



Citrix SD-WAN WANOP 11.0

Contents

About Citrix SD-WAN WANOP	7
Get started with Citrix SD-WAN WANOP	16
Select an appliance based on capacity	17
Select the deployment mode based on datacenter topology	20
Sites with one WAN router	22
Sites with multiple WAN routers	24
Appliance failure handled in various deployment modes	27
Supported mode and feature matrix	27
Configure Citrix SD-WAN WANOP plug-in with Access Gateway VPNs	29
Deploy SD-WAN WANOP VPX on Microsoft Azure	31
SD-WAN WANOP upgrading procedure	37
Initial Configuration	39
Prerequisites	40
Deployment Worksheet	41
Configuring the Appliance	44
Assigning a Management IP Address through the Ethernet Port	44
Assigning a Management IP Address through the Serial Port	46
Provisioning the Appliance	47
Deployment Modes	51
Customizing the Ethernet ports	53
Port Parameters	54
Accelerated Bridges (apA and apB)	55
Motherboard Ports	56

VLAN Support	57
Customizing the Ethernet ports	57
Ethernet Bypass and Link-Down Propagation	58
Accelerating an Entire Site	59
Partial-Site Acceleration	59
WCCP Mode	60
WCCP Mode (Non-Clustered)	64
WCCP Clustering	70
Virtual Inline Mode	77
Configuring Packet Forwarding on the Appliance	78
Router Configuration	78
Virtual Inline for Multiple-WAN Environments	82
Virtual Inline Mode and High-Availability	82
Monitoring and Troubleshooting	83
Group Mode	83
When to Use Group Mode	84
How Group Mode Works	85
Enabling Group Mode	85
Forwarding Rules	87
Monitoring and Troubleshooting Group Mode	88
Customizing the Ethernet ports	89
How High-Availability Mode Works	89
Cabling Requirements	91
Other Requirements	91

Management Access to the High-Availability Pair	92
Configuring the High-Availability Pair	92
Updating Software on a High-Availability Pair	93
Saving/Restoring Parameters of an high availability Pair	94
Troubleshooting High Availability Pairs	94
Two box mode	95
FAQs	99
Acceleration	100
CIFS and MAPI	100
Compression	102
RPC over HTTPS	104
SCPS	106
Secure peering	106
SSL Acceleration	107
Citrix SD-WAN WANOP plug-in	108
Traffic shaping	113
Upgrade (OS) Process	115
Video caching	122
Office 365 Acceleration	127
Compression	129
HTTP acceleration	135
How HTML5 works	137
Internet Protocol version 6 (IPv6) acceleration	139
Link definitions	144

Manage link definitions in traffic shaping	146
Configure link definitions	147
Manage and monitor using Citrix Application Delivery Management	152
Citrix Cloud Connector	153
Configure cloud connector tunnel	157
Configure cloud connector tunnel between two datacenters	160
Configure cloud connector tunnel between a datacenter and AWS/Azure	165
Office 365 acceleration	170
SCPS support	183
Secure traffic acceleration	184
Secure peering	184
CIFS, SMB2, and MAPI	189
Configure Citrix SD-WAN WANOP appliance to optimize secure Windows traffic	191
Configure CIFS and SMB2/SMB3 acceleration	207
Configure MAPI acceleration	214
SSL compression	216
How SSL compression works	217
Configure SSL compression	220
SSL Compression with Citrix SD-WAN WANOP plug-in	228
RPC over HTTP	228
TCP Flow-Control acceleration	231
Lossless and transparent flow control	232
Speed optimization	233
Auto-discovery and auto-configuration	235

TCP flow control modes	236
Firewall considerations	238
Traffic classification	239
Application classifier	240
Service classes	242
Traffic shaping	247
Weighted fair queuing	248
Traffic shaping policies	250
Video caching	253
Video caching scenarios	255
Configure video caching	258
Video prepopulation	263
Verify video caching	270
Manage video caching sources	272
WAN insight	274
Asymmetric routing	278
Citrix SD-WAN WANOP client plug-in	280
Hardware and software requirements	281
How WANOP plug-in works	283
Deploy appliances for use with plug-ins	291
Customize plug-in's MSI file	294
Deploy plug-ins on Windows	297
Citrix SD-WAN WANOP plug-in GUI	301
Update Citrix SD-WAN WANOP plug-in	305

Citrix Virtual Apps and Desktops acceleration	306
Configure Virtual Apps acceleration	307
Optimize Citrix Receiver for HTML5	308
Deployment modes	311
Adaptive transport interoperability	318
Citrix Hypervisor 6.5 upgrade	319
Maintenance	319
Diagnostics	323
Troubleshooting	330
CIFS and MAPI	330
Citrix SD-WAN WANOP plug-in	333
RPC over HTTPS	334
Video caching	335
Citrix Virtual Apps and Desktops acceleration	336

About Citrix SD-WAN WANOP

March 12, 2021

Citrix SD-WAN WANOP appliances optimize your WAN links, giving your users maximum responsiveness and throughput at any distance. A Citrix SD-WAN WANOP appliance is easy to deploy, because it works transparently. A twenty minute installation accelerates your WAN traffic with no other configuration required. You do not have to change your applications, servers, clients, or network infrastructure. You can, however change them after Citrix SD-WAN WANOP installation without affecting traffic acceleration. A Citrix SD-WAN WANOP appliance needs reconfiguration only when your WAN links change.

Citrix SD-WAN WANOP appliances support a full range of optimizations, including:

- Multi-session compression with compression ratios of up to 10,000:1.
- Protocol acceleration for Windows network file systems (CIFS), Virtual Apps (ICA and CGP, including the new *multi-session ICA* standard), Microsoft Outlook (MAPI), and SSL.
- Traffic shaping to ensure that high-priority and interactive traffic takes precedence over low-priority or bulk traffic.
- Advanced TCP protocol acceleration, which reduces delays on congested or high-latency links.
- Video caching.

How Citrix SD-WAN WANOP works?

Citrix SD-WAN WANOP products work in pairs, one at each end of a link, to accelerate traffic over the link. The transformations done by the sender are reversed by the receiver.

However, one appliance (or virtual appliance) can handle many links, so you do not have to dedicate a pair to each connection.

An enterprise typically has one Citrix SD-WAN WANOP appliance per site (larger appliances at larger sites, smaller ones at smaller sites), though a company with numerous branch offices might have multiple appliances at its central data center.

A link from a site with a Citrix SD-WAN WANOP appliance to a site that does not have a Citrix SD-WAN WANOP appliance functions normally, but its traffic is not accelerated.

Citrix SD-WAN WANOP features include robust compression for brisk performance over relatively slow links, and lossless flow control to deal with congestion. TCP optimizations overcome the main limitations of problematic links, and application optimization does away with the limitations of applications designed for high-speed, local networks. An autodetection feature makes deployment quick and easy.

Citrix SD-WAN WANOP features and benefits

Any time workers spend waiting for their computers to respond is lost time, resulting in lost productivity. When users work remotely or use off-site resources, their productivity depends on the responsiveness of their network connections. Safeguarding the responsiveness of their connections requires advanced network acceleration.

The Citrix SD-WAN WANOP product line protects your productivity by providing reliable WAN and Internet link performance through a set of multiple, interlocking optimizations, each reinforcing the others. To provide maximum productivity across your entire enterprise, there are Citrix SD-WAN WANOP products for every need, from the largest data center through the smallest branch office and even the individual laptop.

Citrix SD-WAN WANOP provides robust usability even with undersized or degraded links.

Features at a glance:

For more information, see the [table](#)

Features and benefits:

The following are some of the key benefits of our Citrix SD-WAN WANOP product line.

Compression overcomes low link speeds. The most obvious problem with wide-area network (WAN) links and Internet links is their low bandwidth compared to local-area networks (LANs). A 1 Mbps WAN has only 1% of the throughput of a 100 Mbps LAN. How do you overcome low link bandwidth? With compression. A compression ratio of 100:1 enables a 1 Mbps link to transfer data as quickly as a 100 Mbps. This speedup factor is achieved whenever the following criteria are met:

- The compression algorithm must be able to deliver high compression ratios.
- The compression algorithm must be very fast (much faster than the link bandwidth, and ideally as fast as the LAN).
- The LAN segments of the link must have flow control that is independent of the WAN segment, because the different segments handle data at different rates.
- Multiple compression engines must be used to handle the different needs of different kinds of traffic. Interactive traffic requires relatively little bandwidth but is very sensitive to delay, while bulk-transfers are very sensitive to bandwidth but are insensitive to delay.

TCP protocol acceleration overcomes congestion. Any attempt to send traffic faster than the link speed results in congestion, which results in many problems caused by high packet losses and high queuing latency.

Lossless flow control. The TCP/IP protocol has no flow control to slow senders down directly, and the absence of this necessary control mechanism makes packet losses and excessive queuing delays

normal, even on mission-critical links. (If anything, this problem is getting worse over time, as papers on the phenomenon of **bufferbloat** attest.)

A Citrix SD-WAN WANOP appliance solves this problem by providing the flow control that was omitted from the TCP/IP protocol. Unlike ordinary quality of service (QoS) solutions, which simply reallocate packet loss, Citrix SD-WAN WANOP provides lossless flow control that controls the rate at which the endpoint senders transmit data, instead of allowing senders to transmit data at any speed they like, and dropping packets when they send too much. Each sender transmits only as much data as Citrix SD-WAN WANOP allows it to send, without ever dropping a packet, and this data is placed on the link at exactly the right rate to keep the link full without overflowing. By eliminating excess data, Citrix SD-WAN WANOP is not forced to discard it. Without Citrix SD-WAN WANOP, the dropped packets have to be sent again, causing unnecessary delays. Lossless flow control also eliminates delays caused by excessive buffering. Lossless flow control is the key to maximum responsiveness on a busy link, enabling a link that was once congested to the point of unusability at 40% utilization to remain productive and responsive at 95% utilization.

Eliminating distance-based unfairness. Links with high latency or packet losses are difficult to use at full bandwidth, especially with ordinary TCP variants such as TCP Reno. The consequences are excessive delays and difficulty in getting the bandwidth that you are paying for. The longer the link distance, the worse the problem becomes.

Citrix SD-WAN WANOP TCP protocol acceleration minimizes these effects, allowing intercontinental and even satellite links to run at full speed.

Traffic shaping manages bandwidth automatically. On the output side, a fair-queuing-like algorithm ensures that each connection is independently queued and given its fair share of the link bandwidth. Traffic-shaping policies allow different services to be given higher or lower precedence.

Applications and protocols designed for use on local-area networks are notorious for poor performance over wide-area networks, because the designers did not consider the effects of long speed-of-light delays on their protocols. For example, a simple Windows file system (CIFS) operation can take up to 50 round trips as messages pass back and forth across the network. In a wide-area network with a 100 ms round-trip time, 50 round trips cause a delay of five seconds.

Although speed-of-light delays are a fundamental limitation, application optimizations can perform the same operations in a smaller number of round-trips, usually through speculative operations. Where the original application would issue one command at a time and wait for it to complete before issuing the next one, it is often perfectly safe to issue a series of commands without waiting. In addition, data transfers can be accelerated through a combination of pre-fetching, read-ahead, and write-behind operations. By packing as many operations as possible into a single round trip, performance can be increased tenfold or more.

Citrix SD-WAN WANOP optimizations are especially effective on CIFS/SMB (the Windows file system),

MAPI (the Outlook/Exchange protocol), and HTTP.

Multiple optimizations enhance Virtual Apps/Virtual Desktops (Citrix HDX) performance. Because Citrix SD-WAN WANOP appliances are Citrix products, they are especially effective at accelerating Citrix protocols, such as Citrix Virtual Apps and Desktops. Every aspect of Citrix SD-WAN WANOP acceleration comes into play with these protocols to make the remote user experience as productive as possible.

Citrix SD-WAN WANOP appliances negotiate session options with Citrix Virtual Apps and Desktops servers. This allows the Citrix SD-WAN WANOP appliance to apply the following enhancements:

- It replaces the server's native compression with higher-performance Citrix SD-WAN WANOP compression.
- It bases the connection's traffic-shaping priority on the priority bits embedded in every Citrix Virtual Apps and Desktops connection. This allows the priority of the connection to vary according to the type of traffic. For example, interactive tasks are high-priority tasks and print jobs are low-priority tasks.
- It gathers and reports statistics based on the Virtual Apps or Virtual Desktops applications being used.
- It maintains the end-to-end encryption of the original connection.

Auto detection for minimal configuration. Because the solution is double-ended, requiring that a Citrix SD-WAN WANOP product be present at both ends of the link, deployment would seem to impose a burden on remote offices, especially ones without dedicated IT staff. However, Citrix SD-WAN WANOP is designed to be very easy to install and maintain. A typical installation takes about twenty minutes. The only parameters needed are the usual network parameters (such as IP address and subnet mask), the address of a Citrix license server, and the send and receive speed of the link.

Requiring only a minimal level of configuration is possible because of autodetection, through which a Citrix SD-WAN WANOP determines which connections can be accelerated (and which cannot), without any manual configuration. A Citrix SD-WAN WANOP at the other end of the link is automatically detected, and the connection is then accelerated. You can add Citrix SD-WAN WANOP appliances to your network in an ad hoc fashion. You do not even have to inform the existing appliances of the arrival of a new one. They discover it for themselves.

A Citrix SD-WAN WANOP uses TCP header options to report its presence and to negotiate acceleration parameters with the remote Citrix SD-WAN WANOP because TCP header options are part of the TCP standard, this method works very well, except in cases where firewalls are programmed to reject all but the most common options. Such firewalls exist, but they can be configured to allow the options used by Citrix SD-WAN WANOP to pass through.

Citrix SD-WAN WANOP operations are transparent to both the sender and receiver. The other devices in your network are not aware that Citrix SD-WAN WANOP exists. They continue working just as they

did before Citrix SD-WAN WANOP installation. This transparency also eliminates any need to install special software on your servers or clients in order to benefit from Citrix SD-WAN WANOP acceleration. Everything works transparently.

Product line capabilities:

Every product in the Citrix SD-WAN WANOP product line provides basic Citrix SD-WAN WANOP acceleration features. Most models have additional features as well, such as:

- Video caching
- Multiple accelerated bridges with Ethernet bypass feature
- Monitoring and management through the GUI, CLI, SNMP, AppFlow, and Citrix ADM.

Different Citrix SD-WAN WANOP products have different capabilities. Products that support higher WAN bandwidths also support more users and typically have more resources: more power CPU, more memory, larger disk, and more accelerated bridges.

The capabilities of products that run on your own hardware, such as the Citrix SD-WAN WANOP Plug-in and Citrix SD-WAN WANOP VPX, depend on the speed of the hardware and the amount of system resources that you dedicate to acceleration.

For up-to-date specifications, see the Citrix [SD-WAN Product Data Sheet](#).

Citrix SD-WAN WANOP architecture

Citrix SD-WAN WANOP appliances accelerate the traffic over your WAN links. To accelerate a WAN, you need at least two Citrix SD-WAN WANOP appliances, one for each site you wish to accelerate.

The sender-side Citrix SD-WAN WANOP appliance applies a series of optimizations and transformations to your traffic, such as compression and encryption. Many operations require that the receiver-side Citrix SD-WAN WANOP perform an inverse operation, such as decompression or decryption, to restore the traffic to its original state.

Thus, most optimizations require that the traffic pass through two Citrix SD-WAN WANOP appliances. Some optimizations are single-ended, and are performed by the local appliance acting alone. These optimizations include traffic shaping and video caching.

Citrix SD-WAN WANOP appliances are largely transparent to the network. The appliance itself appears to be a bridge, not a router, gateway, or proxy. This invisibility allows the appliance to be installed without configuring any other hardware. The appliance optimizations are also transparent, detected only by the partner appliance at the other end of the link.

Citrix SD-WAN WANOP appliances can be added to the network at will, because their auto-detection and auto-negotiation features ensure that a new appliance on the network is immediately detected by other appliances, and acceleration begins at once.

Although the diagram above shows a network with just two appliances, a single Citrix SD-WAN WANOP appliance can communicate with any number of partner sites. Point-to-point, hub-and-spoke, and mesh networks are all supported.

In addition to stand-alone appliances, Citrix SD-WAN WANOP acceleration products include virtual machines (the Citrix SD-WAN WANOP VPX series) and an installable acceleration service for Windows systems (the Citrix SD-WAN WANOP Plug-in).

What acceleration means

In Citrix SD-WAN WANOP terminology, “acceleration” is the reduction of transaction time, which reduces the time users spend waiting. Because the time that users spend waiting represents a direct productivity loss, acceleration’s main benefit is increased productivity.

In network traffic, a transaction ranges from very small—a single byte of data in a telnet or SSH terminal session—to very large, as with FTP transfers, which often exceed a gigabyte in size. A practical accelerator has to accelerate the entire range of transaction sizes, from interactive traffic to bulk traffic, giving the best performance and user experience across the board. Citrix SD-WAN WANOP technology achieves this in a variety of ways.

How acceleration works: The pipeline

To see how the Citrix SD-WAN WANOP appliance works, take a close look at the diagram of the traffic-flow pipeline. As you can see, there are two pipelines:

1. The sending pipeline, which accelerates data entering the WAN from the local LAN.
2. The receiving pipeline, which accelerates data exiting the WAN and entering the local LAN.



Send pipeline

To understand the appliance, consider the sending pipeline one unit at a time.

1. **Input buffer.** Packets from the LAN are received by the appliance. Because non-TCP/IP traffic is optimized only by the traffic shaper, non-TCP packets are diverted directly to the traffic shaper. The TCP/IP traffic (called TCP traffic from now on) traverses the rest of the pipeline.
2. **Video Cache.** If the TCP traffic matches the settings for the video cache, the request is handed off to the video cache unit.
3. **LAN-side auto-detection.** Other than traffic shaping, sender-side optimizations require that there be a remote appliance as well as the local appliance. Any connections that don't pass through a remote appliance are diverted to the traffic shaper. This action is performed by the LAN-side auto-detection logic. The actual test for a remote appliance is done by the WAN-side auto-detection unit.
4. **LAN-side flow control.** Citrix SD-WAN WANOP acts as a transparent TCP proxy, receiving and acknowledging packets from the endpoint sender on behalf of the endpoint receiver. This allows the appliance to accept large amounts of data from the local sender very quickly, at full LAN speeds, regardless of how slowly traffic is moving over the WAN. (Normal TCP uses end-to-end speed control, which is not agile enough to allow maximum performance.) In addition, Citrix SD-WAN WANOP flow control is lossless, meaning that the local sender never sees a dropped packet, increasing reliability and efficiency.
5. **Application engines.** Citrix SD-WAN WANOP performs specific optimizations for several protocols, including:
 - Citrix Virtual Apps and Desktops, using the ICA and CGP protocols.
 - Windows Filesystem (CIFS, including the SMB1 and SMB2 versions)
 - Outlook/Exchange (MAPI)

These optimizations reduce transaction time. This is done through rewriting, combining, and reordering commands, using read-ahead and write-behind, using a knowledge of the protocol for more advanced traffic shaping, and compression hinting.

6. **Compression engine.** Compression makes the transactions smaller, reducing the time it takes to transfer the data over the link. The Citrix SD-WAN WANOP compressor uses multiple compression algorithms, some very efficient for small transactions, some optimized for bulk transactions, and some for midsize transactions. Compression ratios of 10,000:1 are readily achieved by the Citrix SD-WAN WANOP compressor. The compressor is very fast, allowing high compression ratios to be maintained at full WAN speeds. With Citrix SD-WAN WANOP processing, a file that compresses at a 100:1 ratio can easily be sent over a 1 Mbps link with an overall throughput of 100 Mbps.

7. **Security engine.** Some Citrix SD-WAN WANOP features require that the two appliances enter a secure peer relationship with each other, and with the origin server. The security engine authenticates this peer relationship and encrypts the accelerated data connections between them. A secure peer relationship allows the use of SSL compression and the acceleration of encrypted Virtual Apps/Virtual Desktops (ICA/CGP), Windows Filesystem (CIFS), and Outlook/Exchange (MAPI) traffic.
8. **WAN-side flow control and auto-detection.** The WAN link is where traffic slowdowns occur, and if the link is congested, packets are lost and must be retransmitted. Retransmitting packets always causes a significant delay, sometimes lasting more one second. The WAN-side flow-control unit uses advanced retransmission elements and an advanced TCP/IP protocol for maximum performance in both “clean” and “troubled” links. The auto-detection unit identifies the presence of a partner Citrix SD-WAN WANOP unit on a connection-by-connection basis, which prevents optimizations from being used where they are not wanted, and allows new appliances to be detected by the existing ones as soon as they are added to the network. Auto-detection uses options in the TCP header field. This is normally transparent but might be blocked by some firewalls, which need to be reconfigured.
9. **Application classifier.** This unit examines all the traffic flowing through Citrix SD-WAN WANOP and identifies which application or protocol it belongs to. This information is used in reporting and by the traffic shaper.
10. **Traffic shaper.** To avoid congestion, excessive queuing, and other sources of avoidable delays, the traffic shaper injects traffic onto the WAN at slightly less than the WAN’s data rate, to ensure that the WAN is never overrun. A weighted fair queuing algorithm is used to ensure that all traffic gets its fair share of the link bandwidth. Traffic-shaping policies allow different traffic types to receive different weights, so that some traffic gets more bandwidth than others.

Receive pipeline

The pipeline in the receiving direction is similar to the sending direction, except that instead of encrypting, it decrypts, and instead of compressing, we have decompresses. Also, note that there is a traffic shaper in the receiving direction as well, applying traffic-shaping policies to incoming WAN traffic, so that both directions are regulated.

Auto-detection and packet-level transformation

The auto-detection algorithm inserts TCP header options to announce the presence of a Citrix SD-WAN WANOP appliance and to facilitate negotiation. These options are in the range of 24-31. The following packet-level transformations are used:

- On the initial packet of the connection (the SYN packet), the sending appliance attaches header options identifying itself as a Citrix SD-WAN WANOP appliance, and also declaring other capabilities, such as compression. This is called a “tagged SYN packet.”
- Upon receiving a tagged SYN packet, the receiving appliance attaches header options to the SYN-ACK packet, identifying itself in turn and announcing its capabilities.
- Once the sending appliance receives the tagged SYN-ACK packet, the connection can be accelerated according to whatever capabilities are shared by both appliances. For example, the connection is compressed if both appliances declared support for compression.
- The TCP initial sequence numbers (ISNs) in both directions are altered by adding 2,000,000,000 to the original values. This is a precaution that prevents the connection from continuing if one appliance fails or has a routing change that prevents it from seeing all the traffic in the connection. Once a connection is accelerated, it must remain accelerated throughout its lifetime.
- The MSS value is reduced, typically to 1380 bytes, to ensure that each packet has room for the inserted Citrix SD-WAN WANOP TCP header options.
- The IP addresses and port numbers of the connection remain unchanged.

Pre-acknowledgement

The SYN and SYN-ACK packets flow from end to end:

- The SYN packet flows from the endpoint client, through the client-side appliance, over the WAN, through the server-side appliance, and finally to the server.
- The SYN-ACK packet flows from the server, through the server-side appliance, over the WAN, through the client-side appliance, and finally to the client.

The same is true for the final packets of the connection, the FIN, FIN-ACK, and RST packets.

Other packets, however, are pre-acknowledged. For example, when the server-side appliance receives a packet from the server, it acknowledges it over the LAN right away, and buffers it for eventual transmission over the WAN. This allows the server-side appliance’s buffers to be filled very quickly, so it always has plenty of data to use for compression and other optimizations. (This is very different from normal TCP operation, where all acknowledgements come from the opposite side of the WAN, making acknowledgement very slow, and forcing every segment of the connection to move no faster than the slowest segment, greatly reducing the effectiveness of acceleration.)

Move traffic into and out of the appliance

Citrix SD-WAN WANOP appliances have a number of “forwarding modes.” A forwarding mode is a method of getting traffic into and out of the appliance. The most common is inline mode, where the

Citrix SD-WAN WANOP appears to be a bridge device. Packets entering on one bridge port appear to exit the other one. Of course, Citrix SD-WAN WANOP transforms data in a variety of ways, so in many cases the packet exiting the second port is not identical to the one that entered the first port, but that is how it appears to the rest of the network.

Where inline mode is not practical, several other methods are available, most notably WCCP mode. These are “one-arm” modes, using a single interface cable.

Tip

You can manage and monitor your Citrix SD-WAN WANOP appliances using Citrix ADM, for more information, see [Managing Citrix SD-WAN instances using Citrix ADM](#)

Get started with Citrix SD-WAN WANOP

March 12, 2021

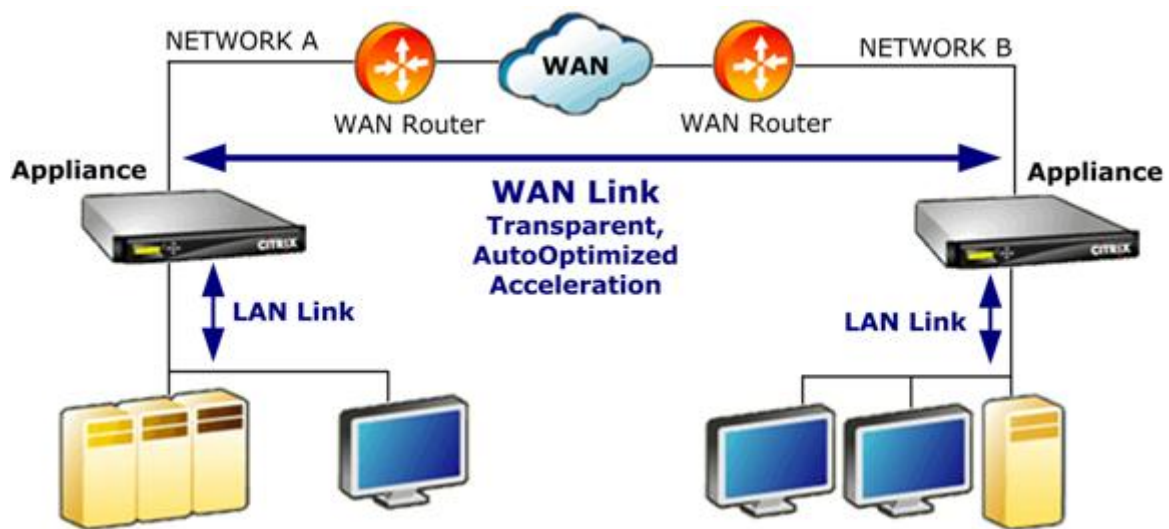
Deploying Citrix SD-WAN WANOP appliances successfully is not difficult, but improper deployments can cause problems and provide inadequate acceleration. Be sure to select appliances with sufficient capacity for the links that you want them to accelerate. Product selection is also one of the factors to consider when deciding how best to fit the appliances into your topology.

The most basic deployment criteria are:

- All packets in the TCP connection must pass through a supported combination of two *acceleration units* (Citrix SD-WAN WANOP appliances or Plug-ins).
- Traffic must pass through the two acceleration units in both directions.

When these criteria are met, acceleration is automatic.

Acceleration Enhances Performance when Traffic Passes through Two Appliances



For sites with only one WAN network, these criteria can be met by placing the Citrix SD-WAN WANOP appliance inline with the WAN. In more complex sites, other options are available. Some, such as WCCP support, are available on all models. Others are available on certain models only. Therefore, the needs of a more complex site might limit your choice of appliances.

When evaluating your options, consider the importance of keeping various segments of your network up and running in the event that a device fails or has to be disabled. For inline deployments, Citrix recommends an *Ethernet bypass card*. This card, which is optional on Citrix SD-WAN WANOP appliances, has a relay that closes if the appliance fails, allowing packets to pass through even if power is lost or removed.

Redundancy is a consideration for all types of deployments. Citrix SD-WAN WANOP appliances offer different types of redundancy:

- SD-WAN WANOP 4000/5000 appliances have dual power supplies.
- SD-WAN WANOP 4000/5000 appliances have redundant disk drives.
- Appliances can be used in high-availability mode (two redundant appliances with automatic failover). This mode is supported on all models.

Note

For more information on Citrix SD-WAN WANOP appliances and deployment modes, see the [SD-WAN WANOP platform documentation](#)

Select an appliance based on capacity

March 12, 2021

For proper operation, your Citrix SD-WAN WANOP appliance must have adequate resources to support the number of WAN links that you want to accelerate, and to support all of the users of those links. Three capacities are important when selecting a Citrix SD-WAN WANOP appliance: link capacity (bandwidth), user capacity, and disk capacity.

Link capacity

When selecting a Citrix SD-WAN WANOP appliance, the most important factor is that it support your WAN links. If your site has a single WAN link, your appliance should support your link speed. For example, a Citrix SD-WAN WANOP 2000-010 can supports links of up to 10 Mbps, which would be suitable for an 8 Mbps link but not a 12 Mbps link. If your site has multiple links that are to be accelerated by a single appliance, the appliance should support the total speed of all these WAN links added together.

The maximum supported speed is determined by a combination of the appliance hardware and the product license. The licensed bandwidth limit is the maximum link speed that is supported by the license.

Product	Licensed WAN BW Range
Current Products	
SD-WAN WANOP Plug-in	N/A
SD-WAN WANOP 400	2-6 Mbps
SD-WAN WANOP 800	2-10 Mbps
SD-WAN WANOP 2000 , 2000WS	10-50 Mbps
SD-WAN WANOP 3000	50-155
SD-WAN WANOP 4000	310-1,000 Mbps
SD-WAN WANOP 5000	1,500-2,000 Mbps
SD-WAN WANOP VPX	1-45 Mbps

Table 1. Licensed Bandwidth Limits by Product Line

Virtual Apps/Virtual Desktops user capacity

Each appliance is rated for a maximum number of XenApp or Virtual Desktops users. This value should not be exceeded when your deployment uses Virtual Apps or Virtual Desktops. If you are not using

Virtual Apps or Virtual Desktops, consider this number a rough guide to the number of users of other applications.

Product	Maximum Users
SD-WAN WANOP Plug-in	1
SD-WAN WANOP 400	10-30
SD-WAN WANOP 800	20-100
SD-WAN WANOP 2000 , 2000WS	100-300
SD-WAN WANOP 3000	300-500
SD-WAN WANOP VPX	20-350
SD-WAN WANOP 4000	750-2,500
SD-WAN WANOP 5000	3,500-5,000

Table 2. Virtual Apps/Virtual Desktops User Capacity

Disk size

Disk space is used mostly for compression history, and more disk space results in greater compression performance.

The SD-WAN WANOP 4000/5000 series offers from 1.8 TB to 2.4 TB of disk capacity. That compares to 2.1 TB for the SD-WAN WANOP 3000, 470 GB for the SD-WAN WANOP 2000, 80 GB for the SD-WAN WANOP 800, and 40 GB for the SD-WAN WANOP 400. SD-WAN WANOP VPX has a disk capacity of 100-500 GB. Ideally, an appliance should have a disk capacity larger than the cycle time of the link's data. For example, a link carrying mostly daily update traffic should have 24 hours of disk capacity or more. With a link carrying mostly user sessions, this window can be smaller. (A 1 Mbps link can transfer about 10 GB per day at full speed.)

Table 3. Examples of Data Lifetime for Disk Sizes

Appliance Model	Link Speed-1 Mbps	Link Speed-10 Mbps	Link Speed-100 Mbps	Link Speed-1000 Mbps
Data lifetime at 33% link utilization				
SD-WAN WANOP 800	23 days	2.3 days	NA	NA

Appliance Model	Link Speed-1 Mbps	Link Speed-10 Mbps	Link Speed-100 Mbps	Link Speed-1000 Mbps
SD-WAN WANOP 2000, 2000WS	141 days	14 days	NA	NA
SD-WAN WANOP 5000	717 days	72 days	7.2 days	17 hours
Data lifetime at 100% link utilization				
SD-WAN WANOP 800	8 days	19 hours	NA	NA
SD-WAN WANOP 2000, 2000WS	47 days	4.7 days	NA	NA
SD-WAN WANOP 5000	239 days	24 days	2.4 days	6 hours

Select the deployment mode based on datacenter topology

March 12, 2021

The appliance can be placed in line with your WAN link. The appliance uses two bridged Ethernet ports for inline mode. Packets enter one Ethernet port and exit through the other. This mode puts the appliance between your WAN router and your LAN. For the rest of the network, it is as if the appliance were not there at all. Its operation is completely transparent.

Inline mode has the following advantages over the other deployment modes:

- Maximum performance.
- Very easy configuration, using only the Quick Installation page.
- No reconfiguration of your other network equipment.

Other modes (WCCP, virtual inline, redirector) are less convenient to set up, generally requiring that you reconfigure your router, and they have somewhat lower performance.

A basic deployment consideration is whether your site has a single WAN router or multiple WAN routers. You also have to think about which features can be used in which modes. A requirement to support VPNs affects the placement of the appliance in your network.

Access Gateway appliances support Citrix SD-WAN WANOP TCP optimizations, enabling accelerated VPN connections when Citrix SD-WAN WANOP appliances are deployed with Access Gateway.

Overview of deployment modes

The appliance can be deployed in the following modes:

Forwarding modes

- **Inline mode**—Highest-performance, most transparent mode. Data flows in on one accelerated Ethernet port and out on the other. Requires no router reconfiguration of any kind.
- **Inline with dual bridges**—Same as inline, but with two independent accelerated bridges.
- **WCCP mode**—Recommended when inline mode is not practical. Supported by most routers. Requires only three lines of router configuration. To use WCCP mode on a Cisco router, the router should be running at least IOS version 12.0(11)S or 12.1(3)T. (WCCP stands for Web Cache Communications Protocol, but the protocol was greatly expanded with version 2.0 to support a wide variety of network devices.)
- **Virtual Inline mode**—Similar to WCCP mode. Uses policy based routing. Generally requires a dedicated LAN port on the router. Not recommended on units without an Ethernet bypass card. To use virtual inline mode on a Cisco router, the router should be running IOS version 12.3(4)T or later.
- **Group mode**—Used with two or more inline appliances, one per link, within a site. Recommended only when multiple bridges, WCCP, and virtual inline modes are all impractical.
- **High-availability mode**—Transparently combines two inline or virtual inline appliances into a primary/secondary pair. The primary appliance handles all the traffic. If it fails, the secondary appliance takes over. Requires no router configuration. Requires an appliance with an Ethernet bypass card.
- **Transparent Mode**—The recommended mode for communication with the Citrix SD-WAN WANOP Plug-in. In transparent mode, the Plug-in initiates connections in essentially the same way as the Citrix SD-WAN WANOP appliance, keeping the original IP address and port number of the connection and adding Citrix SD-WAN WANOP options to the TCP/IP headers of selected packets. By contrast, in redirector mode (not recommended), the Plug-in alters the destination IP and port numbers of the packets to match the signaling IP (and port) of the appliance.
- **Redirector mode** (not recommended)—Used by the Citrix SD-WAN WANOP Plug-in to forward traffic to the appliance. Can be used as a stand-alone mode or combined with one of the other deployments. Requires no router configuration.

Acceleration modes

- **Softboost mode**—A high-performance TCP variant that is recommended for most links. Although it provides less performance than hardboost mode, it works with any deployment. Acts like normal TCP, but faster.
- **Hardboost mode**—A highly aggressive, bandwidth-limited TCP variant useful for high-speed links, intercontinental links, satellite links, and other fixed-speed links for which achieving full link speed is difficult. Recommended for fixed-speed, point-to-point links where traffic shaping is not required.

Note






For more information on Citrix SD-WAN WANOP appliances and deployment modes, see the [Citrix SD-WAN WANOP platform documentation](#).

Sites with one WAN router

March 12, 2021

For a site with only one WAN router, the main issue in deployment is to allow the Citrix SD-WAN WANOP appliance to work in harmony with the router. The following figure shows the recommended deployment modes for a single router. Compare it to your router cabling to find the best mode for your environment.

Recommended Deployment Modes, Based on WAN Router Topology

A. Single LAN, Single WAN**B. Single LAN, Redundant WANs****C. Single LAN, WANs to Two or More Sites****D. Dual LANs, Single WAN****E. Multiple LANs, Multiple WANs**

Comments about the recommended deployment modes:

1. **Single LAN, Single WAN: Inline mode.** The router has a single active LAN interface and a single active WAN interface. The recommended mode for this case is inline mode, which provides the simplest installation, the most features, and the highest performance of any mode.
2. **Single LAN, Redundant WANs: Inline mode.** Inline mode is best for this configuration as well.
3. **Single LAN, Multiple WANs: Inline or WCCP.** This topology falls into two categories: hub-and-spoke or multihop. In a hub-and-spoke deployment, connections are mostly between a spoke site and the hub site. In a multihop deployment, many connections are between two spoke sites, with the data passing through the hub site. A single multihop connection can thus involve

as many as three appliances, depending on the details of where the hub site's appliance is positioned in the traffic flow.

For proper traffic shaping in multihop deployments, all WAN traffic on the hub site's WAN router must also pass through the appliance, instead of being passed by the router directly between WAN interfaces. In this case, WCCP is the preferred mode. If the deployment is hub-and-spoke, with most traffic terminating on the hub site, an inline deployment is preferable.

4. **Dual LANs, single WAN: Inline (with dual bridges) or WCCP.** This mode is supported by dual accelerated bridges, WCCP mode, or virtual inline mode.
5. **Multiple LANs, multiple WANs: Inline (dual bridges) or WCCP.** This is similar to Case C, but complicated by the presence of multiple LAN interfaces as well as multiple WANs. WCCP can always be used here. In the two-LAN case, an appliance with dual bridges can also be used in inline mode.

For more information, see the [table](#)

Sites with multiple WAN routers


March 12, 2021

More than one WAN router at the same site raises the possibility of *asymmetric routing*. Normally, IP networks are not affected by what path the packets take, so long as they arrive at their destination. However, the appliance relies on seeing every packet in the connection. “End-around” packets are not acceptable.

In a site with only one WAN router, asymmetric routing is not a problem, because the appliance can be placed in the path between the router and the rest of the site, so that traffic into or out of the router also passes through the appliance. But with two WAN routers, asymmetric routing can become an issue.

Asymmetric routing problems can appear during installation or later, as a result of failover to a secondary link, or other forms of dynamic routing and load balancing. The following figure shows an example sites that might suffer from asymmetric routing. If sites C and D always use the direct path, C-D or D-C, when sending traffic to each other, everything is fine. However, packets that take the longer path, C-E-D or D-E-C, bypass the appliances, causing new connections to be unaccelerated and existing connections to hang.

Asymmetric Routing




Asymmetric routing can be addressed by router configuration, appliance placement, or appliance configuration.

If the router is configured to ensure that all packets of a given connection always pass through the appliance in both directions, there is no asymmetry.

If the appliance is positioned after the point where all the WAN streams are combined, asymmetry is avoided, and all traffic is accelerated, as shown in the following figure.

Avoiding Asymmetric Routing through Proper Placement of the Appliance




Configuring the appliance to use one of the following asymmetry-resistant forwarding modes can eliminate the problem:

- *Multiple Bridges.* An appliance with two accelerated bridges, or *accelerated pairs*, (for example, apA and apB), allows two links to be accelerated in inline mode. The two links can be fully independent, load-balanced, or primary/backup links.
- *WCCP mode* allows a single appliance to be shared between multiple WAN routers, allowing it to handle all the WAN traffic regardless of which link it arrives on.

- *Virtual inline mode* allows a single appliance to be shared between multiple WAN routers, allowing it to handle all the WAN traffic regardless of which link it arrives on.
- *Group mode* allows two or more inline appliances to share traffic with each other, ensuring that traffic that arrives on the wrong link is handed off properly. Because group mode requires multiple appliances, it is an expensive solution that is best suited to installations where the accelerated links have wide physical separation, making the other alternatives difficult. For example, if the two WAN links are on different offices in the same city (but the campuses are connected by a LAN-speed link), group mode might be the only choice.


Eliminating Asymmetric Routing by Using Group Mode or Virtual Inline Mode



Note

One end of the link can use virtual inline mode while the other end uses group mode. The two ends of a link do not have to use the same forwarding mode.

Sites with Only One WAN Link Cannot Have Asymmetric Routing Problems



Appliance failure handled in various deployment modes

March 12, 2021

Citrix SD-WAN WANOP appliances have safeguards against loss of connectivity in case of software, hardware, and power failures. These safeguards are mode-dependent.

In **inline mode**, appliances maintain network continuity in the event of hardware, software, or power failure. If present, the bypass relay in the appliance closes if power is lost or some other failure occurs. Inline appliances without a bypass card usually block traffic in the event of a serious failure, but they continue to forward traffic under some conditions, namely, when the network stack is running but the acceleration software has been disabled or has shut itself down because of persistent errors.

Existing accelerated connections usually become unresponsive after a failure and are eventually terminated by the application or the network stack at one of the end points. Some accelerated connections might continue as unaccelerated connections after the failure. New connections run in unaccelerated mode.

When the appliance comes back online, existing connections continue as unaccelerated connections. New connections are accelerated.

In **WCCP mode**, the router bypasses an appliance that stops responding, and reopens the connection when the appliance begins responding again. The WCCP protocol has integral health-checking.

If the “verify-availability” option is used with **virtual inline mode**, the router behaves like it does with WCCP mode, bypassing the appliance when it is not available and reconnecting when it is. If “verify-availability” is not used, all packets forwarded to the appliance are dropped if the appliance is not available.

In **group mode**, an appliance can be configured to fail “open”(bridging disabled) or “closed”(bridging or bypass relay enabled).

In **high availability** mode, if one HA appliance fails, the other takes over automatically. The appliances’ bypass cards are disabled in HA mode, so if the HA appliances are in inline mode and both appliances fail, connectivity is lost.

In **redirector mode**, the Citrix SD-WAN WANOP Plug-in performs health checking on redirector-mode appliances and bypasses unresponsive appliances, sending traffic directly to endpoint servers instead.

Supported mode and feature matrix

March 12, 2021

In general, all modes are simultaneously active. However, some combinations should not be used together, as shown in the following table.

**Supported
Combi-
nations,
Units
WITH
Ethernet
Bypass
Cards**

Config.	Inline	Virtual Inline	WCCP- GRE	WCCP- L2	Multiple Bridges	High Avail.	Group Mode
Citrix SD-WAN WANOP Plug-in	Y	Y	Y	Y	Y	Y	N
Inline	Y	N	N	N	Y	Y	Y
Virtual Inline		Y	Y	Y	Y	Y	N
WCCP- GRE			Y	Y	Y	Y	N
WCCP- L2				Y	Y	Y	N
Multiple Bridges					Y	Y	N
High Avail.						Y	Y

**Supported
Combi-
nations,
Units
WITH-
OUT
Ethernet
Bypass
Cards**

**Supported
Combi-
nations,
Units
WITH
Ethernet
Bypass
Cards**

Config.	Inline	Virtual Inline	WCCP- GRE	WCCP- L2	Multiple Bridges	High Avail.	Group Mode
Citrix SD-WAN WANOP Plug-in	N	N	N	N	N	N	N
Inline	Y	N	N	N	N	N	N
Virtual Inline		Y	Y	Y	N	N	N
WCCP- GRE			Y	Y	N	N	N
WCCP- L2				Y	N	N	N
Multiple Bridges					N	N	Y
High Avail.						N	N

Y = Yes, supported. N = Not supported.

Configure Citrix SD-WAN WANOP plug-in with Access Gateway VPNs

March 12, 2021

The Access Gateway Standard Edition VPN supports Citrix SD-WAN WANOP Plug-in acceleration, provided that a Citrix SD-WAN WANOP appliance is deployed with the Access Gateway appliance and the Access Gateway appliance is configured to support it.

For Citrix SD-WAN WANOP Plug-in support with other VPNs, see your VPN documentation or contact your Citrix representative.


To configure Citrix SD-WAN WANOP support, use the Access Gateway administration tool, as follows:

1. On the Global Cluster Policies page, under Advanced Options, select the **Enable TCP optimization with Citrix SD-WAN WANOP Plug-in** check box.
2. Make sure that the IP addresses used by the Citrix SD-WAN WANOP (redirector IP and management IP) have access enabled in the Network Resources section on the Access Policy Manager page.
3. For each of these addresses, enable all protocols (TCP, UDP, ICMP) and enable Preserve TCP Options.
4. Make sure that these same addresses are included under User Groups: Default: Network Policies on the Access Policy Manager page.

VPN support options


VPN support is simply a matter of putting the appliance on the LAN side of the VPN, as shown in the following figure. This placement ensures that the appliance receives and transmits the decapsulated, decrypted, plain-text version of the link traffic, allowing compression and application acceleration to work. (Application acceleration and compression have no effect on encrypted traffic. However, TCP protocol acceleration works on encrypted traffic.)

VPN Cabling for an Inline VPN




The following figure shows one option for accelerating one-arm VPNs. The appliance is on the server side of the VPN. All VPN traffic with a local destination is accelerated. VPN traffic with a remote destination is not accelerated. Non-VPN traffic can also be accelerated.

One-Arm VPN Acceleration, Option A



The following figure shows another option for accelerating one-arm VPNs. The appliance is on the server side of the VPN. All VPN traffic with a local destination is accelerated. VPN traffic with a remote destination is not accelerated. Non-VPN traffic can also be accelerated.

One-Arm VPN Acceleration, Option B



Important

For acceleration to be effective, the VPN must preserve TCP header options. Most VPNs do so.


