thing to be configured on the appliance is the forwarding method, and the default method is recommended.

Like WCCP, Virtual inline deployment requires no rewiring and no downtime, and it provides a solution for asymmetric routing issues faced in a deployment with two or more WAN links. Unlike WCCP, it contains no built-in status monitoring or health checking, making troubleshooting difficult. WCCP is thus the recommended mode, and virtual inline is recommended only when inline and WCCP modes are both impractical.

The following figure shows a simple network in which all traffic destined for or received from the remote site is redirected to the appliance. In this example, both the local site and remote site use virtual inline mode.

Figure 1. Virtual Inline Example

Following are some configuration details for the network in this example:

- Endpoint systems have their gateways set to the local router (which is not unique to virtual inline mode).
- Each router is configured to redirect both incoming and outgoing WAN traffic to the local appliance.
- Each appliance processes the traffic received from its local router and forwards it back to the router.
- PBR rules configured on the router prevent routing loops by allowing packets to make only one trip to and from the appliance. The packets that the appliance forwards back to the router are sent to their original (local or remote) destination.
- Each appliance has its default gateway set to the address of the local router, as usual (on the **Configuration: Network Adapters** page). The options for forwarding packets back to the router are Return to Ethernet Sender and Send to Gateway.

Configuring Packet Forwarding on the Appliance

Aug 09, 2017

Virtual inline mode offers two packet-forwarding options:

Return to Ethernet Sender (default)—This mode allows multiple routers to share an appliance. The appliance forwards virtual inline output packets back to where they came from, as indicated by the Ethernet address of the incoming packet. If two routers share a single appliance, each gets its own traffic back, but not the traffic from the other router. This mode also works with a single router.

Send to Gateway (not recommended)—In this mode, virtual inline output packets are forwarded to the default gateway for delivery, even if they are destined for hosts on the local subnet. This option is usually less desirable than the Return to Ethernet Sender option, because it adds an easily forgotten element of complexity to the routing structure.

To specify the packet-forwarding option—On the Configuration: Optimization Rules: Tuning page, next to Virtual Inline, select Return to Ethernet Sender or Send to Gateway.

Router Configuration

Aug 09, 2017

The router has three tasks when supporting virtual inline mode:

1. It must forward both incoming and outgoing WAN traffic to the SD-WAN appliance.
2. It must forward SD-WAN traffic to its destination (WAN or LAN).
3. It must monitor the health of the SD-WAN appliance so that the appliance can be bypassed if it fails.

In virtual inline mode, the packet forwarding methods can create routing loops if the routing rules do not distinguish between a packet that has been forwarded by the appliance and one that has not. You can use any method that makes that distinction.

A typical method involves dedicating one of the router's Ethernet ports to the appliance and creating routing rules that are based on the Ethernet port on which packets arrive. Packets that arrive on the interface dedicated to the appliance are never forwarded back to the appliance, but packets arriving on any other interface can be.

The basic routing algorithm is:

- Do not forward packets from the appliance back to the appliance.
- If the packet arrives from the WAN, forward it to the appliance.
- If packet is destined for the WAN, forward to the appliance.
- Do not forward LAN-to-LAN traffic to the appliance.
- Traffic shaping is not effective unless all WAN traffic passes through the appliance.

Note: When considering routing options, keep in mind that returning data, not just outgoing data, must flow through the appliance. For example, placing the appliance on the local subnet and designating it as the default router for local systems does not work in a virtual inline deployment. Outgoing data would flow through the appliance, but incoming data would bypass it. To force data through the appliance without router reconfiguration, use inline mode.

If the appliance fails, data should not be routed to it. By default, Cisco policy based routing does no health monitoring. To enable health monitoring, define a rule to monitor the appliance's availability, and specify the "verify-availability" option for the "set ip next-hop" command. With this configuration, if the appliance is not available, the route is not applied, and the appliance is bypassed.

Important: Citrix recommends virtual inline mode only when used with health monitoring. Many routers that support policy-based routing do not support health-checking. The health-monitoring feature is relatively new. It became available in Cisco IOS release 12.3(4)T.

Following is an example of a rule for monitoring the availability of the appliance:

```
!- Use a ping (ICMP echo) to see if appliance is connected track  123 rtr 1 reachabilit y ! rtr 1  type echo protocol Iplcmpecho 192.168.1.200 schedule 1 life forever start-time now
```

This rule pings the appliance at 192.168.1.200 periodically. You can test against 123 to see if the unit is up.

Routing Examples

Aug 09, 2017

The following examples illustrate configuring Cisco routers for the local and remote sites shown in Virtual inline example. To illustrate health monitoring, the configuration for the local site includes health monitoring, but the configuration for the remote site does not.

Note: The configuration for the local site assumes that a ping monitor has already been configured.
The examples conform to the Cisco IOS CLI. They might not be applicable to routers from other vendors.

Local Site, Health-Checking Enabled

```
!
! For health-checking to work, do not forget to start
! the monitoring process.
!
! Original configuration is in normal type.
! appliance-specific configuration is in bold.
!
ip cef
!
interface FastEthernet0/0
ip address 10.10.10.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 172.68.1.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.1.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.1.1
!
ip access-list extended client_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
!- Now set the appliance as the next hop, if it's up.
set ip next-hop verify-availability 192.168.1.200 20 track 123
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop verify-availability 192.168.1.200 10 track 123
```

Remote Site (No Health Checking)

```
! This example does not use health-checking.
! Remember, health-checking is always recommended,
! so this is a configuration of last resort.
!
!
ip cef
!
interface FastEthernet0/0
ip address 20.20.20.5 255.255.255.0
ip policy route-map client_side_map
!
interface FastEthernet0/1
ip address 171.68.2.5 255.255.255.0
ip policy route-map wan_side_map
!
interface FastEthernet1/0
ip address 192.168.2.5 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 171.68.2.1
!
ip access-list extended client_side
permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
ip access-list extended wan_side
permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
route-map wan_side_map permit 20
match ip address wan_side
set ip next-hop 192.168.2.200
!
route-map client_side_map permit 10
match ip address client_side
set ip next-hop 192.168.2.200
!_
```

Each of the above examples applies an access list to a route map and attaches the route map to an interface. The access lists identify all traffic originating at one accelerated site and terminating at the other (A source IP of 10.10.10.0/24 and destination of 20.20.20.0/24 or vice versa). See your router's documentation for the details of access lists and route-maps.

This configuration redirects all matching IP traffic to the appliances. If you want to redirect only TCP traffic, you can change the access-list configuration as follows (only the remote side's configuration is shown here):

```
!
 ip access-list extended client_side
 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
 ip access-list extended wan_side
 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
!
```

Note that, for access lists, ordinary masks are not used. Wildcard masks are used instead. Note that when reading a

wildcard mask in binary, "1" is considered a "don't care" bit.

Virtual Inline for Multiple-WAN Environments

Aug 09, 2017

Enterprises with multiple WAN links often have asymmetric routing policies, which seem to require that an inline appliance be in two places at once. Virtual inline mode solves the asymmetric routing problem by using the router configuration to send all WAN traffic through the appliance, regardless of the WAN link used. The below figure shows a simple multiple-WAN link deployment example.

The two local-side routers redirect traffic to the local appliance. The FE 0/0 ports for both routers are in the same broadcast domain as the appliance. The local appliance must use the default virtual inline configuration (Return to Ethernet Sender).

Figure 1. Virtual Inline Mode With Two WAN Routers

Virtual Inline Mode and High-Availability

Aug 09, 2017

Virtual Inline mode can be used in a high availability (HA) configuration. The below figure shows a simple HA deployment. In virtual inline mode, a pair of appliances acts as one virtual appliance. Router configuration is the same for an HA pair as with a single appliance, except that the Virtual IP address of the HA pair, not the IP address of an individual appliance, is used in the router configuration tables. In this example, the local appliances must use default virtual inline configuration (Return to Ethernet Sender).

Figure 1. High-availability Example

Monitoring and Troubleshooting

Aug 09, 2017

In virtual inline mode, unlike WCCP mode, the appliance provides no virtual inline-specific monitoring. To troubleshoot a virtual inline deployment, log into the appliance and use the Dashboard page to verify that traffic is flowing into and out of the appliance. Traffic forwarding failures are typically caused by errors in router configuration.

If the Monitoring: Usage or Monitoring: Connections pages show that traffic is being forwarded but no acceleration is taking place (assuming that an appliance is already installed on the other end of the WAN link), check to make sure that both incoming WAN traffic and outgoing WAN traffic are being forwarded to the appliance. If only one direction is forwarded, acceleration cannot take place.

To test health-checking, power down the appliance. The router should stop forwarding traffic after the health-checking algorithm times out.

High-Availability Mode

Aug 09, 2017

Two identical appliances on the same subnet can be combined as a *high-availability pair.* The appliances each monitor the other's status by using the standard *Virtual Router Redundancy Protocol (VRRP)* heartbeat mechanism. The pair has a common virtual IP address for management, in addition to each appliance's management IP address. If the primary appliance fails, the secondary appliance takes over. Failover takes approximately five seconds.

High availability mode is a standard feature.

How High-Availability Mode Works

Aug 09, 2017

In a high availability (HA) pair, one appliance is primary, and the other is secondary. The primary monitors its own and the secondary's status. If it detects a problem, traffic processing fails over to the secondary appliance. Existing TCP connections are terminated. To ensure successful failover, the two appliances keep their configurations synchronized. In a WCCP mode high availability configuration, the appliance that is processing traffic maintains communication with the upstream router.

Status monitoring—When high availability is enabled, the primary appliance uses the VRRP protocol to send a heartbeat signal to the secondary appliance once per second. In addition, the primary appliance monitors the carrier status of its Ethernet ports. The loss of carrier on a previously active port implies a loss of connectivity.

Failover If the heartbeat signal of the primary appliance should fail, or if the primary appliance loses carrier for five seconds on any previously active Ethernet port, the secondary appliance takes over, becoming the primary. When the failed appliance restarts, it becomes the secondary. The new primary announces itself on the network with an ARP broadcast. MAC spoofing is not used. Ethernet bridging is disabled on the secondary appliance, leaving the primary appliance as the only path for inline traffic. Fail-to-wire is inhibited on both appliances to prevent loops.

Caution: The Ethernet bypass function is disabled in HA mode. If both appliances in an inline HA pair lose power, connectivity is lost. If WAN connectivity is needed during power outages, at least one appliance must be attached to a backup power source.

Note: The secondary appliance in the HA pair has one of its bridge ports, port apA.1, disabled to prevent forwarding loops. If the appliance has dual bridges, apB.1 is also disabled. In a one-arm installation, use port apA.2. Otherwise, the secondary appliance becomes inaccessible when HA is enabled.

Primary/secondary assignment—If both appliances are restarted, the first one to fully initialize itself becomes the primary. That is, the appliances have no assigned roles, and the first one to become available takes over as the primary. The appliance with the highest IP address on the interface used for the VRRP heartbeat is used as a tie-breaker if both become available at the same time.

Connection termination during failover—Both accelerated and unaccelerated TCP connections are terminated as a side effect of failover. Non-TCP sessions are not affected, except for the delay caused by the brief period (several seconds) between the failure of the primary appliance and the failover to the secondary appliance. Users experience the closing of open connections, but they can open new connections.

Configuration synchronization—The two appliances synchronize their settings to ensure that the secondary is ready to take over for the primary. If the configuration of the pair is changed through the browser based interface, the primary appliance updates the secondary appliance immediately.

HA cannot be enabled unless both appliances are running the same software release.

HA in WCCP mode—When WCCP is used with an HA pair, the primary appliance establishes communication with the router. The appliance uses its management IP address on apA or apB, not its virtual IP address, to communicate with the router. Upon failover, the new primary appliance establishes WCCP communication with the router.

Cabling Requirements

Aug 09, 2017

The two appliances in the high availability pair are installed onto the same subnet in either a parallel arrangement or a one-arm arrangement, both of which are shown in the following figure. In a one-arm arrangement, use the apA.2 port (and, optionally, the apB.2 port), not the apA.1 port. Some models require a separate management LAN, whether deployed in inline or one-armed mode. This is depicted only in the middle diagram.

Figure 1. Cabling for High-Availability Pairs

Do not break the above topology with additional switches. Random switch arrangements are not supported. Each of the switches must be either a single, monolithic switch, a single logical switch, or part of the same chassis.

If the spanning-tree protocol (STP) is enabled on the router or switch ports attached to the appliances, failover will work, but the failover time may increase to roughly thirty seconds. Without STP, failover time is roughly five seconds. Thus, to achieve the briefest possible failover interval, disable STP on the ports connecting to the appliances.

Figure 2. Ethernet Port Locations (Older Models)

Other Requirements

Aug 09, 2017

Both appliances in an HA pair must meet the following criteria:

- Have identical hardware, as shown by on the System Hardware entry on the Dashboard page.
- Run exactly the same software release.
- Be equipped with Ethernet bypass cards. To determine what is installed in your appliances, see the Dashboard page.

Appliances that do not support HA display a warning on the Configuration: High Availability page.

Management Access to the High-Availability Pair

Aug 09, 2017

When configuring a high-availability (HA) pair, you assign the pair a virtual IP (VIP) address, which enables you to manage the two appliances as if they were a single unit. After you enable high-availability mode, managing the secondary appliance through its IP address is mostly disabled, with most parameters grayed out. A warning message displays the reason on every page. Use the HA VIP for all management tasks. You can, however, disable the secondary appliance's HA state from its management UI.

Configuring the High-Availability Pair

Aug 09, 2017

You can configure two newly installed appliances as a high-availability pair, or you can create an HA pair by adding a second appliance to an existing installation.

Prerequisites: Physical installation and basic configuration procedures

To configure high availability

1. Make sure that no more than one appliance is connected to the traffic networks (on the accelerated bridges). If both are connected, disconnect one bridge cable from the active bridges on the second appliance. This will prevent forwarding loops.
2. On the Features page of the first appliance, disable Traffic Processing. This disables acceleration until the HA pair is configured.
3. Repeat for the second appliance.
4. On the first appliance, go to the Configuration: Advanced Deployments: High Availability tab, show below.
5. Select the Enabled Check box.
6. Click the Configure HA Virtual IP Address link and assign a virtual IP address to the apA interface. This address will be used later to control both appliances as a unit.
7. Return to the High Availability page and, in the VRRP VRID field, assign a VRRP ID to the pair. Although the value defaults to zero, the valid range of VRRP ID numbers is 1 through 255. Within this range, you can specify any value that does not belong to another VRRP device on your network.
8. In the Partner SSL Common Name field, type the other appliance's SSL Common Name, which is displayed on that appliance's Configuration: Advanced Deployments: High Availability tab, in the Partner SSL Common Name field. The SSL credentials used here are factory-installed.
9. Click Update.
10. Repeat steps 3-8 on the second appliance. If you are managing the appliance via an accelerated bridge (such as apA), you may have to reconnect the Ethernet cable that you removed in step 1 to connect to the second appliance. If so, plug this cable in and disconnect the corresponding cable on the first appliance.
11. With your browser, navigate to the virtual IP address of the HA pair. Enable Traffic Processing on the Features page. Any further configuration will be performed from this virtual address.
12. Plug in the cable that was left disconnected.
13. On each appliance, the Configuration: Advanced Deployments: High Availability page should now show that high availability is active and that one appliance is the primary and the other is the secondary. If this is not the case, a warning banner appears at the top of the screen, indicating the nature of the problem.

Figure 1. High-availability configuration page

Updating Software on a High-Availability Pair

Aug 09, 2017

Updating the SD-WAN software on an HA pair causes a failover at one point during the update.

Note: Clicking the Update button terminates all open TCP connections.

To update the software on an HA pair

1. Log on to both appliances.
2. On the secondary appliance, update the software and reboot. After the reboot, the appliance is still the secondary. Verify that the installation succeeded. The primary appliance should show that the secondary appliance exists but that automatic parameter synchronization is not working, due to a version mismatch.
3. On the primary appliance, update the software, and then reboot. The reboot causes a failover, and the secondary appliance becomes the primary. When the reboot is completed, HA should become fully established, because both appliances are running the same software.

Saving/Restoring Parameters of an HA Pair

Aug 09, 2017

The System Maintenance: Backup/Restore function can be used to save and restore parameters of an HA pair as follows:

To back up the parameters

Use the backup feature as usual. That is, log on to the GUI through the HA VIP address (as is normal when managing the HA pair) and, on the System Management: Backup/Restore page, click Download Settings.

To restore the parameters

1. Disable HA on both appliances by clearing the Enabled check box on the Configuration: Advanced Deployments: High Availability (HA) tab.
2. Unplug a network cable from the bridge of one appliance. (Call it “Appliance A.”)
3. Unplug the power cord from Appliance A.
4. Restore the parameters on the other appliance (Appliance B), by uploading a previously saved set of parameters on the System Maintenance: Backup/Restore page and clicking Restore Settings. (Completing this operation requires a restart, which reenables HA).
5. Wait for Appliance B to restart. It becomes the primary.
6. Restart Appliance A.
7. Log on to Appliance A’s GUI and reenable HA on the Configuration: Advanced Deployments: High Availability (HA) tab. The appliance get its parameters from the primary.
8. Plug in the network cable removed in step 2.

Both appliances are now restored and synchronized.

Troubleshooting High Availability Pairs

Aug 09, 2017

If the appliances report any failure to enter high-availability mode, the error message will also note the cause. Some issues that can interfere with high-availability mode are:

- The other appliance is not running.
- The HA parameters on the two appliances are not identical.
- The two appliances are not running the same software release.
- The two appliances do not have the same model number.
- Incorrect or incomplete cabling between the appliances does not allow the HA heartbeat to pass between them.
- The HA/Group Mode SSL Certificates on one or both appliances are damaged or missing.

NetScaler SD-WAN VPX

Aug 09, 2017

Citrix NetScaler SD-WAN Standard Edition or WANOP VPX is a virtual Citrix NetScaler SD-WAN appliance that can be hosted on Citrix XenServer, VMware ESX or ESXi, Microsoft Hyper-V, and Amazon AWS- virtualization platforms. A SD-WAN VPX appliance supports most of the features of a physical Standard Edition or WANOP appliances.

Because SD-WAN SE/WAN OP Edition VPX is a virtual machine, you can deploy on your choice of hardware, exactly where you need it, and in combination it with other virtual machines -- servers, VPN units, or other appliances -- to create a unit that precisely suits your needs.

SD-WAN WANOP VPX software is available as:

- A Xen virtual machine running under XenServer 5.5 and later.
- A VMware vSphere virtual machine running under ESX/ESXi 4.1-6.0.
- A Hyper-V virtual machine under 64-bit Windows 2008 R2 SP1 - 2012.
- An Amazon EC2 instance.

Note: XenServer and VMware vSphere support VLAN trunking, but Hyper-V does not.
SD-WAN Standard Edition VPX software is available as:

- A Xen virtual machine running under XenServer 6.5 SP1.
- A VMware vSphere virtual machine running under ESX/ESXi 5.5 and 6.0.

When a newly installed SD-WAN SE/WANOP VPX virtual machine is up and running, you configure as you would configure a physical SD-WAN SE/WANOP appliance, using the same configuration screens.

Differences between WANOP VPX and Physical SD-WAN WANOP Appliances

A SD-WAN WANOP VPX virtual appliance is similar to a SD-WAN Repeater 8500 series appliance, including support for the SD-WAN Plug-in and links of up to 45 mbps. Following are the key differences:

- Except for Amazon EC2 instances, licensing via remote license servers is mandatory for retail licenses. Local licensing is available for non-retail licenses, such as evaluation and VPX Express licenses. For Amazon EC2 instances, you can use either Citrix licensing or select a product with built-in licensing for the bandwidth limit you desire (2, 10, 20, or 45 Mbps).
- SD-WAN VPX obtains its SD-WAN Plug-in licenses from the remote license server (except for SD-WAN VPX for Amazon AWS, which does not support Plug-ins). Plug-ins connecting to multiple virtual appliances consume only a single Plug-in license, not one license per appliance, provided that all virtual appliances use the same license server.
- The SD-WAN LCD front-panel display is not supported.
- The RS-232 serial command interface is not supported.
- Multiple accelerated bridges are not supported.
- Ethernet bypass cards are not supported.
- Group mode is not supported.
- SD-WAN High-availability mode is not supported. (XenServer HA and vSphere HA are supported.)
- Three ports are supported (apA.1, apA.2, and Primary), except for Amazon AWS instances, which support only a single port.

SD-WAN VPX Usage Scenarios

Oct 24, 2017

You can deploy VPX to accelerate the traffic to or from a branch office, to and from a particular server, or in the cloud. In the data center, you can create a flexible and powerful configuration by assigning a separate VPX instance to each server. Or, at any location, you can assign multiple VPX instances to one server, for different types or levels of acceleration services within the same server.

For employees connecting through VPNs, VPX can accelerate their connections.

As with a physical appliance, inline mode is the most common type of configuration, but WCCP and virtual inline modes can provide an effective deployment.

VPX Usage Scenarios

You can deploy VPX to accelerate the traffic to or from a branch office, or to and from a particular server. In the data center, you can create a flexible and powerful configuration by assigning a separate VPX instance to each server. Or, at any location, you can assign multiple VPX instances to one server, for different types or levels of acceleration services within the same server.

For employees connecting through VPNs, VPX can accelerate their connections.

As with a physical appliance, inline mode is the most common type of configuration, but WCCP mode can provide an effective failover mechanism.

Branch-office accelerator

A VPX image can be installed on the server of your choice and deployed just like a SD-WAN/SD-WAN appliance. VPX has all the functionality of a SD-WAN/SD-WAN appliance, and in addition has advantages provided by virtualization. Group mode and high-availability modes are not supported.

Figure 1. VPX use case #1: Branch-office accelerator

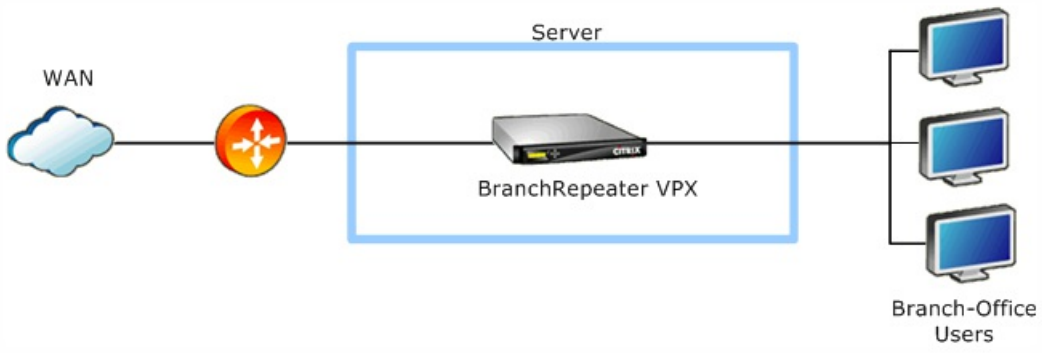

Accelerated branch-office server

If you add a virtual server to the simple branch-office accelerator configuration, you have an accelerated branch-office server, as shown in the figure below. If you assign the virtual networks within the appliance hosting the virtual machines so that the path to the WAN passes through the virtual SD-WAN/SD-WAN, all WAN traffic is accelerated automatically. For example, all web traffic, backups, remote applications, database queries, and operations that require network-file-system

access are accelerated.

The virtual environment allows you to add the desired functionality to the server unit, including the operating system and features of your choice. This configuration accelerates all the WAN traffic from every system in the branch office. You can even deploy multiple virtual servers on the same machine, consolidating your branch-office rack down to a single unit running multiple virtual machines.

Figure 2. VPX use case #2: Accelerated branch-office server

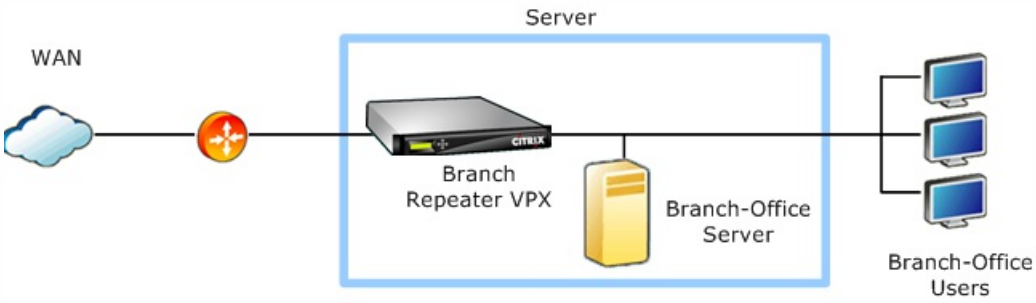

Accelerated datacenter servers

Installing VPX VMs on every server in the data center creates a solution that scales perfectly as you add server capacity, while minimizing the number of servers by adding acceleration to the servers themselves. Once you have more than a few accelerated servers, the aggregate acceleration provided by multipleVPX VMs exceeds anything that can be provided with a single appliance.

VPX accelerates all types of network applications, including XenApp, XenDesktop, Citrix Merchandising Server, network file systems, databases, web servers, and more.

Figure 3. VPX use case #3: Accelerated Datacenter Servers

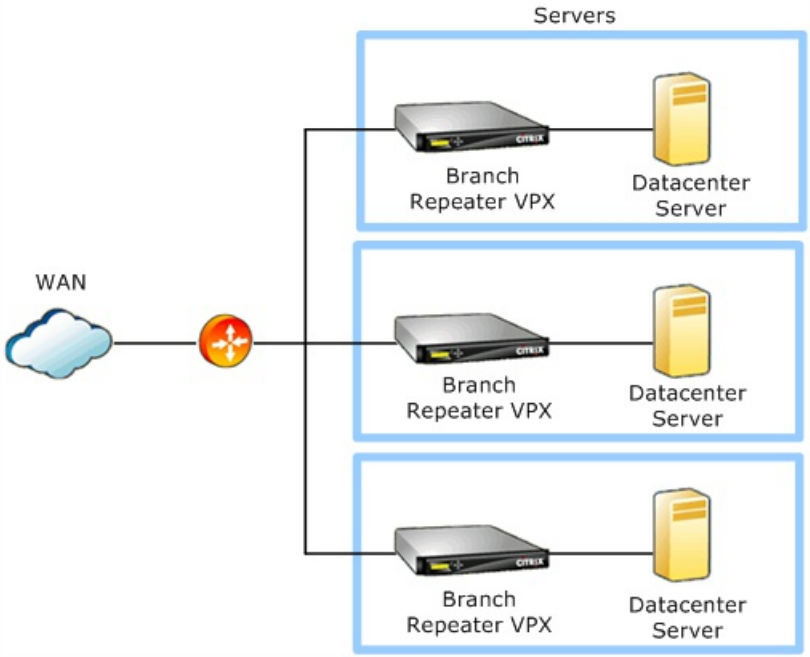

VPN accelerator

By installing the VPN of your choice with VPX, you have an accelerated VPN.

Note: Unlike other configurations, the VPN virtual machine is on the WAN side and the VPX virtual appliance is on the LAN side, because the VPN traffic must be decrypted for compression and application acceleration.

Figure 4. VPX use case #4: VPN accelerator

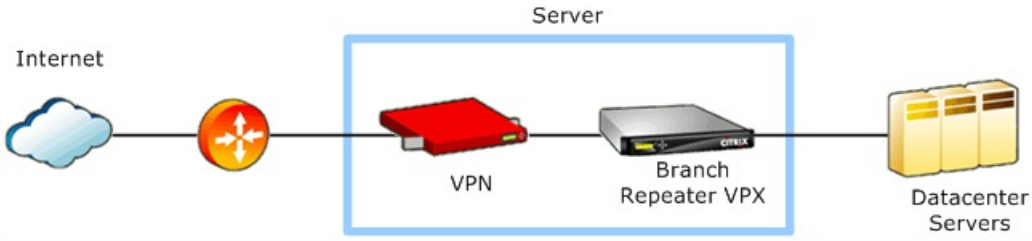

Multiple VPX Instances on the Same Server

By putting multiple VPX VMs on the same server, you can create different types or levels of acceleration services within the same unit. One VPX instance might be dedicated to a critical application, or each instance dedicated to an individual remote site or customer. Use VLAN switches to direct traffic to the appropriate VPX instance.

Figure 5. VPX use case #5: Multiple Instances for Dedicated Acceleration Resources

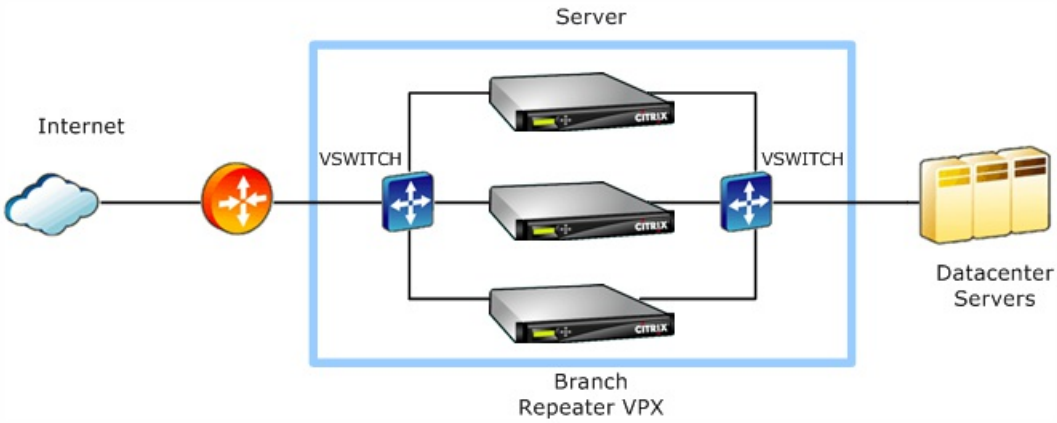

WCCP and virtual inline deployment

WCCP and virtual inline modes are suitable for one-arm deployments, which use only one port. The Amazon AWS version of VPX uses only a single port, and is thus always deployed in a one-armed mode.

Figure 6. VPX use case #6: WCCP or virtual inline deployment

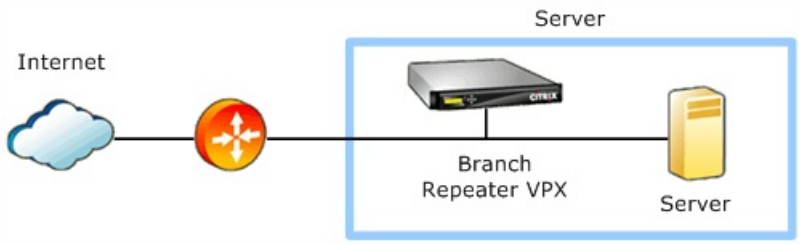

In cases where an Ethernet bypass card would be desirable, using WCCP instead of inline mode provides effective fault-tolerance, because WCCP has built-in health-checking. Instead of forwarding traffic through an unresponsive WCCP device,

the routers send the traffic directly to the end point.

Branch-office acceleratorSD-WAN/SD-WAN VPX can be installed on the server of your choice and deployed just like any other SD-WAN/SD-WAN appliance. VPX has the same functionality as the SD-WAN/SD-WAN appliance along with the additional features provided by virtualization. The group mode and high-availability mode are not supported.

SD-WAN/SD-WAN VPX Features

VPX supports Citrix Command Center release 4.0 or later. SD-WAN also supports SD-WAN/SD-WAN VPX Express licenses, which support a maximum accelerated sending rate of 512 kbps, 10 accelerated connections, and 5 SD-WAN/SD-WAN Plug-ins.

- VPX for XenServer special features include:
  - XenServer Essentials Support
  - XenMotion Live Migration
  - XenServer High Availability
  - Workload Balancing
  - Performance Monitoring and Alerts
- VPX for VMware vSphere special features include:
  - VMware vCenter Server (remote management).
  - VMware vSphere HA (high availability).
  - VMware vSphere vMotion (migrate SD-WAN VPX to a different server with identical processors).
  - VMware Guest Customization (replicate VPX with different per-instance parameters).

System Requirements and Provisioning

Mar 15, 2018

SD-WAN VPX runs on XenServer 5.5 or later, VMware vSphere ESX/ESXi 4.1 or later, Hyper-V under 64-bit Windows Server 2008 R2 SP1, and Amazon AWS. SD-WAN VPX supports four configurations, from 2 GB to 8 GB of RAM and 100 GB to 500 GB of disk space. The intermediate, 4 GB RAM/250 GB disk configuration is similar to the Repeater 8500 series appliance.

Supported Configurations

The following tables list all supported SD-WAN VM configurations. (Amazon AWS configurations are preselected and are different.)

Type	vCPUs	RAM	Disk	Maximum WAN Speed	Maximum Accelerated Connections	Maximum SD-WAN/SD-WAN Plug-ins
2 GB production config.	2	2 GB	100 GB	2 mbps	1,000	50
4 GB production config.	2	4 GB	250 GB	10 mbps	10,000	250
4 GB production config.*	2	4 GB	250 GB	45 mbps	15,000	400
8 GB production config.	4	8 GB	500 GB	45 mbps	25,000	500
* With 45mbps license						

Other configurations (not for production networks)

Type	vCPUs	RAM	Disk	Maximum WAN Speed	Maximum Accelerated Connections	Maximum SD-WAN/SD-WAN Plug-ins
VPX Express	2	1 GB	60 GB	512 kbps	10	5
Min. evaluation config.	2	1 GB	60 GB	2 mbps	1,000	5

Minimum Resource Requirements

An SD-WAN VPX virtual machine has the following minimum hardware requirements for a production environment

- 2 GB RAM
- 100 GB disk (local disks provide the best performance)
- 2 virtual NICs (Ethernet ports), except for Amazon AWS, which requires only one virtual NIC
- 2 virtual CPUs
- A modern CPU (Intel Nehalem or newer or AMD Family 10 h or newer, both of which were introduced in 2008). Older CPUs can run at reduced performance due to the use of emulated x86 TSC (timestamp counter) functionality. When

clock states higher than C1 are not used and SpeedStep/PowerNow modes are disabled in the BIOS of older processors, TSC emulation will not be used and the system runs at normal speed.

The server hosting VPX must have RAM, CPU, and disk resources greater than those required by the VPX VM. (VPX does not support VMware hardware over-commit.) Obviously, the server must have enough resources to run the hypervisor in addition to the virtual appliance. However, having as many physical Ethernet ports as virtual ones is not mandatory if one of a VPX VM's Ethernet ports is connected to another virtual machine on the same server. Possible Ethernet options include:

- Mapping the VPX VM's two virtual ports to two physical ports, rendering its operation equivalent to that of a stand-alone SD-WAN.
- Mapping one of VPX VM's virtual ports to a physical port, and the other to a virtual network containing one or more virtual machines on the same server, thus creating an accelerated server.
- Mapping each of VPX VM's virtual ports to a virtual network, thus chaining the VPX VM between two sets of VMs on the same server.

The following figure shows a VPX VM in a one-arm deployment for traffic that ends on another virtual machine on the same server. Only one physical port is required in this case, but both virtual ports are used.

Figure 1. Ethernet (Network) Port Assignments, One-Arm Operation

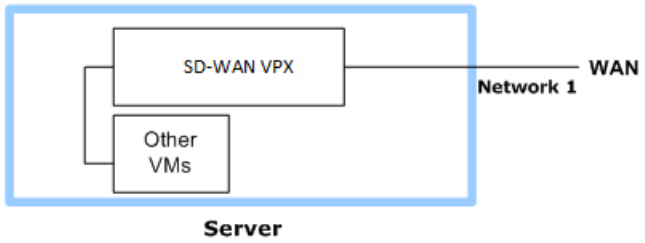

Maximum Usable Resources

Following are the maximum amount of resources that a single VPX virtual machine can use effectively

- 4 virtual CPUs
- 8 GB RAM
- 500 GB disk
- 4 virtual NICs

Server resources not allocated to VPX VMs are available to other VMs on the same server, but be careful to avoid overcommitting resources.

Disk and RAM

While the amounts of RAM and disk space are increased, the additional resources are allocated primarily to the compression subsystem. Increased memory also allows more connections and acceleration partners to be supported.

The SD-WAN compression system makes heavy demands on the disk subsystem. In general, local disk storage outperforms network disk storage and reduces resource contention on both the LAN and the network disk.

The relationship between disk or memory resources and link speed is indirect. Memory and disk sizes have no effect on the speed at which packets are sent more than the link (bps). Providing more memory and disk space improves compression performance by increasing the amount of compression history that can be used for pattern matching.

Virtual NICs

Except for Amazon AWS, two virtual network interfaces are required. They are bridged and used for both acceleration and the browser based user interface. These interfaces must be attached to different virtual networks. For one-arm operation, the second interface can be a stub, attached only to a VPX VM.

A third virtual network interface provides an independent interface to the VPX VM, which is the equivalent to the Primary port on a physical appliance. It can be used for the browser based interface, but not for acceleration.

Other Virtual Machines

- Server resources beyond those allocated to VPX are available for other virtual machines on the same server.
- Resource usage by other VMs affects VPX performance, and conversely. Acceleration makes intensive use of CPU, memory, disk, and network.

Virtual network routing can be used to connect other VMs on the server to VPX VMs, but the simplest method of connecting such VMs is to attach them to the server's LAN-side Ethernet port. WAN-bound packets then pass through the VPX VM's bridge and are accelerated automatically, when they originate inside or outside the server hosting VPX.

Figure 2. An Inline Deployment that Accelerates External Traffic and Traffic from Local VMs

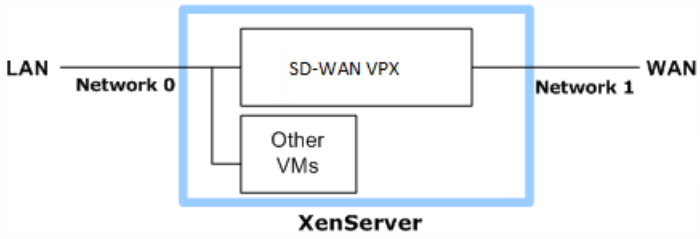

Installing SD-WAN Virtual Appliances on XenServer

Oct 24, 2017

To install Citrix NetScaler SD-WAN virtual appliances on Citrix XenServer, you must first install XenServer on a machine with adequate system resources. To perform the SD-WAN VPX installation, you use Citrix XenCenter, which must be installed on a remote machine that can connect to the XenServer host through the network.

Before you begin installing a virtual appliance, do the following:

- Install a supported version of XenServer® on hardware that meets the minimum requirements. See the SD-WAN release notes for the supported versions of XenServer.
- Install XenCenter® on a management workstation that meets the minimum system requirements.
- Obtain VPX license files.

With the prerequisites met, you are ready to import the virtual appliances and configure them.

To import a SD-WAN virtual appliance to XenServer by using XenCenter

1. Start XenCenter on your workstation.
2. On the **Server** menu, click **Add**.
3. In the **Add New Server** dialog box, in the Server text box, enter the IP address or DNS name of the XenServer server that you want to connect to.
4. In the **User Name** and **Password** text boxes, type the administrator credentials, and then click **Add**. The XenServer name appears in the navigation pane with a green circle, which indicates that the XenServer is connected.

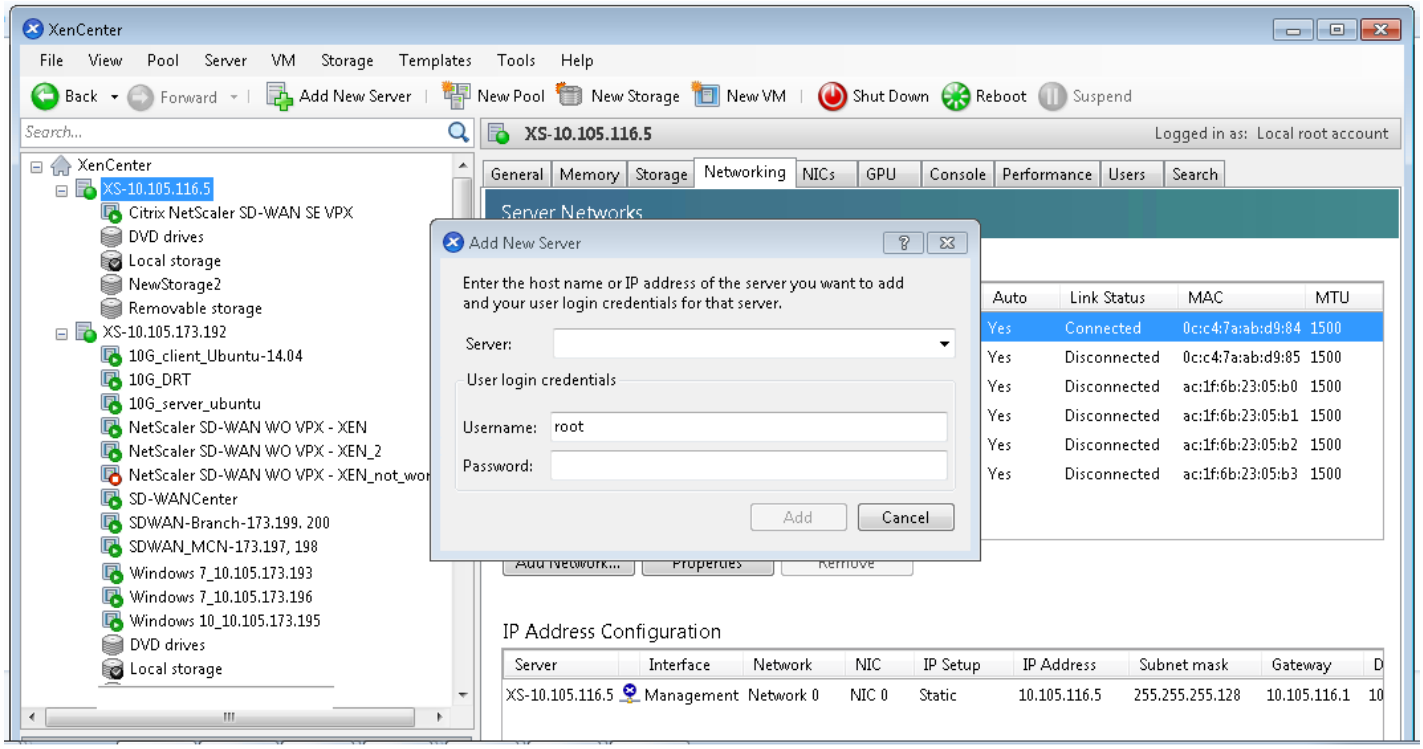

5. In the navigation pane, click the name of the XenServer server on which you want to install SD-WAN VPX-SE.
6. Right-click the XenServer to import the VM and click **Import**.
7. In the Import dialog box, in Import file, browse to the location at which you saved the SD-WAN VPX-SE .xva image

file. Make sure that the Exported VM option is selected, and then click **Next**.

8. Select the XenServer server on which you want to install the virtual appliance, and then click **Next**.

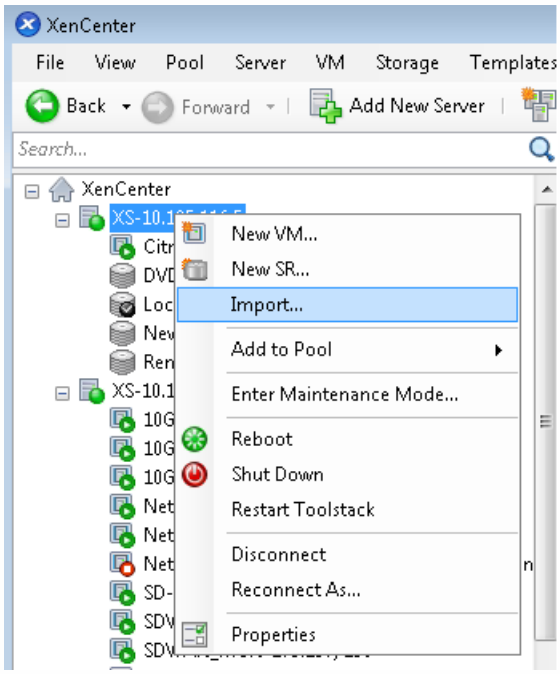

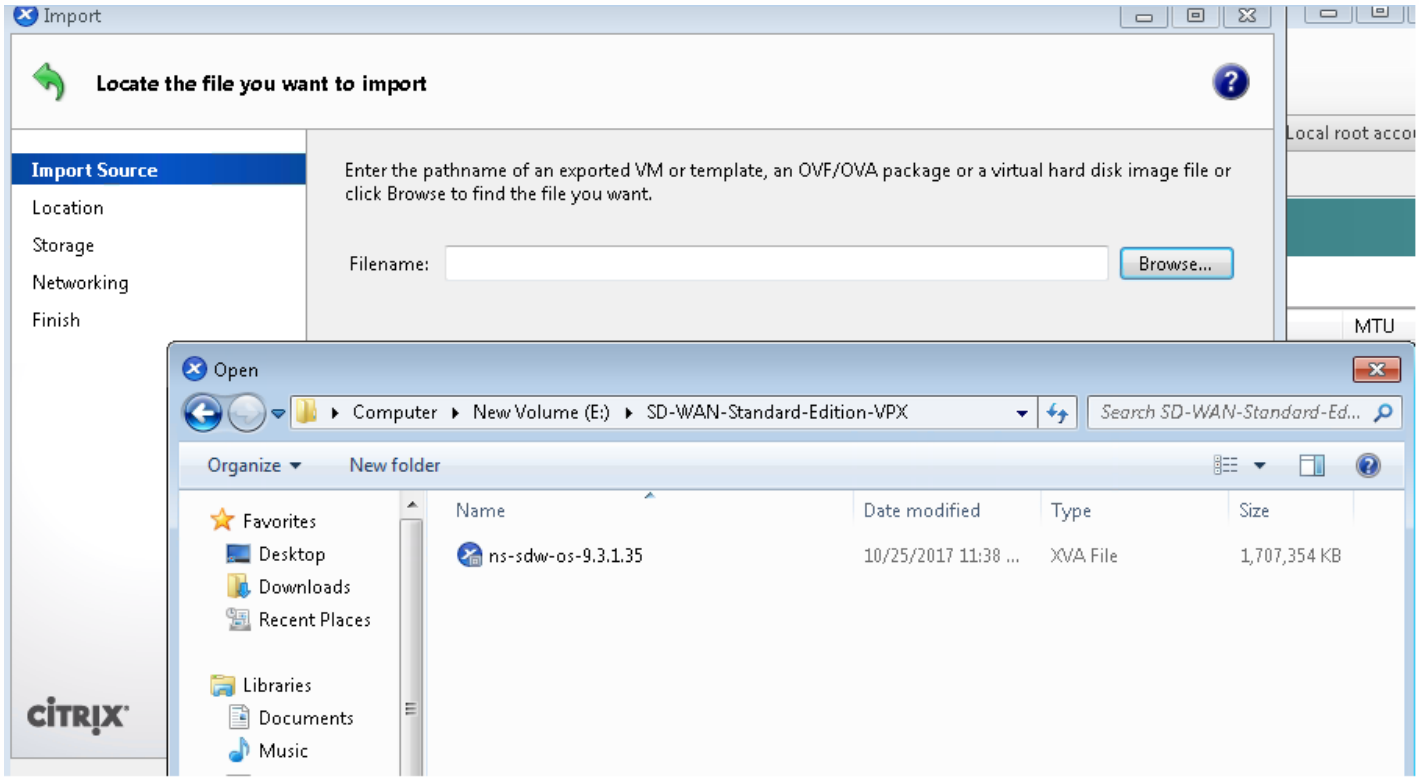

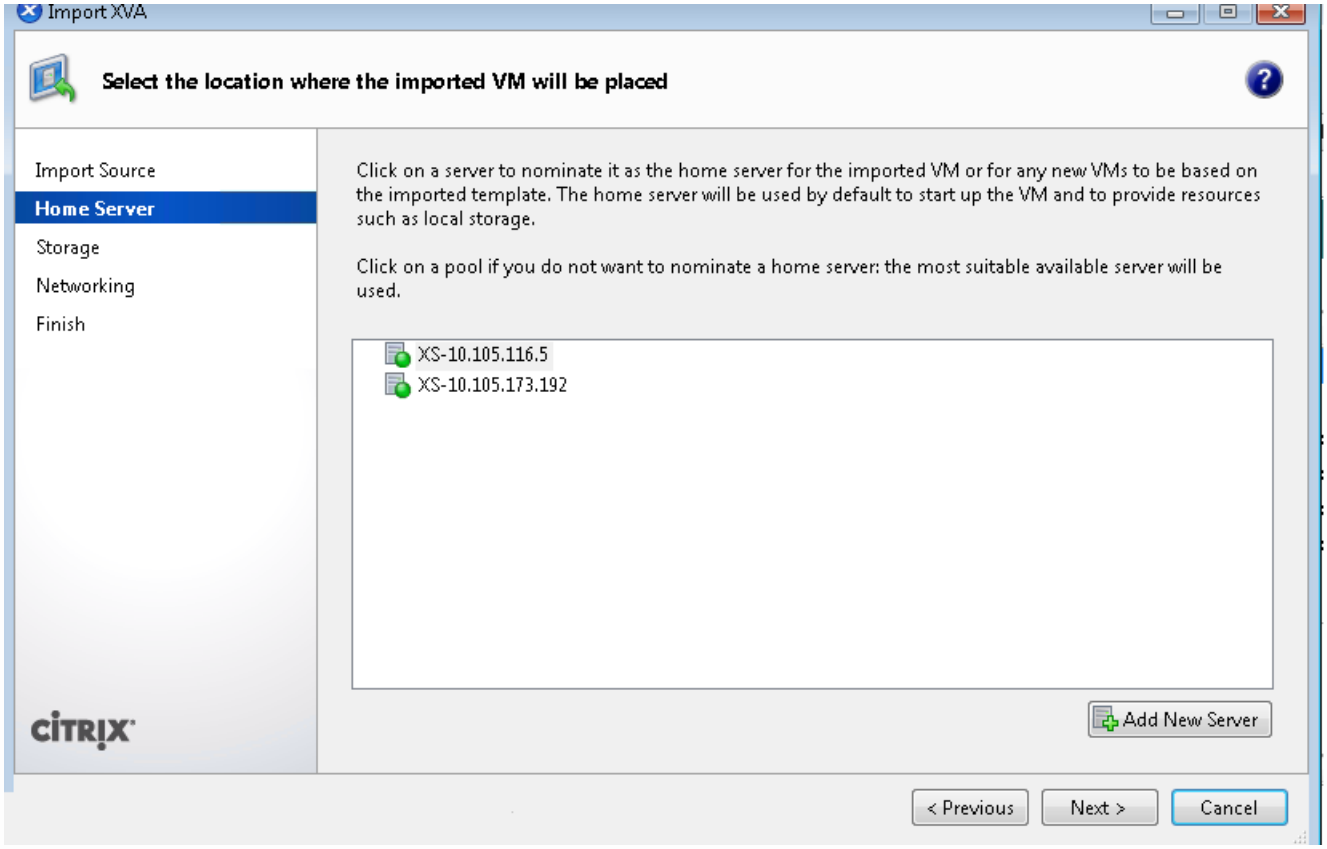

9. Select the local storage repository in which to store the virtual appliance, and then click **Import** to begin the import process.

10. **Add**, **modify,** or **delete** virtual network interfaces as required. Attach virtual network interfaces, interface 0 and interface 1 to the two different virtual adapters (called Networks on this screen). These two interfaces are used as the accelerated bridge of the virtual appliance. If virtual network interface interface 2 exists, it can be assigned as well, and used as a management interface (equivalent to the Primary port). When finished, click **Next**.

11. Click the Start VM after Import check box.

12. Click **Finish** to complete the import process. The newly created virtual machine appears under the server list in the XenCenter interface.

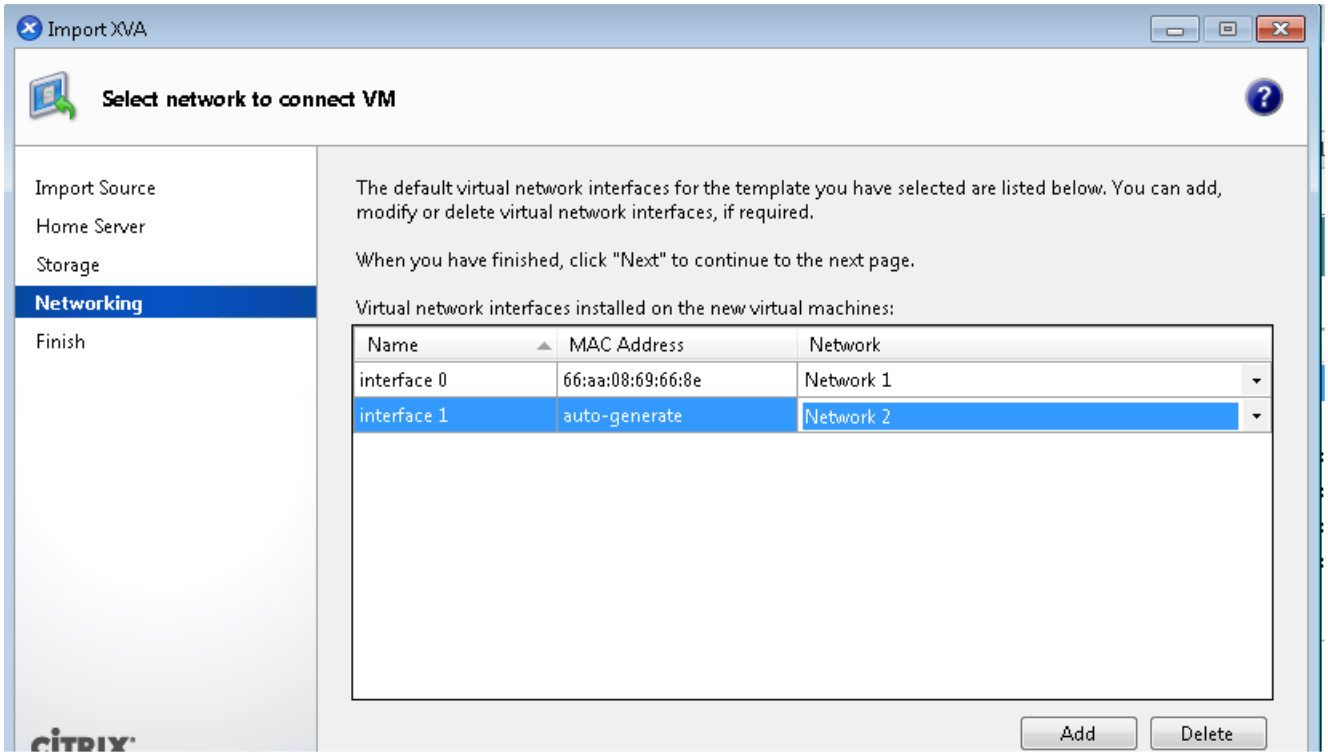

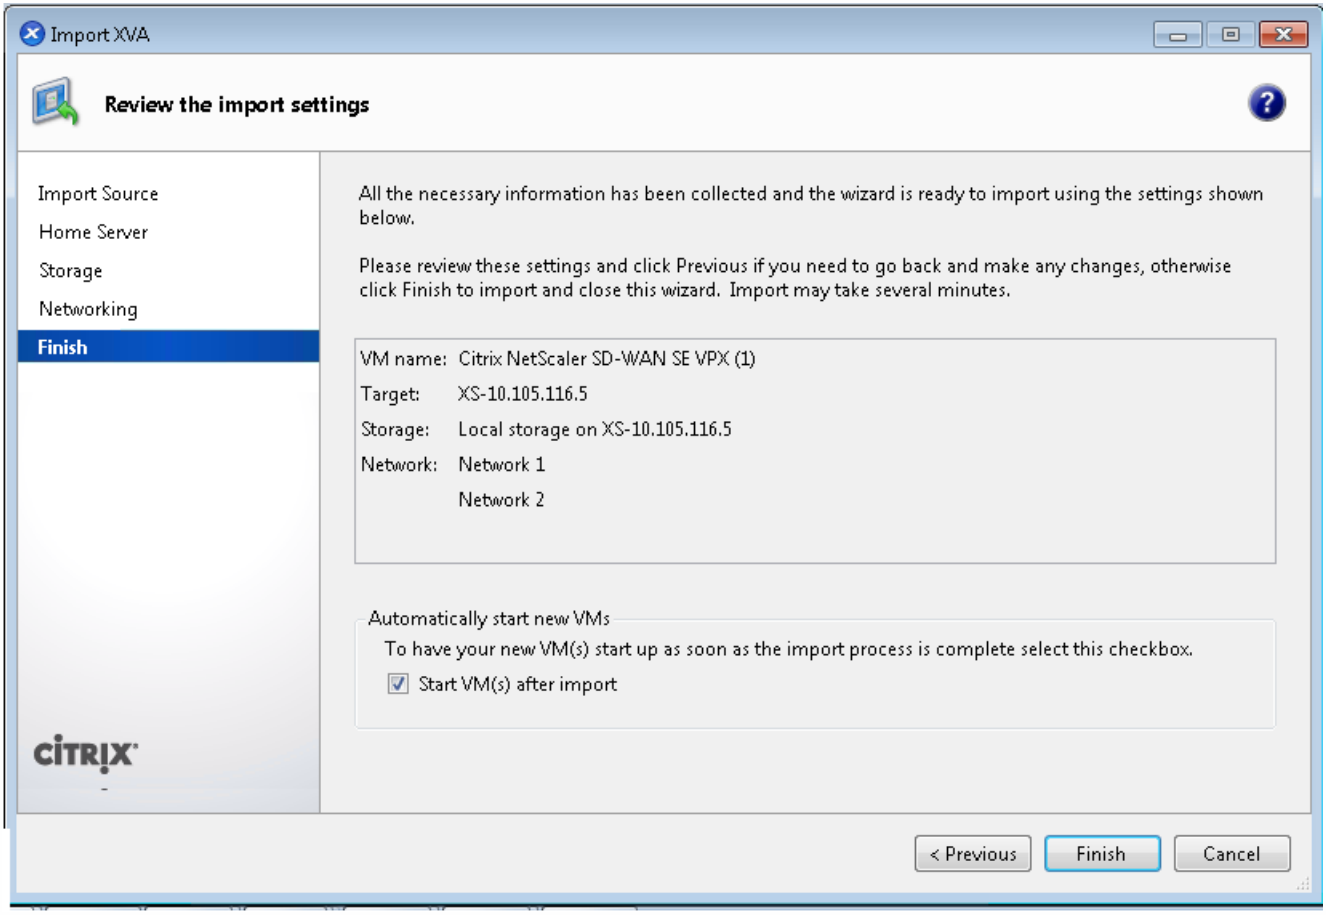

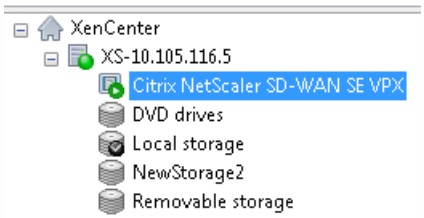

Important

Do not attach both virtual adapters to the same network . Doing so creates forwarding loops, which can cause network outages. Also, do not attach the two physical Ethernet ports associated with SD-WAN VPX to the same Ethernet switch.