Deploy SD-WAN WANOP VPX on Microsoft Azure

March 12, 2021

Citrix SD-WAN WANOP Edition is now available in the Azure marketplace, enabling WAN optimization between enterprise datacenter/branch and Azure cloud. Since L2 mode support is not available on cloud infrastructures, you cannot deploy Citrix SD-WAN WANOP as a standalone VPX in Azure Cloud.

However, you can deploy Citrix SD-WAN WANOP VPX along with Citrix ADC VPX in Azure cloud infrastructure. The Citrix ADC uses cloud connector to create an IPsec tunnel, while the Citrix SD-WAN WANOP VPX accelerates the connections, providing LAN-like performance for applications.

Citrix SD-WAN WANOP in Azure cloud topology



The topology diagram shows a Citrix SD-WAN 4000/5000 deployed in the data center or branch premises. You could also deploy Citrix SD-WAN WANOP and Citrix ADC appliance in two-box mode or it could both be VPX. On the Azure cloud VNET, the Citrix SD-WAN WANOP VPX is deployed in one-arm (PBR) mode with the Citrix ADC VPX.

Deployment overview

To deploy SD-WAN WANOP on Microsoft Azure:

1. Deploy a Citrix ADC VPX instance on the Azure cloud. For more information, see [Deploy a Citrix ADC VPX instance on Microsoft Azure](#). Configure four network interfaces in four different subnets and enable IP forwarding on all the network interfaces. The four network interfaces are used as:
 - Management interface
 - WAN side interface, for IPsec tunnel
 - LAN side interface, to connect to the server
 - WANOP communication interface, to communicate with the Citrix SD-WAN WANOP VPX on the Azure cloud.
2. Deploy a Citrix SD-WAN WANOP VPX on Azure cloud. For more information, see the deployment procedure below.

Note: Enable IP forwarding on WANOP interface.
3. Configure an IPsec tunnel between the on-premise appliance and the Citrix ADC VPX on Azure cloud, using the public IP address of Citrix ADC WAN interface. For more information on configuring IP tunnels see, [IP Tunnels](#).

4. Configure Citrix ADC VPX to redirect the packets to Citrix SD-WAN WANOP VPX. Use the private IP address of WANOP communication interface and create a load balancing virtual server. For more information, see [Create a load balancing virtual server](#).
5. Configure the following route tables on Azure:
 - Route table for WANOP facing interface on Citrix ADC VPX –Route table entries should have source and destination address as client and server subnets respectively. The Citrix ADC VPX's WANOP facing interface IP address is the next hop.
 - Route table for Citrix SD-WAN WANOP interface - Route table entries should have source and destination address as client and server subnets respectively. The Citrix SD-WAN WANOP interface IP address is the next hop.

In the above example, when the source tries to access an application on the cloud destination, the packets flow through the established IPsec tunnel. At the Azure cloud VNET end, the Citrix ADC VPX receives the packets, decrypts, and forwards it to the Citrix SD-WAN WANOP VPX. The Citrix SD-WAN WANOP VPX processes the packets, optimizes it, and sends it back to Citrix ADC VPX. The Citrix ADC VPX sends the packet to the destination. On the return path, the Citrix ADC VPX forwards the packets to Citrix SD-WAN WANOP VPX for optimization. The optimized packets are transmitted back to the source through the established IPsec tunnel.

Deploy Citrix SD-WAN WANOP VPX on Microsoft Azure

To deploy Citrix SD-WAN WANOP VPX on Microsoft Azure:

1. In Microsoft Azure, navigate to **Home > Marketplace > Networking**, search for **Citrix SD-WAN WANOP** and install it.
2. On the Citrix SD-WAN WAN OP page, from the drop-down list select **Resource Manager** and click **Create**. The **Create Citrix SD-WAN WAN Optimization** page appears.
3. In the **Basics** section, select the subscription type, resource group, and location. Click OK.

Note: You can choose to create a resource group. 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.

The screenshot shows the 'Create Citrix SD-WAN WANOP' wizard in the Azure portal. The left sidebar contains navigation links: 'Create a resource', 'All services', 'FAVORITES', 'Dashboard', 'All resources', 'Resource groups', 'App Services', 'Function Apps', 'SQL databases', 'Azure Cosmos DB', 'Virtual machines', 'Load balancers', 'Storage accounts', 'Virtual networks', 'Azure Active Directory', 'Monitor', and 'Advisor'. The main area displays a five-step process: 1. Basics (Configure basic settings), 2. Administrator settings (Configure deployment settings), 3. Citrix SDWAN WANOpt myappl... (Configure Citrix SD-WAN WAN...), 4. Summary (Citrix SD-WAN WAN Optimisat...), and 5. Buy. The 'Basics' step is currently active. The right pane shows the configuration for 'Basics': 'Subscription' is set to 'Enterprise Dev/Test'; 'Resource group' is set to 'surya_wanpt-test' with 'Use existing' selected; and 'Location' is set to 'East US 2'. An 'OK' button is highlighted with a red box at the bottom right.

Create Citrix SD-WAN WANOP Opti... X Basics X

1 Basics
Configure basic settings >

2 Administrator settings
Configure deployment settings >

3 Citrix SDWAN WANOpt myappl...
Configure Citrix SD-WAN WAN... >

4 Summary
Citrix SD-WAN WAN Optimisat... >

5 Buy >

Subscription
Enterprise Dev/Test v

* Resource group ⓘ
☐ Create new ☒ Use existing
surya_wanpt-test v

* Location
East US 2 v

OK

4. In the **Administrator** section, enter the name and credentials for the Citrix SD-WAN WANOP virtual machine. Click **OK**.

The screenshot shows the 'Create Citrix SD-WAN WAN Opti...' window with the 'Administrator settings' section selected. The left sidebar lists various services, and the main pane shows a progress bar with five steps: 1. Basics (Done), 2. Administrator settings (selected), 3. Citrix SDWAN WANOpt myappl..., 4. Summary, and 5. Buy. The 'Administrator settings' section contains four fields: 'Virtual Machine name' (citrixwanopt), 'Username' (suryaprakp), 'Password' (masked), and 'Confirm password' (masked). All fields have a green checkmark indicating they are valid. An 'OK' button is highlighted with a red box at the bottom right.

Step	Section	Action
1	Basics	Done
2	Administrator settings	Configure deployment settings
3	Citrix SDWAN WANOpt myappl...	Configure Citrix SD-WAN WAN...
4	Summary	Citrix SD-WAN WAN Optimisat...
5	Buy	

Administrator settings

- * Virtual Machine name
- * Username
- * Password
- * Confirm password

OK

5. In the **Citrix SD-WAN WANOP settings section**, configure the setting for the Citrix SD-WAN WANOP VPX as per your requirements. Click **OK**.

1 Basics Done ✓

2 Administrator settings Done ✓

3 Citrix SDWAN WANOpt myappl... >
Configure Citrix SD-WAN WAN...

4 Summary >
Citrix SD-WAN WAN Optimisat...

5 Buy >

* Virtual machine size ⓘ
1x Standard D3 v2 >

OS Disk Size(GB) ⓘ
50

* Storage account ⓘ
suryausregion >

* Public IP address for management... ⓘ
(new) sdwanwanopt-mgmt >

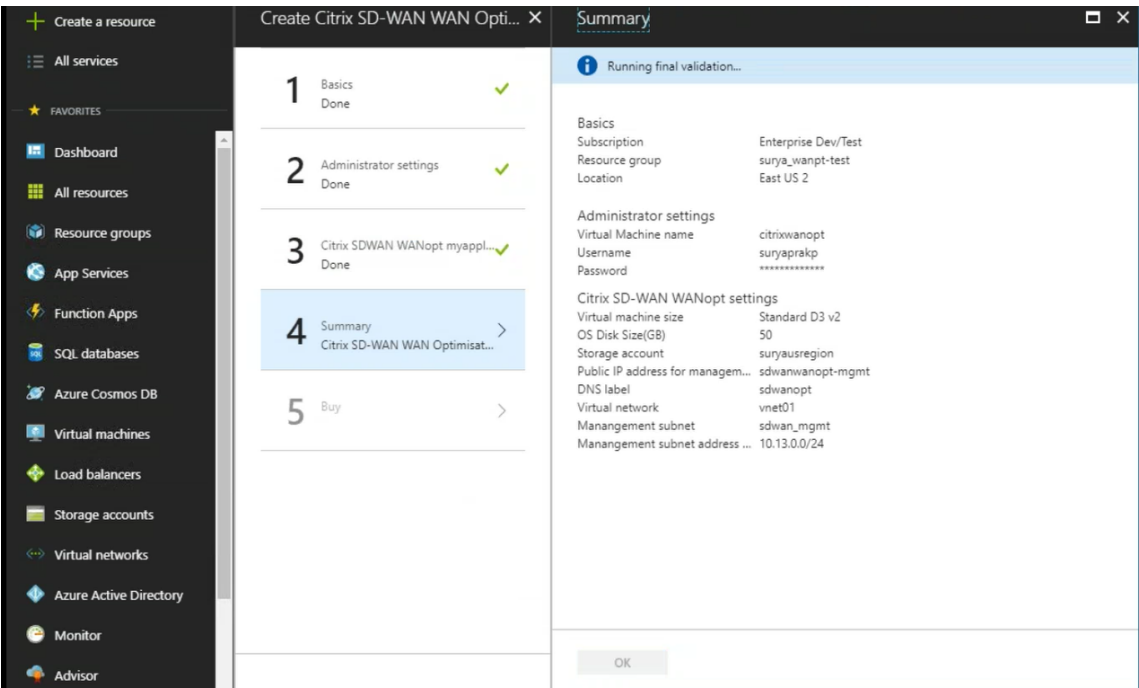
* DNS label ⓘ
sdwanopt ✓
eastus2.cloudapp.azure.com

* Virtual network ⓘ
(new) vnet01 >


* Subnets ⓘ
Review subnet configuration >

OK

6. The configuration that you provided in previous steps is validated and applied. If you have configured correctly, the validation passed message appears. Click **OK**.



7. After successful deployment, navigate to **Virtual Networks** to view the Citrix SD-WAN WANOP VPX. You can further configure the virtual machine parameters using the settings option.



SD-WAN WANOP upgrading procedure

March 12, 2021

This section provides information about downloading and upgrading the Citrix SD-WAN WAN Optimization (WANOP) software packages.

Note:

Before you download the software, you must obtain and register a Citrix SD-WAN software license. For information, see [Licensing](#).

Download the software packages

To download the Citrix SD-WAN WANOP software packages, go to the URL; [product downloads](#). Instructions for downloading the software are provided on this site.

To download the Citrix SD-WAN WANOP software package:


1. Log on to [citrix.com](#) with your credentials.
2. Go to [Downloads](#) page and select the product (Citrix SD-WAN) from the drop-down list.
3. Expand the **Citrix SD-WAN WANOP edition** and select the required software release.
4. The following download options are available. Download the required software.
 - Download .upg upgrade file for SD-WAN WANOP 400/800/1000/1000WS/2000/2000WS/3000/4000/4100 appliances.
 - Download .bin upgrade file for SD-WAN WANOP VPX appliances.

For more information on the SD-WAN WANOP supported platforms, see [SD-WAN platform models and software packages](#).


Upgrade procedure

Perform the following procedure to update software:


1. Navigate to **Configuration > Maintenance > System Administration** > click **Update Software**.



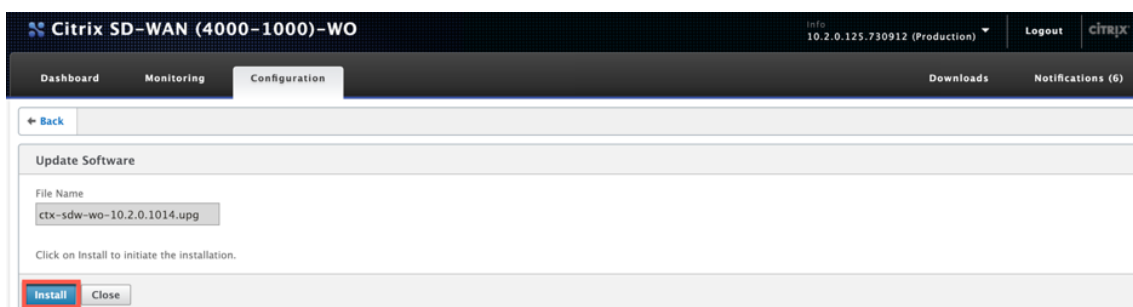
2. Click **Browse** to provide the **ctx-sdw-wo-10.2.X.upg** file. Click **OK**.



You can see the uploading status bar.



3. When a message announces that the upload was successful, click **Install**.



4. The appliance performs the upgrade, which takes 10–40 mins based on the platform model. It displays a series of status messages, starting with **Preparing to upgrade** and ending with **Upgrade completed Successfully**.
5. Click **OK** to display the updated user interface.

Initial Configuration

March 12, 2021

After checking the connections, you are ready to deploy the SD-WAN 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

March 12, 2021

To deploy a Citrix SD-WAN 4100 or 5100 appliance, you must complete the following prerequisite setup before configuring the appliance.

Software versions

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.

License File

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.

Installing the hardware

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

March 12, 2021

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 retains 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 high availability. 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.

	Parameter	Example	Your Value	Description
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.
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.

	Parameter	Example	Your Value	Description
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 receives acceleration.
T8.	GRE Router Host ID	NA		WCCP-GRE only. Host ID of GRE router.
T9.	Traffic Port	10/1		Port used for accelerated traffic.
T10+.	(Inline) more 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.

	Parameter	Example	Your Value	Description
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

March 12, 2021

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

March 12, 2021

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
1. Start the configuration wizard by clicking **Get Started**.
2. On the **Platform Configuration** page, enter respective values from your worksheet, as shown in the following example:
3. Click **Done**. A screen showing the Installation in Progress...message appears. This process takes approximately 2 to 5 minutes, depending on your network speed.
4. A Redirecting to new management IP message appears.
5. Click **OK**.
6. Unplug your computer from the Ethernet port and connect the port to your management network.
7. Reset the IP address of your computer to its previous setting.

8. From a computer on the management network, log on to the appliance by entering the new management service IP address, such as https://<Management_IP_Address>, in a web browser.
9. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
10. Log on to the appliance by using the **nsroot** user name and the password from your [worksheet](#).
11. To complete the configuration process, see [Provisioning the Appliance](#).

Assigning a Management IP Address through the Serial Port

March 12, 2021

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
1. At the **\$** prompt, run the following command to switch to the shell prompt of the appliance:
`$ ssh 169.254.0.10`
2. Enter **Yes** to continue connecting to the management service.
3. Log on to the shell prompt of the appliance with the following default credentials:
Password: nsroot.
4. At the logon prompt, run the following command to open the Management Service Initial Network Address Configuration menu: # `networkconfig`
5. Type **1** and press **Enter** to select option 1, and specify a new management IP address for the management service.

6. Type **2** and press **Enter** to select option 2, and specify a new management IP address for the Citrix Hypervisor.
7. Type **3** and press **Enter** to select option 3, and specify the network mask for the IP addresses.
8. Type **4** and press **Enter** to select option 4, and specify the default gateway for the management service IP address.
9. Type **8** and press **Enter** to save the settings and exit.
10. 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.
11. To continue the configuration, accept the certificate and continue. The option to continue varies according to the web browser you are using.
12. To complete the configuration process, see [Provisioning the Appliance](#).

Provisioning the Appliance

March 12, 2021

After assigning an IP address to the management service, you are ready to provision the NetScaler and accelerator instances. When 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 screenshot displays sample values used in this documentation. Enter values as follows:

- **Citrix Hypervisor IP Address**—(Item M4 on your worksheet, or H4 if this is the second appliance in a high availability pair.) The management address of the built-in Citrix Hypervisor 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 a high availability 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:
 - a) **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.

- b) **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 is 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.

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

March 12, 2021

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

Destination IP Address	Destination Ethernet Address	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.

SD-WAN 4100/5100 appliances have two recommended deployment modes: WCCP and inline. These modes are commonly used without high availability (high availability), and less commonly with high availability.

Currently, Citrix recommends WCCP mode, with a single router and without high availability, for most deployments. Use inline mode when WCCP is not available.

Although not all of the following modes are recommended currently, 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 high availability
- Inline mode
- Inline mode in high availability
- Virtual inline mode
- Virtual inline mode in high availability

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.

Customizing the Ethernet ports

March 12, 2021

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:

Ethernet Port	Name
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)

Table 1. Ethernet Port Names

Port Parameters

March 12, 2021

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)

March 12, 2021

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

March 12, 2021

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

March 12, 2021

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.

Customizing the Ethernet ports

March 12, 2021

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:

Ethernet Port	Name
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)

Table 1. Ethernet Port Names

Ethernet Bypass and Link-Down Propagation

March 12, 2021

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, high availability heartbeat, and other tasks.
- If the Primary port is enabled, the appliance assumes (without checking) that the Primary port is used for management, high availability 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

March 12, 2021

[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

March 12, 2021

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.

WCCP Mode

March 12, 2021

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 an 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 an 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 must 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 must 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 must be sent to it, and also which WCCP features must 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 10 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 high availability 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 high availability 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. Note that apA and apB need to be on different subnet.

WCCP Mode (Non-Clustered)

March 12, 2021

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

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.

Configure the Router

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:

```

1 config term
2 ip wccp version 2
3 ! We will configure the appliance to use group 51 for TCP and 52 for
  UDP.
4 ip wccp 51
5 ip wccp 52
6
7 ! Repeat the following three lines for each WAN interface
8 ! you wish to accelerate:
9 interface your_wan_interface
10 ! If Reverse Path Forwarding is enabled (with an ip verify unicast
11 ! source reachable " statement), delete or comment out the statement:
12 ! ip verify unicast source reachable-via any
13 ! Repeat on all ports.
14
15 ip wccp 51 redirect out
16 ip wccp 51 redirect in
17 ip wccp 52 redirect out
18 ip wccp 52 redirect in
19
20 ! If the appliance is inline with one of the router interfaces
21 ! (NOT SUPPORTED), add the following line for that interface
22 ! to prevent loops:
23 ip wccp redirect exclude in
24 ^Z
25 <!--NeedCopy-->

```

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:

```

1 config term
2 ip wccp version 2
3 ip wccp 51 group-address 225.0.0.1
4
5 ! Repeat the following three lines for each WAN interface
6 ! you wish to accelerate:
7 interface your_wan_interface
8 ! If Reverse Path Forwarding is enabled (with an ip verify unicast

```

```

 9 ! source reachable " statement), delete or comment out the statement:
10 ! ip verify unicast source reachable-via any
11
12 ip wccp 51 redirect out
13 ip wccp 51 redirect in
14 !
15 ! The following line is needed only on the interface facing the other
    router,
16 ! if there is another router participating in this service group.
17 ip wccp 51 group-listen
18
19 !If the appliance is inline with one of the router interfaces,
20 !(which is supported but not recommended), add
21 !the following line for that interface to prevent loops:
22 ip wccp redirect exclude in
23 ^Z
24 <!--NeedCopy-->

```

Basic Configuration Procedure for WCCP Mode on the SD-WAN Appliance

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 more 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 preceding 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

In a service group, a WCCP router and an 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 an 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

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:

```
1 Router>enable
2 Password:
3 Router#show ip wccp
```

```
4 Global WCCP information:
5   Router information:
6     Router Identifier:          172.16.2.4
7     Protocol Version:          2.0
8
9   Service Identifier: 51
10    Number of Cache Engines:    0
11    Number of routers:         0
12    Total Packets Redirected:   19951
13    Redirect access-list:      -none-
14    Total Packets Denied Redirect: 0
15    Total Packets Unassigned:   0
16    Group access-list:         -none-
17    Total Messages Denied to Group: 0
18    Total Authentication failures: 0
19 <!--NeedCopy-->
```

Verify WCCP Mode

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

March 12, 2021

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

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 Virtual Apps users.

As shown on the [SD-WAN Datasheet](#), an 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, as much 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 Deployments

Using multiple WAN routers is 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

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. 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 cannot be used.

Deployment worksheet and cluster limitations

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 Citrix Virtual Apps and Desktops Users Per Appliance	$U_{spec} =$	From data sheet
Citrix Virtual Apps and Desktops Users on WAN Link	$U_{wan} =$	—
User overload Factor	$U_{overload} = U_{wan}/U_{spec} =$	—
Supported BW Per Appliance	$BW_{spec} =$	From data sheet
WAN Link BW	$BW_{wan} =$	—
BW Overload Factor	$BW_{overload} = BW_{wan}/BW_{spec}$ $=$	—
Number of appliances required	$N = \max(U_{overload}, BW_{overload}) + 1 =$	Includes one spare
Min number of buckets	$B_{min} = N$, rounded up a power of 2 =	—
If SD-WAN 4000 or 5000,	$B_{min} = 2N$, rounded up to a power of 2 =	—
Recommended value	$B = 4 \setminus B_{min}$ if $B_{min} \leq 16$, else 2 $\setminus B_{min} =$	—
Number of “one”bits in address mask	$M = \log_2(B)$	If $B=16$, $M=4$.

Mask value: The mask value is a 32-bit address mask with several “one”bits equal to M in the worksheet provided earlier. 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 $\log_2(\text{mask size})$: if mask size is 16, set 4 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 few bits in the network field (only 2 bits). If so 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 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 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 cannot be used.

Testing and Troubleshooting

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.

Configure WCCP Clustering

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](#)

Virtual Inline Mode

March 12, 2021

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

March 12, 2021

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

March 12, 2021

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:

““ pre codeblock

```
!- 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
```

```
1 This rule pings the appliance at 192.168.1.200 periodically. You can
  test against 123 to see if the unit is up.
2
3 ## Routing Examples
4
5 The following examples illustrate configuring Cisco routers for the
  local and remote sites shown in [Virtual inline example](/en-us/
  citrix-sd-wan-wanop/11/cb-deployment-modes-con/br-adv-virt-inline-
  mode-con.html). To illustrate health monitoring, the configuration
  for the local site includes health monitoring, but the configuration
  for the remote site does not.
6
7 Note: The configuration for the local site assumes that a ping monitor
  has already been configured.
8
9 The examples conform to the Cisco IOS CLI. They might not be applicable
  to routers from other vendors.
10
11 Local Site, Health-Checking Enabled:
12
13 ``` pre codeblock
14 !
15 ! For health-checking to work, do not forget to start
16 ! the monitoring process.
17 !
18 ! Original configuration is in normal type.
19 ! appliance-specific configuration is in bold.
20 !
21 ip cef
22 !
23 interface FastEthernet0/0
24 ip address 10.10.10.5 255.255.255.0
25 ip policy route-map client_side_map
26 !
27 interface FastEthernet0/1
28 ip address 172.68.1.5 255.255.255.0
29 ip policy route-map wan_side_map
30 !
31 interface FastEthernet1/0
32 ip address 192.168.1.5 255.255.255.0
33 !
34 ip classless
35 ip route 0.0.0.0 0.0.0.0 171.68.1.1
36 !
37 ip access-list extended client_side
38 permit ip 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
39 ip access-list extended wan_side
40 permit ip 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
41 !
42 route-map wan_side_map permit 20
43 match ip address wan_side
44 !- Now set the appliance as the next hop, if it's up.
45 set ip next-hop verify-availability 192.168.1.200 20 track 123
```

```
46 !
47 route-map client_side_map permit 10
48 match ip address client_side
49 set ip next-hop verify-availability 192.168.1.200 10 track 123
50 <!--NeedCopy-->
```

Remote Site (No Health Checking):

“pre codeblock

! 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

!_

```
1 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.
2
3 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):
4
5 ``` pre codeblock
6 !
7 ip access-list extended client_side
8 permit tcp 10.16.20.0 0.0.0.255 10.10.10.0 0.0.0.255
9 ip access-list extended wan_side
10 permit tcp 10.10.10.0 0.0.0.255 10.16.20.0 0.0.0.255
11 !
12 <!--NeedCopy-->
```

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

March 12, 2021

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

March 12, 2021

Virtual Inline mode can be used in a high availability (high availability) configuration. The below figure shows a simple high availability deployment. In virtual inline mode, a pair of appliances acts as one virtual appliance. Router configuration is the same for an high availability pair as with a single appliance, except that the Virtual IP address of the high availability 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

March 12, 2021

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

March 12, 2021

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

March 12, 2021

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

March 12, 2021

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

March 12, 2021

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 high availability Secondary SSL Common Name is blank. If the other group member is a high-availability pair, specify the SSL Common Name of the high availability secondary appliance in the high availability 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:
 - a) 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.

- b) 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

March 12, 2021

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

March 12, 2021

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.

Customizing the Ethernet ports

March 12, 2021

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:

Ethernet Port	Name
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)

Table 1. Ethernet Port Names

How High-Availability Mode Works

March 12, 2021

In a high availability (high availability) 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 high availability mode. If both appliances in an inline high availability 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 high availability 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 high availability 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.

high availability cannot be enabled unless both appliances are running the same software release.

high availability in WCCP mode—When WCCP is used with an high availability 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

March 12, 2021

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.

Other Requirements

March 12, 2021

Both appliances in an high availability 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 high availability display a warning on the Configuration: High Availability page.

Management Access to the High-Availability Pair

March 12, 2021

When configuring a high-availability (high availability) 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 high availability VIP for all management tasks. You can, however, disable the secondary appliance's high availability state from its management UI.

Configuring the High-Availability Pair

March 12, 2021

You can configure two newly installed appliances as a high-availability pair, or you can create an high availability 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 high availability 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 high availability 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 high availability 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

March 12, 2021

Updating the SD-WAN software on an high availability 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 high availability 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, high availability should become fully established, because both appliances are running the same software.

Saving/Restoring Parameters of an high availability Pair

March 12, 2021

The System Maintenance: Backup/Restore function can be used to save and restore parameters of an high availability pair as follows:

To back up the parameters

Use the backup feature as usual. That is, log on to the GUI through the high availability VIP address (as is normal when managing the high availability pair) and, on the System Management: Backup/Restore page, click Download Settings.

To restore the parameters

1. Disable high availability on both appliances by clearing the Enabled check box on the Configuration: Advanced Deployments: High Availability (high availability) 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 high availability).
5. Wait for Appliance B to restart. It becomes the primary.
6. Restart Appliance A.
7. Log on to Appliance A’s GUI and reenables high availability on the Configuration: Advanced Deployments: High Availability (high availability) 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

March 12, 2021

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 high availability 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 high availability heartbeat to pass between them.
- The high availability/Group Mode SSL Certificates on one or both appliances are damaged or missing.

Two box mode

March 12, 2021

Two box mode is a WCCP one-arm based deployment where the SD-WAN SE appliance acts as a WCCP router and the SDWAN-WANOP (4000/5000) 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 by 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
- SD-WAN WANOP appliances –4000, 4100, 5000, and 5100

Note

High Availability and WCCP deployment modes are not accessible when Two Box mode is enabled. However, these deployment modes are 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.

Citrix SD-WAN standard edition

To configure two-box mode solution in the Standard Edition appliance at the DC or Branch site:

1. In the SD-WAN SE web management interface, go to **Configuration > Virtual WAN > Configuration Editor**. Open an existing configuration package or create a 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 signify which virtual network interface the appliance is 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. Do not enable the WCCP Listener 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.

Citrix SD-WAN WANOP configuration

To configure two-box deployment mode in the SD-WAN WANOP appliance web GUI:

1. In the 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 **OK**.
3. Enable the **Two Box Enabled** checkbox.
4. Enter the **Peer IP**. Peer IP is the 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 SDWAN-SE **Configuration** page under **APPFLOW** node. This enables SDWAN-SE to act as a single pane for configuration of APPFLOW and other data processing configuration attributes such as Service Class, Application Classifiers. The configuration done on the SDWAN-SE reflects on the SDWAN-WO configuration, maintaining seamless APPFLOW functionality support.

SD-WAN WANOP already discovered by Citrix Application Delivery Management (ADM), if used in Two Box Mode, should be isolated and not configured using Citrix ADM until this mode is turned off. This is because the configuration of WANOP for traffic processing is managed by the SD-WAN SE appliance in the Two Box Mode.

Advanced Optimizations or Secure Acceleration should be directly configured on the SDWAN-SE appliance like we would configure on the SDWAN-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 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 SDWAN-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.

1. Go to **Configuration > Virtual WAN > WAN Optimization**
2. Appflow node under **Configuration > Appliance Settings**
3. WAN Optimization 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 Premium (Enterprise) edition appliance.

Management IP Address Change for SD-WAN WANOP Appliance in Two Box Mode

To change the management IP address of SDWAN-WANOP appliance in Two box mode:

1. Execute command *clear_wo_sync* on the SD-WAN SE appliance. It 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.

The management IP address is used for peer IP address 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:

1. Disable the Two Box mode from SD-WAN WANOP appliance.
2. It is expected to 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*.

FAQs

March 12, 2021

- [Acceleration](#)
- [Compression](#)
- [CIFS and MAPI](#)
- [RPC over HTTP](#)
- [SCPS](#)
- [Secure Peering](#)
- [SSL Acceleration](#)
- [CitrixSD-WAN WANOP Plug-in](#)
- [Traffic Shaping](#)
- [Upgrade](#)
- [Video Caching](#)
- [Office 365](#)

Acceleration

March 12, 2021

Does acceleration use a tunnel?

No, acceleration is transparent, using the same IP addresses and port numbers as the original connection. This allows your current monitoring methods to continue to work normally.

How does acceleration change the packet stream?

With non-compressed connections, acceleration adds options to the packet's TCP header, but leaves the packet payload intact. These options allow the Citrix SD-WAN WANOP devices at each end of the connection to communicate with each other. In addition, the TCP sequence number is adjusted to prevent routing issues or appliance failure from mixing accelerated packets and non-accelerated packets in the same connection.

With compressed connections, the payload is compressed, of course, and the output of the compressor is accumulated into full-sized packets. The upshot is that, for example, 3:1 compression results in one-third as many packets being transmitted, rather than the same number of packets, each reduced to one-third size. Compression also uses Citrix SD-WAN WANOP TCP header options and sequence number adjustment.

What are the basic requirements of acceleration?

Acceleration requires a Citrix SD-WAN WANOP device at both ends of the connection, the connection must use the TCP protocol, and all packets for the connection must pass through both the Citrix SD-WAN WANOP devices.

CIFS and MAPI

March 12, 2021

What pre-requisites are required to before configuring MAPI and Signed SMB on a Citrix SD-WAN WANOP appliance?

You must satisfy the following conditions before you configure MAPI and Signed SMB on a Citrix SD-WAN WANOP appliance:

- The Secure Peer option should be set to True on client as well as server side appliance.
- A delegate user must be added to the data center side appliance and its status should be marked as "Success."

- The data center side appliance must successfully join the domain.
- The DNS IP addressed configured on the server side appliance must be reachable.

For more information, see [Configure a Citrix SD-WAN WANOP appliance to optimize secure Windows traffic](#).

What do I need to configure on domain controller for a delegate user?

You must create a user on the domain controller before configuring delegation for the user on a Citrix SD-WAN WANOP appliance.

Do I need to configure anything on DNS server?

Yes. On the DNS server, you must configure forward and reverse lookups for all IP address of the domain controllers.

What do I need to verify before making the Citrix SD-WAN WANOP appliance to join the domain?

Before making the appliance to join the domain, verify the following:

- IP addresses configured to primary or secondary DNS servers should be reachable.
- Domain should be reachable.
- Resolved domain IP addresses should be reachable.
- Optionally, the status of the Pre Domain Join Check utility should pass.

How can I verify if the Citrix SD-WAN WANOP appliance is ready to add a user as a delegate user?

You can verify the user by using the Check delegate user utility on the Windows domain page. If the status for all the parameters does not have any error messages, the appliance is ready to add the user as a delegate user.

If the utility displays any failures, you must address these before you add a user as a delegate user. You can refer to the log to understand the test results.

Are there any requirements for hostname and hostname length of the server side Citrix SD-WAN WANOP appliance?

On the server side Citrix SD-WAN WANOP appliance, make sure that the host name is unique within the network. Additionally, the length of the host name must not be more than 15 characters.

Can I configure one-way trust in the domain?

No. the client and the server must be the members of a domain that has two-way trust with the domain of the server side Citrix SD-WAN WANOP appliance. The appliance does not support one-way trust.

Can I use Macintosh Outlook client and get acceleration benefits of the Citrix SD-WAN WANOP appliance?

No. Macintosh Outlook does not use MAPI as the communication protocol. Therefore, you cannot use Macintosh Outlook in this setup.

Do I need make the branch side Citrix SD-WAN WANOP appliance join the domain for accelerating encrypted MAPI?

No. You do not need to make the branch side Citrix SD-WAN WANOP appliance join the domain for accelerating encrypted MAPI.

Can I configure a Citrix SD-WAN WANOP 2000 appliance with Windows-Server on a data center side for encrypted MAPI?

Yes. You can configure a Citrix SD-WAN WANOP 2000 appliance with Windows-Server on a data center side for encrypted MAPI.

When I make a Citrix SD-WAN WANOP appliance to join a domain and an NTP server configured with a different time zone exists on the network, does the appliance synchronize time with the domain controller or the NTP server?

When you make the Citrix SD-WAN WANOP appliance join a domain, the appliance always synchronized its time with the domain controller and not the NTP server.

On the Citrix SD-WAN WANOP appliance, what is the default duration to clear the black listed connection?

By default, the black listed connections are cleared in 900 seconds.

Which Outlook authentication mechanisms are supported on a Citrix SD-WAN WANOP appliance?

Starting with release 6.2.4, the appliance supports Negotiate (default) and NTLM v2 Outlook authentication, but Kerberos authentication is not supported. However, release 6.2.3 and earlier releases support only Negotiate Outlook authentication.

Does Citrix SD-WAN WANOP support Outlook Anywhere, RPC over HTTPS?

Yes, starting with release 7.3.

Compression

March 12, 2021

What is the benefit of Citrix SD-WAN WANOP Compression?

While the basic mechanism of compression is to make data streams smaller, the benefit of this is to make things faster. A smaller file (or a smaller transaction) takes less time to transfer. Size doesn't matter: the point of compression is speed.

How is compression benefit measured?

There are two ways of measuring compression benefit: time and compression ratio. The two are related when the WAN link is the dominant bottleneck. Because the Citrix SD-WAN WANOP compressor is very fast, compressing data in real time, a file that compresses by 5:1 transfers in one-fifth the time. This holds true until a secondary bottleneck is encountered. For example, if the client is too slow to handle a full-speed transfer, a 5:1 compression ratio delivers less than a 5:1 speedup.

How does compression work?

The compression engine retains data previously transferred over the link, with the more recent data retained in memory and a much larger amount on disk. When a string that was transferred before is encountered again, it is replaced with a reference to the previous copy. This reference is sent over the WAN instead of the actual string, and the appliance on the other end looks up the reference and copies it into the output stream.

What is the maximum achievable compression ratio?

The maximum achievable compression ratio on a Citrix SD-WAN WANOP appliance is approximately 10,000:1.

What is the expected compression ratio?

Overall compression ratio is the average of all attempts to compress the data streams on the link. Some compresses better than others, and some never compress at all. The appliance uses service classes to prevent sending obviously uncompressible streams to the compressor. The effect of compression on different types of data varies as follows:

One-time compressed or encrypted data –streams that are never to be seen again and have already been compressed or encrypted, such as encrypted SSH tunnels and real-time video camera monitoring –are not compress, since their data streams are never the same twice.

Compressed binary data or encrypted data that is seen more than once compresses extremely well on the second and subsequent transfers, with compression ratios in the range of hundreds to thousands to one on these later transfers. On the first transfer, they do not compress. The average compression ratio for such data is dependent on how frequently data is seen more than once. While individual transfers sometimes show compression ratios over 1,000:1, averages for the compressed binary data on the link averages between 1.5:1 and 5:1 on most links, with averages over 10:1 on some links, depending on the nature of the traffic.

Text streams and uncompressed/unencrypted binary data compress even on the first pass. Text streams compress well because even unrelated texts have many substrings in common. This is true of documents, source code, HTML pages, and so on. First-pass compression on the order of 1.5:1 to 4:1

are common. On the second and subsequent passes, they compress almost as well as compressed binary data (100:1 or more). Uncompressed binary data is variable, but often compresses better than text. Examples of uncompressed binary data include CD images, executable files, and uncompressed image, audio, and video formats. On the second and subsequent passes, they compress about as well as compressed binary data.

Citrix Virtual Apps and Desktops data compresses especially well with file transfers, printer output, and video, provided that the same data streams have traversed the link before. Because of protocol overhead, peak compression is approximately 40:1, and average compression is likely to be in the neighborhood of 3:1. Interactive data streams, such as screen updates), give compression results on the order of 2:1.

What is the difference between caching and compression?

Caching saves entire, named objects on the client-side appliance. The name may be a path and file-name in the case of Filesystem caching, or a URL in the case of Web caching. If you transfer an identical object with a different name, the cache provides no benefit. If you transfer an object with the same name as a cached object, but with slight differences in content, the cache provides no benefit. If the object can be served from the cache, it is not fetched from the server.

Compression, on the other hand, has no concept of object names, and provided benefit whenever a string in the transfer matches one that is already in compression history. This means that if you download a file, change 1% of its content, and upload the new file, you might achieve 99:1 compression on the upload. If you download a file and then upload it to a different directory on the remote site, you might achieve a high compression ratio as well. Compression does not require file locking and does not suffer from “staleness.” The object is always fetched from the server and is thus always byte-for-byte correct.

RPC over HTTPS

March 12, 2021

Is it mandatory to create a service class to accelerate RPC over HTTPS connections?

Creating a new service class is an optional task. You can use an existing HTTPS service class. However, to create reports specifically for RPC over HTTPS connections, you must create a new service class and bind the SSL profile to it. If you do not want to create a service class for RPC over HTTPS connections, you can bind the SSL profile you have created to the Web (Private-Secure) service class.

I have not created any service class for the RPC over HTTPS applications. How will this affect the reporting of the RPC over HTTPS connections?

When you upgrade the appliance to release 7.3, the RPC over HTTPS applications that are created do not belong to any service class. As a result, all RPC over HTTPS connections are listed as the TCP Other connections in the reports. If you want to classify these connections as RPC over HTTPS connections, you must create a service class for these applications.

Is there a default service class for RPC over HTTPS on the appliance?

No. The appliance has only default applications, and not default service classes. You must create the service class for an application.

Does the appliance provide any SSL compression benefits to the RPC over HTTPS connections?

No. The appliance does not provide any SSL compression benefits to the RPC over HTTP connections. Compression benefits are available only for encryption and decryption of HTTPS traffic.

Similar to MAPI, does the appliance optimize latency for RPC over HTTPS connections?

No. The appliance does not optimize latency for RPC over HTTPS.

Is MAPI over HTTP different from RPC over HTTPS?

Yes. MAPI over HTTP is a new protocol supported on Microsoft Exchange Server 2013 SP1 or later.

What is the difference between RPC over HTTPS settings on client-side and server-side Citrix SD-WAN WANOP appliances?

Except for creating a service class and adding RPC over HTTPS applications to it, you do not need any additional configuration on a client-side Citrix SD-WAN WANOP appliance.

What happens if I configure the SSL profile in transparent proxy mode?

Some Exchange servers require TLS session ticket support. To accelerate connections to these servers, you need to create an SSL profile with split proxy, because transparent proxy mode does not support TLS session tickets.

If a load balancing setup is used for the Microsoft Exchange Server, which destination IP address should I add to the filter rule when creating an RPC over HTTPS service class?

If you are using a load balancing appliance, add its virtual IP (VIP) address to the filter rule when creating an RPC over HTTP service class.

How can I differentiate between the MAP and RPC over HTTPS traffic in the Outlook (MAPI) page?

You can differentiate the traffic based on applications shown on the Outlook (MAPI) page. For example, MAPI and RPC over HTTPS are used for the following applications:

- **MAPI:** MAPI and eMAPI
- **RPC over HTTPS:** HTTP MAPI, HTTP eMAPI, HTTPS MAPI, and HTTPS eMAPI

SCPS

March 12, 2021

What is SCPS protocol?

Space Communications Protocol Standard (SCPS) protocol is a variant of the TCP protocol.

What is the use of SCPS protocol?

SCPS protocol is used in satellite communication and similar applications.

Is SCPS protocol supported on a Citrix SD-WAN WANOP appliance?

Yes. The Citrix SD-WAN WANOP appliance supports SCPS protocol and accelerate data transferred using this protocol.

Can I use an SCPS-enabled appliance with a non-SCPS-enabled appliance?

Yes. If you must mix SCPS-enabled appliances with non-SCPS-enabled appliances, deploy them in such a way that mismatches do not occur. You can either use IP-based service class rules or arrange the deployment so that each path has matching appliances.

What happens if I use an SCPS-enabled appliance at one end non-SCPS-enabled appliance on the other end of the link?

If the appliance on one end of the connection has SCPS enabled and one does not, retransmission performance suffers. This condition also causes an “SCPS Mode Mismatch” alert.

What is the difference between the behavior of a SCPS-enabled appliance and default appliance?

The main difference between a SCPS-enabled and the default appliance behavior is that SCPS-style “selective negative acknowledgements”(SNACKs) are used instead of standard selective acknowledgements (SACKs).

Secure peering

March 12, 2021

Which Citrix SD-WAN WANOP features required secure peering?

You need to establish secure peering between Citrix SD-WAN WANOP appliances at two ends of the link when you intend to use any of the following features:

- SSL compression

- Signed CIFS support
- Encrypted MAPI support

Do I need to consider anything before configuring a secure tunnel?

Yes. You must order and receive a crypto license before you can configure a secure tunnel between the Citrix SD-WAN WANOP appliances at to ends of the link.

What happens when you enable secure peering on an appliance at one end of the link?

When you enable secure peering on a Citrix SD-WAN WANOP appliance at one end of the link, the other appliance detects it and attempts to open an SSL signaling tunnel. If the two appliances successfully authenticate each other over this tunnel, the appliances have a secure peering relationship. All accelerated connections between the two appliances are encrypted, and compression is enabled.

What happens when I do not enable secure peering on the partner appliance?

When an appliance has secure peering enabled, connections with a partner for which it does not have a secure peer relationship are not encrypted or compressed, though TCP flow-control acceleration is still available. Compression is disabled to ensure that data stored in compression history from secured partners cannot be shared with unsecured partners.

Why do I need a keystore password?

You need a keystore password to access the security parameters. This password is different from the administrator's password and allows security administration to be separated from other tasks. If the keystore password is reset, all existing encrypted data and private keys are lost.

To protect data even if the appliance is stolen, the keystore password must be reentered every time the appliance is restarted. Until this is done, secure peering and compression are disabled.

Does the Citrix SD-WAN WANOP appliance I received from Citrix contain keys and certificate to set up secure tunnel?

No. Citrix SD-WAN WANOP products are shipped without the required keys and certificates for the SSL signaling tunnel. You must generate them yourself.

SSL Acceleration

March 12, 2021

Does acceleration use a tunnel?

No, acceleration is transparent, using the same IP addresses and port numbers as the original connection. This allows your current monitoring methods to continue to work normally.

How does acceleration change the packet stream?

With non-compressed connections, acceleration adds options to the packet's TCP header, but leaves the packet payload intact. These options allow the Citrix SD-WAN WANOP devices at each end of the connection to communicate with each other. In addition, the TCP sequence number is adjusted to prevent routing issues or appliance failure from mixing accelerated packets and non-accelerated packets in the same connection.

With compressed connections, the payload is compressed, of course, and the output of the compressor is accumulated into full-sized packets. The upshot is that, for example, 3:1 compression results in one-third as many packets being transmitted, rather than the same number of packets, each reduced to one-third size. Compression also uses Citrix SD-WAN WANOP TCP header options and sequence number adjustment.

What are the basic requirements of acceleration?

Acceleration requires a Citrix SD-WAN WANOP device at both ends of the connection, the connection must use the TCP protocol, and all packets for the connection must pass through both Citrix SD-WAN WANOP devices.

Citrix SD-WAN WANOP plug-in

March 12, 2021

What methods can I use to install the Citrix SD-WAN WANOP plug-in on my computer?

You can use any of the following methods to install the Citrix SD-WAN WANOP plug-in on your computer:

- Standalone installation: Run the Microsoft Installer (msi) file.
- Silent installation: Run the following command:

```
*\> msixexec.exe /i path\\CitrixSD-WANWANOPPluginReleasex64-\\<
Release\\_Nunmer\> /qn*
```
- Remote installation: Install the Citrix SD-WAN WANOP plug-in remotely from Citrix Receiver. This installation is done using the merchandising server.

Can I customize the Citrix SD-WAN WANOP plug-in installer?

Yes. You can customize the signaling IP address and disc based compression (DBC) size with the msi file for the Citrix SD-WAN WANOP plug-in.

What are the minimum hardware requirements for installing the Citrix SD-WAN WANOP plug-in?

For the Citrix SD-WAN WANOP plug-in, your computer should meet the following requirements:

- Pentium 4 class CPU
- Minimum 4 GB of RAM
- Minimum 2 GB for free hard disk space

On which operating systems can I install the Citrix SD-WAN WANOP plug-in?

You can install the Citrix SD-WAN WANOP plug-in on the following operating systems:

Operating System	Edition	Version
Windows XP	Home, Professional	32-bits
Windows Vista	Home Basic, Home Premium, Business, Enterprise, Ultimate	32-bits
Windows 7	Home Basic, Home Premium, Business, Enterprise, Ultimate	32-bits, 64-bits
Windows 8	Professional, Enterprise	32-bits, 64-bits
Windows 10	Professional, Enterprise	32-bits, 64-bits

What precautions should I take before installing the Citrix SD-WAN WANOP plug-in?