To configure the virtual SD-WAN appliance

1. In XenCenter, select the icon for the SD-WAN VPX virtual machine. Then, on the Storage tab, select **Properties** and, in the **Properties** dialog box , adjust the disk allocation to the desired level.
2. Right-click the SD-WAN VPX icon and select the **Properties** option. Under CPU, select the number of **VCPUs** and under memory select the amount of VM memory corresponding to a supported configuration.
3. Set the basic network parameters. Depending on which release you are running, do one of the following:
4. After the virtual machine starts, go to the SD-WAN VPX console, login with admin user account.
5. Type the command *management_ip*.
6. Type the command *set interface <ip_address> <subnet mask> <gateway>*.
7. Type the command *"apply"*. Verify/confirm the configuration.
8. Ater the SD-WAN VPX has restarted, log on to the browser-based web management interface (GUI) (Default credentials: admin and password) at the IP address that you assigned to apA.
9. Complete the configuration.

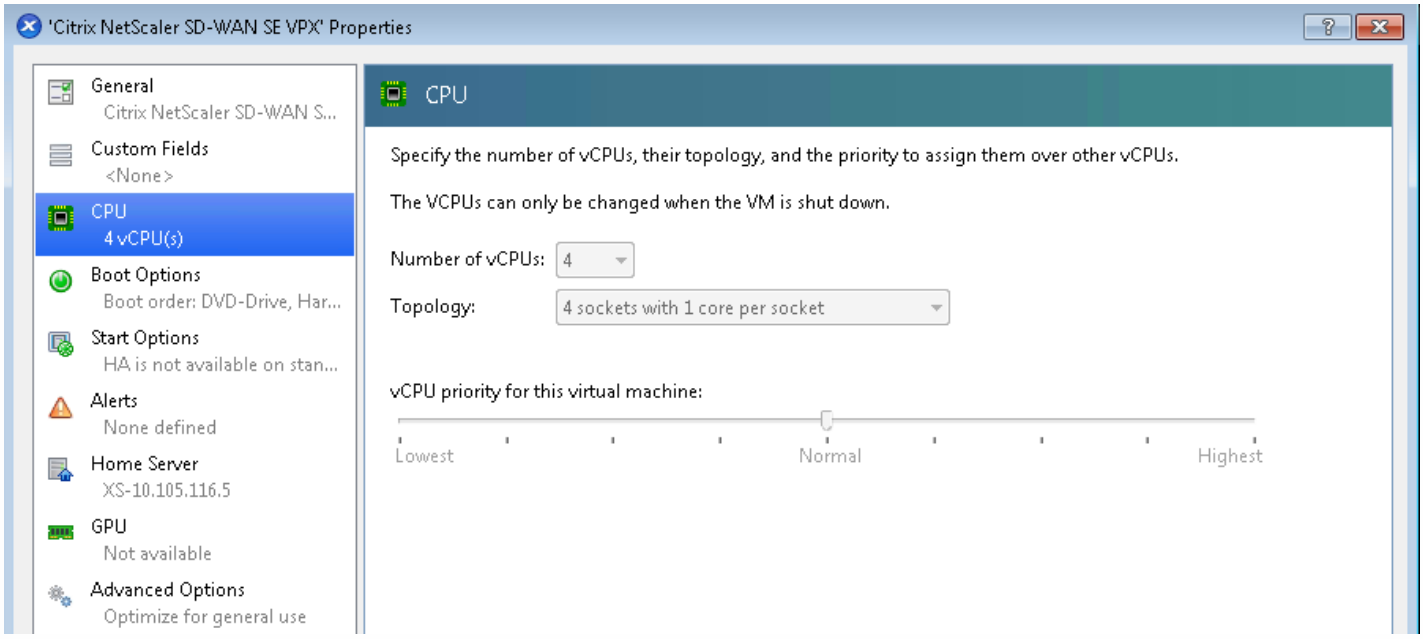

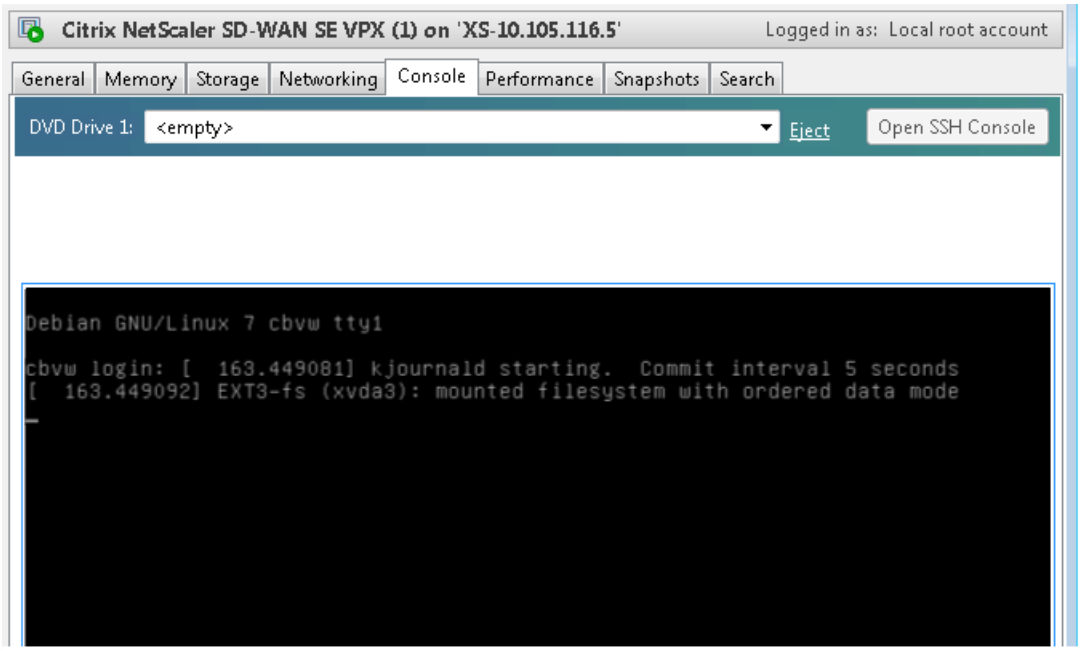

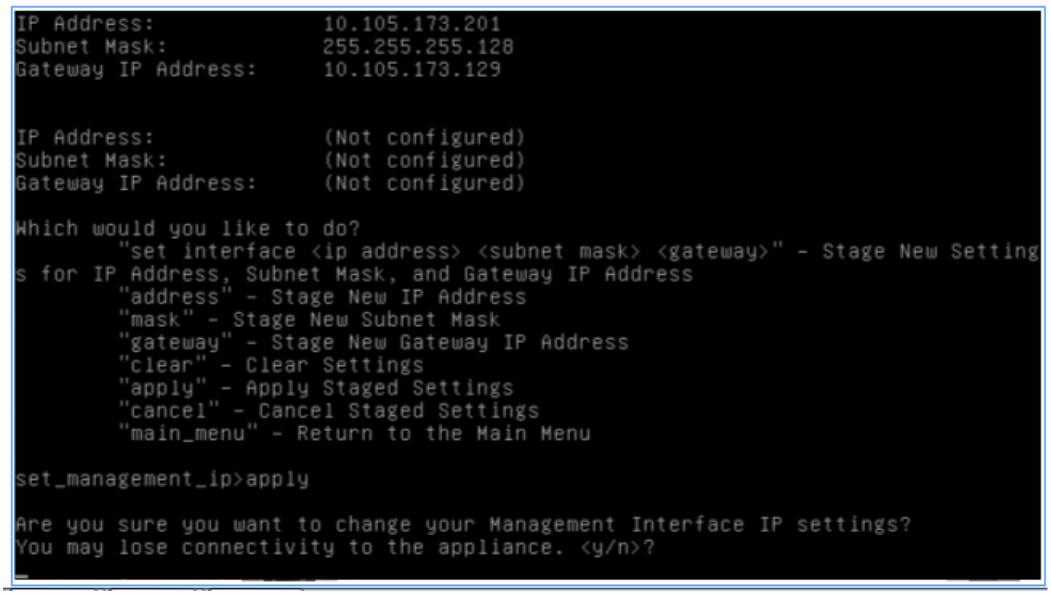

Note

- Changing the disk allocation on the SD-WAN VPX virtual machine resizes and reinitializes the compression history. Any accumulated history is lost.
- Do not attempt to change resource allocation while SD-WAN VPX is running.
- Do not use the Force Shutdown or Force Reboot commands. They might not work and can cause problems. Use the Shutdown and Reboot commands instead.

Installing SD-WAN Virtual Appliances on VMware ESX

Oct 24, 2017

Warning

Ensure that you enable the promiscuous mode on VM Network only. Do not enable the promiscuous mode on the Virtual Switch setting.

Note

VMware vSphere Client operation details might change with new releases of the vSphere software. For the most complete and current vSphere Client installation and operation instructions, please see also your VMware documentation. The instructions in this chapter are intended to provide the most basic and essential guidelines, only, for installing an SD-WAN VPX-SE Virtual appliance on the ESXi platform.

The following summarizes the top-level steps for installing and deploying an SD-WAN VPX-SE. Perform these procedures in the exact order listed.

1. Install the VMware vSphere Client.
2. Install and deploy the SD-WAN VPX-SE OVF Template.
3. Configure the SD-WAN VPX-SE Management IP Address.
4. Connect and test the deployment.

This chapter provides step-by-step instructions for installing, configuring, and deploying the SD-WAN VPX-SE. This includes basic instructions for installing the VMware vSphere Client, which you use to create and deploy the SD-WAN VPX-SE virtual machine.

Before you begin installing a virtual appliance, do the following:

- Install VMware ESX version 5.5 or ESXi 6.0, or later, on hardware that meets the minimum requirements.
- Install the VMware vSphere client on a management workstation that meets the minimum system requirements.
- Download the SD-WAN VPX-SE set up files.
- Obtain SD-WAN VPX-SE license files.

Also, before installing an SD-WAN VPX-SE virtual appliance, label all the interfaces that you plan to assign to VPX virtual appliances, in a unique format. In large deployments, labeling these interfaces in a unique format helps in quickly identifying them between other interfaces used by other virtual machines, such because Windows and Linux virtual machines. Such labeling is especially important if different types of virtual machines share interfaces.

SD-WAN VPX-SE requires non-default networking options. Between other things, you create two new virtual switches (vswitch0 and vswitch1) for the accelerated bridge, which must be assigned to two different virtual switches.

Installing the VMware vSphere Client

This section provides basic instructions for downloading and installing the VMware vSphere client you use to create and deploy the SD-WAN VPX-SE virtual machine.

> ### Note
>
> See also your VMware vSphere Client documentation for additional information.

1. Open a browser and navigate to the ESXi server that hosts your vSphere Client and SD-WAN VPX-SE virtual machine (VM) instance. The **VMware ESXi Welcome** page displays.

2. Choose the **Download vSphere Client** link to download the vSphere Client installation file.

3. Install the vSphere Client.

   Run the vSphere Client installer file that you just downloaded, and accept each of the default options whether prompted.

4. After the installation completes, start the vSphere Client program.

   The **VMware vSphere Client** login screen displays, prompting you for the ESXi server login credentials.

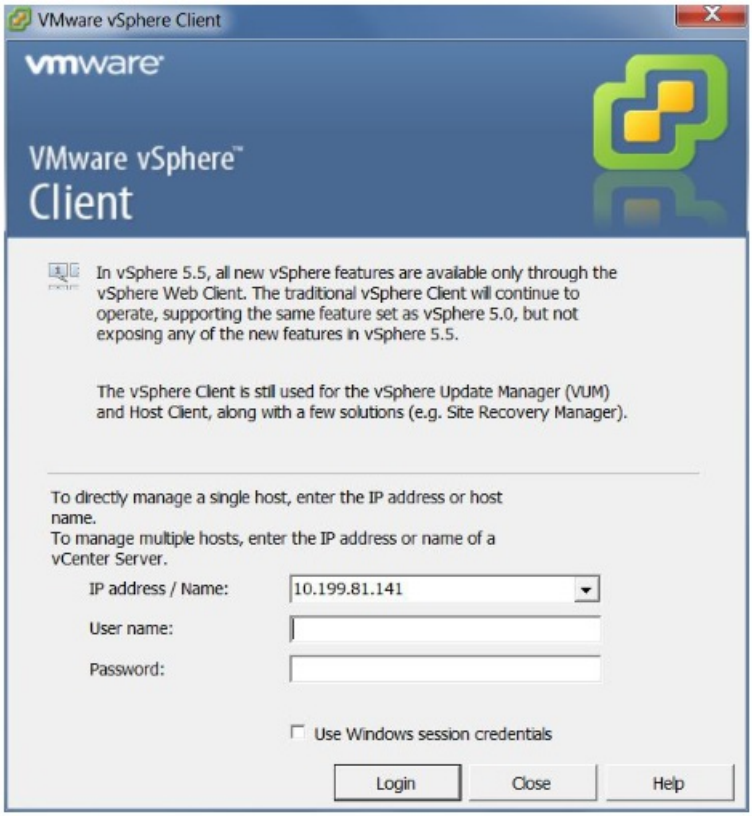

5. Type the ESXi server login credentials.

   Type the following:

- IP **address / Name**: Type the IP Address or Fully Qualified Domain Name (FQDN) for the ESXi server that hosts your SD-WAN VPX-SE VM instance.

- User **name:** Type the server administrator account name. The default is root.

- Password: Type the password associated with this administrator account.

6. Choose **Login.**

   This appears the **vSphere Client** main page.

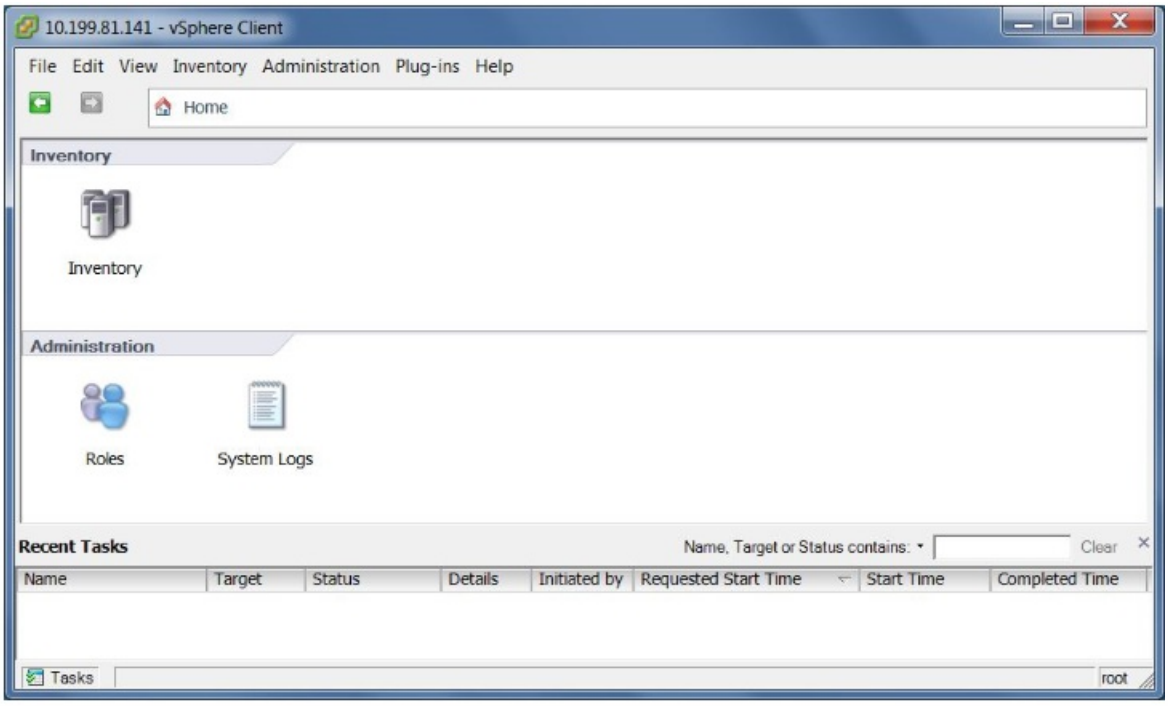

The next step is to install and deploy the SD-WAN VPX-SE OVF template and setup the virtual machine. The following section provides instructions for these procedures.

Installing and Deploying the SD-WAN VPX-SE OVF Template

This section provides instructions for installing the SD-WAN VPX-SE OVF template and creating the SD-WAN VPX-SE virtual machine.

1. When you have not already done so, download the SD-WAN VPX-SE OVF template file (.ova file) to the local PC.

   Download or copy the SD-WAN VPX-SE OVF template to the local PC you are using to connect to the ESXi server that hosts your SD-WAN VPX-SE. The OVF template file has a file name using the following naming convention:

   cb-vwc-version_number-vmware.ova

   Where:

   version_number is the SD-WAN VPX-SE release version number.

   .ova is the file name suffix indicating that this is an OVF template file.

> **Note**
>
> For additional information, please see downloading the Software Packages in System Requirements section.

2. Continuing in the vSphere Client, choose **File** and then choose **Deploy OVF Template...** from the drop-down menu.

   This appears the first page of the **Deploy OVF Template** wizard, the **Source** page.

3. Choose the SD-WAN VPX-SE OVF template *(.ova* file) you want to install.

   Browse to the location of the *.ova* file you downloaded earlier to the local PC, and choose it.

4. Choose **Next**.

   This imports the selected *.ova* file and appears the **OVF Template Details** page.

5. The next page appears some basic information regarding the OVF template you just imported.

6. Choose **Next**.

   This proceeds to the **EULA** page.

7. Choose **accept**, and then choose **Next**.

   This proceeds to the **Name and Location** page.

8. Type a unique name for the new VM (or accept the default).

   The name must be unique within the current **Inventory** folder, and can be up to 80 characters in length**.**

9. Choose **Next**.

   This proceeds to the **Storage** page.

10. Choose a datastore that has sufficient space available for the VM.

    The SD-WAN VPX-SE virtual machine requires 39.1 GB of disk space.

11. Choose **Next**.

    This appears the **Disk Format** page.

12. Accept the default settings, and choose Next.

    This proceeds to the **Network Mapping** page.

13. Accept the default (**VM Network**) and choose **Next**.

    This proceeds to the Ready to Complete page.

14. Choose **Finish** to create the VM. This appears the **Deploying Citrix NetScaler SD-WAN VPX-SE** status dialog box. Depending on the conditions present on your server, the deployment can take from several minutes to a few hours to complete. When the SD-WAN VPX-SE virtual machine has been successfully created, a success message displays.

15. Choose **close**. This closes the **Deploy OVF Template** wizard and returns to the vSphere Client main window. When this is the first VM you have created using this vSphere Client, the vSphere Client **home page** displays. When you have previously created one or more VMs, the **Inventory** page displays.

The next step is to configure the SD-WAN VPX-SE Management IP Address. The following section provides instructions for this procedure.

Configuring the Management IP Address for the SD-WAN VPX-SE

There are two methods for assigning the Management IP Address to the SD-WAN VPX-SE virtual machine:

- When you are not using DHCP: When you are not using DHCP, you must manually assign a static Management IP Address for the SD-WAN VPX-SE Virtual Appliance. For instructions, see Manually Configuring a Static Management IP Address for the VPX.
- When you are using DHCP: By default, all SD-WAN -VW Virtual Appliances use DHCP to acquire the Management IP Address. To use DHCP, the DHCP server must be present and available in the Virtual WAN. For instructions on identifying the acquired Management IP Address, see appearing the DHCP-assigned Management IP Address for the VPX.

Manually Configuring a Static Management IP Address for the VPX

When you are not using DHCP, or want to set a static Management IP Address for the SD-WAN VPX-SE Virtual Appliance VM, you must do this manually. To do so, you use the console for the virtual machine you just created, in the vSphere Client.

Also see, Setting up the SD-WAN Virtual WAN Appliances.

To set the Management IP Address manually, do the following:

> **Note**
>
> DHCP is enabled by default for the SD-WAN VPX-SE Management IP Address.

1. Continuing in the vSphere client **Inventory** page, choose the new SD-WAN VPX-SE VM in the Inventory tree (left pane).

   This appears the **Inventory** page for the new VM, with the **Getting Started** tab preselected.

2. Power on the new virtual machine.

   In the **Basic Tasks** section of the **Getting Started** tab page, choose **Power on the virtual machine** (green play mouse button) to power on the new SD-WAN VPX-SE VM.

3. Click the **Console** tab in the **Inventory** page tab bar.

   The Console tab is located in **Inventory** page tab bar at the top of the main page area. Selecting this tab appears and enables access to the CLI console for the VM.

Because the new VM starts up, a series of status messages are displayed in the console. When the startup process completes, the console login prompt displays.

4. Choose anywhere inside the console area to type console mode.

   This turns control of your pointing device cursor more than to the VM console, and enables console mode.

5. Log into the VM console.

   The default login credentials for the new SD-WAN VPX-SE VM are because follows:

   Login: *Administrator*

Password: *password*

This appears the console **Welcome** screen.

6. Type the following command line at the console prompt:

management_ip

This switches to the *management_ip* CLI in the console, and appears the *set_management_ip* prompt.

7. Configure the interface settings for the VM.

Type the following command line at the *set_management_ip* prompt:

set interface <ipaddress> <subnetmask> <gateway>

Where:

- *<ip>* is the Management IP Address for the SD-WAN VPX-SE Virtual Appliance.

- *<subnetmask>* is the subnet mask used to define the network in which the SD-WAN VPX-SE Virtual Appliance resides.

- *<gateway>* is the Gateway IP Address the SD-WAN VPX-SE Virtual Appliance uses to communicate with external networks.

This stage but does not apply the interface settings.

8. Apply the staged settings for the VM interface.

Do the following:

a. Type the following command at the *set_management_ip* prompt:

Apply

b. When prompted to confirm the *apply* operation, type *Y*.

This applies the staged interface settings for the VM, and appears the results.

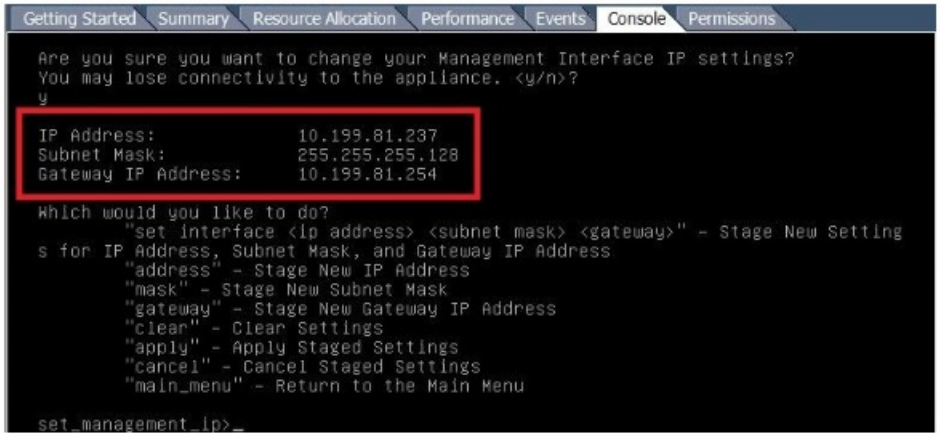

9. Type exit and press Return at the prompt to exit the management_ip CLI.

10. Exit the console.

Type **exit and press Return** at the console prompt, and then press **Ctrl+Alt to regain control of the cursor.**

11. Shutdown and start the VM.

Do the following:

a. Choose the **Getting Started** tab to display the **Basic Tasks** options.

b. In the **Basic Tasks** section, choose **shutdown the virtual machine** (red check box icon). You are prompted to confirm that you want to end the guest operating system for the VM.

c. Choose **Yes** to confirm. This shuts down the guest operating system and powers off the VM. When the shutdown completes, the **Power on the virtual machine** option (green play mouse button) becomes available.

12. Start the virtual machine. Choose **Power on the virtual machine** (green right-arrow) to start the VM. You can view the progress of the start-up process in the **Console** tab page for the VM.

When the startup process completes, the login prompt displays. You can now proceed to the final step, Connecting to the SD-WAN VPX-SE and Testing the Deployment,

Installing SD-WAN Appliances on the Microsoft Hyper-V Platform

Aug 09, 2017

To install Citrix SD-WAN virtual appliances on Microsoft Windows Server, you must first install Windows Server, with the Hyper-V role enabled, on a machine with adequate system resources. While installing the Hyper-V role, be sure to specify the network interface cards (NICs) on the server that Hyper-V will use to create the virtual networks. You can reserve some NICs for the host. Use Hyper-V Manager to perform the SD-WAN VPX installation.

SD-WAN VPX for Hyper-V is delivered in virtual hard disk (VHD) format. It includes the default configuration for elements such as CPU, network interfaces, and hard-disk size and format. After you install a SD-WAN VPX instance, you can configure its network adapters, add virtual NICs, assign the SD-WAN IP address, subnet mask, and gateway, and complete the basic configuration of the virtual appliance.

Microsoft Server Hardware Requirements

- The server's processor must support Intel Virtualization Technology.
- The server must run 64-bit Windows 2008 R2 SP1 (Standard, Enterprise, or DataCenter Editions), or 2012 (Standard or DataCenter Editions) with a full installation (not a Core installation), and the Hyper-V component enabled.
- Minimum system configuration is 4 GB RAM, 200 GB hard drive, and 2 physical CPU.
- Two physical Ethernet NICs are required; three are recommended.
  Note: The procedure below uses three NICs.
  For more information about Windows Server 2008 R2 system requirements, see http://www.microsoft.com/windowsserver2008/en/us/system-requirements.aspx(the exact location is subject to change by Microsoft at any time).

  For information about installing Microsoft Server 2008 R2, see http://technet.microsoft.com/en-us/library/dd379511(WS.10).aspx(the exact location is subject to change by Microsoft at any time).

Prerequisites for Installing SD-WAN Virtual Appliances on the Microsoft Hyper-V platform

Before you begin installing a virtual appliance, do the following:

- Enable the Hyper-V role on Windows Server 2008 R2 or 2012. For more information, see http://technet.microsoft.com/en-us/library/ee344837(WS.10).aspx(the exact location is subject to change by Microsoft at any time).
- Download the VPX setup files. If you do not have a My Citrix account, access the home page at http://www.mycitrix.com, click the New Users link, and follow the instructions to create a new My Citrix account.

To download the SD-WAN VPX setup files

1. In a Web browser, go to http://www.citrix.com/ and click My Citrix.
2. Type your user name and password.
3. Click Downloads.
4. In Search Downloads by Product, select NetScaler SD-WAN.
5. Under Virtual Appliances, select and download the required SD-WAN VPX distribution.
6. Copy the compressed file to your server.

To configure virtual NICs on the SD-WAN VPX

1. Log on to the Windows Server as an Administrator, either at a keyboard or VGA console, or through a NIC that you plan to use for managing the virtual appliance (not at one of the ports that you will use for the accelerated bridge).
2. To start Hyper-V Manager, click Start, point to Administrative Tools, and then click Hyper-V Manager
3. In the navigation pane, under Hyper-V Manager, select the server on which you want to install SD-WAN VPX.
4. On the Actions menu, click Virtual Network Manager...
5. In the Virtual Network Manager window, in the navigation pane, under Virtual Networks, click New virtual network.
6. Choose External as type of virtual network, and then click Add.
7. Name the new virtual network as apA Network 1 and select the physical NIC to map it to.
8. Click OK to apply the changes.
9. The Apply Networking Changes popup displays a caution indicating that pending changes might disrupt network connectivity. Click Yes.
10. Repeat steps 5 to 9 for the second accelerated bridge port, but name it as apA Network 2 and connect it to a different physical port.
11. Click Apply to apply the networking changes.

Installing SD-WAN VPX on Microsoft Server by using Hyper-V Manager

After you have enabled the Hyper-V role on Microsoft Server and extracted the VPX files, you can use Hyper-V Manager to install SD-WAN VPX. After you import the virtual machine, you must configure the virtual NICs by associating them with the virtual networks created by Hyper-V. Based on the Microsoft server you are using, follow the procedure to complete the installation.

- Microsoft Server 2008 R2
- Microsoft Server 2012

Installing SD-WAN VPX on Microsoft Server 2008 R2

Aug 09, 2017

Performing the Installation Procedures

After you have enabled the Hyper-V role on Microsoft Server 2008 R2 and extracted the VPX files, you can use Hyper-V Manager to install SD-WAN VPX. After you import the virtual machine, you must configure the virtual NICs by associating them with the virtual networks created by Hyper-V.

Note: You cannot change any settings while the virtual appliance is running. Shut down the virtual appliance and then make changes.

To install SD-WAN VPX on Microsoft Server 2008 R2 by using Hyper-V Manager

1. Unzip the SD-WAN distribution that you downloaded from My Citrix.
2. Start **Hyper-V Manager**.
3. In the navigation pane, under **Hyper-V Manager**, select the server on which you want to install SD-WAN VPX.
4. On the Actions menu, click **Virtual Switch Manager**.
5. In the **Import Virtual Machine** dialog box, in **Location**, specify the path to the folder that contains the Branch VPX SD-WAN files.
   Note: If you received a compressed file, make sure that you extract the files into a folder before you specify the path to the folder.
6. Click Import.
7. Verify that the virtual appliance that you imported is listed under Virtual Machines.
8. Right-click the imported virtual machine, and then click **Settings**.
9. In the **Settings** window's navigation pane, under Hardware, select the first network adapter in the list.
   Figure 1. Configuring Ethernet ports using Hyper-V Manager

10. In the Network drop down menu, select apA Network 1 . This is the LAN interface for apA1.
11. Make sure the Enable spoofing of the MAC addresses box is selected. If it is not, select it and apply the changes.
12. In the Settings window's navigation pane, under Hardware, select the second network adapter in the list. Repeat the step 10 and step 11, and assign the adapter to apA Network 2. This is the WAN interface for apA2.
    Important: Do not configure the same Network for both the network adapters. Incorrect configuration creates packet loops, which can bring down the network.
13. Optionally, change the virtual hard disk size:
    - In the Settings window navigation pane, under IDE Controller 0, select Hard Drive.
    - Click Edit.
    - Follow the steps in the Edit Virtual Hard Disk Wizard to increase the allocation to one of the supported sizes, using the Expand option in the wizard.
      Figure 2. Configuring disk and RAM allocation

14. Optionally, change the memory size.
    - In the Settings window's navigation pane, under Hardware, select Memory.
    - Allocate the RAM space by adjusting the memory to one of the supported sizes.

- Click OK.

15. Optionally, define the management port.
    - Right-click the virtual machine, and then click Settings.
    - In the Settings window navigation pane, under Hardware, select Add Hardware.
    - Select Network Adapter from the list of devices, and then click Add.
    - Name the new virtual network as Primary Network 3.
    - Make sure the Enable spoofing of MAC addresses check box is selected.
    - Click OK to apply the changes.
16. Right-click the Branch Repeater VPX virtual machine and select Connect.
17. In the file menu, click Action, and then click Start to start the virtual machine.
    Figure 3. Starting the VPX virtual machine

18. When a SD-WAN VPX virtual machine is started for the first time, it automatically starts the Deployment Wizard. This wizard asks questions about the deployment mode: Inline, WCCP, or PBR (virtual inline), or Setup Using Web UI. Select Setup Using Web UI. On the next screen, enter the IP, netmask, and gateway for the apA interface, and click Finish.
19. After SD-WAN VPX has restarted, log on to the browser based UI ((user name: admin, password: password) at the IP address that you assigned to apA, for example:
    https://172.16.0.213

Additional Configuration

For additional configuration instructions, see the documentation for physical SD-WAN and SD-WAN appliances.

Upgrading to a Previous Release

The software upgrade mechanism built into physical SD-WAN appliances is also supported by SD-WAN VPX. Alternatively, you can install a new virtual machine running the desired release.

Installing SD-WAN VPX on the Microsoft Server 2012

Aug 09, 2017

Performing the Installation Procedures

After you have enabled the Hyper-V role on Microsoft Server and extracted the VPX files, you can use Hyper-V Manager to install SD-WAN VPX. After you import the virtual machine, you must configure the virtual NICs by associating them with the virtual networks created by Hyper-V.

Note: You cannot change any settings while the virtual appliance is running. Shut down the virtual appliance and then make changes.

To install SD-WAN VPX on Microsoft Server 2012 by using Hyper-V Manager

1. Unzip the Sd-WAN distribution that you downloaded from My Citrix.
2. Start Hyper-V Manager.
3. In the navigation pane, under Hyper-V Manager, select the server on which you want to install SD-WAN VPX.
4. On the Actions menu, click Import Virtual Machine .
5. In the Import Virtual Machine dialog box, in Location box, specify the path to the folder that contains the SD-WAN VPX files.
   Note: If you received a compressed file, make sure that you extract the files into a folder before you specify the path to the folder.
6. Click Import.
7. Verify that the virtual appliance that you imported is listed under Virtual Machines.
8. Right-click the imported virtual machine, and then click Settings.
9. In the Settings window's navigation pane, under Hardware, select the first network adapter in the list.
10. In the Network drop down menu, select apA1 Network. This is the LAN interface for apA1.
11. Make sure the Enable MAC address spoofing box is selected. If it is not, select it and apply the changes.
12. In the Settings window's navigation pane, under Hardware, select the second network adapter in the list. Repeat the step 10 and step 11, and assign the adapter to apA2 Network. This is the WAN interface for apA2.
    Important: Do not configure the same Network for both the network adapters. Incorrect configuration creates packet loops, which can bring down the network.
13. Optionally, change the virtual hard disk size:
    - In the Settings window navigation pane, under IDE Controller 0, select Hard Drive.
    - Click Edit.
    - Follow the steps in the Edit Virtual Hard Disk Wizard to increase the allocation to one of the supported sizes, using the Expand option in the wizard.
14. Optionally, change the memory size.
    - In the Settings window's navigation pane, under Hardware, select Memory.
    - Allocate the RAM space by adjusting the memory to one of the supported sizes.
    - Click OK.
15. Optionally, define the management port.
    - Right-click the virtual machine, and then click Settings.
    - In the Settings window navigation pane, under Hardware, select Add Hardware.
    - Select Network Adapter from the list of devices, and then click Add.
    - Name the new virtual network as Primary Network 3.

- Make sure the **Enable spoofing of MAC addresses** check box is selected.
- Click OK to apply the changes.

16. Right-click the SD-WAN VPX virtual machine and select Connect.
17. In the file menu, click Action, and then click Start to start the virtual machine.
18. When a SD-WAN VPX virtual machine is started for the first time, it automatically starts the Deployment Wizard. This wizard asks questions about the deployment mode. Select Setup Using Web UI. On the next screen, enter the IP address, netmask, and gateway for the apA interface, and click Finish.
19. After SD-WAN VPX has restarted, log on to the browser based UI ((user name: admin, password: password) at the IP address that you assigned to apA, for example:
    https://172.16.0.213

Additional Configuration

For additional configuration instructions, see the documentation for physical SD-WAN/SD-WAN appliances.

Downgrading to a Previous Release

The software upgrade mechanism built into physical SD-WAN/SD-WAN appliances is also supported by Sd-WAN/SD-WAN VPX. Alternatively, you can install a new virtual machine running the desired release.

Installing SD-WAN Standard Edition AMI on Amazon AWS

Aug 09, 2017

The NetScaler SD-WAN SE appliances bond multiple network paths in single virtual path. The virtual paths are monitored so that critical application paths are always routed through optimal paths. This solution enables customers to deploy applications in the cloud and utilize multiple service provider networks for seamless delivery of applications to the end-users.

To create a SD-WAN SE-VPX on Amazon AWS, you go through the same process as with creating any other instance, setting a few instance parameters to non-default settings.

Instantiating a SD-WAN Virtual Appliance (AMI) on AWS:

To install a SD-WAN virtual appliance in an AWS VPC, you need an AWS account. You can create an AWS account at http://aws.amazon.com/. SD-WAN is available as an Amazon Machine Image (AMI) in AWS Marketplace.

Note: Amazon makes frequent minor changes to its AWS pages, so the following instructions may not be up-to-date.

To instantiate a SD-WAN virtual appliance (AMI) on AWS:

1. In a web browser, type http://aws.amazon.com/.
2. Click My Account/Console, and then click My Account to open the Amazon Web Services Sign in page.
3. Use your Amazon AWS account credentials to sign in. This will take you to the Amazon Web Services page.

NetScaler SD-WAN SE appliances offers the following AWS service instances:

VPC Dashboard - isolated portion of the AWS cloud populated by AWS objects, such as EC2 instances

– Enabled by creating a VPC in AWS. See below for configuration steps.

EC2 Dashboard - elastic compute cloud, resizable virtual services / instances

– Enabled by creating NetScaler SD-WAN AMI. See below for configuration steps.

–CIDR – Classless Inter-Domain Routing block, consisting of continuous IP address range, used to specify your VPC (cannot be larger then 16 regions).

SD-WAN Web Interface

– Configure NetScaler SD-WAN AMI

The following are the requirements and limitations for deploying SD-WAN SE-VPX AMI in AWS:

• **Minimum requirements**:

–AWS EC2 Instance Type: c3.2xlarge

–Virtual CPU: 8

–RAM: 15 GB

– Storage: 160 GB

– Network Interfaces: minimum of 2 (one management, one for LAN/WAN), maximum of #

– BYOL – bring-your-own-license and subscription

• **Limitations**:

– AWS does not allow bridging of interface so Fail-to-Wire is not an option for configuring Interface Groups

NetScaler SD-WAN with AWS

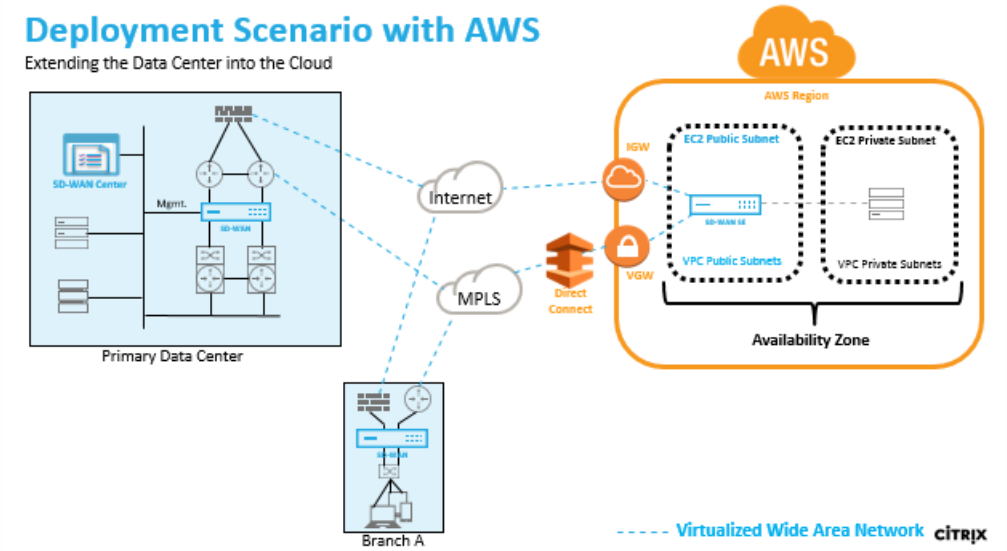

Deploying an AWS region with a specified Availability Zone. Within that Virtual Private Cloud (VPC) infrastructure, SD-WAN Standard Edition AMI (Amazon Machine Image) is deployed as the VPC Gateway.

- The VPC private has routes towards the VPC Gateway.

- SD-WAN instance has a route towards the AWS VGW (VPN Gateway) for direct connect and another route towards IGW (Internet Gateway) for internet connectivity.

- Connectivity between Data Center, Branch, and Cloud leveraging different transport modes utilizing multiple WAN paths simultaneously.

-Automatic Route Learning with OSPF and BGP.

-Single IPSec tunnel across multiple paths where security re-negotiation is not required upon any link failure occurrences.

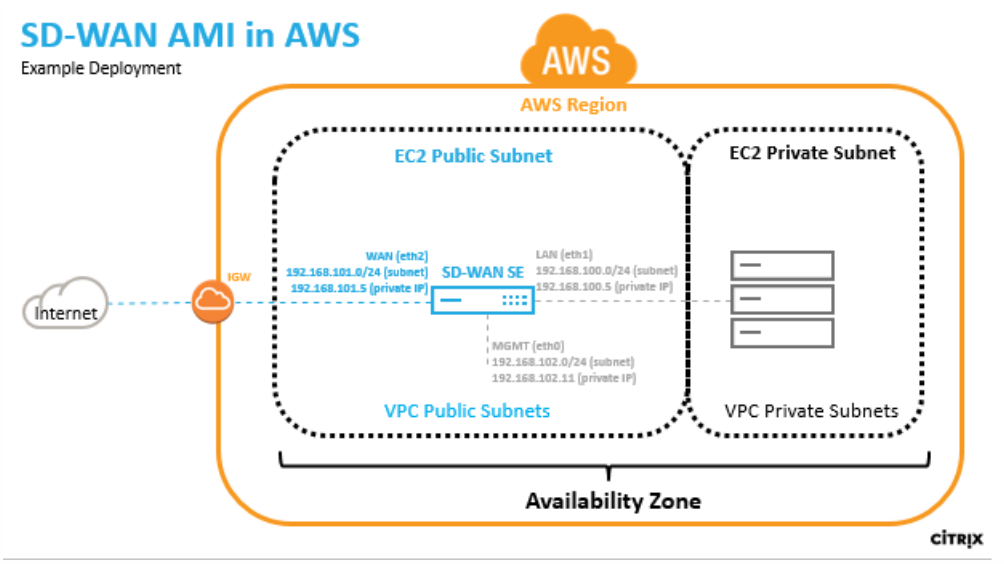

In AWS a subnet and IP address must be defined for each SD-WAN AMI interface. The number of interfaces utilized depends on the deployment use case. If the goal is to reliably access application resources that are on the LAN side of the VPX (inside the same Region), the VPX can be configured with three Ethernet interfaces; one for management on eth0, one for LAN on eth1, and one for WAN on eth2.

Alternatively, if the goal is to hair-pin traffic through the VPX to some other region or to the public internet, the VPX can be configured with two Ethernet interfaces; one for management on eth0, and a second for LAN/WAN on eth1.

SD-WAN SE AMI in AWS Overview

1. **Create VPC in AWS using VPC Dashboard**

To get started with Amazon Virtual Private cloud you need to create a VPC, which is a virtual network dedicated to your AWS account.

- Define CIDR blocks/Subnets and assign to VPC - for identifying the device in the network. For example; 192.168.100.0/22 is selected for the VPC in the example network diagram encompassing the WAN, LAN and Management subnets – 192.168.100.0 – 192.168.103.255) - 192.168.100.0/22
- Define an Internet Gateway for the VPC – for communicating with outside the cloud environment
- Define routing for each defined subnet - for communication between the subnets and Internet
- Define Network ACLs (Access Control List) - for controlling the inflow/outflow of the traffic from/to the subnet for security purposes
- Define Security Group - for controlling the inflow/outflow of the traffic from/to each instance of network device

2. **Create an NetScaler SD-WAN AMI**

- Define the Network Interfaces for the EC2 instance
- Create Elastic IP addresses for the EC2 instance
- Define Security for the EC2 instance and network interfaces

3. **Connect to the SD-WAN web interface**

- License
- Install identify using Local Change Management

Create a VPC in AWS - Virtual Private Cloud (VPC)

To create VPC:

1. From the AWS management console tool bar, select **Services** -> **VPC** (Networking & Content Delivery).

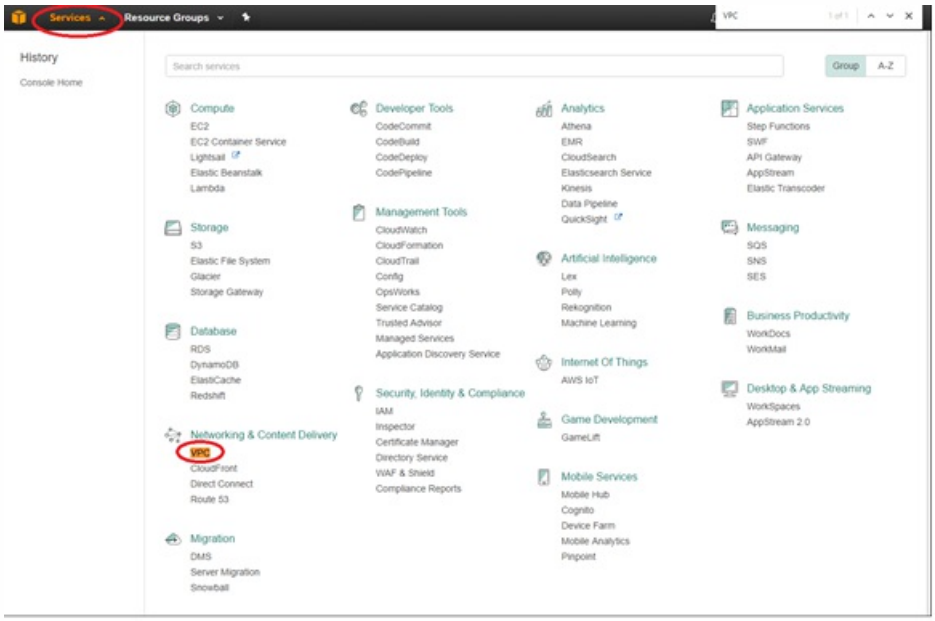

2. Select your **VPC**s, then click the **Create VPC** button.

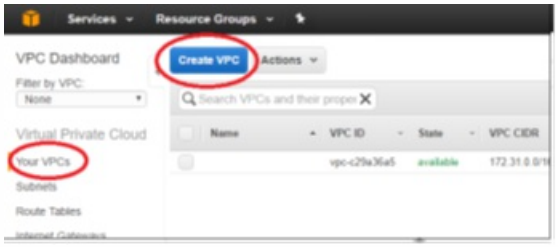

3. Add **Name** tag, CIDR block according to your network diagram and Tenancy = default, and click **Yes, Create**.

Define an Internet Gateway for the VPC

To define internet gateway for the VPC:

1. From the AWS management console, select **Internet Gateways** - > **Create Internet Gateway**. The Internet Gateway traffic matching the 0.0.0.0/0 route needs to be configured in the route table. It is also required for external

access to the SD-WAN AMI web interface for further configuration.

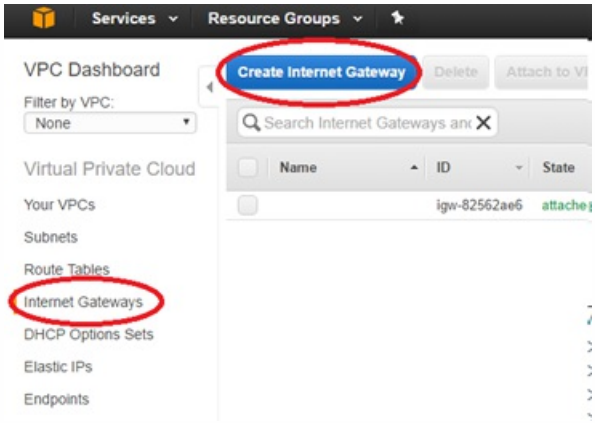

2. Give the IGW a Name tag, and click **Yes, Create**.

3. Select newly created IGW and click **Attach to VPC**.

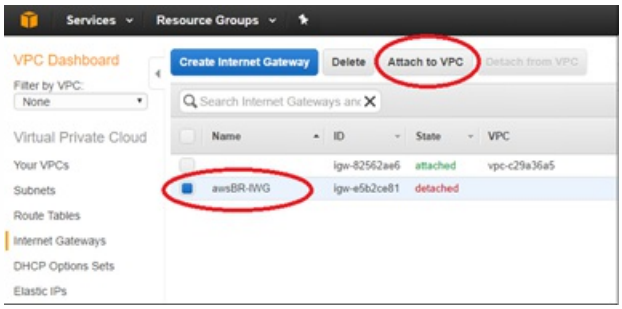

4. Select the previously created VPC and click **Yes, Attach**.

Define Subnets for the VPC to differentiate Mgmt, LAN, and WAN

To define subnets for VPC:

1. From the AWS management console, select **Subnets** > **Create Subnets** to create Mgmt, LAN and WAN subnets. Use the defined subnets to distinguish between the LAN, WAN and Mgmt subnets defined in the SD-WAN configuration.

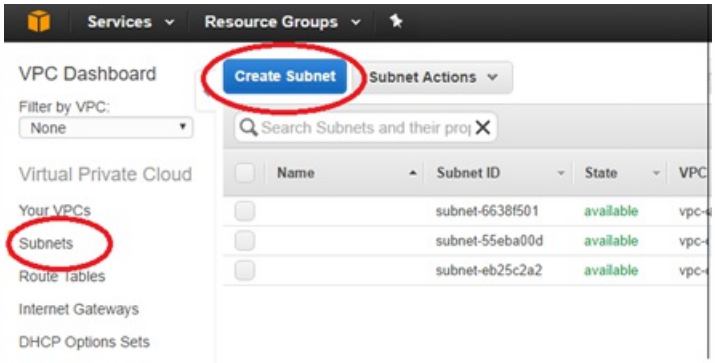

2. Enter the details specific for the Mgmt subnet of the VPX, then create it using the **Yes, Create** button.

    - Name tag: name to identify different subnets (Mgmt, LAN or WAN)

    - VPC: <the VPC previously created>

    - Availability Zone: <set at discretion>

    - CIDR block: subnet specific to the defined name (Mgmt, LAN or WAN) that is a smaller subset of the CIDR previously defined

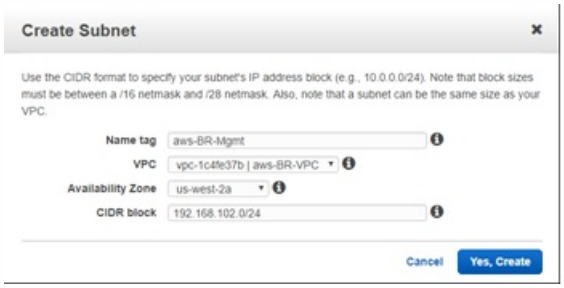

3. Repeat the process until you have created a subnet for the Mgmt, LAN, and WAN networks.

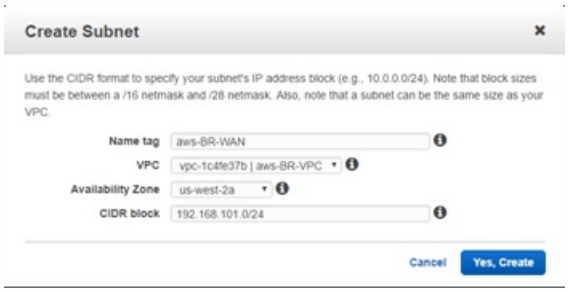

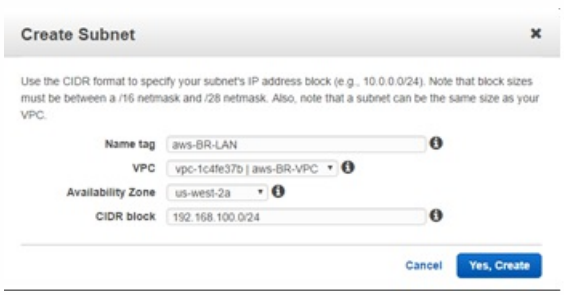

Define Route Tables for the Management Subnet

To define route tables:

1. From the AWS management console, select **Route Tables** > **Create Route Table** to create route tables for the

Mgmt, LAN and WAN subnets.

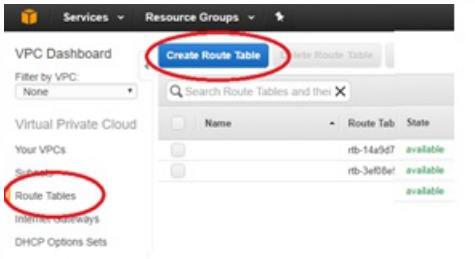

2. Enter the detail for the Mgmt subnet

   • Name tag: name to identify different subnets (Mgmt, LAN or WAN)

   • VPC: The previously created VPC

3. With the newly created route table still highlighted, select **Subnet Association** > **Edit**.

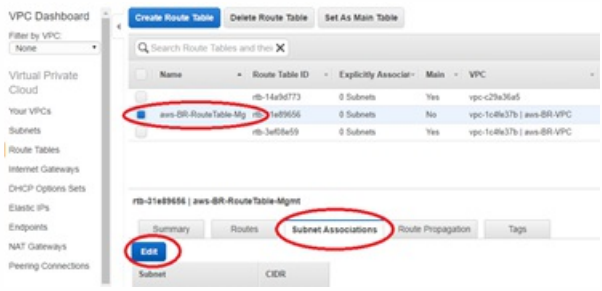

4. Make the association with the desired subnet, then click **Save**.

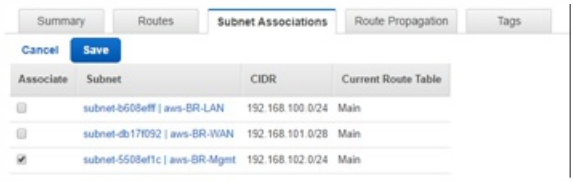

5. With the newly created route table still highlighted, select **Routes** > **Edit**.

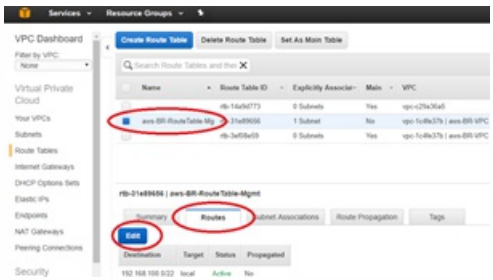

6. Click **Add** another route button (only required for the Mgmt, and WAN subnets), then **Save**.

   • Destination: 0.0.0.0/0

• Target: The Internet Gateway (igw-xxxxxxx previously defined)

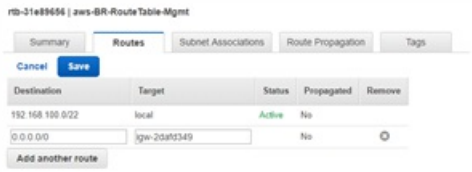

Note

AWS provides a global route table in the EC2 instance but the NetScaler SD-WAN AMI will use local route tables so that the user can control traffic forwarding to the Virtual Path.

Define Route Tables for the WAN Subnet

To define route tables:

1. From the AWS management console, select **Route Tables** > **Create Route Table** to create route tables for the Mgmt, LAN and WAN subnets.

2. Enter the details for the WAN subnet:

   • Name tag: name to identify different subnets (Mgmt, LAN or WAN)

   • VPC: The previously created VPC

3. With the newly created route table still highlighted, select **Subnet Association** > **Edit**.

4. Make the association with the desired subnet, then click **Save**.

5. With the newly created route table still highlighted, select **Routes** > **Edit**.

6. Click **Add** another route button (only required for the Mgmt, and WAN subnets), then **Save**.

   •Destination: 0.0.0.0/0

   •Target: <The Internet Gateway (igw-xxxxxxx previously defined)

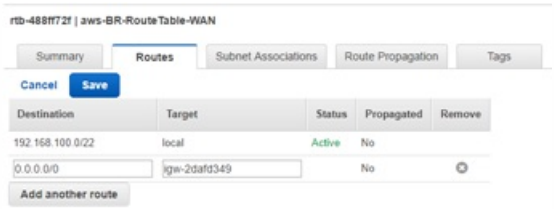

Define Route Tables for the LAN Subnet

To define route tables for the LAN subnet:

1. From the AWS management console, select **Route Tables** > **Create Route Table** to create route tables for the Mgmt, LAN and WAN subnets.

2. Enter the details for the LAN subnet:

- Name tag: name to identify different subnets (Mgmt, LAN or WAN)
- VPC: The previously created VPC

3. With the newly created route table still highlighted, select **Subnet Association** > **Edit**.

4. Make the association with the desired subnet, then click **Save**.

Create an SD-WAN SE AMI

To create EC2 instance:

1. From the AWS management console tool bar, select **Services** -> **EC2 (Compute)**.

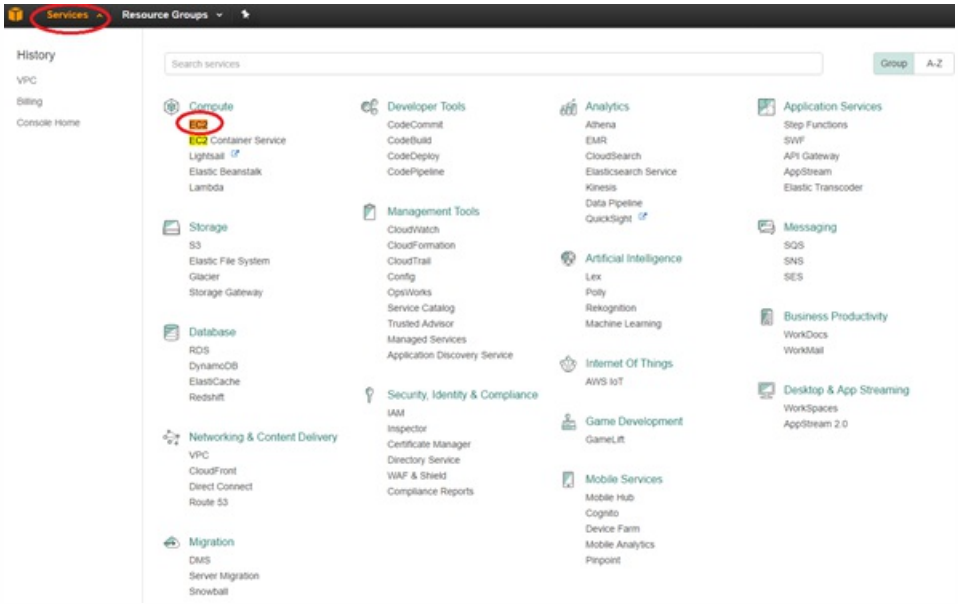

2. Select the **EC2** dashboard tool bar, select **Instances** > **Launch Instance**.

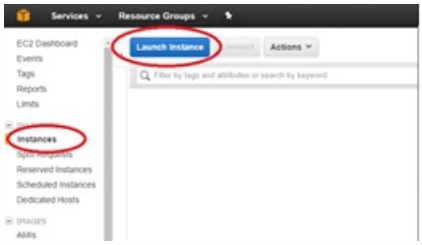

3. Use the **AWS Marketplace** tab to search for the SD-WAN Amazon Machine Image (AMI) or use the **My AMIs** tab to locate an owned or shared SD-WAN AMI, locate **Citrix NetScaler SD-WAN Standard Edition** and then click **Select**.

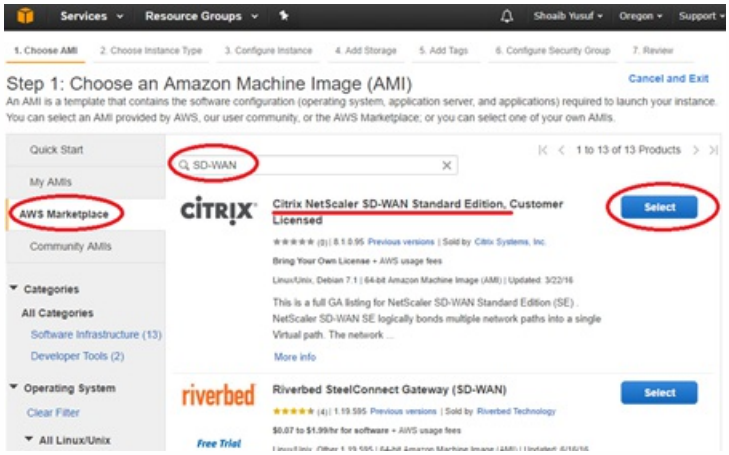

4. Confirm the selection with **Continue**.

5. On the **Choose Instance Type** screen, select the **EC2 Instance Type** that was identified during preparation, then select **Next: Configure Instance Details**.

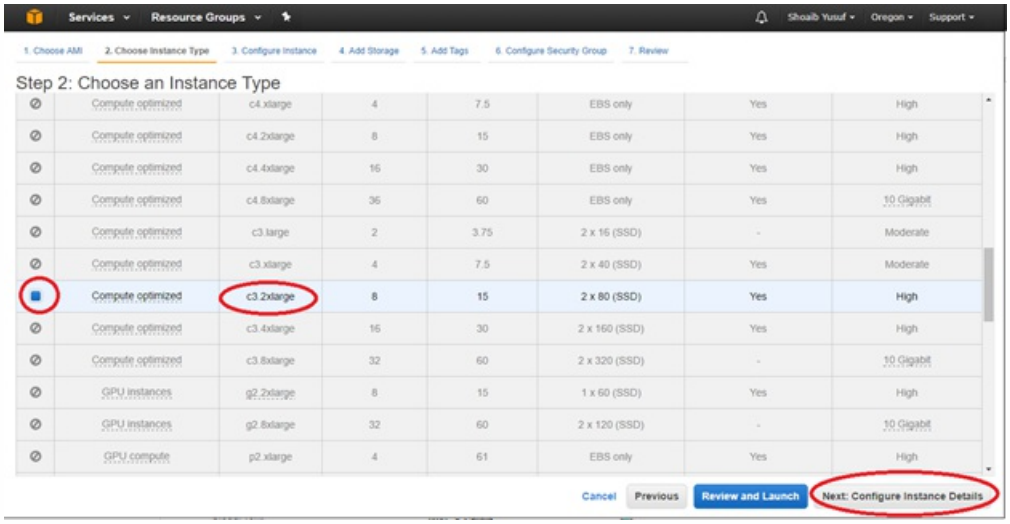

6. Enter instance details (anything not specified should be left unset/default):

- Number of Instances: 1
- Network: <select VPC previously created>
- Subnet: <select Mgmt Subnet previously defined>
- Auto-assigning Public IP: enabled
- Network interfaces > Primary IP: <enter predefined Mgmt IP>

7. Click **Next: Add Storage**

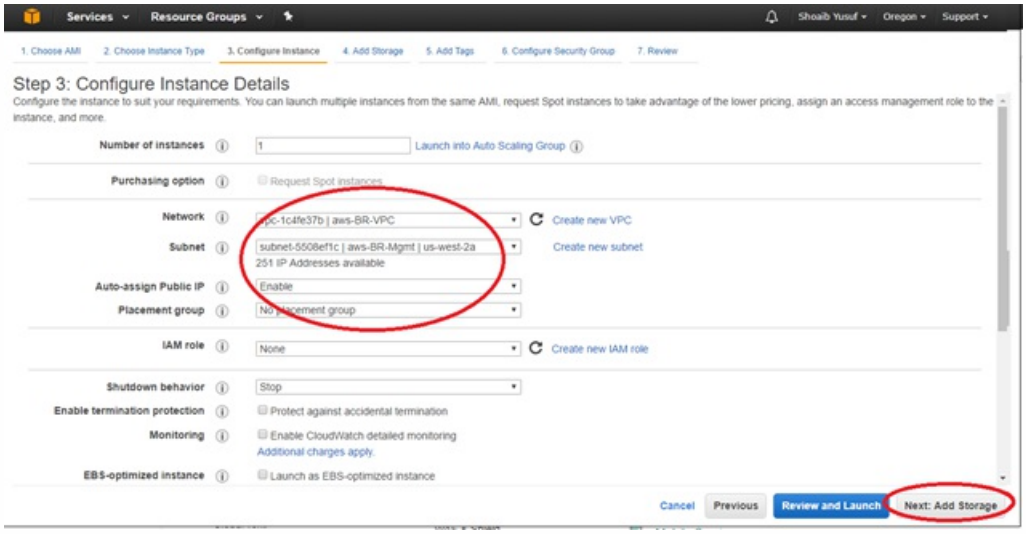

Note

You must associate the EC2 instance with the Mgmt Subnet in order to associate the first EC2 interface (eth0) with the SD-WAN Mgmt interface. If eth0 is not associated with the SD-WAN Mgmt interface, connectivity will be lost following a reboot.

8. Enter the following information for the Root Storage:

- Volume Type: General Purpose (SSD) GP2

9. Then select **Next: Tag Instance**

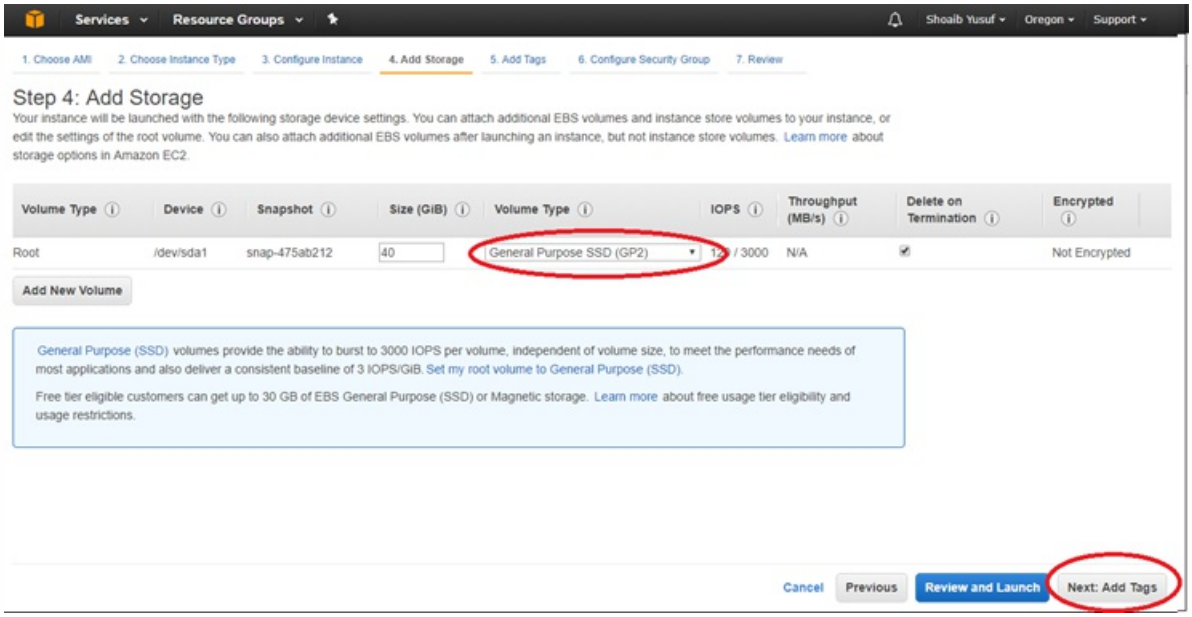

10. Give the EC2 instance a name by specifying a value for the default **Name** Tag. Optionally create other desired Tags.

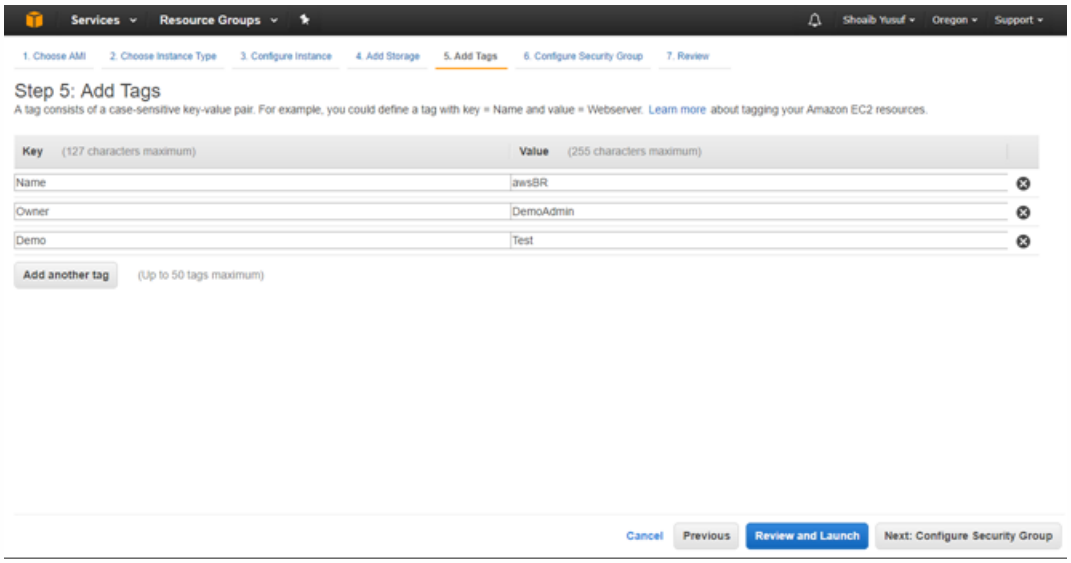

11. Then select **Next: Configure Security Group**.

12. Select an existing **Security Group** or create a new Security Group:

- Default security group generated will include HTTP, HTTPS, SSH

Click the Add Rule button to add two more:

- All ICMP with Source: Custom 0.0.0.0/0
- Custom UDP Rule with Port Range: 4980 and Source: custom <known IP addresses from partner SD-WAN>

13. Select **Review** and **Launch**.

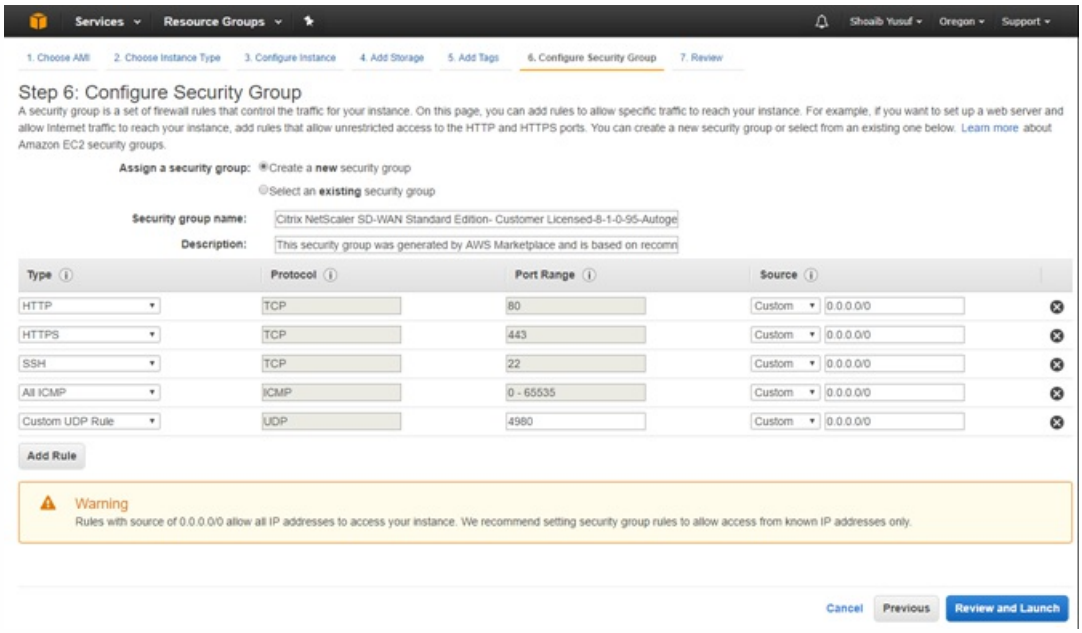

14. After complete with reviewing, select **Launch**.

15. In the **Key Pair** pop-up, either select an existing key pair or create a new key pair, then select **Launch Instance**.

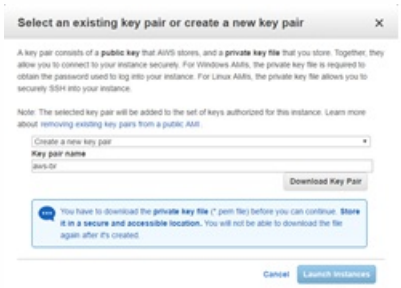

Important

If a new key pair is created, be sure to download and store it in a safe location.

16. NetScaler SD-WAN SE AMI should now be launched successfully.

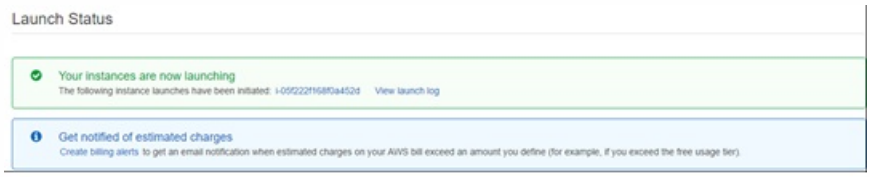

Note

A Security Group is a set of firewall rules that controls traffic for an EC2 Instance. Inbound and outbound rules can be edited during and after EC2 launch. Each EC2 Instance must have a Security Group assigned. Additionally, each Network Interface must have a Security Group assigned. Multiple Security Groups can be used to apply distinct sets of rules to individual Interfaces. The default Security Group added by AWS only allow traffic within a VPC.

The Security Group(s) assigned to the NetScaler SD-WAN AMI and its interfaces should accept SSH, ICMP, HTTP, and HTTPS. The

Security Group assigned to the WAN interface must also accept UDP on port 4980 (for Virtual Path support). Refer to AWS help for additional detail on Security Group configuration information.

Important

Wait two hours if provisioned from a new account and then retry

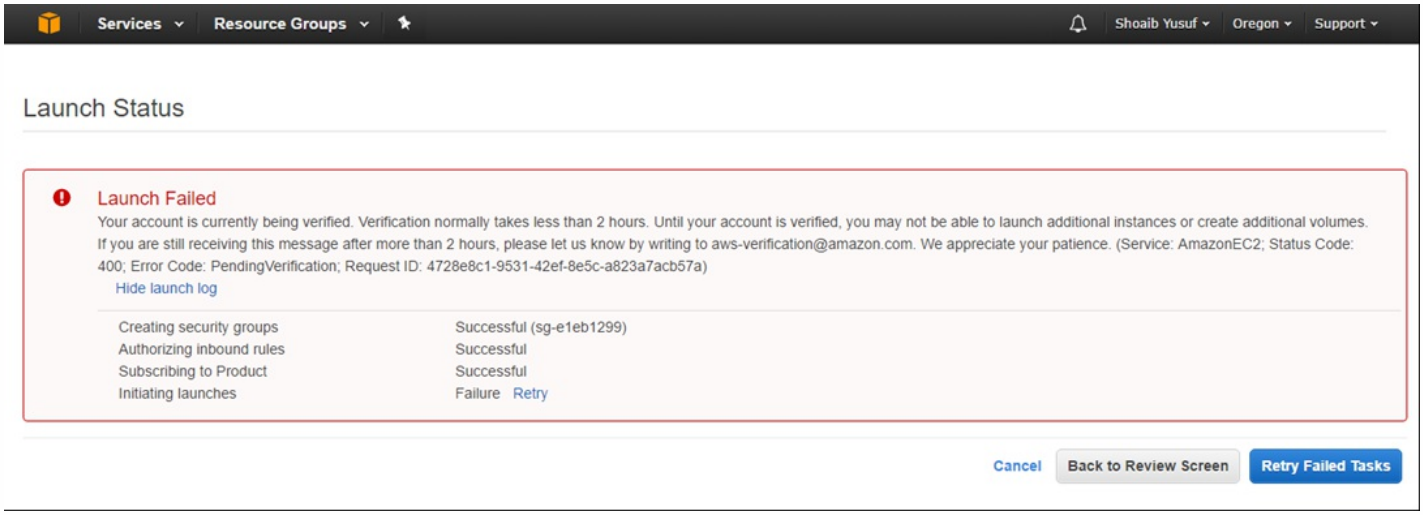

17. Navigate back to your **AWS Console: EC2 Dashboard**.

18. From the tool bar, under **Network & Security** select Network Interfaces, highlight the Mgmt interface and Edit the Name tag to give the interface a useful name.

19. Then click **Create Network Interface** to create the LAN interfaces:

- Description: <a user-defined description for the interface>
- Subnet: <the subnet previously defined for the interface>
- Private IP: <the private IP for the interface previously defined during preparation>
- Security Group: <the appropriate security group for the interface>

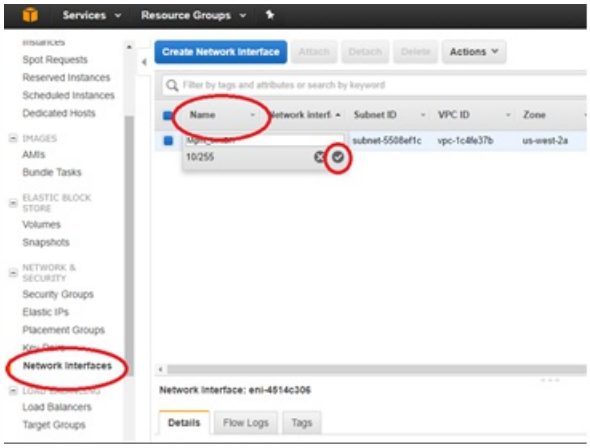

20. Repeat and click **Create Network Interface** to create the WAN interface.

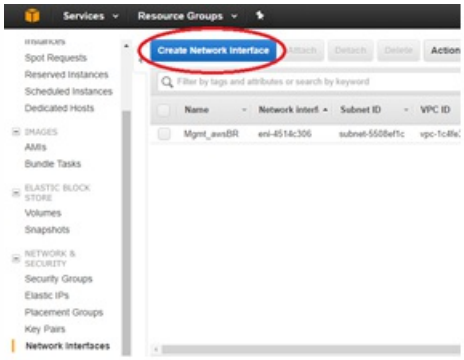

21. Edit the Name tag for each new interface and give a useful name.

22. Highlight the **Mgmt Interface** and select **Actions** > **Change Source/Dest**. Check to disable **Source/Dest**. Check, then select **Save**.

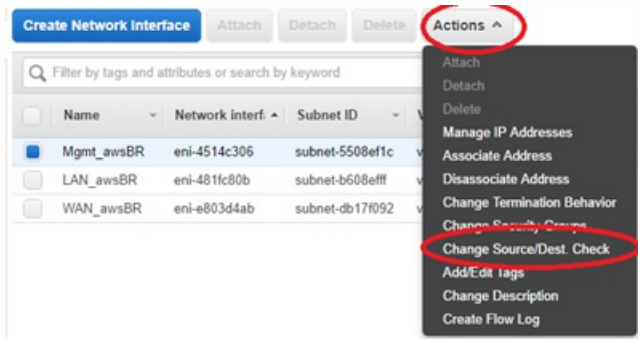

23. Repeat for LAN and WAN interfaces.

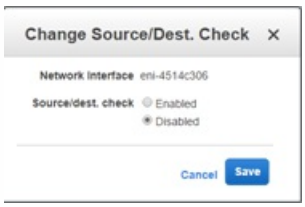

24. At this point all the Network Interfaces: **Mgm**t., **LAN**, and **WAN** each are configured with a **Name**, **Primary private IP**, and disabled for **Source/Dest. Check** attribute. Only the Mgmt. Network Interface will have a Public IP associated with it.

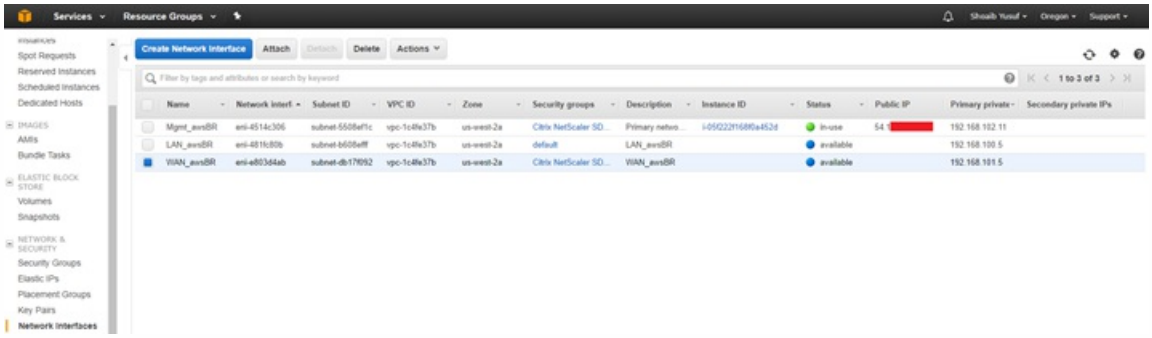

Important

Disabling the Source/Dest. Check attribute enables the interface to handle network traffic that is not specifically destined for the EC2 instance. As the NetScaler SD-WAN AMI acts as a go-between for network traffic, the Source/Dest. Check attribute must be disabled for proper operation.

The Private IPs defined for these Network Interfaces, ultimately, must match the IP addresses in your SD-WAN configuration. It may be necessary to define more than one Private IP for the WAN Network Interface if that interface will be associated with more than one WAN Link IP in the SD-WAN configuration for this site node. This can be accomplished by defining Secondary Private IPs for the WAN Interface as needed.

25. From the **EC2 Dashboard** tool bar, select **Instances**.

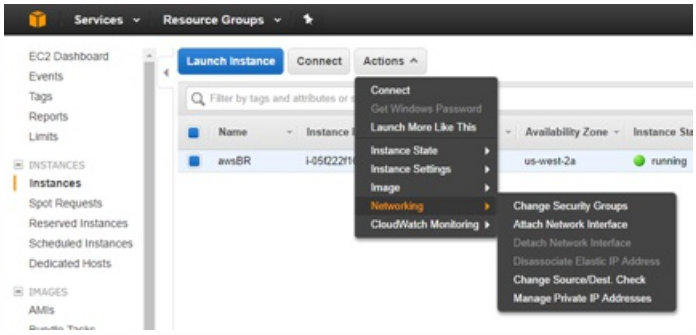

26. Highlight the newly created instance, then select **Actions** > **Networking** > **Attach Network Interface**.

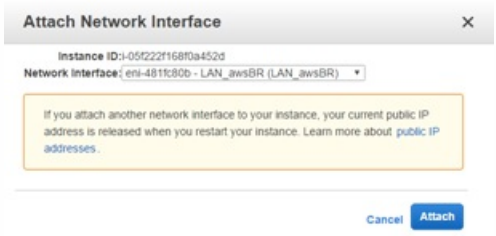

27. Attach first the LAN network interface and then the WAN network interface to the SD-WAN SE AMI.

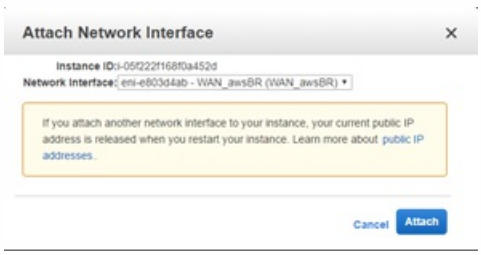

> **Note**
>
> Attaching the Mgmt, LAN, and WAN in that order attaches to eth0, eth1, eth2 in the SD-WAN AMI. This aligns with the mapping of the provisioned AMI and will ensure that interfaces are not reassigned incorrectly in the event of AMI reboot.

28. From the **EC2 Dashboard** tool bar, select **Elastic IPs (EIP)**, then click **Allocate new address**.

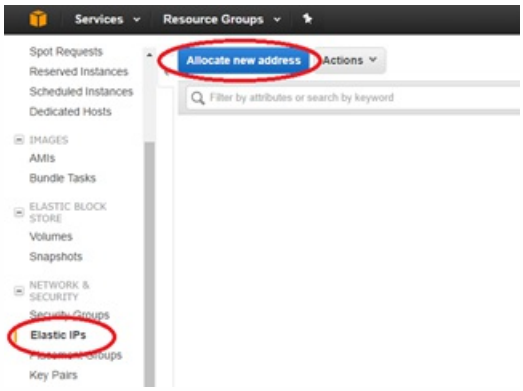

29. Click **Allocate** to allocate a new IP address, then **Close** after the New address request succeeded.

30. Highlight the new EIP and select **Action** > **Associate address** to associate the EIP with the Mgmt. Interface, then click **Associate**.

- Resource type: <network interface>
- Network interface: <previously created Mgmt. Network Interface>
- Private IP: <previously defined private IP for Mgmt>

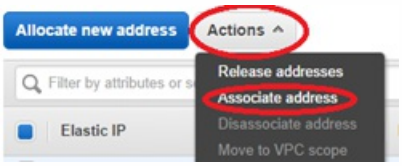

31. Repeat the process to associate another new EIP with the WAN interface.

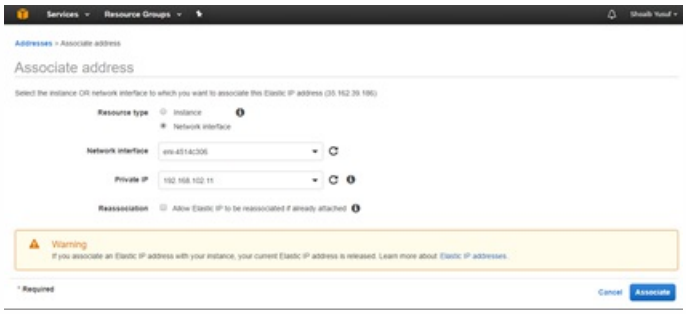

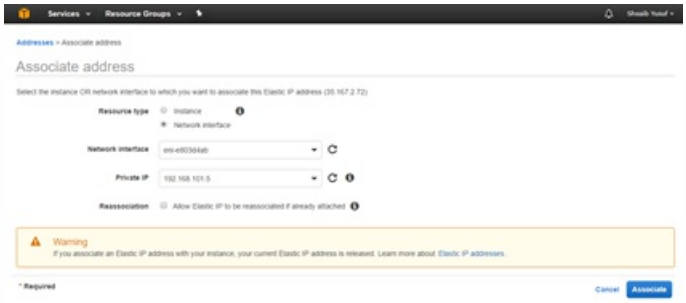

Configure SD-WAN SE AMI - SD-WAN Web Management Interface

To configure SD-WAN SE AMI:

1. At this point, you should be able to connect to the SD-WAN SE AMI's management interface using a web browser.

2. Enter the **Elastic IP (EIP)** associated with the Mgmt. Interface. You may need to create a security exception, if the security certificate is not recognized.

3. Login to the SD-WAN SE AMI using the following credentials:

- Username: *admin*
- Password: *<aws-instance-id>* (example; i-00ab111abc2222abcd)

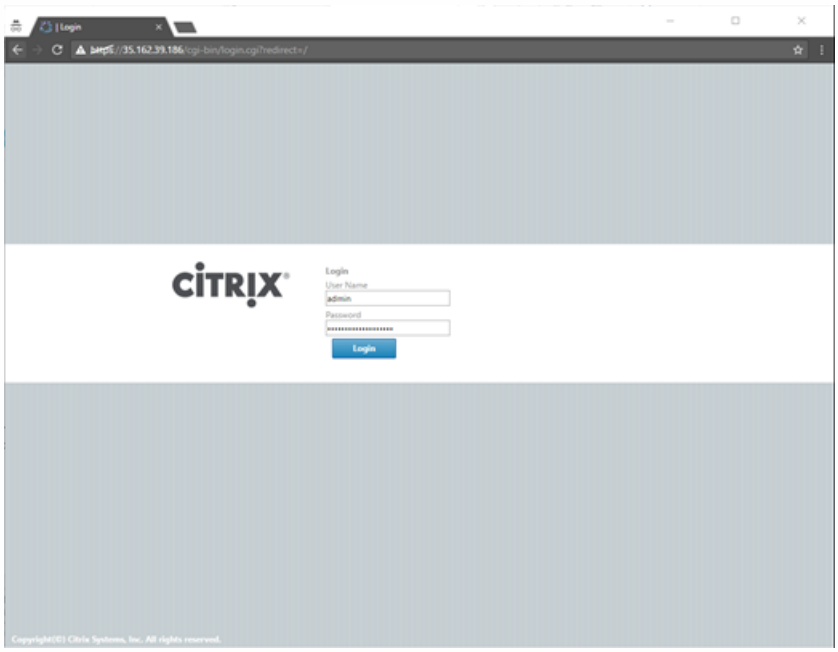

Note

If the Mgmt. Interface cannot be reached, check the following:

- Make sure the EIP is correctly associated with the Mgmt. interface
- Make sure the EIP responds to ping
- Makes sure the Mgmt. interface Route Table includes an Internet Gateway route (0.0.0.0/0)

- Make sure the Mgmt. interface Security Group is configured to allow HTTP/HTTPS/ICMP/SSH

Starting with release 9.1 SD-WAN AMI, users may also login to the SD-WAN AMI console using *ssh admin@<Mgmt. EIP>*, assuming that the key pair for the EC2 Instance has been added to the user's SSH key chain.

4. For SD-WAN SE **bring-your-own-license (BYOL) AMI**, a software license must be installed:

- On the SD-WAN web interface, navigate to **Configuration** > **Appliance Settings** > **Licensing**
- From **License Configuration**: **Upload License for this Appliance**, select **Choose File**, browse and open the SD-WAN SE AWS license, then click **Upload and Install**
- After successful upload, License Status will indicate State: Licensed

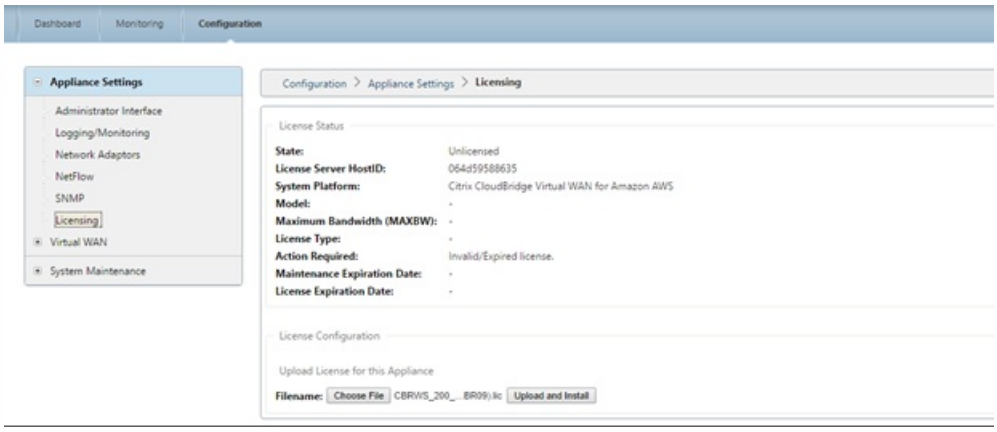

5. Set the appropriate **Data/Time** for the new AMI:

- On the SD-WAN web interface, navigate to **Configuration** > **System Maintenance** > **Date/Time Settings**
- Set the correct date and time using **NTP**, **Date/Time Setting** or **TImezone**

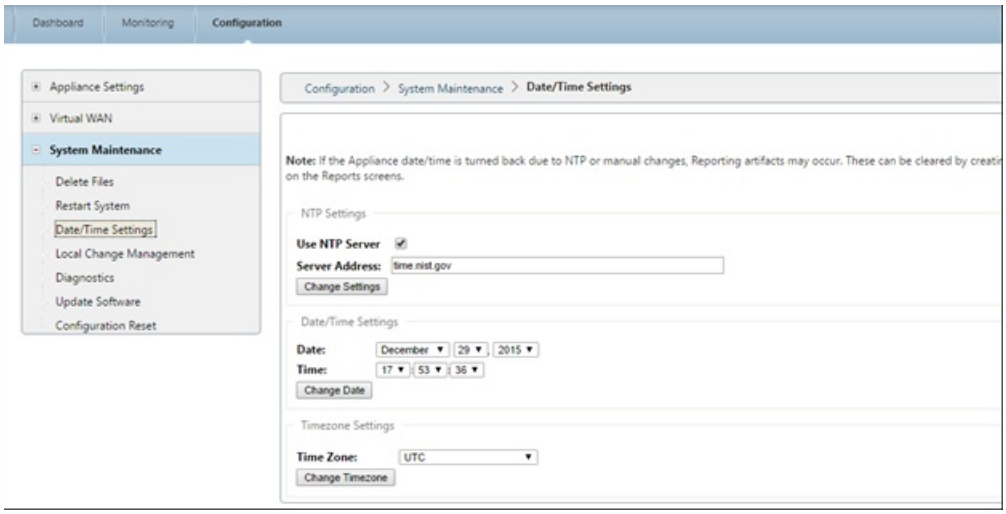

Note

The SD-WAN SE AMI Virtual WAN Service will remain disabled until an appliance package (Software + Configuration) is installed on the AMI.

Add SD-WAN SE AMI to your SD-WAN Environment

To add SD-WAN AMI to your SD-WAN envinroment:

1. Navigate to your **SD-WAN Center** or **Master Control Node** for your SD-WAN environment.

2. Add a **new site** node using the **Configuration Editor**:

- Add site: model CBAWS, Mode: client
- Interface Groups: awsLAN = eth1, awsWAN = eth2 (untrusted)
- Virtual IP Address: 192.168.100.5 = awsLAN, 192.168.101.5 = awsWAN

*with awsLAN virtual IP address being configured, the SD-WAN will advertise the LAN subnet of 192.168.100.5/24 as a local route to the SD-WAN Environment (refer to the **Connections** > <**AWSnode**> > **Routes**).

WAN Links:

- AWSBR-WAN with Access Type Public Internet, Autodetect Public IP if client node or configure the EIC for WAN link if MCN node, Access Interfaces: awsWAN 192.168.101.5 with gateway 192.168.101.1 (#.#.#.1 is typically the AWS reserved gateway).

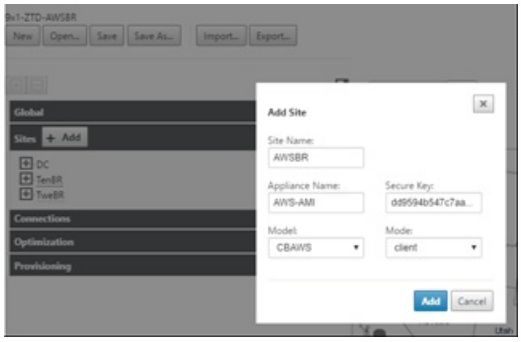

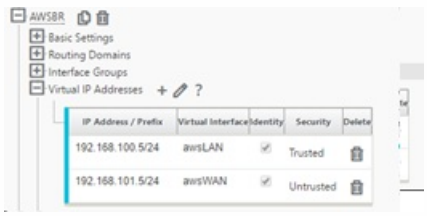

3. In the Configuration Editor validate the path association under **Connections** > **DC** > **Virtual Paths** > **DC-AWS** > **Paths**.

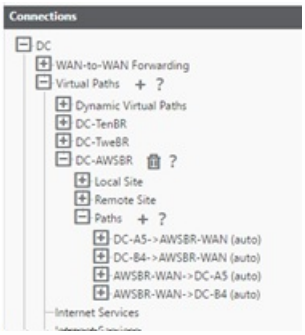

Note

The Virtual Path will be used across the AMI WAN interface to push software and configuration updates to the SD-WAN AMI instead of via direct connection to the Mgmt. interface.

Private IP addresses must be defined on EC2 WAN Network Interface for every WAN Link IP in the Configuration Editor. This can be accomplished by defining one or more Secondary Private IPs for the Network Interface as necessary.

Important

Recall the assigned mapping in AWS EC2 dashboard assigning Mgmt. to eth0, LAN to eth1 and WAN as eth2

Amazon reservices the first four IP addresses and the last IP address in each subnet CIDR block and can not be assigned to an instance. For example, in a subnet with CIDR block 192.168.100.0/24, the following five IP addresses are reserved:

- 192.168.100.0 : Network address
- 192.168.100.1 : Reserved by AWS for the VPC router
- 192.168.100.2 : Reserved by AWS for the DNS server
- 192.168.100.3 : Reserved by AWS for furture use
- 192.168.100.255: Network broadcast address, which is not supported in a VPC

4. **Save and Export** the newly created SD-WAN configuration and export to the **Change Management Inbox**.

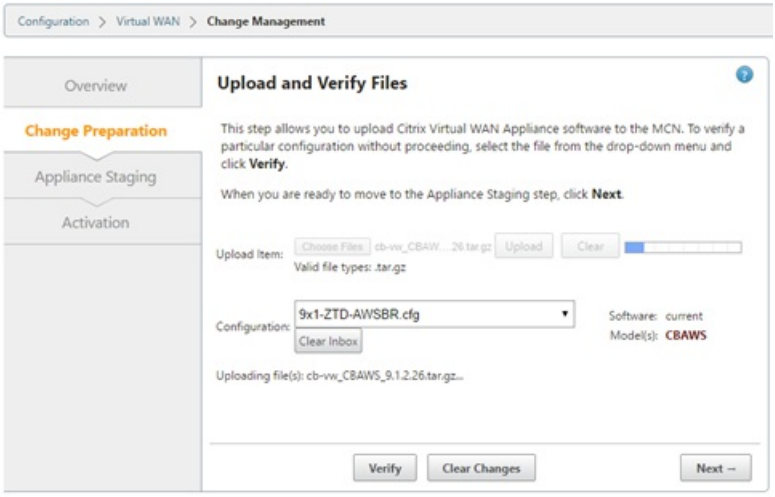

5. Navigate to the MCN **Change Management** and run through the change management process to push the latest configuration to the SD-WAN environment informing all existing SD-WAN nodes of the newly added AWS node and the subnets (virtual interfaces) associated with it. Make sure to upload the software package specific to CBAWS in the Change Preparation step that matches the current software used by the existing SD-WAN environment.

6. From the **Change Management** page, download the package generated specifically for the new AWS node using the **active** link.

Site-Appliance	Model	State	Currently Active		Currently Staged		Traffic Interruption		Download Package
			Software	Config	Software	Config	Expected	Actual	
DC-VPX	CBVPX		9.1.2.26.561201	13:17 on 12/29/16	9.1.2.26.561201	21:15 on 12/28/16	<1 min	582 ms	active / staged
AWSBR-AWS-AMI	CBAWS		Not Connected				Loc Chg Mgt		active / none
TenBR-EE1000	CB1000		9.1.2.26.561201	13:17 on 12/29/16	9.1.2.26.561201	21:15 on 12/28/16	<1 sec	170 ms	active / staged
TweBR-SE400	CB400		9.1.2.26.561201	13:17 on 12/29/16	9.1.2.26.561201	21:15 on 12/28/16	<1 sec	44 ms	active / staged

7. Navigate back to the SD-WAN SE AMI's management interface using the assigned EIP for the Mgmt. interface.

8. Navigate to **Configuration** > **System Maintenance** > **Local Change Management**.

9. Click **Choose File** to browse and **Upload** the active AWS software/config package recently downloaded.

10. After successful **Local Change Management**, the web interface should auto-refresh with the latest installed software, with the **Virtual WAN Service** still disabled.

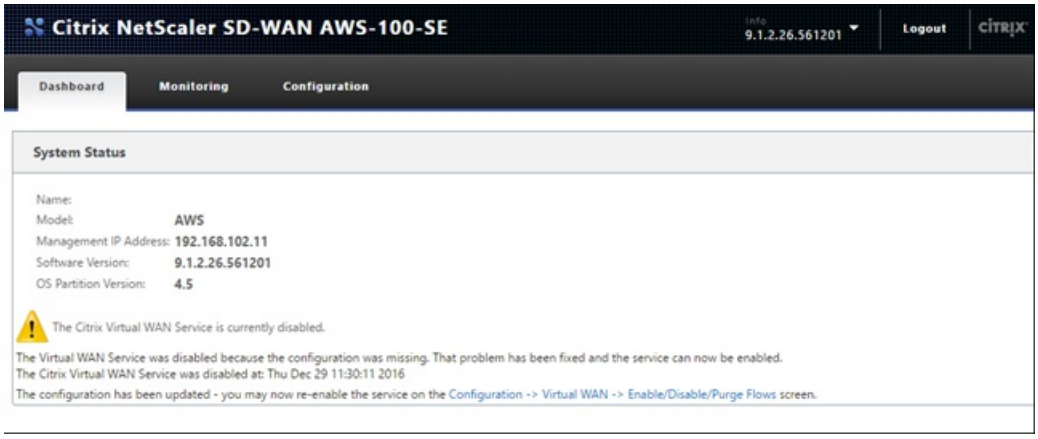

11. On the SD-WAN SE AMI navigate to **Configuration** > **Virtual WAN Enable/Disable/Purge Flows** and enable the service using the **Enable** button.

12. Upon successful connectivity on the WAN interface, the SD-WAN will report Good Path State on the **Monitoring** > **Statistics** > **Paths** page.

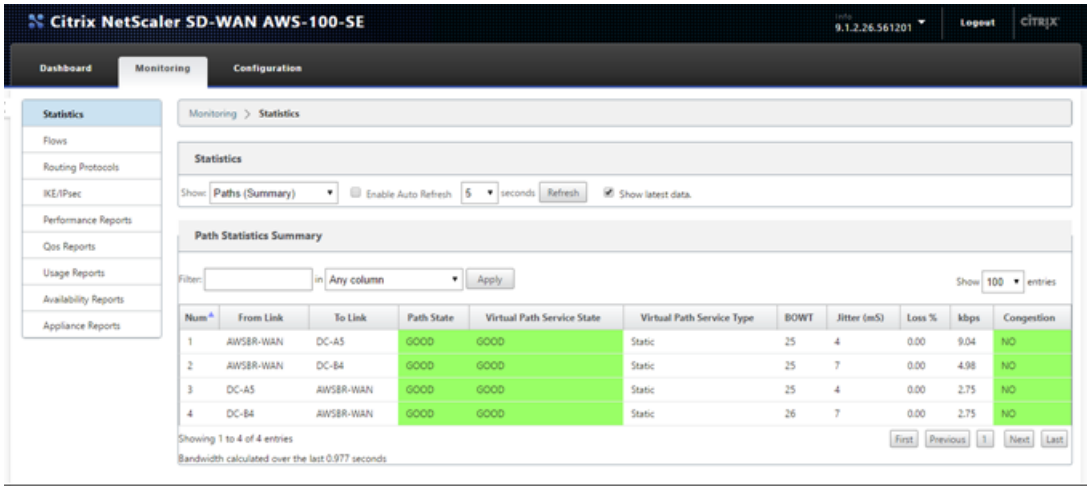

Troubleshooting

The correct private Internet Web Gateway (IWG) IP must be used in the SD-WAN Access Interface configuration

- If an incorrect IWG is used in the Configuration Editor to define the WAN Link for the AWS Site (Virtual IP Address and the correct Gateway) then Virtual Path will fail to establish.
- A quick way to check if the IWG is incorrectly configured is to check the SD-WAN ARP table.

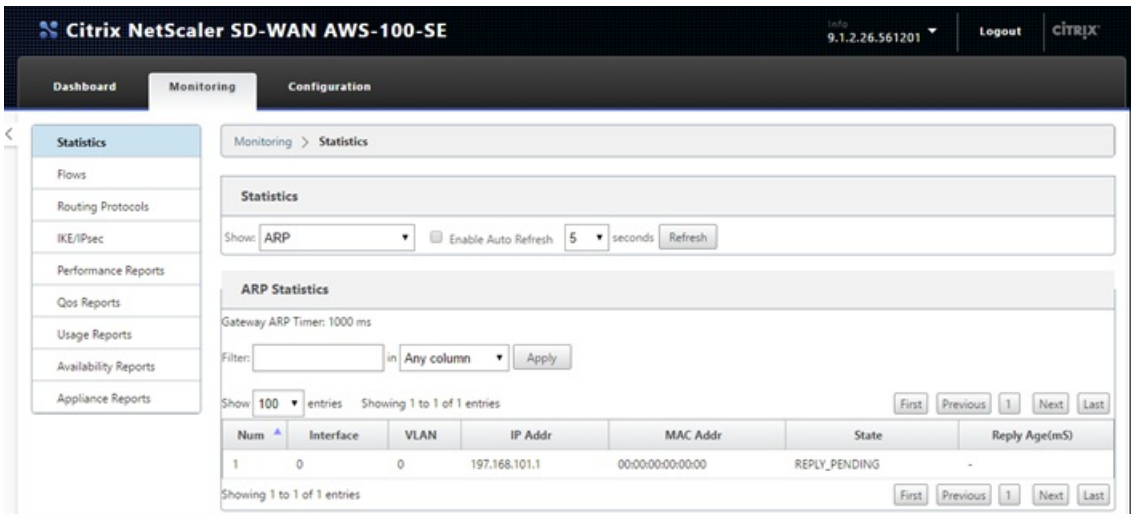

SD-WAN built in Packet Capture tool can help confirm proper packet flow

1. Navigate to the **Configuration** > **System Maintenance** > **Diagnostic** page of the SD-WMA AMI.

2. Select the **Packet Capture** tab, and set the following settings, then click **Capture**:

- Interfaces: 2 //in order to capture on eth2 which was associated with the WAN interface

3. The capture output on the web page should show the UDP probe packets leaving the SD-WAN SE AMI with the WAN VIP / Private IP as the source, with a destination of the Static Public IP(s) used for the MCN, also the returning UDP packet with source of the MCN Static Public IP and destination of the local VIP/Private IP (which was NAT'd by the IWG).

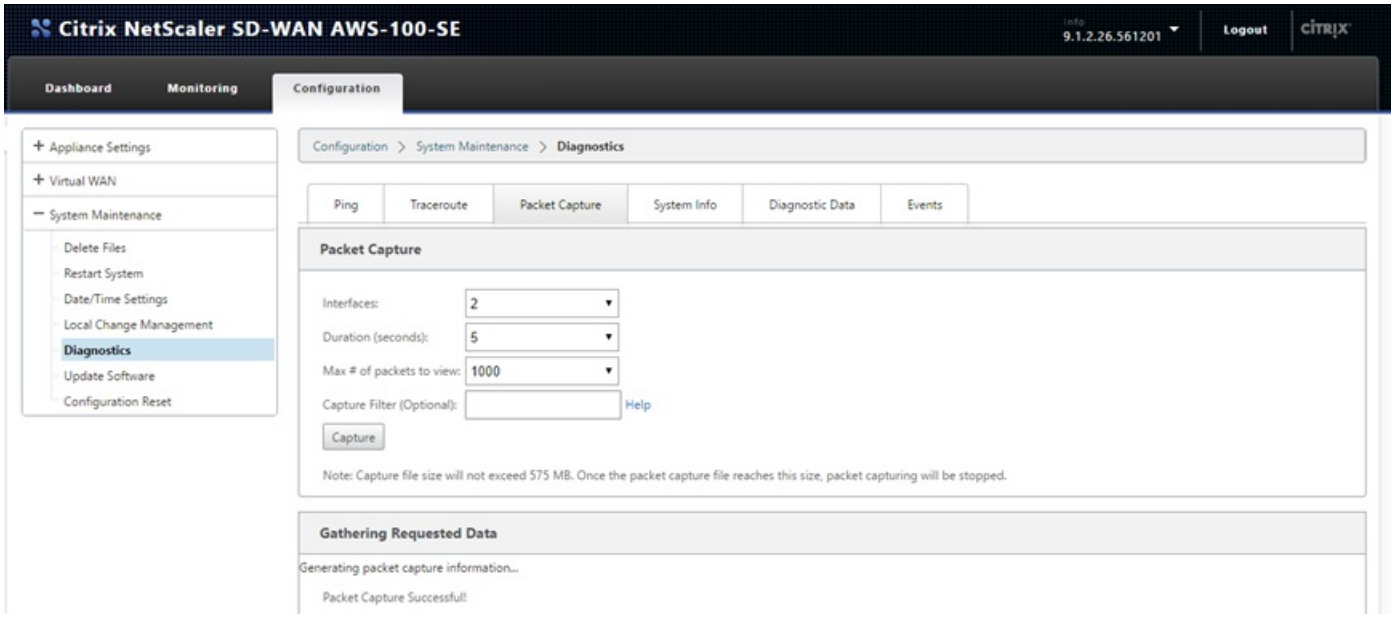

Note

This can typically occur when an IP address is accidently created outside of the CIDR block assigned to the VPC.

Installing SD-WAN VPX Standard Edition on Azure

May 03, 2018

NetScaler SD-WAN Standard Edition is now available in Azure marketplace on BYOL basis. NetScaler SD-WAN Standard Edition for Azure logically bonds multiple network links into a single secure logical virtual path. The solution enables organizations to use connections from different service providers including Broadband, MPLS, 4G/LTE, Satellite, and point-to-point links to get high resiliency virtual WAN paths. NetScaler SD-WAN for Azure enables organizations to have a direct secure connection from each branch to the applications hosted in Azure eliminating the need to backhaul cloud bound traffic through a data center. Some of the benefits of using NetScaler SD-WAN in Azure are:

- Create direct connections from every location to Azure.
- Ensure an always on connection to Azure.
- Extend your secure perimeter to the cloud.
- Evolve to a simple, easy to manage branch network.

Topology - SD-WAN in Azure

NetScaler SD-WAN Standard Edition can only be deployed in Gateway deployment mode in Azure. A Public IP address (Static/Dynamic) can be assigned to both the management and WAN facing interface of SD-WAN in Azure. The new template in Azure allows multiple Public IP assignments, one for the SD-WAN management interface and another for the WAN facing interface.

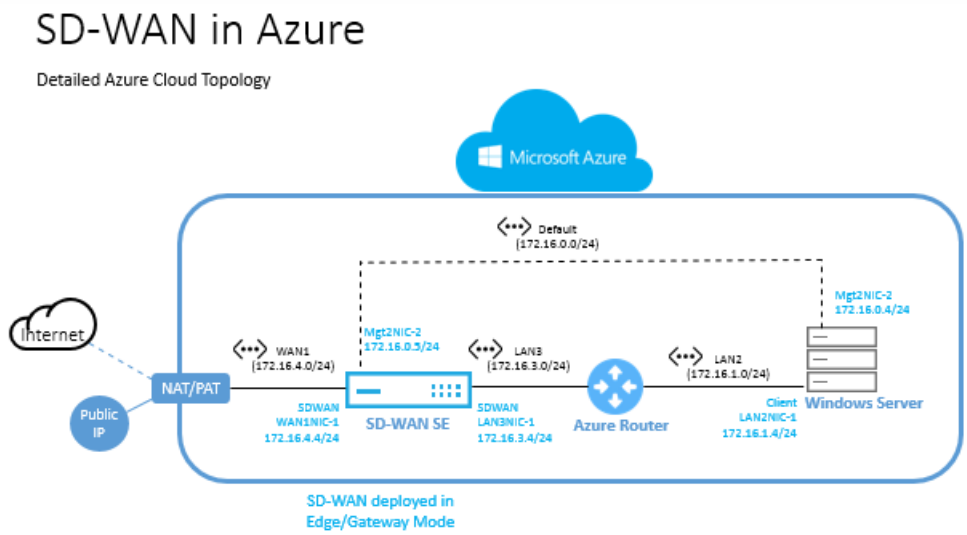

Use Case

An Azure VM is deployed within a specified region and can be connected to multiple branch locations through MPLS, Internet, or 4G/LTE. Within a Virtual Network (VNET) infrastructure, SD-WAN Standard Edition VHD (Virtual Hard Drive) is deployed in Gateway mode. The VNET has routes towards the Azure Gateway. SD-WAN instance has a route towards the Azure Gateway for internet connectivity.