Before you install the Citrix SD-WAN WANOP plug-in on your computer, take the following precautions:

- Depending on your operating system version, download either 32-bit or 64-bit Citrix SD-WAN WANOP installer version.
- You cannot install the Citrix SD-WAN WANOP plug-in on a compressed drive or folder.
- Make sure that the computer has sufficient free disk space.
- You cannot downgrade the Citrix SD-WAN WANOP plug-in release. If you want to use an earlier Citrix SD-WAN WANOP release, you must uninstall the current release and then install an earlier release.

Which Citrix SD-WAN WANOP appliances support the Citrix SD-WAN WANOP plug-in?

The following Citrix SD-WAN WANOP appliances support the Citrix SD-WAN WANOP plug-in:

- SD-WAN WANOP 2000
- SD-WAN WANOP 2000 appliance with Windows Server
- SD-WAN WANOP 3000

- SD-WAN WANOP 4000
- SD-WAN WANOP 5000

Which Citrix SD-WAN WANOP appliances do not support the Citrix SD-WAN WANOP plug-in?

The following Citrix SD-WAN WANOP appliances do not support the Citrix SD-WAN WANOP plug-in:

- SD-WAN WANOP 400
- SD-WAN WANOP 700
- SD-WAN WANOP 800
- SD-WAN WANOP 1000 with Windows Server

Do I need to install a Concurrent (CCU) license on Citrix SD-WAN WANOP 2000, 3000, and VPX appliances to use the Citrix SD-WAN WANOP plug-in?

Yes. You must install a CCU license on Citrix SD-WAN WANOP 2000, 3000, and VPX appliances to use the Citrix SD-WAN WANOP plug-in.

Do I need install a CCU license on Citrix SD-WAN WANOP 4000 and 5000 appliances to use the Citrix SD-WAN WANOP plug-in?

No. You do not need to install a CCU license on Citrix SD-WAN WANOP 4000 and 5000 appliances to use the Citrix SD-WAN WANOP plug-in. The appliance base license is sufficient for the Citrix SD-WAN WANOP plug-in to connect to these appliances.

What are the Citrix recommendations for accelerating subnets?

Citrix recommends following for accelerating subnets:

- Never use ALL/ALL for acceleration configuration. Specify the subnets on the basis of the requirements.
- Do not configure acceleration for the Citrix Gateway VIP address.

Is the Citrix SD-WAN WANOP plug-in supported on Windows thin clients?

No. The Citrix SD-WAN WANOP plug-in is not supported on Windows thin clients.

Which Citrix Receiver and Citrix Gateway releases are supported with the Citrix SD-WAN WANOP plug-in?

The Citrix SD-WAN WANOP plug-in supports Citrix Receiver 4.1 and Citrix Gateway 10.5 releases.

Which Citrix SD-WAN WANOP features are not supported with the Citrix SD-WAN WANOP plug-in?

The Citrix SD-WAN WANOP plug-in does not support the following Citrix SD-WAN WANOP features:

- Video Caching

- Traffic Shaping
- IPv6

Do I need to configure acceleration rules on a Citrix SD-WAN WANOP 4000 or 5000 appliance for the Citrix SD-WAN WANOP plug-in to work with it?

Yes. You must configure acceleration rules on a Citrix SD-WAN WANOP 4000 or 5000 appliance for the Citrix SD-WAN WANOP plug-in to work with it.

What is the significance of signaling-channel source filtering?

By using signaling-channel source filtering, you can either allow or deny a specific subnet or IP address the ability to connect to the appliance and fetch acceleration rules. The denied source subnet cannot establish signaling connections and accelerate the traffic.

What is the significance of LAN detection?

When you enable LAN detection, it prevents traffic acceleration when the Citrix SD-WAN WANOP plug-in and appliance are on the same LAN. Local acceleration is not desirable, because applying the bandwidth limit of the appliance to the local connection might reduce the speed of the local traffic.

To accelerate traffic, what is the minimum recommended RTT value between the Citrix SD-WAN WANOP plug-in and appliance?

Citrix recommends that you configure an RTT value that is greater than any RTT (ping time) on the local LAN, but less than the RTT for any remote user. The default value of 20 milliseconds is adequate for most networks.

What conditions should I consider when defining acceleration rules for the Citrix SD-WAN WANOP plug-in?

Consider the following conditions when defining acceleration rules for the Citrix SD-WAN WANOP plug-in:

- Define acceleration rules for all subnets that are local to the appliance. These subnets are the LAN subnets at the site where the appliance is installed.
- If there are any destination IP addresses that are not part of the LAN, add exclude rules for these IP addresses. Make sure that the rules for excluding IP addresses precede the rules for accelerating traffic for subnets. This includes subnets at remote sites with IP addresses that appear local.
- If you have installed the appliance in inline mode with a VPN and it is operating in transparent mode, you can configure the appliance to accelerate all enterprise traffic, not just the traffic originated by or destined to the local site. In this case, the only accelerated connections are between the Citrix SD-WAN WANOP plug-in and VPN. Acceleration of the traffic between the Citrix SD-WAN WANOP plug-in and the VPN is optimal.

Where are the Citrix SD-WAN WANOP plug-in crash and trace files stored on the computer?

The crash and trace files of the Citrix SD-WAN WANOP plug-in are stored in the following folders:

- Crash files: C:/ProgramFiles/Citrix/Citrix SD-WAN WANOP
- Trace files: C:/Users/admin/AppData/Local/Temp

How does the Citrix SD-WAN WANOP plug-in connect to a high availability pair?

The Citrix SD-WAN WANOP plug-in always connects to the same signaling IP address. The signaling IP address is bound to only the primary appliance of the high availability pair, not to the secondary appliance. Therefore, the Citrix SD-WAN WANOP plug-in always connects to the primary appliance of the high availability pair.

Which deployment modes does the Citrix SD-WAN WANOP plug-in support?

The Citrix SD-WAN WANOP plug-in supports the following deployment modes:

- Inline.
- WCCP.
- High Availability.
- Citrix SD-WAN WANOP plug-in with NAT deployment.
- Citrix SD-WAN WANOP plug-in with Citrix SD-WAN WANOP appliance in WCCP mode using ICA proxy.
- Citrix SD-WAN WANOP plug-in with Citrix SD-WAN WANOP 4000 or 5000 appliance. In this deployment, the management port (0/1) is connected to the management network, and the signaling IP address is on a different network.

How do packets flow in transparent and redirector modes?

In transparent mode, the Citrix SD-WAN WANOP appliance does not change the source IP address of the packet. In redirector mode, the Citrix SD-WAN WANOP appliance proxies servers and changes the IP address of the packets.

Note

Citrix recommends transparent mode for the production deployment.

How can I establish a secure tunnel between the Citrix SD-WAN WANOP plug-in and appliance?

To establish a secure tunnel between the Citrix SD-WAN WANOP plug-in and appliance, complete the following procedure:

1. On the Citrix SD-WAN WANOP plug-in user interface, open the **Certificates** tab.
2. Select the **CA Certificate** option.

3. Click **Import** and upload the relevant CA certificate.
4. Select a Certificate Store where you want to store the certificate.
5. Select the **Client Certificate** option.
6. Click **Import**.
7. Select appropriate certificate formats and upload the relevant certificates.
8. Store the certificates in a Certificate Store.
9. If the private key is password protected, enter the password to decrypt the private key.
10. You must upload the same CA certificate and key pair to the appliance to establish a secure tunnel.

How can I verify that a secure tunnel is established?

To verify that a secure tunnel is established, complete the following procedure:

1. The computer where you have installed the Citrix SD-WAN WANOP plug-in, run the following command:

```
*\> telnet localhost 1362*
```

2. On the console, run the following command:

```
*\> showtunnels*
```

Following is sample output of the command. If the output includes the text *secure* in the *Connected Available* section, a secure tunnel has been established. If a secure tunnel is not established, the text reads *cleartext*.

```
1  ```
2  Showtunnels
3  Message Tunnels:
4    Connected Available:
5      172.16.9.100 auto,secure,client,initiator,configured
6      CN: mike.199.130
7
8
9  Connected Available : 1
10 Clients: 1 peers: 0
11 <!--NeedCopy--> ```
```

For more information on Citrix SD-WAN WANOP plug-in, see [Citrix SD-WAN WANOP Plug-in](#).

Traffic shaping

March 12, 2021

What is Citrix SD-WAN WANOP Traffic Shaping?

Citrix SD-WAN WANOP traffic shaping uses a group of policies to set the priority of different link traffic and send traffic onto the link at a rate close to, but no greater than, the link speed. Unlike acceleration, which applies only to TCP/IP traffic, the traffic shaper handles all traffic on the link.

What is the benefit of traffic shaping?

Traffic shaping uses scarce link resources according to the policies you set, so that traffic that is known to be important will receive more bandwidth than traffic that is known to be unimportant.

How does the traffic shaper interact with Citrix Virtual Apps and Desktops traffic?

The Citrix SD-WAN WANOP device parses the Virtual Apps/Virtual Desktops data stream and is aware of the different types of traffic and its priorities, favoring high-priority traffic. It is the only product that can prioritize encrypted ICA streams and provide native support for MultiStream ICA, which divides a user's session into up to four connections with different priorities.

What is weighted fair queuing?

A Citrix SD-WAN WANOP appliance uses weighted fair queuing, which provides a separate queue for each connection. With fair queuing, a too-fast connection can overflow only its own queue. It has no effect on other connections.

What is the difference between weighted and non-weighted fair queuing?

Weighted fair queuing includes the option of giving some traffic a higher priority (weight) than others. Traffic with a weight of two receives twice the bandwidth of traffic with a weight of one. In a Citrix SD-WAN WANOP configuration, the weights are assigned in traffic-shaping policies.

What is a link definition?

A link definition specifies which traffic is associated with the defined link, the maximum bandwidth to allow for traffic received on the link, and the maximum bandwidth for traffic sent over the link. The definition also identifies traffic as inbound or outbound and as WAN-side or LAN-side traffic.

What are the benefits for link definition?

Link definitions enable the appliance to prevent congestion and loss on your WAN links and to perform traffic shaping. The definition also identifies traffic as inbound or outbound and as WAN-side or LAN-side traffic. All traffic flowing through the appliance is compared to your list of link definitions, and the first matching definition identifies the link to which the traffic belongs.

I have not configured any service class with Default Policy. However, the traffic shaping reports displays a large amount of traffic represented by Default Policy. Have I configured something incorrectly?

No. There is no issue with your configuration. Traffic shaping is only applicable to the WAN link. Traffic on the LAN or any other link is represented by Default Policy.

For example, consider a configuration where you create a service class, such as Management_Service_Class, that has the management subnet as the destination IP address and you bind a custom traffic shaping policy to this service class. In this case, when there is no traffic on WAN, you can notice that the management traffic is classified as Management_Service_Class in the service class report. However, in the Traffic Shaping Policy report, entries for Default Policy still exist that you might expect to exist as custom traffic shaping policy.

In the Traffic Shaping Policy report, the appliance does not use customized traffic shaping policy for the Management_Service_Class policy and applies Default Policy. To avoid this confusion, you can clear the All other option or define the LAN type link for the management interface.

Upgrade (OS) Process

March 12, 2021

The new WANOP OS Kernel upgrade is supported from which SD-WAN release?

Citrix SD-WAN release 10.1 and later.

Is the new OS supported on all SD-WAN platforms?

Yes. The OS upgrade is supported on all SD-WAN WANOP (VPX, Physical, Cloud), and Premium/Enterprise edition appliances.

What are the WANOP VPX profiles (RAM/Disk/vCPU) that are supported with release 10.1?

- 6GB RAM, 100GB Disk and 2 vCPUs
- 6GB RAM, 250GB Disk and 2 vCPUs
- 8GB RAM, 500GB Disk and 4 vCPUs
- 16GB RAM, 500GB Disk and 4 vCPUs

What are the key feature differences between WANOP running with release 10.0 or lower versus 10.1?

Feature	10.0 or earlier	10.1 or later	Comments
Video Caching support on WANOP	supported	Not supported	none
Minimum RAM requirement for WANOP VPX	4GB RAM	6GB RAM	none
WANOP VPX deployment wizard	supported	Not supported	none

Feature	10.0 or earlier	10.1 or later	Comments
Primary / apA Adapter Management IP address for WANOP VPX	DHCP is disabled by default	DHCP is enabled by default	none
Upgrade support on existing standalone WANOP VPX on Citrix Hypervisor	supported	not supported. Fresh SD-WAN 10.1 XVA image should be imported	none
Upgrade support on Physical WANOP platform that have Citrix Hypervisor 6.0 Hypervisor version (Platforms that are shipped with 7.2.2 or earlier factory base image version would have Citrix Hypervisor 6.0 version) release 10.1	supported	You have to upgrade Citrix Hypervisor to 6.5 version (using WANOP Citrix Hypervisor 6.5 upgrade bundle) and then perform WANOP 10.1 upgrade	Clicking on “Configuration” GUI, would display Citrix Hypervisor hypervisor version

Upgrade of WANOP VPX running on standalone Citrix Hypervisor (with WO build 10.0 or earlier) to 10.1 version is that supported, If not, why?

This upgrade is not supported because of PV to HVM conversion. You have to provision a fresh SD-WAN release on 10.1 Citrix Hypervisor WANOP VPX using the XVA image.

Upgrade of WANOP VPX running on standalone ESXi / Hyper-V (with WO build 10.0 or earlier) to 10.1 version is that supported, If not, why?

This upgrade is supported. Before upgrade, please be aware of the new RAM resource requirement changes.

Upgrade of WANOP on physical appliance (with WANOP build 10.0 or earlier) to 10.1 version is that supported, If not, why?

This upgrade is supported. Prerequisite for this upgrade is to have the hosting Citrix Hypervisor Hypervisor (on physical SD-WAN appliance) to have Citrix Hypervisor version 6.2 / 6.5 or higher version. This can be verified by using the **Configuration** tab.

Dashboard

Monitoring

Configuration

Downloads

Notifications (2)

+ Appliance Settings

+ Optimization Rules

+ Secure Acceleration

+ Diagnostics

+ Maintenance

Configuration Overview

Current Versions

Management Service	Version: 11.1, Build: 51.143
XenServer	Version: 6.5, Build: 90233c
Supplemental Pack	Version: 6.5.0-3.10.0-2-2.0.0-1020-1020
Hotfixes	XS65E001,XS65ESP1002,XS65E015,XS65ESP1005,XS65E008,XS65ESP1020,XS65E013,XS65E014,XS65ESP1023,XS65ESP1008,XS65ESP1012,XS65E00
NetScaler SD-WAN WO	Version: 10.1.0, Build: 147

Hypervisor Information

Uptime	29 minutes
Edition	Citrix XenServer
Version	6.5
iSCSI IQN	iqn.2018-07.com.example:3cd59988
Kernel Version	3.10.0+2


System Information

Platform	800
Product	Citrix NetScaler SD-WAN
Build	11.1: Build 51.143, Date: May 30 2018, 01:37:04
IP Address	10.106.133.156
System ID	450150
Serial Number	FT29C2EACM
System Time	Fri Jul 27 15:02:01 IST 2018

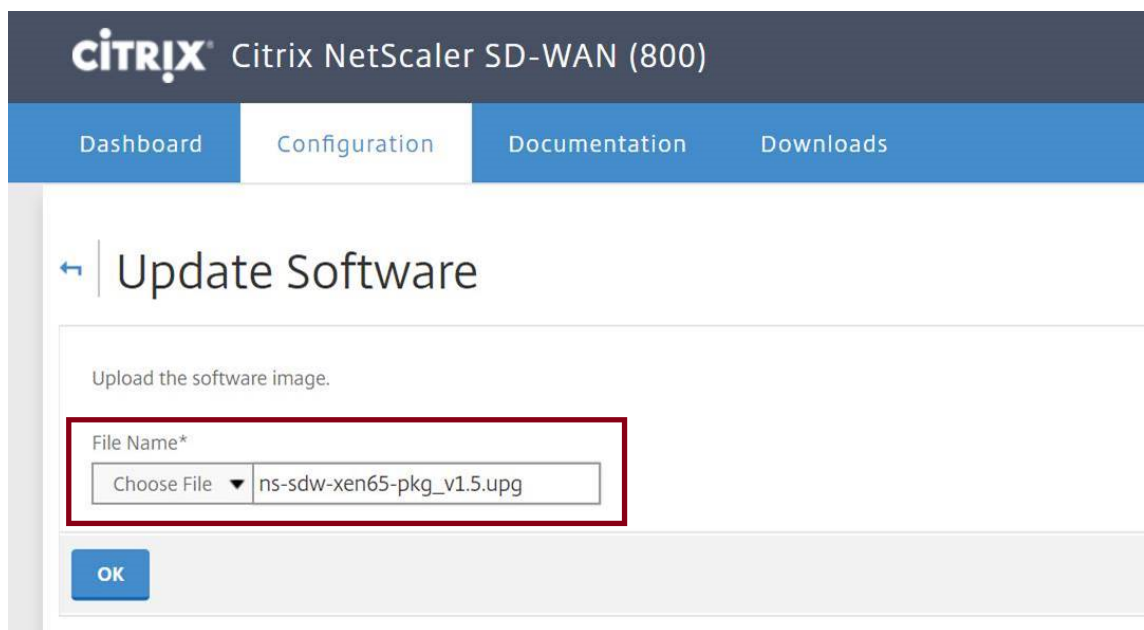
If the physical WANOP appliance is not running with Citrix Hypervisor 6.2 / 6.5 or higher, what does the user have to perform?

You have to upgrade Citrix Hypervisor, before upgrading SD-WAN WO version. For example, in this below use case, let us consider planning to upgrade SD-WAN 800 WANOP platform running with 7.2.2 (that has Citrix Hypervisor 6.0 version).

1. While upgrading this appliance to SD-WAN 10.1 release, the following error message would occur.



2. Upgrade Citrix Hypervisor to 6.5, using “ns-sdw-xen65-pkg_v1.5.upg”(this can be downloaded from Citrix Download website).



CITRIX Citrix NetScaler SD-WAN (800)

Dashboard Configuration Documentation Downloads

Update Software


Upload the software image.

File Name*

Choose File ▼ ns-sdw-xen65-pkg_v1.5.upg

OK

3. If SD-WAN WO does not have 9.0 or later version, then upgrade to Citrix Hypervisor 6.5 would not happen. The below error message would appear.



4. Let us assume, the user have upgraded the WO version to 10.0.2 now.

Citrix NetScaler SD-WAN 800 Series-WO

10.0.2.37.686956 (Production)Logout

DashboardMonitoringConfigurationDownloadsNotifications (3)

Appliance Settings

Optimization Rules

Video Caching

Secure Acceleration

Diagnostics

Maintenance

Configuration Overview

Current Versions

Management Service	Version: 11.1, Build: 51.143
XenServer	Version: 6.0, Build: 50762p
Supplemental Pack	Version: 2.0.0-1023
Hotfixes	XS60E055,XS60E001,XS60E045,XS60E058,XS60E014,XS60E050,XS60E047,XS60E035,XS60E040,XS60E024,XS60E052,XS60E034,XS60E020,XS60E010,XS60E009,XS60E008,XS60E007,XS60E006,XS60E005,XS60E004,XS60E003,XS60E002,XS60E001
NetScaler SD-WAN WO	Version: 10.0.2, Build: 37

Hypervisor Information

Uptime	17 hours 24 minutes
Edition	Citrix XenServer
Version	6.0
iSCSI IQN	iqn.2018-07.com.example:3cd59988
Kernel Version	2.6.32-12-0.7.1.xs6.0.0.533.170664xen

System Information

Platform	800
Product	Citrix NetScaler SD-WAN
Build	11.1: Build 51.143, Date: May 30 2018, 01:37:04
IP Address	10.106.133.156
System ID	450150
Serial Number	FT29C2EACM

5. Now, upgrade Citrix Hypervisor to 6.5, using “ns-sdw-xen65-pkg_v1.5.upg”.

Update Software

Upload the software image.

File Name*

Choose File

ns-sdw-xen65-pkg_v1.5.upg

OK

Upgrade in progress...

1/1

Upgrading XEN...

Time remaining 20 minutes

✓

Upgrade successfully completed.

Ok

Citrix NetScaler SD-WAN 800 Series-WO

10.0.2.37.686956 (Production)

Logout

CITRIX

Dashboard

Monitoring

Configuration

Downloads

Notifications (2)

+ Appliance Settings

+ Optimization Rules

+ Video Caching

+ Secure Acceleration

+ Diagnostics

+ Maintenance

Configuration Overview

Current Versions

Management Service

Version: 11.1, Build: 51.143

XenServer

Version: 6.5, Build: 90233c

Supplemental Pack

Version: 6.5.0-3.10.0-2-2.0.0-1020-1020

Hotfixes

XS65E001,XS65ESP1002,XS65E015,XS65ESP1005,XS65E008,XS65ESP1020,XS65E013,XS65E014,XS65ESP1023,XS65ESP1008,XS65ESP1012,XS65

NetScaler SD-WAN WO

Version: 10.0.2, Build: 37

Hypervisor Information

System Information

Uptime

5 minutes

Edition

Citrix XenServer

Version

6.5

iSCSI IQN

iqn.2018-07.com.example:3cd59988

Kernel Version

3.10.0+2

Platform

800

Product

Citrix NetScaler SD-WAN

Build

11.1: Build 51.143, Date: May 30 2018, 01:37:04

IP Address

10.106.133.156

System ID

450150

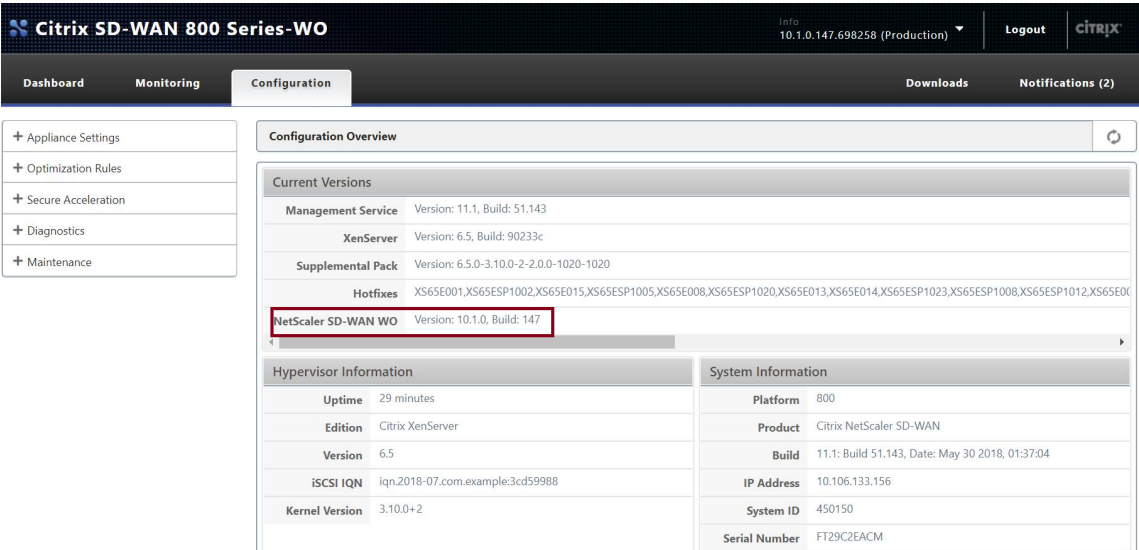
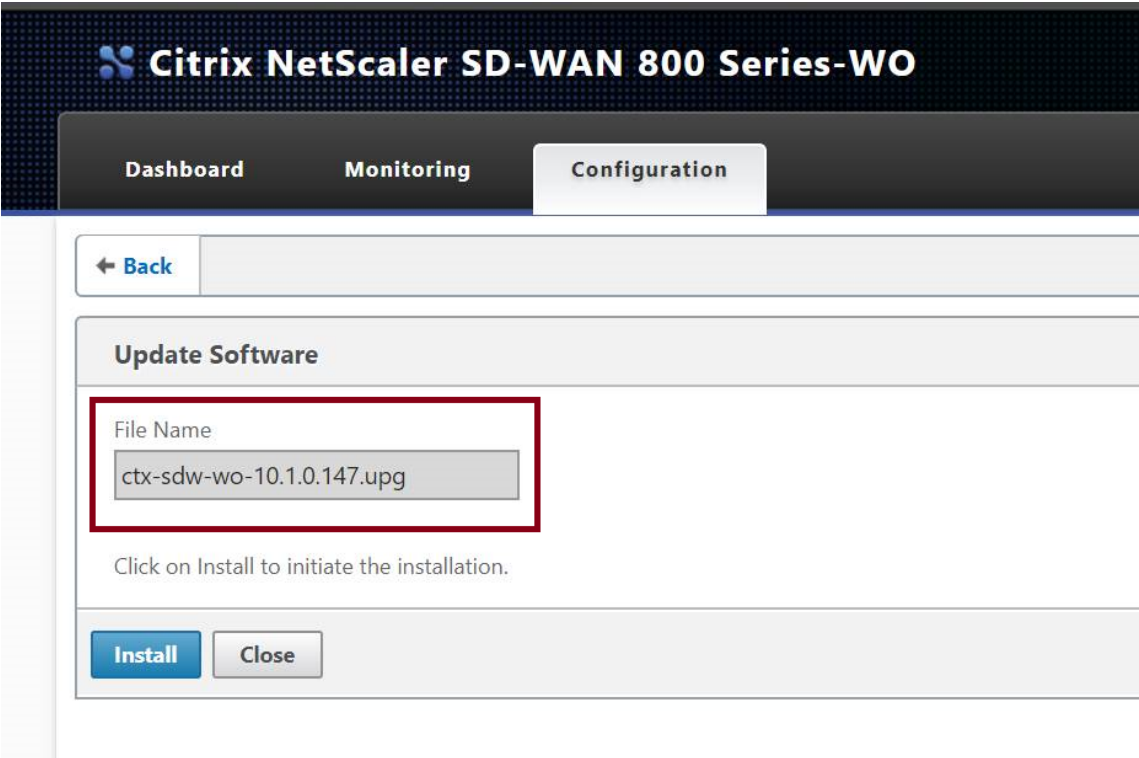
Serial Number

FT29C2EACM

System Time

Fri Jul 27 14:38:00 IST 2018

6. Now, upgrade SD-WAN to 10.1 release.



Client to Server ICMP Ping is working fine, but TCP traffic is not going through the WANOP VPX appliance (disabling WANOP traffic processing works fine)?

Check the firewall settings on the Client, Server and Router.

When WANOP VPX or Client/Server are hosted as VM, make sure that checksum is disabled on the end hosts VM.

```
1 Example Linux Commands:  
2     ethtool -K eth0 tx off  
3     ethtool -K eth0 rx off
```

```
4 ethtool --offload eth0 tx off
5 ethtool --offload eth0 rx off
```

Enable the “Checksum.SendForceSW” parameter on both WO VPXs should be “ON”.

```
1 Example:
2 Checksum.SendForceSW on
```

Is there any change to SDWAN SE/EE/WO Appliance upgrade process due to new WO OS Kernel?

No.

Video caching

March 12, 2021

How is video caching different from Disk Based Compression?

With caching, a local copy of the cached object is served by the local appliance, without downloading it again from the remote server. Caching does not require an appliance on both ends of the link, just on the local end. With compression, a remote copy of the object is served by the remote server. The remote (server-side) appliance compresses it, reducing its size, and therefore, increasing its transmission speed, and the local (client-side) appliance decompresses it.

Compression works on both modified and unmodified objects. If a file changes by 1% on the server, the next transfer achieves up to 99:1 compression.

Caching works only on unmodified objects. If a file changes by 1% on the server, the new version must be downloaded in its entirety. Caching and compression are complementary technologies, because anything that is not cached, is compressed, achieving the benefits of both.

Can I partition the appliance’s total memory between the video cache and other Citrix SD-WAN WANOP features?

No. Cache partition and memory required are not configurable.

What are the supported video container formats?

Video caching is independent of codec format and supports all major container formats.

Can I activate caching for internal and external enterprise videos on my own sites?

Yes. If access to these videos is through HTTP, you can configure these sites for caching.

Can I configure the maximum size for a cached object?

Yes. An object larger than the limit that you configure is not be cached. To set this limit, navigate to **Configuration > Optimization Rules > Video Caching** and select the value from the available limits.

How does video caching improve the user experience?

Caching improves the user experience for videos that are viewed more than once, especially on slower links. The first viewer of a given video stream does not benefit from the video caching feature, but subsequent views are delivered at the LAN speed from the Citrix SD-WAN WANOP appliance, with the additional benefit of reduced WAN usage.

In addition, if a second user requests the same video while it is still being streamed for the first user, the second user will receive the cached copy.

Unlike normal Citrix SD-WAN WANOP TCP operation, where the appliance preserves the original source and destination IP addresses, the appliance replaces the client's source address with IP address assigned to the accelerated bridge, so all HTTP traffic passing through the appliance appears to originate from the appliance itself.

Which Citrix SD-WAN WANOP appliances support Video Caching?

The following appliances support the video caching feature:

- SD-WAN WANOP 800 appliance with all bandwidth license models.
- SD-WAN WANOP 1000 appliance with Windows Server, with all bandwidth license models.
- SD-WAN WANOP 2000 appliance with all the bandwidth license models.
- SD-WAN WANOP 2000 appliance with Windows Server, with all bandwidth license models.
- SD-WAN WANOP 3000 appliance with all the bandwidth license models.

For video caching, which deployment modes are the supported on a Citrix SD-WAN WANOP appliance?

- Supported deployment - Inline Virtual Inline, VLAN, and WCCP
- Not supported features - Citrix SD-WAN WANOP high availability, Group Modes, and Daisy Chaining

Which file extensions are supported for video caching?

The video file name must have one of the following extensions: .3gp, .avi, .dat, .divx, .dvx, .dv-avi, .flv, .fmv, .h264, .hdmov, .m15, .m1v, .m21, .m2a, .m2v, .m4e, .m4v, .m75, .moov, .mov, .movie, .mp21, .mp2v, .mp4, .mp4v, .mpe, .mpeg, .mpeg4, .mpg, .mpg2, .mpv, .mts, .ogg, .ogv, .qt, .qtm, .ra, .rm, .ram, .rmd, .rms, .rmvb, .rp, .rv, .swf, .ts, .vfw, .vob, .webm, .wm, .wma, .wmv, and .wtv.

Can I enable the video caching feature on an unsupported Citrix SD-WAN WANOP platform?

No. The Video Caching feature cannot be used on unsupported platforms.

What are the minimum configuration and other prerequisites for enabling the video caching feature?

To enable the video caching feature, you must:

- Assign a valid IP address and gateway to the apA interface and, if present, to the apB interface.
- On the appliance, configure a valid DNS server that can resolve to www.citrix.com.
- Have at least one application in the Selected Video Caching Applications list.
- Check the Citrix SD-WAN WANOP GUI alerts/notification of existing configuration alerts.

Can the Citrix SD-WAN WANOP plug-in use the Video Caching feature?

No. You cannot use the Video Caching feature with Citrix SD-WAN WANOP plug-in.

What are the supported browsers and devices?

Video caching supports the Internet Explorer, Firefox, and Chrome browsers. Videos can be viewed on Windows 7 or 8, Apple iPad, and Android iOS devices.

Does the Citrix SD-WAN WANOP appliance support video caching for all video websites?

No. The video website is available and added from the Supported Application list on the Video Caching configuration page. By default the supported applications include YouTube, Vimeo, Youku, Dailymotion and Metacafe. You can add other websites by specifying their IP addresses, if they do not use caching avoidance mechanisms, such as adding random characters to URLs.

Is the SNMP monitoring supported for video caching?

Yes. You can use SNMP MIBs to monitor video caching specific tasks.

Is video caching supported for non-HTTP traffic?

No. Video Caching is not supported for non-HTTP traffic, such as HTTPs, RTSP, and RTMP.

Can I use video caching with HTTP traffic sent to a port other than port 80?

Yes. For video caching, you can add customized ports to the appliance. To add customized ports for video caching, navigate to the **Configuration > Optimization Rules > Video Caching** page and click the **Global Settings** link on the **Settings** tab.

Can Citrix SD-WAN WANOP compression (using an HTTP Service Class policy) be used with Video Caching?

Yes. When the cached objects are present in both Citrix SD-WAN WANOP compression history and the video cache, the content is served from the cache on a cache hit, and fetched from the server (and compressed) on a cache miss.

Does an existing HTTP Application which requires IP address configuration when there is a transparent proxy, require any changes?

Yes. Citrix SD-WAN WANOP performs HTTP transparent proxying, in which it replaces the Source IP address of the packet. Therefore, if the existing HTTP application has certain policies (such as to block certain IP addresses or Proxy mechanisms), those policies have to be changed.

What are the system memory and connection limits for the HTTP proxy connection?

To determine the limits, check the graphs and statistics on the Video Caching Debug page (support.html). Additionally, verify that the Videocaching.cmd stats info command shows the following information.

	SD-WAN WANOP 800	SD-WAN 1000 with Widows Server	SD-WAN 2000 with Widows Server	SD-WAN 2000	SD-WAN 3000
Disk	25 GB	25 GB	50 GB	50 GB	99 GB
RAM	375 MB	375 MB	700 MB	700 MB	1024 MB
Total HTTP Connections limit	1000	1000	1500	1500	3000
Maximum HTTP Write limit	200	200	300	300	600

After the above HTTP connection limits are reached, new connections are bypassed.

Note

Make sure that you do not change the above configuration.

Does the Monitoring page for video caching include only video traffic?

Yes. Non-video HTTP traffic (even though it is intercepted by the proxy), is not included in the video caching GUI statistics.

Do I need to configure apA as well as apB interfaces with a valid IP address on a Citrix SD-WAN WANOP appliance?

No. You do not need to assign a valid IP address to both the interfaces. HTTP packets received from the apA interface are proxied with the apA IP address, and HTTP packets received from the apB interface are proxied with the apB IP address.

If you do not configure an IP address for an interface, the HTTP packets received on that interface do not get the caching benefit.

What is the minimum and maximum limit for the size of a video file that can be cached?

- Minimum: 100 KB
- Maximum: 300 MB
- Default: 100 MB

How is the video caching disk cleared?

Cached objects are cleared as specified by the Least Recently Used algorithm.

What happens when I upgrade the Citrix SD-WAN WANOP appliance from release 6.x to 7.y and video caching is enabled?

The existing Citrix SD-WAN WANOP DBC history is lost and a separate partition for video caching is created.

What happens when I downgrade the Citrix SD-WAN WANOP appliance from release 7.y to 6.x and video caching is enabled?

Citrix SD-WAN WANOP DBC and Video Caching history is preserved. However, the video caching feature is not available with release 6.x.

What happens when I upgrade the Citrix SD-WAN WANOP appliance from release 7.x to 7.y and video caching is enabled?

The Citrix SD-WAN WANOP DBC and video caching history is preserved.

I have a single network in branch office that shares a management as well as data traffic. How should I configure video caching in this network?

If you have single network for management and data traffic, Citrix recommends that you add the primary IP address to the LAN side of the accelerated bridge port.

What is the maximum number of prepopulation tasks I can run at the same time?

One. If you attempt to start multiple prepopulation tasks at the same time, the appliance builds a queue of tasks on a first in first out basis.

What is the maximum number of videos sources I can configure on the appliance?

100

What is the maximum number of prepopulation entries I can add to the appliance?

50

What is the maximum number of video files be downloaded and cached from a directory listed folder?

300

Does the video downloading and caching initiated by the prepopulation feature get the disk based compression (DBC) benefits?

Yes. Because the video file is cached, the attempt to access the video is served from the cache.

Office 365 Acceleration

March 12, 2021

1. Why do we parse the SAN?

It is tedious to create multiple profiles for FQDNS for each of the domains, to overcome this we parse the SAN from the certificates.

2. What is an exclude list?

An error or warning message is displayed If the browser or app does not contain the CA certificate, in such cases the client's IP address will be added to an exclude list after few attempts to connect from browser or app (2-3 times). In the next attempt, connection is not SSL proxied and the page loads without any error or warning. The client IP address will remain in the exclude list for 48hrs. The exclude list is maintained only for split proxy.

3. Where to check for office 365 acceleration connection information?

Navigate to **Monitoring > Connections > Accelerated Connections**, check for the SSL proxy state. For connection details, click the details icon.

Monitoring > Optimization > Connections > Accelerated Connections									
Accelerated Connections									
Action	Details	Initiator	Responder	Duration	Idle	Bytes Transferred	Compression Ratio/Type	Bandwidth Savings (%)	SSL Proxy
		172.16.139.236 : 49713	13.107.6.156 : 443	1m 0s	0m 55s	15.42 KB	1.9 to 1 (Disk)	51.1	True
		172.16.139.236 : 49719	111.221.111.196 : 443	0m 57s	0m 56s	7.41 KB	2.8 to 1 (Disk)	65.4	True
		172.16.139.236 : 49717	23.101.222.248 : 443	1m 0s	0m 58s	21.18 KB	1.1 to 1 (Disk)	8.4	True

4. What happens if exclude list option is not enabled by default as part of SSL profile configuration?

If the browser or app does not contain the CA certificate, it displays an error or warning and the connections from that client or App will be blocked. To avoid such issues, select **Exclude List** option as part of SSL profile configuration.

5. What happens if the required SAN's are not part of the configured/created proxy certificate?

The connections will not be SSL proxied and there will be no acceleration benefits for non-proxied SSL connections.

6. What happens when the client is not part of the domain or if the client does not have the root certificate of the domain?

The connections get blocked if exclude list is not enabled.

7. What happens if the Data Center side Citrix SD-WAN WANOP does not have root or intermediate CA's?

The connections are blocked or the Office 365 application pages which require the missing root or intermediate CA's are partially loaded. To unblock the connections or to have these pages fully loaded, either add the appropriate CA certificates or disable the SSL profile from acceleration.

8. How to know which clients are excluded from acceleration?

Excluded client information can be known from logs or by using the CLI command *show ssl-exclude -list*.

9. What to do when clients are excluded?

By default, exclude list information from the appliance will be cleared after 48 hours. User can forcibly clear the exclude list information using CLI commands `*clear ssl-exclude -list -\<all\>/\<Client_IP\>*`.

10. How to know which SSL connections(SNI's) are not proxied?

From the logs or by using the CLI command *show ssl-non-proxied-sni*, you can know the list of the non-proxied SNI's.

11. How to clear non-proxied SNI's?

Using the CLI command `*clear ssl-non-proxied-sni -\<all\>/\<server name identifier\>*`.

12. What is the default time for client in exclude state?

Client remains in the exclude state for 48 hrs.

13. Can we have multiple profiles applied for a particular service class?

Yes, we can apply service classes with multiple SSL profiles.

To do this, on your Virtual WAN appliance navigate to **Configuration > Service Class > Web (Internet-Secure) > Edit > Edit** (Application) and add the available profiles.

14. How do you check the reason for non-proxied connections?

Check the TCP connection page, for more information check the logs. To debug the non-proxied connection issues, do the following.

- a) If the log shows no valid configuration - Set the valid configuration. For more information on configuring office 365 feature, see [Office 365 Acceleration](#).
- b) If the log shows that certification verification failed - Add valid CA certificates to the data center side Citrix SD-WAN WANOP appliance.

- c) if the log shows client excluded - Information about excluded clients can be cleared from the appliance using the CLI command `*clear ssl-exclude-list -\<all\>/\<Client_IP\>*`.

Additional Notes

- Logging to OneDrive client sometimes displays a warning message “spurious warning”, This is a known issue from Microsoft (<https://support.microsoft.com/en-us/kb/3097938>) and not specific to Citrix SD-WAN WANOP appliance.
- For the office 365 redirected pages to be proxied, it is recommended to create a separate proxy certificate which contains SAN list corresponding to the certificate of the redirected pages. Create another profile with this proxy certificate and apply to the service class. Also add the relevant CA in the Citrix SD-WAN WANOP appliance.
- Sometimes browser doesn't show the correct CA certificates, in such cases use Wireshark or OpenSSL to get the root and Intermediate CA names and get the certificates from 'authentic' source (for example, windows SSL store).
- Difference in browser behavior can be observed in accessing the office 365 applications from different browsers having no required certificates and with Exclude list option disabled.
- When office 365 connections are SSL proxied (that means SSL proxy set to True) and in browser office 365 certificate is displayed instead of the proxy certificate, it is recommended to open the browser in in-cognitive mode and check the behavior or clear the cache and then check the behavior again.
- Microsoft Office 365 includes many components and applications such as OneDrive, Outlook, SharePoint, Word, PPT, Excel, OneNote. All these applications have been tested and is known to work without any problems. Other applications are expected to work without any problems, too; however, this status can change over time, and you might encounter unknown problems.

Compression

March 12, 2021

Citrix SD-WAN WANOP compression uses breakthrough technology to provide transparent multilevel compression. It is true compression that acts on arbitrary byte streams. It is not application-aware, is indifferent to connection boundaries, and can compress a string optimally the second time it appears in the data. Citrix SD-WAN WANOP compression works at any link speed.

The compression engine is very fast, allowing the speedup factor for compression to approach the compression ratio. For example, a bulk transfer monopolizing a 1.5 Mbps T1 link and achieving a 100:1 compression ratio can deliver a speedup ratio of almost 100x, or 150 Mbps, provided that the WAN bandwidth is the only bottleneck in the transfer.

Unlike with most compression methods, Citrix SD-WAN WANOP compression history is shared between all connections that pass between the same two appliances. Data sent hours, days, or even weeks earlier by connection A can be referred to later by connection B, and receive the full speedup benefit of compression. The resulting performance is much higher than can be achieved by conventional methods.

Compression can use the appliance's disk as well as memory, providing up to terabytes of compression history.

How compression works

All compression algorithms scan the data to be compressed, searching for strings of data that match strings that have been sent before. If no such matches are found, the literal data is sent. If a match is found, the matching data is replaced with a pointer to the previous occurrence. In a very large matching string, megabytes or even gigabytes of data can be represented by a pointer containing only a few bytes, and only those few bytes need be sent over the link.

Compression engines are limited by the size of their compression history. Traditional compression algorithms, such as LZS and ZLIB, use compression histories of 64 KB or less. Citrix SD-WAN WANOP appliances maintain at least 100 GB of compression history. With more than a million times the compression history of traditional algorithms, the Citrix SD-WAN WANOP algorithm finds more matches and longer matches, resulting in superior compression ratios.

The Citrix SD-WAN WANOP compression algorithm is very fast, so that even the entry-level appliances can saturate a 100 Mbps LAN with the output of the compressor. The highest-performance models can deliver well over 1 Gbps of throughput.

Only payload data is compressed. However, headers are compressed indirectly. For example, if a connection achieves 4:1 compression, only one full-sized output packet is sent for every four full-sized input packets. Thus, the amount of header data is also reduced by 4:1.

Compression as a general-purpose optimization:

Citrix SD-WAN WANOP compression is application-independent: it can compress data from any non-encrypted TCP connection.

Unlike caching, compression performance is robust in the face of changing data. With caching, changing a single byte of a file invalidates the entire copy in the cache. With compression, changing a single byte in the middle of a file just creates two large matches separated by a single byte of nonmatching

data, and the resulting transfer time is only slightly greater than before. Therefore, the compression ratio degrades gracefully with the amount of change. If you download a file, change 1% of it, and upload it again, expect a 99:1 compression ratio on the upload.

Another advantage of a large compression history is that precompressed data compresses easily with Citrix SD-WAN WANOP technology. A JPEG image or a YouTube video, for example, is precompressed, leaving little possibility for additional compression the first time it is sent over the link. But whenever it is sent again, the entire transfer is reduced to just a handful of bytes, even if it is sent by different users or with different protocols, such as by FTP the first time and HTTP the next.




In practice, compression performance depends on how much of the data traversing the link is the same as data that has previously traversed the link. The amount varies from application to application, from day to day, and even from moment to moment. When looking at a list of active accelerated connections, expect to see ratios anywhere from 1:1 to 10,000:1.

Monitoring > Optimization > Connections > Accelerated Connections

Accelerated Connections

Unaccelerated Connections

Action ▾


Details	Initiator	Responder	Duration	Idle	Bytes Transferred ↑	Compression Ratio/Type
	172.16.0.1 : 55222	172.16.0.71 : 3120	0m 43s	0m 13s	7.39 MB	969.0 to 1 (Disk)
	172.16.0.52 : 58730	208.85.46.23 : 80	1m 41s	1m 37s	1.70 MB	97.9 to 1 (Disk)
	172.16.0.34 : 51869	173.194.33.142 : 443	1m 7s	0m 3s	913.82 KB	N/A (None)

Compress encrypted protocols:

Many connections showing poor compression performance do so because they are encrypted. Encrypted traffic is normally uncompressible, but Citrix SD-WAN WANOP appliances can compress encrypted connections when the appliances join the security infrastructure. Citrix SD-WAN WANOP appliances join the security infrastructure automatically with Citrix Citrix Virtual Apps and Desktops, and can join the security infrastructure of SSL, Windows file system (CIFS/SMB), and Outlook/Exchange (MAPI) servers with manual configuration.

Adaptive, zero-config operation:

To serve the different needs of different kinds of traffic, Citrix SD-WAN WANOP appliances use not one but five compression engines, so the needs of everything from the most massive bulk transfer to the most latency-sensitive interactive traffic can be accommodated with ease. The compression engine is matched dynamically to the changing needs of individual connections, so that compression is automatically optimized. An added benefit is that the compression engine requires no configuration.



Memory based compression

Most of the compression engines use RAM to store their compression history. This is called memory-based compression. Some appliances devote gigabytes of memory to these compression engines. Memory-based compression has a low latency and is often chosen automatically for interactive tasks such as Virtual Apps/Virtual Desktops traffic.

Disk based compression

The disk-based compression engine uses anywhere between tens of gigabytes and terabytes of memory to store compression history, allowing more and better compression matches. The disk-based compression engine is very fast but sometimes has a higher latency than the memory-based engines, and is often chosen automatically for bulk transfers.

Enable or disable compression

Compression is enabled, on a per-service-class basis, on the Configuration: Service Classes page. This page has a pull-down menu for each service class, with the following options:

- **Disk**, meaning that both disk based and memory based compression are enabled. This option should be selected unless you have a specific reason for disabling it.
- **Memory**, meaning that memory based compression is enabled but disk based compression is not. This setting is rarely used, because the appliance automatically selects memory or disk if both types of compression are enabled.

- **Flow-Control Only**, which disables compression but enables flow-control acceleration. Select this option for services that are always encrypted, and for the FTP control channel.
- **None**, meaning that compression and flow-control are both disabled.

For more information, see [Service Classes](#).

Measure disk based compression performance

The Compression Status tab of the

Reports: Compression page reports the system compression performance since the system was started or since the Clear button was used to reset the statistics. Compression for individual connections is reported in the connection closure messages in the system log.

Compression performance varies with a number of factors, including the amount of redundancy in the data stream and, to a lesser extent, the structure of the data protocol.

Some applications, such as FTP, send pure data streams; the TCP connection payload is always byte-for-byte identical to the original data file. Others, such as CIFS or NFS, do not send pure data streams, but mix commands, metadata, and data in the same stream. The compression engine distinguishes the file data by parsing the connection payload in real time. Such data streams can easily produce compression ratios between 100:1 and 10,000:1 on the second pass.

Average compression ratios for the link depend on the relative prevalence of long matches, short matches, and no matches. This ratio is dependent on the traffic and is difficult to predict in practice.

Test results show the effect of multi-level compression as a whole, with memory based and disk based compression each making its contribution.

Maximum compression performance is not achieved until the storage space available for disk based compression is filled, providing a maximum amount of previous data to match with new data. In a perfect world, testing would not conclude until the appliance's disks had not only been filled, but filled and overwritten at least once, to ensure that steady-state operation has been reached. However, few administrators have that much representative data at their disposal.

Another difficulty in performance testing is that acceleration often exposes weak links in the network, typically in the performance of the client, the server, or the LAN, and these are sometimes misdiagnosed as disappointing acceleration performance.

You can use Iperf or FTP for preliminary and initial testing. Iperf is useful for preliminary testing. It is extremely compressible (even on the first pass) and uses relatively little CPU and no disk resources on the two endpoint systems. Compressed performance with Iperf should send more than 200 Mbps over a T1 link if the LANs on both sides use Gigabit Ethernet, or slightly less than 100 Mbps if there is any Fast Ethernet equipment in the LAN paths between endpoints and appliances.

Iperf is preinstalled on the appliances (under the Diagnostics menu) and is available from <http://iperf.sourceforge.net/>. Ideally, it should be installed and run from the endpoint systems, so that the network is tested from end to end, not just from appliance to appliance.

FTP is useful for more realistic testing than is possible with Iperf. FTP is simple and familiar, and its results are easy to interpret. Second-pass performance should be roughly the same as with Iperf. If not, the limiting factor is probably the disk subsystem on one of the endpoint systems.

To test the disk based compression system:

1. Transfer a multiple-gigabyte data stream between two appliances with disk based compression enabled. Note the compression achieved during this transfer. Depending on the nature of the data, considerable compression may be seen on the first pass.
2. Transfer the same data stream a second time and note the effect on compression.

Compression reports in premium edition

Citrix SD-WAN Premium (Enterprise) edition does not have a view for showing compression reports on a per protocol or application basis through WANOP service classes, which have the protocol or application association. If you are using a Premium (Enterprise) edition appliance then the only report available for compression is a connection level compression report which does not give visibility into the extent to which a protocol has been optimized or compressed. Compression reports are available in the WAN Optimization GUI which displays a break-up of all unique protocols and how reports have been optimized over a period of time.

In the Citrix SD-WAN Premium (Enterprise) Edition appliance GUI, for WAN Optimization, the following widgets have been added under the WAN Optimization Dashboard.

- Consolidated compression ratio—all traffic passing through WANOP appliance and total number of accelerated and un-accelerated connections. This allows you to monitor total traffic transmitted from LAN to WAN.
- Compression Ratio - top 10 Service Classes.
- Aggregated Link Throughput –LAN and WAN.

Consolidated compression ratio:

This report displays consolidated compression ratio for all traffic transmitted to WANOP and total number of accelerated and un-accelerated connections. It also shows the up-time of the WANOP service in the appliance.

Monitoring > WAN Optimization > Dashboard			
Up Time 1 hr 17 min	Compression Ratio 12,283 to 1 (91.859%)	Accelerated Connections 12	Unaccelerated Connections 2

Aggregated link throughput:

This report displays the total traffic that is transmitted to WANOP and the total traffic that transmits out with break-ups in categories of optimized and unoptimized data on both ends.

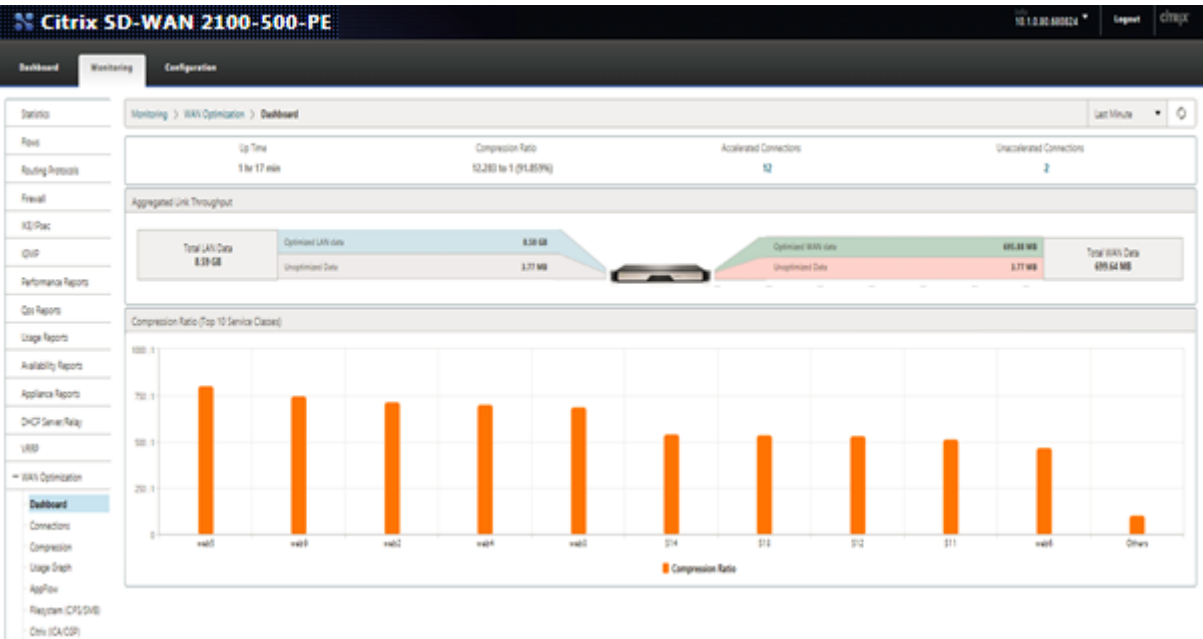


Compression ratio (Top 10 service classes):

In the Citrix SD-WAN appliance GUI, you can check the connection details and the compression ratio (per service-class dashboard) by navigating to **Monitoring > WAN Optimization**. This auto selects the Dashboard node and provides an overview in the form of dashboard.

The graph displays the top 10 values of compression ratio for traffic categorised by service classes.

An extra “others” bar is displayed, which shows the compression ratio for all the other accelerated connections that are part of the system in addition to the Top 10 service classes compression ratio reports.



HTTP acceleration

March 12, 2021

The Citrix SD-WAN WANOP accelerator uses a variety of zero-config optimizations to speed up HTTP traffic. This in turn accelerates Web pages and any other applications using the HTTP protocol (file downloads, video streaming, automatic updates, and so on).

Optimizations that accelerate HTTP include compression, traffic shaping, flow control, and caching.

Compression

HTTP is an ideal application for Citrix SD-WAN WANOP multi-level compression.

Static content, including standard HTML pages, images, video, and binary files, receives variable amounts of first-pass compression, typically 1:1 on pre-compressed binary content, and 2:1 or more on text-based content. Starting with the second time the object is seen, the two largest compression engines (memory-based compression and disk-based compression) deliver extremely high compression ratios, with larger objects receiving compression ratios of 1,000:1 or more. With such high compression ratios, the WAN link stops being the limiting factor, and the server, the client, or the LAN becomes the bottleneck.

The appliance switches between compressors dynamically to give maximum performance. For example, the appliance uses a smaller compressor on the HTTP header and a larger one on the HTTP body.

Dynamic content, including HTTP headers and dynamically generated pages –pages that are never the same twice but have similarities to each other –are compressed by the three compression engines that deal with smaller matches. The first time a page is seen, compression is good. When a variant on a previous page is seen, compression is better.

Traffic shaping

HTTP consists of a mix of interactive and bulk traffic. Every user's traffic is a mix of both, and sometimes the same connection contains a mix of both. The traffic shaper seamlessly and dynamically ensures that each HTTP connection gets its fair share of the link bandwidth, preventing bulk transfers from monopolizing the link at the expense of interactive users, while also ensuring that bulk transfers get any bandwidth that interactive connections do not use.

Flow control

Advanced retransmission algorithms and other TCP-level optimizations retain responsiveness and maintain transfer rates in the face of latency and loss.

Video caching

HTTP caching for video files was introduced in release 7.0. Caching involves saving HTTP objects to local storage and serving them to local clients without reloading them from the server.

What is the difference between caching and compression? While caching provides speedup that is similar to compression, the two methods are different, making them complementary.

- Compression speeds up transfers from the remote server, and this higher data rate can place a higher load on the server if compression were not present. Caching prevents transfers from the server, and reduces the load on the server.
- Compression works on any data stream this is similar to a previous transfer –if you change the name of a file on the remote server and transfer it again, compression will work perfectly. Caching works only when the object being requested by the client and the object on the disk are known to be identical –if you change the name of a file on the remote server and transfer it again, the cached copy is not used.
- Compressed data cannot be delivered faster than the server can send it. Cached data is dependent only on the speed of the client-side appliance.
- Compression is CPU-intensive; caching is not.

How HTML5 works

March 12, 2021

HTML5 uses HTTP, which is a request/response protocol for communication between clients and servers. A client initiates a TCP connection and uses it to send HTTP requests to the server. The server responds to these requests by granting access rights for the available resources. After the client and server establish a connection, the messages exchanged between them contain only WebSocket headers, not HTTP headers.

The infrastructure of HTML5 consists of WebSockets, which further use the existing HTTP infrastructure to provide a lightweight mechanism for communication between a client and a web server. You typically implement the WebSocket protocol in a browser and web servers. However, you can use this protocol with any client or server application.

When a client attempts to make a connection using WebSockets, web servers treat the WebSocket handshake as an upgrade request, and the server switches to the WebSocket protocol. The WebSocket protocol enables frequent interaction between the browser and the web servers. Therefore, you can use this protocol for live updates, such as stock indexes and score cards, and even live games. This is possible because of a standardized way for the server to send unsolicited responses to the client while maintaining an open connection for two-way ongoing communication between the client browser and the server.

Note

You can also achieve this effect, in non-standardized ways, by using various other technologies, such as Comet. For more information about Comet, see [http://en.wikipedia.org/wiki/Comet_\(programming\)](http://en.wikipedia.org/wiki/Comet_(programming)).

The WebSocket protocol communicates over TCP ports 80 and 443. This facilitates communication in environments that use firewalls to block non-web Internet connections. Additionally, WebSocket has its own fragmentation mechanism. A WebSocket message can be sent as multiple WebSocket frames.

Note

You cannot use WebSocket if the web applications on the servers do not support it.

How HTML5 establishes a WebSocket session

A browser supporting HTML5 uses JavaScript APIs to perform the following tasks:

- Open a WebSocket connection.
- Communicate over the WebSocket connection.
- Close the WebSocket connections.

To open a WebSocket connection, the browser sends an HTTP upgrade message to the server for switching to the WebSocket protocol. The server either accepts or rejects this request. Following are snippets of a sample client request and server response:


- Sample client request

```
pre codeblock GET /HTTP/1.1 Upgrade: websocket Sec-websocket-protocol: <List of protocols that the client supports over this websocket session, such as an application level protocol, for example ICA.> Sec-websocket-extensions: <List of extensions client wants applied to this session, such as compression.> Sec-WebSocket-version: <Version of websocket protocol that the client intends to use.> <!--NeedCopy-->
```

- Sample server response

```
pre codeblock HTTP/1.1 101 Switching Protocols Upgrade: websocket Connection: Upgrade Sec-WebSocket-Protocol: <One from the list of protocols in the client request.> Sec-WebSocket-extensions: <List of extensions server accepts for session.> Sec-WebSocket-version: <Version of websocket protocol that the server supports .> <!--NeedCopy-->
```

The following figure shows the sequence of messages exchanged between a client and a server:



During an HTML5 connection, the following messages are exchanged between the client and the server:

- Client sends an HTTP request to upgrade WebSocket.
- Server responds to the client request and switches to WebSocket protocol.
- Server sends WebSocket frames to the client.
- Client sends a request to close the WebSocket.
- Server closes the WebSocket.

Internet Protocol version 6 (IPv6) acceleration

March 12, 2021

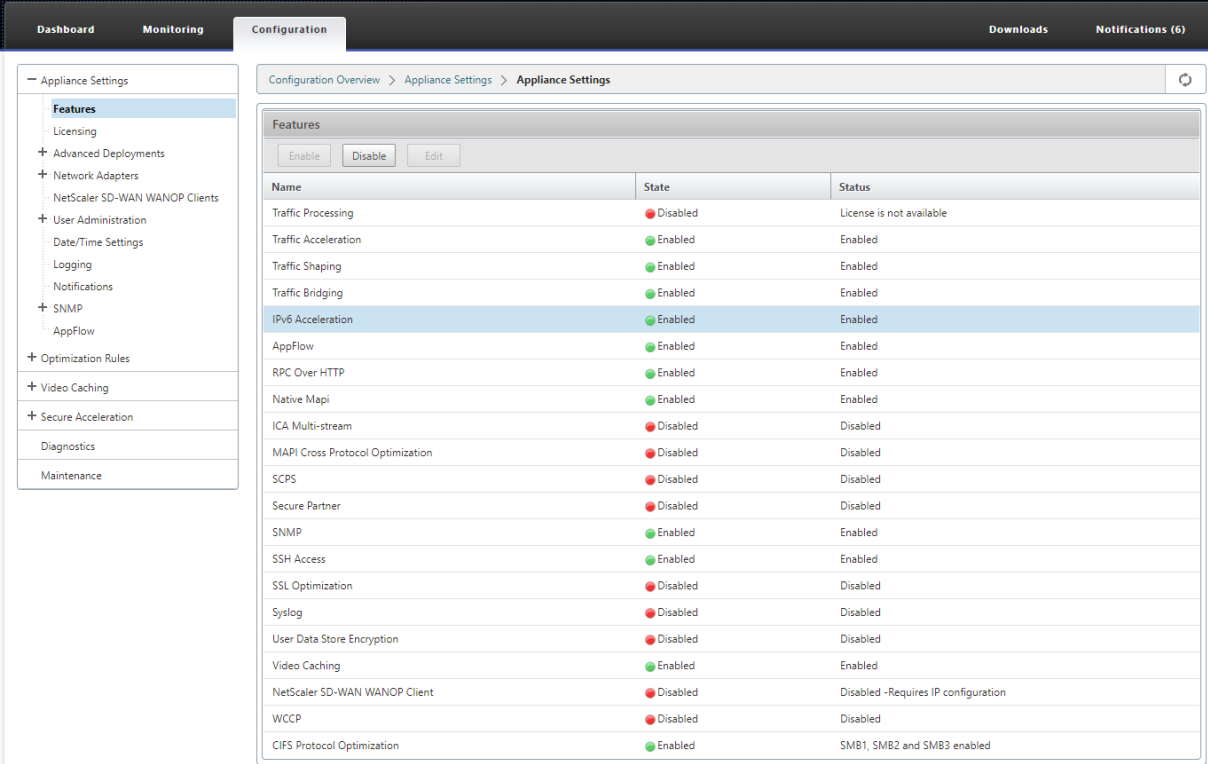
When you connect to the Internet through a device, the device is assigned an IP address. The IP address identifies the appliance and indicates its location. The number of devices connecting to the Internet is rapidly increasing. As a result, it is difficult to manage the request for the IP addresses 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 Citrix SD-WAN WANOP appliance, you can enable the IPv6 Acceleration feature.

By default, IPv6 is disabled on the appliance. To enable IPv6 acceleration on a Citrix SD-WAN WANOP appliance, navigate to **Configuration > Appliance Settings > Feature** page and enable the **IPv6 Acceleration** feature.



Verify IPv6 connections

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 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.

Unaccelerated Connections											
Action	Initiator	Responder	Duration	Idle	Bytes Transferred	Compression Ratio/Type	SSL Proxy	Service Class	State	Partner Unit	CloudBridge Instance
	2000:10:60730	4000:10:5001	6m 33s	0m 0s	34.29 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60717	4000:10:5001	6m 33s	0m 0s	34.27 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60725	4000:10:5001	6m 33s	0m 0s	33.63 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	192.168.1.10:33688	172.16.1.10:5001	2m 19s	0m 0s	26.03 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	192.168.1.10:33689	172.16.1.10:5001	2m 19s	0m 0s	25.73 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60718	4000:10:5001	6m 33s	0m 0s	31.32 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60722	4000:10:5001	6m 33s	0m 0s	31.07 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60728	4000:10:5001	6m 33s	0m 0s	30.92 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60729	4000:10:5001	6m 33s	0m 0s	30.55 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60715	4000:10:5001	6m 33s	0m 0s	30.29 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60727	4000:10:5001	6m 33s	0m 0s	29.36 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60721	4000:10:5001	6m 33s	0m 0s	26.23 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60713	4000:10:5001	6m 33s	0m 0s	24.67 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60714	4000:10:5001	6m 33s	0m 0s	23.58 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60726	4000:10:5001	6m 33s	0m 0s	23.08 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60711	4000:10:5001	6m 33s	0m 0s	22.99 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60729	4000:10:5001	6m 33s	0m 0s	22.95 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60723	4000:10:5001	6m 33s	0m 0s	22.71 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A
	2000:10:60712	4000:10:5001	6m 33s	0m 0s	22.55 MB	N/A (None)	False	Iperf	Open	10.105.145.125	N/A


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 Citrix 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.

Top Applications						
Application	Sent Rate (bps)	Received Rate (bps)	Total Bytes Sent	Total Bytes Received	Total %	Order
IPv6 ICMP	719.56 K	719.56 K	5.4 M	5.4 M	98.3	1
HTTPS	10.57 K	9.64 K	79.3 K	72.35 K	1.38	2
Microsoft DirectX Gaming	416	416	3.14 K	3.14 K	0.06	4
Other TCP	312	312	2.35 K	2.35 K	0.04	5
Other UDP	128	1.7 K	984	12.73 K	0.12	3
ARP	24	488	232	3.66 K	0.04	6
SNMP	0	496	0	3.76 K	0.03	7
ICMP	0	376	0	2.84 K	0.03	8

- **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.

Application Groups				
Application	Application Group	Parent Application	Bytes Sent	Bytes Received
IPv6 ICMP	IP Protocols	IPv6	5.4 M	5.4 M
HTTPS	Web, Security Protocols	TCP	79.3 K	72.35 K
Microsoft DirectX Gaming	Games	TCP	3.14 K	3.14 K
Other TCP	N/A	N/A	2.35 K	2.35 K
Other UDP	N/A	N/A	984	12.73 K
ARP	Legacy Or Non-IP	N/A	232	3.66 K
SNMP	Network Management, Infrastructure	UDP	0	3.76 K
ICMP	Infrastructure, IP Protocols	IPv6	0	2.84 K

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.

Link definitions

March 12, 2021

Link definitions enable the appliance to prevent congestion and loss on your WAN links and to perform traffic shaping. A link definition specifies which traffic is associated with the defined link, the maximum bandwidth to allow for traffic received on the link, and the maximum bandwidth for traffic

sent over the link. The definition also identifies traffic as inbound or outbound and as WAN-side or LAN-side traffic. All traffic flowing through the appliance is compared to your list of link definitions, and the first matching definition identifies the link to which the traffic belongs.

By performing the Quick Installation procedure, you customize the appliance's default link definitions. You have then defined the appliance's link to the WAN and its link to the LAN. For a simple inline deployment, no further configuration of link definitions is necessary. Other types of deployments require additional configuration of link definitions.

Every link has two bandwidth limits, representing the sending speed and the receiving speed. Only when the link speed is known can the appliance inject traffic into the link at exactly the right speed, thus eliminating the congestion and packet loss that result from attempting to send too much, or the loss of performance that results from sending too little. When placed between a fast LAN and a slower WAN and acting as a *virtual gateway*, the appliance has the ability to receive traffic faster than the WAN can accept it, creating a backlog of traffic. The existence of this backlog enables the appliance to choose which packet to send next, and this choice in turn makes traffic shaping possible. Unless there are packets from multiple streams to choose from, there is no ability to favor one stream over the other. Traffic shaping is therefore dependent on the existence of the virtual gateway and correctly set bandwidth limits.

Note

Link definitions normally apply to connections to the accelerated pair of bridge ports. The two motherboard ports, Primary and Aux1, can also be defined as links, but doing so rarely serves any purpose, because they are used for management and as a back-channel for high-availability and group modes, not for WAN traffic.

Important

Important: For link-definition purposes, a *link* is a physical link, with its own bandwidth capacity. It is typically a cable that leaves the building. Remember the following points:

- A VLAN is not a link.
- A virtual link is not a link.
- A tunnel is not a link.

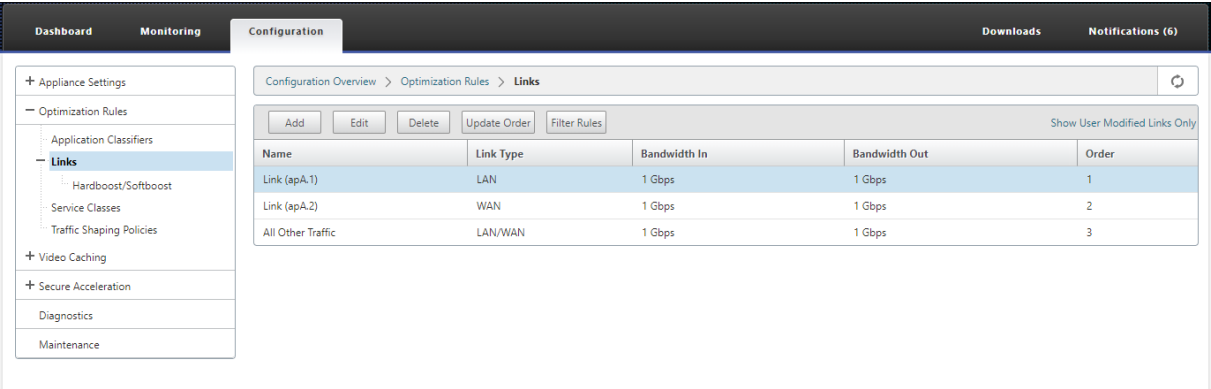
Default link definitions

Navigate to **Configuration > Optimization Rules > Links** to view the currently defined links. The following links are defined by default.

1. **apA.1**, one of the two ports on the accelerated bridge.
2. **apA.2**, the other port on the accelerated bridge.
3. If the system has dual accelerated bridges, apB.1 and apB.2 also exist.

- 4. All Other Traffic, which is not a true link, but is a catch-all for traffic that does not match any actual link definitions.

The order in which the links are shown on this page is significant. When deciding which link a packet belongs to, the Appliance tests the links in order, and the first matching link is selected. This means that overlapping definitions are allowed, and the last definition in the link can match all traffic, serving as a default link. To change the order click **Update Order**.



Manage link definitions in traffic shaping


March 12, 2021

To manage a link, the traffic shaper needs the following information:

- The speed of the link in both the send and receive directions.
- Whether the link is a WAN link or a LAN network.
- A way of distinguishing link traffic from other traffic.
- The direction in which traffic is flowing over the link.

Link Speed—*Link speed* always refers to the speed of the physical link. In the case of a WAN link, it is the speed of the WAN segment that terminates in the building with the Citrix SD-WAN WANOP appliance. The speed of the other end of the link is not considered. For example, the following figure shows a network of four appliances. Each appliance has its incoming and outgoing bandwidths set to 95% of the speed of its own, local WAN segment, without regard to the speed of the remote endpoints.

Figure 1. Local bandwidth limits track local link speeds



The reason for setting the bandwidth limits to 95% of the link speed instead of 100% is to allow for link overhead (few links can carry data at 100% of their published speeds) and to ensure that the appliance is slightly slower than the link, so that it becomes a slight bottleneck. Traffic shaping is not effective unless the traffic shaper is the bottleneck in the connection.

Distinguishing different types of traffic—In each link definition, you must declare whether the definition applies to a WAN link or a LAN network.

The traffic shaper needs to know whether a packet is traveling on the WAN, and, if so, in which direction. To provide this information:

- For simple inline deployments, you declare that one port of the accelerated bridge belongs to the WAN link and that the other port belongs to the LAN.
- In other deployment modes, the appliance examines IP addresses, MAC addresses, VLANs, or WCCP service groups. (Note that testing for WCCP service groups is not yet supported.)
- If a site has multiple WANs, the local link definitions must include rules that enable the appliance to distinguish traffic from different WANs.

Configure link definitions

March 12, 2021

Link definitions are arranged in an ordered list, one entry per link, which is tested from top to bottom for every packet entering or leaving the appliance. The first matching definition determines which link the packet belongs to. Within each link definition is an ordered list of rules, which is also tested from top to bottom. Each packet is compared to these rules, and if it matches one of them, the packet is considered to be traveling over that link.

Within a single rule, the fields are all ANDed together, so all specified values have to match. All fields default to Any, a wildcard entry that always matches. When a field consists of a list, such as a list of IP subnets, the list entries are ORed together. That is, if any element matches, the list as a whole is considered to be a match.

Links can be based on the Ethernet adapter associated with the traffic, the source and destination IP addresses, VLAN tag, WCCP service group (for WCCP-GRE only), and the source and destination Ethernet MAC address. A simple inline deployment might identify only the LAN-side and WAN-side accelerated bridge ports (apA.1 and apA.2), while a complex datacenter deployment might need to use most of the options provided to disambiguate traffic.

Defining a link in terms of its IP addresses is possible except when redundant links are used. Since a given packet may go over either link in an active-standby or active-active dual-link deployment, some other method must be used to determine which link the packet is using. If dual bridges are used, then the traffic for one link can go over apA and the other over apB, and the links can be defined in terms of adapters. If the two links are served by different routers, the MAC addresses of the routers can be used to tell the traffic apart. When all else fails, WCCP-GRE can be used, and the router can use a different service group for each WAN link, allowing the Citrix SD-WAN WANOP unit to tell the link traffic apart in by service group.

Citrix recommends port based link definitions for simple inline deployments, and IP based link definitions for all other deployments.

To configure link definitions:

1. Navigate to **Configuration > Optimization Rules > Links** and click **Add**.

The screenshot shows the 'Create Links' form in the Citrix SD-WAN WANOP Configuration page. The form has the following fields:

- Name***: Text input field with the value 'WAN-side link'.
- Link Type***: Dropdown menu with the value 'WAN'.
- Bandwidth In***: Text input field with the value '67' and a unit dropdown menu with the value 'mbps'.
- Bandwidth Out***: Text input field with the value '950' and a unit dropdown menu with the value 'mbps'.
- Filter Rules**: A table with columns for Adapter, Source IP Address, Dest IP Address, VLAN, WCCP Service Group, Source MAC Address, and Destination MAC Address. The table contains one row with the value 'Any' for all columns.

Below the form, there are buttons for 'Add', 'Edit', and 'Delete'. At the bottom, there are buttons for 'Create' and 'Close'.

2. Enter values for the following parameters:
 - **Name**: A descriptive name of the link, that can also describe if it is a LAN side link or a WAN side link.
 - **Link Type**: The link type, either LAN or WAN.
 - **Bandwidth In**: The incoming bandwidth limit.
 - **Bandwidth Out**: The outgoing bandwidth limit.

3. In the **Filter Rules** section, click **Add** and enter values for the following parameters:

- **Adapter:** This specifies a list of adapters (Ethernet ports). When links can be identified by ethernet adapter, this simplifies configuration.
- **Source IP Address:** The Source IP rules are considered for packets entering the unit (packets exiting the unit are ignored). On these packets, the rules in the Src IP field are compared against the Source Address field in the IP header. The rule specifies a list of IP addresses or subnets. Negative matches, such as “Exclude 10.0.0.1” are also supported.
- **Destination IP Address:** The Destination IP rules are considered for packets exiting the unit (packets entering the unit are ignored). On these packets, the rules in the Dst IP field are compared against the Destination Address field in the IP header. The rule specifies a list of IP addresses or subnets. Negative matches, such as “Exclude 10.0.0.1” are also supported.
- **VLAN:** The VLAN rules are applied to the VLAN headers of packets entering or exiting the unit.
- **WCCP Service Group:** The WCCP Service Group rules are applied to GRE-encapsulated WCCP packets entering or leaving the unit. (This does not work with L2 WCCP.)
- **Source MAC Address:** The Source MAC address used as a filter criteria.
- **Destination MAC Address:** The destination MAC address used as a filter criteria.


4. Click **Create**.

The traffic classifier uses the Src IP and Dest IP fields in a specialized way (the same applies to Src MAC and Dst MAC):

- The Src field is only examined on packets entering the appliance.
- The Dst is only examined on packets exiting the appliance.

Inline links

Most Citrix SD-WAN WANOP appliances use a simple inline deployment, where each accelerated bridge serves just one WAN link. This is the simplest mode to configure.

Simple inline link

In the above figure, all the traffic passing through the accelerated bridge is assumed to be WAN traffic. The link is an ADSL link with different send and receive speeds (6.0 mbps down, 1.0 mbps up). The WAN is connected to accelerated bridge port apA.1, and the LAN is connected accelerated bridge port apA.2.


The tasks for defining the WAN-side link (apA.1) are:

1. Give the WAN a descriptive name, such as “WAN to HQ (apA.1).”
2. Set the type to “WAN.”
3. Set the incoming and outgoing bandwidth limits to 95% of the nominal link speed.
4. Verify that a rule has been defined that specifies the WAN Ethernet adapter, which in this example is apA.
5. Click Create.

The tasks for the LAN-side link (apA.2) are similar:

1. Give it a descriptive name, such as “Local LAN (apA.2).”
2. Set the type to “LAN.”
3. Set the incoming and outgoing bandwidth limits to 95% of the nominal Ethernet speed (95 mbps or 950 mbps).
4. Verify that a rule exists that specifies the LAN Ethernet adapter, which in this example is apA.2.
5. Click Create.

Inline deployment with dual bridges



The configuration is similar to the simple inline link configuration, but the site has a second link, a T1 link to the corporate WAN, in addition to the ADSL Internet link. The Citrix SD-WAN WANOP appliance has two accelerated bridges, one for each WAN link.

Configuration is almost as simple as the single-bridge case, with the following additional steps:

1. Edit a second WAN link on apB, which in this case is apB.1. Set the type to “WAN.” Set the link bandwidth to 95% of the 1.5 mbps T1 speed, and give the link a new name, such as “WAN to HQ.”
2. Add a rule specifying apB.2 to the “LAN” definition and delete the default link definition for apB.2. (Alternatively, you can edit the default link definition for apB.2 to specify it as a LAN link, as was done for apA.2.)

Non-inline links

For other than simple inline deployments (which serve only one WAN per accelerated bridge), use IP subnets instead of bridge ports to distinguish LAN traffic from WAN traffic. This approach is essential for one-arm deployments, which use only a single bridge port. IP subnets are sometimes useful for inline deployments as well, especially when the appliance serves more than one WAN. For simple inline deployments, however, port based links are easier to define.

The traffic classifier applies a specialized convention when examining the Src IP and Dst IP:


- The Src IP field is examined only in packets entering the appliance.
- The Dst IP field is examined only in packets exiting the appliance.

This convention can sometimes be confusing, but it allows the direction of packet travel to be implicitly considered as part of the definition.

Use IP address in link definitions

To Configure simple inline LAN definition using IP-based rules, you can define the LAN and WAN links without specifying the Ethernet ports at all, using the LAN subnet instead:

- Create a rule for the LAN link definition and specify the LAN subnet in the Src IP field.
- Create a rule for the WAN link definition and specify the LAN subnet (not the WAN subnet) in the Dst IP field.

WCCP and Virtual inline modes

Configuration WCCP or virtual inline deployment using IP based rules is the same as using IP address in link definition, because the LAN and WAN IP subnets are identical.

When WCCP-GRE is used, the GRE headers are ignored and the IP headers within the encapsulated data packets are used. Therefore, this same link definition works for WCCP-L2, WCCP-GRE, inline, and virtual inline modes.

(WCCP and virtual inline modes require configuration of your router. WCCP also requires configuration on the Configuration: Advanced Deployments page.)

Manage and monitor using Citrix Application Delivery Management

March 12, 2021

Citrix SD-WAN WANOP AppFlow support enables flexible, customized monitoring of your Citrix SD-WAN WANOP appliances.

The AppFlow interface works with Citrix Application Delivery Management (ADM). Citrix ADM receives detailed information from the appliance, using the AppFlow open standard (<http://www.appflow.org>). Citrix ADM allows you to monitor, manage, and view analytics of the Citrix SD-WAN appliances in your network.

Citrix ADM supports a wide range of devices and can present a more complete view of your network. The Citrix SD-WAN WANOP appliance has an extensive view of WAN traffic, including detailed statistics about Virtual Apps/Virtual Desktops traffic, it provides key insights into the WAN user experience.

For more information, see [Managing Citrix SD-WAN instances using Citrix Application Delivery Management](#).

Virtual Apps/Virtual Desktops example

In a Citrix Virtual Apps and Desktops environment, if a branch user encounters low performance, the administrator might have to monitor the network, the users, and applications hosted on Virtual Apps or Virtual Desktops. The administrators might need to ask the following questions:

- Which part of the network is causing a bad user experience?
- What is an easy way to identify the slowness in published applications?
- Which virtual channels are consuming the most bandwidth over a given time period?
- Which Virtual Desktops or Virtual Apps users are consuming the most bandwidth over a given time period?
- For a given Virtual Desktops user, what is the average client and server-side latency, and the average jitter?
- What are the top applications across all Virtual Apps users, by up-time and total number of launches over a given time period?
- What is the Datacenter latency?

The Citrix SD-WAN WANOP AppFlow support provides answers to all of the above questions, allowing, for example, a congested WAN link to be distinguished from a slow server or a slow client.

Citrix Cloud Connector

March 12, 2021

The Citrix Cloud Connector feature of the Citrix SD-WAN WANOP appliance connects enterprise datacenters to external clouds and hosting environments, making the cloud a secure extension of your enterprise network. Cloud-hosted applications appear as though they are running on one contiguous enterprise network. With Citrix Cloud Connector, you can augment your datacenters with the capacity and efficiency available from cloud providers.

The Citrix Cloud Connector enables you to move your applications to the cloud to reduce costs and increase reliability.

The WAN optimization feature of the Citrix SD-WAN WANOP appliance accelerates traffic, providing LAN-like performance for applications running across enterprise datacenters and clouds.

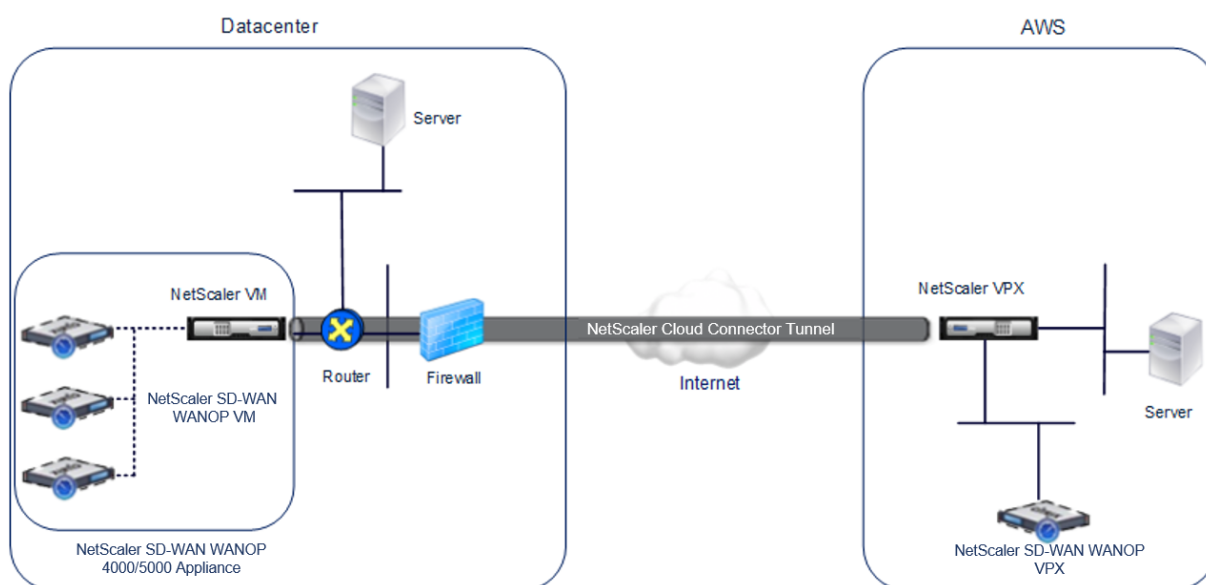
In addition to using Citrix Cloud Connector between a datacenter and a cloud, you can use it to connect two datacenters for a high-capacity secure and accelerated link.

To implement the Citrix Cloud Connector solution, you connect a datacenter to another datacenter or an external cloud by setting up a tunnel called the Citrix Cloud Connector tunnel.

To connect a datacenter to another datacenter, you set up a Citrix Cloud Connector tunnel between two appliances, one in each datacenter.

To connect a datacenter to an external cloud (for example, Amazon AWS cloud), you set up a Citrix Cloud Connector tunnel between a Citrix SD-WAN WANOP appliance in the datacenter and a virtual appliance (VPX) that resides in the Cloud. The remote end point can be a Citrix Cloud Connector or a Citrix VPX with platinum license.

The following illustration shows a Citrix Cloud Connector tunnel set up between a datacenter and an external cloud.



The appliances between which a Citrix Cloud Connector tunnel is set up are called the *end points* or *peers* of the Citrix Cloud Connector tunnel.

A Citrix Cloud Connector tunnel uses the following protocols:

- Generic Routing Encapsulation (GRE) protocol
- Open-standard IPSec Protocol suite, in transport mode

The GRE protocol provides a mechanism for encapsulating packets, from a wide variety of network protocols, to be forwarded over another protocol. GRE is used to:

- Connect networks running non-IP and non-routable protocols.
- Bridge across a wide area network (WAN).
- Create a transport tunnel for any type of traffic that needs to be sent unchanged across a different network.

The GRE protocol encapsulates packets by adding a GRE header and a GRE IP header to the packets.

The Internet Protocol security (IPSec) protocol suite secures communication between peers in the Citrix Cloud Connector tunnel.

In a Citrix Cloud Connector tunnel, IPSec ensures:

- Data integrity
- Data origin authentication
- Data confidentiality (encryption)
- Protection against replay attacks

IPSec uses the transport mode in which the GRE encapsulated packet is encrypted. The encryption is done by the Encapsulating Security Payload (ESP) protocol. The ESP protocol ensures the integrity of the packet by using a HMAC hash function, and ensures confidentiality by using an encryption algorithm. After the packet is encrypted and the HMAC is calculated, an ESP header is generated. The ESP header is inserted after the GRE IP header and, an ESP trailer is inserted at the end of the encrypted payload.

Peers in the Citrix Cloud Connector tunnel use the Internet Key Exchange version (IKE) protocol (part of the IPSec protocol suite) to negotiate secure communication, as follows:

- The two peers mutually authenticate with each other, using one of the following authentication methods:
 - **Pre-shared key authentication.** A text string called a pre-shared key is manually configured on each peer. The pre-shared keys of the peers are matched against each other for authentication. Therefore, for the authentication to be successful, you must configure the same pre-shared key on each of the peers.

- **Digital certificates authentication.** The initiator (sender) peer signs message interchange data by using its private key, and the other receiver peer uses the sender's public key to verify the signature. Typically, the public key is exchanged in messages containing an X.509v3 certificate. This certificate provides a level of assurance that a peer's identity as represented in the certificate is associated with a particular public key.
- The peers then negotiate to reach agreement on:
 - An encryption algorithm.
 - Cryptographic keys for encrypting data in one peer and decrypting the data in the other.

This agreement upon the security protocol, encryption algorithm and cryptographic keys is called a Security Association (SA). SAs are one-way (simplex). For example, when two peers, CB1 and CB2, are communicating through a Connector tunnel, CB1 has two Security Associations. One SA is used for processing out-bound packets, and the other SA is used for processing inbound packets.

SAs expire after a specified length of time, which is called the *lifetime*. The two peers use the Internet Key Exchange (IKE) protocol (part of the IPSec protocol suite) to negotiate new cryptographic keys and establish new SAs. The purpose of the limited lifetime is to prevent attackers from cracking a key.

Also, Citrix SD-WAN WANOP instances on the Citrix Cloud Connector tunnel end-points provide WAN optimization over the tunnel.

Prerequisites to configure Citrix Cloud Connector tunnel

Before setting up a Citrix Cloud Connector tunnel between AWS Cloud and a Citrix SD-WAN WANOP appliance configured for one-arm mode in the data center, verify that the following tasks have been completed:

1. Make sure that the Citrix SD-WAN WANOP appliance in the datacenter is set up correctly. For more information on deploying a Citrix SD-WAN appliance in one-arm mode that uses WCCP/Virtual Inline protocol, see [Sites with One WAN Router](#).
2. Install, configure, and launch a Citrix virtual appliance (VPX instance) on AWS cloud. For more information, see [Installing NetScaler VPX on AWS](#).
3. Install, configure, and launch an instance of Citrix SD-WAN WANOP virtual appliance (VPX) on AWS cloud. For more information, see [Installing SD-WAN VPX S AMI on Amazon AWS](#).
4. On AWS, bind the Citrix SD-WAN WANOP VPX instance on AWS to a load balancing virtual server in the Citrix VPX instance on AWS. This binding is required for sending traffic through the Citrix SD-WAN WANOP VPX instances, to achieve WAN optimization over the Citrix Cloud Connector tunnel.

To create a load balancing virtual server by using the command line interface:

At the command prompt, type:

- **enable ns mode l2**
- **add lb vsrver** <cbvpxonaws_vs_name> ANY * * **-l2Conn ON -m MAC**

To add the Citrix SD-WAN WANOP VPX instance on AWS as a service and bind it to the load balancing virtual server by using the command line interface:

At the command prompt, type:

- **add service** < cbvpxonaws_service_name> <cbvpxonaws_IP> ANY * **-cltTimeout 14400 -svrTimeout 14400**
- **bind lb vsrver** <cbvpxonaws_vs_name> <cbvpxonaws_service_name>

Configure cloud connector tunnel

March 12, 2021

To configure the Citrix Cloud Connector tunnel, use the configuration utility of both the Citrix VPX appliances to perform the following tasks:

- **Create an IPSec profile**—An IPSec profile entity specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, hash algorithm, and PSK, to be used by the IPSec protocol in the Citrix Cloud Connector tunnel.
- **Create an IP tunnel and associate the IPSec profile with it**—An IP tunnel specifies the local IP address, remote IP address, protocol used to set up the Citrix Cloud Connector tunnel, and an IPSec profile entity. The created IP tunnel entity is also called the Citrix Cloud Connector tunnel entity.
- **Create a PBR rule and associate the IP tunnel with it**—A PBR entity specifies a set of conditions and an IP tunnel (Citrix Cloud Connector tunnel) entity. The source IP address range and the destination IP range are the conditions for the PBR entity. You must set the source IP address range and the destination IP address range to specify the subnet whose traffic is to traverse the Citrix Cloud Connector tunnel. For example, consider a request packet that originates from a client on the subnet in the datacenter and is destined to a server on the subnet in the AWS cloud. If this packet matches the source and destination IP range of the PBR entity on the Citrix virtual appliance on the Citrix SD-WAN WANOP appliance in the datacenter, it is considered for Citrix SD-WAN WANOP processing, which sends the packet across the Citrix Cloud Connector tunnel associated with the PBR entity.