Connectivity between Data Center, Branch, and Cloud is achieved by using different transports methods utilizing multiple WAN paths simultaneously.

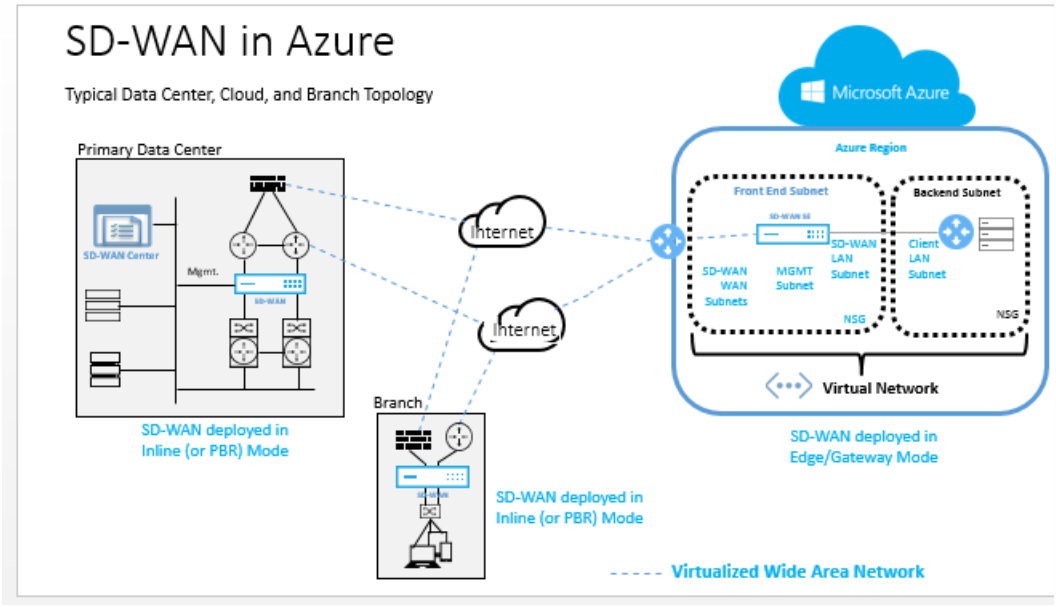

To deploy SD-WAN Standard Edition in Microsoft Azure:

1. In a web browser, type https://portal.azure.com. Log into Microsoft Azure account. Then, search for Citrix SD-WAN Standard Edition.

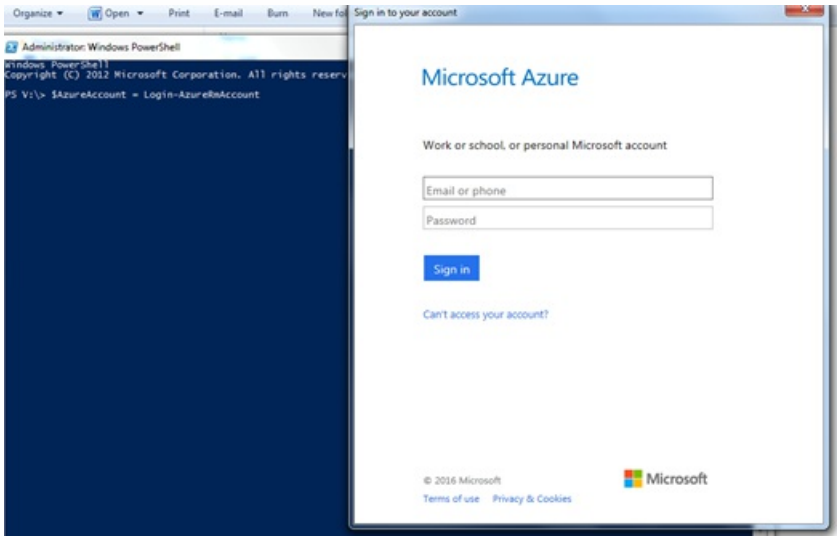

2. In the search results window, choose the solution as shown below:

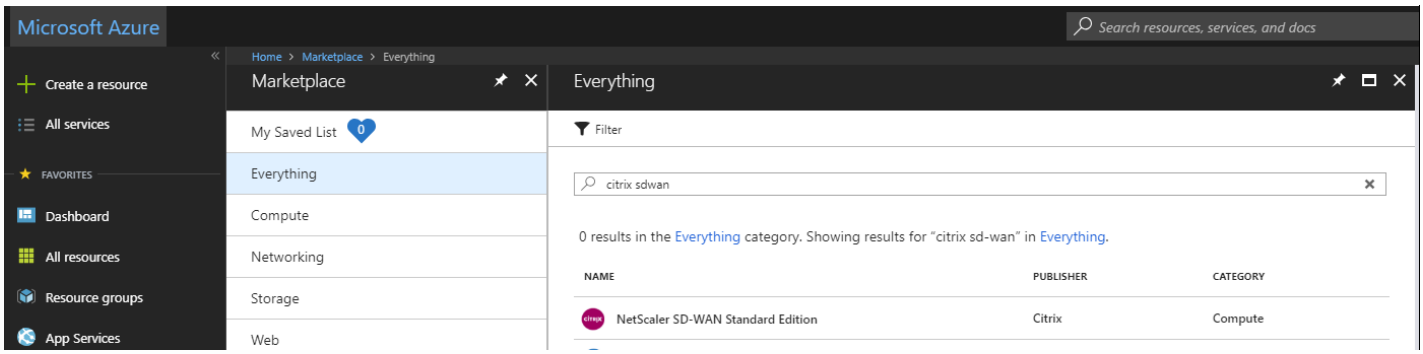

3. Click **create** after going through the description and making sure the solution chosen is correct.

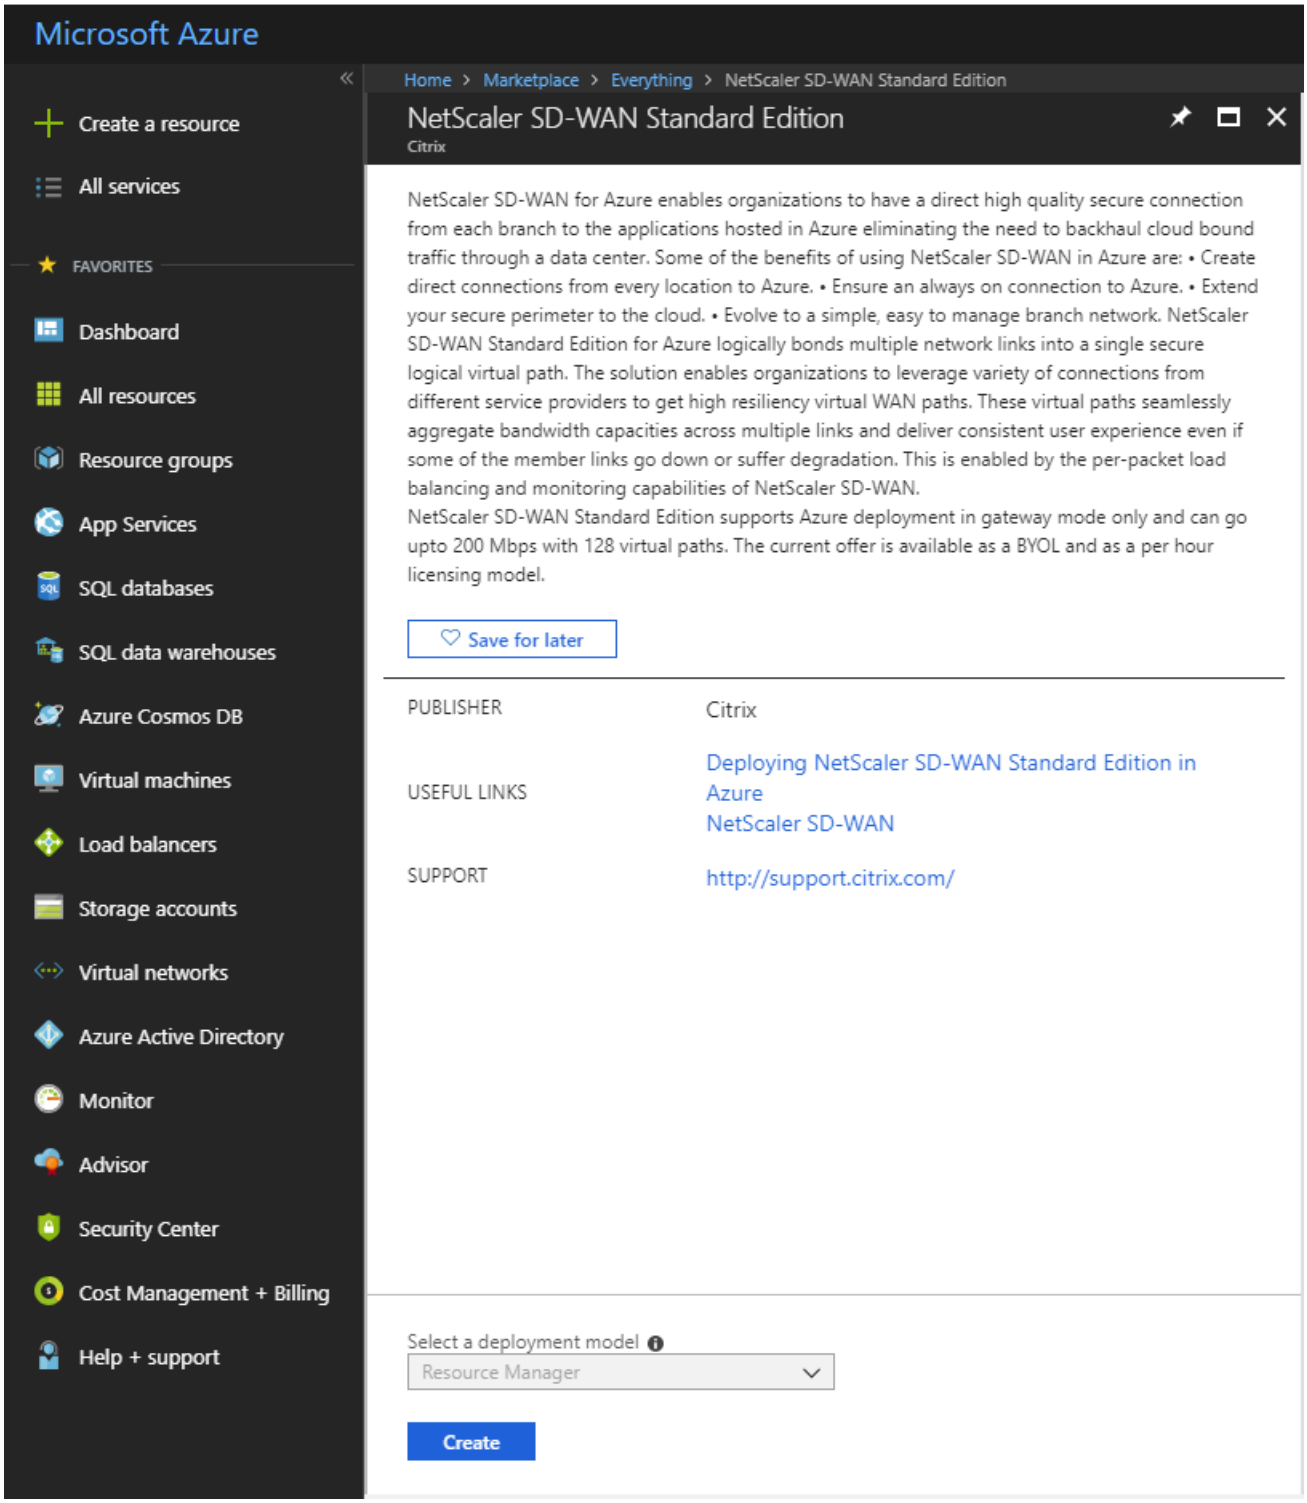

4. After you click **Create,** a wizard prompting for details necessary to create the virtual machine in Azure appears. In the first step, choose the resource group in which you like to deploy the solution. A resource group is a container that holds related resources for an Azure solution. The resource group can include all the resources for the solution, or only those resources that you want to manage as a group. You can decide how you want to allocate resources to resource groups based on your deployment. Some important points to consider when defining your resource group are:

- a) All the resources in your group should share lifecycle. You deploy, update, and delete them together.
- b) If one resource, such as a database server, needs to exist on a different deployment cycle it should be in another resource group.

- c) Each resource can only exist in one resource group.
- d) You can add or remove a resource to another resource group at any time.
- e) You can move a resource from one resource group to another resource group
- f) A resource group can contain resources that reside in different regions.
- g) A resource group can be used to scope access control for administrative actions.
- h) A resource can interact with resources in other resource groups. This interaction is common when the two resources are related but do not share lifecycle (for example, web apps connecting to a database).

In the following image, choose **Create New is** chosen and a name for the service group is provided.

Under **Location,** choose the region in which you want to deploy the solution. When creating a resource group, you need to provide a location for that resource group. The resource group stores metadata about the resources that you are creating. Therefore, when you specify a location for the resource group, you are specifying where that metadata is stored.

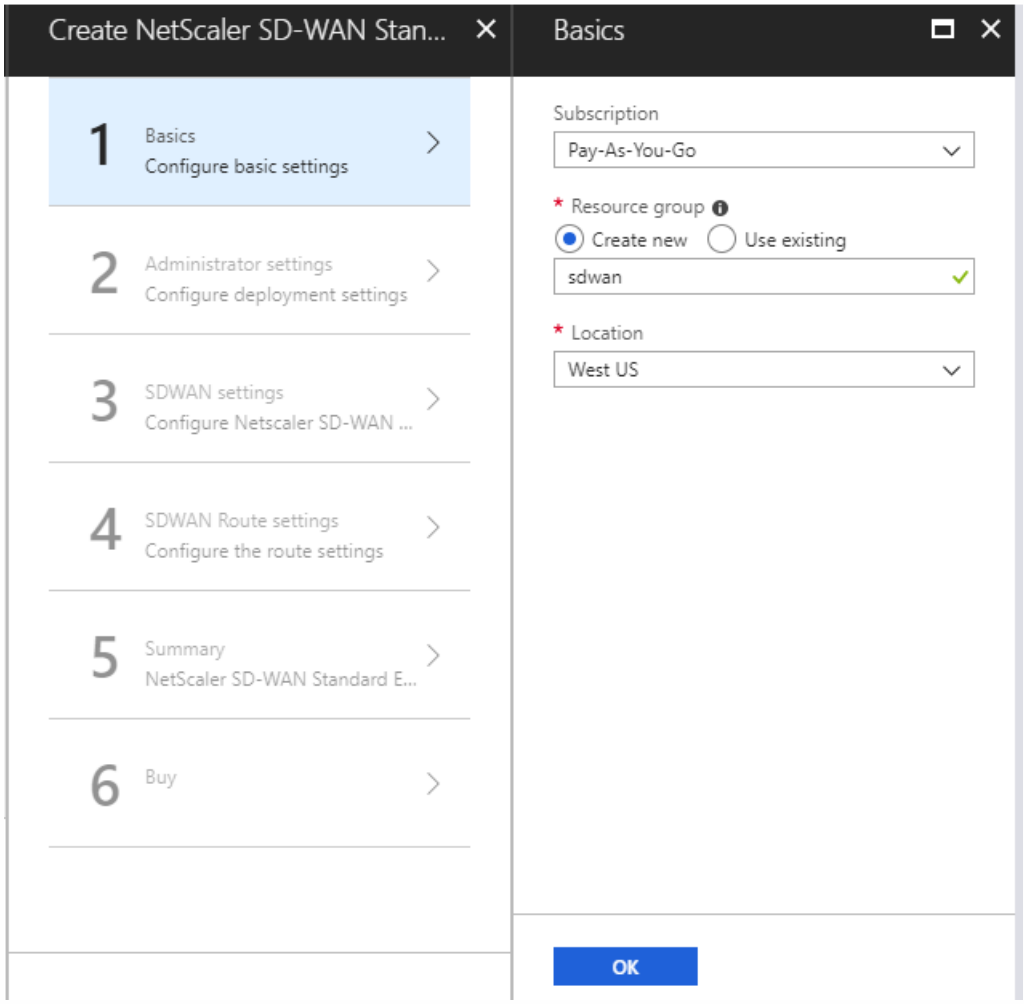

5. Provide a name for the Virtual Machine. Leave **HA Deployment Mode** disabled. See, SD-WAN Standard Edition Virtual Appliance (VPX) HA Support for Microsoft Azure documentation for HA deployment instructions. Choose a username and strong password. The password must consist of an upper case letter, special character and must be more than nine characters. Click **OK**.

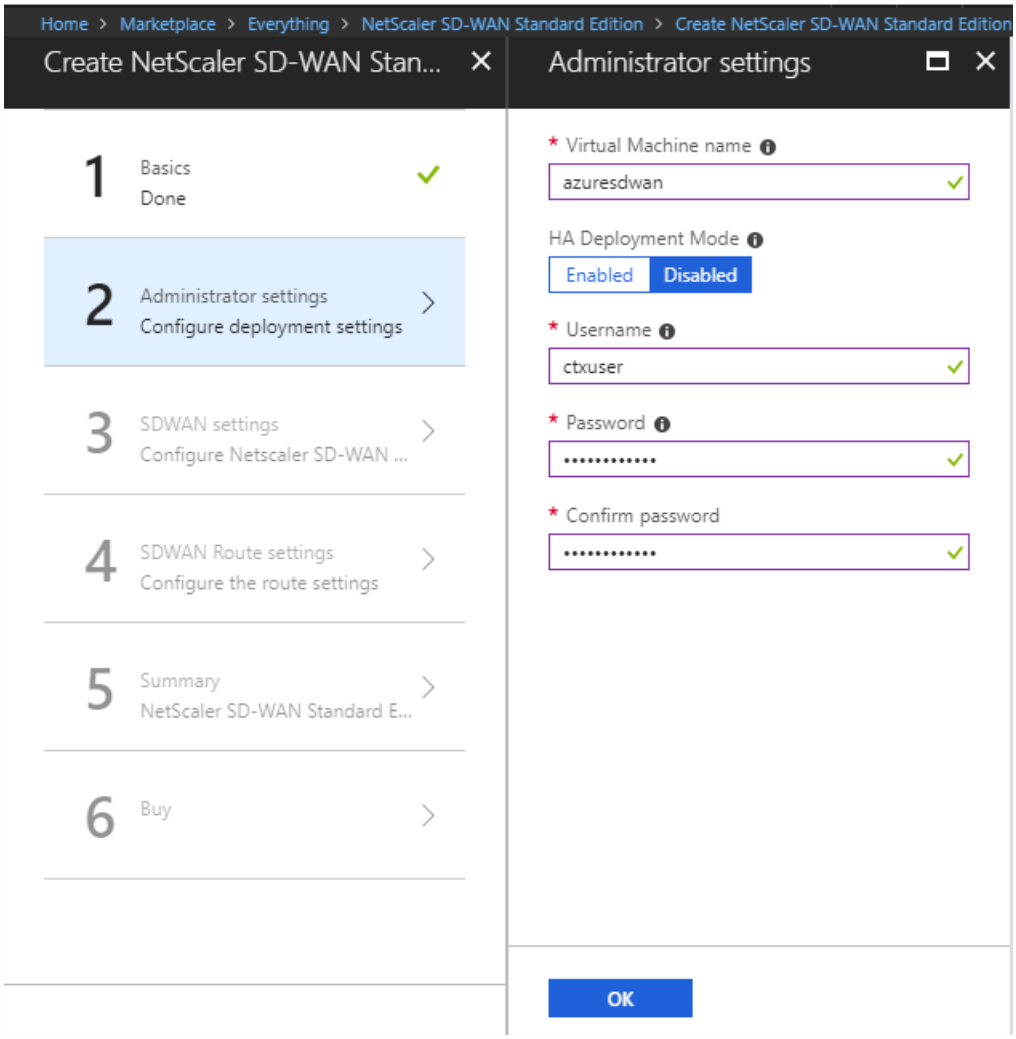

6. Choose the VM size in which you want to run the image. The Standard_D3_V2 instance is applicable to the VPX-SE appliance only. The Standard_D4_V2, F8, F16 instance types are supported for the VPXL-SE appliance. The selected site device in the SD-WAN configuration should match with this selection. If the Azure instance is intended to be a client node, then the VPX-SE (Standard D3 V2) is sufficient. If you are looking to promote the Azure instance as the MCN, then the larger VPX-L is required enabling up to 128 Virtual Paths.

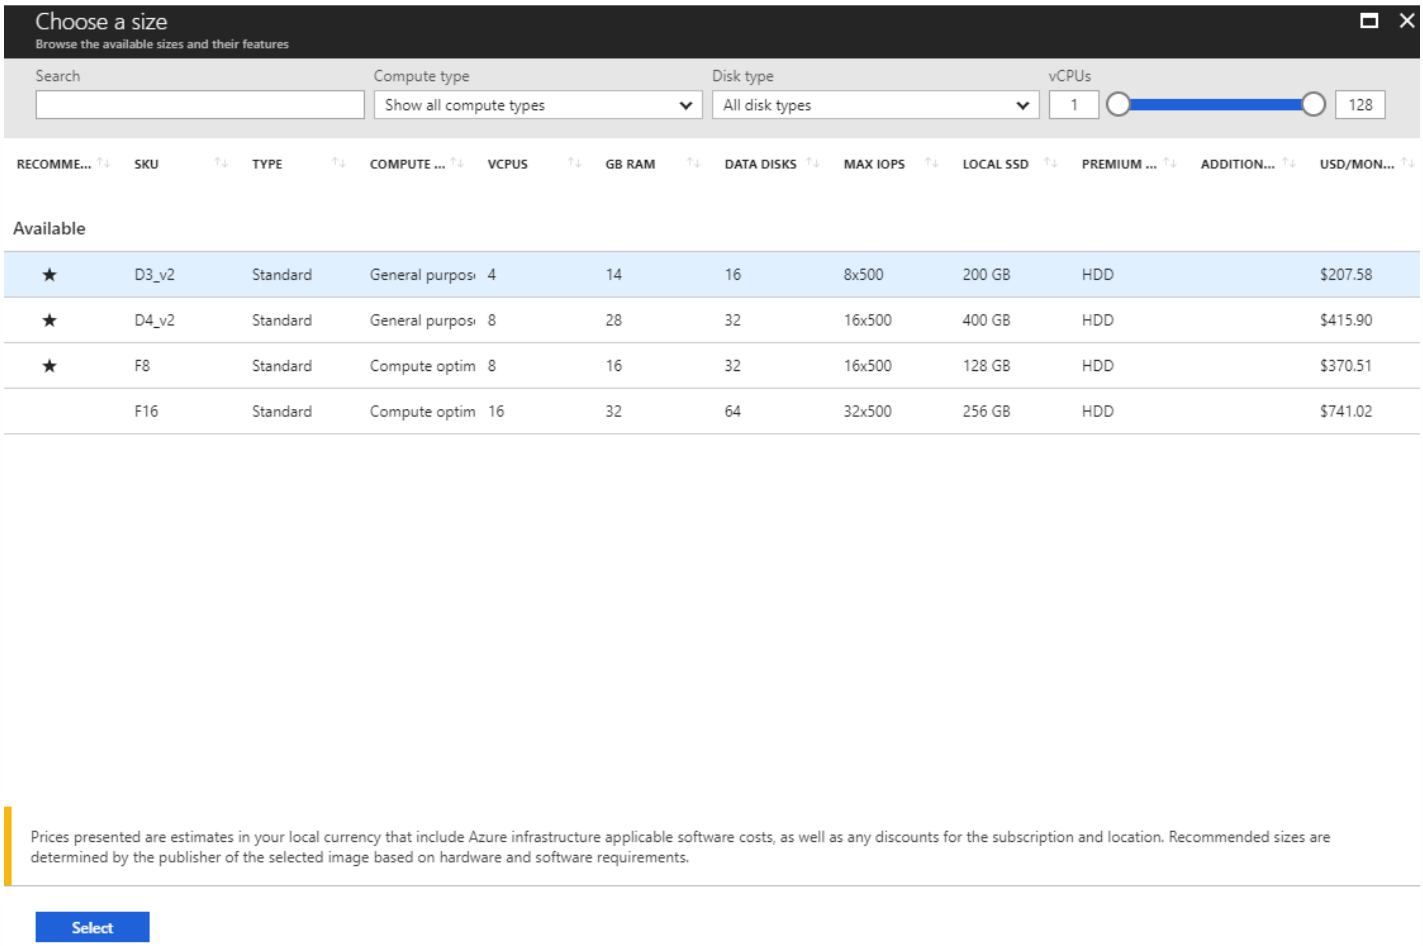

Next select the SD-WAN SKU. Below the “Bring Your Own License” option is selected.

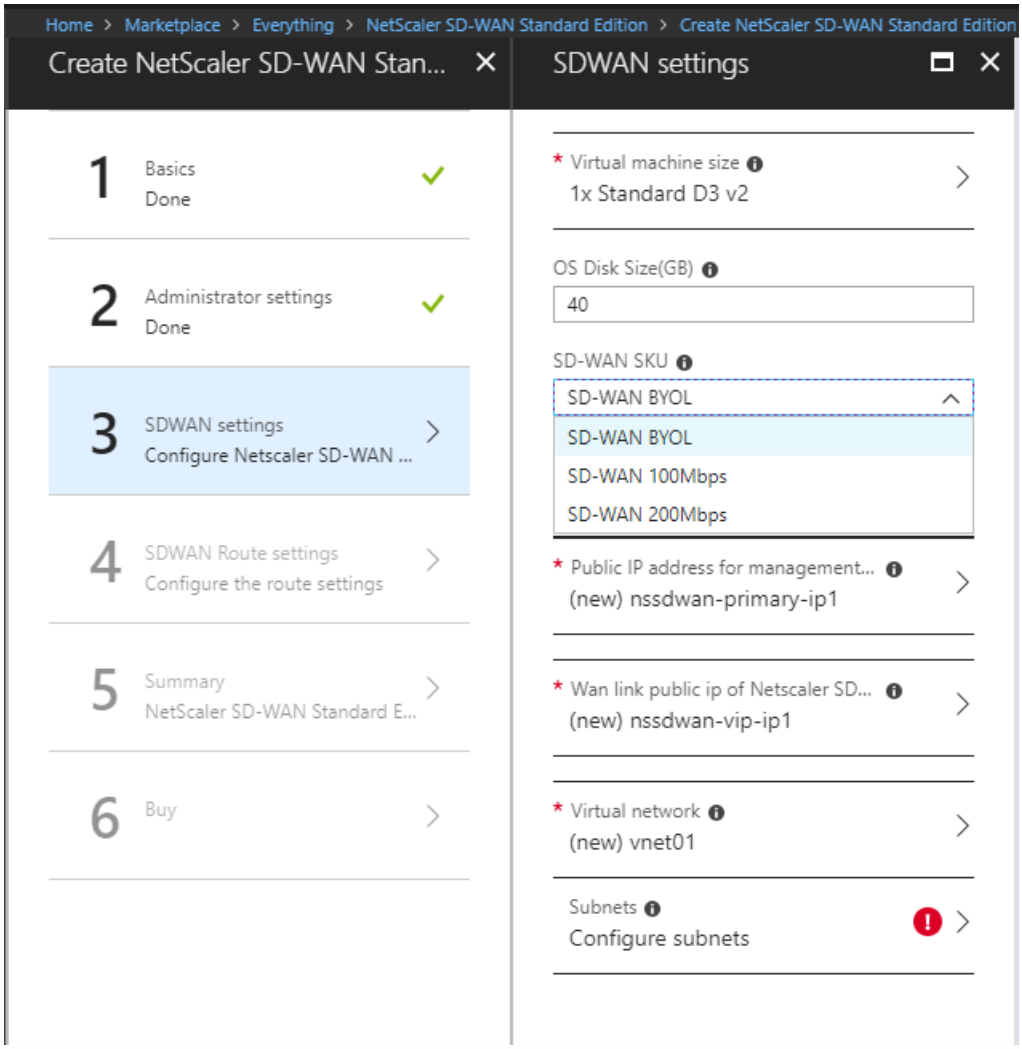

7. After choosing the image, choose the storage account. If you have an existing storage account, then you may choose that. In this step, you are creating a storage account as seen below. A storage account stores the VHDs for the operating system's temporary and more data disks.

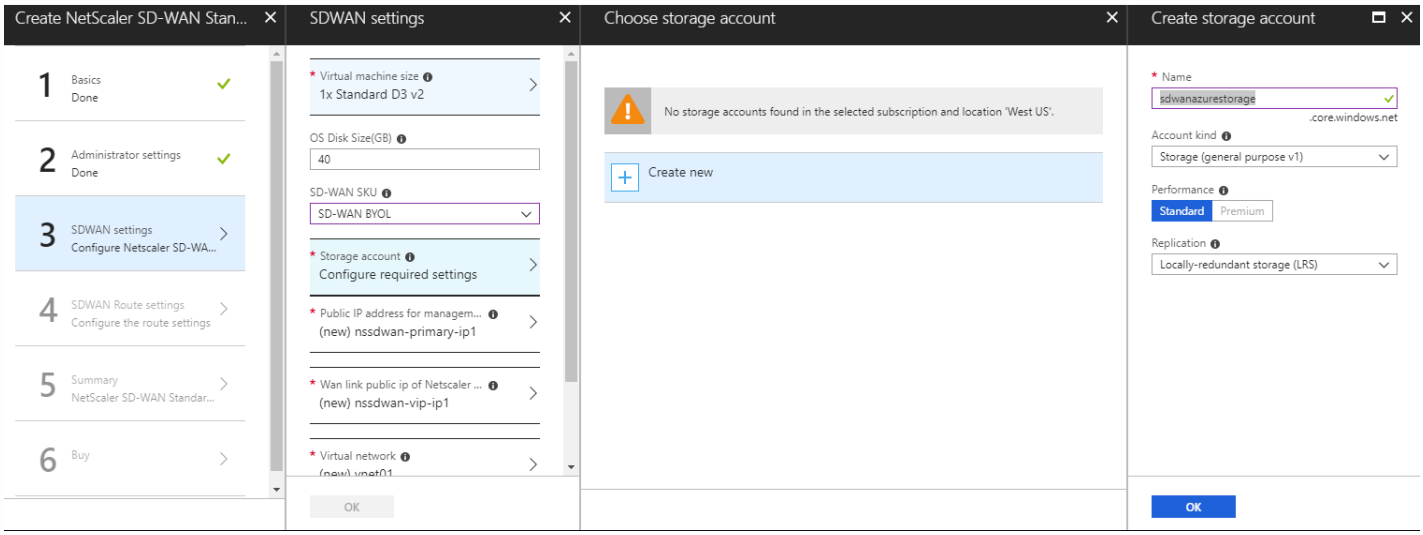

8. Next, choose a Public IP for the Virtual Machine management interface. You may choose either a static or dynamic IP. However, it is recommended that you choose a static IP.

Then, select the WAN link public IP. This is the public IP that is assigned to WAN interface of the appliance and is the IP used to establish Virtual Path connectivity.

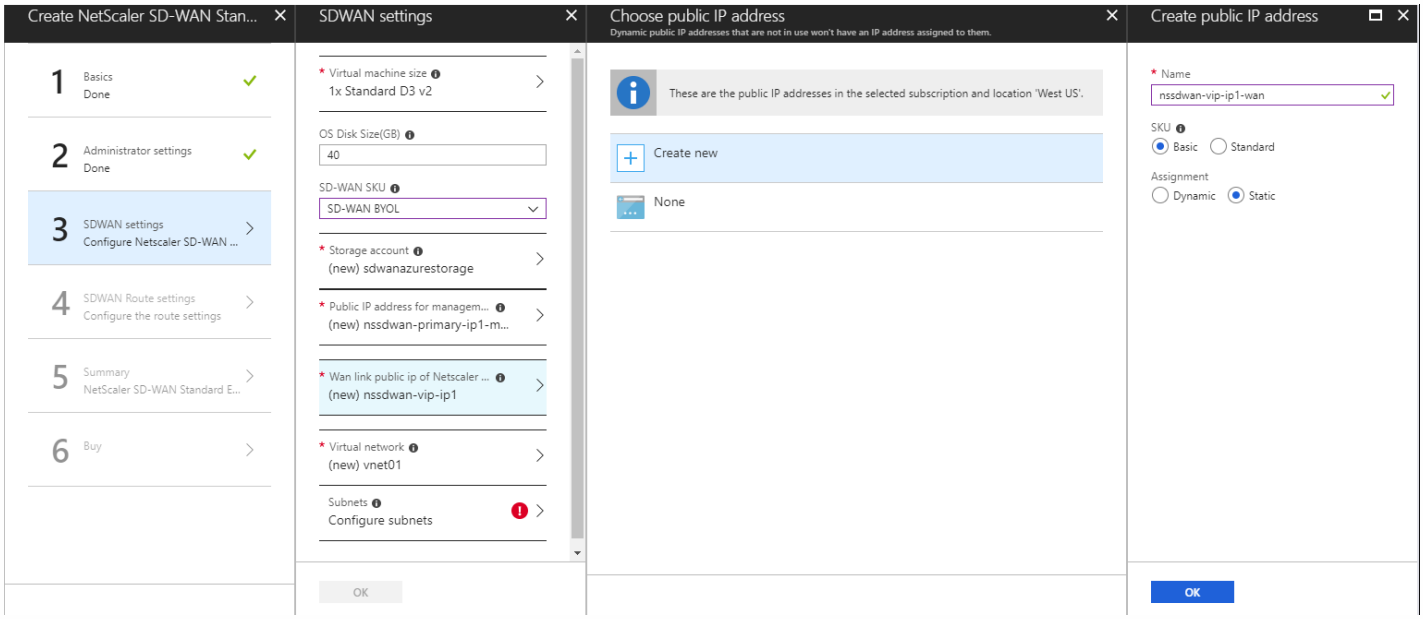

9. Create a new Virtual Network (VNET) or use an existing VNET. This is the most critical step for deployment as this step chooses the subnets to be assigned to the interfaces of the SD-WAN VM.

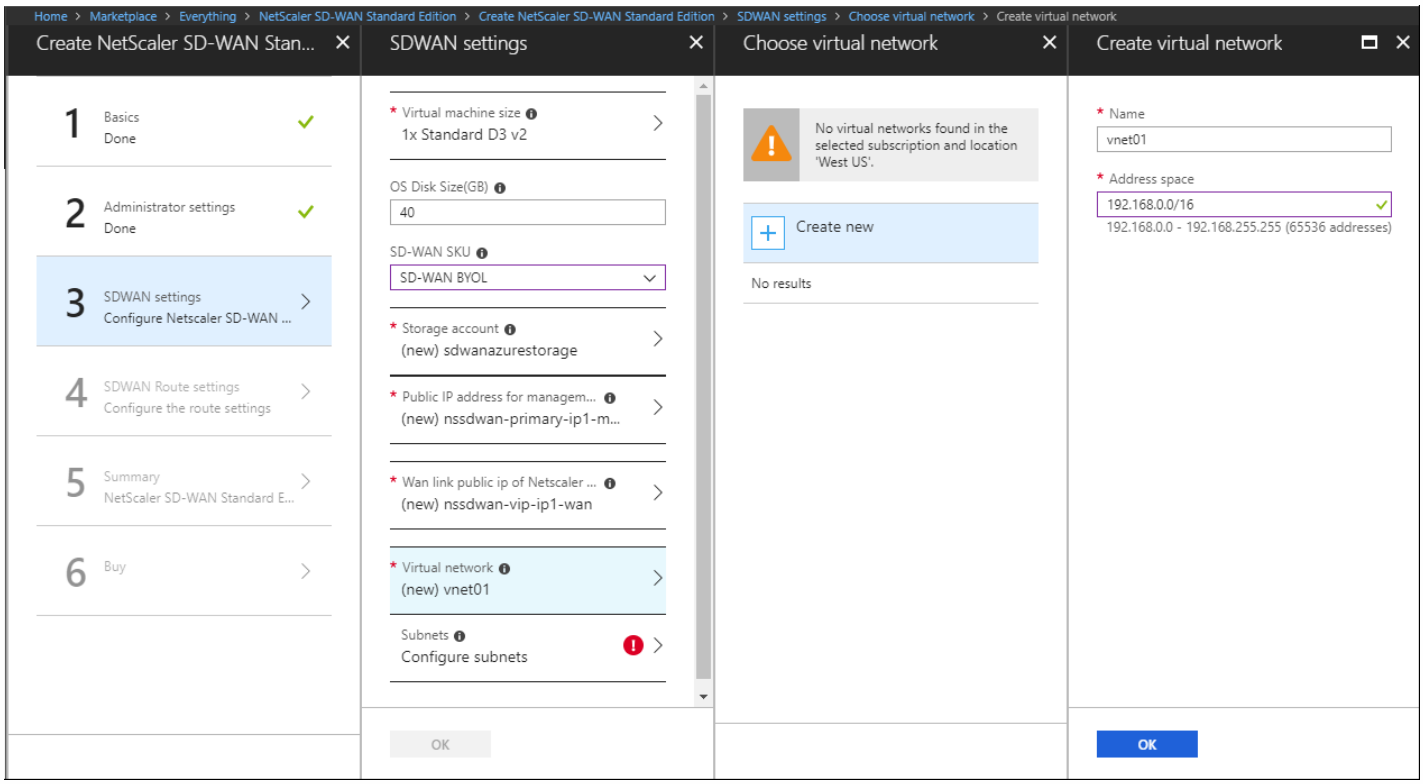

10. You can now configure subnets to assign the required subnets to each of the interfaces in the VM. Configure as required to match the running configuration in the defined SD-WAN site associated with the Azure deployment, and click **OK**. In the following image, you are assigning the subnets to each of the interfaces.

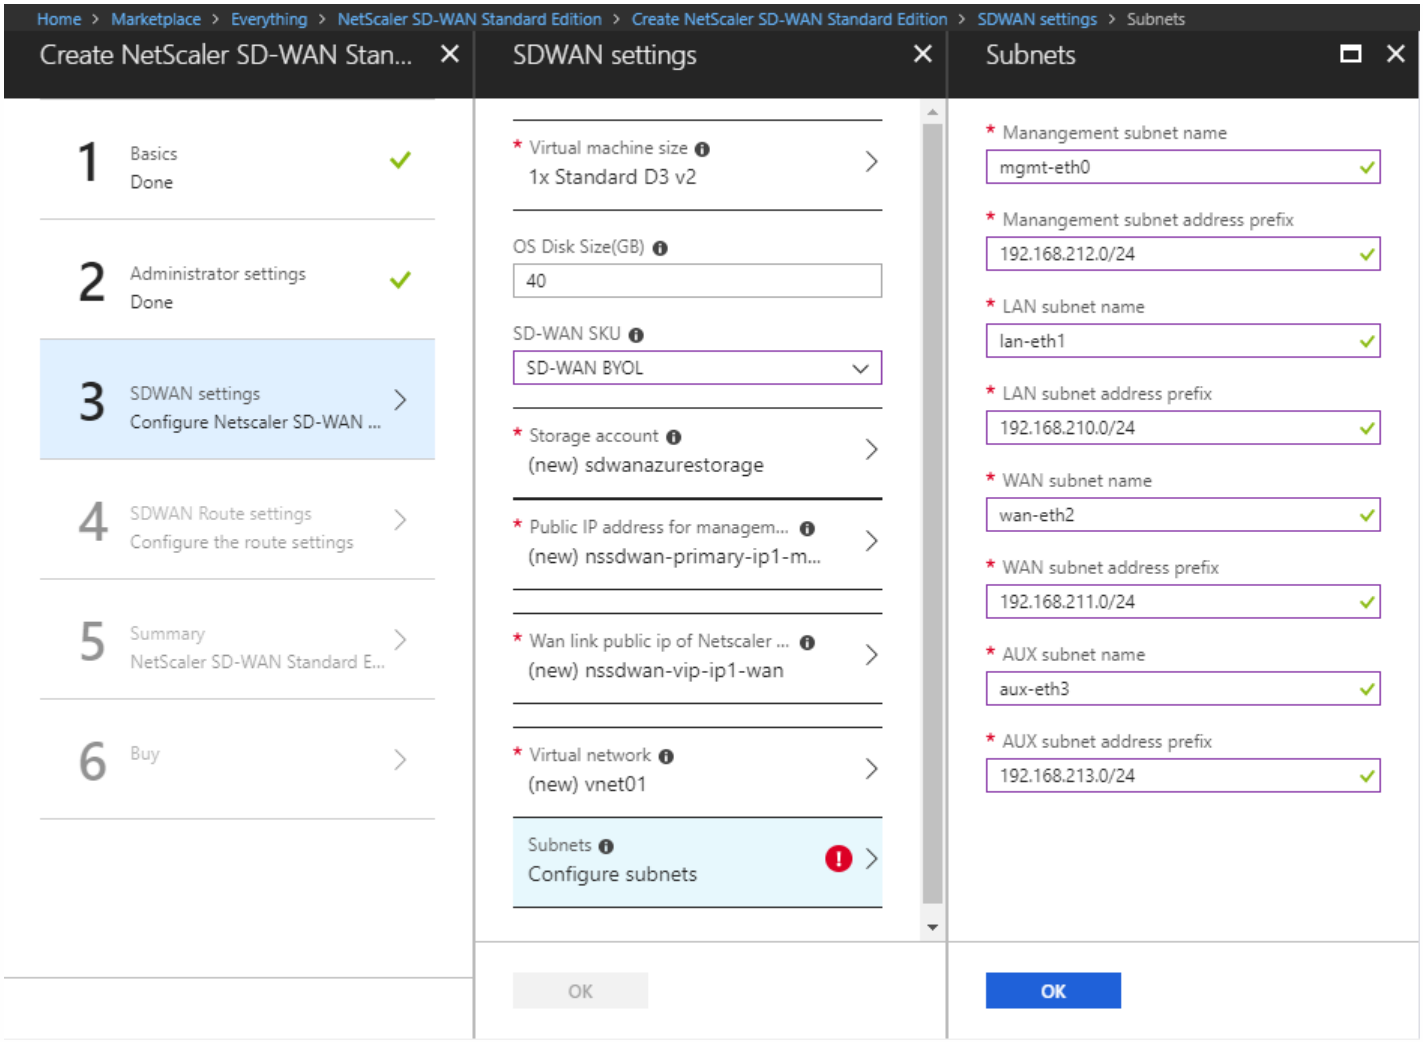

11. Verify all the configuration details and click **OK**.

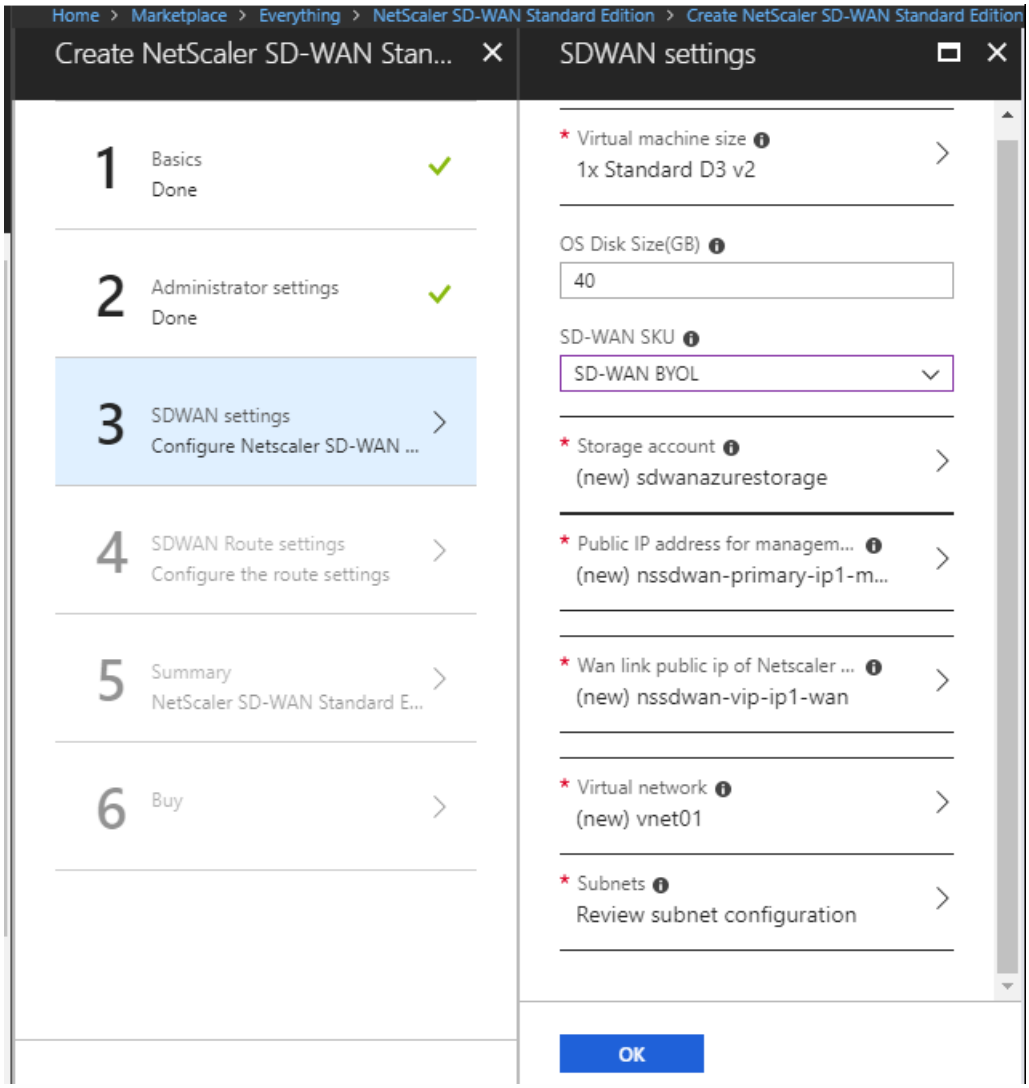

12. Configure route settings.

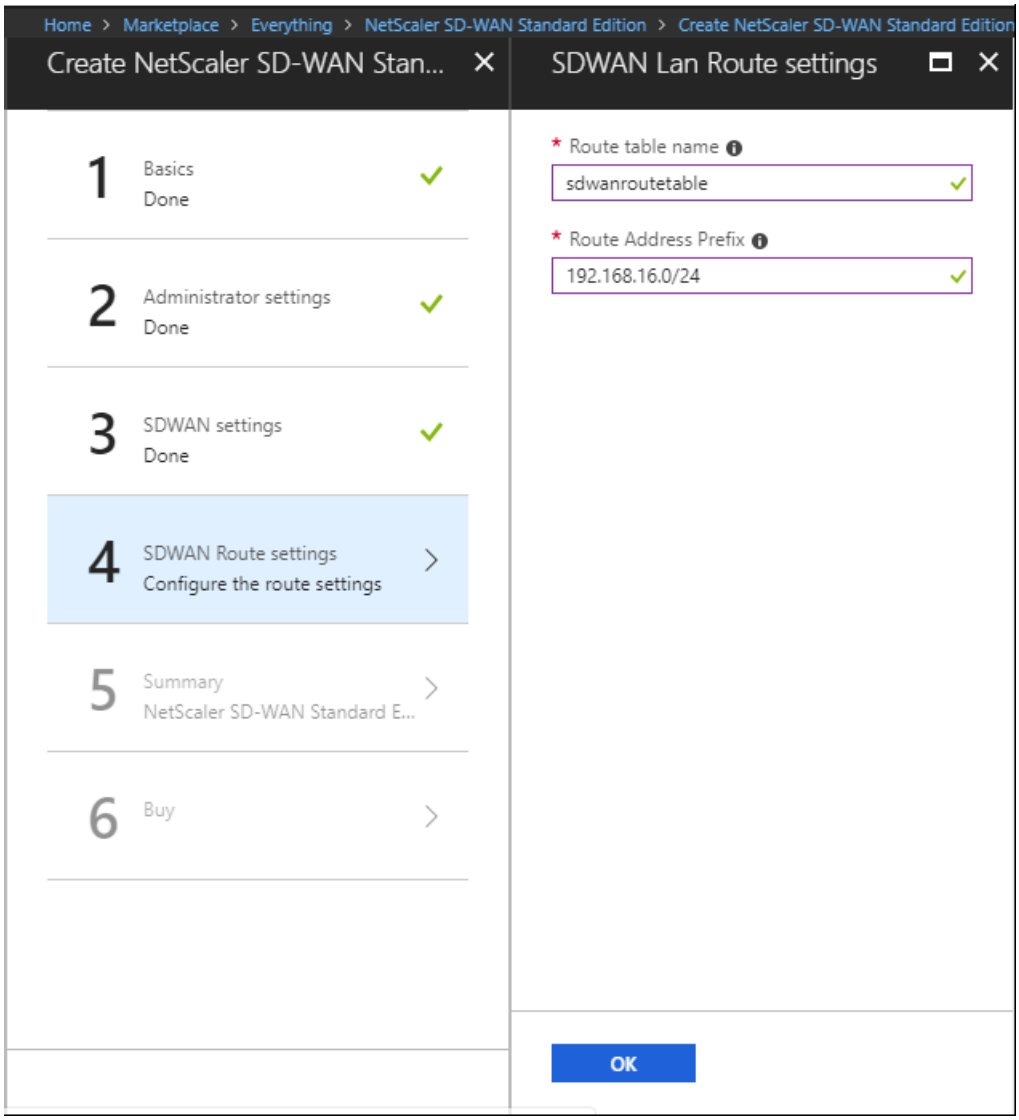

13. All the configuration that you provided in previous steps is validated and applied. If you have configured correctly, then you should see that the validation passed message as shown below. Click **OK**.

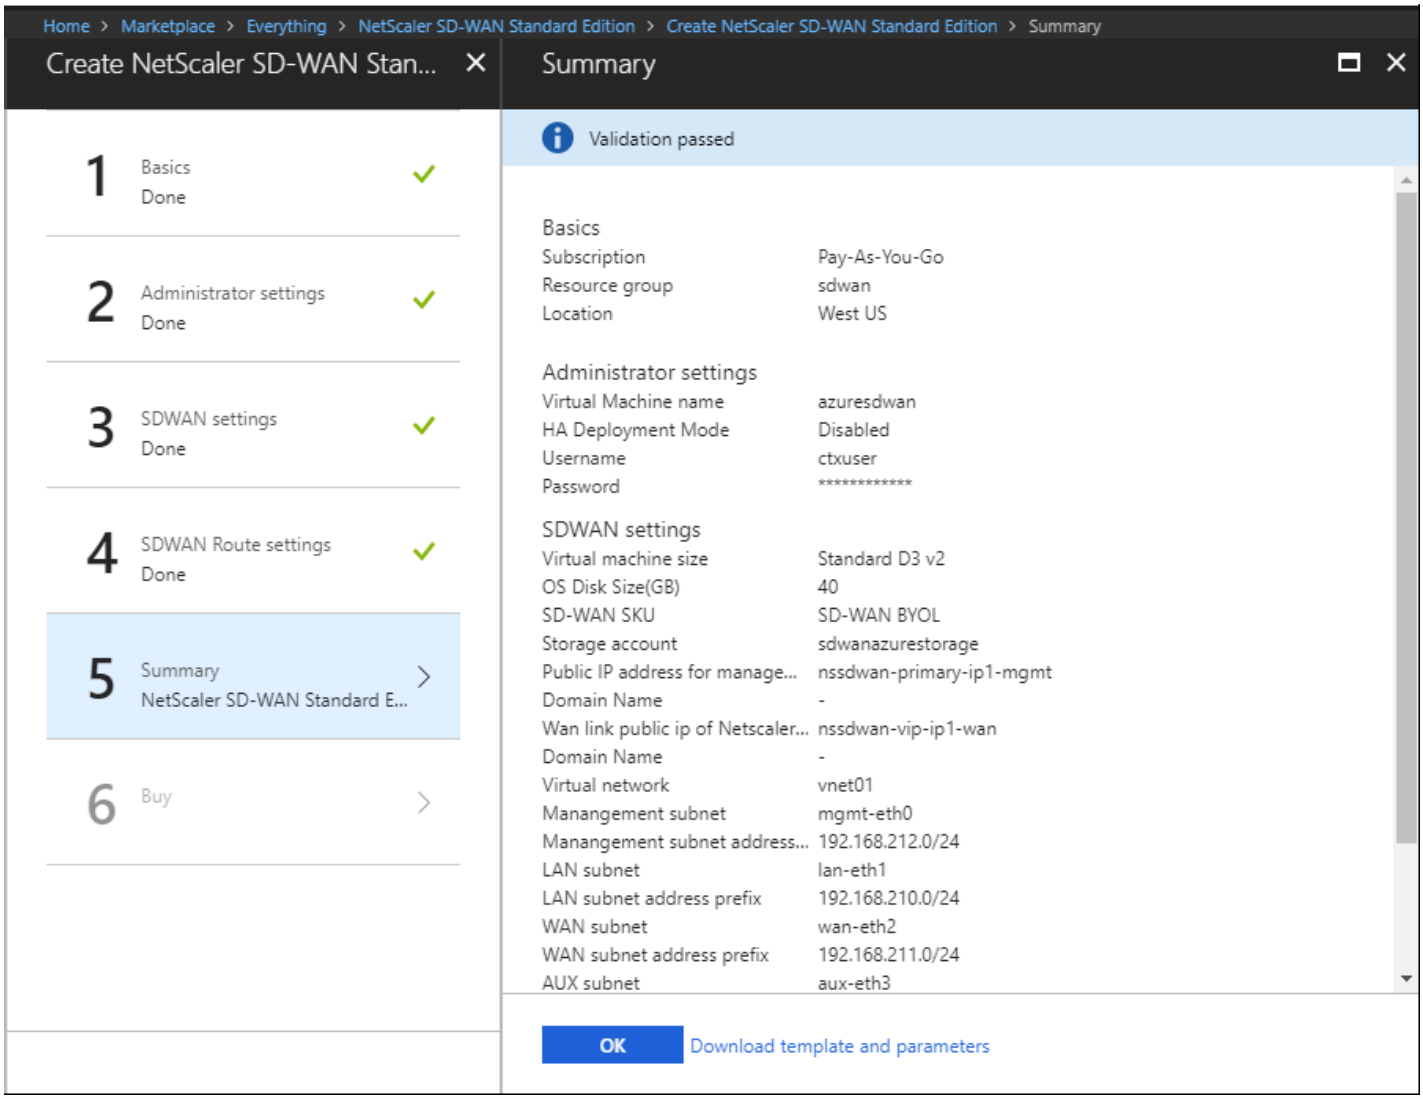

14. After validation is passed, click on **Create** to purchase and create the image.

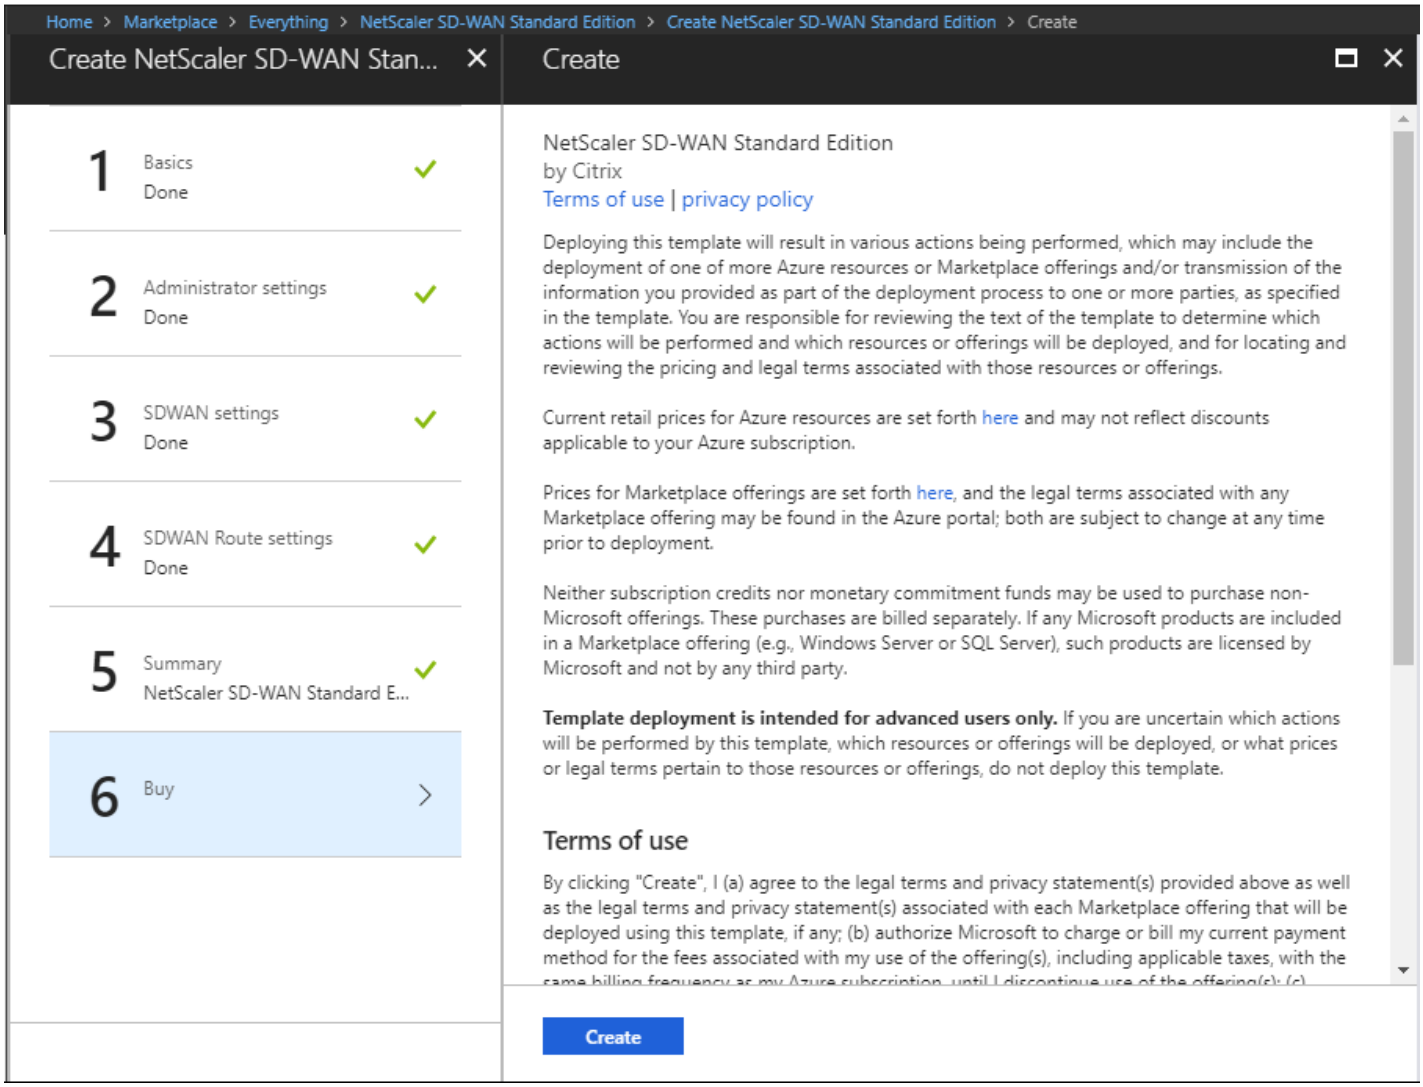

15. After you make the purchase, the deployment starts and you can view the status in the notifications section.

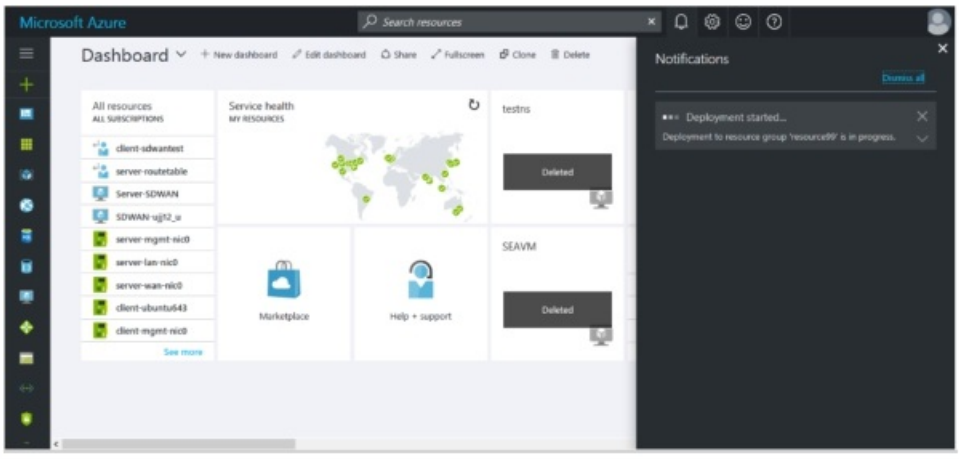

16. You can get more details of your deployment progress by selecting the Resource Group, then selecting Deployments under Settings.

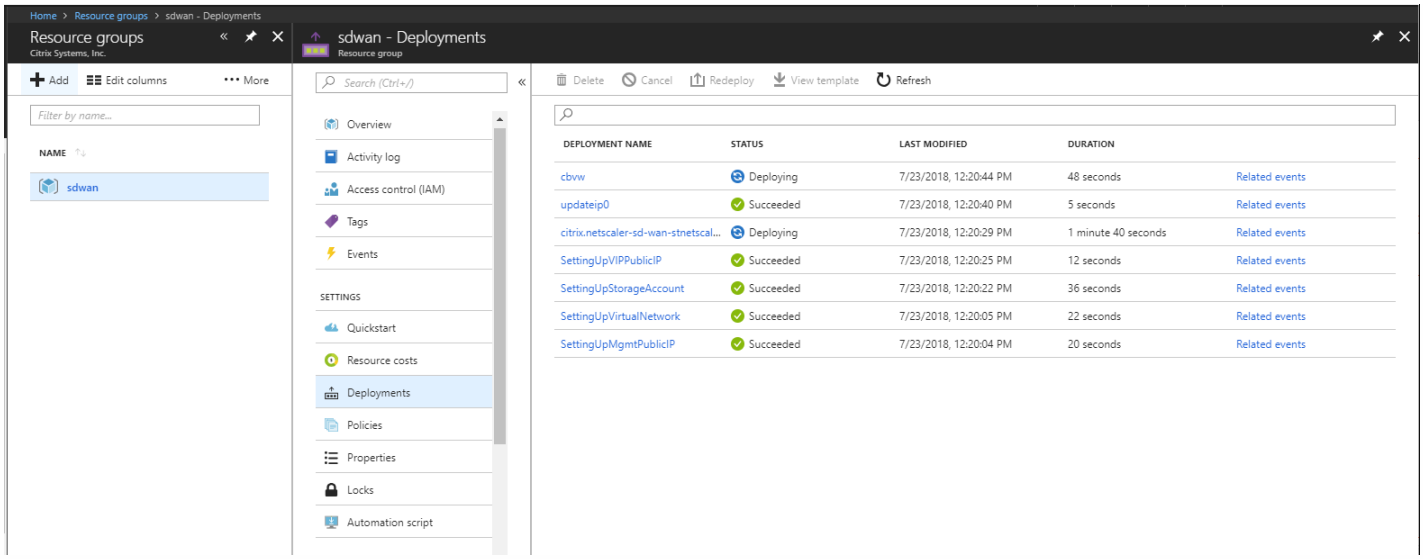

17. After the deployment is successful, try to access the GUI through the assigned public IP address for management. This can be identified by selecting the newly created instance under Virtual Machines, then identifying the assigned Public IP address in the Overview detail. Use the default credentials (admin/password) to log in. Modify the default password for security purposes.

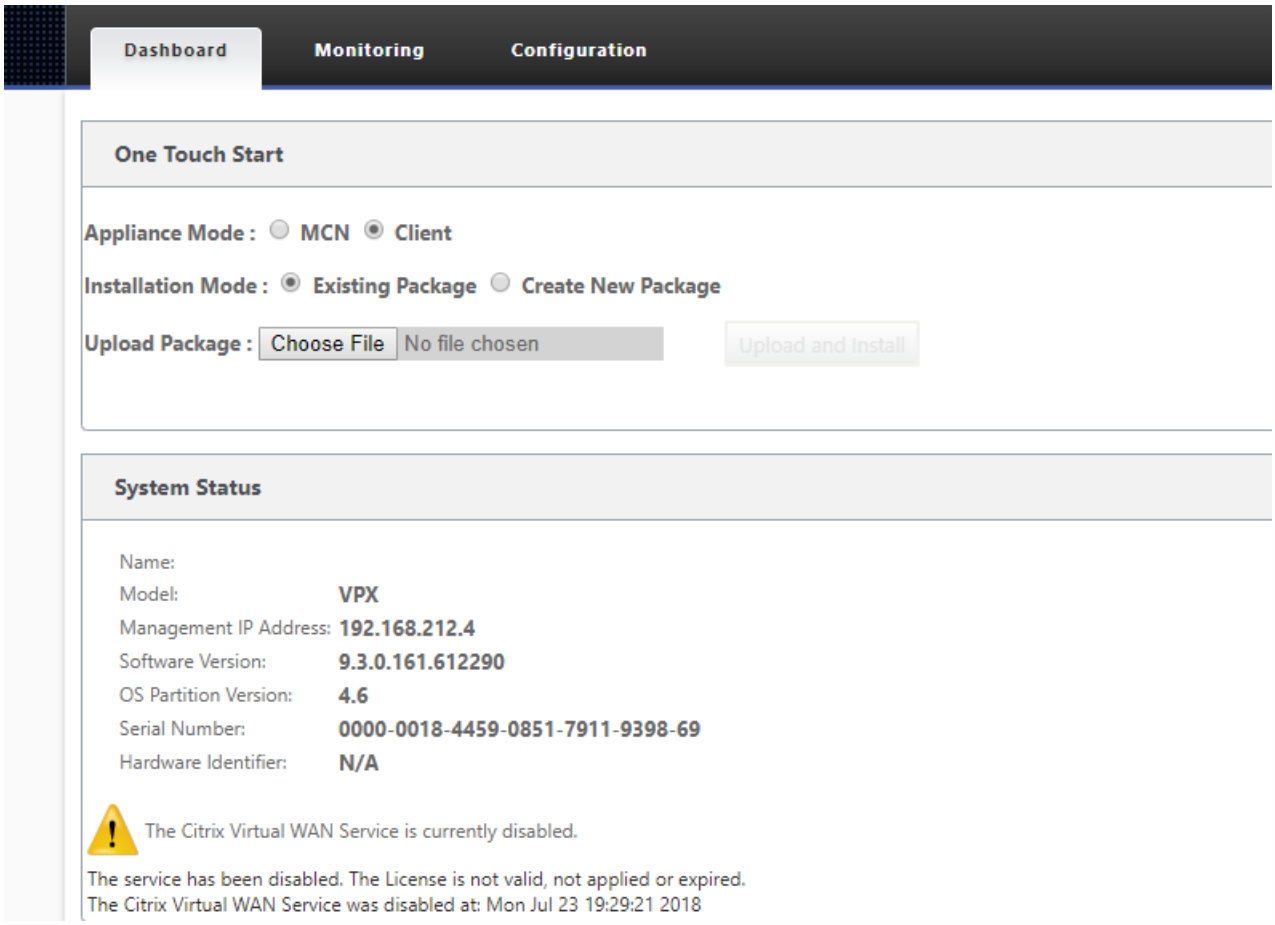

19. After you log into the GUI, notice that the virtual service is disabled. This is because the instance does not have configuration and license is not installed.

If this instance is intended to be a client node, upload package that is intended for this site with the matching subnet IP addresses defined.

If this instance is intended to be the Master Control Node (MCN), begin by promoting to an MCN and building the configuration.

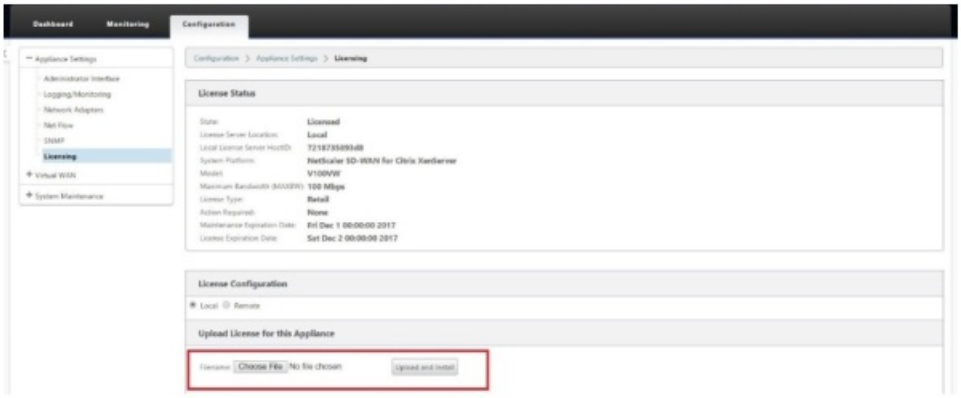

For more information about configuring SD-WAN service, see SD-WAN configuration.

Limitations - Microsoft Azure VMs

- After a VM is created and booted in Azure, the interfaces cannot be added or deleted. The VM profile (RAM/HD/CPUs) can be changed.
- Azure does not allow two network interfaces NIC on a VM to have IP address on same subnet. There is no L2 Support and bridging is not allowed. SE-VPX on Azure has to be deployed in Gateway mode.
- There is no concept of MAC address spoofing in Azure Cloud. The LAN subnet of the SE-VPX and the LAN subnet of the Client/Server Host have to be different. This requires more routing configuration to be done in two places.
- PCI Enumeration causes the order of NICs in an Azure VM to get switched on reboots. This might cause Management Subnet unreachability.

  – Routes have to be added in the Virtual WAN Configuration file directing all Virtual WAN Data traffic coming from the WAN to the Client/Server LAN Subnet.

Microsoft Azure supports only gateway mode for deployments. Refer to the following article for more information about gateway mode: https://docs.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vpngateways

Supported Modes

Aug 09, 2017

Table 1. Features Table for Citrix NetScaler SD-WAN VPX, and Citrix NetScaler SD-WAN VPX for Amazon

	Citrix NetScaler SD-WAN VPX	Citrix NetScaler SD-WAN VPX for Amazon
AutoConfiguration	N	N
SD-WAN/SD-WAN Plug-In	Y	N
Compression	Y	Y
RPC over HTTPS	Y	Y
SSL Compression	Y	Y
TCP Acceleration	Y	Y
Traffic Shaping	Y	N
Video Caching	N	N
Windows File System Acceleration	Y	Y
Windows Outlook Acceleration	Y	Y
XenApp/ XenDesktop Acceleration	Y	Y
Group Mode Mode	N	N
High Availability	N	N
Inline Mode	Y	Y**
Virtual Inline Mode	Y	N
WCCP Mode	N (SD-WAN)	N
VLANs	Y/Y/N***	N

*Depends on configuration of user-provided hardware.

**See WAN Optimization for SD-WAN.

***The three values are for is for XenServer, VMware, and Hyper-V, respectively. In columns showing only one value, the value applies to all three hypervisors.

Managing the Appliance

Aug 09, 2017

SD-WAN 4000/5000 is an appliance containing multiple virtual SD-WAN WAN accelerators controlled by a single virtual NetScaler load balancer. This combination provides a high-performance WAN accelerator for busy datacenters.

Automatically Configuring SD-WAN WANOP Appliances

Aug 09, 2017

If you are using Citrix Command Center to manage your Citrix SD-WAN WANOP appliances, the AutoConfiguration feature enables a SD-WAN WANOP appliance to automatically register itself with Citrix Command Center. After you have specified a DNS IP address in the setup wizard, the appliance performs reverse and forward lookups to identify the Command Center IP address. If you opt for having the appliance automatically configured by the Command Center server, the server starts configuring the appliance automatically soon after the appliance is registered with it. The Command Center server uses configuration profiles selected for the appliance to run configuration commands on the appliance. Additionally, you can use Citrix Command Center to manage and monitor the appliance remotely.

> **Note**
>
> This feature is supported with Citrix Command Center release 5.2 build 41 and later.

The autoconfiguration feature is supported on the following appliances

• SD-WAN 400 WANOP appliance

• SD-WAN 800 WANOP appliance

• SD-WAN 1000/2000 WANOP appliances

• SD-WAN 3000 WANOP appliance

• SD-WAN 1000/2000 WANOP appliances with Windows Server

Registering a SD-WAN WANOP Appliance with Citrix Command Center

Before you can use Citrix Command Center to manage a SD-WAN WANOP appliance, you must register the appliance with it.

To configure a SD-WAN WANOP appliance for autoconfiguration:

1. In the SD-WAN WANOP setup wizard (System > Configuration > System > Setup Wizard), specify the IP address of the DNS server used by the SD-WAN WANOP device.

2. Enter the registration password that you specified in Configure SD-WAN Registration Settings on the Command Center server, and click OK. Leave this field blank if you have not changed the password.

The SD-WAN WANOP appliance sends a registration request to the Command Center server, which automatically discovers the device and runs the commands available in the configuration profile.

Alternatively, you can navigate to **Configuration** > **Appliance Settings** > **Logging/Monitorin**g page and specify the Command Center details on the Command Center tab.

To register the appliance with Citrix Command Center:

1. Navigate to the **Configuration** > **Appliance Settings** > **Logging/Monitoring** page.

2. Click the **Command Center** tab.

3. In the **IP Address** field, type the IP address of the Citrix Command Center appliance with which you want to register the SD-WAN appliance.

4. In the Port field, type the port number used for Citrix Command Center. 8443 is the default port number used for Citrix Command Center.

5. In the Registration Password field, type the password that the Command Center administrator has set for a SD-WAN appliance to log on to Citrix Command Center.

Do not specify any password if the Command Center administrator has accepted the default password for registering the appliance.

6. Click Update.

7. Select AutoConfiguration By Citrix Command Center option to automatically configure the appliance through Command Center. The Status field changes from Disabled to Initiated registration, which later changes to Request accepted if the registration of the appliance with Citrix Command Center is successful.

Your SD-WAN WANOP appliance is registered with Citrix Command Center and you can now manage the appliance remotely by using Citrix Command Center.

NetScaler SD-WAN VPXL

Nov 21, 2017

NetScaler SD-WAN VPXL-SE is an enhanced version of the SD-WAN VPX-SE platform. Depending on the RAM/CPU/Disk configuration, the VPX platform can be operated either as VPX-SE or VPXL-SE. It is available on all VPX-SE platforms including Azure and AWS.

The NetScaler SD-WAN VPXL-SE platform can handle up to 128 sites when provisioned as an MCN.

Following are the configuration requirements for the VPXL-SE platform.

- 16 GB memory and 16 CPU cores.
- 250 GB of HDD.

Configuring NetScaler SD-WAN VPXL-SE

To configure SD-WAN VPXL-SE:

1. Import the SD-WAN VPX-SE base image (.ova or .xva template). Do not **Power ON** the Virtual machine.
2. Modify the VM resources for Memory to 16GB RAM, CPU to 16vCPUs and hard disk size to 250 GB.
3. Add the required NIC interfaces to the VM (for LAN and WAN interfaces).
4. **Power ON** the VM.
5. Now, SD-WAN VPX-SE would operate as VPXL-SE platform model.
6. In case the VM is already Powered ON, before modifying the VM resources, you need to perform **Re-image Virtual WAN Appliance Software** under **Configuration** > **System Maintenance** > **Update Software**, and use the **cb-vw_CBVPXL_version.tar.gz** image file.

> **Note**
>
> Ensure to uncheck options that mention to Power on the Virtual machine, after VM provisioning/import process is complete.

How to Upgrade VPX-SE to VPXL-SE

To upgrade VPX-SE to VPXL-SE:

1. The VPX-SE should have been installed with SD-WAN base release version of 9.3.0 or higher.
2. Backup and save your existing configuration, if the VPX-SE you are upgrading is an MCN.

> **Important**
>
> A VPX-SE appliance provisioned with a SD-WAN release version 9.2 image cannot be upgraded to a VPXL-SE appliance.

You want to use VPXL-SE platform, upgrade SD-WAN release version 9.2. to 9.3 using the change management procedure in SD-WAN web GUI.

For more information, see upgrading to SD-WAN 9.3.

Deploying VPXL-SE in Microsoft Azure and AWS

The steps to configure and deploy VPXL-SE in Microsoft Azure are similar to the steps to deploy VPX-SE appliance. The only difference in the configuration steps is to choose hard disk space as 250 GB and instance as F8 for the VPXL-SE appliance to be successfully deployed in Azure.

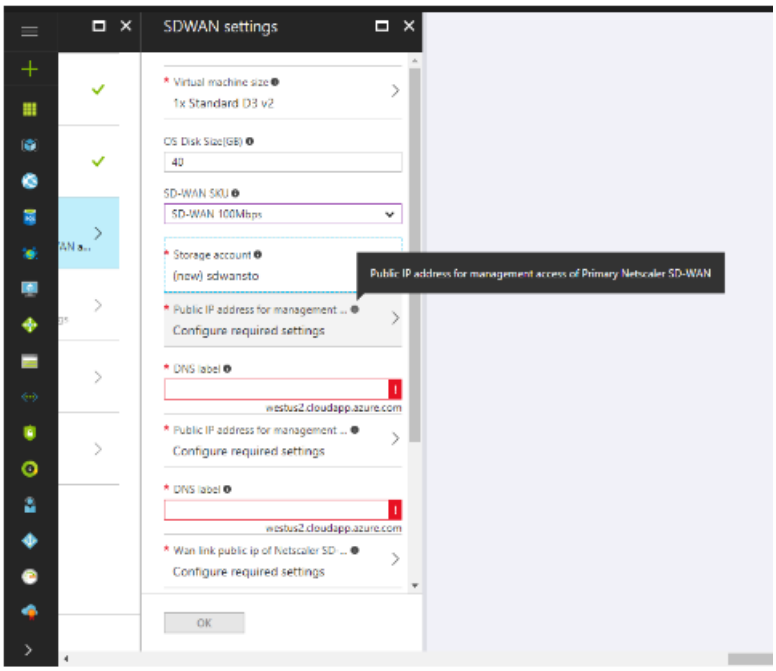

Similarly, the steps to configure and deploy VPXL-SE in AWS is similar to the steps for deploying VPX-SE appliance in AWS with the only difference of selecting disksize of 40 GB and instance of m4.4x large.

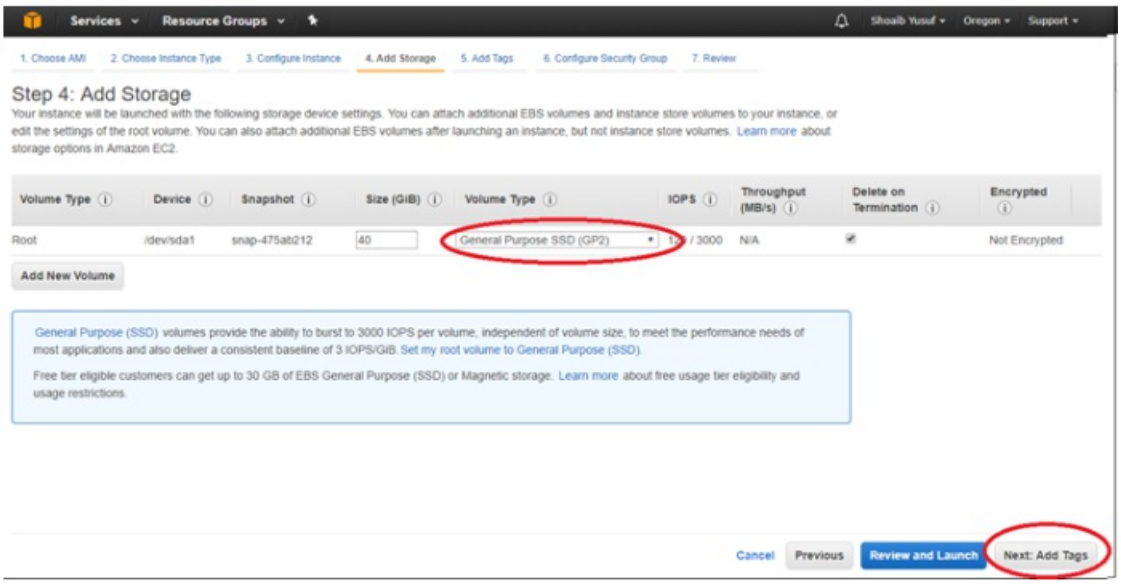

Related information

- Installing SD-WAN Virtual Appliances on VMware ESX
- Installing SD-WAN Virtual Appliances on XenServer
- Installing SD-WAN Standard Edition AMI on Amazon AWS
- Installing SD-WAN VPX Standard Edition on Azure
- SD-WAN Standard Edition Virtual Appliance (VPX) HA Support for AWS
- SD-WAN Standard Edition Virtual Appliance (VPX) HA Support for Microsoft Azure

Installing SD-WAN WANOP Edition AMI on Amazon AWS

Aug 09, 2017

The NetScaler SD-WAN VPX for Amazon AWS brings acceleration support to the Amazon cloud.

Note: At the time of the 7.1.0 software release, the newest supported release of SD-WAN (now SD-WAN) WANOP-VPX for Amazon AWS is 7.0.1. Use this version in conjunction with release 7.1.0 on other appliances.

Five variations are supported, four of which have hardwired licensing, and one of which uses ordinary SD-WAN licensing:

- 2 Mbps
- 10 Mbps
- 20 Mbps
- 45 Mbps
- "Bring your own license," which uses a standard Citrix license to determined the licensed bandwidth.

Besides the hardwired licensing, the major difference between SD-WAN WANOP-VPX for Amazon AWS is that it supports only a single port for both management and acceleration. This means that the appliance cannot be used in inline mode.

To create a SD-WAN WANOP-VPX on Amazon AWS, you go through the same process as with creating any other instance, setting a few instance parameters to non-default settings.

Instantiating a SD-WAN Virtual Appliance (AMI) on AWS

To install a SD-WAN virtual appliance in an AWS VPC, you need an AWS account. You can create an AWS account at http://aws.amazon.com/. SD-WAN is available as an Amazon Machine Image (AMI) in AWS Marketplace.

Note: Amazon makes frequent minor changes to its AWS pages, so the following instructions may not be exact.

To instantiate a SD-WAN virtual appliance (AMI) on AWS

1. In a web browser, type http://aws.amazon.com/.
2. Click My Account/Console, and then click My Account to open the Amazon Web Services Sign in page.
3. Use your Amazon AWS account credentials to sign in. This will take you to the Amazon Web Services page.
4. Click EC2 in the Compute & Networking section, then click Launch Instance.
5. In the Create a New Instance dialog box, select AWS Marketplace, and then click Continue to open the Request Instance Wizard.
6. In the Request Instance Wizard dialog box, click AWS Marketplace tab.
7. In the Search text field, type SD-WAN to search for the SD-WAN AMI, and click Search.

   On the search result page, select one of the Citrix SD-WAN offerings On the Citrix SD-WAN page, click Continue.
8. On the Launch with EC2 Console tab, click the Accept Terms button, if present, then click Launch with EC2 Console for the region where you want to launch Citrix SD-WAN AMI.
9. On the Request Instance Wizard page, type 1 in the Number of Instances text box, and from the Instance Type drop-down list, select Large (m1.large, 7.5GIB).

10. From the Subnet drop-down list, select the private network subnet, and then click Continue.
11. On the next page, in the Advanced Instance Options section, you can change values from their defaults if you choose, and then click Continue.
    Note: SD-WAN AMI is not supported with more than one network interface. Therefore, the value of Number of Network Interfaces field is set to 1.
12. On the Request Instances Wizard page, enter a name for the EC2 instance in the Value text box, and then click Continue.

    On the Request Instances Wizard page, select one of the three Kay Pair options and then click Continue.
13. Verify the EC2 instance configuration details, and then click Launch to launch the EC2 instance.
14. Click Close to close the Launch Instance Wizard dialog box. The new EC2 instance is launched successfully.

Disabling the Source/Destination Check Feature

Aug 09, 2017

You must disable the Source/Destination check feature of SD-WAN AMI instance for it to work properly on AWS.

To disable the Source/Destination check feature

1. On the Amazon EC2 Console Dashboard page, in the navigation pane, click instances. The new EC2 instance should appear in the My Instances list.
2. Select the new EC2 instance. The instance details appear in the EC2 Instances pane.
3. Right-click the new EC2 instance and then select Change Source/Dest Check from the popup menu.
4. In the Change Source / Dest. Check dialog box, click Yes, Disable to disable the feature.

Configuring SNMP Monitoring for the SD-WAN WANOP Edition AMI on AWS

Aug 09, 2017

You must enable SNMP monitoring on the SD-WAN AMI on AWS. Also, you must grant SNMP monitoring access to the paired NetScaler VPX or SD-WAN Connector on AWS by adding its NSIP on the SD-WAN AMI instance.

To configure SNMP monitoring on the SD-WAN Connector AMI by using the SD-WAN graphical user interface

1. In the navigation pane, expand Configuration, and then click Logging/Monitoring.
2. In the details pane, click the SNMP tab.
3. In the System Information section, in the SNMP Status row, click Enable. This action enables SNMP monitoring on the SD-WAN AMI instance.
4. In the Access Configuration section, add SNMP monitoring access to SD-WAN VPX appliance by setting the following parameters:
   - Community String (set to the string public)
   - Management Station IP (set to the NSIP of the SD-WAN VPX on AWS)
5. Click Add.

Limitations and Usage Guidelines for the SD-WAN WANOP Edition AMI Instances on AWS

Aug 09, 2017

- High Availability setup for SD-WAN AMI instances is not supported.
- SD-WAN AMI instance in Group Mode is not supported.
- SD-WAN plug-ins are not supported.
- Tagged VLAN is not supported because of the inherent limitation of AWS.
- Traffic shaping is not supported.
- You may create only an m1.large SD-WAN AMI instance on AWS.
- IP address/gateway/subnet assignment using the SD-WAN management user interface is not supported.
- Console access is not available for SD-WAN AMI instance on AWS.
- While configuring the SD-WAN instance, you may not change the disk size, which has a default value of 250 GB. A higher capacity disk does not increase the available Disk Based Compression (DBC) cache size.

SD-WAN Licensing Overview

Aug 09, 2017

The process of allocating your SD-WAN licenses has been greatly simplified. The new licensing framework allows you to focus on getting maximum value from Citrix products.

In the SD-WAN configuration utility (GUI), you can use your hardware serial number (HSN) or your license activation code (LAC) to allocate your licenses. Alternatively, if a license is already present on your local computer, you can upload it to the appliance.

For all other functionality, such as returning or reallocating your license, you must use the licensing portal. Optionally, you can still use the licensing portal for license allocation. For more information about the licensing portal, see "http://support.citrix.com/article/CTX131110".

Note:

- On a SD-WAN appliance, you can use the HSN or LAC to allocate your license or upload the license to the appliance from a local computer. On a SD-WAN VPX appliance, you can only upload the license to the appliance from a local computer.
- You must purchase separate licenses for each appliance in a high availability (HA) pair. Make sure that the same type of licenses are installed on both the appliances. For example, if you purchase a platinum license for one appliance, you must purchase another platinum license for the other appliance.

Note: You must purchase separate licenses for each appliance in a high availability (HA) pair. Make sure that the same type of licenses are installed on both the appliances. For example, if you purchase a platinum license for one appliance, you must purchase another platinum license for the other appliance.

Prerequisites

Aug 09, 2017

To use the hardware serial number or license activation code to allocate your licenses:

- The management service virtual machine of the appliance must have an access to public domain.
- Your license must be linked to your hardware, or you must have a valid license activation code (LAC). Citrix sends your LAC by email when you purchase a license.

Allocating your License by using the Configuration Utility

Aug 09, 2017

If your license is already linked to your hardware, the license allocation process can use the hardware serial number. Otherwise, you must type the license activation code (LAC).

To allocate your license

1. In a web browser, type the IP address of the management service of the SD-WAN appliance (for example, http://192.168.100.1).
2. In User Name and Password, type the administrator credentials.
3. Navigate to the System > Configuration tab.
4. On the Configuration tab, navigate to System > Licenses.
5. In the details pane, click Manage Licenses.

6. Click Update Licenses, and then select one of the following options:
   - **Use Hardware Serial Number**—The software internally fetches the serial number of your appliance and uses this number to display your license(s).

   - **Use License Activation Code**—Citrix emails the LAC for the license that you purchased. Enter the LAC in the text box.

7. Click Get Licenses. Depending on the option that you selected, one of the following dialog boxes appears.
   - The following dialog box appears if you selected Hardware Serial Number.

   - The following dialog box appears if you selected License Activation Code.

8. Select the license that you want to allocate, and then click Get.
9. Click Apply for the license to take effect.

Installing the License

Aug 09, 2017

If you downloaded your license file to your local computer by accessing the licensing portal, you must upload the license to the appliance.

To install a license file by using the configuration utility

1. In a web browser, type the IP address of the SD-WAN applaince (for example, http://192.168.100.1).
2. In User Name and Password, type the administrator credentials.
3. On the Configuration tab, navigate to System > Licenses .
4. In the details pane, click Manage Licenses.
5. Click Update Licenses, and then select **Upload License Files**.

6. Click Add New License, then select **Upload license files from a local computer**.
7. Click Browse. Navigate to the location of the license files, select the license file, and then click Open.

Verifying SD-WAN Licenses

Aug 09, 2017

Before using a feature, make sure that your license supports the required number of accelerators.

To verify the number of licensed accelerators by using the configuration utility

1. In a Web browser, type the IP address of the appliance, such as http://192.168.100.1.
2. In User Name and Password, type the administrator credentials.
3. Navigate to the **System** > **Configuration** tab. The License Information section displays the number of accelerators licensed.
4. Alternately, navigate to **SD-WAN** > **Instances**. The Licensed column displays the status of licensed accelerates as green.

Internet Protocol Version 6 (IPv6) Acceleration - SD-WAN WANOP Appliances

Aug 09, 2017

> **Note**
>
> IPv6 Acceleration is supported only on the SD-WAN WANOP appliances. This feature is not supported on the SD-WAN Standard Edition and Enterprise Edition appliances.

When you connect to the Internet through a device, the device is assigned an IP address. The IP address identifies and indicates the location of the appliance. The number of devices connecting to the Internet is rapidly increasing. As a result, it is difficult to manage the request for the IP addresses with the existing version of Internet Protocol (IP), IPv4, which uses 32-bit addresses. By using IPv4, approximately 4.3 billion addresses can be assigned to the devices connecting to the Internet.

IPv6 addresses this issue by using 128-bit addresses and a hexadecimal label to identify the network interfaces of devices on an IPv6 network. Because IPv6 supports far more IP addresses than does IPv4, organizations and applications are gradually introducing support for the IPv6 protocol.

The IPv4 and IPv6 protocols are not interoperable, which makes the transition difficult. To accelerate the increasing IPv6 traffic from various applications supported on the SD-WAN appliance, you can enable the IPv6 Acceleration feature.

By default, IPv6 is disabled on the appliance.

To enable IPv6 acceleration on a SD-WAN appliance

1. Navigate to the System > Configuration > System page.
2. Click enable the Enable IPv6 for Data Traffic link in the System Settings section.
3. Select the Enable IPv6 for Data Traffic option in the dialog box
4. Click OK, as shown in the following screen shot.

Verifying IPv6 Connections

Aug 09, 2017

After enabling IPv6 acceleration on the appliance, the appliance starts accelerating traffic for the applications using IPv6 protocol. To make sure that the appliance is accelerating the IPv6 traffic, you can monitor such connections on the appliance.

To monitor the IPv6 connections, navigate to the SD-WAN > Monitoring tab. The **Connections** page of the **Monitoring** tab display IPv6 protocols traffic related statistics:

- **Connections**: The Connections page lists details of all the connections established with the appliance. This page consists of two tabs, Accelerated Connections and Unaccelerated Connections. The Accelerated Connections tab lists all connections that the appliance is accelerating. You can identify IPv6 traffic in this tab by referring to the Initiator and Responder column of each entry. If these columns contain hexadecimal IP address values, the entry represents an IPv6 connection, as shown in the following screen shot.

  IPv6 connections that are not accelerated, are listed on the Unaccelerated Connections tab. If you want to accelerate these connections, you might need to troubleshoot and fine tune the application parameters on the appliance. As on the **Accelerated Connections** tab, you can identify the IPv6 connections on this tab by referring to the **Initiator** and **Responder** columns of each entry.

- **Top Applications**: The Top Applications page provides granularity in the time frame that you can use to graphically represent the traffic throughput of various applications served by the SD-WAN appliance. By default, traffic throughput is displayed by the last minute. However, you can change the time frame by selecting Last Minute, Last Hour, Last Day, Last Week, or Last Month from the list available on the Title bar of the page. This page has three tabs, **Top Applications Graphs**, **Since Last Restart**, and **Active Applications (Since Last Restart)**. The Top Applications Graphs tab contains the following statistics:
  - **Total Application Link Throughput Percentage (Sent)**: This is a pie chart depicting the percentage of traffic that the appliance has sent to each application. If the appliance has sent a significant percentage of traffic for an application using IPv6 protocol, the application has its percentage of traffic depicted in this graph.
  - **Total Application Link Throughput Percentage (Received)**: This is a pie chart depicting the percentage of traffic that the appliance has received from each application. If the appliance has received a significant percentage of traffic from an application using IPv6 protocol, the graph displays the percentage of traffic generated by the application.
  - **Sent Rate**: This is a stacked graph of series of data depicting the rate, in bits per second, at which the appliance has sent traffic to each application. If the appliance has sent data to an application using IPv6 protocol, a series depicting each application using IPv6 protocol is also plotted on this graph.
  - **Received Rate**: This is a stacked graph of series of data depicting the rate, in bits per second, at which the appliance has received traffic from each application. If the appliance has received data from an application using IPv6 protocol, a series depicting each application using IPv6 protocol is also plotted on this graph.
  - **Top Applications table**: This is a table of statistics for each application. The table lists all applications for which the

appliance has served traffic, along with sent and received rates in bits per second, total bytes sent and received, percentage of the traffic for the application, and the rate at which the appliance has served traffic for the application. If the appliance has served traffic for an application using IPv6 protocol, the application is listed in this table, along with its statistics.

- **Application Groups**: This is a table of statistics for each application, along with its application group and parent application, if any. The table lists bytes sent and received for the application. Each application, and its application group and parent application are displayed as hyperlinks. If you click the hyperlink, granular details of the statistics are displayed for the link you have clicked. If the appliance has served traffic for an application using IPv6 protocol, the application is listed in this table, along with its statistics.

The **Since Last Restart** tab contains statistics on the application traffic since the time you restarted the appliance. The tab contains the Total Application Link Throughput Percentage (Sent) and Total Application Link Throughput Percentage (Received) graphs, and Top Applications and Application Groups tables, depicting statistics similar to the Top Applications Graphs tab but with data since the appliance was restarted. The **Active Applications (Since Last Restart)** tab contains a table listing all active applications since the appliance was restarted. This table contains details about sent and receive rate, total bytes sent and received, and total packets sent and received for the applications.

Standard Edition in AWS for Cloudwatch Support

Aug 09, 2017

NetScaler SD-WAN Standard Edition in AWS now supports basic CloudWatch for monitoring your SD-WAN instance running on AWS infrastructure. CloudWatch alarms send notifications or automatically make changes to the resources you are monitoring. Under basic monitoring, seven pre-selected metrics at five minute frequency and three status check metrics at one minute frequency are available for your SD-WAN instance for no additional charge. You can view the following metrics for your SD-WAN image.

a) CPU Utilization - The percentage of allocated compute units that are currently being used for the instance. This metric identifies the processing power required to run an application upon a selected instance.

b) Diskread operations - Read operations from all instance volumes available for the specified period.

c) DiskWrite operations – Write operations for all instance store volume available for the instance in a specified duration of time.

d) DiskReadBytes – Bytes written to all instance store volumes available to the running instance

e) NetworkIn – This metric identifies the volume of incoming traffic for a single instance.

f) Network Out - This metric identifies the volume of outgoing traffic for a single instance.

g) Networkpacketsout – Number of packets sent out on all network interfaces by the instance, this is only available for basic monitoring.

Installing SD-WAN SE Virtual Appliances (VPX) in Linux-KVM Platform

Apr 12, 2018

Installing SDWAN VPX-SE Appliances in KVM Hypervisor platform:

1) To set up NetScaler SDWAN VPX-SE for the Linux-KVM platform:

a. Use the graphical Virtual Machine Manager (Virtual Manager) application.

or

b. Use the virsh program Linux-KVM command line.

2) The host Linux operating system must be installed on suitable hardware by using virtualization tools such as KVM Module and QEMU. The number of virtual machines (VMs) that can be deployed on the hypervisor depends on the application requirement and the chosen hardware.

3) The .qcow2 file has to be unique for each of the NetScaler VPX instance provisioned. It is a virtual hard disk (VHD) that is attached to VM.

Prerequisites:

1) Install Ubuntu 16.04 on the bare metal appliance which supports Virtualization. Follow the below steps to check if the bare metal appliance supports Virtualization.

2) 64-bit x86 processors with the hardware virtualization features included in the AMD-V and Intel VT-X processors.

a. To test whether your CPU of Linux host supports virtualization, enter the following command at the host Linux shell prompt:

egrep -c '(vmx|svm)' /proc/cpuinfo, this output must be more than 0.

3) Alternative to step 2, install a package/tool called “cpu-checker” (sudo apt-get install cpu-checker), enter the following command :

kvm-ok, the output must be “KVM acceleration can be used”.

4) Minimum hardware requirements:

As the SDWAN-Virtual WAN (guest OS) requires 4 V CPUs, 4GB RAM and 40 GB (VHD). You must have a host with these specifications which can satisfy this.

5) Software requirements:

Ubuntu 16.04.2 LTS (GNU/Linux 4.4.0-78-generic x86_64)

Install qemu-kvm, libvirt-bin, virt-manager : sudo apt-get install qemu-kvm libvirt-bin virt-manager bridge-utils. Execute this command to obtain all the required packages/software.

Provisioning the SD-WAN VPX appliances by using Virtual Machine Manager (VMM):

1) Open the Virtual Machine Manager. Go to **Application** > **System Tools** > **Virtual Machine Manager**, and provide the logon credentials in the **Authenticate** window.

2) Once the VMM opens, you should see QEMU/KVM. This indicates that the VMM is not connected to the QEMU Virtualization.

NIC ordering for SD-WAN VPX-SE provisioning must be in the following order; **Management**, **LAN** and **WAN**.

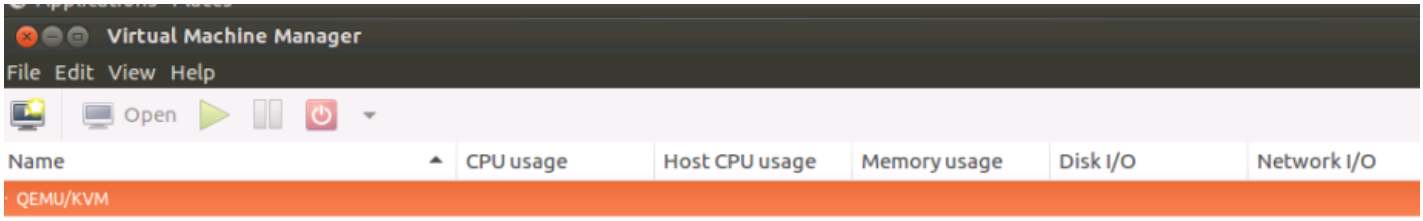

3) Select **New Virtual Machine**.

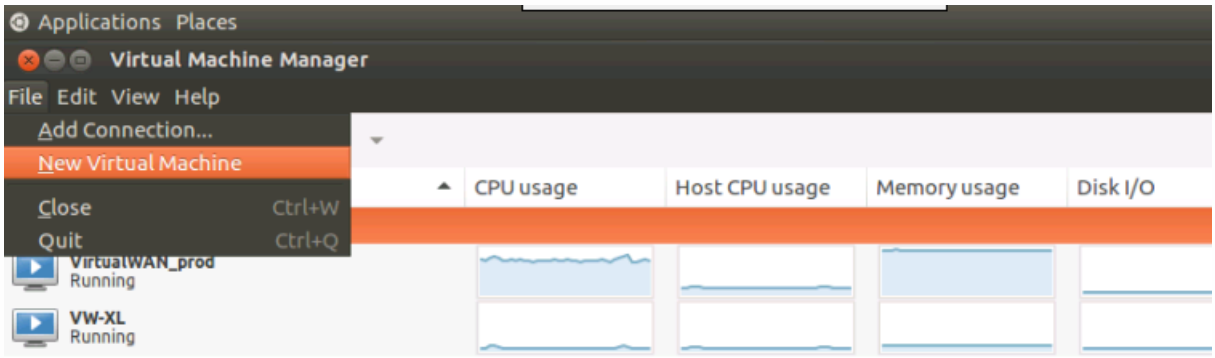

4) Select the VHD, the VHD used by one machine cannot be shared. Unique VHD is required for every Virtual Machine.

Browse the image and select the path where it is downloaded.

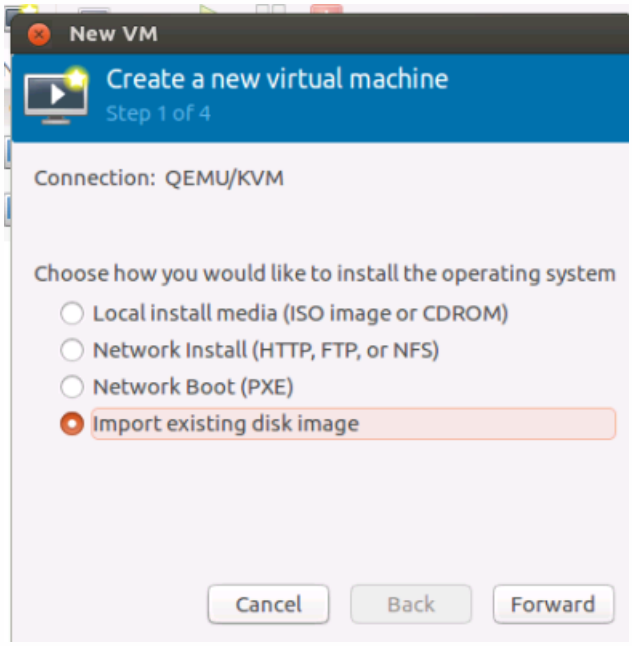

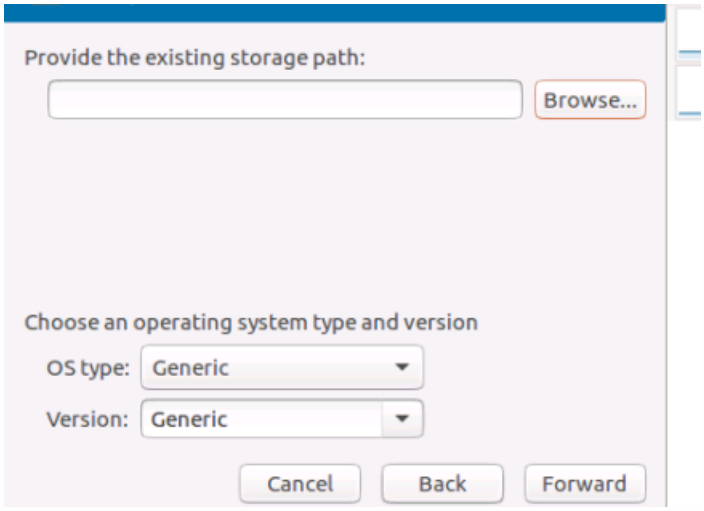

5) Provide RAM as 4096 MB and CPU as 4.

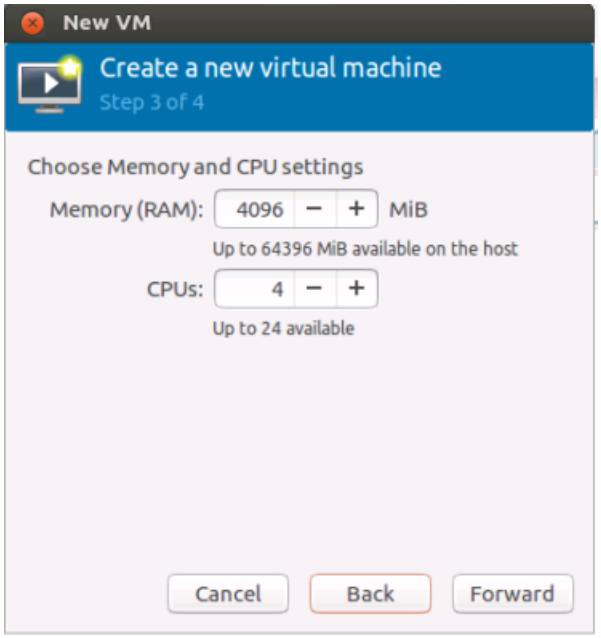

6) Name the VM as needed and select **Customize configuration before Install**. As by default one NIC gets selected to the Virtual Machine, you can see the **Network selection** option.

In this setup **enp4s0f0** is the Management Network for the Host machine, and if you want to use this NIC, sharing same NIC between guests and host for Management access. Source Mode is Bridge since it is shared between VMs.

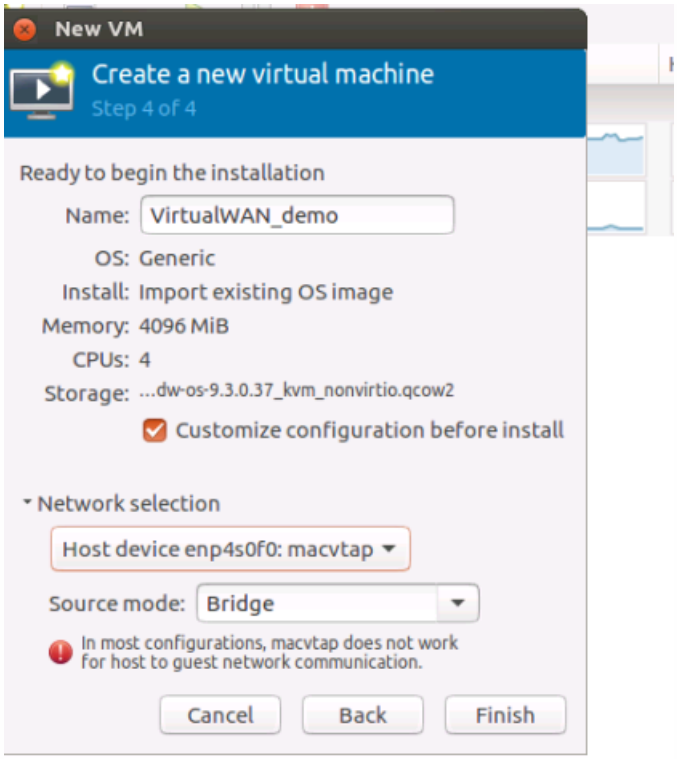

7) After clicking Finish, ensure you select **customize configuration before install** for further configuration.

For the NIC that is assigned, in this example “enp4s0f0:macvtap” you need to select the Device model as “virtio”. The model that is supported for communication.

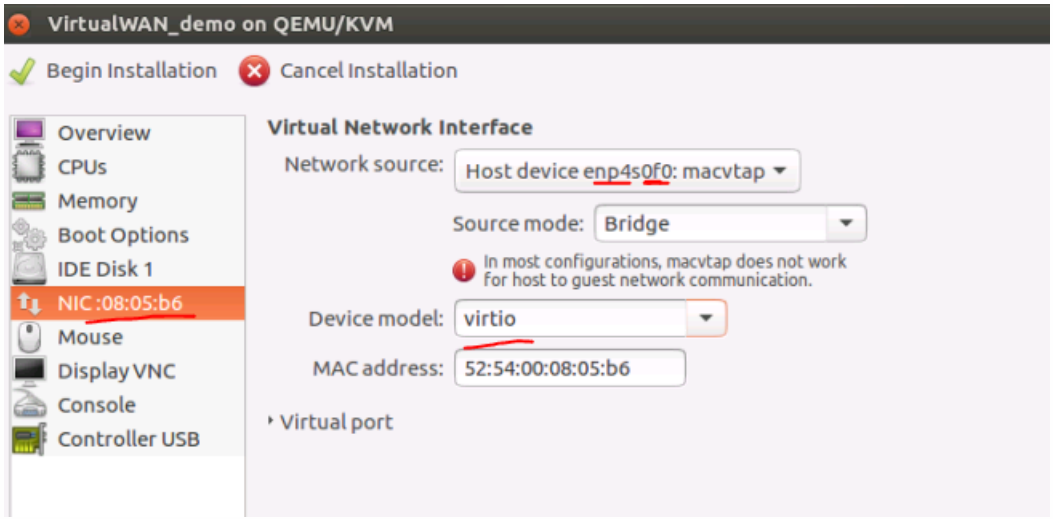

8) Add additional NICs for LAN and WAN with **Add Hardware** at the bottom left side corner.

For good Performance, it is recommended to use Source Mode as Pass-through (Only one VM can use the Lower NIC and hence it cannot be shared across VMs). For LAN and WAN interfaces use “Pass-through” Mode and Device Model should be “virtio”.

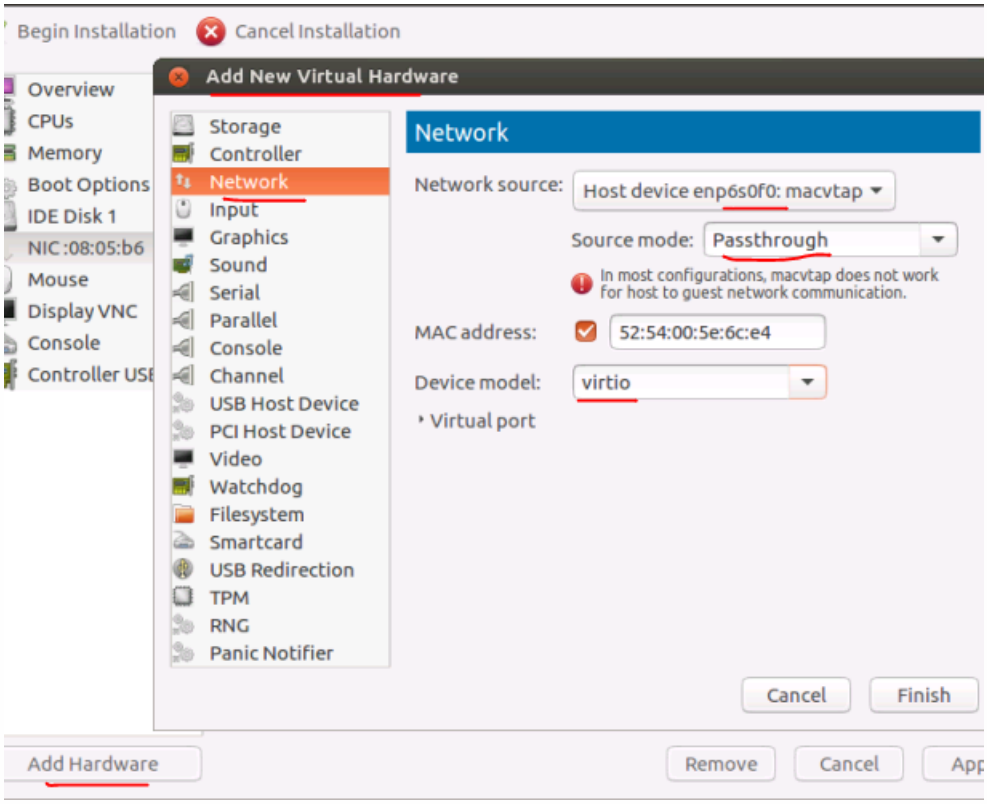

9) Select **Begin Installation** for the installation process to start and you can see the console of the appliance.

10) Use **management_ip command** to set the IP address.

How to Deploy SD-WAN Appliances in Linux-KVM Hypervisor Platform Instance on the same Host

Deploying SD-WAN appliances in HA mode on the same host requires sharing the same physical interface across SD-WAN VPX appliances. For example; the eth3 of physical hypervisor (host) is used for WANLink-1 for Primary VM, the same interface should be used for secondary appliance, so that if primary appliance becomes inactive, the secondary appliance can respond to the ARP requests for shared MAC.