To create an IPSEC profile by using the command line interface:

At the command prompt, type:

- `**add ipsec profile** \<ipsec_profile_name\> -**encAlgo** AES -**hashAlgo** HMAC_SHA1 -**lifetime** 500 -**psk** \<password \>`

To create an IP tunnel and bind the IPSEC profile to it by using the command line interface:

At the command prompt, type:

- `**add iptunnel** \<tunnel_name\> \<Remote CBC Public IP\> \<remote_cbs_Netmask\> \<lan_subnet_IP\> -**protocol** GRE -**ipsecProfileName** \<ipsec_profile\>`

To create a PBR rule and bind the IPSEC tunnel to it by using the command line interface:

At the command prompt, type:

- `**add ns pbr** \<pbr_name\> ALLOW -**srcIP** = \<local_lan_subnet\> -**destIP** = \<remote_lan_subnet\> -**ipTunnel** \<tunnel_name\>`
- **apply ns pbrs**

To create an IPSEC profile by using the configuration utility:

1. Navigate to **System > Citrix Cloud Connector > IPSec Profile**.
2. In the details pane, click **Add**.
3. In the Add IPSec Profile dialog box, set the following parameters:
 - Name
 - Encryption Algorithm
 - Hash Algorithm
 - IKE Protocol Version (select V2)
4. Use one of the following IPSec authentication methods to be used by the two peers to mutually authenticate.
 - For Pre-shared key authentication method, set the Pre-Shared Key Exists parameter.
 - For Digital certificates authentication method , set the following parameters:</div>

- Peer Public Key

5. Click **Create**, and then click **Close**.

To create an IP tunnel and bind the IPSEC profile to it by using the configuration utility:

1. Navigate to **System > Citrix Cloud Connector > IP Tunnels**.
2. On the IPv4 Tunnels tab, click **Add**.
3. In the Add IP Tunnel dialog box, set the following parameters:
 - Name
 - Remote IP
 - Remote Mask
 - Local IP Type (In the Local IP Type drop down list, select Subnet IP).
 - Local IP (All the configured IPs of the selected IP type will be populated in the Local IP drop down list. Select the desired IP from the list.)
 - Protocol
 - IPsec Profile
4. Click **Create**, and then click **Close**.

To create a PBR rule and bind the IPSEC tunnel to it by using the configuration utility:

1. Navigate to **System > Network > PBR**.
2. On the PBR tab, click **Add**.
3. In the create PBR dialog box, set the following parameters:
 - Name
 - Action
 - Next Hop Type (Select IP Tunnel)
 - IP Tunnel Name
 - Source IP Low
 - Source IP High
 - Destination IP Low
 - Destination IP High

4. Click **Create**, and then click **Close**.

The new Citrix Cloud Connector tunnel configuration on the Citrix SD-WAN WANOP appliance in the datacenter appears on the Home tab of the Management Service user interface.

The corresponding new Citrix Cloud Connector tunnel configuration on the Citrix VPX appliance in the AWS cloud appears on the configuration utility.

The current status of the Citrix Cloud Connector tunnel is indicated in the Configured Citrix SD-WAN WANOP pane. A green dot indicates that the tunnel is up. A red dot indicates that the tunnel is down.

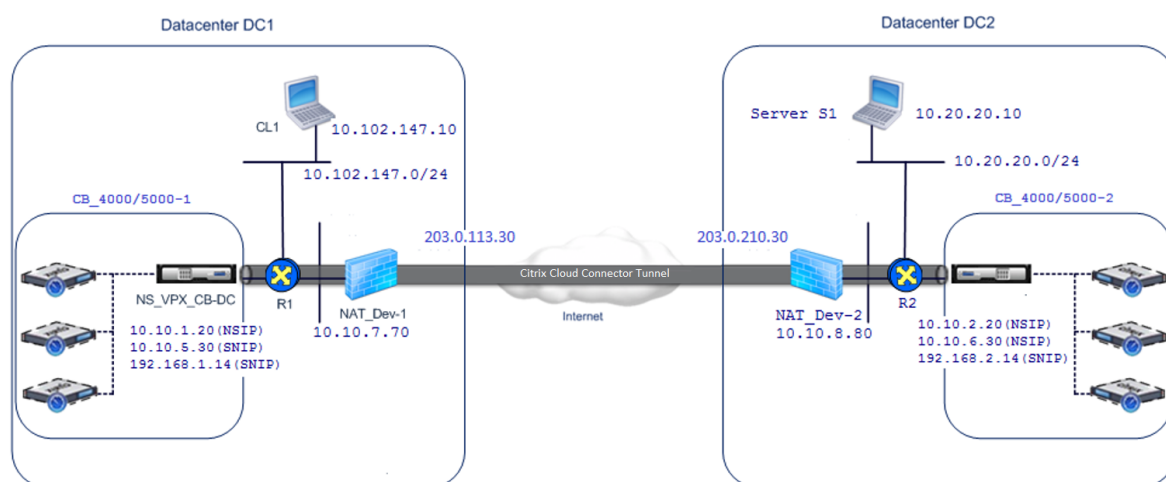
Configure cloud connector tunnel between two datacenters

March 12, 2021

You can configure a Citrix Cloud Connector tunnel between two different datacenters to extend your network without reconfiguring it, and leverage the capabilities of the two datacenters. A Citrix Cloud Connector tunnel between the two geographically separated datacenters enables you to implement redundancy and safeguard your setup from failure. The Citrix Cloud Connector tunnel helps achieve optimal utilization of infrastructure and resources across two datacenters. The applications available across the two datacenters appear as local to the user.

To connect a datacenter to another datacenter, you set up a Citrix Cloud Connector tunnel between a SD-WAN WANOP 4000/5000 appliance that resides in one datacenter and another SD-WAN WANOP 4000/5000 appliance that resides in the other datacenter.

To understand how a Citrix Cloud Connector tunnel is configured between two different datacenters, consider an example in which a Cloud Connector tunnel is set up between Citrix appliance CB_4000/5000-1 in datacenter DC1 and Citrix appliance CB_4000/5000-2 in datacenter DC2.



Both CB_4000/5000-1 and CB_4000/5000-2 function in one arm mode (WCCP/PBR). They enable communication between private networks in datacenters DC1 and DC2. For example, CB_4000/5000-1 and CB_4000/5000-2 enable communication between client CL1 in datacenter DC1 and server S1 in datacenter DC2 through the Citrix Cloud Connector tunnel. Client CL1 and server S1 are on different private networks.

For proper communication between CL1 and S1, L3 mode is enabled on NS_VPX_CB_4000/5000-1 and NS_VPX_CB_4000/5000-2, and routes are configured as follows:

- Router R1 has a route for reaching S1 through NS_VPX_CB_4000/5000-1.
- NS_VPX_CB_4000/5000_1 has a route for reaching NS_VPX-CB_4000/5000-2 through R1.
- S1 should have a route reaching CL1 through NS_VPX-CB_4000/5000-2.
- NS_VPX-CB_4000/5000-2 has a route for reaching NS_VPX_CB_4000/5000-1 through R2.

The following table lists the settings on CB_4000/5000-1 in datacenter DC1.

Entity	Name	Details
IP address of Client CL1		10.102.147.10
Settings on NAT device		
NAT-Dev-1		
NAT IP address on public side		203.0.113.30*
NAT IP address on private side		10.10.7.70
Settings on CB_4000/5000-1		
Management service IP address of CB_4000/5000-1		10.10.1.10
Settings on NS_VPX_CB_4000/5000-1 running on CB_4000/5000-1		
The NSIP address		10.10.1.20
SNIP address		10.10.5.30
Cloud Connector tunnel	Cloud_Connector_DC1-DC2	Local endpoint IP address of the Citrix Cloud Connector tunnel = 10.10.5.30, Remote endpoint IP address of the Citrix Cloud Connector tunnel = 203.0.210.30*
GRE Tunnel Details		

Entity	Name	Details
		Name = Cloud_Connector_DC1-DC2
		IPSec Profile Details
		Name = Cloud_Connector_DC1-DC2, Encryption algorithm = AES, Hash algorithm = HMAC SHA1
Policy based Route	CBC_DC1_DC2_PBR	Source IP range = Subnet in datacenter1 = 10.102.147.0-10.102.147.255, Destination IP range = Subnet in datacenter2 = 10.20.20.0-10.20.20.255, Next hop type = IP Tunnel, IP tunnel name = CBC_DC1_DC2

*These should be public IP addresses.

The following table lists the settings on CB- 4000/5000-2 in datacenter DC2.

Entity	Name	Details
IP address of Server S1		10.20.20.10
Settings on NAT device		
NAT-Dev-2		
NAT IP address on public side		203.0.210.30*
NAT IP address on private side		10.10.8.80
Settings on CB_4000/5000-2		
Management service IP address of CB_SDX-1		10.10.2.10
Settings on		
NS_VPX_CB_4000/5000-2		
running on CB_4000/5000-2		
The NSIP address		10.10.2.20
SNIP address		10.10.6.30

Entity	Name	Details
Citrix Cloud Connector tunnel	Cloud_Connector_DC1-DC2	<p>Local endpoint IP address of the Citrix Cloud Connector tunnel = 10.10.6.30, Remote endpoint IP address of the Citrix Cloud Connector tunnel = 203.0.113.30*</p> <p>GRE Tunnel Details</p> <p>Name = Cloud_Connector_DC1-DC2</p> <p>IPSec Profile Details</p> <p>Name = Cloud_Connector_DC1-DC2, Encryption algorithm = AES, Hash algorithm = HMAC SHA1</p>
Policy based Route	CBC_DC1_DC2_PBR	<p>Source IP range = Subnet in datacenter2 = 10.20.20.0-10.20.20.255, Destination IP range = Subnet in datacenter1 = 10.102.147.0-10.102.147.255, Next hop type = IP Tunnel, IP tunnel name = CBC_DC1_DC2</p>

*These should be public IP addresses.

Following is the traffic flow in the Citrix Cloud Connector tunnel:

1. Client CL1 sends a request to server S1.
2. The request reaches the Citrix virtual appliance NS_VPX_CB_4000/5000-1 running on Citrix SD-WAN WANOP appliance CB_4000/5000-1.
3. NS_VPX_CB_4000/5000-1 forwards the packet to one of the SD-WAN WANOP instances running on the Citrix SD-WAN WANOP appliance CB_4000/5000-1 for WAN optimization. After processing the packet, the SD-WAN WANOP instance returns the packet to NS_VPX_CB_4000/5000-1.
4. The request packet matches the condition specified in PBR entity CBC_DC1_DC2_PBR (configured in NS_VPX_CB_4000/5000-1), because the source IP address and the destination IP address of the request packet belong to the source IP range and destination IP range, respectively, set in CBC_DC1_DC2_PBR.

5. Because the tunnel CBC_DC1_DC2_PBR is bound to CBC_DC1_DC2_PBR, the appliance prepares the packet to be sent across the Cloud_Connector_DC1-DC2 tunnel.
6. NS_VPX_CB_4000/5000-1 uses the GRE protocol to encapsulate each of the request packets by adding a GRE header and a GRE IP header to the packet. In the GRE IP header, the destination IP address is the address of the cloud connector tunnel (Cloud_Connector_DC1-DC2) end point in datacenter DC2.
7. For Cloud Connector tunnel Cloud_Connector_DC1-DC2, NS_VPX_CB_4000/5000-1 checks the stored IPsec security association (SA) parameters for processing outbound packets, as agreed between NS_VPX_CB_4000/5000-1 and NS_VPX_CB_4000/5000-2. The IPsec Encapsulating Security Payload (ESP) protocol in NS_VPX_CB_4000/5000-1 uses these SA parameters for outbound packets, to encrypt the payload of the GRE encapsulated packet.
8. The ESP protocol ensures the packet's integrity and confidentiality by using the HMAC hash function and the encryption algorithm specified for the Citrix Cloud Connector tunnel Cloud_Connector_DC1-DC2. The ESP protocol, after encrypting the GRE payload and calculating the HMAC, generates an ESP header and an ESP trailer and inserts them before and at the end of the encrypted GRE payload, respectively.
9. NS_VPX_CB_4000/5000-1 sends the resulting packet NS_VPX_CB_4000/5000-2.
10. NS_VPX_CB_4000/5000-2 checks the stored IPsec security association (SA) parameters for processing inbound packets, as agreed between CB_DC-1 and NS_VPX-AWS for the Cloud Connector tunnel Cloud_Connector_DC1-DC2. The IPsec ESP protocol on NS_VPX_CB_4000/5000-2 uses these SA parameters for inbound packets, and the ESP header of the request packet, to decrypt the packet.
11. NS_VPX_CB_4000/5000-2 then decapsulates the packet by removing the GRE header.
12. NS_VPX_CB_4000/5000-2 forwards the resulting packet to CB_VPX_CB_4000/5000-2, which applies WAN-optimization-related processing to the packet. CB_VPX_CB_4000/5000-2 then returns the resulting packet to NS_VPX_CB_4000/5000-2.
13. The resulting packet is the same one that was received by CB_VPX_CB_4000/5000-2 in step 2. This packet has the destination IP address set to the IP address of server S1. NS_VPX_CB_4000/5000-2 forwards this packet to server S1.
14. S1 processes the request packet and sends out a response packet. The destination IP address in the response packet is the IP address of client CL1, and the source IP address is the IP address of server S1.

Configure cloud connector tunnel between a datacenter and AWS/Azure

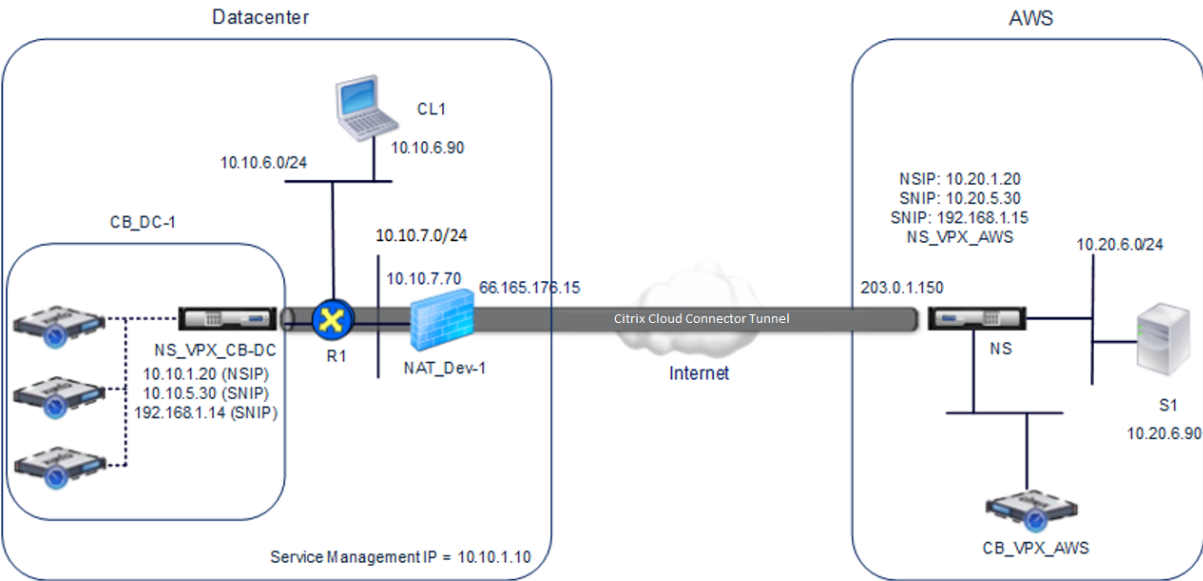
March 12, 2021

You can configure a cloud connector tunnel between a datacenter and AWS, or Azure cloud.

Consider an example in which a Citrix Cloud Connector tunnel is configured between Citrix SD-WAN WANOP appliance CB_DC-1, which is deployed in WCCP/PBR one-arm mode in a datacenter, and AWS cloud. CB_DC-1 is connected to router R1. A NAT device is also connected to R1 for connections between the datacenter and Internet.

Note: The settings in the example would also work for any type of Citrix SD-WAN WANOP deployment. This setting in this example includes policy based routes instead of netbridge for allowing the desired subnet’s traffic to pass through the Citrix Cloud Connector tunnel.

As shown in the following figure, the Citrix Cloud connector tunnel is established between Citrix virtual appliance NS_VPX_CB-DC, running on the Citrix SD-WAN WANOP appliance CB_DC-1, and Citrix virtual appliance NS_VPX-AWS running on AWS cloud. For WAN optimizing the traffic flow over the Citrix Cloud Connector tunnel, NS_VPX_CB-DC is paired to Citrix SD-WAN WANOP instances running on CB_DC-1, and on the AWS side, Citrix SD-WAN WANOP virtual appliance CB_VPX-AWS running on AWS is paired to NS_VPX-AWS.



The following table lists the settings in the datacenter in this example.

Entity	Name	Details
IP address of client CL1		10.10.6.90

Entity	Name	Details
Settings on NAT device		
NAT-Dev-1		
NAT IP address on public side		66.165.176.15 *
NAT IP address on private side		10.10.7.70
Settings on CB_DC-1		
Management service IP address of CB_DC-1		10.10.1.10
Settings on NS_VPX_CB-DC running on CB_DC-1		
The NSIP address		10.10.1.20
SNIP address		10.10.5.30
IPSec profile	CBC_DC_AWS_IPSec_Profile	IKE version = v2, Encryption algorithm = AES, Hash algorithm = HMAC SHA1
Cloud Connector tunnel	CBC_DC_AWS	Local endpoint IP address of the Cloud Connector tunnel = 10.10.5.30, Remote endpoint IP address of the Cloud Connector = Public EIP address mapped to Cloud Connector endpoint address (SNIP) on NS_VPX-AWS on AWS = 203.0.1.150*, Tunnel protocol = GRE and IPSEC, IPSec profile name = CBC_DC_AWS_IPSec_Profile
Policy based route	CBC_DC_AWS_PBR	Source IP range = Subnet in the datacenter = 10.10.6.0-10.10.6.255, Destination IP range = Subnet in AWS = 10.20.6.0-10.20.6.255, Next hop type = IP Tunnel, IP tunnel name = CBC_DC_AWS

*These should be public IP addresses.

The following table lists the settings on AWS cloud in this example.

Entity	Name	Details
IP address of server S1		10.20.6.90
Settings on NS_VPX-AWS		
NSIP address		10.20.1.20
Public EIP address mapped to the NSIP address		203.0.1.120*
SNIP address		10.20.5.30
Public EIP address mapped to the SNIP address		203.0.1.150*
IPSec profile	CBC_DC_AWS_IPSec_Profile	
IKE version = v2, Encryption algorithm = AES, Hash algorithm = HMAC SHA1		
Cloud Connector tunnel	CBC_DC_AWS	Local endpoint IP address of the Cloud Connector tunnel = 10.20.5.30, Remote endpoint IP address of the Cloud Connector tunnel = Public NAT IP address of NAT device NAT-Dev-1 in the datacenter = 66.165.176.15*, Tunnel protocol = GRE and IPSEC, IPSec profile name = CBC_DC_AWS_IPSec_Profile
Policy based route	CBC_DC_AWS_PBR	Source IP range = Subnet in the AWS = 10.20.6.0-10.20.6.255, Destination IP range = Subnet in datacenter = 10.10.6.0-10.10.6.255, Next hop type = IP Tunnel, IP tunnel name = CBC_DC_AWS

*These should be public IP addresses.

Both NS_VPX_CB-DC, on CB_DC-1, and NS_VPX-AWS function in L3 mode. They enable communication between private networks in the datacenter and AWS cloud. NS_VPX_CB-DC and NS_VPX-AWS enable communication between client CL1 in the datacenter and server S1 in the AWS cloud through

the Cloud Connector tunnel. Client CL1 and server S1 are on different private networks.

Note: AWS does not support L2 mode. Therefore, it is necessary to have only L3 mode enabled on both the endpoints.

For proper communication between CL1 and S1, L3 mode is enabled on NS_VPX_CB-DC and NS_VPX-AWS, and routes are configured as follows:

- R1 has a route for reaching S1 through NS_VPX_CB-DC.
- NS_VPX_CB-DC has a route for reaching NS_VPX-AWS through R1.
- S1 should have a route reaching CL1 through NS_VPX-AWS.
- NS_VPX-AWS has a route for reaching NS_VPX_CB-DC through an upstream router.

The following are the routes configured on various network devices in the datacenter for the Cloud Connector tunnel to work properly:

Routes	Network	Gateway
Routes on router R1		
Route for reaching server S1	10.20.6.X/24	Tunnel endpoint SNIP address of NS_VPX_CB-DC = 10.10.5.30
Route for reaching remote end point of the Cloud Connector tunnel	EIP address mapped to the Cloud connector SNIP address of NS_VPX-AWS = 203.0.1.50	Private IP address of the NAT device = 10.10.7.70
Routes on NS_VPX_CB-DC		
Route for reaching NS_VPX-AWS	EIP address mapped to the Cloud connector SNIP address of NS_VPX-AWS = 203.0.1.50	IP address of R1 = 10.10.5.1

The following are the routes configured on various network devices on AWS cloud for the Cloud Connector tunnel to work properly:

Routes	Network	Gateway
Routes on server S1		
Route for reaching client CL1	10.10.6.X/24	Tunnel endpoint SNIP address of NS_VPX-AWS = 10.10.6.1
Routes on Citrix virtual appliance NS_VPX-AWS		

Routes	Network	Gateway
Route for reaching NS_VPX_CB-DC	Public IP address of NAT_Dev-1 in the datacenter = 66.165.176.15*	IP address of the upstream router on AWS

Following is the traffic flow of a request packet from Client CL1 in the Cloud Connector tunnel:

1. Client CL1 sends a request to server S1.
2. The request reaches the Citrix virtual appliance NS_VPX_CB-DC running on Citrix SD-WAN WANOP appliance CB_DC-1.
3. NS_VPX_CB-DC forwards the packet to one of the Citrix SD-WAN WANOP instances running on the Citrix SD-WAN WANOP appliance CB_DC-1 for WAN optimization. After processing the packet, the Citrix SD-WAN WANOP instance returns the packet to NS_VPX_CB-DC.
4. The request packet matches the condition specified in PBR entity CBC_DC_AWS_PBR (configured in NS_VPX_CB-DC), because the source IP address and the destination IP address of the request packet belong to the source IP range and destination IP range, respectively, set in CBC_DC_AWS_PBR.
5. Because the Cloud connector tunnel CBC_DC_AWS is bound to CBC_DC_AWS_PBR, the appliance prepares the packet to be sent across the CBC_DC_AWS tunnel.
6. NS_VPX_CB-DC uses the GRE protocol to encapsulate each of the request packets by adding a GRE header and a GRE IP header to the packet. The GRE IP header has the destination IP address set to the IP address of the Cloud connector tunnel (CBC_DC-AWS) end point on AWS side.
7. For Cloud Connector tunnel CBC_DC-AWS, NS_VPX_CB-DC checks the stored IPsec security association (SA) parameters for processing outbound packets, as agreed between NS_VPX_CB-DC and NS_VPX-AWS. The IPsec Encapsulating Security Payload (ESP) protocol in NS_VPX_CB-DC uses these SA parameters for outbound packets, to encrypt the payload of the GRE encapsulated packet.
8. The ESP protocol ensures the packet's integrity and confidentiality by using the HMAC hash function and the encryption algorithm specified for the Cloud Connector tunnel CBC_DC-AWS. The ESP protocol, after encrypting the GRE payload and calculating the HMAC, generates an ESP header and an ESP trailer and inserts them before and at the end of the encrypted GRE payload, respectively.
9. NS_VPX_CB-DC sends the resulting packet to NS_VPX-AWS.
10. NS_VPX-AWS checks the stored IPsec security association (SA) parameters for processing inbound packets, as agreed between CB_DC-1 and NS_VPX-AWS for the Cloud Connector tunnel

CBC_DC-AWS. The IPSec ESP protocol on NS_VPX-AWS uses these SA parameters for inbound packets, and the ESP header of the request packet, to decrypt the packet.

11. NS_VPX-AWS then decapsulates the packet by removing the GRE header.
12. NS_VPX-AWS forwards the resulting packet to CB_VPX-AWS, which applies WAN optimization related processing to the packet. CB_VPX-AWS then returns the resulting packet to NS_VPX-AWS.
13. The resulting packet is the same packet as the one received by CB_DC-1 in step 2. This packet has the destination IP address set to the IP address of server S1. NS_VPX-AWS forwards this packet to server S1.
14. S1 processes the request packet and sends out a response packet. The destination IP address in the response packet is the IP address of client CL1, and the source IP address is the IP address of server S1.

Office 365 acceleration

March 12, 2021

Citrix SD-WAN WANOP optimizes WAN to provide consistent user experience for business applications across branch offices and remote sites.

Microsoft Office 365 is a software-as-a-service (SaaS) application, which provides the Microsoft's Office suite of enterprise-grade productivity applications. This application is hosted on the cloud and is delivered on demand to users.

The Office 365 acceleration feature allows the branch offices to gain the optimization benefits that Citrix SD-WAN WANOP provides for Microsoft Office 365 application.

Use case


When the WAN segment is considerably slower than the internet segment, and Microsoft's Office 365 servers are closer to the larger office than the branch-office.

Topology

The branch-office Office 365 traffic is sent over the WAN to the main office, and then forwarded to Office 365 servers through the Internet. The segment between the branch office and main office is accelerated.

Note

The segment between the main office and Microsoft Office 365 servers is not accelerated. It is advised that the main office connects to the closest Office 365 server.

**How it works?**

Citrix SD-WAN WANOP SSL acceleration can decrypt and accelerate Office 365 traffic, providing compression. In short, Office 365 branch-office acceleration can be thought of as a special case of RPC-over-HTTPS acceleration.

Procedure

1. Create secure peering between the branch and main office Citrix SD-WAN WANOP appliances.
2. Generate proxy certificates / private key in domain certification authority (CA).
3. Add all required CA's in Citrix SD-WAN WANOP.
 - a) CA, Intermediate CA's, root CA of the Microsoft certificates.
 - b) Proxy certificates/Private keys generated for office 365 URL's.

Note

To avoid security alerts on your browsers, the proxy certificates must be signed by your Windows domain's CA server, which makes it acceptable to any domain user.

4. Create SSL split proxy profile and bind the split proxy to service class (web (internet- secure)).
5. Initiate the office 365 connection and check the Accelerated connections.

Warning

Branch office devices that are not part of the domain will display security warnings unless you install the certificates manually. Firefox users also have to install the certificates manually, since Firefox does not honor the device's certificate store.

Configure Office 365 acceleration

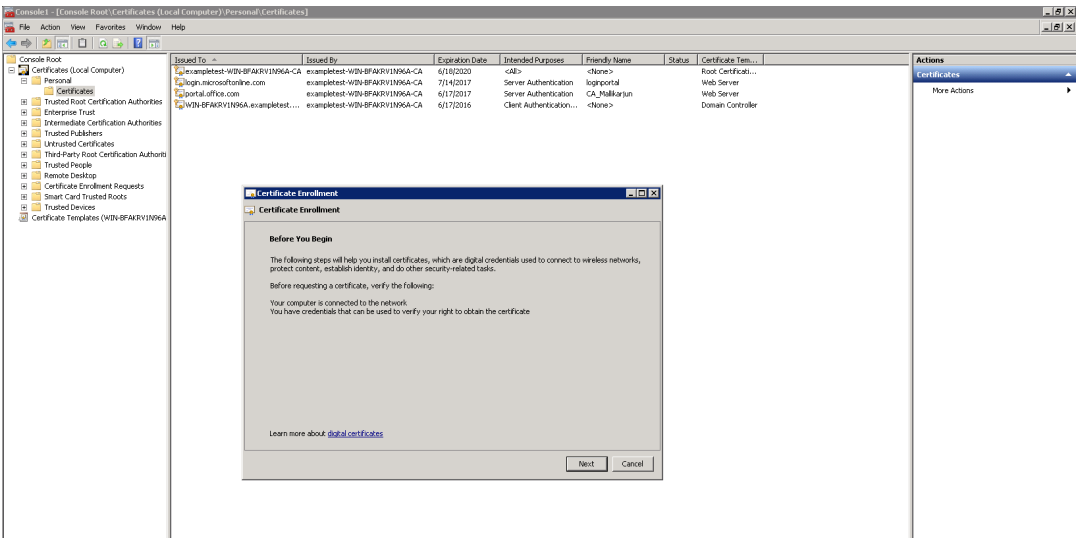
To configure office 365 acceleration:

1. Set up a secure peering relationship between the two Citrix SD-WAN WANOP appliances, as described in [Secure Peering](#)
2. Create a new certificate.

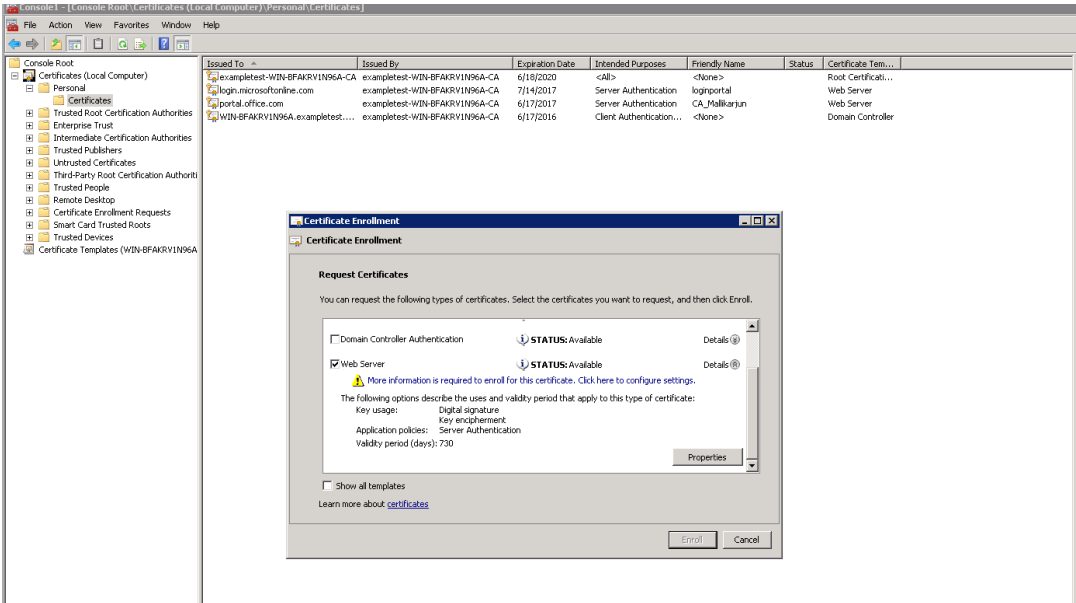
Note

The server-side Citrix SD-WAN WANOP appliance serves as an intermediary between Office 365 and the clients, so these certificates will be signed by the server-side domain controller but it refers to the Office 356 domains.

- a) Log on to the **Certificate Authority Server** for your Windows domain.
- b) If necessary, add the snap-ins for **Certification Authority**, **Certificate Template** and **Certificates**.
- c) Navigate to **Certificate Templates > Web Server Properties > Security** and select all the options.
- d) Navigate to **Certificates > Personal > Certificates (Computer) > All Tasks > Request New Certificate**.



- e) In the **Certificate Enrolment** window, click **Next**.
- f) In the **Select Certificate Enrolment Policy** window, select **Active directory enrolment policy**.
- g) In the **Active Directory Enrolment Policy** window, select **Web Server > Details > Properties**.



- 3. Copy information from Office365 certificates into your new certificates. You will end up with a single certificate from three Office365 certificates. Proceed as follows:
 - a) In a browser, such as Chrome, enter the url - <https://login.microsoftonline.com>.

Note


Do not log in.

- b) Click the padlock icon on the URL bar and select **Connection > Certificate Information > Details**.


Note

These instructions are for the Chrome browser; the procedure is the same for other browsers also.


- c) Click **Subject Alternative Name**, this will reveal a list of DNS names such as “login.microsoftonline.com.” Copy the information in the text box below it.



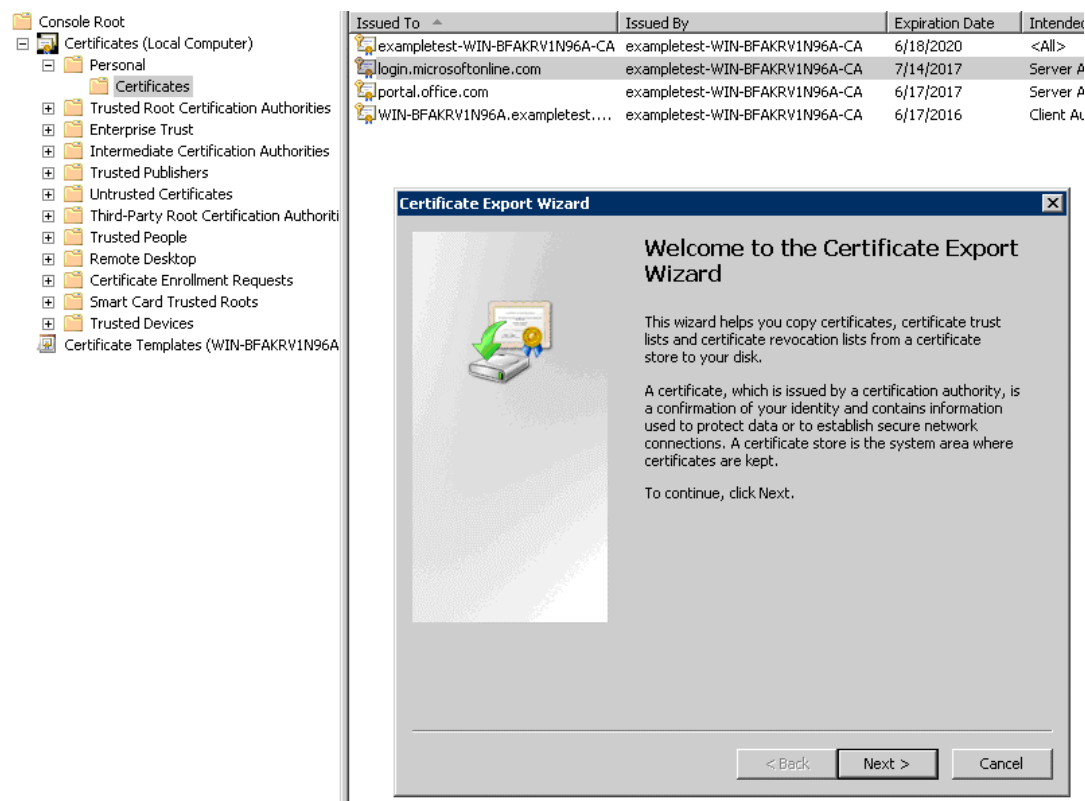
- d) Return to your new certificate's **Certificates Properties** window. Add the alternative names in the **Value** field with **Type** as **DNS** to match each alternative name in the Microsoft certificate.



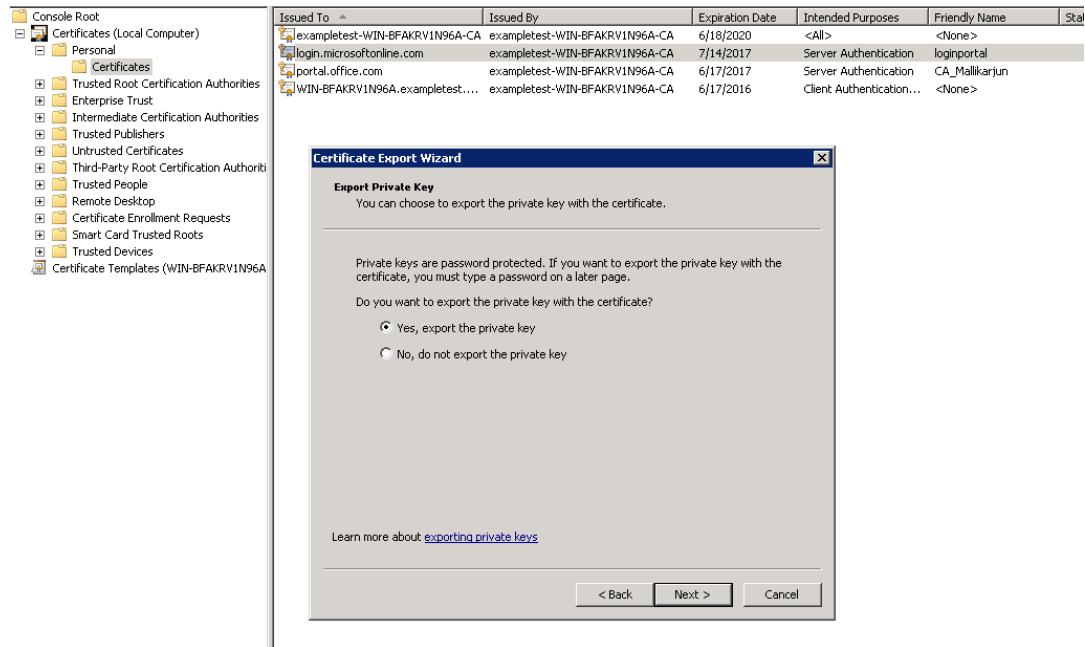
- e) Repeat the process of discovering Subject Alternate Names and adding them to your certificate for <https://outlook.office365.com>, <https://portal.office.com>, <https://office.live.com>, and <https://sharepoint.com> (the SharePoint URL is customer-specific).
- f) Create a Common Name for your new certificate. The example above shows a common name as “Office365 proxy.”



- g) In the **Private Key** tab, select **Make private key exportable**.
 - h) Click **OK**, **Enroll**, and **Finish**.
4. Export the certificate.
 - a) Under **Certificates > Personal > Certificates**, select the above created proxy certificate, and then select **All Tasks > Export**.



- b) The **Certificate Export Wizard** appears. Click **Next**.
- c) In **Export Private Key**, select the option **Yes, export the private key** and click **Next**.




- d) Retain the default values for the export file format.
- e) Type and confirm the password, export the private key, and save the certificate as *login-*


portal.pfx.

5. Export your certificates.

- a) In the **Certificate Export Wizard**, click **Next**. In **Export Private Key**, select the option **No, do not export the private key**. Click **Next**.



- b) Retain the default values for the export file format.
- c) Type and confirm the password, and export the private key and certificate, saving the file to a file to a file name such as office365_keys.pfx.
- ## 6. Download the public keys of the root CA and Intermediate CAs of the Microsoft certificates.
- a) From the browser, navigate to <https://login.microsoftonline.com>. Click the padlock icon in the browser. Navigate to **Connection > Certificate Information > Certification Path**.
- b) Select the root certificate (the one at the top of the list), and then click **View Certificate > Details > Copy to File**. The **Certificate Export Wizard** appears. Click **Next**.



c) Enter the file name and save the file.

Note

Alternatively, you can use Wireshark or OpenSSL to get the root and intermediate CA names and get the certificates from 'AUTHENTIC' source (for example, Windows SSL store).

d) Repeat step 6 to save the root and intermediate CA's of the following domains:

- i. login.microsoftonline.com
- ii. portal.office.com
- iii. outlook.office365.com
- iv. *.sharepoint.com
- v. office.live.com

7. Add all the Office 365 server CA's, proxy certificate/key pairs, and private keys to the server-side Citrix SD-WAN WANOP appliance. The CA's are added using the **CA Certificates** tab on the **Certificates and Keys** page. Certificates and certificate/key pairs are added on the **Certificate/Key Pairs** tab.

The screenshot shows the 'Configuration' tab in the Citrix SD-WAN WANOP 11.0 interface. The left sidebar has 'Certificate and Keys' highlighted. The main content area shows the 'CA Certificates' section, which is also highlighted. Below the section title, there are tabs for 'CA Certificates' and 'Certificate Key Pairs'. The 'Add' button is highlighted. Below the buttons is a table of CA Certificates.

Name	Expiration Date
Symantec_root_ca	Oct 30 23:59:59 2023 GMT
Verisign	Jul 16 23:59:59 2036 GMT
ca	Feb 25 01:39:42 2032 GMT
login_Portal_root_ca	Feb 1 23:59:59 2017 GMT
office_Portal_root_ca	Apr 22 19:47:55 2016 GMT

The screenshot shows the 'Configuration' tab in the Citrix SD-WAN WANOP 11.0 interface. The left sidebar has 'Certificate and Keys' highlighted. The main content area shows the 'Certificate Key Pairs' section, which is also highlighted. Below the section title, there are tabs for 'CA Certificates' and 'Certificate Key Pairs'. The 'Add' button is highlighted. Below the buttons is a table of Certificate Key Pairs.

Certificate Key Pair Names	Expiration Date
login_Portal_pri	2017-07-14 09:07:33
office_portal_private_key	2017-06-17 12:09:27
pri	2033-07-18 20:01:18

8. Create an SSL split-proxy profile and bind the split proxy to the Web (Internet-Secure) service class.
 - a) Navigate to **Configuration > Secure Acceleration > SSL Profile > Add Profile**.
 - b) Enter the profile name of your choice. Select **Profile Enabled**, **Parse Subject Alternative Names**, and **Split Proxy**.
 - c) Under **Server-Side Proxy Configuration > Verification Store**, select **Use all configured CA stores**.
 - d) Under **Client-side Proxy Configuration > Certificate/Private Key**, select the cert/private key pair you created and exported previously (the one shown in the example as loginportal.pfx). Select **Build Certificate Chain**. Select the CA associated with the certificate/key pair under **Certificate Chain Store**.

The screenshot shows the 'SSL Profile' configuration page in the Citrix SD-WAN WANOP 11.0 management console. The page is divided into three main sections: 'SSL Profile', 'Server-Side Proxy Configuration', and 'Client-Side Proxy Configuration'. At the top, there is a 'Back' button. The 'SSL Profile' section includes a 'Profile Name' field set to 'Office365_Profile', checkboxes for 'Profile Enabled' and 'Parse Subject Alternative Names', a 'Proxy Type' dropdown set to 'Split', an 'Enable Exclude List' checkbox, and a 'Certificate Verification' dropdown set to 'None - allow all requests'. The 'Server-Side Proxy Configuration' section includes a 'Verification Store' dropdown set to 'Use all configured CA stores', an 'Authentication Required' checkbox, a 'Protocol Version' dropdown set to 'SSL Version 2.3 or TLS 1.0', a 'Cipher Specification' field set to 'TADH:HIGH:MEDIUM:85STRENGTH', and a 'Renegotiation Type' dropdown set to 'Old Style Renegotiation Disabled'. The 'Client-Side Proxy Configuration' section includes a 'Certificate/Private Key' dropdown set to 'single_cert_private', checkboxes for 'Disable Session Re-use' and 'Build Certificate Chain', a 'Certificate Chain Store' dropdown set to 'Use all configured CA stores', a 'Protocol Version' dropdown set to 'SSL Version 2.3 or TLS 1.0', a 'Cipher Specification' field set to 'TADH:HIGH:MEDIUM:85STRENGTH', and a 'Renegotiation Type' dropdown set to 'Old Style Renegotiation Disabled'. At the bottom, there are 'Create' and 'Close' buttons.

Back

SSL Profile

Profile Name*
Office365_Profile

☒ Profile Enabled

☒ Parse Subject Alternative Names

Proxy Type
☒ Split ☐ Transparent

☒ Enable Exclude List

Certificate Verification*
None - allow all requests

Server-Side Proxy Configuration

Verification Store
Use all configured CA stores

☐ Authentication Required

Protocol Version*
SSL Version 2.3 or TLS 1.0

Cipher Specification*
TADH:HIGH:MEDIUM:85STRENGTH

Renegotiation Type*
Old Style Renegotiation Disabled

Client-Side Proxy Configuration

Certificate/Private Key*
single_cert_private

☒ Disable Session Re-use

☒ Build Certificate Chain

Certificate Chain Store
Use all configured CA stores

Protocol Version*
SSL Version 2.3 or TLS 1.0


Cipher Specification*
TADH:HIGH:MEDIUM:85STRENGTH

Renegotiation Type*
Old Style Renegotiation Disabled

Create Close

9. Bind the created SSL profile to the Internet (Web-Secure) service class. Navigating to **Configure > Optimization Rules > Service Classes** and add the SSL profile to the SSL profile list.
10. Enable acceleration and disk-based compression for the **Internet (Web-Secure)** service class.
11. Initiate an Office 365 session from your browser.

The connection is accelerated. In the browser, the certificate should display your root CA, not the actual Office 365 certificate, as the server-side appliance's CA certificate.



12. On the appliance **Monitoring > Connections** page, verify that the Office 365 connections are compressed and are receiving SSL acceleration.

The screenshot shows the 'Accelerated Connections' page in the Citrix SD-WAN WANOP interface. A left sidebar contains navigation links like 'Optimisation', 'Connections', 'Compression', etc. The main area has a table with columns: 'Details', 'Initiator', 'Responder', 'Duration', 'Idle', 'Bytes Transferred', 'Compression Ratio/Type', 'Bandwidth Savings (%)', and 'SSL Proxy'. The table lists several connections, all of which are compressed (1.0 to 1 (Disk)) and have SSL Proxy set to 'True'.