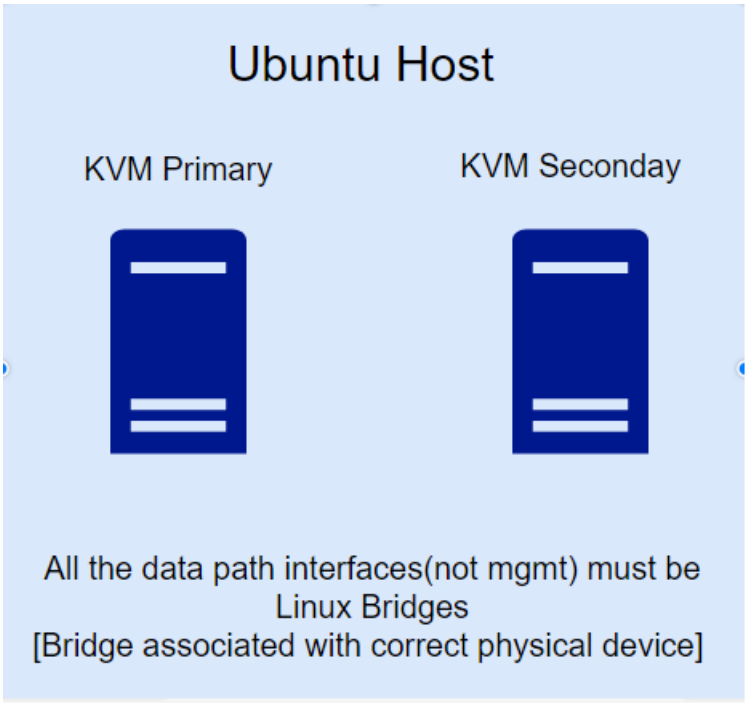

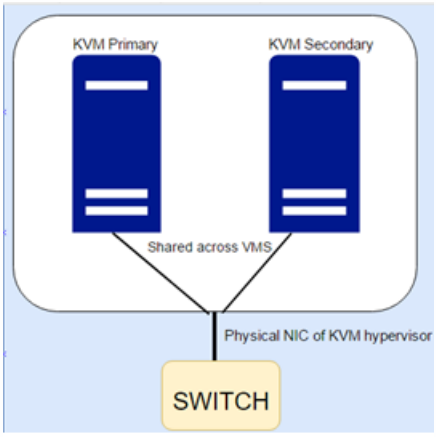

For sharing of the Physical NIC between the VMs which are on the same host, the source modes that can be used according to KVM networking is **MACVTAP Bridge** or **Linux Bridge**.

How to use Linux Bridge

- Create Bridge using *"brctl"* on the Host (KVM Hypervisor level).
- Associate the required Physical NIC to the bridge created (using brctl commands).
- These bridges created at Hypervisor level should be now be associated to the SD-WAN VM.
- Primary and Secondary VMs are now associated with the Linux bridges created.

To create Linux Bridge and associate it with Virtual Machine:

- Adding bridge,

  "brctl addbr ha-brwan1"

- Associating physical nic to the bridge "ha-brwan1"

‘brctl addif ha-brwan1 eth3”

- Associating the bridge “ha-brwan1” to the SD-WAN-SE (Virtual WAN) (both Physical and Secondary)

  1. When adding network interface, select Network source as “Specify shared device name”.

  2. Under Bridge Name, provide the name of the bridge created.

  3. Device Model should always be “virtio”.

Create bridges for LAN and WAN interfaces. The below snapshot depicts the way to associate interface to SDWAN-SE using Virtual Machine Manager.

> ## Note
>
> These steps should be followed only when both Primary and Secondary HA node are present on the same KVM Hypervisor/Host. In case, if HA nodes are present on different Hypervisors then **MACVTAP: Passthrough source** mode can be used.

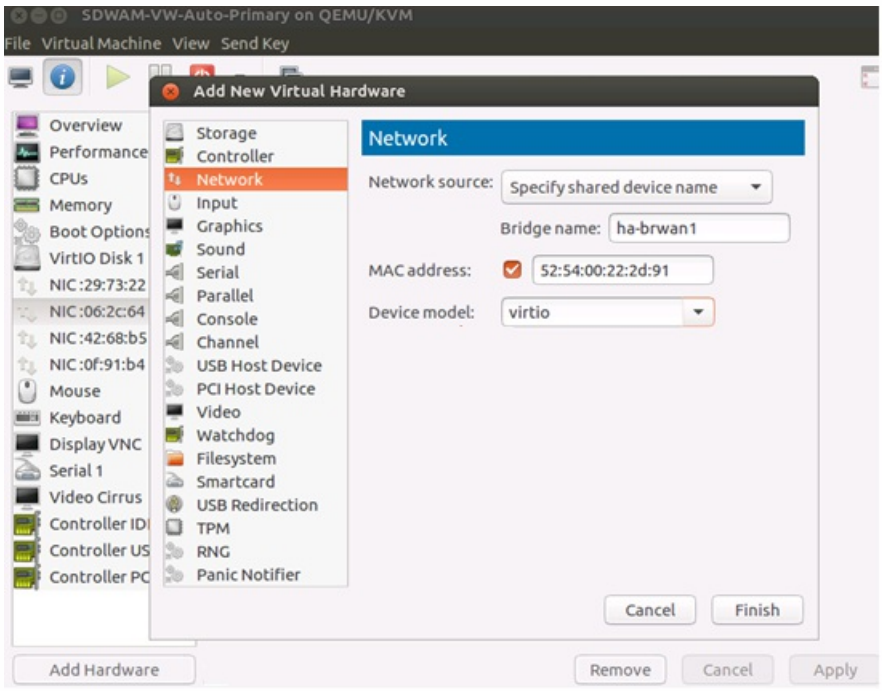

Limitation with MACVTAP Bridge mode type

With interface associated to Virtual Machines as MACVTAP Bridge mode type there are issues with shared MAC communication. SD-WAN Virtual WAN uses shared MAC (AA: AA: AA: 00:00: XX). When MACVTAP Bridge mode is used, ARP resolution does not occur for shared Mac. So MACVTAP Bridge is not a recommended option.

SD-WAN Standard Edition Virtual Appliance (VPX) HA Support for AWS

Aug 09, 2017

This procedure below describes how to deploy NetScaler SD-WAN virtual (VPX) appliances in high-availability mode on Amazon AWS cloud.

Points to consider when deploying SD-WAN VPX HA appliances in AWS Cloud:

1. AWS does not support GARP (Generic Attribute Registration Protocol), VLAN or L2 related functionality, such as promiscuous mode and bridging. This is because two VMs belonging to different customers can be scheduled on the same host sharing NICs.
2. L2 requires the switch appliance to be configured and these are not exposed to AWS users.
3. SD-WAN appliance HA model depends on GARP. When fail-over occurs, the new primary appliance sends GARPS out for VIP addresses.
4. AWS takes a new approach for HA failover. A new concept of ENI (Elastic Network Interface) is introduced. ENI is an entity which stands for Network Interface which has attributes like the IP address, MAC address, Security Group, and Port Rules.
5. You can move ENIs from active or inactive Instance to another active/inactive Instance.
6. The Instance needs to be capable to handling hot plug of interfaces.
7. Each Instance type has limitations on number of ENIs associated and number of IPs per ENI.
8. When an ENI moves all the attributes of the ENI MAC address, the IP address and Firewall Rules move with ENI.
9. AWS design for HA fail-over involves Instances communicating with external server to call Query API AWS servers.
10. The AWS servers are traditional HTTP servers. A request is sent from an instance to Query API server to get or post information regarding an Instance/subnet/VPC or any other attribute on the AWS.
11. For the cloud platform setup, the shared base MAC address configuration is ignored and has no significance.

How to create HA solution template in AWS

To create HA solution template in AWS:

1. Go to www.aws.amazon.com and log on with AWS credentials. After successful login, navigate to the **Services** > **Management Tools** > **CloudFormation**.

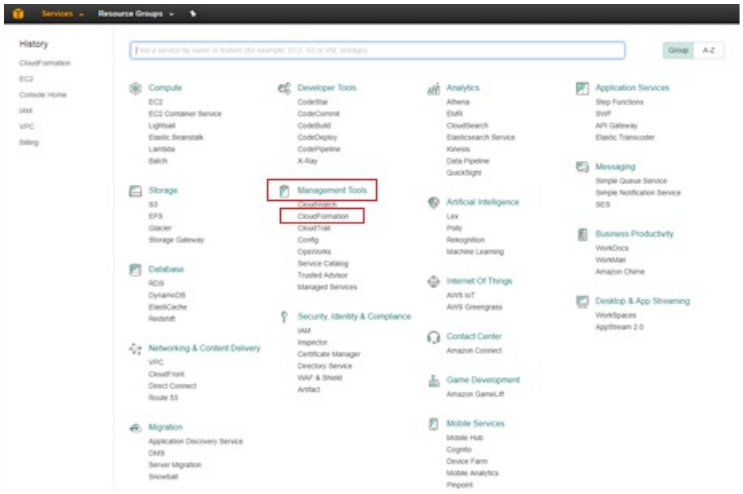

2. On the **CloudFormation Stacks** page, select the **Region** in which you want to deploy the NetScaler SD-WAN VPX instance, and then in the Create a stack section, choose **Create new stack** to create a new AWS CloudFormation stack.

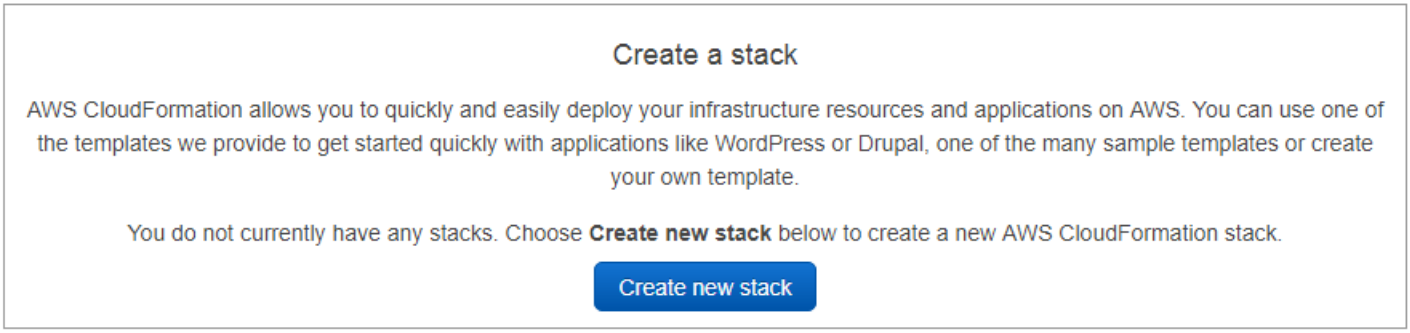

3. In the **Select Template** section, choose a template by:

- Uploading a template using **Upload a template to Amazon S3** option. (or)
- Specifying Amazon s3 template URL using **Specify an Amazon S3 template URL** option.

In both the cases, you will provide the Template or URL.

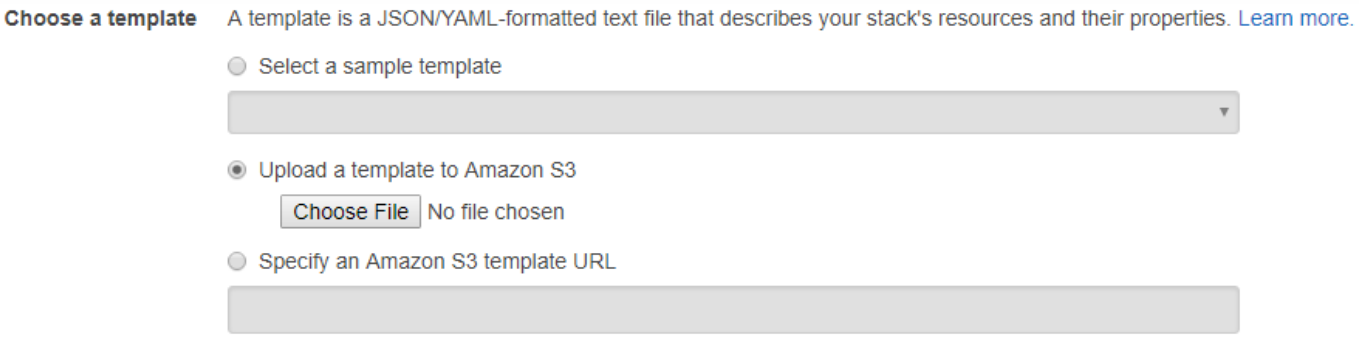

4. In the Specify Details section, specify a **Stack name**.

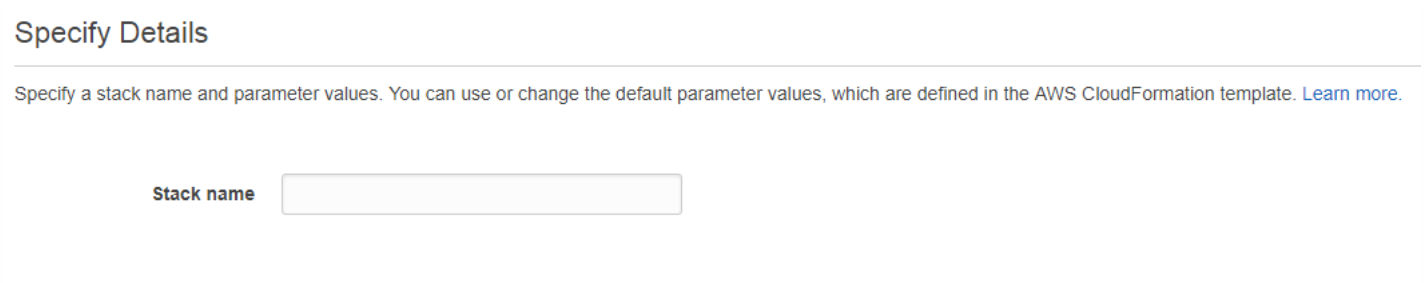

5. Configure **Virtual Private Network Configuration**. Fill in the below details as suggested. You can also find the tool tips besides each field.

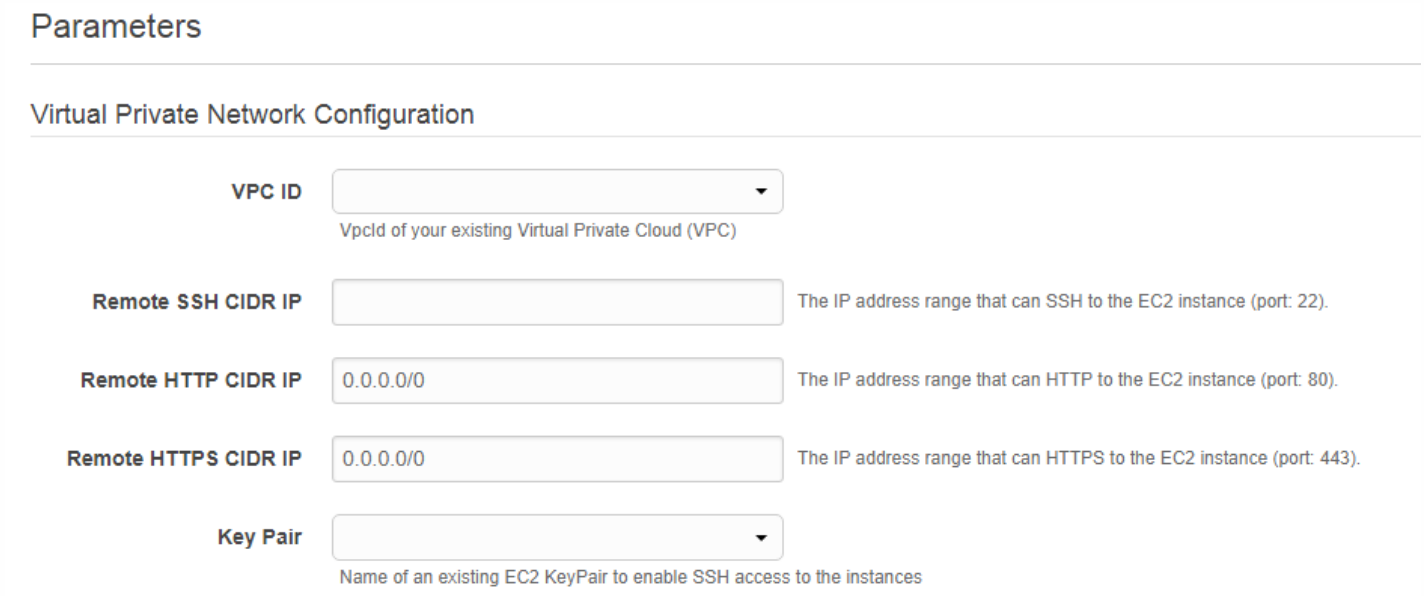

6. Configure Network Interfaces which should be attached to the instances created. Please note that the Primary IP's are for primary instance of HA pair and Secondary IP's are configured for secondary instance of the HA pair.

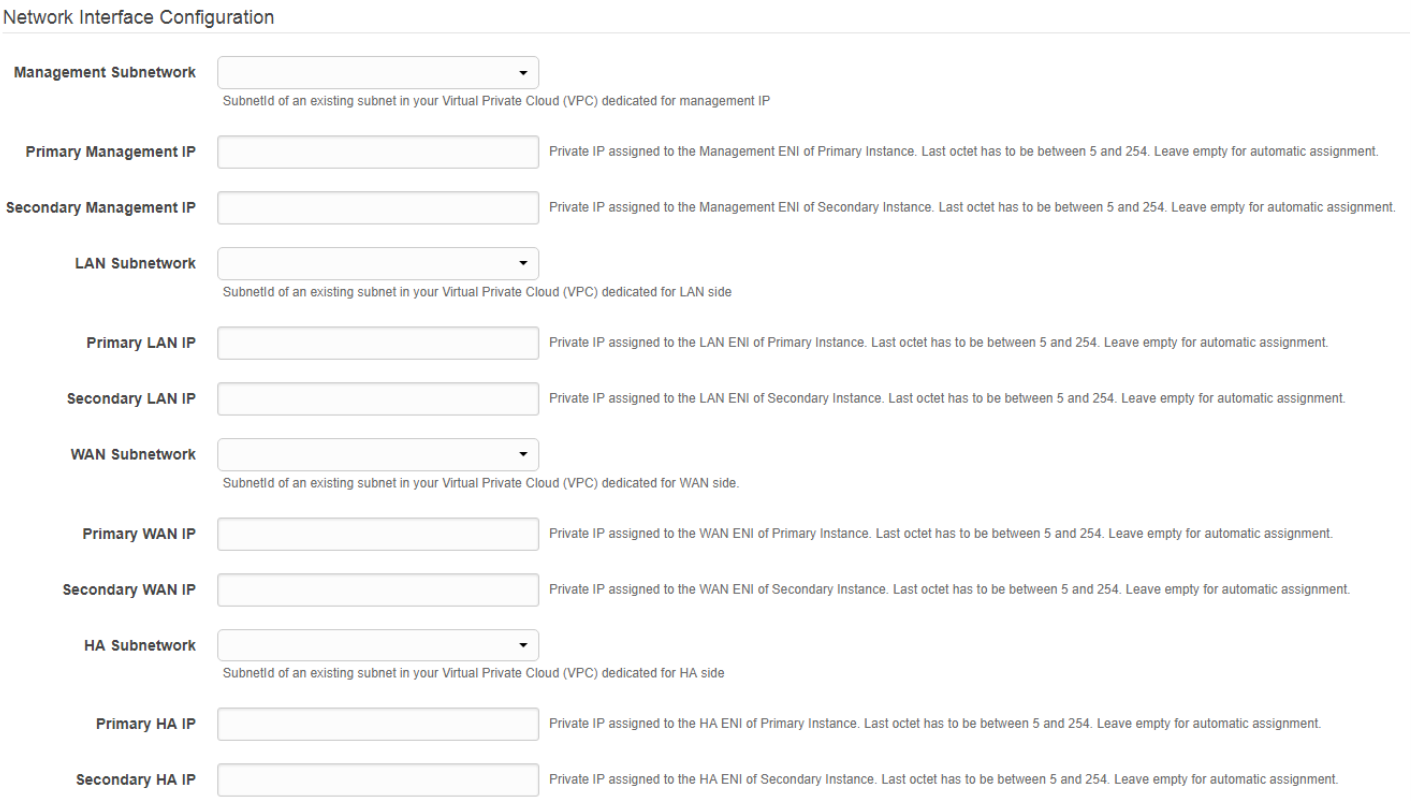

7. Configure Other Parameters such as **Instant Type** and **Tenancy Type** and click on **Next**.

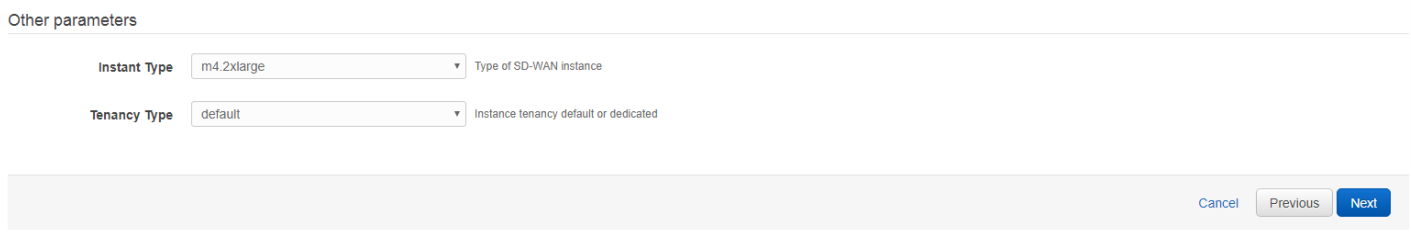

> ## Note
>
> If any validations fail, AWS will notify you and would not let you proceed until the errors are resolved.

8. Set Tags. These are AWS specific options which are user configurable.

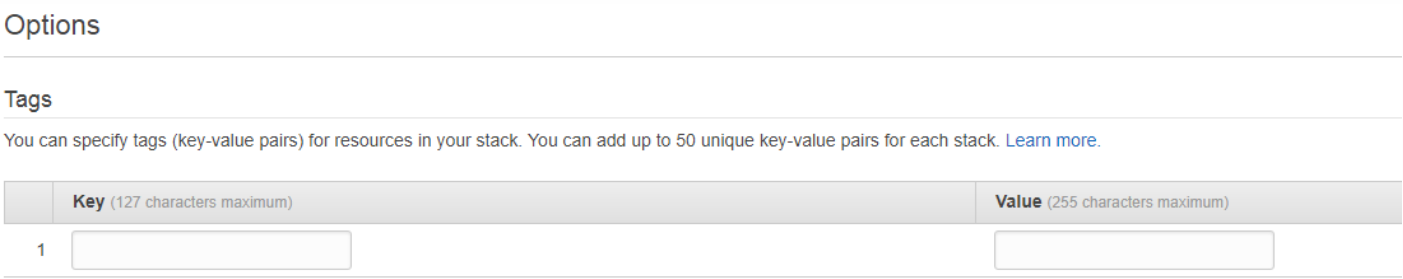

9. Configuring IAM role is not recommended. This is already created by the customized IAM role, which is done through

the cloud Formation Template.

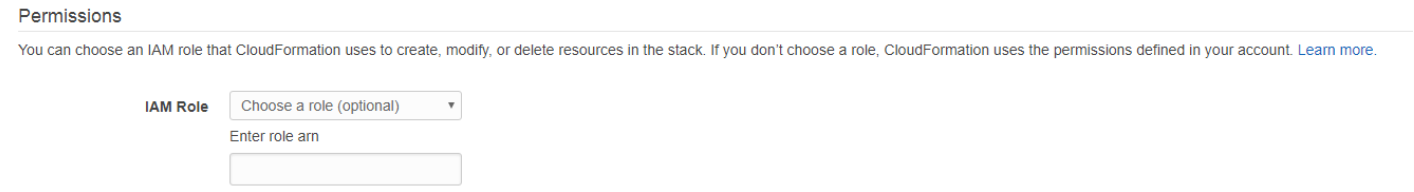

10. After clicking next, Review the template and acknowledge the custom IAM role which has been created by CloudFormation template. Proceed with **Create**.

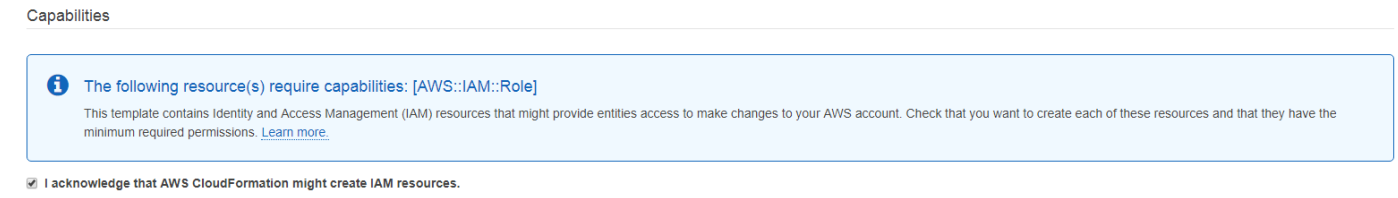

11. The new stack that you created appears on the CloudFormation Stacks page. After successful template upload, Monitor the status of the template.

12. Monitor the events of all the resources created by the CloudFormation template. In case of any failure, detailed description of events are generated by AWS which helps in debugging the issue. The Events appear as follows:

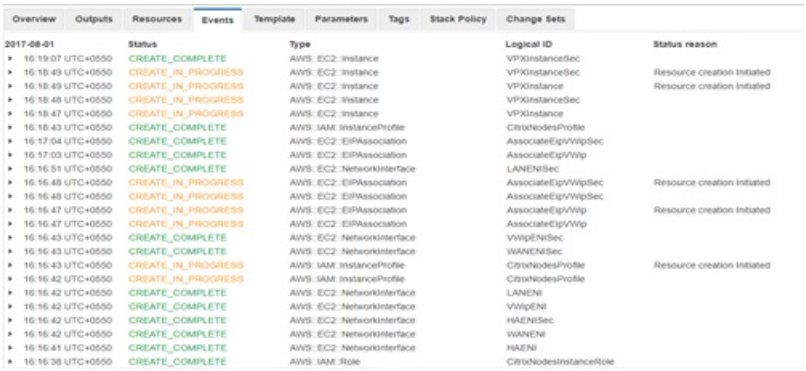

13. After successful stack creation, the status of the template should appear as **Create_Complete**.

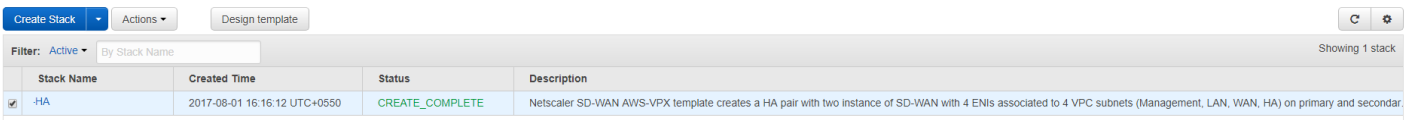

14. Navigate from AWS console to **Services** > **EC2** > **Instances**. You should see two instances **SDWANPrimary** and **SDWANSecondary** instances created, up and running with Elastic IP's associated with the instances.

15. Select **SDWANPrimary** instance. You should notice all the resources rightly assigned to instance, Security groups, Elastic IP, IAM role and four Network Interfaces. Failed to create any HA functionality may not work as expected.

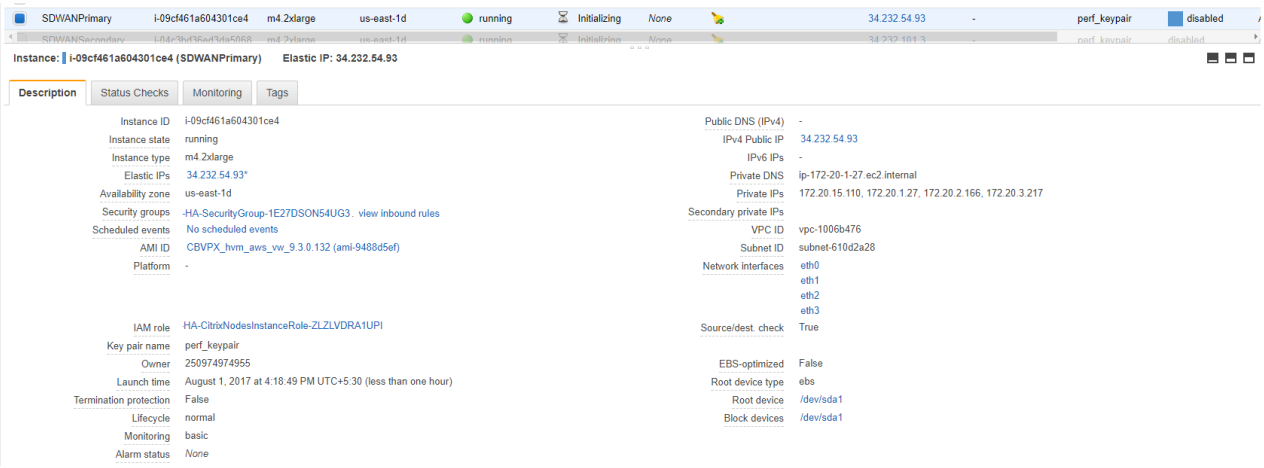

16. Similarly select **SDWANSecondary** instance and verify the above resources.

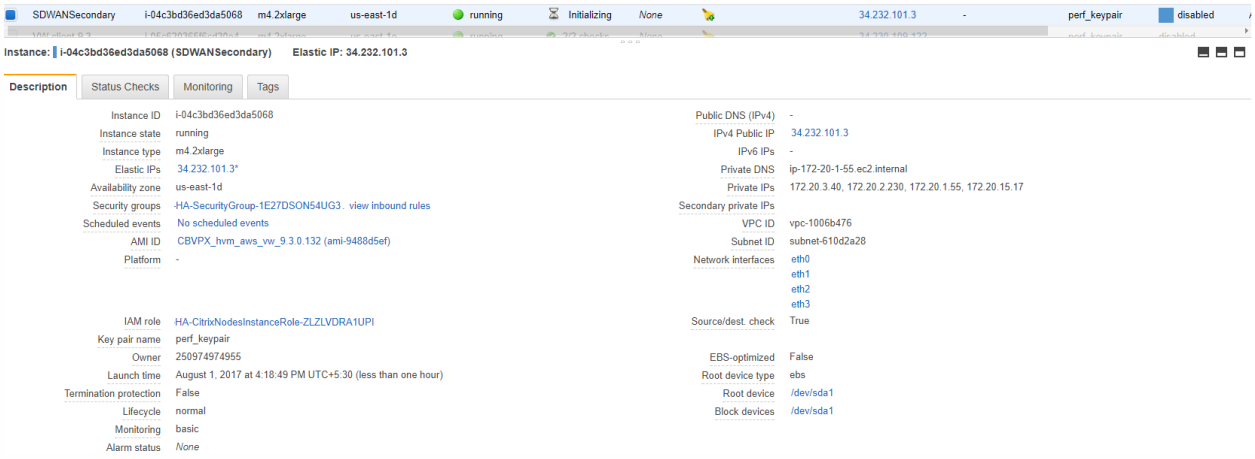

How to Configure HA Fail-Over for any SD-WAN Instance Running on AWS

Set up HA peers with one HA peer with three or more ENIs, and 1 HA peer with equal number of ENI's. In both Peers, the first ENI is dedicated to Management. One HA peer owns all Traffic ENI's. During a Failover, the traffic ENIs move from the failing instance to the new Primary instance.

For example; it can take upto or more than 20 secs to move two traffic ENIs. AWS do not have SLAs on API response and you cannot have one for HA failover time.

> ### Note
>
> The AWS design has a limitation of instances dependent on the AWS servers to respond for attach and detach. The failover time is unpredictable.

Configuration Steps

1. Acquire information about your HA Peer Instance about information on the number of ENIs associated and details of ENIs associated using REST API.
2. Detect the condition of the failing instance.
3. Call Detach of ENIs from failing instance using REST APIs.
4. Ensure all ENIs associated are detached.
5. Attach ENIs to the current Primary instance.
6. Ensure All ENIs are attached.
7. Trigger upper layers to detect that new ENIs are in place.

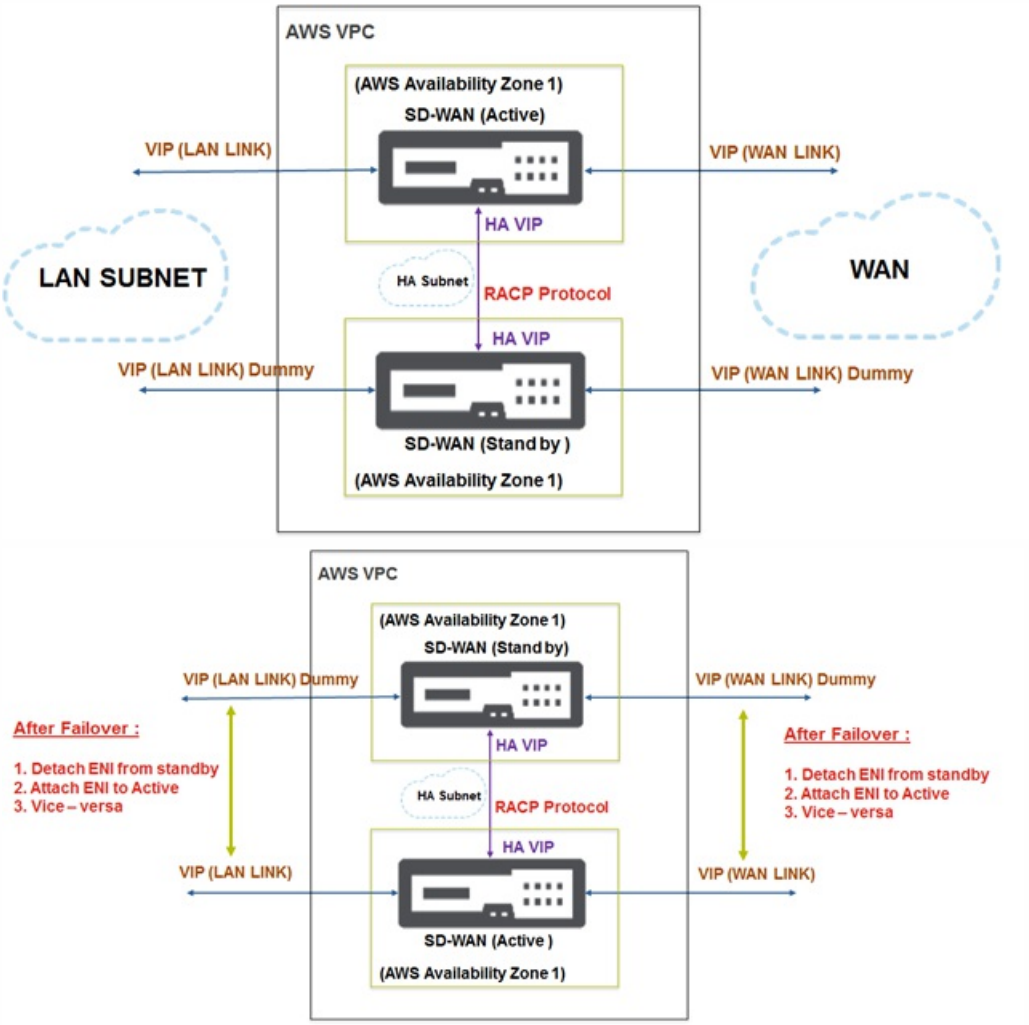

How to Configure SD-WAN VPX-SE in a Single AWS Virtual Private Cloud (VPC) Subnet or Between Regions with Public WAN Link IP Address

In AWS VPC, for an active SD-WAN instance, another high available SD-WAN instance running in the same VPC is released.

1. The links configured are the same between active and stand-by SD-WAN appliances.
2. For AWS, you can create a new subnet and a dedicated link for RACP protocol to communicate between the SD-WAN appliances.

3. In the SD-WAN GUI, configure the following:

a. Create an interface group. Name it as HA-LINK. Add the interface used for HA.

b. Create a Virtual IP address for Interface group.

c. In High Availability Node, Enable HA and add control Virtual IP's which RACP protocol uses for communication. Ensure that the IP addresses are same as the configured IP addressed while creating network interfaces in AWS.

d. Perform Change Management and download the active configuration for the stand-by SD-WAN appliance.

e. After applying configuration through local change management on the stand-by SD-WAN appliance, you will see heartbeats exchanged between active and stand-by SD-WAN HA appliances.

f. When failover occurs, you will see SD-WAN appliance transitioning from stand-by to active modes and/or vice-versa without any configuration loss.

Note

1. AWS supports HA mode with features such as Elastic Load balancing and auto-scaling where the challenge is to sync configuration within the SD-WAN appliances. In this deployment you leverage the already existing RACP protocol for efficient HA.
2. Both MCN and branch site appliances can be made available in the cloud environment.

SD-WAN Standard Edition Virtual Appliance (VPX) in Hypervisor on HyperV 2012 R2 and 2016

Apr 12, 2018

To install SD-WAN VPX-SE in hypervisor on HyperV 2012 R2 and 2016:

1. Install **HyperV Manager**. For more information, see documentation at Microsoft.com.

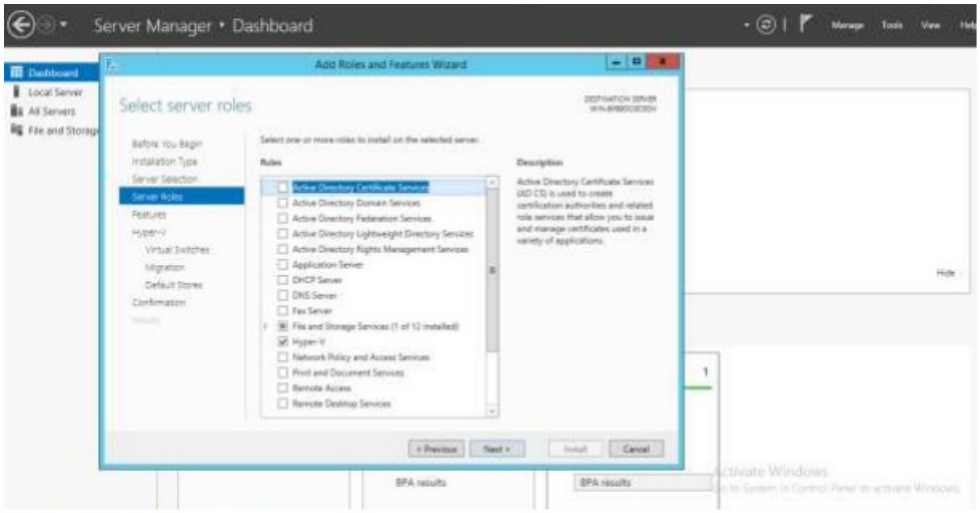

2. In the **HyperV Manager** window, go to **Virtual Switch Manager,** and configure interfaces in the following order; management, LAN, and WAN.

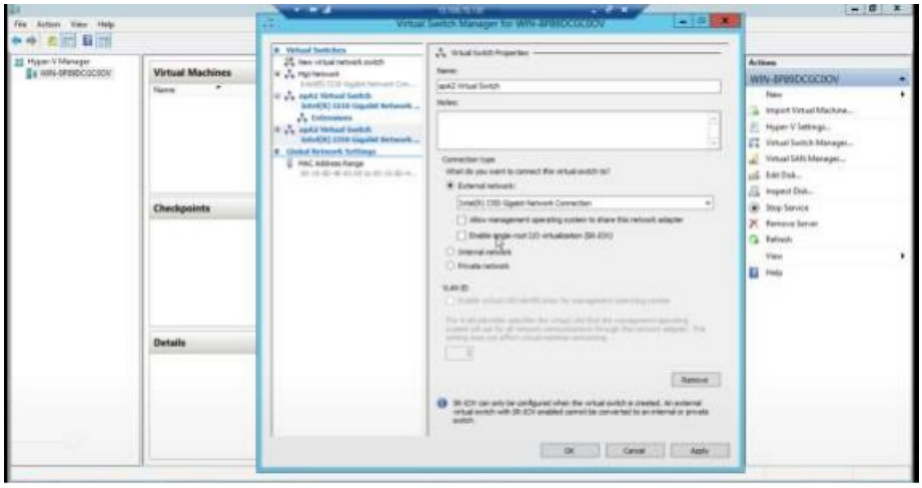

3. Download the ***hyperv.tgz*** file and untar it.

4. Import VM using the above extracted *vhd* and assign the number of CPUs and memory accordingly. Add interfaces in the order (management, LAN, and WAN). Enable Mac spoofing on LAN and WAN interfaces. Go to **Settings** > **Interface** >**Advanced features**.

Adding interfaces

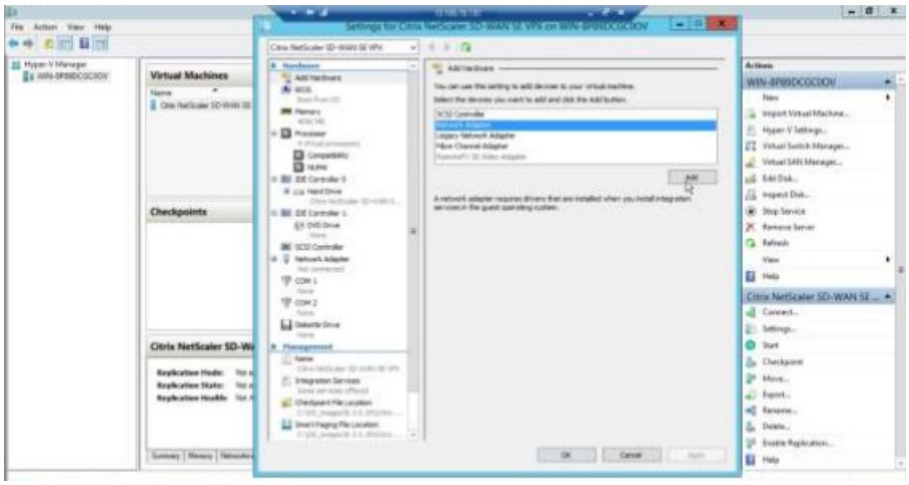

Enabling Mac spoofing on the interfaces (LAN and WAN)

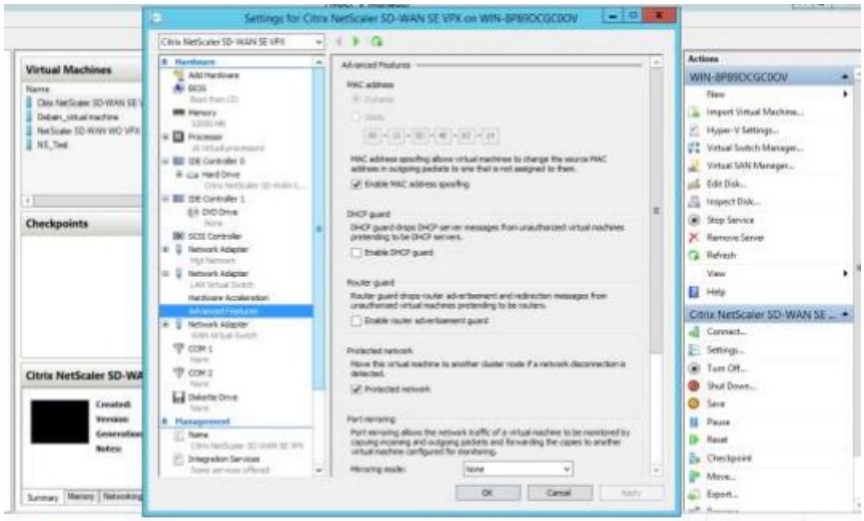

5. After the VM is up, assign free IP address. The VM can be accessed after assigning the IP address.

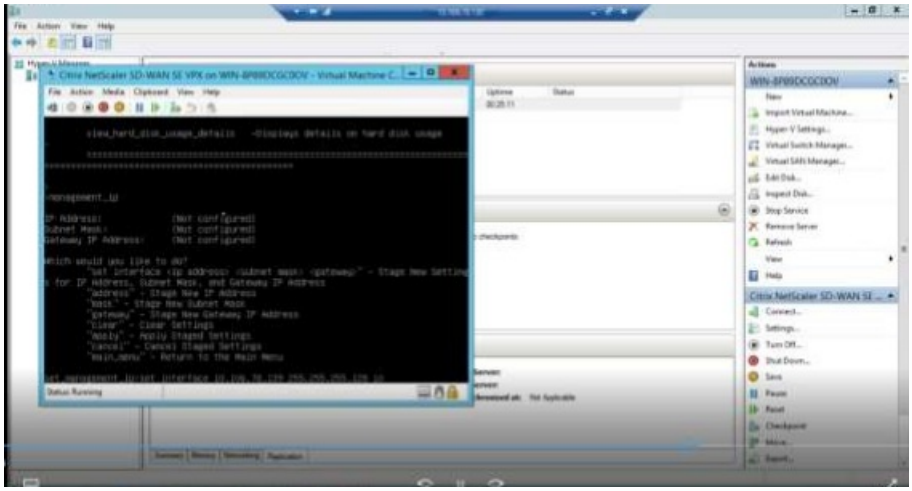

Note

The qcow image downloaded must be present under the default folder /var/lib/libvirt/images. If this image is downloaded, and used in different folder in KVM, there can be issues when disk size expansion is performed.

Limitations for deploying SD-WAN VPX-SE in HyperV 2012 R2 and 2016

- VLAN Tagged Trunk Deployment is not supported.

SD-WAN Standard Edition Virtual Appliance (VPX) HA Support for Microsoft Azure

May 03, 2018

The SD-WAN solution template is a unified template in Microsoft Azure that allows users to deploy both Netscaler SD-WAN Standard Edition appliance or an HA cluster of Netscaler SD-WAN appliances. For HA to work and to create the HA cluster creation using solution template, the user or administrator should create a registered application with the role of an owner. The user then obtains the key for registered application with the application ID. After the application is registered, the KEY for the application is displayed only once after creation. The user has to store the key since it needs to be uploaded as input for the HA solution template. The Application ID and the Subscription ID can be obtained anytime which are also potential inputs to create SD-WAN solution template for HA.

The registered application is used to populate the LAN routing table based on HA convergence to make sure that LAN always points to the latest active appliance as the next-hop for reaching remote sites through WAN.

The following sections describe the workflow to create solution template in Microsoft Azure and configure HA in SD-WAN GUI.

- Register application - obtain application ID, application Key, and the Directory ID which is provided to create solution template for HA deployment.

- Create solution template.

- Configure HA in SD-WAN GUI - Assign Virtual IP addresses and interfaces as required for LAN, WAN, and HA control exchange.

How to Register Application in Microsoft Azure

To register the application:

1. Log into the **Azure Active Directory**.

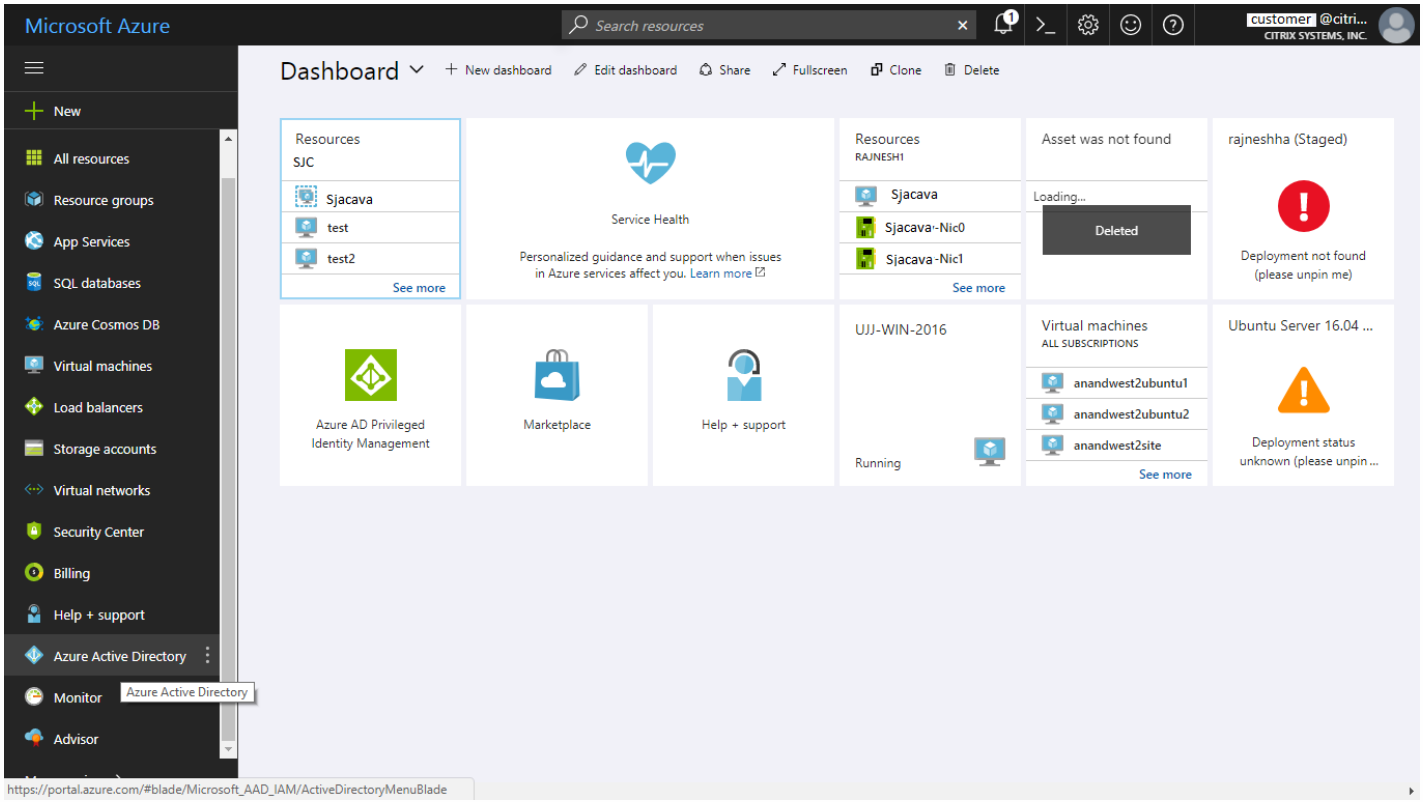

2. Select **App registrations**.

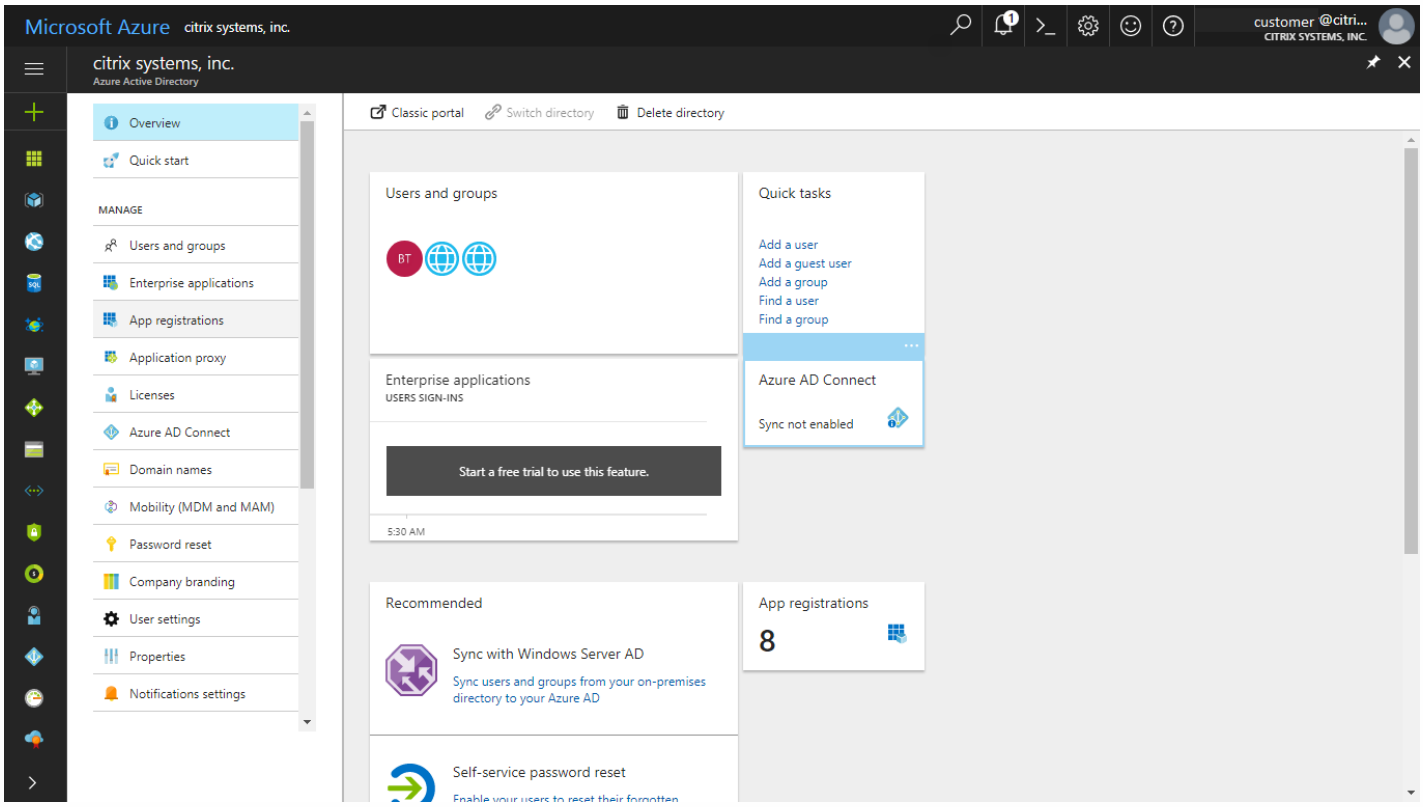

3. Click on **New application registration**.

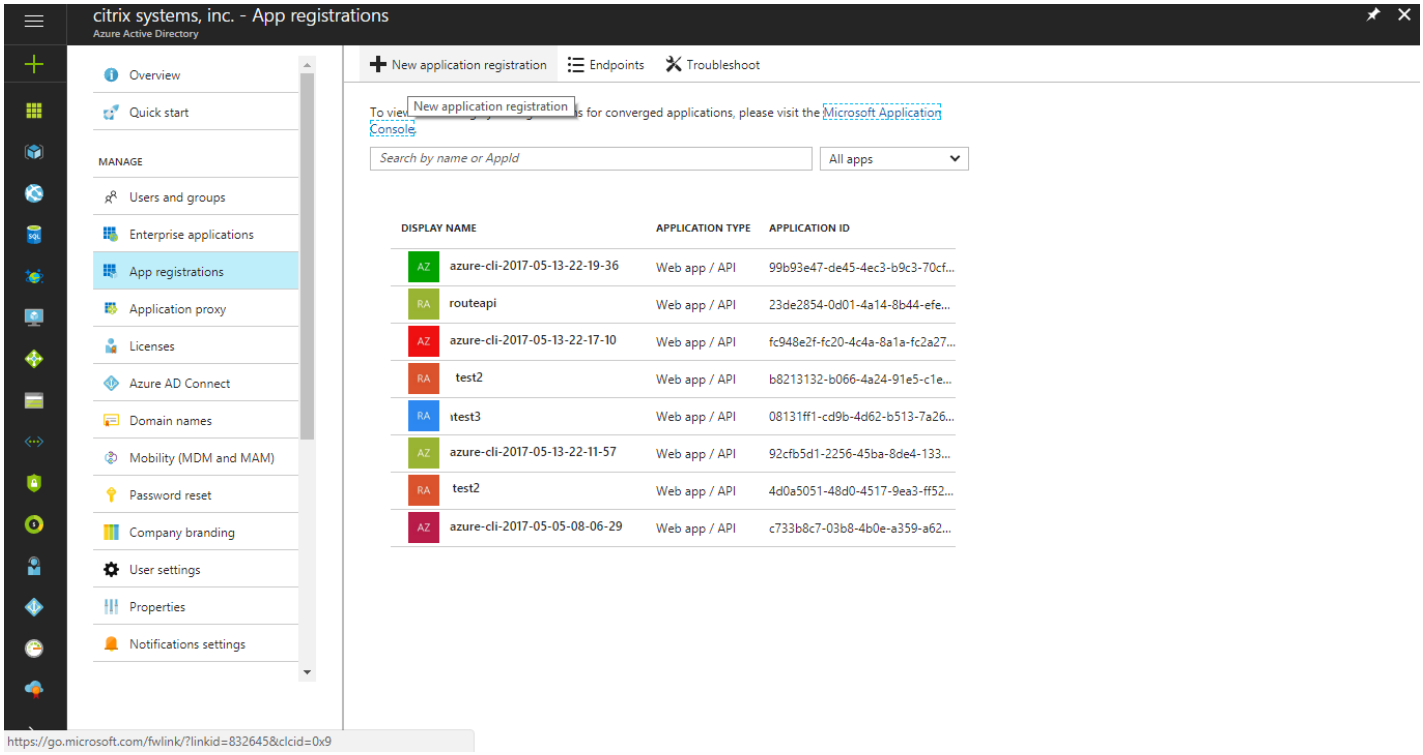

4. Provide a name for the application. Select **Application type** as **Web app/API** and the **Sign-on URL** can be any dummy URL.

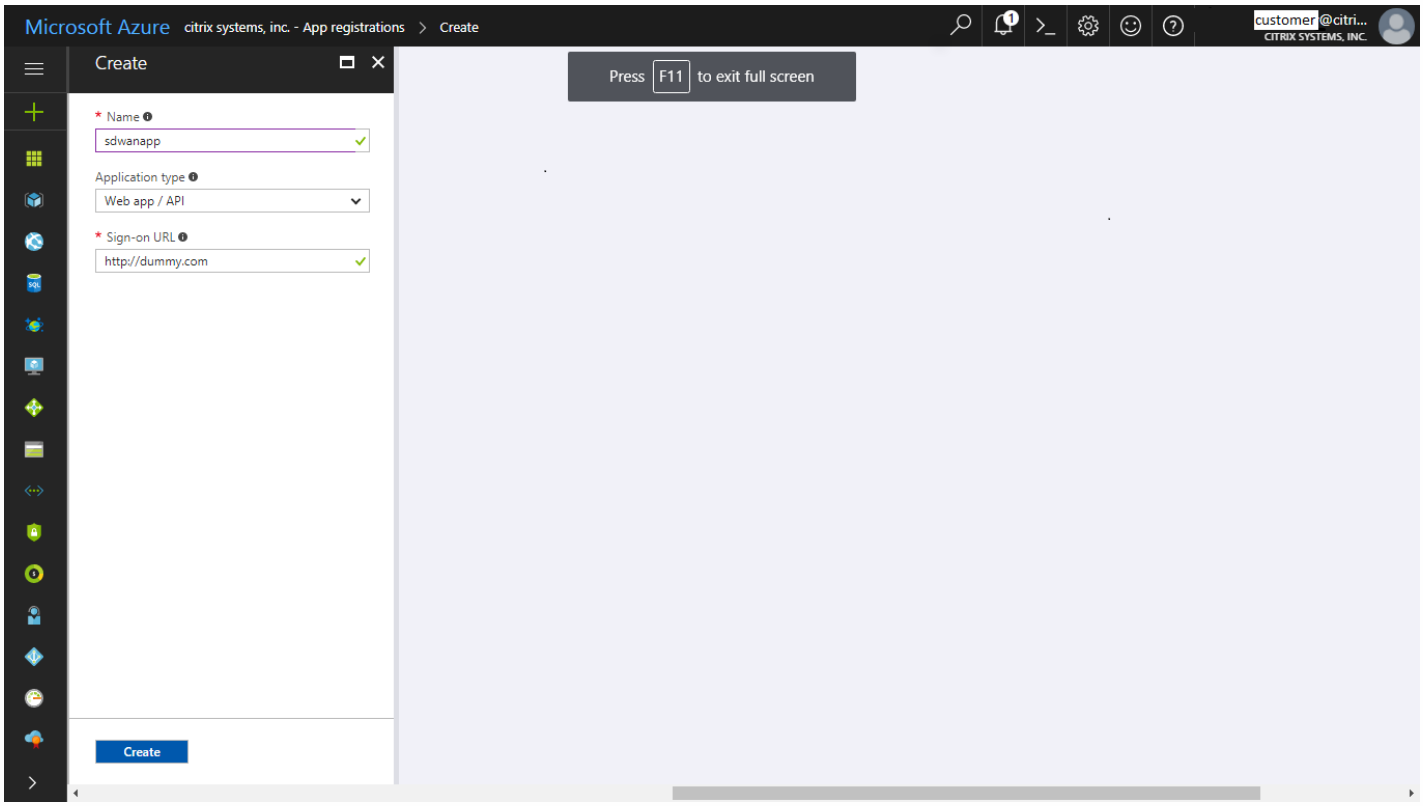

5. Observe that after step 4, the new application is created and registered. An application ID is generated for the newly created application. You should store the application ID so that it can be presented when creating the solution template for HA creation.

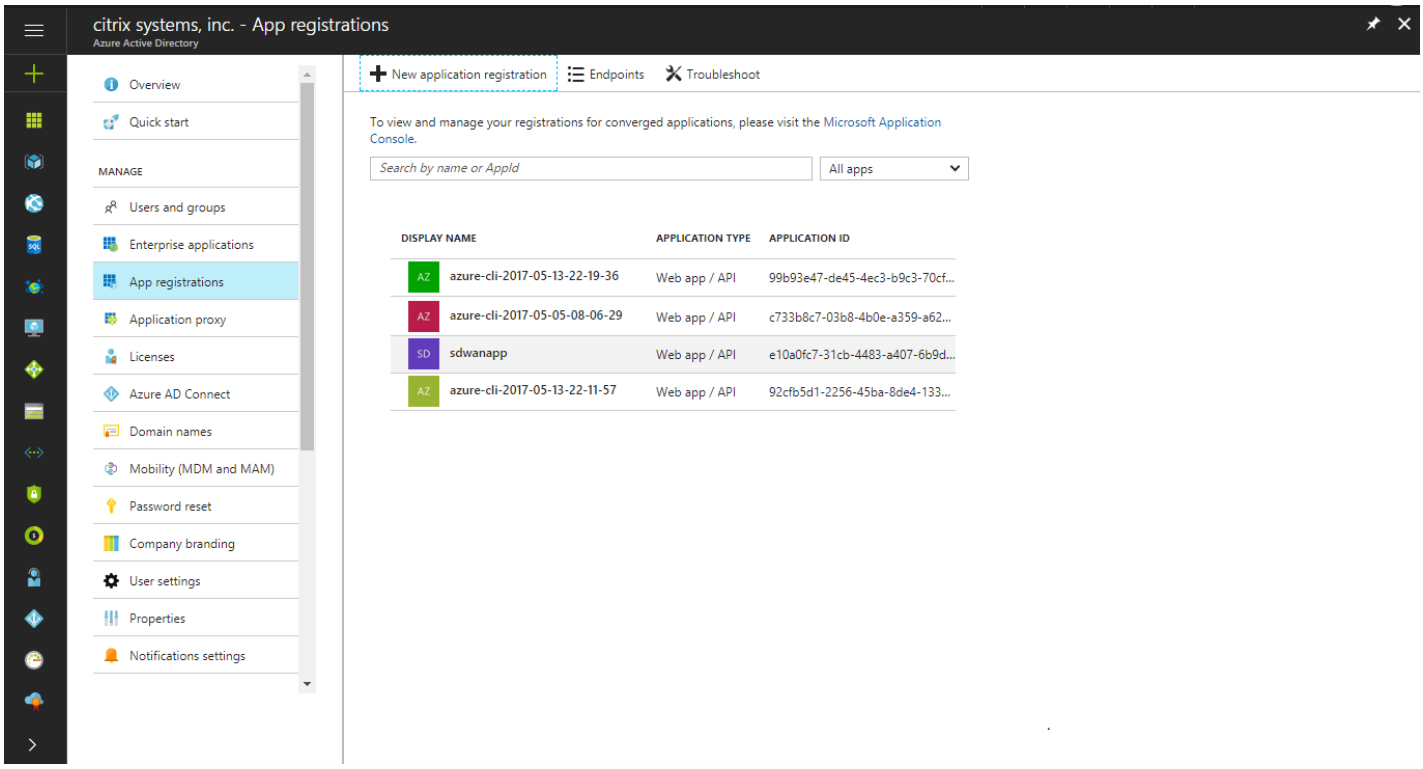

6. Go to **App Registrations** > **Newly created app** > **All settings** > **Keys.** Then create a key description and select **Never expires**. Save your selections.

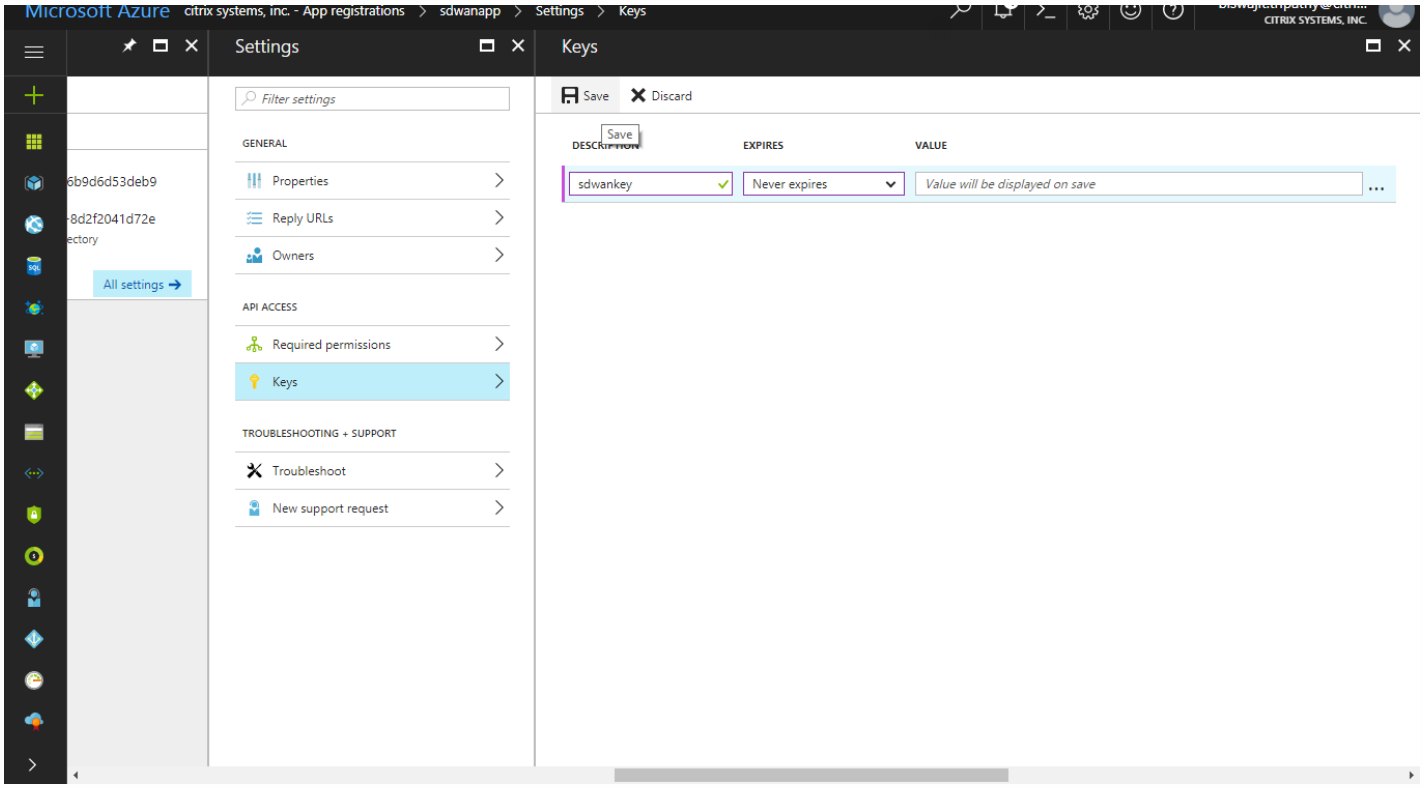

7. A KEY value that should be kept safe is displayed. You would need this as an input to be presented while creating solution template for HA deployment.

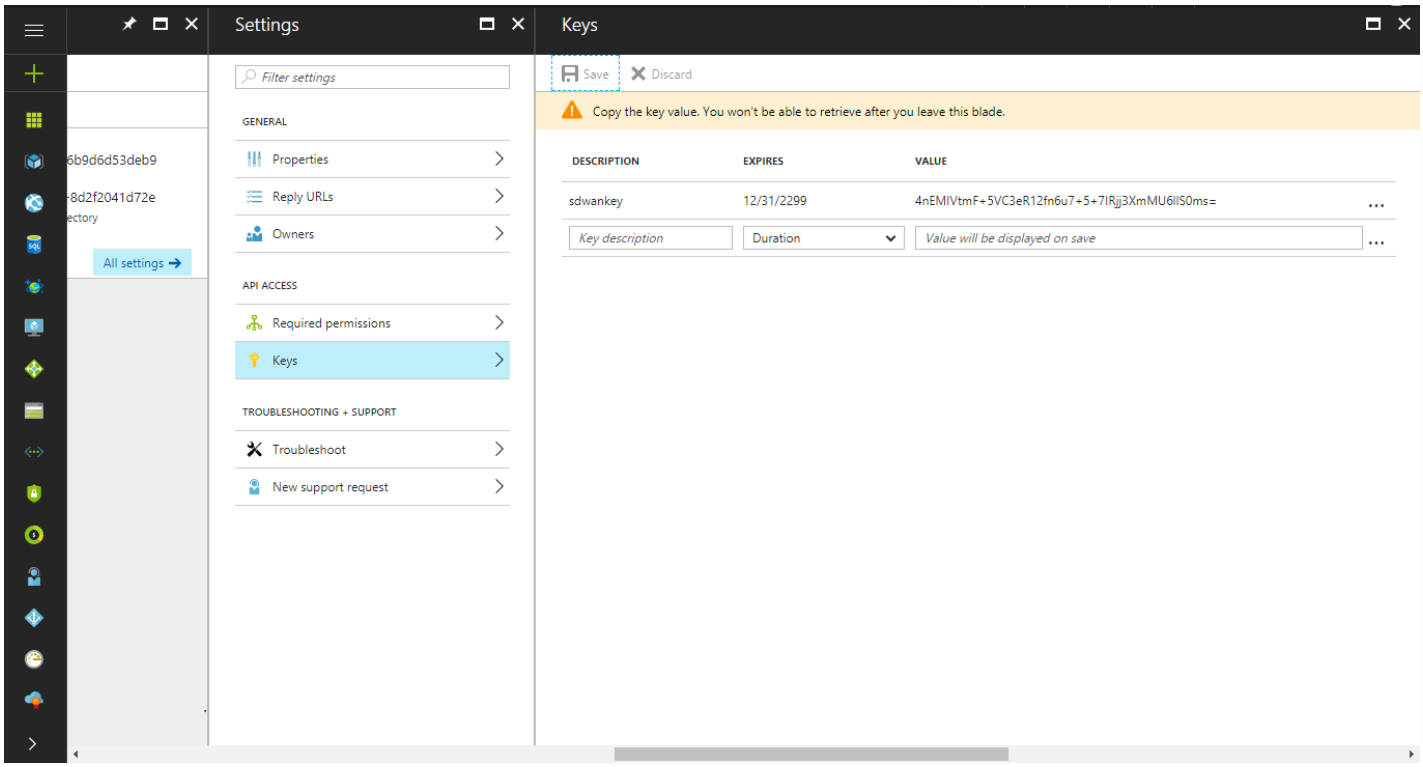

8. Go to **Azure Active Directory** > **Properties** for the **Directory ID** attribute. Store the directory ID and provide it while creating solution template for HA deployment.

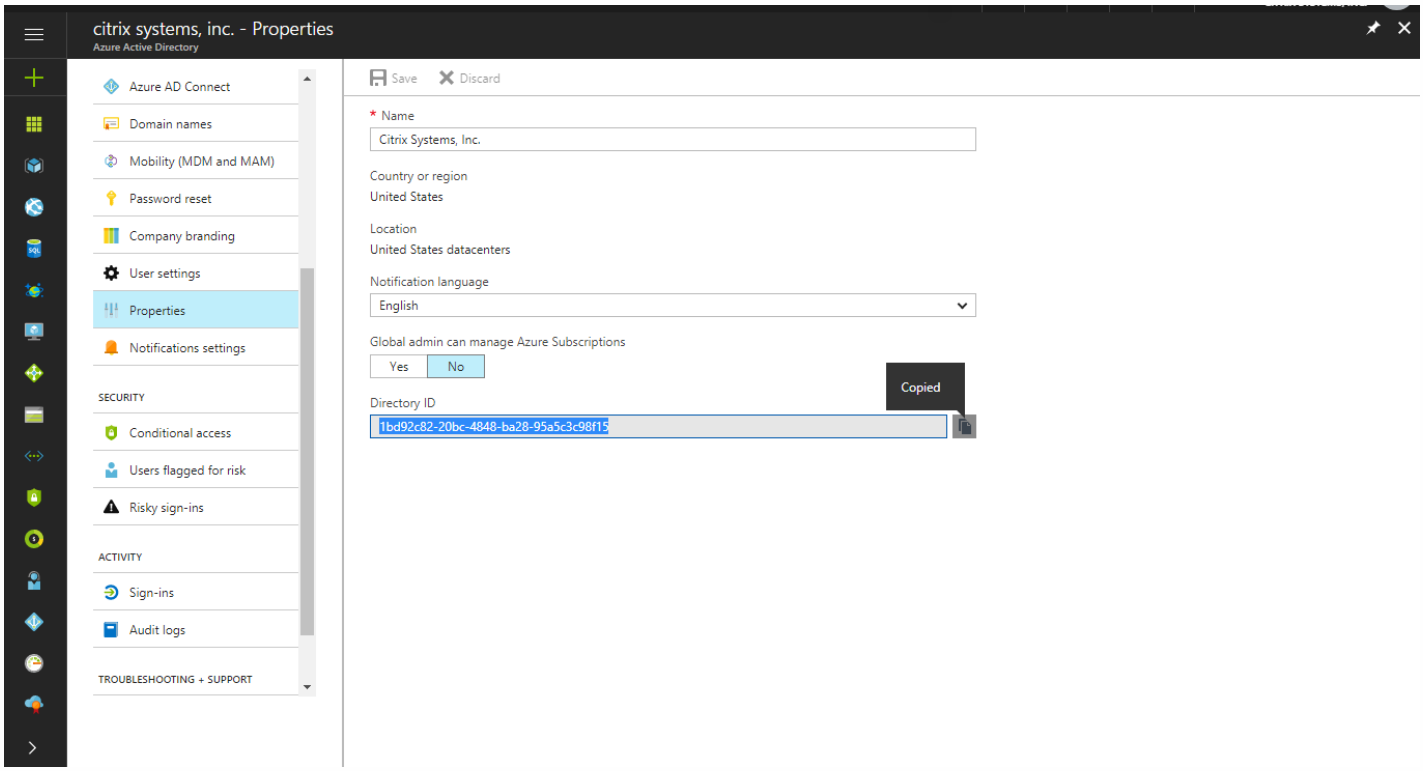

How to Create Solution Template for HA Deployment

To create solution template:

1. In the **Configure basic settings** page, provide the **Resource group** name and the **Location** where you want the

resource group to be created.

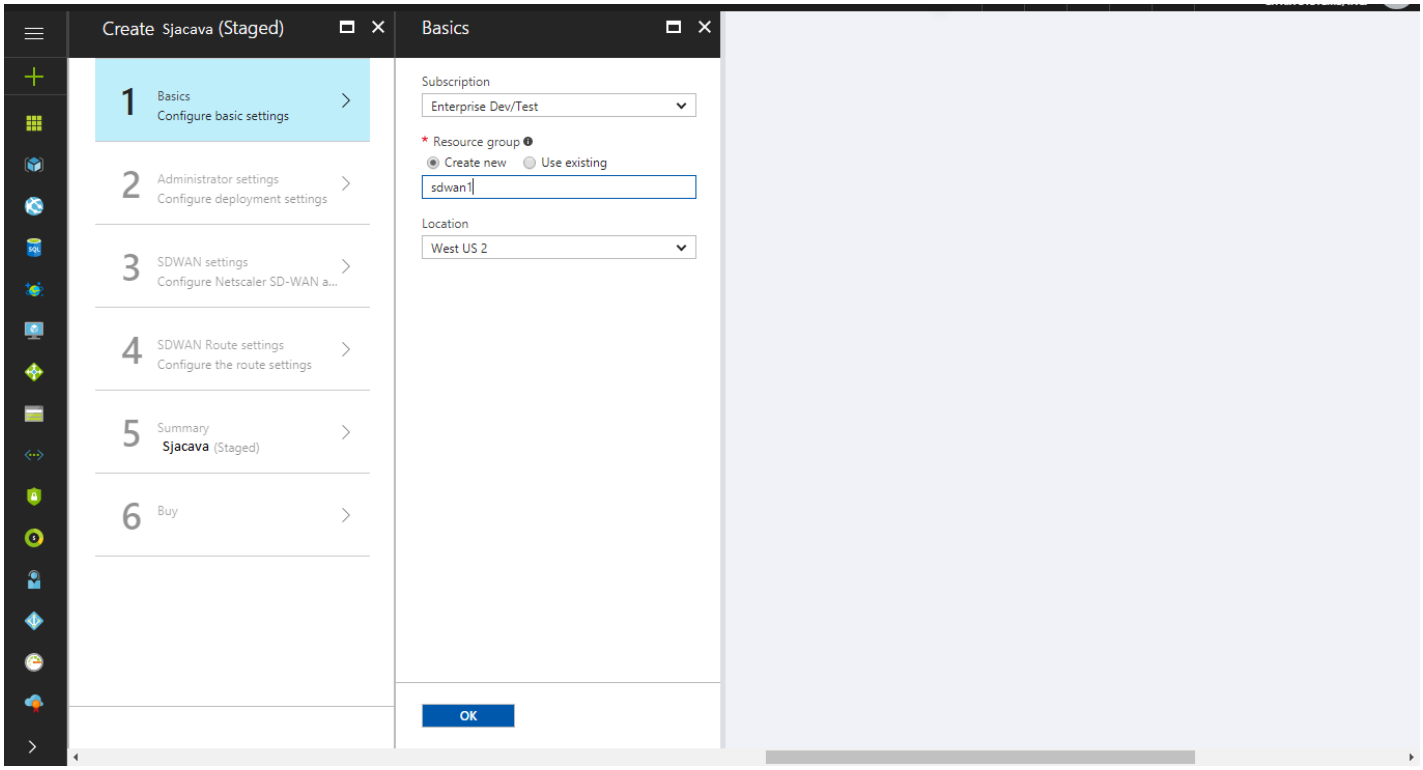

2. Navigate to **Administrator settings** page to configure deployment settings. Provide a name for creating the HA virtual machine. On the Virtual Machine name page, the primary instance is created, and the secondary instance is auto-created and suffixed with Ha. For HA deployment, you need to enable it in **HA Deployment Mode**. Provide a name for **Availability Set**. Create **Username** and **Password** of choice. **Confirm password** as shown in the figure below.

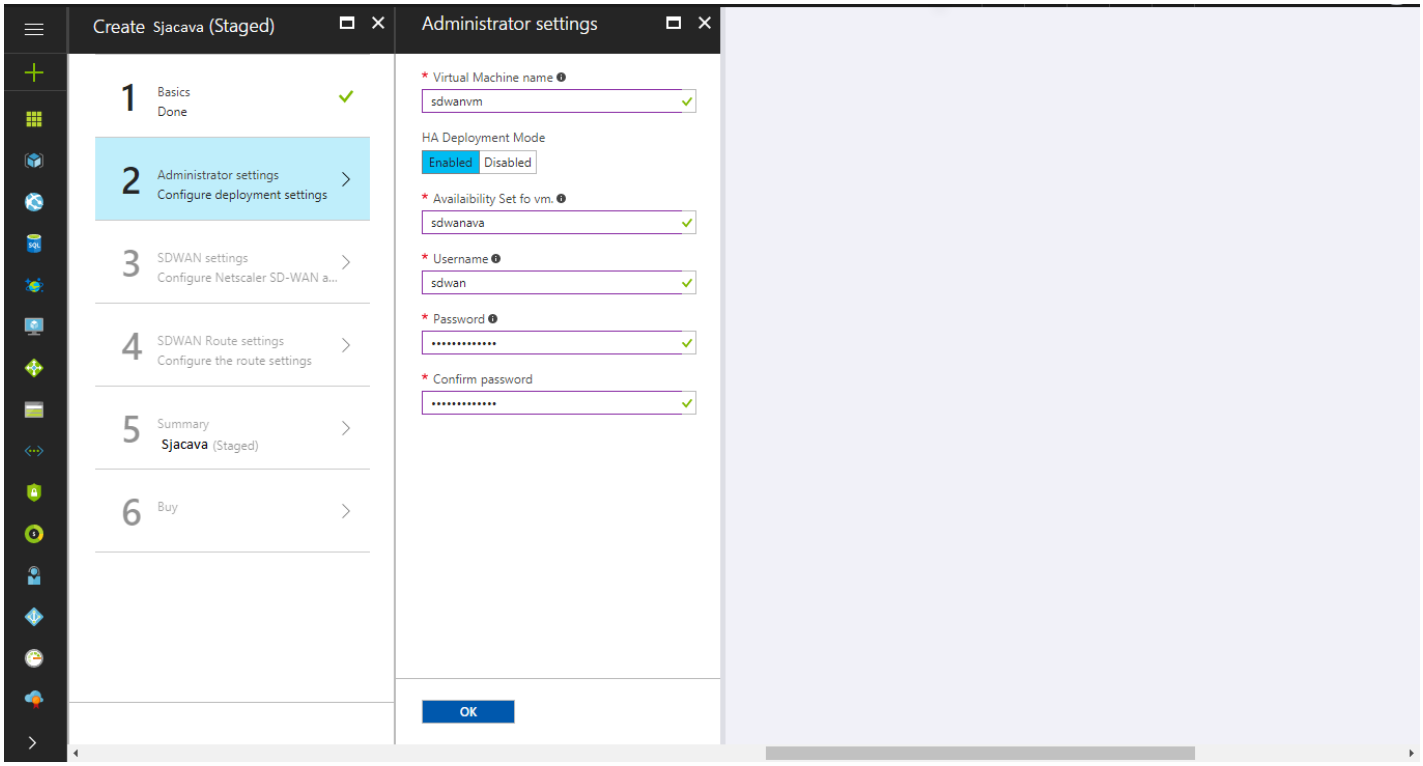

3. Go to **SDWAN settings** to configure NetScaler SD-WAN. This allows you to use existing storage or create a storage in the resource group.

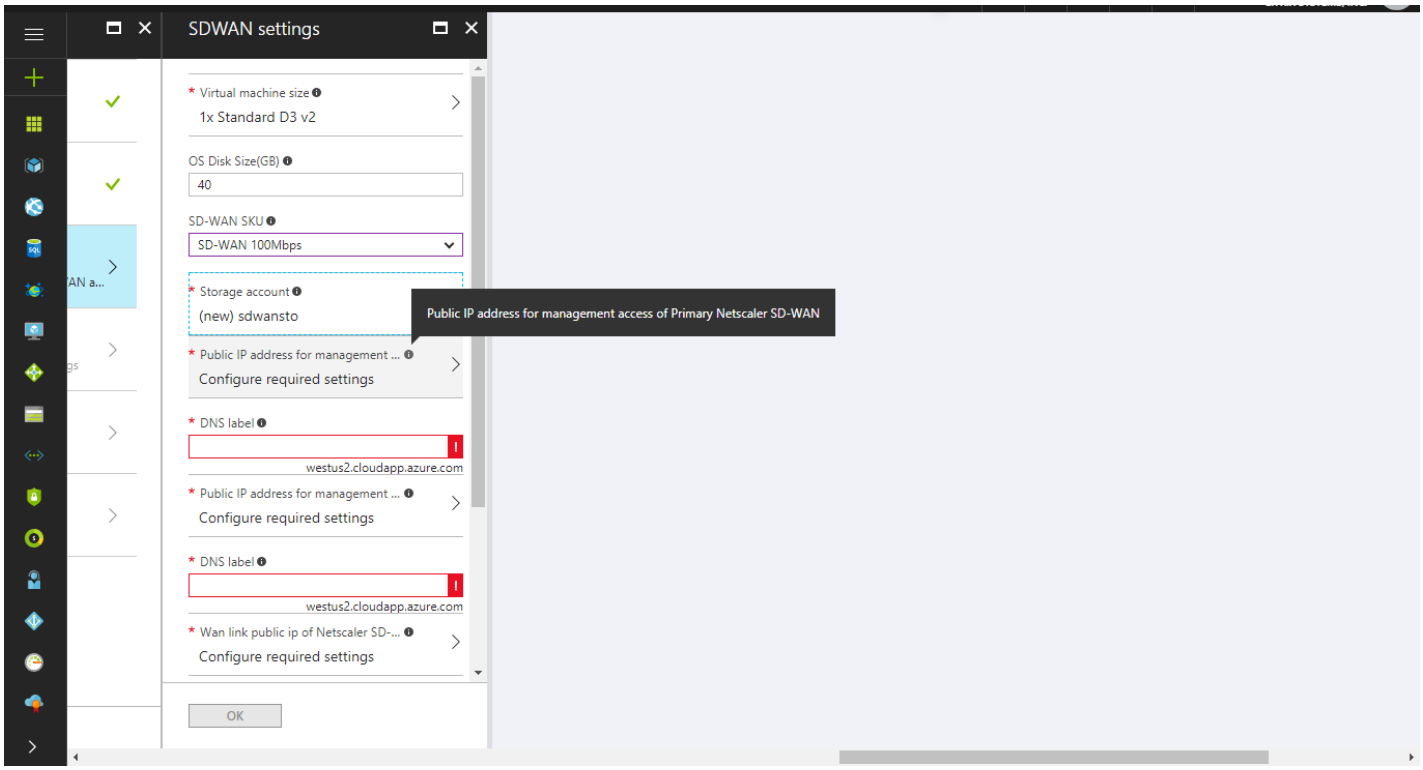

4. After creating storage account, the public IPs and the DNS (optional and can be any text but unique across all DNS sections) for Primary HA appliance, Secondary HA appliance, and the Azure Load Balancer should be provided. The figure below displays how to create Public IP by providing a unique string. Create the Assignment as "**STATIC"** so that the IP is retained even after reboot. This is recommended for HA.

a. Provide the DNS name as some unique string after providing the public IP for management of Primary Netscaler

SD-WAN.

b. The next immediate two sections are used to provide public IP for management of secondary Netscaler SD-WAN appliance and a unique DNS name for it. This DNS name should be different from the other two DNS names asked to be entered by the administrator.

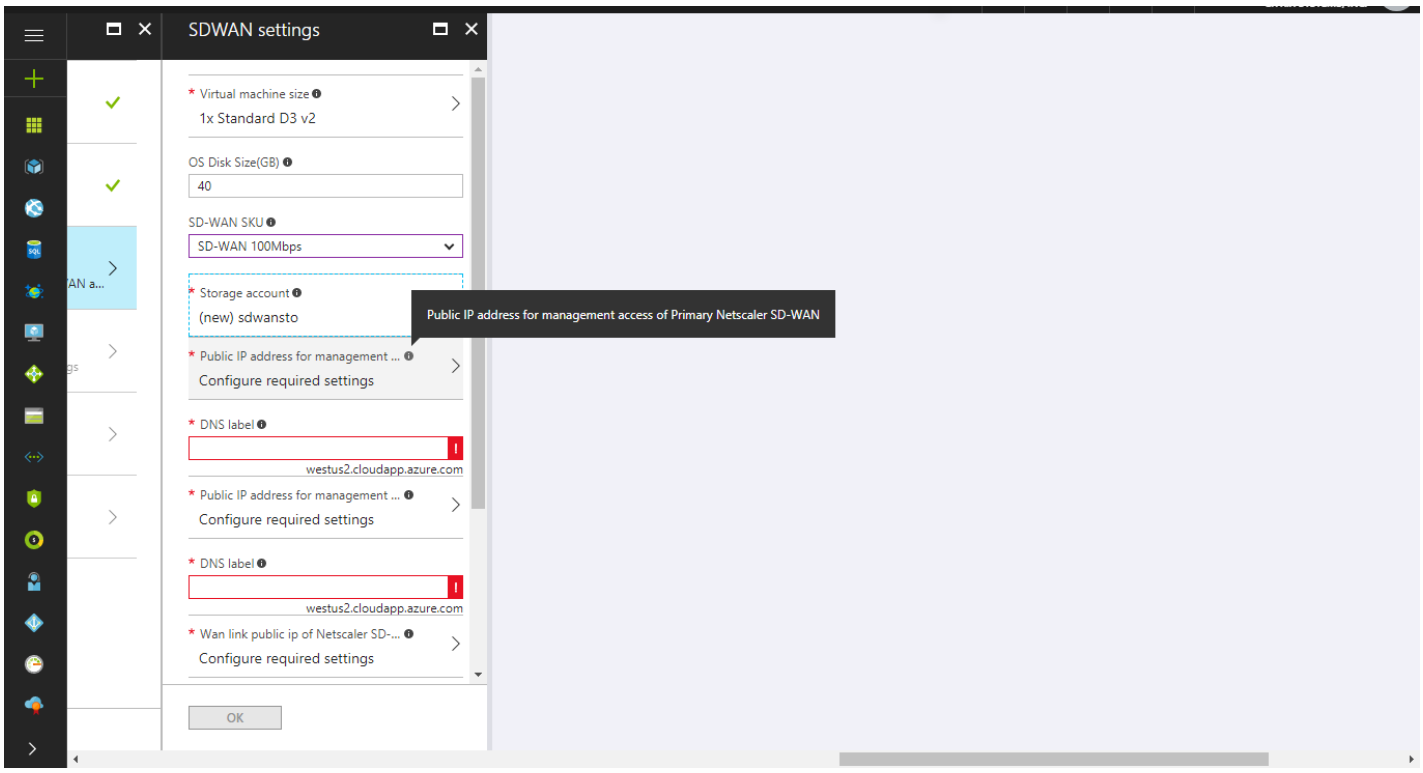

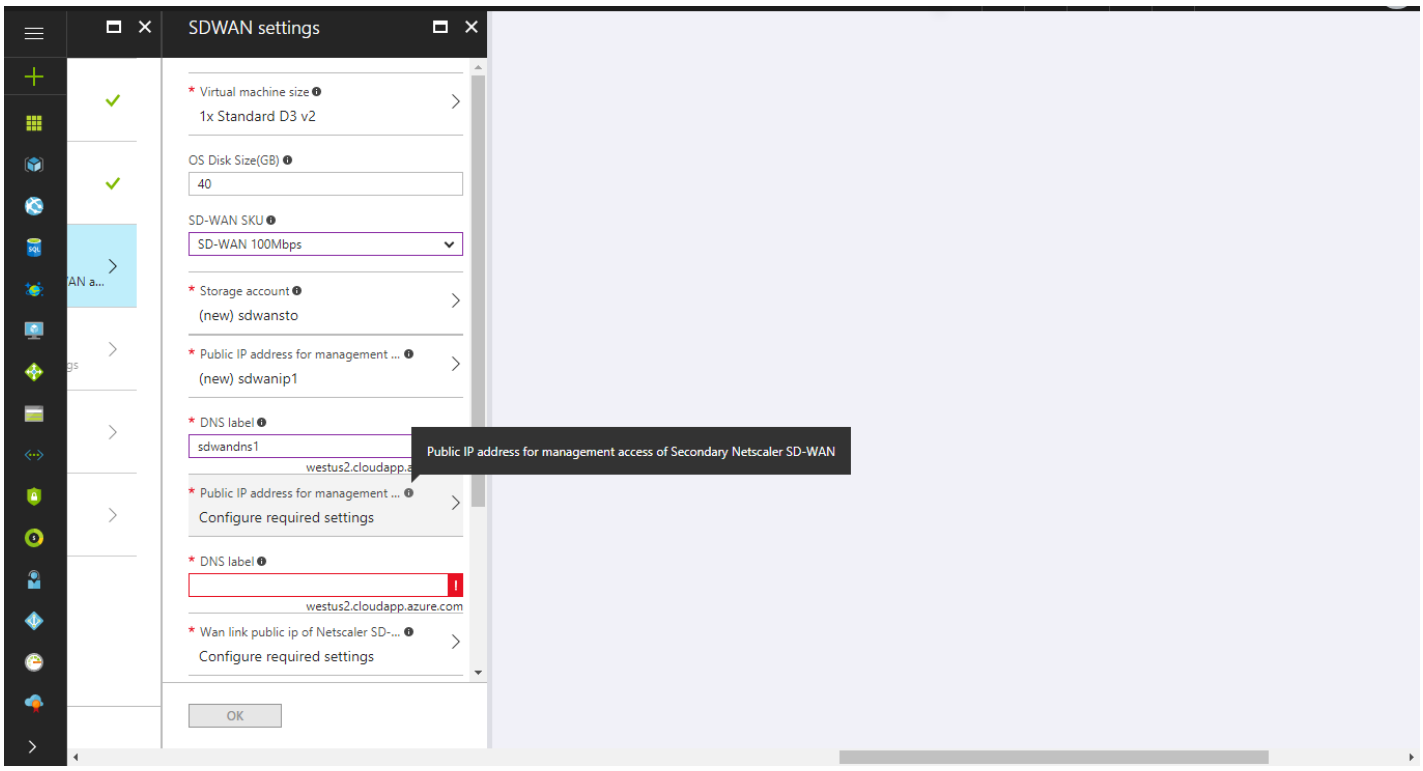

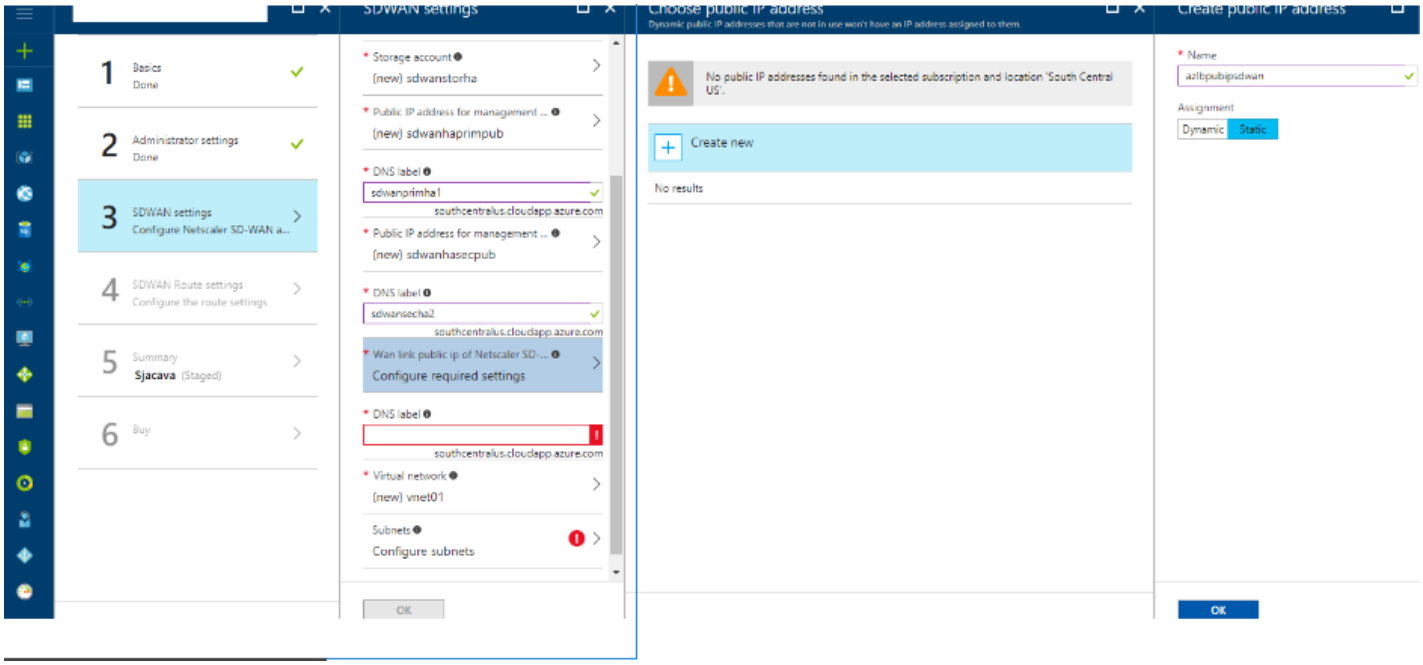

5. Create public IP for the Azure Load Balancer that governs the WAN side of the HA cluster. This public IP is what is known by the remote sites connecting to the hosts or the network behind the HA cluster of Netscaler SD-WAN appliances.

6. Choose the subnetwork to be assigned to the network interfaces which will be used for the HA cluster of the NetScaler SD-WAN appliances. This is automatically populated with 10.11.0.0/16. It can be changed and administered the way you would want it. The network IPs for various NICs of Management, LAN, WAN and for HA control traffic can be chosen and created automatically as part of the solution template.

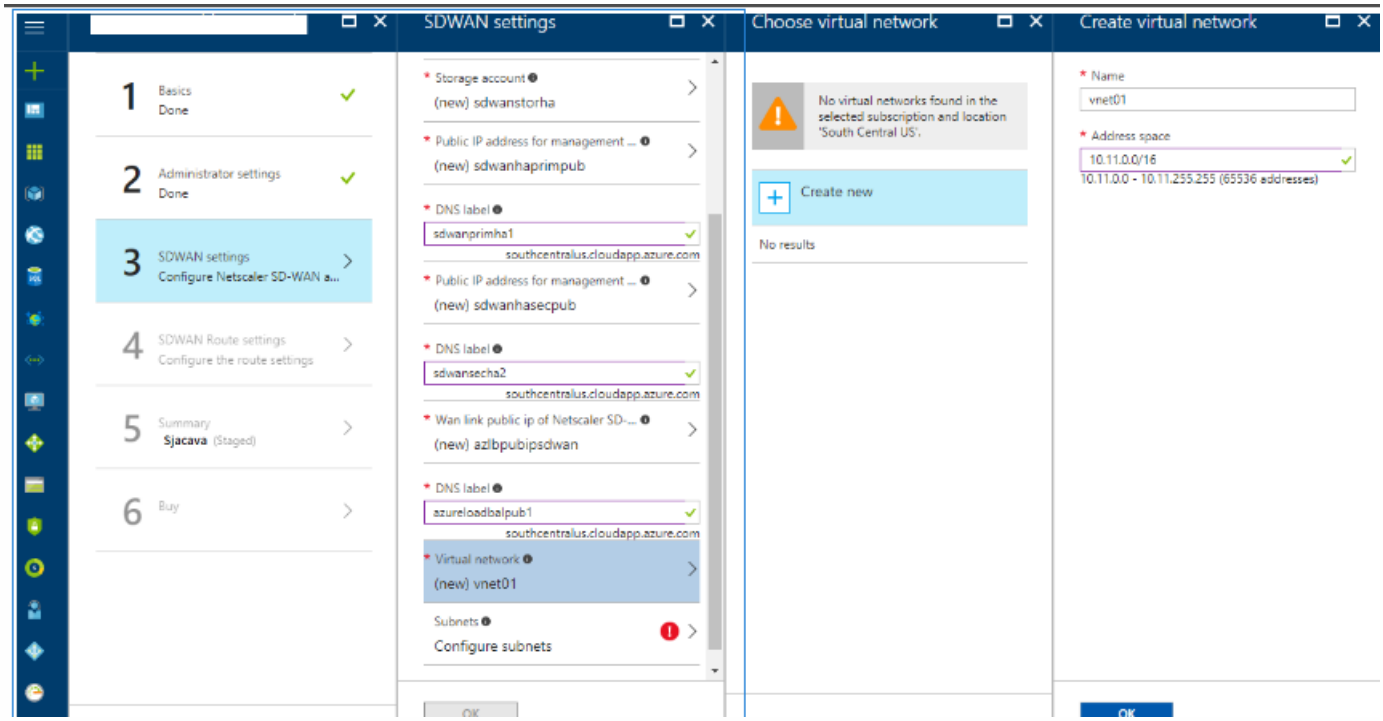

7. Name the various NICs to be created for the HA cluster of NetScaler SD-WAN appliances. The order should be as follows; MANAGEMENT, LAN, WAN, and the AUX subnet which is used for HA control packets exchange for achieving HA convergence and state association of Active/Standby.

a. Provide a name for the Subnet Names. The subnetworks are auto-populated for using the VNET created. You

can change these networks as per your requirement.

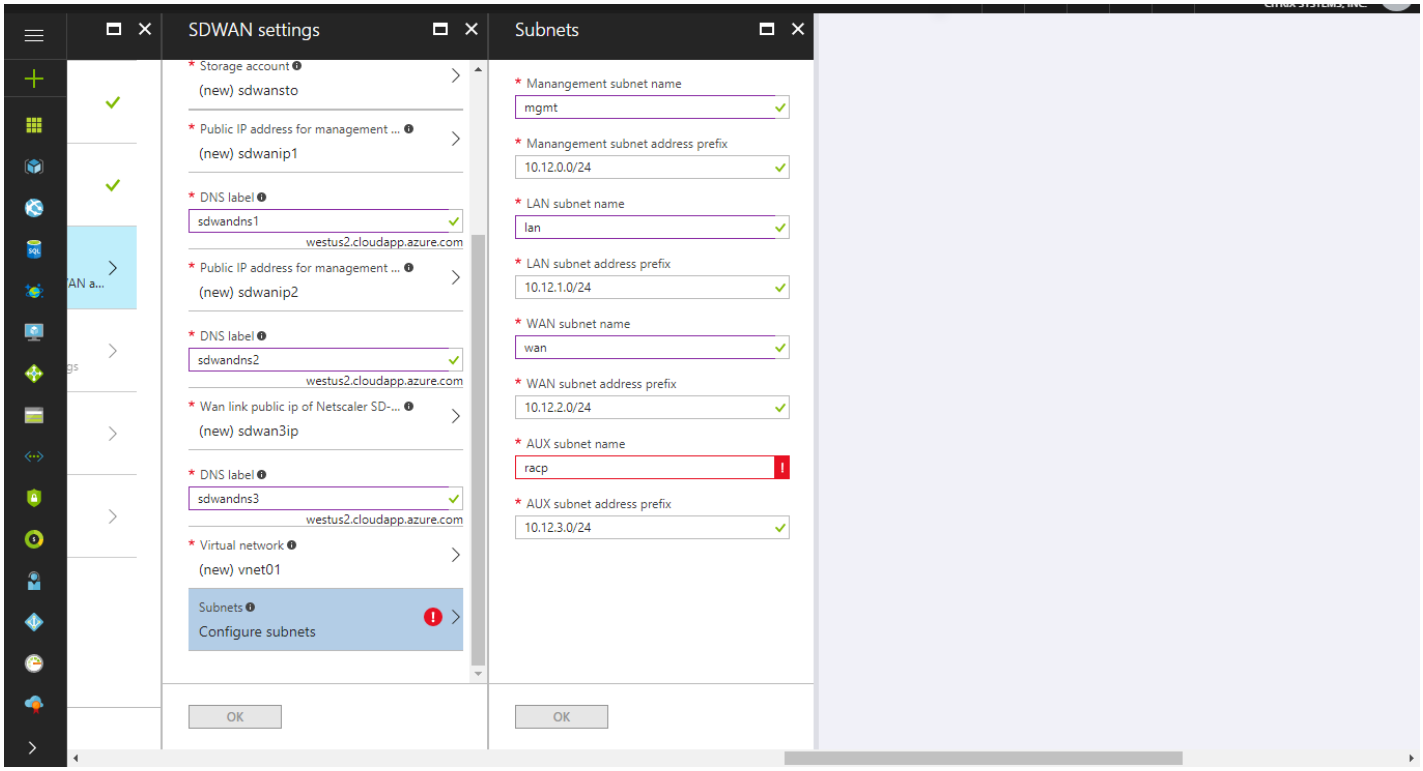

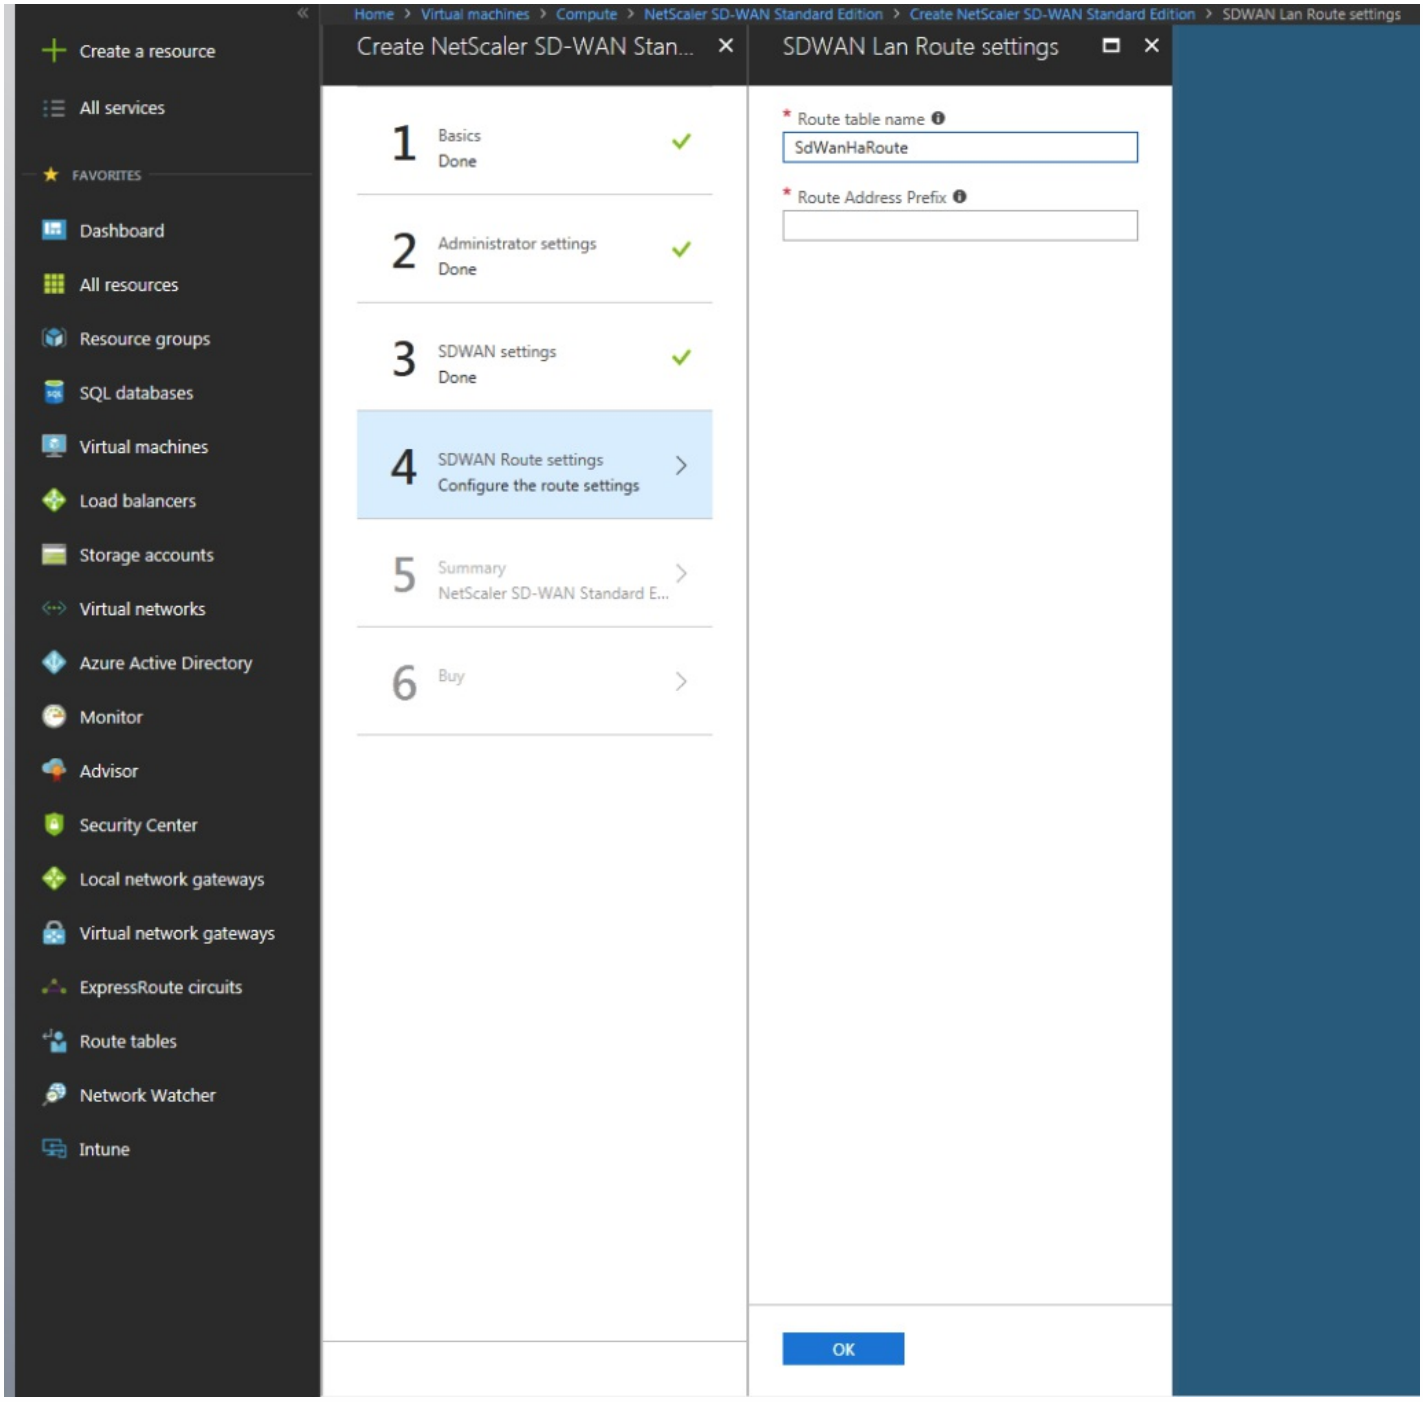

8. After all of above steps are completed, the final step is to run the parameters for validation and check for errors in deployment to determine success or failure. The notifications section in Azure provides details on the latest status of the deployment creation and whether or not the deployment succeeded or failed. If failed, Azure indicates comprehensive output on failure that can be addressed.

9. After successful deployment, you can go to resource groups from the Azure icon pane and check for the resource group you created. This resource group hosts all the types created as part of the solution template for administrative reference.