Details	Initiator	Responder	Duration	Idle	Bytes Transferred	Compression Ratio/Type	Bandwidth Savings (%)	SSL Proxy
	172.16.139.221 : 50454	132.245.163.178 : 443	3m 31s	0m 11s	6.67 KB	1.1 to 1 (Disk)	29.6	True
	172.16.139.221 : 50453	132.245.163.178 : 443	3m 32s	0m 31s	6.19 KB	1.2 to 1 (Disk)	35.9	True
	172.16.139.221 : 50456	191.236.88.160 : 443	2m 2s	0m 53s	6.08 KB	1.6 to 1 (Disk)	46.8	True
	172.16.139.221 : 50459	132.245.165.130 : 443	1m 33s	1m 32s	3.15 KB	1.9 to 1 (Disk)	27.1	True
	172.16.139.216 : 11745	172.229.161.125 : 443	3m 25s	3m 4s	54 bytes	1.0 to 1 (Disk)	0	True
	172.16.139.216 : 11744	132.245.164.34 : 443	3m 25s	3m 21s	0 bytes	1.0 to 1 (Disk)	0	True
	172.16.139.216 : 11747	132.245.164.226 : 443	3m 24s	3m 21s	0 bytes	1.0 to 1 (Disk)	0	True


Note

Firefox does not accept the device's certificates by default, but has its own certificate store. Therefore, credentials accepted in the normal Windows domain behavior by other browsers, and by the device as a whole, must be installed manually into Firefox. To install certificates into Firefox, follow the procedure in the section, Installing certificates to Firefox.

Install the certificates to Firefox

To Install the server-side appliance's proxy certificate to the Firefox certificate store:

1. In the Firefox browser navigating to **Options > Advanced > Certificate > View Certificates > Authorities > Import**.
2. Upload the local CA proxy certificate, select all the options in the **Downloading Certificate** wizard and click **OK**.



SCPS support

March 12, 2021

Citrix SD-WAN WANOP supports the SCPS (Space Communications Protocol Standard) TCP variant. SCPS is widely used for satellite communication.

See <http://www.scps.org> for general SCPS information.

SCPS is a TCP variant used in satellite communication and similar applications. The appliance can accelerate SCPS connections if the **SCPS** option is selected on the Configuration: Tuning page.

The main practical difference between SCPS and the default appliance behavior is that SCPS-style “selective negative acknowledgements” (SNACKs) are used instead of standard selective acknowledgements (SACKs). These two methods of enhancing data retransmissions are mutually exclusive, so if the appliance on one end of the connection has SCPS enabled and one does not, retransmission performance suffers. This condition also causes an “SCPS Mode Mismatch” alert.

If you must mix SCPS-enabled appliances with non-SCPS-enabled appliances, deploy them in such a way that mismatches do not occur. You can either use IP-based service class rules or arrange the deployment so that each path has matching appliances.

Secure traffic acceleration

March 12, 2021

Secure traffic acceleration is achieved by secure peering. Several advanced functions require that the Citrix SD-WAN WANOP appliances at the two ends of the link establish a *secure peer relationship* with each other, setting up an SSL signaling tunnel (also called a *signaling connection*). These functions are SSL compression, signed CIFS support, and encrypted MAPI support.

When secure peering is enabled, compression is automatically disabled for all partner appliances (and computers running the Citrix SD-WAN WANOP Plug-in) that have not established a secure peer relationship with the local appliance.

To establish a secure peer relationship, you have to generate security keys and certificates, and configure a securing signaling tunnel between the appliances. Before configuring the tunnel, order a crypto license from Citrix.

Secure peering

March 12, 2021

When an appliance has secure peering enabled, connections with a partner for which it does not have a secure peer relationship are not encrypted or compressed, though TCP flow-control acceleration is still available. Compression is disabled to ensure that data stored in compression history from secured partners cannot be shared with unsecured partners.

When the appliance at one end of a connection detects that the other appliance has secure peering enabled, it attempts to open an SSL signaling tunnel. If the two appliances successfully authenticate each other over this tunnel, they have a secure peering relationship. All accelerated connections between the two appliances are encrypted, and compression is enabled.

Note

An appliance with secure peering enabled does not compress connections to unsecured partners,

using the same appliance successfully with a mix of secured and unsecured partners is difficult. Keep this point in mind when designing your accelerated network.

A keystore password is required to access the security parameters. This keystore password is different from the administrator's password, to allow security administration to be separated from other tasks. If the keystore password is reset, all existing encrypted data and private keys are lost.

To protect data even if the appliance is stolen, the keystore password must be reentered every time the appliance is restarted. Until this is done, secure peering and compression are disabled.

Generate security keys and certificates

Citrix SD-WAN WANOP products are shipped without the required keys and certificates for the SSL signaling tunnel. You must generate them yourself. You can generate keys and certificates through your normal process for generating credentials, or with the “openssl” package from <http://www.openssl.org>.

For testing purposes, you can generate and use a self-signed X509 certificate based on a private key (which you also generate). In production, use certificates that refer to a trusted certifying authority. The following example calls openssl from the command line on a PC to generate a private key (my.key) and self-signed certificate (my.crt):

```
1 pre codeblock
2 # Generate a 2048-bit private key
3 openssl genrsa -out my.key 2048
4 # Now create a Certificate Signing Request
5 openssl req -new -key my.key -out my.csr
6 # Finally, create a self-signed certificate with a 365-day expiration
7 openssl x509 -req -days 365 -in my.csr -signkey my.key -out my.crt
8 <!--NeedCopy-->
```

For production use, consult your organization's security policies.

Configure secure peering

There are two ways to establish secure peering:


1. Using credentials generated by the appliances.
2. Using credentials you provide yourself.

Because an appliance with secure peering enabled will only compress connections with partner appliances with which it has a secure peering relationship, this procedure should be applied at the same time to all your appliances.

To prepare the appliances for secure peering:

Perform the following procedure on each appliance in your network.

1. Install a crypto license on the appliance. Without a crypto license, secure acceleration is not available.
 - a) If you have not done so already, acquire crypto licenses from Citrix.
 - b) If you are using a network license server, go to the **Configuration > Appliance Settings > Licensing**. On the **Add License** section click **edit**, and select the Remote license server and set Crypto License On.
 - c) If you are using local licensing, go to **Configuration > Appliance Settings > Licensing**. In the **Add License** page, click the Local license server option, and click **Add** to upload a local crypto license.
 - d) Verify successful license installation on the **Configuration > Appliance Settings > Licensing** page. Under Licensing Information, a crypto license should be shown as active and with an expiration date in the future.
2. Go to the **Configuration > Secure Acceleration** page. If the page has a button labeled Secure, click it.



3. If you are taken to a Keystore Settings screen automatically, do the following:
 - a) Enter a keystore password twice and click Save.
 - b) When the screen updates to show the Secure Peering Certificates and Keys section, click Enable Secure Peering and CA Certificate, then click Save.

- c) Skip to Step 6.
- 4. If you were not taken to the Keystore Settings screen automatically, click the pencil icon under **Secure Peering**, then click the pencil icon under **Keystore Settings**. Open on the Keystore Status pulldown menu, and enter a keystore password twice. Click **Save**.
- 5. Enable secure peering by going to the **Configuration > Secure Acceleration** page and clicking the **Enable** button. Ignore any warnings at this stage. This setting enables secure peering when the required additional configuration is complete.
- 6. Enable encryption of compression history by going to **Configuration > Secure Acceleration User Data Store** and clicking the pencil icon. Click **Enable Disk Encryption**, then click **Save**. User data store encryption prevents unauthorized reading of the disk based compression history, in case the appliance is stolen or returned to the factory. The security of disk data encryption relies on the keystore password. This feature uses AES-256 encryption. (Disk data encryption does not encrypt the entire disk, just the compression history.)
- 7. If you are using appliance-generated credentials, skip to the next step. If you are using your own credentials, do the following:
 - a) Go to **Configuration > Secure Acceleration** and click the pencil icon under Secure Peering, then click the pencil icon under **Secure Peering Certificates and Keys**. Click **Enable Secure Peering and Certificate Configuration > CA Certificate**. The credential specification fields appear.
 - b) Under **Certificate/Key Pair Name**, click the + icon and upload or paste the cert/key pair for this appliance. If required by the credentials, also enter the key password or file password. Click **Create**.
 - c) Under **CA Certificate Store Name**, click the + icon and upload or paste the CA certificate for this appliance.
 - d) Keep the default values for the Certificate Verification and SSL Cipher Specification fields unless your organization requires otherwise.
 - e) Click **Save**.

Secure Peering

Keystore Settings

Keystore Status
Opened

Secure Peering Certificate and Keys

Secure communications with the CloudBridge partner appliance requires that you generate OpenSSL credentials, including CA Certificate and a Certificate/Key pair, and select a verification method. You can optionally change the OpenSSL cipher specification. If PrivateCA is selected, certificates and keys are generated automatically.

☒ Enable Secure Peering

Certificate Configuration

☐ Private CA ☒ CA Certificate

Certificate/Key Pair Name
private_172_16_0_243

CA Certificate Store Name
PrivateRootCA

Certificate Verification
Signature/Expiration

SSL Cipher Specification
[ADH:AECDH:MD5:DH@STRENGTH]

☐ Edit Cipher Specification

Save Cancel

8. Repeat for the rest of your appliances.
9. If you are using credentials that you provided yourself, secure peering configuration is complete.
10. If you are using appliance-generated credentials, perform the following procedure.

To use secure peering with appliance-generated credentials:

1. Use the “Prepare the appliances for securing peering” procedure, above, to prepare your appliances for this procedure.
2. On one datacenter appliance, go to **Configuration > Secure Acceleration** and click the **Enable** button, if present, to enable secure peering.
3. Click the pencil icon under Secure Peering. The keystore should be open. If it isn’t, open it now.
4. Click the pencil icon under **Secure Peering Certificate and Keys**. Click the **Enable Secure Peering and Private CA** options, then click **Save**. This will generate a local self-signed CA certificate and a local certificate-key pair.
5. Click **+** under **Connected Peers**. Enter the IP address, administrator’s user name, and administrator’s password for one of your remote appliances and click **Connect**. This issues a CA certificate and certificate-key pair for the remote appliance, and copies it to the remote appliance.

Note

For SD-WAN WANOP appliances, the IP address could be the IP address of any of the interface where web access is enabled. For SD-WAN PE appliances, the IP address is the management IP address.

6. Repeat this process for your other remote appliances.

7. On the datacenter appliance, verify connectivity by going to **Monitoring > Partners and Plugins > Secure Partners**. For each remote appliance, the content of the Secure field should be True, and the Connection Status should be Connected Available.

CIFS, SMB2, and MAPI

March 12, 2021


Windows is one of the common operating systems deployed on the network. The Windows operating system supports distributed resources shared across locations. For example, you can make resources in your datacenter accessible from various branch offices. For access over the network, Windows uses Common Internet File System (CIFS) protocol for accessing shared files, and Messaging Application Programming Interface (MAPI) protocols for accessing email through Microsoft Outlook. That is, Windows uses CIFS protocol for CIFS-based (Windows and Samba) file transfer and directory browsing, and Microsoft Outlook uses the MAPI protocol to access Outlook data.

You can use a Citrix SD-WAN WANOP appliance to optimize the CIFS, Server Message Block version 2 (SMB2), and MAPI connections over the network.

In addition to supporting the Windows operating system, Citrix SD-WAN WANOP appliances support CIFS and SMB2 on NetApp and Hitachi storage systems.

The flow chart below shows the complete procedure to configure a Citrix SD-WAN WANOP appliance for optimizing CIFS, SMB2, and MAPI traffic.

Configuring a Citrix SD-WAN WANOP appliance for optimizing CIFS, SMB2, and MAPI traffic



Configure Citrix SD-WAN WANOP appliance to optimize secure Windows traffic

March 12, 2021

You must add the Citrix SD-WAN WANOP appliance to the Windows security infrastructure before you can optimize the signed Windows file system and encrypted MAPI Outlook/Exchange traffic.

As a result of enhancements to the Windows security system in the recent Windows releases, clients and servers secure the traffic by authenticating and encrypting data. This requires the Citrix SD-WAN WANOP appliance be a trusted member of the Windows security infrastructure before it can optimize signed Windows file system and encrypted MAPI Outlook/Exchange traffic.

After you add the appliance to the Windows security infrastructure, the appliance has the following capabilities:

- Acceleration of fileserver traffic for Microsoft Windows servers, NetApp servers, and Hitachi HNAS by using signed SMB and signed SMB2 protocol.
- Acceleration of Microsoft Exchange server traffic when it is accessed by Outlook clients using encrypted MAPI or RPC over HTTPS.

How Citrix SD-WAN WANOP appliance works in a Windows security system

Joining the appliance to a Windows domain requires administrator credentials. When it joins the Windows domain, the appliance becomes a trusted member of the domain. This allows the appliance to be declared a member of the domain's security infrastructure.

After the appliance has become a part of the Windows security infrastructure, users have to be authenticated before they can access resources. To avoid the difficulty of configuring a large number of users in the domain, you can delegate the authentication responsibility to a delegate user.


You create a delegate user in the active directory. This user is similar to a normal user, but with special privileges. After creating the delegate user, you must configure this user on the Citrix SD-WAN WANOP appliance. The appliance uses the delegate user to authenticate on behalf of users when they access authenticated and encrypted data streams using Windows protocols, such as CIFS and MAPI.

For accelerating CIFS and MAPI traffic, the standard Windows delegation mechanism allows you to limit security delegation to the relevant services. This constrained delegation has been available since the release of Windows Server 2003.

After becoming a part of the domain, the appliance accelerates the secure Windows traffic. A data-center appliance that joins a Windows domain must have a secure peer relationship with the remote

appliance or Citrix SD-WAN WANOP Plug-in, but only the datacenter appliance joins the Windows domain. For purposes of CIFS or MAPI acceleration, the remote appliance acts as a slave to the datacenter appliance, being controlled over the secure SSL tunnel between the two. Therefore, the delegate user credentials do not leave the datacenter.

The following figure shows a sample topology diagram for this setup.



In the above figure, a branch office client accesses resources of the datacenter. The branch office client, being in another domain, uses NTLM authentication as a part of the Windows security system. As with all accelerated connections between two Citrix SD-WAN WANOP appliances in a secure peer relationship, the CIFS or MAPI connections and NTLM authentications over the WAN are encrypted. Depending on the Windows domain controller version, the user request from the datacenter Citrix SD-WAN WANOP appliance is authenticated using NTLM or Kerberos authentication protocol. After the domain authenticates the user, subsequent access requests to the Exchange server and file servers use Kerberos authentication protocol. The Citrix SD-WAN WANOP appliance then optimizes the connections established between the client and the server.

If the appliances do not have a secure peer relationship, or if the datacenter appliance has not successfully joined the domain, the connections use TCP flow-control acceleration, which performs no security operations, compression, or data transformations. The connections between the client and server are established as if the Citrix SD-WAN WANOP appliances were not there.

You can configure different client authentication modes on Windows operating systems. The types of connections that the Citrix SD-WAN WANOP appliance optimizes depend on the client authentication mode that you configure.

The following table list the Windows client-authentication modes on Windows, and the corresponding Citrix SD-WAN WANOP optimizations.

Authentication and Optimization Supported for Windows Operating System

Client Operating System	Client Authentication Mode	Optimization	Comments
Windows XP/Windows Vista/Windows 7/Windows 8	Negotiate Authentication (SPNEGO)	TCP flow-control acceleration, Compression, CIFS protocol acceleration	Default setting used for all Windows releases.
Windows XP/Windows Vista/Windows 7/Windows 8	NTLM only or Kerberos only	TCP flow-control acceleration only	Non-default authentication modes

Note: If you use the NTLM only or Kerberos only client authentication modes, the traffic is not accelerated if it is encrypted.

Requirements to add a Citrix SD-WAN WANOP appliance to the Windows security system

To optimize traffic for secured Windows signed SMB and encrypted MAPI traffic, your Citrix SD-WAN WANOP deployment must meet the following requirements before you add the appliance to the Windows security infrastructure:

- Both the client-side and server-side acceleration appliances must have established a secure peer relationship.
- The appliances must use an NTP server that is closely synchronized to the time on the Windows domain server. Ideally, the appliances and the Windows domain server are all clients of the same NTP server.
- Outlook **must not** be configured for the (non-default) **Kerberos only** or **NTLM only** option. The default (negotiated) option is required for acceleration.
- The client and server can be members of any domain that has two-way trust with the server-side appliance's domain. One-way trust is not supported.
- A Kerberos delegate user must be set up on the domain controller, to be used by the appliance participating in the domain's security infrastructure.
- The DNS server IP addresses for the domain must be configured and reachable on the server-side appliance.
- The domain servers must be fully reachable, with both forward and reverse lookups for all the IP addresses of the domain controllers configured on the DNS servers.
- The server-side Citrix SD-WAN WANOP appliance's host name must be unique. Using the default host name of "hostname" is likely to cause problems.

Note

The Macintosh Outlook client does not use the MAPI (Outlook/Exchange) standard and is not accelerated by this feature.


Add a Citrix SD-WAN WANOP appliance to the Windows security infrastructure

To optimize secure Windows traffic, the Citrix SD-WAN WANOP appliance must be a part of the Windows security system and must authenticate itself with the security system or domain. As shown in the below figure, to make the appliance a part of the Windows security system, you must make the appliance join a domain (using administrative credentials). Additionally, you need to configure a new or existing user as a delegate user by associating CIFS and Exchange services with that user. You then have to configure this delegate user on the Citrix SD-WAN WANOP appliance.

You can use the **Pre Domain Check** utility to find out if there are any issues with joining the appliance to a domain.

Note

The Windows security system uses the Exchange service to manage MAPI connections. Configuring the setup to optimize secure Windows traffic



Join a Citrix SD-WAN WANOP appliance to the Windows domain:

When the appliance joins the domain, it exchanges a shared secret with the domain controller, allowing the appliance to remain part of the domain indefinitely. When joining an appliance to a domain, make sure that you have administrator credentials for the domain controller.

To make sure that the Citrix SD-WAN WANOP appliance optimizes the CIFS and MAPI traffic (including traffic encapsulated as RPC over HTTPS), you must make the appliance part of the domain that the Windows fileserver and Exchange server are a part of. You need to join the server-side appliance to the domain.

Note: The domain administration credentials are not saved on the appliance.

To join a Citrix SD-WAN WANOP appliance to a Windows domain:

1. Navigate to the **Configuration > Secure Acceleration > Windows Domain** tab.
2. Click **Join Windows Domain**.
3. Enter the Windows domain name in the Domain Name field.
4. In the User Name field, enter the user name of the domain controller administrator.

5. In the Password field, specify the domain controller administrator password.
6. If necessary, edit the DNS servers for consistency with the Windows domain.
7. Click **OK**.
8. In the Delegate Users section, add a delegate user, as described in the procedures below.

The screenshot shows the Citrix SD-WAN WANOP Configuration page. The left sidebar contains a navigation menu with options: Appliance Settings, Optimisation Rules, Video Caching, Secure Acceleration (selected), Certificate and Keys, User Data Store, Diagnostics, and Maintenance. The main content area is titled 'Configuration Overview > Secure Acceleration'. It features a green status bar indicating 'SSL Optimization status: ACTIVE' with a 'Disable' button. Below this, the 'Secure Peering' section shows 'Keystore Status: Opened' and 'Secure Peering Status: Enabled'. The 'Windows Domain' section is active, displaying instructions to join the server-side CloudBridge appliance to a domain. It includes input fields for 'Domain Name*' (example.com), 'User Name*' (user), and 'Password*' (masked with asterisks). There is a 'Check Domain Join' button, a 'Leave Domain' checkbox, and a 'DNS Servers*' dropdown menu showing '172.16.0.71'. At the bottom are 'OK' and 'Cancel' buttons.

Configure a delegate user:

After you join the appliance to a Windows domain, you must create a user that the appliance can use to authenticate users with the domain. This user is known as the *delegate user*.

Note: To create a delegate user account, you need administrator access to the Windows domain controller and the appliance. If you do not have administrator access to the Windows domain controller, make sure that an authorized administrator performs the required tasks on the domain controller.

Setting up user authentication by using Kerberos delegation involves two tasks—configuring a delegate user on the domain controller and then adding this user to the Citrix SD-WAN WANOP appliance.

Configure a delegate user on a domain controller:

Before you configure a delegate user on a Citrix SD-WAN WANOP appliance, you must configure a delegate user with the required properties on the domain controller. You can either create a delegate user account or use an existing user account as a delegate user account.

After creating an account or selecting an existing account, enable delegation for this user. You then associate the delegate user with the CIFS and Exchange services, so that the traffic for these services

can be accelerated. After you add this user to the Citrix SD-WAN WANOP appliance, the appliance presents delegated credentials for the services associated with this account.

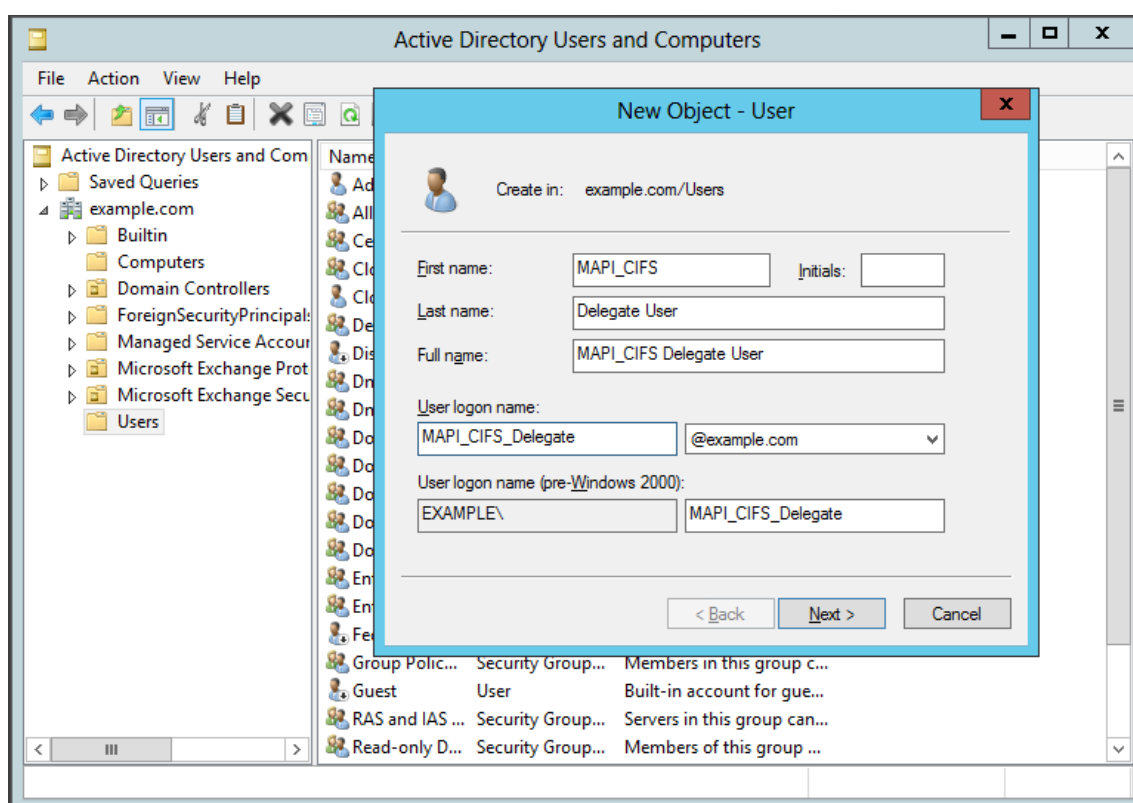
Create a delegate user account:

Create a delegate user account on the Windows domain controller so that the Citrix SD-WAN WANOP appliance can use this account on behalf of the users to authenticate them with the domain controller.

Note: If you want to configure an existing user as a delegate user, skip this procedure.

To create a delegate user account:

1. Log on to the Windows domain controller as an administrator. Make sure that the file server or Exchange server is a member of this domain.
2. From the **Start** menu, open the **Active Directory Users and Computers** Window.
3. Create a delegate user, as shown in the following screen shot:




Enable delegation for a user:

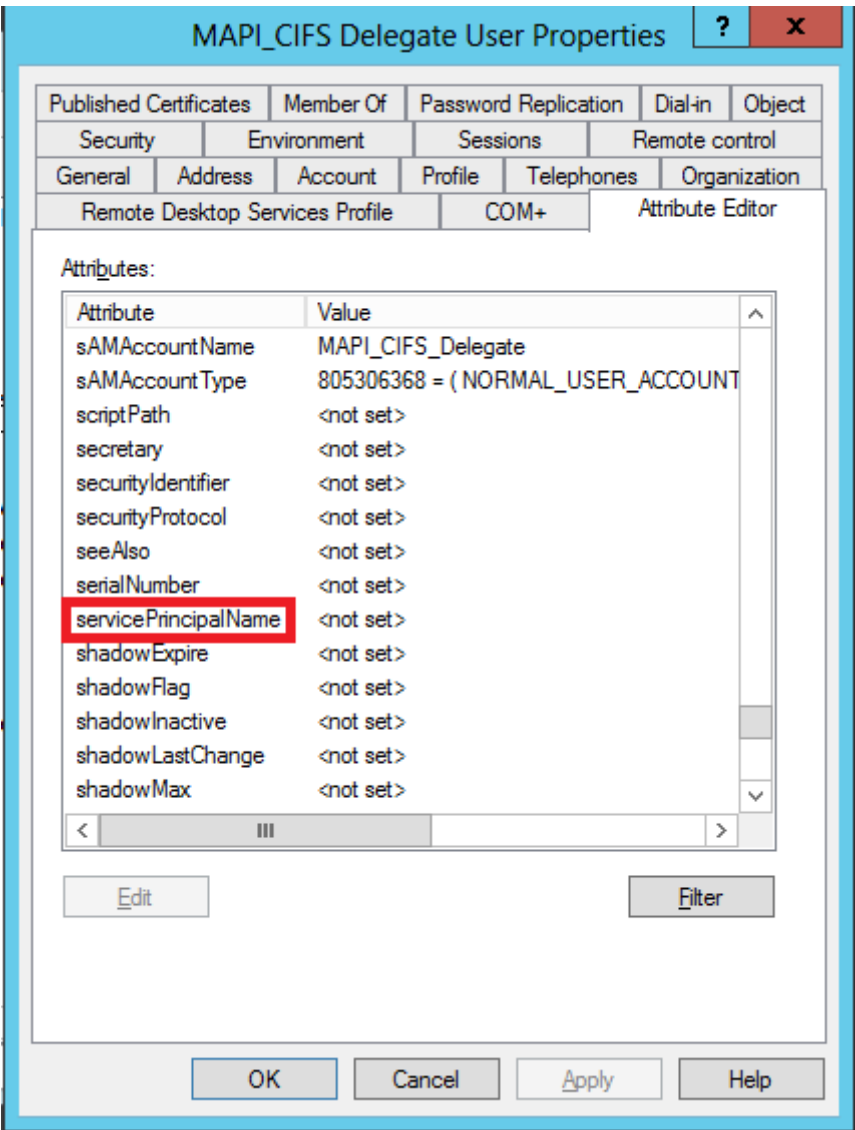
So far, the user you have created is similar to any user you create on the Active Directory server. To enable delegation for the user, you must set the Service Principal Name attribute of the user to *delegate* and associate the delegate user with the required services. This makes the user to have special privileges attached to it and make it a delegate user.

To enable delegation for the user:


1. From the **Start** menu, open the **Active Directory Users and Computers** Window.
2. From the View menu, select **Advanced Features**.
3. Select the **User** node.
4. Right-click the user that you want to make a delegate user.
5. From the shortcut menu, select **Properties** and navigate to the **Attribute Editor** tab, as shown in the following screen shot:



6. From the **Attributes** list, select **servicePrincipalName**, as shown in the following screen shot:



7. Click **Edit**.
8. In the **Multi-valued String Editor** dialog box, in the **Value to add** field, specify **delegate/<User_Name>**, as shown in the following screen shot:



9. Click **Add**.
10. Click **OK**.
11. Click **Apply**.
12. Click **OK**.
13. Open the user's **MAPI-CIFS Delegate User Properties** dialog box and verify that the **Delegation** tab has been added to the dialog box, as shown in the following screen shot:

The screenshot shows the 'MAPI_CIFS Delegate User Properties' dialog box. The 'Delegation' tab is selected and highlighted with a red box. The dialog box contains the following fields and options:

- Organization:** MAPI_CIFS
- Published Certificates:** (empty)
- Member Of:** (empty)
- Password Replication:** (empty)
- Dial-in:** (empty)
- Object:** (empty)
- Security:** (empty)
- Environment:** (empty)
- Sessions:** (empty)
- Remote control:** (empty)
- Remote Desktop Services Profile:** (empty)
- COM+:** (empty)
- Attribute Editor:** (empty)
- General:** (selected)
- Address:** (empty)
- Account:** (empty)
- Profile:** (empty)
- Telephones:** (empty)
- Delegation:** (selected)

The 'MAPI_CIFS Delegate User' section contains the following fields:

- First name:** MAPI_CIFS
- Initials:** (empty)
- Last name:** Delegate User
- Display name:** MAPI_CIFS Delegate User
- Description:** MAPI and CIFS delegate user account
- Office:** (empty)
- Telephone number:** (empty) **Other...**
- E-mail:** (empty)
- Web page:** (empty) **Other...**

The bottom of the dialog box contains the following buttons: OK, Cancel, Apply, and Help.


Associate the delegate user with CIFS and Exchange Services:

After enabling the Delegation tab for the user, you can associate the user with services for which the user can present delegated credentials. When you add this user to the Citrix SD-WAN WANOP appliance, the appliance presents delegated credentials for the services associated with this account.

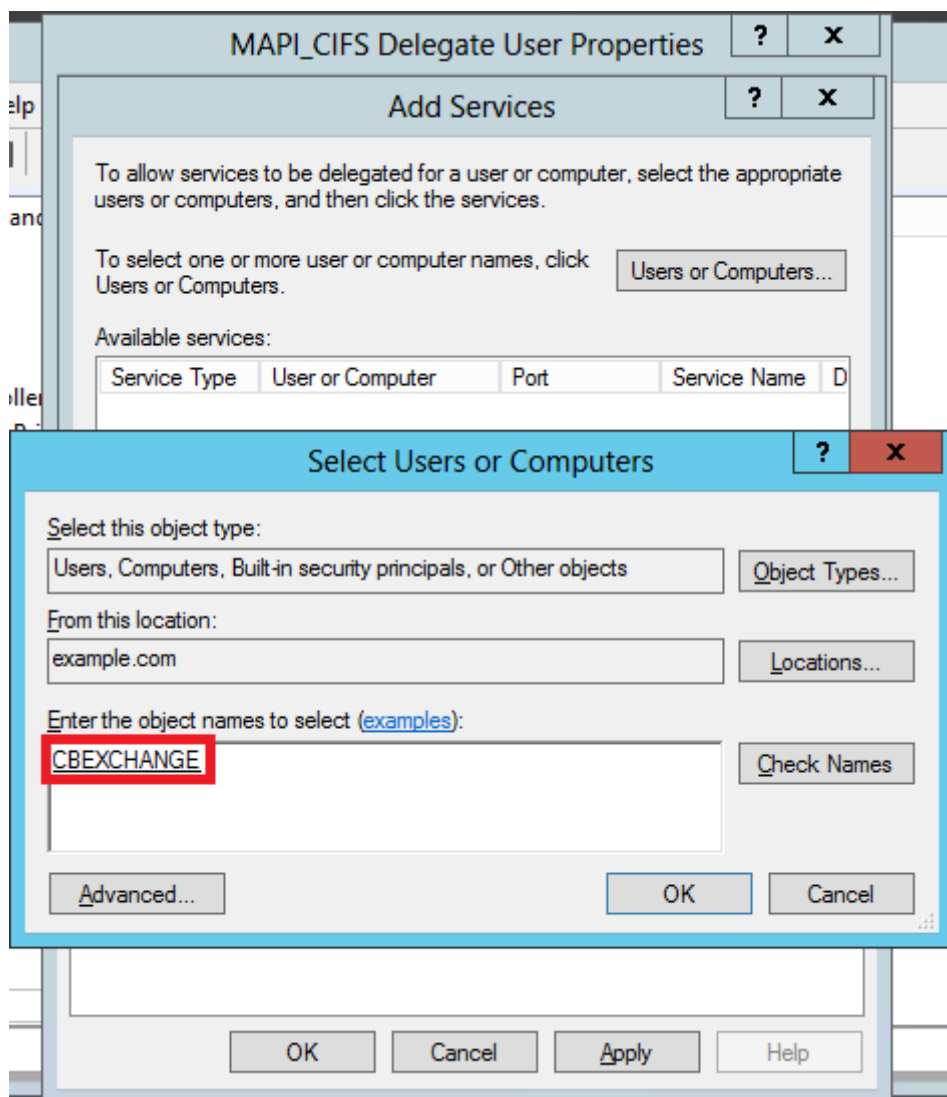
Note: Windows security infrastructure uses the Exchange service to manage MAPI traffic.

To associate the delegate user with CIFS and Exchange services:


1. In the Delegation tab, select the **Trust this user for delegation to specific services only** option.
2. Select the **Use any authentication protocol** option.
3. Click **Add**, as shown in the following screen shot:




4. In the **Add Service** dialog box, click **Users and Computers**.
5. In the **Select Users or Computers** dialog box, add the local computer to be selected, as shown in the following screen shot:



6. Click **OK**.
7. In the Add Services dialog box, from the **Available services** list, select **cifs**, as shown in the following screen shot:



8. If you have to set up MAPI acceleration on the Citrix SD-WAN WANOP appliance, press and hold the **Ctrl** key, and select the **exchangeMDB** service.
9. Click **OK**. The services you have selected are added to the **Services to which this account can present delegated credentials** list, as shown in the following screen shot:



10. Click **OK**.
11. Close the **Active Directory Users and Computers** Window.

Configure a delegate user on a Citrix SD-WAN WANOP appliance:

After configuring the delegate user on the Active Directory server, you must configure this user on the Citrix SD-WAN WANOP appliance, so that the appliance can present this user's delegated credentials to the domain. This enables the appliance to actively optimize the network traffic for the advanced CIFS and MAPI acceleration features.

To add the delegate user to the server-side appliance:

1. Navigate to the **Configuration > Secure Acceleration > Windows Domain** tab.
2. Click the **Join Windows Domain** button, if present.
3. Under **Delegate Users**, click **Add**.

4. In the **Domain Name** field, specify the domain name. This is typically the domain that you specified under the **Windows Domain** section.
5. In the **User Name** field, enter the user name of the delegate user.
6. In the **Password** field, specify the password of the delegate user.
7. Click **Add**.

Delegate Users

Add X Edit Delete Services

Add a delegate user account of the Windows domain controller. The CloudBridge appliance uses this account on behalf of the users, to authenticate them with the domain controller.

Domain Name*
example.com
Check Delegate User

User Name*
delegate_user

Password*
..... ?

Add Cancel

User Name	Domain Name
No items	


Verify that the appliance has joined the domain

If, after adding the appliance to the domain, you notice that the appliance is not optimizing secure Windows traffic, some error might have prevented the appliance from joining the domain. You can use the **Pre Domain Check** utility to find out if there are any issues with the appliance joining the domain. You can even run this utility to identify possible issues before you attempt to join the appliance to a domain.

To check the delegate user:

1. Log on to the server-side Citrix SD-WAN WANOP appliance.
2. Navigate to **Configuration > Secure Acceleration > Windows** tab.
3. Click the **Join Windows Domain** button, if present.
4. Select a delegate user and click **Edit**.
5. Click **Check Delegate User**.

6. Wait for the Delegate User Domain Check to complete and examine the results.



Configure CIFS and SMB2/SMB3 acceleration

March 12, 2021

The CIFS acceleration feature provides a suite of protocol-specific performance enhancements to CIFS-based (Windows and Samba) file transfer and directory browsing, including enhancements to CIFS transport and to related protocols such as DCERPC.

CIFS acceleration has three parts:

- TCP flow-control acceleration—This is performed on all accelerated CIFS connections, regardless of protocol version (SMB1, SMB2, or SMB3) or degree of authentication and encryption.
- CIFS protocol acceleration—These optimizations increase CIFS performance by reducing the number of round trips needed for running a CIFS command. These optimizations are performed

automatically on SMB1 and SMB2 CIFS connections that either do not use CIFS packet authentication (“signing”), or where signing is used and the appliances have joined the Windows domain in a “security delegate” role.

- CIFS compression—CIFS connections are compressed automatically whenever they meet the requirements for CIFS protocol acceleration. In addition, SMB3 connections are compressed when unsigned and unsealed.

On networks where CIFS signing is enabled, CIFS protocol acceleration and compression require that you either disable CIFS packet authentication (signing), or have your datacenter appliances join the Windows domain, and create a secure peer relationship between the datacenter appliances and your remote appliances and Citrix SD-WAN WANOP Plug-ins.

Table 1. CIFS acceleration features, by SMB protocol version and whether the appliance has joined the Windows domain.

SMB Version	TCP Flow Control	Compression	Protocol Acceleration
<i>Signing disabled</i>			
SMB 1.0	Y	Y	Y
SMB 2.0	Y	Y	Y
SMB 2.1	Y	Y	N
SMB 3.0	Y	Y	N
<i>Signing enabled, Citrix SD-WAN WANOP has joined domain **</i>			
SMB 1.0	Y	Y	Y
SMB 2.0	Y	Y	Y
SMB 2.1	Y	Y	Y
SMB 3.0	Y	Y	Y *
<i>Signing enabled, Citrix SD-WAN WANOP has not joined domain</i>			
SMB 1.0	Y	N	N
SMB 2.0	Y	N	N
SMB 2.1	Y	N	N
SMB 3.0	Y	N	N

* SMB 3.0 Support was added in release 7.4.2.

** Citrix SD-WAN WANOP does not support NTLMv2 authentication (default for Windows 7) up with SMB 1/ SMB 2/ SMB 3 and with NetApp server. Enabling Kerberos authentication allows acceleration.

Table 2. Which SMB protocol version is used, by client and server operating system.

Client/Server OS	Windows 8, Windows 10, or Windows Server 2012	Windows 7 or Windows Server 2008 R2	Windows Vista or Windows Server 2008	Earlier versions of Windows
Windows 8, Windows 10, or Windows Server 2012	SMB 3.0	SMB 2.1	SMB 2.0	SMB 1.0
Windows 7 or Windows Server 2008 R2	SMB 2.1	SMB 2.1	SMB 2.0	SMB 1.0
Windows Vista or Windows Server 2008	SMB 2.0	SMB 2.0	SMB 2.0	SMB 1.0
Earlier versions of Windows	SMB 1.0	SMB 1.0	SMB 1.0	SMB 1.0

Supported Versions of CIFS:

Not every CIFS implementation uses request patterns that are recognized by the appliance. These unsupported versions do not achieve acceleration in the full range of cases, as shown in the following table.

Table 3. Citrix SD-WAN WANOP Support for CIFS Servers and Clients.

Product	Server	Client
Windows Server 2003-2012	Yes*	Yes*
Windows XP, Vista, 7, 8, 2000	Yes*	Yes*
NetApp	Yes**	N/A
Hitachi	Yes**	N/A

Product	Server	Client
Windows NT	Yes	No
Windows ME and earlier	No	No
Others	See Note	See Note

* SMB 3.0 support was introduced in release 7.4.2.

** Operation with SMB 3.0 has not been tested as of release 7.4.2.

Note: Most third-party CIFS implementations emulate one of the servers or clients listed above. To the extent that the emulation is successful, traffic is accelerated, or not, as shown in the above table. If the emulation behaves differently from what the CIFS accelerator expects, CIFS acceleration is terminated for that connection.

The behavior of CIFS acceleration with a given CIFS implementation cannot be known for certain until it has been tested.

The modes of CIFS acceleration are:

- Large file reads and writes
- Small file reads and writes
- Directory browsing.

Large file reads and writes—These SMB1 optimizations are for file transfers of at least 640 KB. Safe read-ahead and write-behind techniques are used to stream the data without pauses for every transfer (a transfer is 64 KB or less).

These optimizations are enabled only if the transfer has a BATCH or EXCLUSIVE lock and is “simple.” File copies are always simple. Files opened through applications might or might not be, depending on how they are handled within the application.

Speedup ratios of 10x are readily obtainable with CIFS acceleration, provided that your link and disks are fast enough to accommodate ten times your current transfer speeds. 50x speedup can be obtained if necessary, but is not normally enabled, because of memory consumption. Contact your Citrix representative if 10x is not sufficient.

Small file reads and writes—Small-file enhancements center more around metadata (directory) optimizations than around data streaming. Native CIFS does not combine metadata requests in an efficient way. CIFS acceleration does. As with large-file acceleration, these optimizations are not performed unless they are safe (for example, they are not performed if the CIFS client was not granted an exclusive lock on the directory.) When the SMB2 protocol is used, file metadata is cached locally for even greater improvements.

Directory browsing—Standard CIFS clients perform directory browsing in an extremely inefficient way, requiring an enormous number of round trips to open a remote folder. CIFS acceleration reduces the number of round trips to 2 or 3. When the SMB2 protocol is used, directory data is cached locally for even greater improvements.

CIFS protocol acceleration

CIFS acceleration is supported on all models. CIFS is a TCP based protocol and benefits from flow control. However, CIFS is implemented in a way that is highly inefficient on long-haul networks, requiring an excessive number of round trips to complete an operation. Because the protocol is very sensitive to link latency, full acceleration must be protocol-aware.

CIFS acceleration reduces the number of round-trips through a variety of techniques. The pattern of requests from the client is analyzed and its next action is predicted. In many cases, it is safe to act on the prediction even if it is wrong, and these safe operations are the basis of many optimizations.

For example, SMB1 clients issue sequential file reads in a non-overlapping fashion, waiting for each 64KB read to complete before issuing the next one. By implementing read-ahead, the appliance can safely deliver up to 10x acceleration by fetching the anticipated data in advance.

Additional techniques accelerate directory browsing and small-file operations. Acceleration is applied not only to CIFS operations, but also to the related RPC operations.

Prerequisites

CIFS acceleration is supported on all models. CIFS is a TCP based protocol and benefits from flow control. However, CIFS is implemented in a way that is highly inefficient on long-haul networks, requiring an excessive number of round trips to complete an operation. Because the protocol is very sensitive to link latency, full acceleration must be protocol-aware.

CIFS acceleration reduces the number of round-trips through a variety of techniques. The pattern of requests from the client is analyzed and its next action is predicted. In many cases, it is safe to act on the prediction even if it is wrong, and these safe operations are the basis of many optimizations.

For example, SMB1 clients issue sequential file reads in a non-overlapping fashion, waiting for each 64KB read to complete before issuing the next one. By implementing read-ahead, the appliance can safely deliver up to 10x acceleration by fetching the anticipated data in advance.

Additional techniques accelerate directory browsing and small-file operations. Acceleration is applied not only to CIFS operations, but also to the related RPC operations.

If your network uses CIFS signing, the appliance must be a trusted member of the domain. To make the appliance a trusted member of the domain, see [Adding a Citrix SD-WAN WANOP Appliance to the Windows Security Infrastructure](#).


Configure CIFS protocol acceleration

CIFS acceleration is enabled by default for connections that do not use CIFS signing. If your network uses signing, it can either be disabled or the server-side appliances can [join the Windows domain](#).

Disable CIFS signing

Depending on their security settings, Windows servers or domain servers might need to have their security settings adjusted.

Figure 1. Windows Server Security Options, Windows Server 2003 and Windows Server 2008.



Windows file servers have two security modes: “sealing”and “signing.”

Sealing encrypts the data stream and prevents CIFS protocol acceleration altogether.

Signing adds authentication data to every data packet, without encrypting the data stream. This prevents acceleration unless you have implemented the procedures described in [Adding a Citrix SD-WAN](#)

[WANOP Appliance to the Windows Security Infrastructure](#). When this requirement is met, signing is accelerated automatically. Otherwise, signing must be disabled (if it is not disabled already) for protocol acceleration to take place.

By default, Windows file servers offer signing but do not require it, except for domain servers, which require it by default.

To achieve CIFS acceleration with systems that currently require signing, you must change the system security settings to disable this requirement. You can do so in the local security settings on the file server, or in group policies. The following examples, for Windows Server 2003 and Windows Server 2008, show the local settings. The group-policy changes are, of course, almost identical.

Citrix SD-WAN WANOP

To change the server's setting to allow CIFS acceleration

1. Navigate to the system's Local Security Settings page.
2. Set Domain member: Digitally encrypt or sign secure channel data (always) to Disabled.
3. Set Microsoft network client: Digitally sign communications (always) to Disabled.
4. Set Microsoft network server: Digitally sign communications (always) to Disabled.

Interpret CIFS statistics

The Monitoring: Filesystem (CIFS/SMB) page shows a list of accelerated CIFS connections. These connections are divided into "optimized" and "non-optimized" connections. Because all these connections are accelerated (with flow control and compression), "optimized" connections have CIFS optimizations in addition to flow control and compression, while "non-optimized" connections have flow control and compression only.

CIFS management summary

- CIFS acceleration provides significant improvement even at relatively short link distances.
- CIFS acceleration begins when a file system is first accessed by the client. If acceleration is enabled with the file server and client already up and running, no acceleration occurs for many minutes, until the preexisting CIFS connections are fully closed. CIFS connections are very persistent and last a long time before closing themselves, even when idle. This behavior is annoying during test, but has little importance in normal deployment.
- Dismounting and remounting a file system in Windows does not close the CIFS connections, because Windows does not really dismount the file system fully. Rebooting the client or server

works. For a less invasive measure, use the `NET USE devicename /DELETE` command from the Windows command line to fully dismount the volume. In Linux, `smbmount` and `umount` fully dismount the volume.