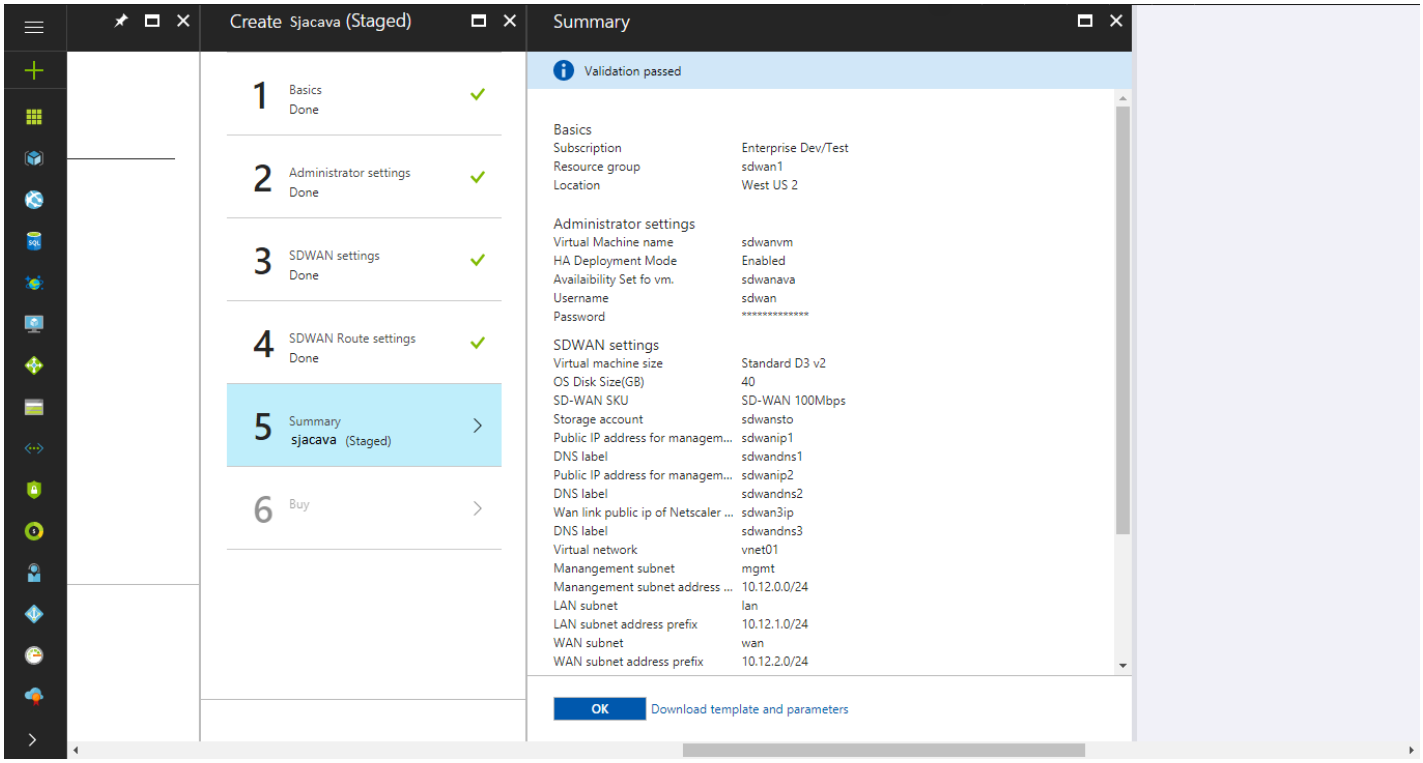

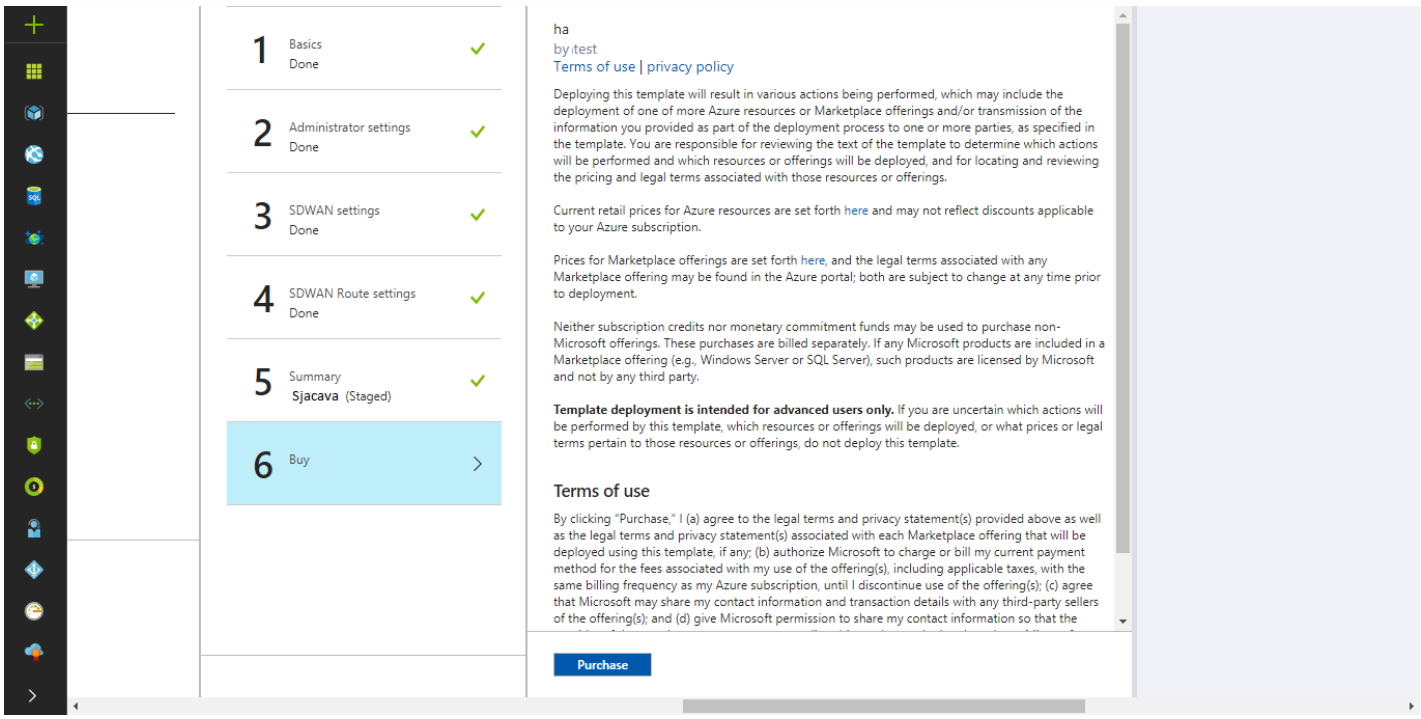

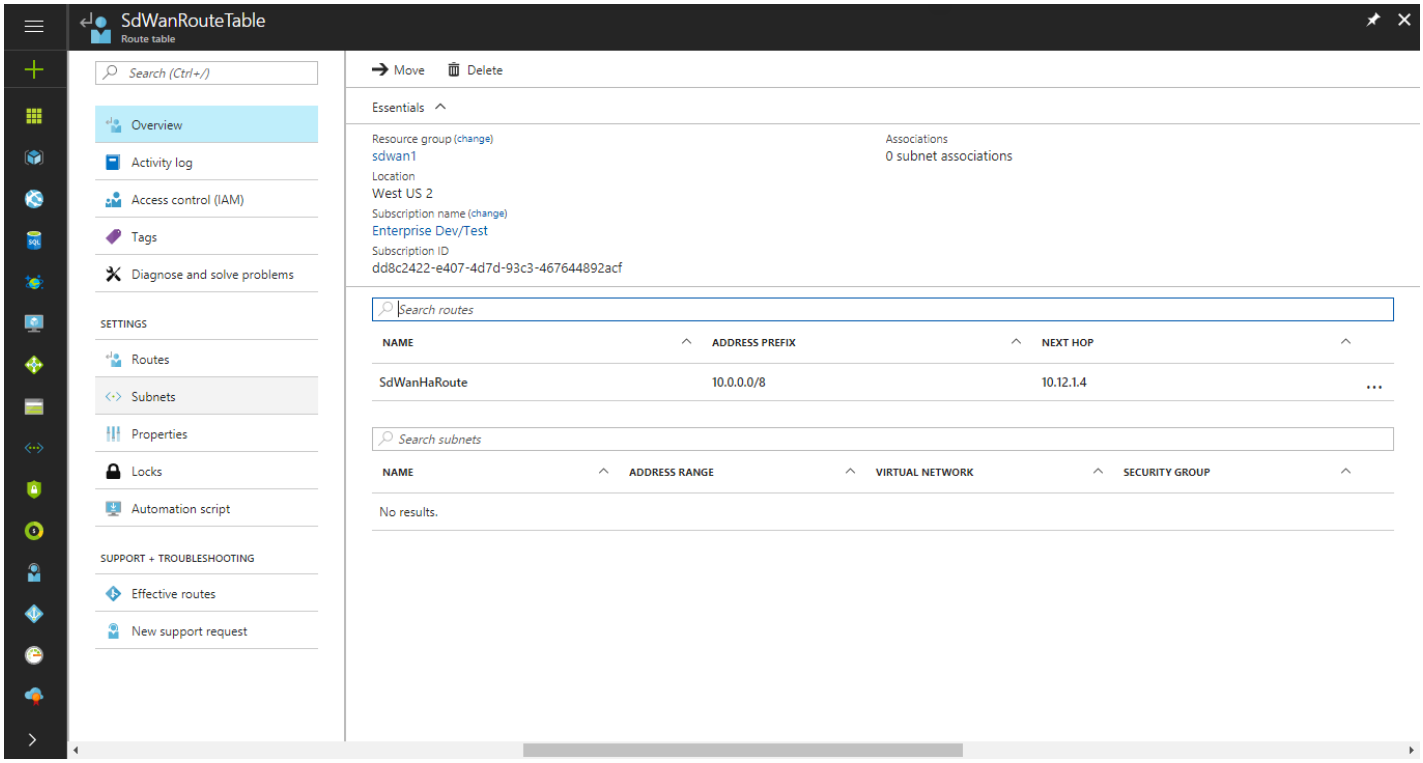

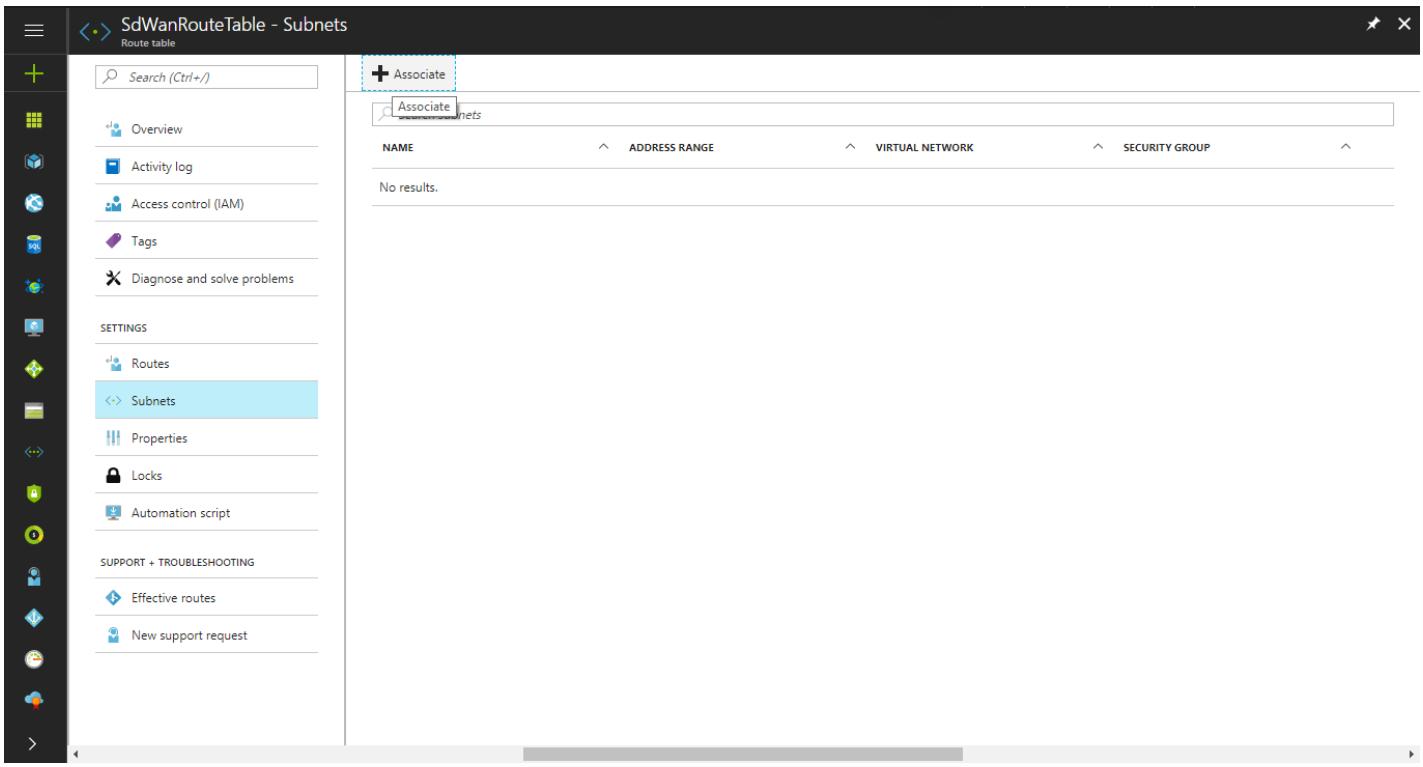

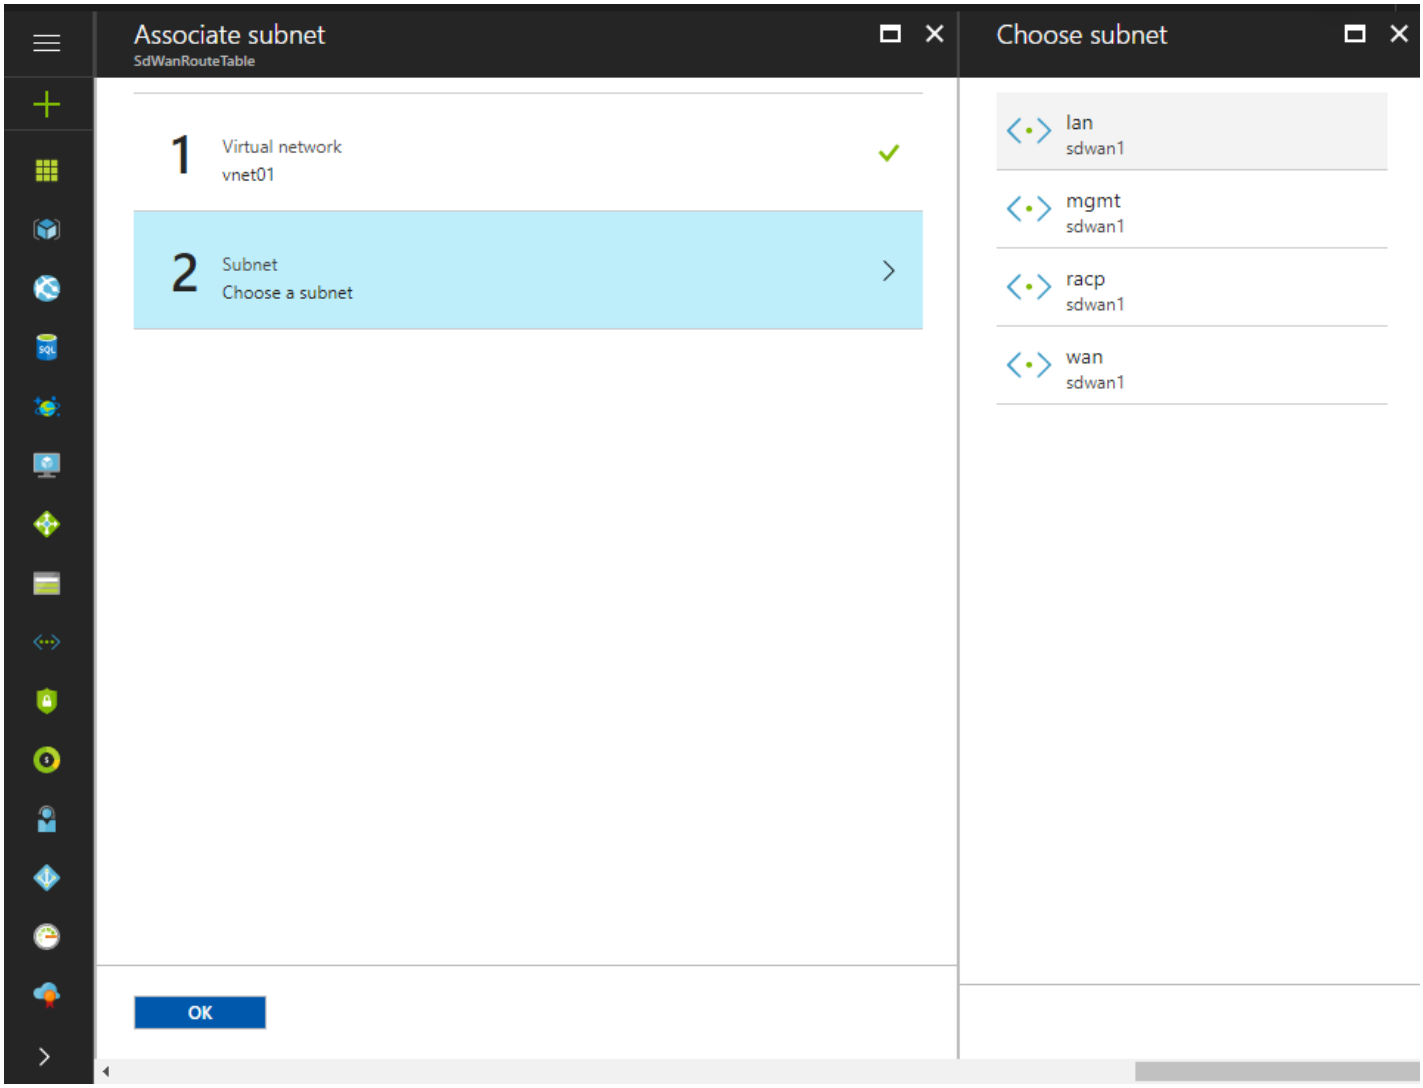

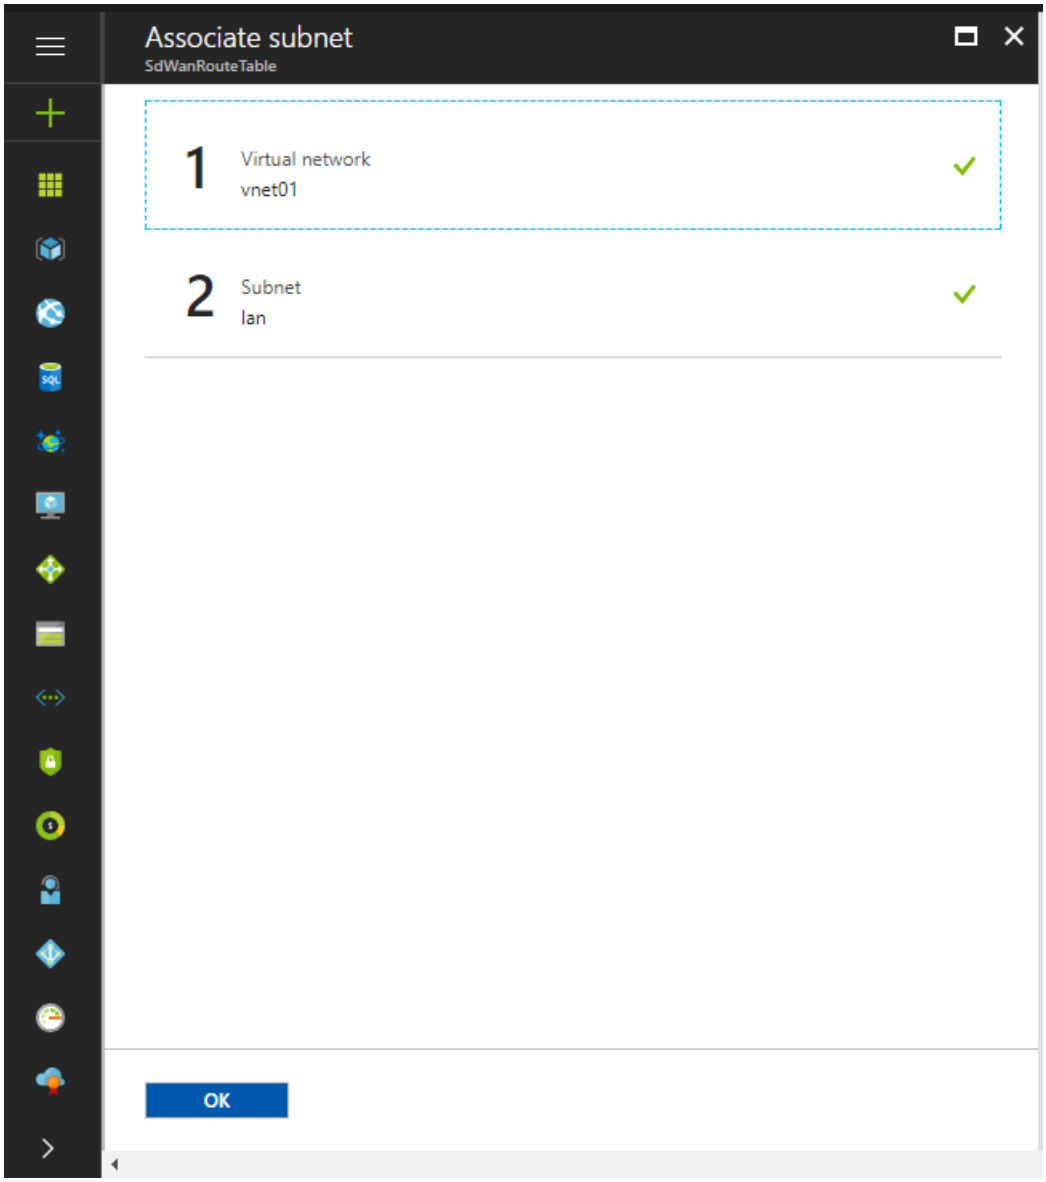

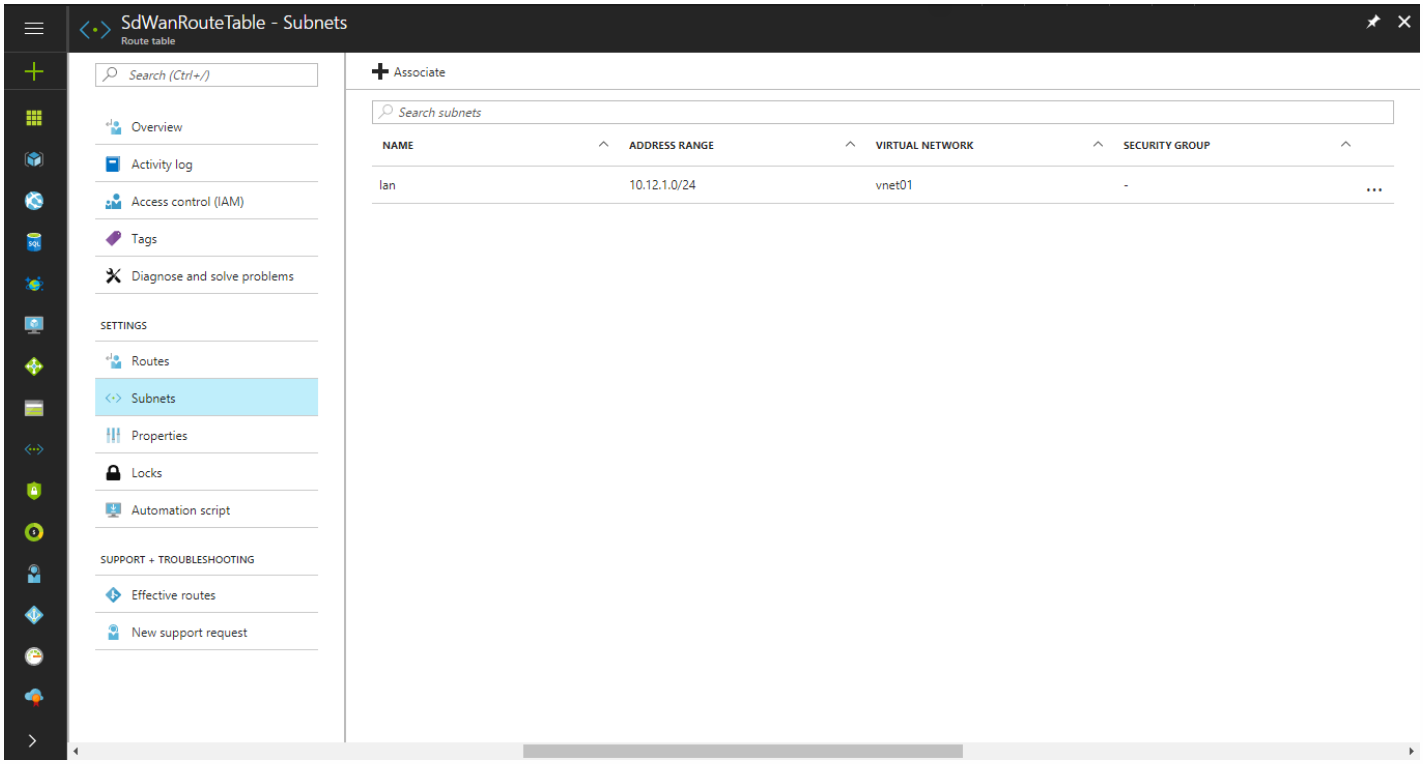

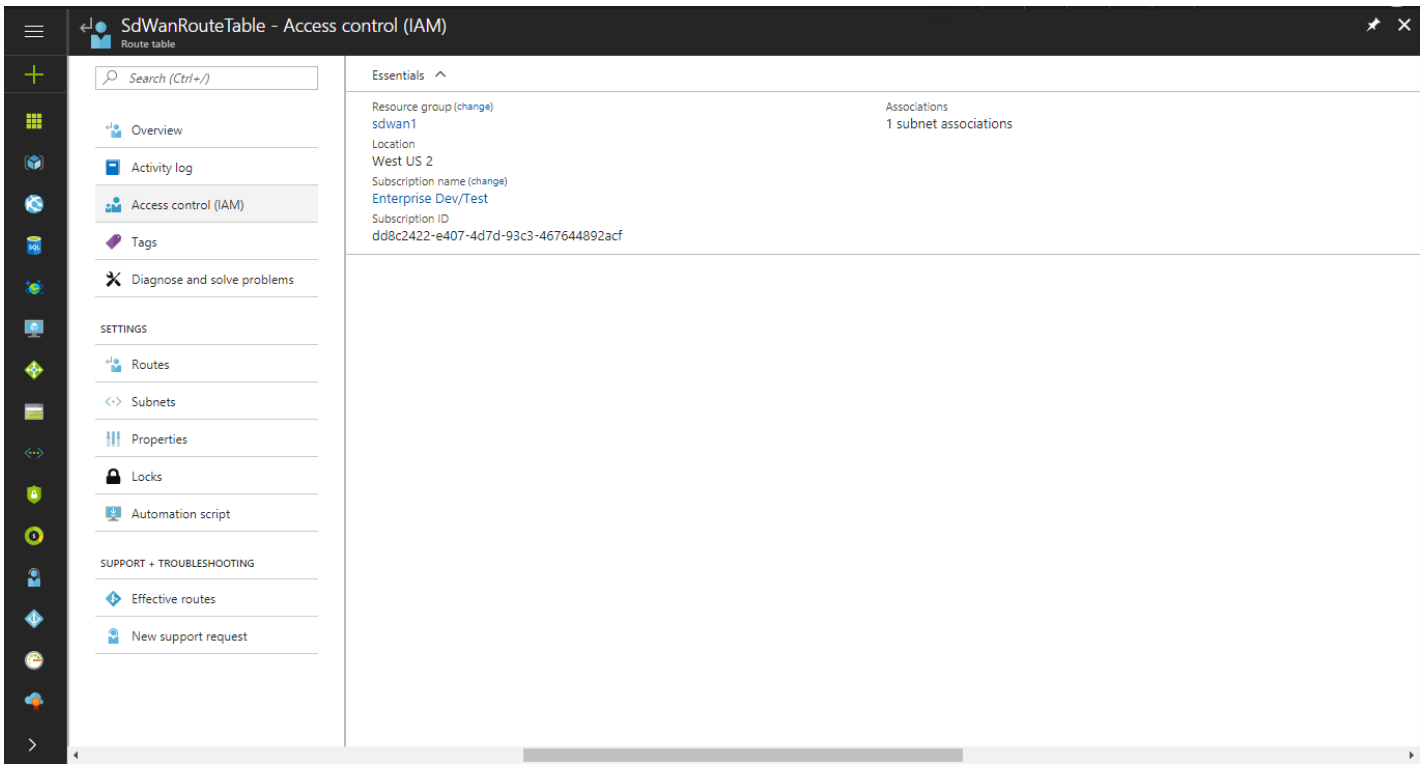

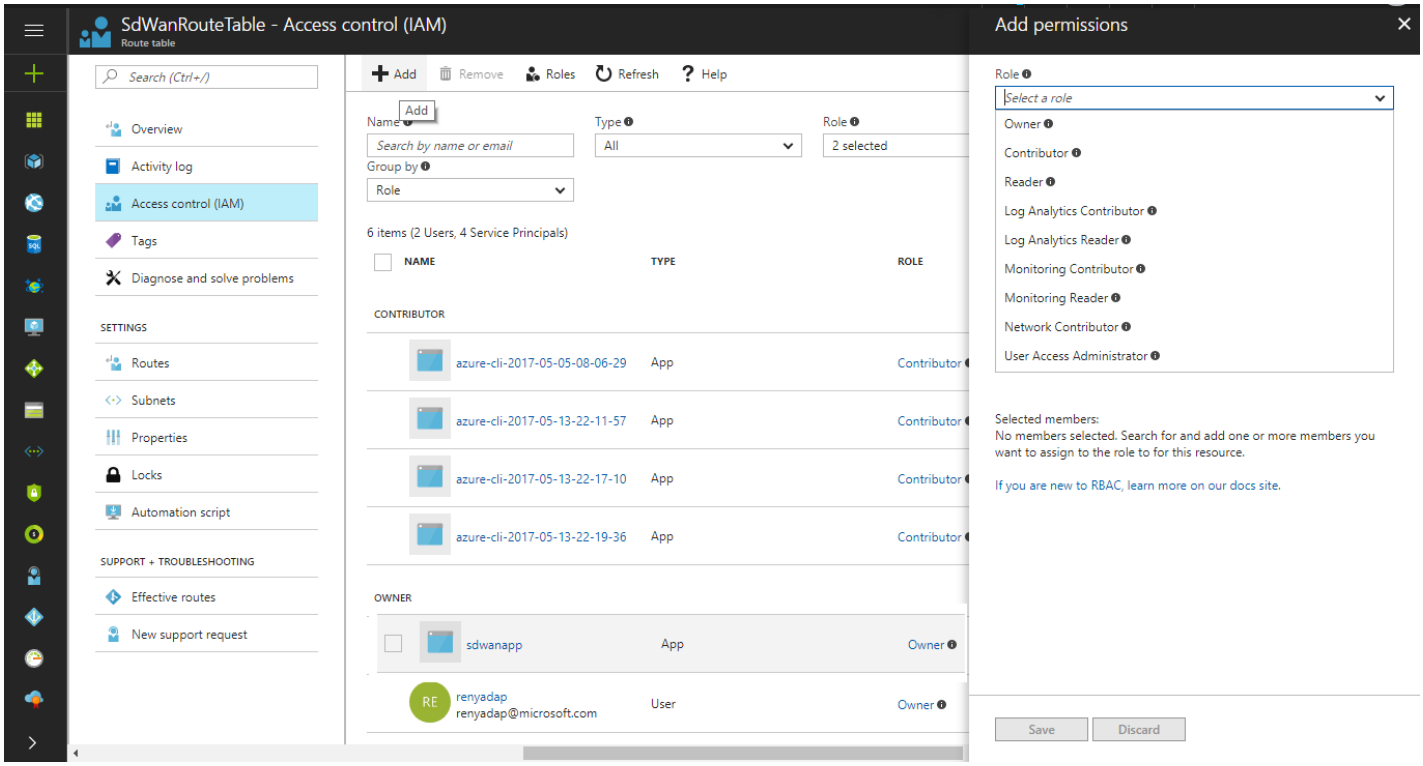

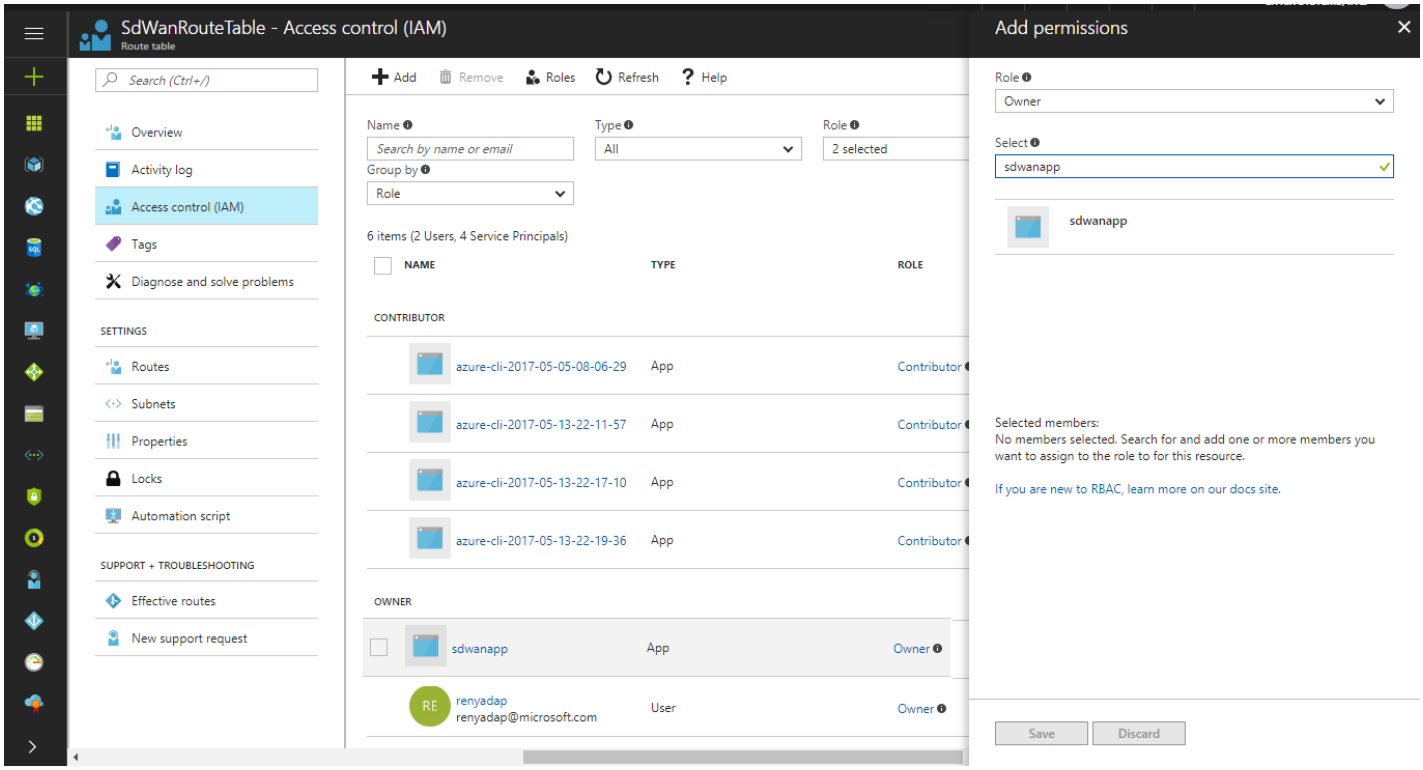

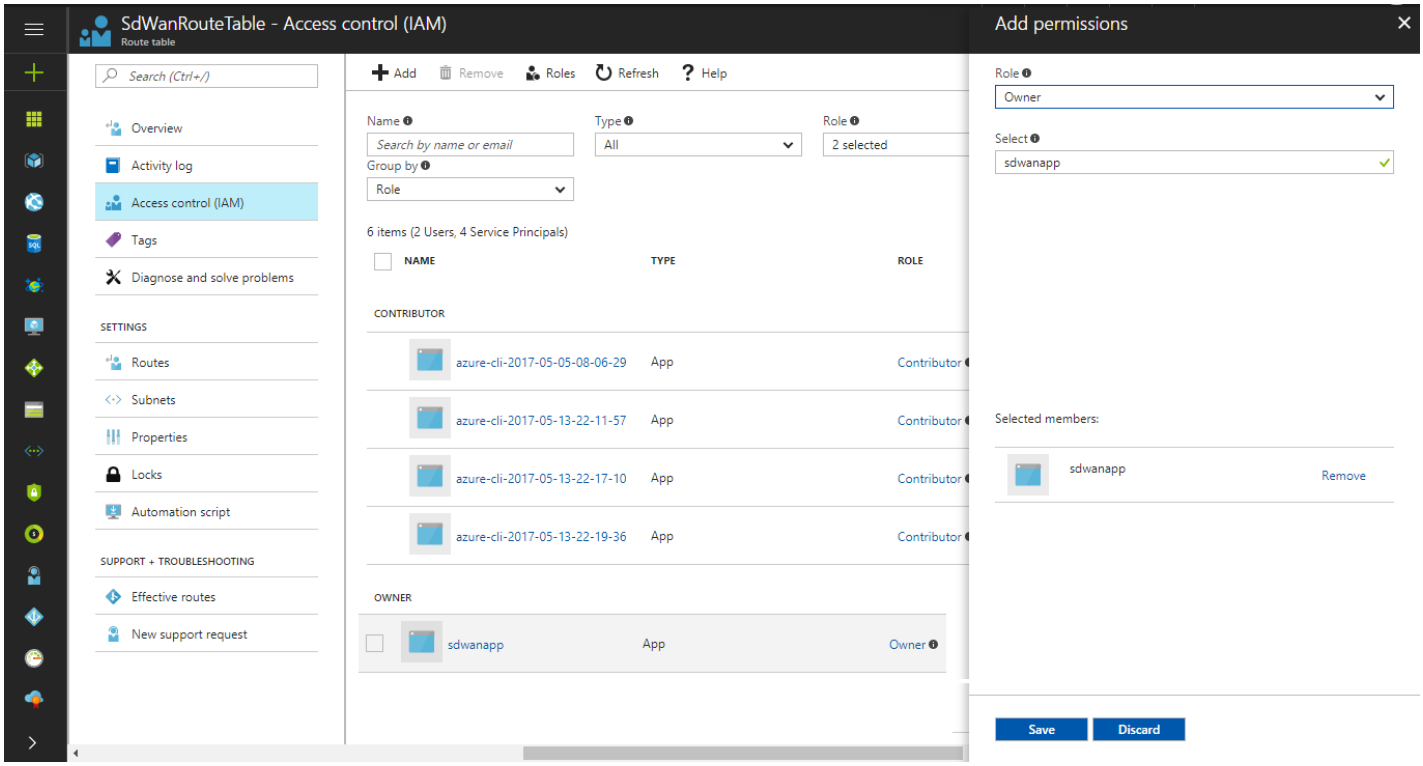

Configuring HA in NetScaler SD-WAN GUI

To configure high-availability for SD-WAN appliances:

1. In NetScaler SD-WAN GUI, go to **Configuration** > **Virtual WAN** > **Configuration Editor**. Expand the **DC** site for which you want to configure interface groups for HA.

2. Go to **Interface Groups**. Configure LAN, WAN, and HA control exchange interfaces as shown below.

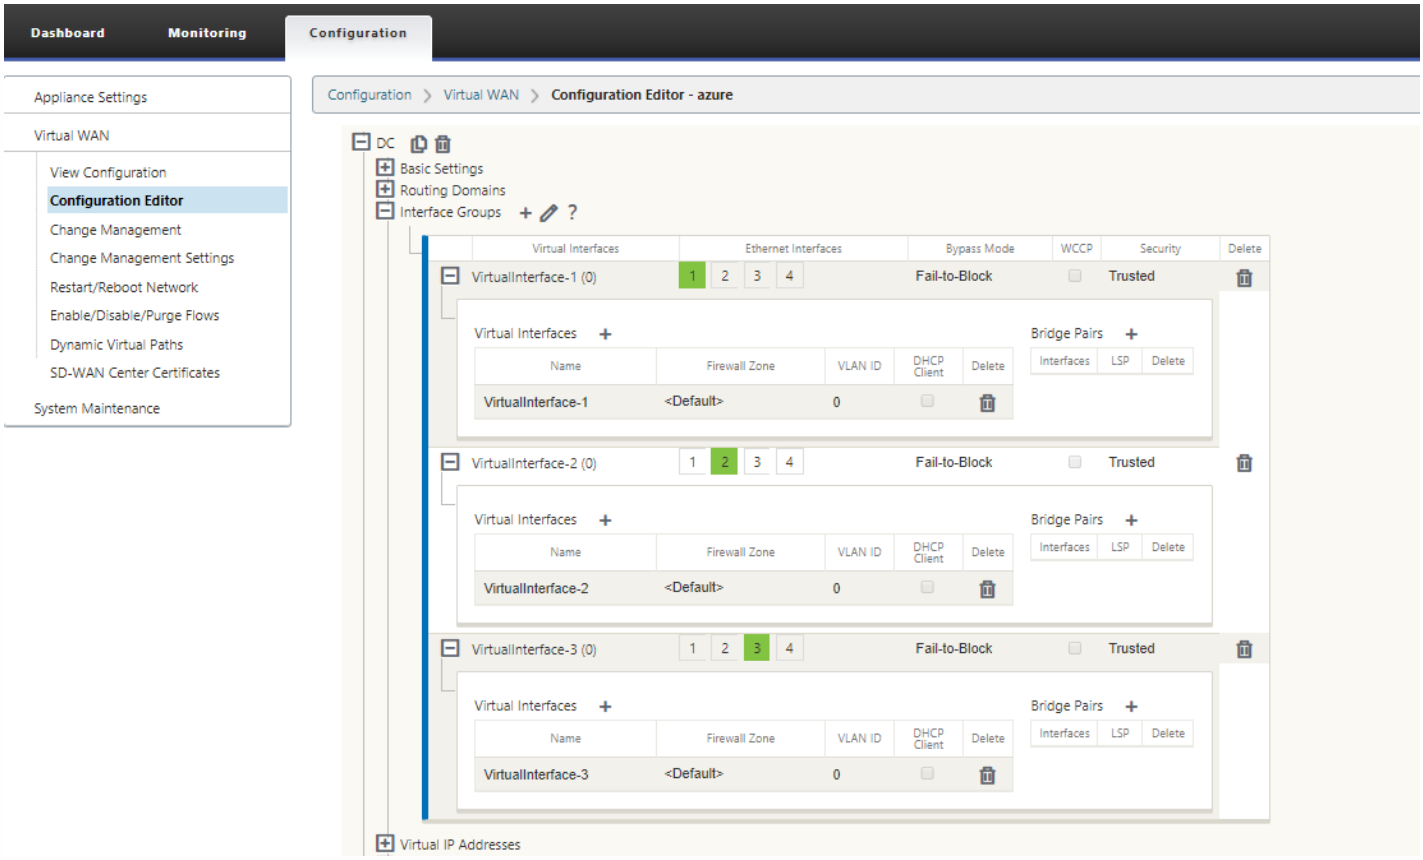

Cloud platforms have associated IP addresses for every interface (LAN/WAN). Define both IP addresses for the LAN/WAN network in Azure instance configuration for both primary and secondary instances. This is configured so that the platform is aware of the correct IP address that becomes primary and is able to respond to ARPs based on whichever instance is active.

For HA control exchange NIC Virtual IP definition, the network interface IP addresses configured as part of Azure solution template configuration should be used for high availability configuration in SD-WAN as the Primary and Secondary address. For Virtual IP definition at interface group level, you need to use one random unused IP address on the same subnet as that of the network.

In the example shown below, the IP address 10.18.3.x/24 is the HA control exchange subnetwork and the actual addresses configured as part of network interfaces are displayed.

3. Configure **Virtual IP Addresses**, which is used for Primary and Secondary instances of LAN, WAN, and HA respectively.

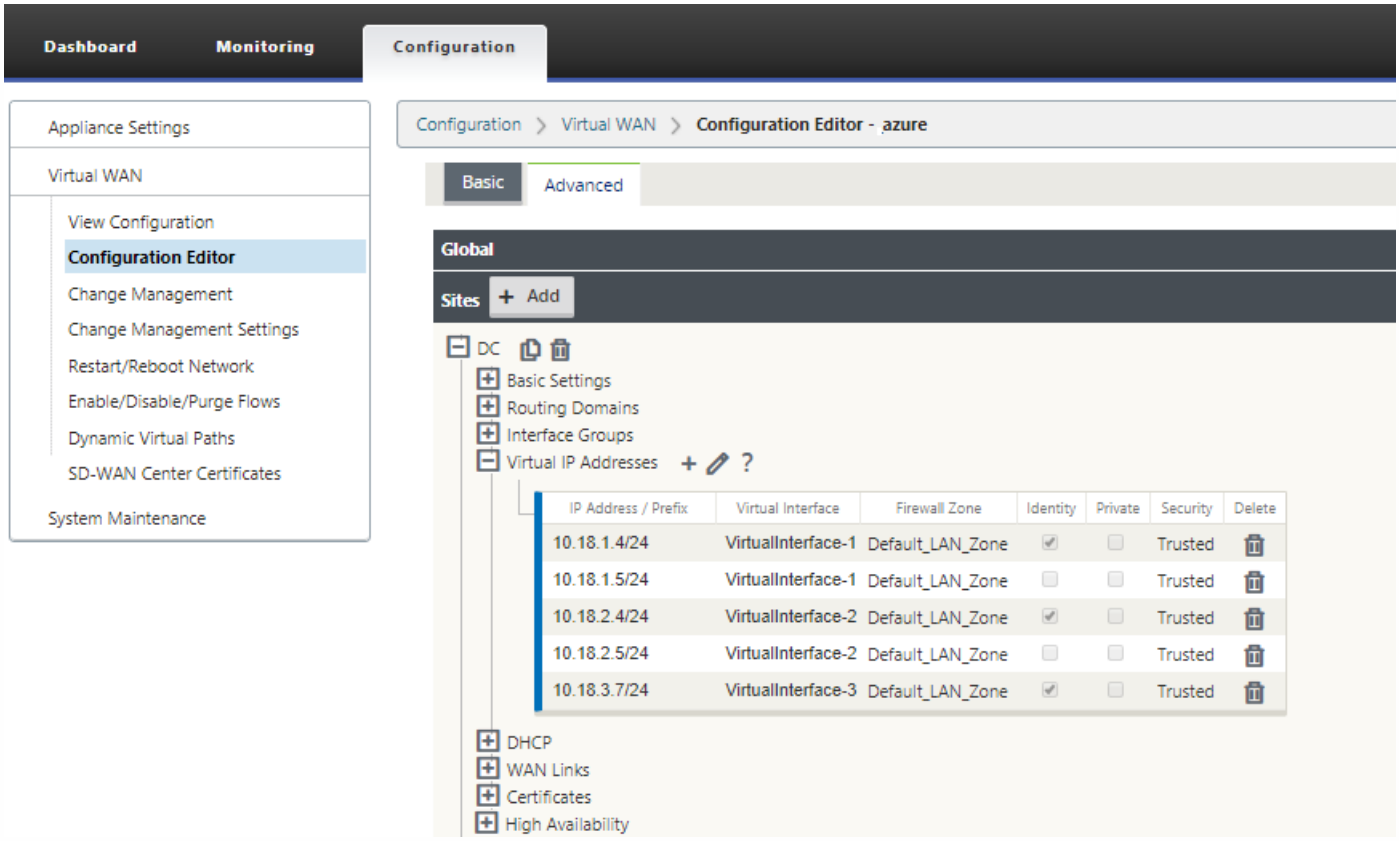

4. Enable **High Availability**. Configure virtual Interfaces as shown below.

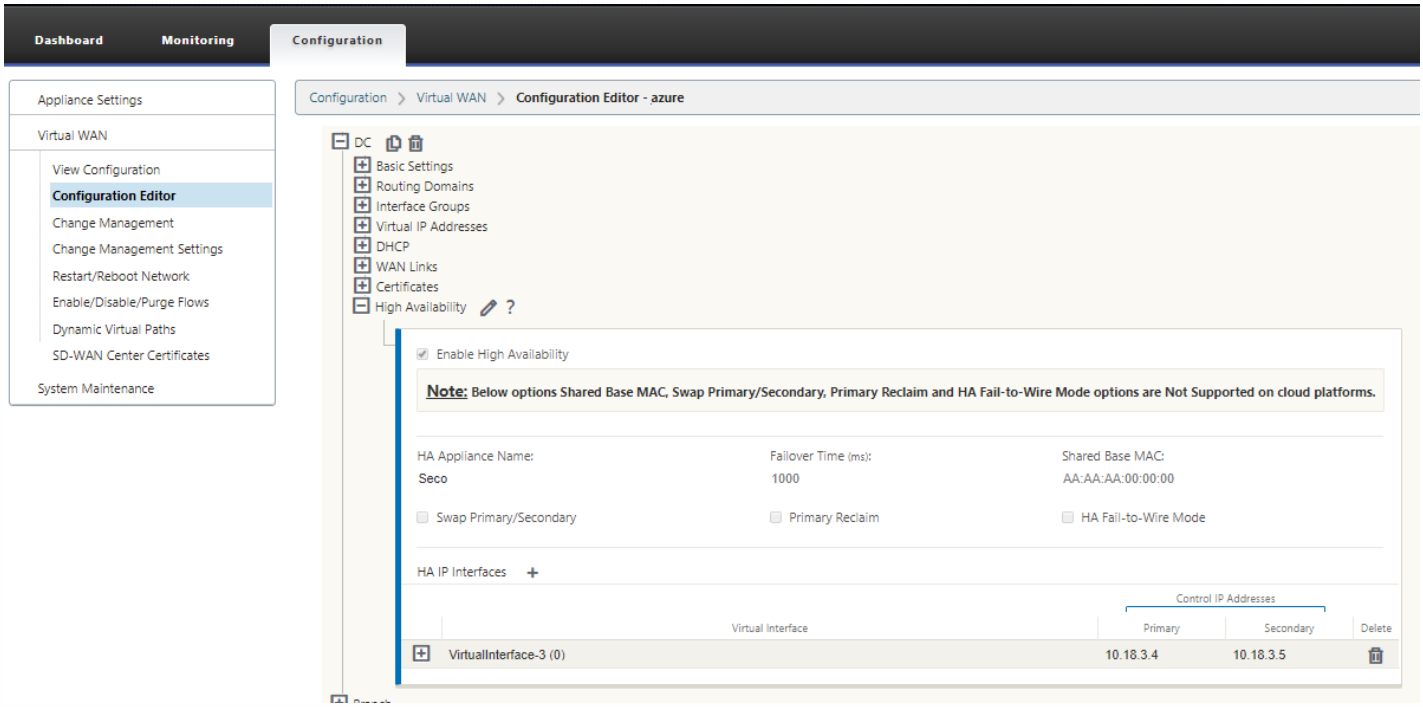

5. View SD-WAN GUI **Dashboard** to validate and confirm the HA configuration status.

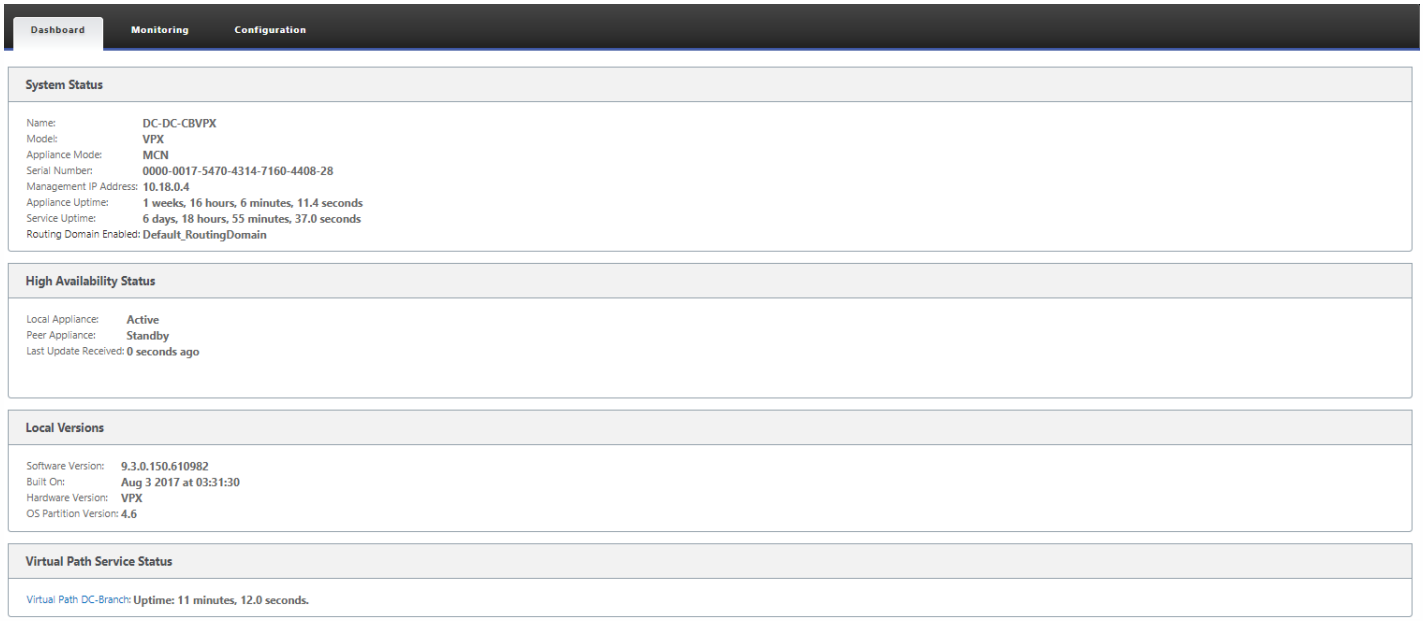

XenServer 6.5 Upgrade for SD-WAN Standard Edition Appliances

Aug 14, 2017

> **Important**
>
> **To upgrade to XenServer version 6.5, the appliances must be running NetScaler SD-WAN software release 9.0.x or later.**

> **Note**
>
> Do not attempt upgrading, if the appliance is running on software version lower than release 9.0.x to prevent upgrade issues.

How to Upgrade to XenServer 6.5

To upgrade to XenServer 6.5 on SD-WAN Standard Edition appliances, ensure that the appliance is running software release version 9.0.x or later. If the appliances are running older software release version, upgrade to the latest software release version first.

1. Upgrade SD-WAN Standard Edition software through the change management procedure. See, the Change Management procedure.

a. In NetScaler SD-WAN SE GUI, go to **Configuration** > **System Maintenance** > **Update Software**. Download the *cb-vw-<Platform_Model>-<Build_No>.tar.gz* file. Then, download *ns-sdw-vw-<Build_No>.upg* file to update operating system software.

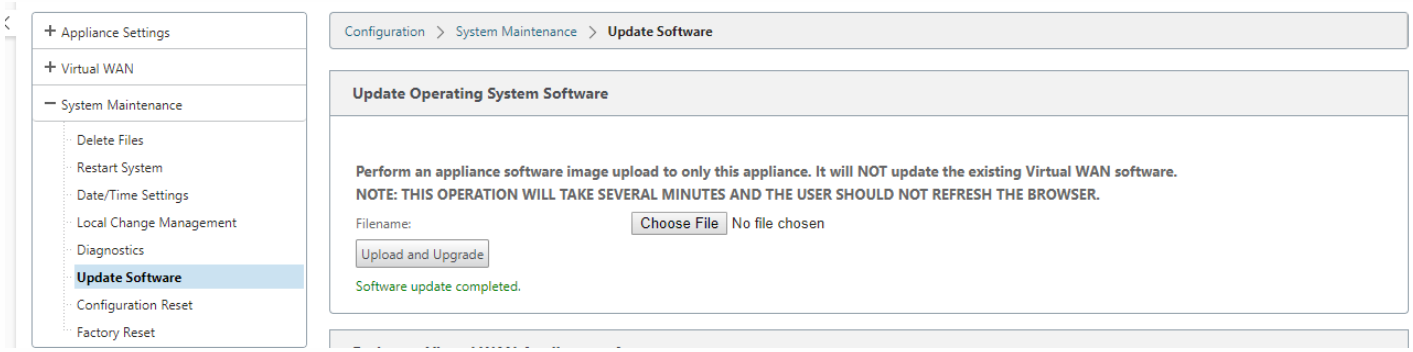

b. Follow Single-Step Upgrade work flow to upgrade SD-WAN software.

2. Perform steps a or b as outlined in step 1 before upgrading to Citrix XenServer 6.5.

3. Navigate to **Update Software** in the SD-WAN GUI.

4. Upload Citrix XenServer6.5 bundle which has been download from download server to **Update Operating System Software** by selecting the downloaded file location.

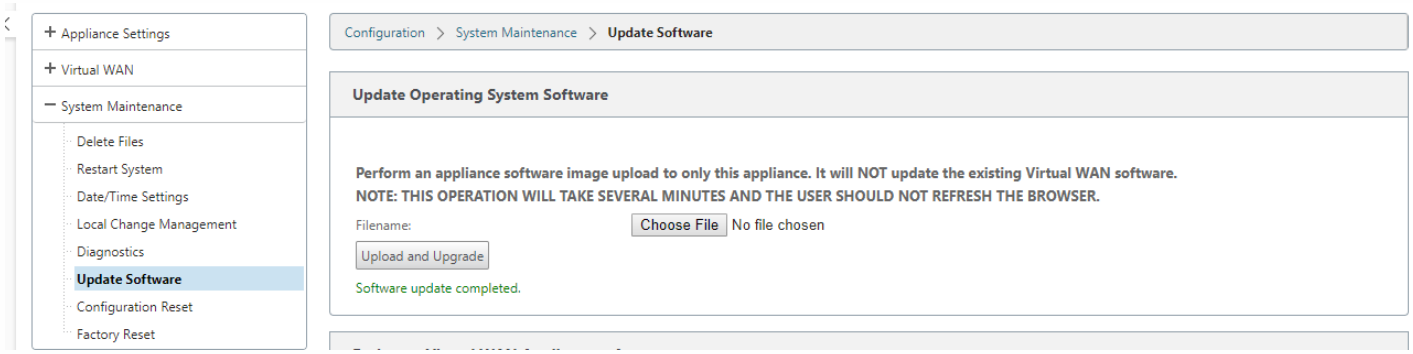

5. Click **Upload and Upgrade**. Wait for approximately 20 mins for the upgrade to complete. The appliance restarts after the upgrade is successfully completed.