- Disabling and then reenabling CIFS read and write optimizations on the appliance raises similar issues. Existing connections do not become accelerated when CIFS is enabled, and the number of “protocol errors detected” on the Monitoring: Filesystem (CIFS/SMB) page increases briefly.
- CIFS statistics can be confusing, because only the appliance farthest from the fileserver reports CIFS acceleration with full statistics. The other appliance sees it as ordinary acceleration.
- CIFS acceleration is not supported in proxy mode.
- If CIFS acceleration does not take place with a Windows server, check the server’s security settings.

Configure MAPI acceleration

March 12, 2021

Microsoft Outlook acceleration provides improved performance for traffic between Microsoft Outlook clients and Microsoft Exchange Servers, increasing throughput with a variety of optimizations, including data prefetching and compression.

This feature is also called “MAPI acceleration,” after the MAPI protocol used between Outlook and Exchange Server.

In networks where the Outlook data stream is unencrypted (the default before Outlook 2007), this feature requires no configuration.

In networks where the Outlook data is encrypted (the default with Outlook 2007 and later), acceleration can be obtained in one of two ways: by disabling encryption in the Outlook clients or by having the appliances [join the Windows domain](#).

Supported outlook exchange versions and modes

Citrix SD-WAN WANOP appliances provide MAPI acceleration for Microsoft Outlook 2003-2016 and Exchange Server 2003-2010, in the following circumstances:

- Any combination of supported clients and servers (using the MAPI protocol) is supported.
- If the server-side appliance has joined a Windows domain, connections with MAPI encryption are accelerated. Otherwise, they are not, and encryption should be disabled in the Outlook clients.

Note

In Exchange Server 2013 the MAPI protocol changed to RPC over HTTP protocol, this protocol is supported. With Exchange Server SP1, the RPC over HTTP protocol changed to MAPI over HTTP protocol, this protocol is currently not supported.

Prerequisites

If your network uses encrypted Outlook data, which is the default setting in Outlook 2007 and later, you must implement one of the following prerequisites to make sure that MAPI connections are accelerated:

- Disable encryption in the Outlook clients.
- Perform the tasks described in [Adding a Citrix SD-WAN WANOP Appliance to the Windows Security Infrastructure](#).

Configuration

Outlook acceleration is a zero-configuration feature that is enabled by default. (If not wanted, it can be disabled by disabling acceleration on the MAPI service class on the

Configuration: Service Class Policy page.) Outlook acceleration takes place automatically if the following conditions are met:

- There is an appliance at the Exchange Server end of the WAN.
- Either there is an appliance at the Outlook end of the WAN, or the system running Outlook is also running the Citrix SD-WAN WANOP Plug-in.
- All Outlook/Exchange traffic passes through the appliances (or appliance and plug-in).
- Either the Exchange Server or Outlook is restarted (acceleration does not begin until existing MAPI connections are closed).
- Either encryption is disabled on Outlook, or the server-side appliance belongs to the Windows domain and has a secure peer relationship with the client-side appliance (or Citrix SD-WAN WANOP Plug-in). In the case where the appliance has joined the Windows domain, authentication on the domain must be kept at the default setting (negotiate), for acceleration to work.

Disable encryption on Outlook 2007 or Outlook 2010

Unless the server-side appliance has joined the Windows domain and has a secure peer relationship with the client-side appliance (or Citrix SD-WAN WANOP Plug-in), encryption between Outlook and Exchange Server must be disabled for acceleration to take place.

Encryption was disabled by default before Outlook 2007. Starting with Outlook 2007, encryption is enabled by default.

Performance note

MAPI uses a different data format from other protocols. This difference prevents effective cross-protocol compression. That is, a file that was first transferred through FTP and then as an email attachment does not receive a compression advantage on the second transfer. If the same data is sent two times in MAPI format, the second transfer receives full compression.

SSL compression

March 12, 2021


Citrix SD-WAN WANOP SSL compression applies multisession compression to SSL connections (for example, HTTPS traffic), providing compression ratios of up to 10,000:1.

Note


SSL compression requires a secure peering (signaling) connection between the two appliances at the ends of the accelerated link.

Encryption is maintained from end to end by splitting the connection into three encrypted segments: client to client-side appliance, client-side appliance to server-side appliance, and server-side appliance to server.

Ordinary SSL Connection



Accelerated SSL Connection



Caution: SSL Compression decrypts the encrypted data stream and, unless the User Data Encryption option is used, the compression histories of both acceleration units retain clear-text records of the decrypted data. Verify that your deployment and settings are consistent with your organization's security policies. Citrix recommends that you enable encryption of the compression history on each unit when you configure the secure peering signaling connection required for SSL acceleration.

Note

- When you enable SSL compression, the appliance stops attempting compression with other appliances with which it does not have a secure peer relationship (whether Citrix SD-WAN WANOP, or Citrix SD-WAN WANOP Plug-in). This feature is thus best-suited for networks where all appliances are configured for SSL compression.
- With SSL compression enabled, you must manually type in the Key Store password each time the appliance is restarted.

How SSL compression works

March 12, 2021

SSL compression has access to the clear-text data of the connection, because the server-side appliance acts as a *security delegate* of the endpoint servers. This behavior is possible because the server-side appliance is configured with copies of the servers' security credentials (private keys and certifi-

ates), allowing it to act on the servers' behalf. To the client, this behavior is equivalent to communicating directly with the endpoint server.

Because the appliance is working as a security delegate of the server, most configuration is on the server-side appliance. The client-side appliance (or plug-in) acts as a satellite of the server-side appliance and does not require per-server configuration.

The server-side and client-side appliances share session status through an *SSL signaling connection*. All accelerated connections between the two appliances are sent over *SSL data connections*, whether the original connections were encrypted or not.

Note: SSL compression does not necessarily encrypt all link traffic. Traffic that was originally encrypted remains encrypted, but unencrypted traffic is not always encrypted. The appliances do not attempt to encrypt unaccelerated traffic. Because there is no absolute guarantee that any given connection will be accelerated (various events prevent acceleration), there is no guarantee that the appliances will encrypt a given unencrypted connection.

SSL compression operates in one of two modes: transparent proxy or split proxy. These two modes support slightly different SSL features. You select the mode that provides the features a given application requires.

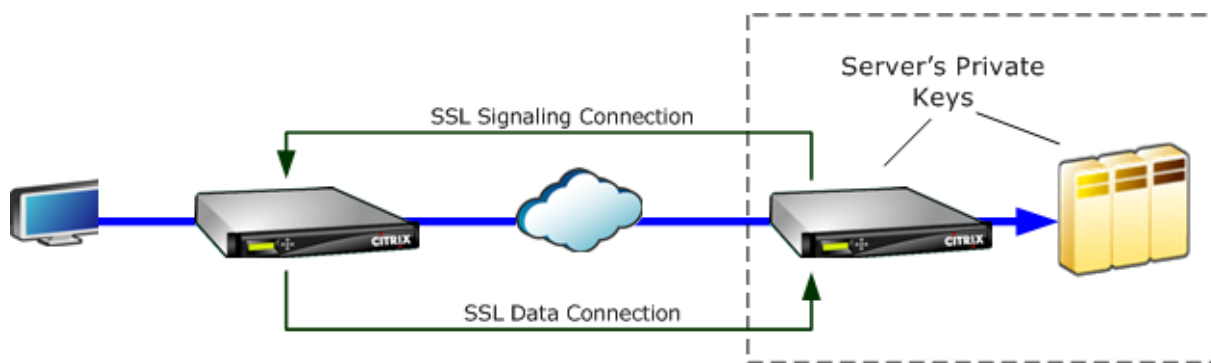
Which SSL proxy mode to use—Use SSL transparent proxy mode *only* if you require true client authentication (that is, authentication that correctly identifies the individual endpoint client) *and* you do not require Diffie-Hellman, Temp RSA, TLS session tickets, SSL version 2, or session renegotiation. Use SSL split proxy for all other deployments.

SSL transparent proxy

In *SSL transparent proxy mode* (not to be confused with transparent mode on the Citrix SD-WAN WANOP Plug-in), the server-side appliance masquerades as the server. The server's credentials (certificate-key pair) are installed on the server-side appliance so that it can act on the server's behalf. The server-side appliance then configures the client-side appliance to handle the client end of the connection. The server's credentials are not installed on the client-side appliance.

True client authentication is supported in this mode, but Temp RSA and Diffie-Hellman are not. SSL transparent proxy mode is suited for applications that require client authentication, but only if none of the following features are required: Diffie-Hellman, Temp RSA, TLS session tickets, SSL version 2. Also, session renegotiation must not be attempted, or the connection terminates.

No configuration is required on the client-side appliance (other than configuring a secure peering relationship with the server-side appliance), and no configuration is required on the client, which treats the connection exactly as if it were communicating directly with the server.



SSL split proxy

SSL split proxy mode is preferred in most instances, because it supports Temp RSA and Diffie-Hellman, which many applications require. In SSL split proxy mode, the server-side appliance masquerades as a server to the client, and as a client to the server. You install server credentials (a certificate-key pair) on the server-side appliance to allow it to act on the server's behalf.

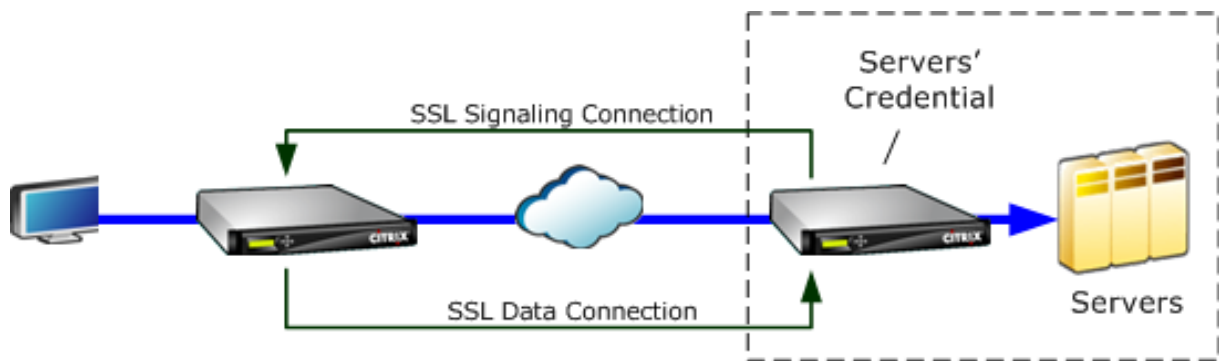
Split proxy mode also supports proxied client authentication if you install optional client credentials, which are presented to the endpoint server application if it requests client authentication. These client credentials will be presented instead of the actual endpoint client's credentials. (Use transparent proxy if the endpoint client credentials are required by the application.)

Because true client authentication is not supported in this mode, the server cannot authenticate the actual endpoint client. If the server-side appliance is not configured with client credentials, all attempts by the server-side application at client authentication fail. If the server-side appliance is configured with client credentials, all requests for client authentication will be answered with these credentials, regardless of the identity of the actual client.

No configuration is required on the client-side appliance (other than configuring a secure peering relationship with the server-side appliance), and no configuration is required on the client, which treats the connection as if it were communicating directly with the server. The server credentials on the server-side appliance are not installed on the client-side appliance.

To support multiple servers, multiple private certificate-key pairs 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 credentials.

In SSL split proxy mode, the CA certificates and certificate-key pairs and CA certificates do not actually have to match those of the servers, though they can. Due to the nature of a split proxy, the server-side appliance can use credentials that are acceptable to the client application (valid credentials issued by a trusted authority). Note that, in the case of HTTPS connections, Web browsers issue a warning if the common name does not match the domain name in the URL. In general, using copies of the server's credentials is the more trouble-free option.



Configure SSL compression

March 12, 2021

The Citrix 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 Citrix 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.

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.

Configure a split proxy SSL profile


To configure a split proxy SSL profile:

1. In the server-side Citrix 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:

In the **Verification Store** field, select an existing server Certificate Authority (CA), or click **+** to upload a server CA.

Choose **Authentication Required** and in the **Certificate/Private Key** field select a certificate key pair, or click **+** to upload a certificate key pair.


In the **Protocol Version** field, select the protocols your server accepts.

Note

Citrix 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.

If necessary, edit the **Cipher Specification** string, using the OpenSSL syntax.

If required, select the type of renegotiation from the **Renegotiation Type** drop-down list to allow client-side SSL session renegotiation.



Server-Side Proxy Configuration

Verification Store
CA

☒ Authentication Required

Certificate/Private Key*
split

☒ Build Certificate Chain

Protocol Version*
TLS 1.0, TLS 1.1 or TLS 1.2

Cipher Specification*
!ADH:HIGH:MEDIUM:@STRENGTH

Renegotiation Type*
Old Style Renegotiation Disabled

7. Perform client-side proxy configuration:

In the **Certificate/Private Key** field, retain the default value.

Choose **Build Certificate Chain** to allow the server-side appliance to build the SSL certificate chain.

If required, select or upload a CA store to use as the Certificate Chain Store.


In the **Protocol Version** field, select the protocol versions you want to support on the client side.

Note

Citrix 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.

If necessary, edit the client-side Cipher Specification.

If required, select the type of renegotiation from the **Renegotiation Type** drop-down list to allow client-side SSL session renegotiation.



Client-Side Proxy Configuration

Certificate/Private Key*
split

☐ Disable Session Re-use

☒ Build Certificate Chain

Certificate Chain Store

Protocol Version*
TLS 1.0, TLS1.1 or TLS 1.2

Cipher Specification*
!ADH:HIGH:MEDIUM:@STRENGTH

Renegotiation Type*
Old Style Renegotiation Disabled

8. Click **Create**.

Configure transparent proxy SSL profile

To configure a transparent proxy SSL profile:

1. In the server-side Citrix 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.

Create SSL Profile

☒ Manually add Profile ☐ Import Profile

Profile Name*
SSL-Server2

☒ Profile Enabled
☐ Parse Subject Alternative Names

Virtual Host Name
Server2

Proxy Type
☐ Split ☒ Transparent

SSL Server's Private Key*
split

Create Close

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**.

Create or modify the service class

To create or modify the service class and attach the SSL Profile:

1. In the Citrix 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

Citrix 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:

```
1  *--name "profile-name" *
2
3  *\[--state {
4    enable, disable }
5    \]*
6
7  *--proxy-type split*
8
9  *\[--virtual-hostname "hostname" \]*
10
11 *--cert-key "cert-key-pair-name" *
12
13 *\[--build-cert-chain {
14   enable, disable }
15   \]*
16
17 *\[--cert-chain-store {
18   use-all-configured-CA-stores, "store-name" }
19   \]*
20
21 *\[--cert-verification {
22   none, Signature/Expiration, Signature/Expiration/*
23
24 *Common-Name-White-List, Signature/Expiration/Common-Name-Black-List }
25   \]*
26
27 *\[--verification-store {
```

```
28 use-all-configured-CA-stores, "store-name" }
29 \]*
30
31 *\[ -server-side-protocol {
32     TLS-1.2, TLS-version-any }
33 \]*
34
35 *\[ -server-side-ciphers "ciphers" \]*
36
37 *\[ -server-side-authentication {
38     enable, disable }
39 \]*
40
41 *\[ -server-side-cert-key "cert-key-pair-name" \]*
42
43 *\[ -server-side-build-cert-chain {
44     enable, disable }
45 \]*
46
47 *\[ -server-side-renegotiation {
48     disable-old-style, enable-old-style, new-style,*
49
50 *compatible }
51 \]*
52
53 *\[ -client-side-protocol-version {
54     TLS-1.2, TLS-version-any }
55 \]*
56
57 *\[ -client-side-ciphers "ciphers" \]*
58
59 *\[ -client-side-renegotiation {
60     disable-old-style, enable-old-style, new-style,*
61
62 *compatible }
63 \]*
```

set ssl-profile:

```
1 *-name "profile-name" \[-state {
2     enable, disable }
3 \]*
4
5 *\[ -proxy-type split\]*
6
7 *\[ -virtual-hostname "hostname" \]*
8
9 *\[ -cert-key "cert-key-pair-name" \]*
10
11 *\[ -build-cert-chain {
12     enable, disable }
13 \]*
14
```

```
15 *\[ -cert-chain-store {
16   use-all-configured-CA-stores, "store-name" }
17   \]*
18
19 *\[ -cert-verification {
20   none, Signature/Expiration, Signature/Expiration/*
21
22   *Common-Name-White-List, Signature/Expiration/Common-Name-Black-List }
23   \]*
24
25 *\[ -verification-store {
26   use-all-configured-CA-stores, "store-name" }
27   \]*
28
29 *\[ -server-side-protocol {
30   TLS-1.2, TLS-version-any }
31   \]*
32
33 *\[ -server-side-ciphers "ciphers" \]*
34
35 *\[ -server-side-authentication {
36   enable, disable }
37   \]*
38
39 *\[ -server-side-cert-key "cert-key-pair-name" \]*
40
41 *\[ -server-side-build-cert-chain {
42   enable, disable }
43   \]*
44
45 *\[ -server-side-renegotiation {
46   disable-old-style, enable-old-style, new-style,*
47
48   *compatible }
49   \]*
50
51 *\[ -client-side-protocol-version {
52   TLS-1.2, TLS-version-any }
53   \]*
54
55 *\[ -client-side-ciphers "ciphers" \]*
56
57 *\[ -client-side-renegotiation {
58   disable-old-style, enable-old-style, new-style,*
59
60   *compatible }
61   \]*
```


SSL Compression with Citrix SD-WAN WANOP plug-in

March 12, 2021

The Citrix SD-WAN WANOP Plug-in is always used as the client-side unit and thus requires no additional SSL configuration other than installing credentials for the SSL signaling (secure peering) connection. The main difference between SSL compression on the plug-in and the appliance is that the plug-in is unable to encrypt the user data in the disk based compression history.

Caution: Because disk based compression history on the Plug-in is not encrypted, it retains a clear-text record of potentially sensitive and ephemeral encrypted communications. This lack of encryption is potentially dangerous on computers for which physical access is not controlled. Therefore, Citrix recommends the following best practices:

- Do not use **Certificate Validation: None** on your appliances. (Note that, in this case, the appliance refuses to allow compression with plug-ins that do not have appropriate certificates.)
- Install certificates only on systems that can be verified to meet your organization's requirements for physical or data security (for example, laptops that use full-disk encryption).

The Citrix SD-WAN WANOP Plug-in supports both SSL split proxy and SSL transparent proxy. The plug-in ships without certificate-key pairs for the SSL signaling connection. If desired, the same credentials can be used by all plug-ins, or each plug-in can have its own credentials.

The plug-in does not attempt SSL compression unless credentials have been installed.

The plug-in inherits its crypto license from the appliance.

RPC over HTTP

March 12, 2021

Microsoft Exchange Server is one of the common email servers used across organizations. As a result of recent enhancements in Microsoft Exchange Server, you can securely connect to it over the Internet. Depending on the available bandwidth, you might experience latency in the email delivered to the Outlook client. In addition to the MAPI protocol, the Citrix SD-WAN WANOP appliance supports Remote Procedure Call over HTTPS (RPC over HTTPS) to optimize Microsoft Exchange traffic. This feature is also known as Outlook Anywhere.

RPC over HTTPS is not a new protocol, but starting with Microsoft Exchange 2013, it replaces MAPI as the default protocol. The main advantage of RPC over HTTPS is that it enables clients to securely connect to the mail server over the Internet.

When you use RPC over HTTPS, the Microsoft Exchange server must use a digital certificate and private key to authenticate itself to the Outlook client. The communication between the client and server uses HTTPS as a transport protocol.

On the Citrix SD-WAN WANOP appliance, RPC over HTTPS is supported for the following the Microsoft Outlook and Exchange Server versions:

- Microsoft Outlook
 - Microsoft Outlook version 2007
 - Microsoft Outlook version 2010
 - Microsoft Outlook version 2013
- Microsoft Exchange Server
 - Microsoft Exchange Server version 2007
 - Microsoft Exchange Server version 2010
 - Microsoft Exchange Server version 2013

Of these, all versions except Microsoft Exchanges Server 2013 support MAPI (over TCP) as well as RPC over HTTPS. However, Microsoft Exchange Server 2013 forces connections to use RPC over HTTPS, regardless of the Microsoft Outlook version you use, to connect to the Exchange server.

Configure RPC over HTTPS

By default, the RPC over HTTPS feature is enabled on the appliance. However, to configure the appliance to accelerate RPC over HTTPS, you must perform the following additional tasks:

- Configure encrypted MAPI.
- Configure an SSL profile with a server certificate.
- Create an RPC over HTTPS service class and bind the SSL profile to it.

Configure Encrypted MAPI

Note

Skip this section if you have already configured encrypted MAPI acceleration on the appliance.

Microsoft Outlook uses Messaging Application Programming Interface (MAPI) connections between Outlook clients and the Microsoft Exchange server. MAPI connections use RPCs, which are encapsulated by an HTTP connection. Therefore, before you configure RPC over HTTPS on a Citrix SD-WAN WANOP appliance, you must configure encrypted MAPI on the appliance.

Prerequisites:

Before you configure encrypted MAPI, make sure that the following prerequisites are met:

- The Secure Peer option should be set to True on the client as well as the server-side appliance. To configure a secure partner, see [Secure Peering](#).
- The DNS IP address configured on the server-side appliance must be reachable.
- The datacenter-side appliance must successfully join the domain.
- A delegate user must be added to the datacenter-side appliance, and its status should be marked as “Success.”

For more information, see [Configure a Citrix SD-WAN WANOP appliance to optimize secure Windows traffic](#).

Configure an SSL profile with a server certificate

The HTTPS connection that encapsulates the MAPI connection is secured by SSL. As a result, RPC over HTTPS requires connectivity through TCP port 443. This port is assigned to HTTPS, which web-server administrators usually keep open in the firewall application. Using SSL-protected communication helps RPC over HTTPS to maintain the security of all communications.

To enable RPC over HTTPS acceleration, you must install a server certificate on the appliance. Using this server certificate, you can configure an SSL profile that RPC over HTTPS uses for secure communication. To configure an SSL profile with an Exchange server certificate, see [Installing Server and Client Certificates](#).

Note

You must configure an SSL profile only on the datacenter-side appliance.

Create an RPC over HTTPS service class and bind the SSL profile to it

To optimize the RPC over HTTP connections, you must create a service class that lists HTTPS and all MAPI applications. You must provide the IP address of the Microsoft Exchange server as a destination IP address for this service class, and then bind the SSL profile you created to this service class. Binding the profile to the service class makes sure that the communication between the Outlook client and Microsoft Exchange server is secured by using this profile.

Note

You must configure and bind an SSL profile to the service class only on the datacenter-side appliance.

Verify accelerated RPC over HTTPS connections

After you have configured RPC over HTTPS on the appliance , you can verify that the appliance is accelerating the RPC over HTTPS connection on the Monitoring page for MAPI. The accelerated RPC over HTTPS connections are listed on the Accelerated MAPI Sessions tab.

Note

You must configure RPC over HTTPS on your client-side appliances as well as your server-side Citrix SD-WAN WANOP appliances to accelerate the RPC over HTTPS connections.

To verify that RPC over HTTPS Connections are being accelerated

1. Navigate to the **Monitoring > Optimization > Outlook (MAPI)**.
2. On the **Accelerated MAPI Sessions** tab, verify that RPC over HTTPS connections are accelerated.

The screenshot shows the Citrix SD-WAN WANOP Monitoring interface. The left sidebar has a menu with 'Optimization' expanded, showing 'Outlook (MAPI)' selected. The main content area shows the 'Accelerated MAPI Sessions' tab. It displays summary statistics: 'Optimized MAPI Session Count' is 58 and 'Accelerated TCP connection count' is 213. Below this is a table with columns: TCP Connection Count, Client, Server, Bytes Sent, Bytes Received, User Name, Encrypted, and Service Class. One row is visible with values: 213, 192.168.10.33, 192.168.20.5, 744.26 MB, 2.70 GB, Administrator, True, and HTTPS eMAPI.

TCP Connection Count	Client	Server	Bytes Sent	Bytes Received	User Name	Encrypted	Service Class
213	192.168.10.33	192.168.20.5	744.26 MB	2.70 GB	Administrator	True	HTTPS eMAPI

Note

The Application has possible values of: HTTPS eMAPI, HTTP eMAPI, HTTPS MAPI, and HTTP MAPI.

TCP Flow-Control acceleration

March 12, 2021

Ordinary WANs have very poor responsiveness at high link utilization and at long distances. A widely used rule of thumb for ordinary, non-accelerated WAN links is, “once link utilization reaches 40%, it is time to add more bandwidth, because performance and reliability have degraded to the point where the link is largely unusable.” Interactive performance suffers, making it hard for people to get work done, and connections frequently time out. Accelerated links do not have this problem. A link with 95% utilization is still perfectly usable.

Citrix SD-WAN WANOP appliances become virtual gateways that control the TCP traffic on the WAN link. Ordinary TCP is controlled on a per-connection basis by the endpoint devices. Optimal control

of link traffic is difficult, because neither the endpoint devices nor individual connections have any knowledge of the link speed or the amount of competing traffic. A gateway, on the other hand, is in an ideal position to monitor and control link traffic. Ordinary gateways squander this opportunity because they cannot supply the flow control that TCP lacks. Citrix SD-WAN WANOP technology adds the intelligence that is missing in the network equipment and the TCP connections alike. The result is greatly improved WAN performance, even under harsh conditions such as high loss or extreme distance.

Citrix SD-WAN WANOP flow control is lossless and transparent, and it implements a broad spectrum of speed optimizations. No configuration is required, because of autodiscovery and autoconfiguration. You might, however, have to tweak your firewalls if they block the TCP options used by the acceleration algorithms.

Lossless and transparent flow control

March 12, 2021

Acceleration operates on any TCP connection passing through two appliances (one at the sending site and one at the receiving site), or a Citrix SD-WAN WANOP appliance and a Citrix SD-WAN WANOP Plug-in. Although the above figure shows a network of two appliances, any appliance can accelerate connections between any number of other appliance-equipped sites simultaneously. This allows a single appliance to be used per site, rather than two per link.

Like any gateway, the Citrix SD-WAN WANOP appliance meters packets onto the link. Unlike ordinary gateways, however, it imposes transparent, lossless flow control on each link segment, including:

- The LAN segment between the sender and the sending appliance
- The WAN segment between the sending and receiving appliances
- The LAN segment between the receiving appliance and the receiver


Flow control can be managed independently for each of these three segments. The segments are partly decoupled, so each can have its speed controlled independently. This is important when a connection's speed needs to be ramped up or down quickly to its fair bandwidth share, and is also important as a means of supporting enhanced WAN algorithms and compression.

The TCP protocol is designed to make every TCP connection attempt to increase its bandwidth usage continuously. However, the link bandwidth is limited. The result is that the links become overrun. Citrix SD-WAN WANOP flow control keeps the TCP connections flowing at just the right speed. The link is filled but is never overrun, so queuing latency and packet losses are minimized, while throughput is maximized.

With ordinary TCP, long-running connections (which have had time to seize all the bandwidth) tend to squeeze out short-running connections. This problem, which ruins interactive responsiveness, does not occur with flow control.

Flow control is a standard feature on all appliances in the Citrix SD-WAN WANOP family.

Figure 1. Acceleration Enhances Performance Transparently



Speed optimization

March 12, 2021

Most TCP implementations do not perform well over WAN links. To name just two problems, the standard TCP retransmission algorithms (Selective Acknowledgments and TCP Fast Recovery) are inadequate for links with high loss rates, and do not consider the needs of short-lived transactional connections.

Citrix SD-WAN WANOP implements a broad spectrum of WAN optimizations to keep the data flowing under all kinds of adverse conditions. These optimizations work transparently to ensure that the data arrives at its destination as quickly as possible.


WAN optimization operates transparently and requires no configuration.

WAN optimization is a standard feature on all Citrix SD-WAN WANOP appliances.

The figure below shows the transfer speeds possible at various distances, without acceleration, when the endpoints use standard TCP (TCP Reno). For example, gigabit throughputs are possible without acceleration within a radius of a few miles, 100 Mbps is attainable to less than 100 miles, and throughput on a worldwide connection is limited to less than 1 Mbps, regardless of the actual speed of the link.

With acceleration, however, the speeds above the diagonal line become available to applications. Distance is no longer a limiting factor.

Figure 1. Non-accelerated TCP Performance Plummet With Distance



Note

Without Citrix acceleration, TCP throughput is inversely proportional to distance, making it impossible to extract the full bandwidth of long-distance, high-speed links. With acceleration, the distance factor disappears, and the full speed of a link can be used at any distance. (Chart based on model by Mathis, *et al*, Pittsburgh Supercomputer Center.)

Accelerated transfer performance is approximately equal to the link bandwidth. The transfer speed is not only higher than with unaccelerated TCP, but is also much more constant in the face of changing network conditions. The effect is to make distant connections behave as if they were local. User-perceived responsiveness remains constant regardless of link utilization. Unlike normal TCP, with which a WAN operating at 90% utilization is useless for interactive tasks, an accelerated link has the same responsiveness at 90% link utilization as at 10%.

With short-haul connections (ones that fall below the diagonal line in the figure above), little or no acceleration takes place under good network conditions, but if the network becomes degraded, performance drops off much more slowly than with ordinary TCP.

Non-TCP traffic, such as UDP, is not accelerated. However, it is still managed by the traffic shaper.

Example

One example of advanced TCP optimizations is a retransmission optimization called *transactional mode*. A peculiarity of TCP is that, if the last packet in a transaction is dropped, its loss not noticed by the sender until a receiver timeout (RTO) period has elapsed. This delay, which is always at least one second long, and often longer, is the cause of the multiple-second delays seen on lossy links—delays that make interactive sessions unpleasant or impossible.

Transactional mode solves this problem by automatically retransmitting the final packet of a transaction after a brief delay. Therefore, an RTO does not happen unless both copies are dropped, which is unlikely.

A bulk transfer is basically a single enormous transaction, so the extra bandwidth used by transactional mode for a bulk transfer can be as little as one packet per file. However, interactive traffic, such as key presses or mouse movements, has small transactions. A transaction might consist of a single undersized packet. Sending such packets twice has a modest bandwidth requirement. In effect, transactional mode provides forward error correction (FEC) on interactive traffic and gives end-of-transaction RTO protection to other traffic.


Auto-discovery and auto-configuration

March 12, 2021

In process called autodiscovery, Citrix SD-WAN WANOP units detect each other's presence automatically. The appliances attach TCP header options to the first packets in each connection: the SYN packet (sent by the client to the server to open the connection), and the SYN-ACK packet (sent by the server to the client to indicate that the connection has been accepted). By tagging the SYN packets and listening for tagged SYN and SYN-ACK packets, the appliances can detect each other's presence in real time, on a connection-by-connection basis.

The main benefit of autodiscovery is that you do not have to reconfigure all of your appliances every time you add a new one to your network. They find each other automatically. In addition, the same process allows autoconfiguration. The two appliances use the TCP header options to exchange operating parameters, including the bandwidth limits (in both the sending and receiving directions), the basic acceleration mode (hardboost or softboost), and the acceptable compression modes (disk, memory, or none). All of the information that each appliance needs about its partner is exchanged with each connection, allowing per-connection variations (for example, per-service-class variations in the allowable compression types).

Figure 1. How autodiscovery works



The autodiscovery process works as follows:

1. The client opens a TCP connection to the server, as usual, by sending it a TCP SYN packet.
2. The first appliance passes the SYN packet through after attaching a set of appliance-specific TCP header options to it and adjusting its window size.
3. The second appliance reads the TCP options, removes them from the packet, and forwards them to the server.
4. The server accepts the connection by responding as usual with a TCP SYN-ACK packet.
5. The second appliance remembers that this connection is a candidate for acceleration and attaches its own acceleration options to the SYN-ACK header.
6. The first appliance reads the options added by the second appliance, strips them from the packet header, and forwards the packet to the client. The connection is now accelerated. The two appliances have exchanged the necessary parameters through the option values, and they store them in memory for the duration of the connection.

The connection is accelerated, and the acceleration is transparent to the client, server, routers, and firewalls.

TCP flow control modes

March 12, 2021

TCP flow control has two modes: softboost and hardboost.

Softboost uses a rate-based sender that sends accelerated traffic at speeds up to the link's bandwidth limit. If the bandwidth limit is set slightly lower than the link speed, packet loss and latency are minimized, while link utilization is maximized. Interactive applications see fast response times while bulk-transfer applications see high bandwidth. Softboost shares the network with other applications in any topology, and it interoperates with third-party QoS systems.

Hardboost is more aggressive than softboost. By ignoring packet losses and other so-called "congestion signals," it performs very well on links plagued with heavy, non-congestion-related losses, such as satellite links. It is also excellent on low-quality, long-haul links with a high background packet loss, such as many overseas links. Hardboost is recommended only for point-to-point links that do not achieve adequate performance with softboost.

Softboost is the default mode and is recommended in most cases.


Note

- Hardboost should be used only on fixed-speed point-to-point links or hub-and-spoke deployments where the hub bandwidth is at least equal to the sum of the accelerated spoke bandwidths.
- Softboost and hardboost are mutually exclusive, which means that all the Appliances that must communicate with each other must be set the same. If one unit is set to hardboost and the other is set to softboost, no acceleration takes place.

To select softboost mode:

Softboost is the default mode and is recommended in most cases.

1. Navigate to **Configuration > Links > Hardboost / Softboost** and click edit.
2. Select **Softboost** as the **WAN Boost Mode**.



The screenshot shows a 'Link Settings' dialog box. Under 'WAN Boost Mode', there are two radio buttons: 'Hardboost' and 'Softboost'. The 'Softboost' radio button is selected. Below this, there is a section for 'WAN Bandwidth Receive Limit' with a text input field containing the number '1' and a dropdown menu set to 'gbps'. At the bottom of the dialog are 'Save' and 'Cancel' buttons.

3. Click **Save**

To select hardboost mode:

Select hardboost mode only on fixed-speed point-to-point links or hub-and-spoke links where the hub bandwidth is greater than or equal to that of the accelerated spoke links.

1. Navigate to **Configuration > Links > Hardboost / Softboost** and click edit.

2. Select **Hardboost** as the **WAN Boost Mode**.
3. Set **WAN Bandwidth Receive Limit** to 95% of the link speed.
4. Click **Save**.

Firewall considerations

March 12, 2021

The Citrix SD-WAN WANOP appliance's use of TCP options puts accelerated traffic at risk from firewalls that have aggressive rules about denying service to connections using less-common TCP options.

Some firewalls strip off the "unknown" options and then forward the packet. This action prevents acceleration but does not impair connectivity.

Other firewalls deny service to connections with unknown options. That is, the SYN packets with Citrix SD-WAN WANOP options are dropped by the firewall. When the appliance detects repeated connection-attempt failures, it retries without the options. This restores connectivity after a delay of variable length, usually in the range of 20-60 seconds, but without acceleration.

Any firewall that does not pass Citrix SD-WAN WANOP options through unmodified must be reconfigured to accept TCP options in the range of 24–31 (decimal).

Most firewalls do not block these options. However, Cisco ASA and PIX firewalls (and perhaps others) with release 7.x firmware might do so by default.

The firewalls at both ends of the link should be examined, because either one might be permitting options on outgoing connections but blocking them on incoming connections.

The following example should work with Cisco ASA 55x0 firewalls using 7.x firmware. Because it globally allows options in the range of 24-31, there is no customized per-interface or per-unit configuration:

```
1  =====
2  CONFIGURATION FOR CISCO ASA 55X0 WITH 7.X CODE TO ALLOW TCP OPTIONS
3  =====
4  hostname(config)# tcp-map WSOptions
5  hostname(config-tcp-map)# tcp-options range 24 31 allow
6  hostname(config-tcp-map)# class-map WSOptions-class
7  hostname(config-cmap)# match any
8  hostname(config-cmap)# policy-map WSOptions
9  hostname(config-pmap)# class WSOptions-Class
10 hostname(config-pmap-c)# set connection advanced-options WSOptions
11 hostname(config-pmap-c)# service-policy WSOptions global
12 <!--NeedCopy-->
```

Configuration for a PIX firewall is similar:

```

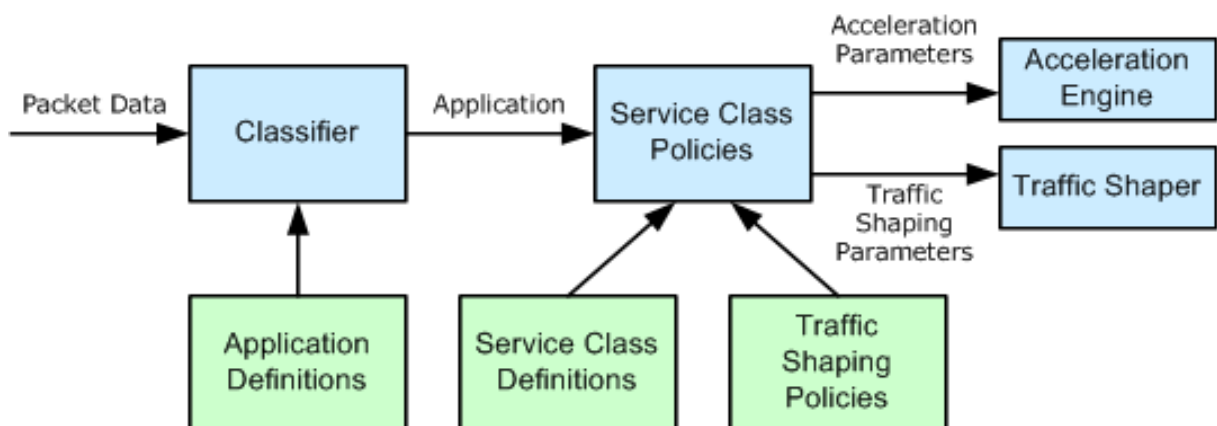
1  =====
2  POLICY MAP TO ALLOW APPLIANCE TCP OPTIONS TO PASS (PIX 7.x)
3  =====
4  pixfirewall(config)#access-list tcpmap extended permit tcp any any
5  pixfirewall(config)# tcp-map tcpmap
6  pixfirewall(config-tcp-map)# tcp-opt range 24 31 allow
7  pixfirewall(config-tcp-map)# exit
8  pixfirewall(config)# class-map tcpmap
9  pixfirewall(config-cmap)# match access-list tcpmap
10 pixfirewall(config-cmap)# exit
11 pixfirewall(config)# policy-map global_policy
12 pixfirewall(config-pmap)# class tcpmap
13 pixfirewall(config-pmap-c)# set connection advanced-options tcpmap
14 <!--NeedCopy-->

```

Traffic classification

March 12, 2021

The two main functions of a Citrix SD-WAN WANOP appliance are traffic shaping, which maximizes link usage for all types of traffic, and acceleration, which applies compression and various optimizations to accelerate TCP traffic. Two basic components of both traffic shaping and acceleration are the application-classifier mechanism and the service-class mechanism. The former identifies the type of traffic, so that the latter can assign the traffic to a service class. Each service class has a traffic shaping policy and an acceleration policy.



Application classifier

March 12, 2021

The application classifier uses application definitions to categorize the traffic by protocol and application. This information is used to create reports, and by the service-class mechanism. Many applications are already defined, and you can define more as needed.

Protocol and port specifications in application definitions

The application classifier uses the official protocol and port specifications from the Internet Assigned Numbers Authority (IANA), <http://www.iana.org>. Sometimes applications other than the official ones use a port. The classifier generally cannot detect such use. If your network uses such applications, you can usually resolve this problem by renaming the application, in the application classifier, to indicate the actual application that uses this port on your network. For example, if you use port 3128 not for its standard use for a Squid web cache, but for a SOCKS proxy, you could rename the Squid (TCP) application to S OCKS (Port 3128) for clarity.

Applications must not have overlapping definitions. For example, if one application on your network uses TCP ports 3120 and 3128, and another application uses port 3120, only one Citrix SD-WAN WANOP application definition can include port 3120.

Configure application definitions

- Dynamic TCP, for applications using dynamic port allocations
- Ether type, for Ethernet packet types
- ICA Published App, for Virtual Apps/Virtual Desktops applications
- IP, for IP protocols such as ICMP or GRE
- TCP, for TCP applications
- UDP, for UDP applications
- Web Address, for specific Web sites or domains.

To configure an application definition:

1. Navigate to the **Configuration > Optimization Rules > Application Classifiers** and click **Add**.

The screenshot shows the 'Create Application' page in the Citrix SD-WAN WANOP 11.0 configuration interface. The page has tabs for Dashboard, Monitoring, and Configuration. The Configuration tab is active. A 'Back' button is at the top left. The form fields include: Name* (Viber), Description (messaging), Application Group* (with 'Available (25)' and 'Configured (2)' lists), Classification Type* (TCP), and Port* (5243). 'Create' and 'Close' buttons are at the bottom.

2. On the **Create Application** page, set the following parameters:

- **Name** - Name of the application classifier. Must begin with an ASCII alphanumeric or underscore (_) character, and must contain only ASCII alphanumeric, underscore, hash (#), period (.), colon (:), at (@), equals (=), and hyphen (-) characters. Maximum length: 31 characters.
- **Description** - Description of the application classifier.
- **Application Group** - The application classifier belongs to this application group. Application groups are a set of predefined groups of applications that are categorized on the basis of their functionality.
- **Classification Type** - The high-level classification you want to use for this application classifier. The high-level classification is mostly done on the basis of the port that an application uses.
- **Port** - The port number to be used. You can enter a range, a list or a number between 0 and 65535.

3. Click **Create**.

The **Application Classifiers** page lists all the applications recognized by the SD-WAN WANOP classifier.

The **Application Classifiers** page lists all the applications recognized by the SD-WAN WANOP classifier.

Tip

Click **Auto Discover** to allow any Citrix published applications seen in the data stream to be

added to the application list automatically. Once discovered, they will show up in reports and can be used for traffic-shaping policies.

Service classes

March 12, 2021

Service classes are assigned traffic-shaping policies and acceleration policies to be used for all connections that match the service-class definition. Service classes can be based on the following parameters:

- Applications
- IP or VLAN addresses
- DSCP bits
- SSL profiles

The default service-class definitions are recommended as a starting point. Modify them if they prove inadequate for your links.

The service classes are defined in an ordered list. The first definition that matches the traffic being processed becomes the service class for the traffic.

Differences between acceleration decisions and traffic shaping policies

To make an acceleration decision, the Citrix SD-WAN WANOP appliance examines the initial SYN packet of each TCP connection to determine whether the connection is a candidate for acceleration. The SYN packet contains no payload, only headers, so the acceleration decision must be based on the contents of the SYN packet's headers, such as the destination port or destination IP address of the connection. Acceleration, once applied, lasts for the duration of the connection.

Unlike acceleration decisions, traffic-shaping policies can be based on the contents of the connection's data stream. Depending on how long it takes for the application classifier to receive enough data for a final classification, a connection might be reclassified during its lifetime.

For example, the first packet in an HTTP connection to <http://www.example.com> is an SYN packet that contains a header but no payload. The header has an IP destination port of 80, which matches the HTTP: Internet service class definition, so the acceleration engine bases its acceleration decision, in this case, none (no acceleration) on that service class.

The traffic shaper uses the traffic-shaping policy from the HTTP: Internet service-class, but this decision is temporary. The first payload packet contains the string GET <http://www.example.com>,

which matches the example application definition in the application classifier. The service class that includes the example application is selected by the traffic shaper, instead the service class that includes HTTP: Internet, and the traffic shaper uses the service-class policy named in that service-class definition.

Note

Regardless of the service class policy, the reporting feature tracks the usage of the example application.

Important

All traffic is associated with an application and a service class, and all service classes have a traffic shaping policy, but only TCP connections have an acceleration policy other than none.

Configure service class definitions

Because service-class definitions are an ordered list, a definition that is an exception to a general case must precede the more general definition on the service-class page. The first definition whose rule matches the traffic is the one that is applied. For example:

- Service classes based on URLs must precede the HTTP service classes in the service-class list, because any URL-based rule also matches the HTTP service class. Therefore, putting the HTTP service class first would prevent the URL-based rules or published application-based rules from ever being used.
- Similarly, service classes based on ICA (Virtual Apps/Virtual Desktops) published applications must precede the Citrix service class.

Because all URL-based rules match the HTTP service class, putting the HTTP service class above them would result in the URL-based rules or published application-based rules never being used.

Configuration Overview > Optimization Rules > Service Classes					
<div> Add Edit Delete Update Order Filter Rules </div> <div>Show User Modified Service Classes Only</div>					
Order	Name	Status	Acceleration Policy	Traffic Shaping Policy	Appflow Reporting Status
1	ICA	Enabled	disk	ICA Priorities	Enabled
2	Web (Private)	Enabled	disk	Default Policy	Enabled
3	Web (Private-Secure)	Enabled	Flow Control Only	Default Policy	Enabled
4	Web (Internet)	Enabled	disk	Default Policy	Enabled
5	Web (Internet-Secure)	Enabled	Flow Control Only	Default Policy	Enabled
6	CIFS	Enabled	disk	Default Policy	Enabled
7	NFS	Enabled	disk	Default Policy	Enabled
8	Microsoft Exchange (MAPI)	Enabled	disk	Default Policy	Enabled
9	Mail (Other)	Enabled	disk	Default Policy	Enabled
10	VOIP and Multimedia	Enabled	None	VOIP Traffic	Enabled
11	VOIP Webcam	Enabled	None	High Priority Traffic	Enabled
12	FTP Data	Enabled	disk	Low Priority Traffic	Enabled
13	FTP Control	Enabled	Flow Control Only	Default Policy	Enabled
14	Instant Messaging	Enabled	disk	Default Policy	Enabled
15	Session Applications	Enabled	Flow Control Only	Default Policy	Enabled
16	Directory and Security	Enabled	Flow Control Only	Default Policy	Enabled
17	Database Applications	Enabled	Flow Control Only	Default Policy	Enabled
18	Secure Applications	Enabled	Flow Control Only	Default Policy	Enabled
19	Iperf	Enabled	Flow Control Only	Low Priority Traffic	Enabled
20	NetApp SnapMirror	Enabled	memory	Default Policy	Enabled
21	Other TCP Traffic	Enabled	None	Default Policy	Enabled
22	Unclassified Traffic	Enabled	None	Default Policy	Enabled

To create an RPC over HTTP service class and bind the SSL profile to it:

1. Navigate to **Configuration > Optimization Rules > Service Classes** and click **Add**.

The screenshot shows the 'Create Service Classes' configuration page in the Citrix SD-WAN WANOP 11.0 interface. The page has a dark header with 'Dashboard', 'Monitoring', and 'Configuration' tabs. The 'Configuration' tab is active. Below the header, there's a 'Back' button. The main content area is titled 'Create Service Classes' and contains several configuration options: 'Name*' (text input with 'RPC over HTTP'), 'Enabled' (checked checkbox), 'Acceleration Policy*' (dropdown menu with 'disk' selected), 'Traffic Shaping Policy' (radio buttons for 'Single Policy' and 'Per Link Policy', with 'Single Policy' selected), 'Enable AppFlow Reporting' (checked checkbox), 'Exclude from SSL Tunnel' (unchecked checkbox), and 'Default Policy' (dropdown menu). Below these options is a 'Filter Rules' section with 'Add', 'Edit', and 'Delete' buttons. At the bottom, there's a table with columns: 'Application', 'Source IP Address', 'Destination IP Address', 'VLANs', 'DiffServ DSCP Bits', 'Direction', and 'SSL Profiles'. The table is currently empty, showing 'No items'. At the very bottom, there are 'Create' and 'Close' buttons.

2. In the **Name** field, enter a name for the service class.
3. Make sure that the **Enabled** option is selected.
4. From the **Acceleration Policy** list, select an acceleration policy. **Memory** and **Disk** specify where to store the traffic history used for compression. **Disk** is usually the best choice, because the appliance automatically selects disk or memory, depending on which is more appropriate for the traffic. **Memory** specifies memory only. Select **Flow Control Only** to disable compression but enable flow-control acceleration. Select this for services that are always encrypted, and for the FTP control channel. **None** is used only for uncompressible encrypted traffic and real-time video.
5. Select **Enable AppFlow Reporting** to enable AppFlow reporting for this Service Class. Information from this service class is included in any AppFlow reports. AppFlow is an industry standard for unlocking application transactional data processed by the network infrastructure. The WAN Optimization AppFlow interface works with any AppFlow collector to generate reports. The collector receives detailed information from the appliance, using the AppFlow open standard.
6. Select **Exclude from the SSL Tunnel** to exclude traffic associated with the Service Class from SSL Tunneling.
7. In the traffic shaping policy list, make sure that **Default Policy** option is selected. Traffic shaping policies have a weighted priority and other attributes that determine how matching traffic will be treated, relative to other traffic. Most service classes are set to Default Policy, but higher-priority traffic can be assigned a higher-priority traffic-shaping policy, and lower-priority traffic can be assigned a lower-priority policy.
8. In the Filter Rules section, click **Add** to create filter rule that has Any as the default value for all parameters. If a rule is evaluated as TRUE for a given connection, the connection is assigned to that service class. Filter rules for most service classes consist solely of a list of applications,

but rules can also include IP addresses, VLAN tags, DSCP values, and SSL profile names. All the fields in a rule default to Any (a wildcard). Fields within a rule are ANDed together.

9. Click **Add** to add filter rules.

Filter Rules

Application Group*
Email and Collaboration

Application Classifiers*

Available (27) Select All

NNTP	+
Novell Groupwise	+
POP3 (secure)	+
POP3 Kerberos	+
SMTP (clear)	+

Configured (2) Remove All

POP3 (clear)	-
Biff	-

Source IP Address
10.102.29.230 + ▼
No items

Direction*
Unidirectional ▼

Destination IP Address
+ ▼
No items

VLANs
+ ▼
No items

DiffServ DSCP Bits*
Best Effort ▼

10. From the **Application Group** list, select **Email and Collaboration**.
11. From the **Available** list, select the required applications.
12. Move the selected applications to the **Configured** list.
13. In the **Source IP Addresses** field, add the client IP addresses.
14. From the **Direction** list, select the direction of the traffic.
15. From the **SSL Profiles** list, select the SSL profile you have created
16. Click **Create**.

Note

- You must configure and bind an SSL profile to the service class only on the datacenter-side appliance.
- Only the service classes that have their filter rules direction set to unidirectional can be associated with SSL profiles.

Traffic shaping

March 12, 2021

Apr 18, 2018

Traffic shaping allows you to regulate the network traffic flow to assure a certain level of quality of service (QoS). You can regulate the flow of packets into a network (bandwidth throttling) or out of a network (rate limiting).

Using traffic shaping policies you can set the priority of different link traffic and send traffic onto the link at a rate close to, but no greater than, the link speed. Unlike acceleration, which applies only to TCP/IP traffic, the traffic shaper handles all traffic on the link.

You can set high bandwidth for traffic flows that are considered more important than the rest of the traffic flows, allowing you to optimally use the scarce link resources.

The traffic shaping is based on weighted fair queuing, which gives each service class its fair share of the link bandwidth. If the link is idle, any connection (in any service class) can use the entire link. When multiple connections are competing for the link bandwidth, the traffic shaper applies traffic shaping policies to determine the right mix of traffic.

For information on Weighted Fair Queuing, see [Weighted Fair Queuing](#).

To configure traffic shaping:

1. Configure the Link definition.

The Link definition is used by the traffic shaper to determine the send and receive link speed and other link related information. For more information on how traffic shaper uses link definition and how to configure link definitions, see [Link Definitions](#).

2. Configure the Application definition.

The traffic flowing through the link is examined by the application classifier to determine which application it belongs to and then the application is looked up in the service class list to determine which service class it belongs to. For more information on application classification and how to configure application definition see [Traffic Classification](#).

3. Create a traffic shaping policy.

You can use the default traffic shaping policies or create a new policy to set the weighted priority and other parameters as per your network requirements. For information on creating traffic shaping policy, see [Traffic Shaping Policies](#).

4. Configure a service class definition and associate the traffic shaping policy to the service class.

For information on configuring service class definition, see [Service Classes](#).

Some highlights of the traffic shaper:

- All WAN traffic is subject to traffic shaping: accelerated connections, unaccelerated connections, and non-TCP traffic such as UDP flows and GRE streams.
- The algorithm is weighted fair queuing, in which the administrator assigns each service class a priority. Each service class represents a bandwidth pool, entitled to a minimum fraction of the link speed, equal to $(\text{my_priority} / \text{sum_of_all_priorities})$. A service class with a weighted priority of 100 gets twice as much bandwidth as a service class with a weighted priority of 50. You can assign weights from 1 through 256.
- Each connection within a service class gets an equal share of the bandwidth allotted to that service class.
- Each connection gets its fair share of the link bandwidth, because priorities are applied to the actual WAN data transferred, after compression. For example, if you have two data streams with the same priority, one achieving 10:1 compression and the other achieving 2:1 compression, users see a 5:1 difference in throughput, even though the WAN link usage of the two connections is identical. In practice, this disparity is desirable, because WAN bandwidth, not application bandwidth, is the scarce resource that needs to be managed.
- Traffic-shaping policies apply equally to both accelerated and unaccelerated traffic. For example, an accelerated Virtual Apps connection and an unaccelerated Virtual Apps connection both receive traffic shaping, so both can have an elevated priority compared to bulk traffic. As another example, time-sensitive non-TCP traffic, such as VoIP (which uses the UDP protocol) can be expedited.
- Traffic shaping is applied to the WAN link in both the sending and receiving directions, to both accelerated and non-accelerated traffic. This feature prevents congestion and increased latency even when the other side of the link is not equipped with a Citrix SD-WAN WANOP appliance. For example, Internet downloads can be prioritized and managed.
- The traffic-shaping policy for a service class can be specified on a per-link basis if desired.
- In addition to shaping the traffic directly, the traffic shaper can affect it indirectly by setting the Differentiated Services Code Point (DSCP) field to inform downstream routers about the type of traffic shaping each packet requires.

Weighted fair queuing

March 12, 2021

In any link, the bottleneck gateway determines the queuing discipline, because data in the non-bottleneck gateways does not back up. Without pending data in the queues, the queuing protocol is irrelevant.


Most IP networks use deep FIFO queues. If traffic arrives faster than the bottleneck speed, the queues fill up and all packets suffer increased queuing times. Sometimes the traffic is divided into a few different classes with separate FIFOs, but the problem remains. A single connection sending too much data can cause large delays, packet losses, or both for all the other connections in its class.

A Citrix SD-WAN WANOP appliance uses *weighted fair queuing*, which provides a separate queue for each connection. With fair queuing, a too-fast connection can overflow only its own queue. It has no effect on other connections. But because of lossless flow control, there is no such thing as a too-fast connection, and queues do not overflow.

The result is that each connection has its traffic metered into the link in a fair manner, and the link as a whole has an optimal bandwidth and latency profile.

The following figure shows the effect of fair queuing. A connection that requires less than its fair share of bandwidth (the bottom connection) gets as much bandwidth as it attempts to use. In addition, it has very little queuing latency. Connections that attempt to use more than their fair share get their fair share, plus any bandwidth left over from connections that use less than their fair share.

Figure 1. Fair Queuing in Action



The optimal latency profile provides users of interactive and transactional applications with ideal performance, even when they are sharing the link with multiple bulk transfers. The combination of lossless, transparent flow control and fair queuing enables you to combine all kinds of traffic over the same link safely and transparently.

The difference between weighted fair queuing and unweighted fair queuing is that weighted fair queuing includes the option of giving some traffic a higher priority (weight) than others. Traffic with a weight of two receives twice the bandwidth of traffic with a weight of one. In a Citrix SD-WAN WANOP

configuration, the weights are assigned in traffic-shaping policies.

Traffic shaping policies

March 12, 2021

Every service class definition is associated to a traffic-shaping policy, which sets parameters for traffic of the associated service class. You can create and configure traffic shaping policies for sites with special needs, but the default policy settings works fine for most installations, providing the following benefits:

- Increased responsiveness for interactive traffic such as Citrix Virtual Apps and Desktops.
- Protection of latency- and jitter-sensitive VoIP traffic.
- No “hitting the wall” during peak periods. You get usable performance even under extreme load.
- Improved bandwidth utilization by allowing bulk transfers to fill the link with whatever bandwidth is left over from interactive tasks.
- Extension of the benefits of fair queuing to all traffic

A Citrix SD-WAN WANOP appliance is shipped with factory-default traffic shaping policies that span a broad range of priorities. These policies are listed in the **Traffic Shaping Policies** page. Apart from the **Default Policy**, the other factory-default policies cannot be edited or deleted. The reason is to ensure that they have the same meaning on all appliances. To make changes, create a new traffic-shaping policy with the new parameters and change the appropriate service-class definitions to refer to the new traffic-shaping policy.

To create a traffic shaping policy:

1. In the SD-WAN WANOP management UI, navigate to **Configuration > Optimization Rules > Traffic Shaping Policies** and click **Add**.

Configuration Overview > Optimization Rules > Traffic Shaping Policies

Buttons: Add, Edit, Delete

Show User Modified Traffic Shaping Policies Only

Name		Voice Optimized	DiffServ/TOS	Maximum Incoming Bandwidth	Maximum Outgoing Bandwidth
VOIP Traffic	Very High (Priority 256)	✓	Expedited Forward...	75 %	75 %
Very High Priority Traffic	Very High (Priority 256)	✗	Disabled	0	0
High Priority Traffic	High (Priority 128)	✗	Disabled	0	0
Medium High Priority Traffic	Medium High (Priority 64)	✗	Disabled	0	0
Medium Priority Traffic	Medium (Priority 32)	✗	Disabled	0	0
Medium Low Priority Traffic	Medium Low (Priority 16)	✗	Disabled	0	0
Low Priority Traffic	Low (Priority 8)	✗	Disabled	0	0
Very Low Priority Traffic	Very Low (Priority 4)	✗	Disabled	0	0
ICA Priorities	Very High (Priority 256)	✗	Disabled	0	0
Default Policy	Medium (Priority 32)	✗	Disabled	0	0
TSP1	High (Priority 128)	✗	Disabled	10 %	10 %

2. In the **Create Traffic Shaping Policy** page, enter values for the following parameters:

- **Name**—The name of the new policy. Must be unique.
- **Weighted Priority**—You can select an existing priority value or can select a custom value between 1 and 256. A connection with a priority of 256 will get 256 times the bandwidth share as a connection with a priority of 1.
- **Optimize for Voice**—If selected, this policy will have effectively infinite priority. This is highly undesirable for most traffic, since it will prevent meaningful traffic shaping and will cause data starvation for other traffic if there is enough “optimized for voice” traffic to fill the link. Use only for VoIP, and always use in conjunction with a bandwidth limit on the policy (for example, 50% of the link speed)

Note

Voice Optimization cannot be configured while ICA Priorities are set.

The screenshot shows the 'Create Traffic Shaping Policy' configuration page. The 'Name' field is 'TSP1'. 'Weighted Priority' is set to 'Very Low'. 'Optimize for Voice' is checked. 'DiffServ/TOS' is set to 'AF12 - Silver' and 'DSCP 12 (binary: 001100)'. 'Bandwidth Limit' is set to 'By Percentage of Link Bandwidth'. 'Maximum Incoming Bandwidth Rate (%)' and 'Maximum Outgoing Bandwidth Rate (%)' are both set to 50. There are sections for 'ICA Priority Settings' and 'ICA DiffServ/TOS Settings', both with checkboxes to set them. The page has a navigation bar with 'Dashboard', 'Monitoring', 'Configuration', 'Downloads', and 'Notifications (1)'. At the bottom, there are 'Add' and 'Cancel' buttons.

- **Diffserv/TOS**—Sets the DSCP bits on output packets to the selected value. Used to control downstream routers.
- **Bandwidth Limit**—Prevents the traffic using this policy from exceeding the specified bandwidth, stated either as a percentage of link speed or as an absolute value. Citrix recommends specifying a percentage, so that the same definition can apply to links of different speeds. This feature can leave bandwidth unused. For example, a policy set to 50% of link speed does not allow the affected traffic to use more than 50% of the link, even if the link is otherwise idle. Throttling traffic in this way is inconsistent with maximum performance, so this feature is rarely used, except with VoIP traffic with the Optimize for Voice

setting.

Note

Configuring **Bandwidth Limit** is applicable only for Citrix SD-WAN WANOP edition. For Citrix SD-WAN PE edition, the **Bandwidth Limit** parameter is disabled by default.

- **Set ICA priorities**—If this policy is used for Citrix Virtual Apps/Virtual Desktops traffic, the traffic’s internal priority for Real-time, Interactive, Bulk Transfer and Background traffic is overwritten by the priority set here.

ICA Priority Settings

☒ Set ICA Priority

0 - Realtime*

High

*

Priority 128

1 - Interactive*

Medium High

*

Priority 64

2 - Bulk Transfer*

Medium Low

*

Priority 16

3 - Background*

Very Low

*

Priority 4

- **Set ICA DiffServ/TOS:** For ICA (Virtual Apps/Virtual Desktops) traffic, each of the four ICA priority values can be tagged with a different DSCP value. This capability is particularly useful with the new Multistream ICA feature, in which the Virtual Apps or Virtual Desktops client uses different connections for different priority levels.

ICA DiffServ/TOS Settings	
<input checked="" type="checkbox"/> Set ICA DiffServ/TOS	
Multi-Stream (0 - Realtime)*	*
AF11 - Gold ▼	DSCP 10 (binary: 001010) ▼
Multi-Stream (1 - Interactive)*	*
AF21 - Gold ▼	DSCP 18 (binary: 0010010) ▼
Multi-Stream (2 - Bulk Transfer)*	*
AF12 - Silver ▼	DSCP 12 (binary: 001100) ▼
Multi-Stream (3 - Background)*	*
AF13 - Bronze ▼	DSCP 14 (binary: 001110) ▼
Single-Stream (All priorities)*	*
AF33 - Bronze ▼	DSCP 30 (binary: 0011110) ▼

3. Click **Add**. The newly Created Traffic Shaping Policy is listed in the Traffic Shaping Policies list.

You can now associate the traffic shaping policy to a service class, for more information see [Service Classes](#).

Video caching

March 12, 2021

Many organizations use video for communications that are not time sensitive, (for example, training sessions and prerecorded messages to employees). Communicating messages through videos is not only cost effective but also convenient when the audience is spread across time zones. However, videos consume a lot of bandwidth when played over the Internet. Insufficient bandwidth causes latency, which affects the user experience and degrades the impact of video communication.

Video caching improves the viewing experience for HTTP video streams, especially on slower links. The video cache is maintained on the local Citrix SD-WAN WANOP appliance. When a local user views a video that has already been cached, the appliance can deliver the cached copy at full LAN speed.

After you configure the appliance to cache videos, it caches the videos viewed by your users. You can also use the pre-population option to fetch selected videos from the local video server in anticipation

of later use.

The video caching feature uses an intercepting proxy cache to examine all HTTP requests. Requests that meet the requirements listed below are cached. Videos are not served from the cache unless they are evaluated as fresh by the cache engine. Otherwise, they are fetched again for the viewer, and the previously cached version is overwritten.

Latest content guaranteed. Every time a video is viewed, the cache checks the origin server, and if the video has changed, the cached content is discarded and the new content is downloaded.

Note

Caching is now transparent. That is, the IP address of both the client and the server are maintained end-to-end. In earlier releases, the IP address of the Citrix SD-WAN WANOP appliance was displayed as the source address.

A video is cached when all of the following criteria are met:

- The protocol used to stream the video is HTTP. By default, port 80 is configured for video caching. However, if you have configured another port, such as 8080 for a web server, you must specify this port for caching videos.
- You have added video sources from which you want to cache videos. By default, YouTube, Vimeo, Youku, Dailymotion, and Metacafe video sources are added to the appliance, but only YouTube and Vimeo are enabled. If you want to cache videos from any of the other default sources, you must enable them. When adding new video sources, you can enable them as you add them.
- Besides YouTube, Vimeo, Metacafe, Dailymotion, and Youku, you can specify additional websites, IP addresses, or subnets as video sources. Note that these websites should not have any avoidance mechanisms, such as adding random characters to a URL.
- The video must be in one of the recognized video formats and have the one of the following file extensions: .3gp, .avi, .dat, .divx, .dvv, .dv-avi, .flv, .fmv, .h264, .hdmov, .m15, .m1v, .m21, .m2a, .m2v, .m4e, .m4v, .m75, .moov, .mov, .movie, .mp21, mp2v, .mp4, .mp4v, .mpe, .mpeg, mpeg4, mpg, mpg2, .mpv, .mts, .ogg, .ogv, .qt, .qtm, .ra, .rm, .ram, .rmd, .rms, rmvb, .rp, rv, .swf, .ts, .vfw, .vob, .webm, .wm, .wma, .wmv, and .wtv.

Platforms supported

The video caching feature is supported by the following appliances:

- SD-WAN WANOP 600 appliance with 1 Mbps and 2 Mbps bandwidth license models.
- SD-WAN WANOP 800 appliance with all bandwidth license models.

- SD-WAN WANOP 1000 appliance with Windows Server, with all bandwidth license models.
- SD-WAN WANOP 2000 appliance with all the bandwidth license models.
- SD-WAN WANOP 2000 appliance with Windows Server, with all bandwidth license models.
- SD-WAN WANOP 3000 appliance with all the bandwidth license models.
- SD-WAN WANOP VPX and SD-WAN WANOP VPX for Amazon

Video server supported

Video caching feature is supported on Adobe Flash Media Server 4.5 or later. Additionally, any video server that serves videos over HTTP as static links are supported for video caching.

Deployment modes supported

Video caching is supported in inline, inline within VLAN trunk ports, virtual inline, and WCCP deployment modes.

Considerations for using the video caching feature


Following are some points to be aware of when using the video caching feature.

- If any of the supported websites change the way they present content, the video caching benefit for those sites might not be achieved until the video caching policy file is updated. For such occasional changes, Citrix provides an updated video caching policy file. To use it, see [Upgrading the Video Caching Policy File](#).
- Some video websites might use different file formats for the same video, depending on the operating system or the browser used to access the video. This might result in a cache miss.
- Some video websites, such as YouTube, adapt to the network conditions. The quality of a video can therefore depend on the network conditions at the time it is cached.

Video caching scenarios

March 12, 2021

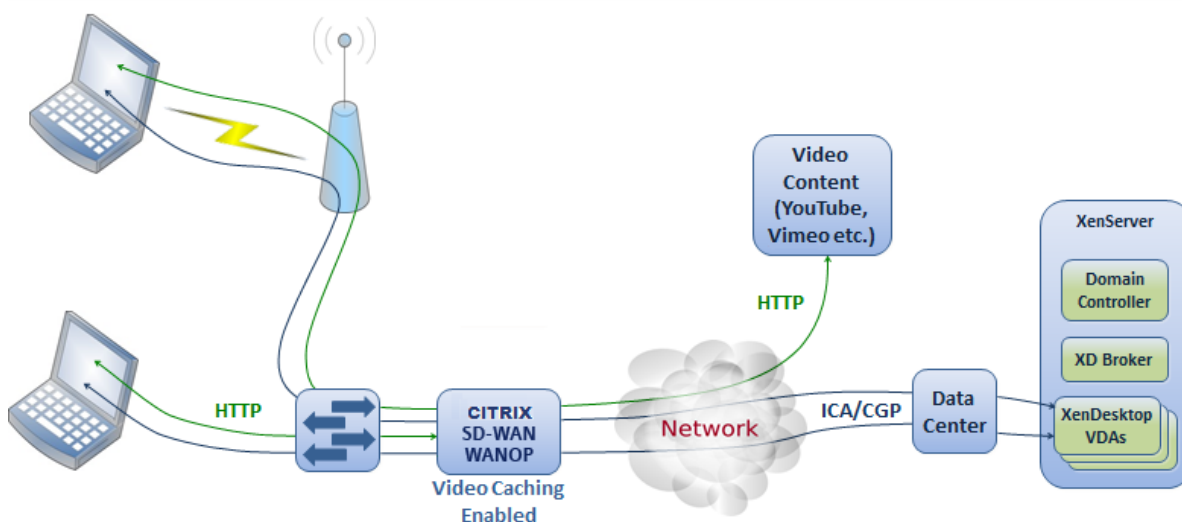
You can deploy video caching on Citrix SD-WAN WANOP appliance under the following scenarios:

Branch office access

In this use case, users access the Internet through the web browsers on their computers. Those requests that involve video content from an enabled site, such as Vimeo, are cached on the local Citrix SD-WAN WANOP appliance. Any subsequent access of the same video results in cache hits on the local appliance, allowing the video to be delivered at LAN speed and without waiting for the remote server.


Unlike other Citrix SD-WAN WANOP features, which accelerate traffic between paired devices, this feature is a single-ended operation that requires only the local appliance, with access to the video website.

Branch office with Citrix Virtual Apps and Desktops users using HDX MediaStream flash redirection



HDX flash redirection is a feature of Citrix Virtual Apps and Desktops. Instead of rendering the video on the remote Virtual Desktops display by using the server-side or datacenter Internet, flash videos are tunneled to the local system through this feature. The video is streamed to the actual client machine and is rendered on the actual client, using the branch office Internet. Enabling the video caching feature on the branch-side Citrix SD-WAN WANOP appliance can give users a significantly improved viewing experience. Additionally, enabling the feature reduces the bandwidth requirement for streaming videos.

Enterprise HTTP video web server



In this use case, users access the video web servers from the datacenter. When you enable the video caching feature on the branch-side Citrix SD-WAN WANOP appliance, the user request is served from cache of the branch-side Citrix SD-WAN WANOP appliance. This helps reduce network traffic to the datacenter Citrix SD-WAN WANOP appliance. As a result, the bandwidth of the datacenter Citrix SD-WAN WANOP appliance can be used to serve traffic for other branches.

Configure video caching

March 16, 2021

You can configure the video caching feature through either the Citrix SD-WAN WANOP graphical user interface or command line interface. By default, the appliance is configured to cache videos from YouTube and Vimeo. Youku, Metacafe, and Dailymotion are also configured on the appliance by default. All you have to do is enable them. You can add video websites, such as an internal website serving video tutorials or other information .

Note

Video caching is an optional feature that is not enabled by default. You need not enable it unless you have a substantial amount of HTTP video traffic .

Prerequisites

To configure video caching on the appliance, make sure that the following prerequisites are met:

- You have configured appropriate IP address for accelerated bridge port you are planning to use for video caching.
- You can ping the apA/apB gateway from the appliance.
- DNS Server details are accurate.
- Appliance can resolve the DNS name `www.Citrix.com`.
- The Citrix SD-WAN WANOP apX IP address has an HTTP access in your corporate network.
- If the appliance is deployed between the trunk ports of two network devices, you must specify the VLAN ID with the IP address to be used by the appliance to send HTTP requests on the Network Configuration page.
- For **Web (Internet)** and **Web (Private) service classes**, the **Acceleration Policy** setting should not be set to **None**.

Enable video caching feature

Before you can start using the video caching feature, you must enable it.

To enable video caching:

1. Navigate to the **Configuration > Appliance Settings > Network Adapters**, under **Management Settings** section verify and ensure that the Primary DNS Server details are accurate and the appliance is able to resolve the DNS name `www.Citrix.com`. Click the edit icon to change the settings.

The screenshot shows the Citrix SD-WAN WANOP 11.0 Configuration page. The left sidebar contains a tree view with the following items: Appliance Settings, Features, Licensing, Advanced Deployments, Network Adapters (selected), Ethernet, NetScaler SD-WAN WANOP Clients, User Administration, Date/Time Settings, Logging, Notifications, SNMP, AppFlow, Optimization Rules, Video Caching, Secure Acceleration, Diagnostics, and Maintenance. The main content area is titled 'Configuration Overview > Appliance Settings > Network Adapters'. It contains two sections: 'Management Settings' and 'Network Adapters'. The 'Management Settings' section has a 'Host Name' field with the value 'vpx-175', a 'DHCP for DNS' checkbox, a 'Primary DNS Server' field with the value '10.102.29.16', and a 'Secondary DNS Server' field with the value '10.102.29.70'. There are 'Save' and 'Cancel' buttons. The 'Network Adapters' section has an 'Edit' button and a table with the following data:

Name	Status	DHCP	IPv4 Address	IPv4 Gateway	IPv6 Address	IPv6 Gateway	SSH	Web	VLAN	VLAN Group
apA	Enabled	Disabled	192.168.10.20/24	192.168.10.1	::	::	Enabled	Enabled	Disabled	0
Primary	Enabled	Disabled	10.102.203.175/24	10.102.203.1	::	::	Enabled	Enabled	Disabled	0

2. Navigate to the **Configuration > Appliance Settings > Network Adapters**. In the **Network Adapters** section, select an acceleration pair (for example apA) and klik **Edit**.

Ensure that the IP addresses, network mask, and default gateway IP addresses specified for the accelerated pair are accurate.

Modify Adapter

Modify Adapter

Name

apA

☒ Enabled

☐ DHCP for IPv4 Address

IPv4 Address/MaskBits*

10.102.29.88/32

IPv4 Gateway

10.102.29.1

IPv6 Address/Prefixlength

::

IPv6 Gateway

::

Management Access

☒ SSH

☒ Web

VLAN

☐ VLAN

Save

Close

3. Navigate to the **Configuration > Appliance Settings > Features** page and enable **Video Caching** feature.

A confirmation dialog box appears, click **Yes**.

Appliance Settings

Features

Licensing

Advanced Deployments

Network Adapters

- Ethernet
- NetScaler SD-WAN WANOP Clients

User Administration

Date/Time Settings

Logging

Notifications

SNMP

AppFlow

Optimization Rules

Video Caching

Secure Acceleration

Diagnostics

Maintenance

Configuration Overview > Appliance Settings > Appliance Settings

Features

Enable Disable Edit

Name	State	Status
Traffic Processing	Disabled	License is not available
Traffic Acceleration	Enabled	Disabled - due to disabled traffic processing
Traffic Shaping	Enabled	Disabled - due to disabled traffic processing
Traffic Bridging	Enabled	Enabled
IPv6 Acceleration	Enabled	Disabled - due to disabled traffic processing
AppFlow	Enabled	Enabled
RPC Over HTTP	Enabled	Disabled - due to disabled traffic processing
Native Mapi	Enabled	Disabled - due to disabled traffic processing
ICA Multi-stream	Disabled	Disabled
MAPI Cross Protocol Optimization	Disabled	Disabled
SCPS	Disabled	Disabled
Secure Partner	Disabled	Disabled
SNMP	Enabled	Enabled
SSH Access	Enabled	Enabled
SSL Optimization	Enabled	Disabled - due to disabled traffic processing
Syslog	Enabled	Enabled
User Data Store Encryption	Disabled	Disabled
Video Caching	Disabled	Disabled
NetScaler SD-WAN WANOP Client	Enabled	Disabled -Requires IP configuration
WCCP	Enabled	Disabled - due to disabled traffic processing
CIFS Protocol Optimization	Disabled	Disabled - due to disabled traffic processing

Note

The service restarts and a new caching partition is created. If you are enabling the feature for the first time on the appliance, a new partition is created by reducing the disk space allocated to other disk based compression. The disk based compression history is reset, and existing connections are terminated.

4. Alternatively, you can navigate to **Configuration > Optimization Rules > Video Caching** and click **Enable**.

Dashboard Monitoring Configuration Downloads Notifications (8)

Configuration Overview > Video Caching

Video Caching is Disabled

Enable

Video Caching

Settings

Set Global Parameters

Clear Video Cache

Add video websites

The appliance is configured to cache videos from YouTube and Vimeo, and is partially configured to cache videos from Youku, Metacafe, and Dailymotion. To cache videos from any of the latter three sites, you must enable the site. A video from an enabled website is cached as soon as a user accesses it. You can configure additional video websites that do not require URL rewrite by adding their host


names or IP addresses to the Video Source list on the appliance. You can also include custom sites that do not have any cache avoidance mechanisms.

You must enable these video sources before the appliance can cache videos from them.

The video caching feature uses video sources for the configuration workflow. If you configure any of the video sources with a host name or a website/hostname, the appliance proxies all HTTP traffic that flows through the appliance. However, if you configure all video sources with IP addresses only, the appliance proxies and caches only these IP addresses. Regardless of whether you use host names or IP addresses, if your organization does not permit access to the YouTube, Vimeo, Dailymotion, Metacafe, and Youku websites, make sure that you disable these video sources.

To enable a video source:

1. Navigate to **Configuration > Optimization Rules > Video Caching > Video Sources**.
2. Select a video source from the list click **Modify**.




The screenshot shows a 'Modify Cache Status' dialog box. The title bar and header both read 'Modify Cache Status'. The 'Video Source' field contains 'youku.com'. The 'Cache Status*' dropdown menu is set to 'Enabled'. At the bottom, there are 'Save' and 'Close' buttons.

3. In the **Cache Status** drop-down box, select **Enable** and click **Save**.

To add a video source:

1. Navigate to **Configuration > Optimization Rules > Video Caching > Video Sources**, click **Add**.
2. In the **Video Source** field, type the website name or IP address of the web server that you want to add to the video source list.
3. In the **Cache Status** list, ensure that **Enabled** is selected. You can select **Disabled** from this list if you want to enable video caching for this site at a later time.



Add Video Source

Add Video Source

Video Source

youtube.com

Cache Status *

Enabled ▼

Create **Close**

4. Click **Create**.

To delete a video source, select it from the **Video Sources** list and click **Delete**.

Video prepopulation

March 12, 2021

A Citrix SD-WAN WANOP appliance can download and cache videos from your internal video server before anyone views them. This feature is useful when you want to make sure that all users get the same benefits (for example when playing a self-training video scheduled at a specific time). You can schedule static URLs from which you want to fetch videos.

The fetched videos are stored in the video cache. As soon as a user sends a request for the URL, the video is served from the cache, even for the first access of the video.

To fetch videos in advance, you can perform the following tasks:

- Specify a URL from which you want to cache videos in advance.
- Schedule date and time at which to cache the videos.
- Schedule an interval at which you want to cache the videos.
- Manage the entries you have added to the list.

To download and cache a video in advance, you must specify the absolute path for the URL of a specific video or a video folder on which directory indexing is enabled.

Note

If you just add an entry to the video prepopulation tasks, the related video is downloaded and cached. However, when a client accesses the video, it is served from the video server and does not get caching benefits. To make sure that the client gets caching benefits, you must add the video server or IP address used in the prepopulation task to the video sources list.

To add a URL to cache videos in advance:

1. Navigate to **Configuration > Video Caching > Prepopulation** and click **Add**.

Add Prepopulation Entry

Add Prepopulation Entry

Name*

Example

URL*

http://example.com/ ?

Interface*

apA ▼

State

☒ Enable ☐ Disable

Schedule

☒ Now ☐ Later

Repeat*

Only Once ▼

Create Close

2. In the **Name** field, specify a name that you can use to identify the prepopulation entry.
3. In the **URL** field, specify the URL from which you want to cache one or more videos. The URL can be for a specific video or a video server. Make sure that you specify a complete URL or a video folder.

4. In the **Interface** field, select the accelerated bridge port to download videos from the URL.
5. Set **State** to **Enable** to receive state information. The various states and their description is provided in the table below.
6. You can either start downloading and caching videos from the URL to the appliance immediately, or download them at a scheduled time.
7. Click **Create**.

The following table describes the status messages:

Status	Description
Configured	Fetching video for caching before the first view is configured for the URL and a new task is added.
Connection timeout error	Connection to the server has timed out and there is no response from the server.
Error 301 - Moved Permanently	The video to be downloaded and cached has been permanently moved to another location.
Error 403 - Forbidden	Access to the video to be downloaded and cached is denied.
Error 404 - Not Found	The video to be downloaded and cached is not available at the link provided.
Error 504: Server unreachable	The URL you have specified is not reachable.
Successfully downloaded “x”file(s)	Download successful for the URL, and “x” number of media files are downloaded to the cache.
Failed to download “x”out of “y”files	Download failed for some of the media files from the URL.
Failed to download x files(s)	Failed to download any media file from the URL.
Download completed	Processing of all URLs for this entry is complete.
Download in progress	The download is in progress.
Starting	The appliance has started downloading media files from the URL.
Deleting this entry	The entry is being deleted from the list of URLs.
Failed to get Directory listing	Failed to get listing from the remote directory you specified.
Entry removed by clear cache operation	The entry has been purged by the clear cache operation.
Updating Status	The appliance is updating the status of the entry.

Status	Description
Schedule time elapsed	The scheduled time at which to download the remote object is past.
In-cache “x”/”y”files	On refreshing the status of an entry, the appliance has found that “x”number of files out of “y”number of files exist in the cache.
Interface ap”X”disabled for Video Caching	The bridge interface ap”X”is not enabled for Video Caching.
Refreshing status	The status of the entry is being refreshed.
Error 0	An unknown error has occurred while downloading the videos. Contact Citrix Technical Support team to resolve the issue.

Manage video caching prepopulation

You can manage video caching prepopulation to control how you want to download and cache videos from the URLs. You can perform the following tasks to manage video caching prepopulation:

- Start downloading videos before or after the scheduled date and time.
- Update the URL of an entry.
- Disable caching of videos from a URL entry.
- Schedule caching of videos from a URL entry.
- Update an interface for a URL entry.
- Refresh the status of a URL entry.
- Delete a URL entry.

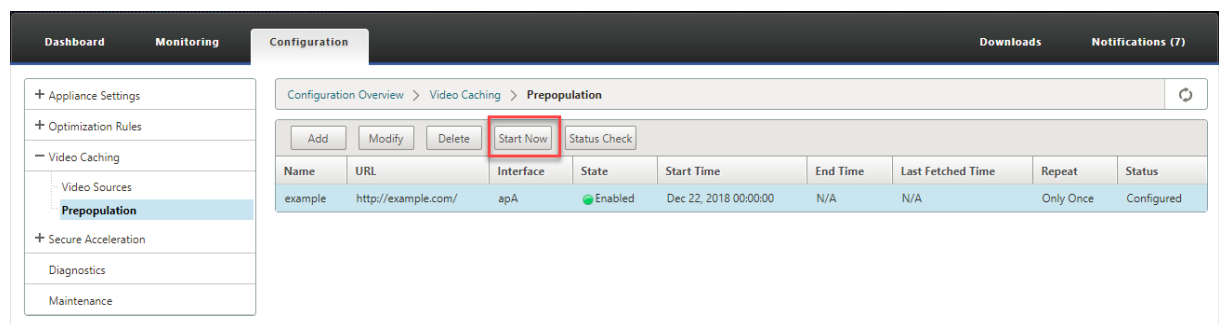
The following flowchart shows the flow control of the processes followed when managing various activities of the video prepopulation feature.



Download videos

If technical issues with a website or the URL that you have added interfere with scheduled downloading and caching, you can start downloading and caching videos when required at any time.

To download and cache a video immediately, navigate to **Configuration > Video Caching > Prepopulation**, select the entry for the video you want to cache, and then click **Start Now**. Updating the status of the video takes approximately one minute.



After you click Start Now, the Status column displays the status of the video downloads from the URL.

Update the URL of a prepopulation entry

After you have added a URL from which to download and cache video in advance, you can fine tune the URL for optimum results, such as reconfiguring the URL when the location of videos changes or the name of the media file is changed in the source.

To update a URL:

1. Navigate to the **Configuration > Video Caching > Prepopulation** page.
2. Select the entry that you want to update and click **Modify**.
3. In the URL field, specify the new URL.
4. Click **OK**.

Disable caching of videos from a URL in a prepopulation entry

If you want to periodically prepopulate the cache with videos from a given URL, you need not delete the entry. You can disable it, and then enable it when needed.

To disable an entry:

1. Navigate to **Configuration > Video Caching > Prepopulation** page.
2. Select the entry that you want to update and click **Modify**.
3. From State, select the **Disable** option.
4. Click **OK**.

Schedule caching of videos from a URL in a prepopulation entry

You can schedule the date and time at which you want to start downloading and caching videos from the URL to the appliance. For example, you might want to fetch videos just before you expect users to

start accessing them. That not only saves disk space, but also puts the latest versions of the videos in the cache.

To schedule caching from a URL:

1. Navigate to **Configuration > Video Caching > Prepopulation** page.
2. Select the entry that you want to update and click **Modify**.
3. From **Schedule**, select the **Later** option.
4. In the **Start** field, specify the date and time at which you want to videos from the URL to be downloaded. The format for the date and time is YYYY-MM-DD HH:MM:SS.
5. From the **Repeat** list, select the frequency of downloading and caching the videos. The available options are:
 - **Only Once**: Download videos from the URL only once, at the scheduled date and time.
 - **Daily**: Download videos from the URL every day, starting with the scheduled date and time. The download starts every day at the start time that you specify.
 - **Weekly**: Download videos from the URL once in a week, starting with the scheduled date and time. The download starts every week on the day and time that you specify.
 - **Monthly**: Download videos from the URL once in a month, starting with the scheduled date and time. The download starts every month on the day and time that you specify.
6. Click **OK**.

Update an interface in a URL entry

If you have configured multiple links on the network, you might want to use a particular link to download videos, because of better network connectivity. To configure multiple links, you use the available bridge ports, such as apA and apB bridged ports. You can use these ports to download videos for a URL entry.

To update an interface for a URL entry:

1. Navigate to **Configuration > Video Caching > Prepopulation**.
2. Select the entry that you want to update. and click **Modify**.
3. From the **Interface** list, select the interface that you want to use for the URL entry. The list displays the interfaces that are available and configured on the appliance.
4. Click **OK**.

Refresh the status of a URL entry

Over time, the status of the cached videos might change. Checking the status of the entry periodically makes sure that users do not get unexpected results when accessing videos.

To check the latest status of the videos cached from a URL:

1. Navigate to **Configuration > Video Caching > Prepopulation**.
2. Select the entry for which you want to refresh the status of cached videos.
3. Click **Status Check**.

Delete a URL entry

If you do not need a URL entry, you can delete it from the list. To delete a URL entry, select the entry and click **Delete**.

Note

When you delete a video prepopulation task from the list, it also removes the related video objects from the cache.

Verify video caching


March 12, 2021

Graphs and data on the Monitoring page, Dashboard page, and Usage page help you evaluate the benefits provided by your video caching configuration. The data-reduction ratio resulting from video caching (similar to the overall compression ratio) is displayed on the Dashboard, on the video caching monitoring page, and on the Usage graph page. Also, hovering over the Data Reduction ratio on the Dashboard page displays the caching benefit percentage along with compression benefit percentage on the supported platforms.

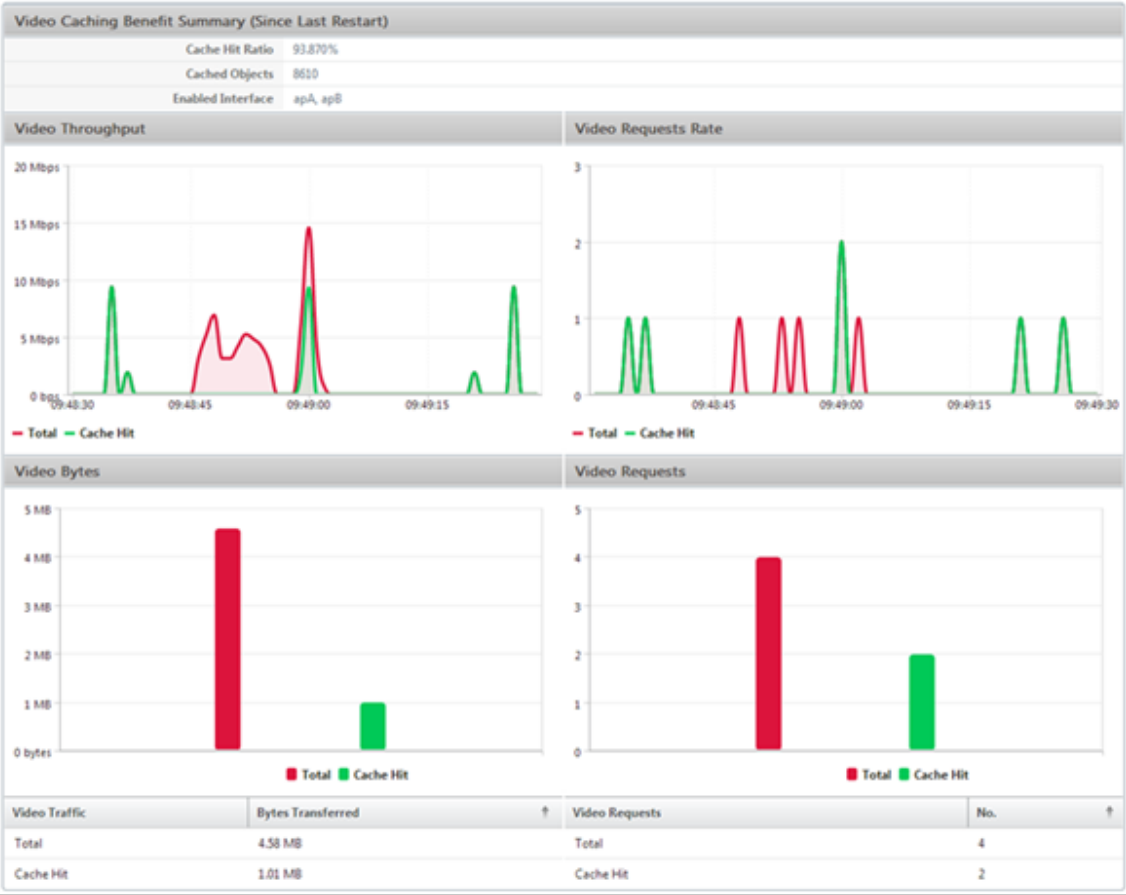
The purpose of caching is not just to save bandwidth, but also to increase performance, decrease load on the video servers, and lessen the impact of network congestion.

The estimated WAN bandwidth savings resulting from video caching are displayed as follows:


- On the Dashboard page, you can view the caching benefit, as a percentage, by hovering the cursor over the Data Reduction field on the Dashboard. You can also view the bytes served from the cache (Cached Data) under Aggregated Link Throughput.



- On the **Monitoring > Video Caching** page, you can view the number of objects cached, and the Cache Hit Ratio (as a percentage). The bar and the time graphs display the number of requests and bytes served from cache over 1 minute, 1 hour, 1 day, 1 week, and 1 month. This data is also displayed in a tabular format below the graph.



- On the **Monitoring > Optimization > Usage Graph** page, you can view the cached data in the LAN Monitoring graph.



- On the **Monitoring > Video Caching > HTTP State List** page, you can monitor the improved cache behavior. This page reports the state of HTTP connections with respect to video caching.
- On the **Monitoring > Optimization > Connections** page, you can view the cached connections on the Accelerated Connections tab. Both cache hits and cache misses are displayed here. The cache connections are displayed here even if they are not accelerated. That is, the cached connections are displayed here even if a partner Citrix SD-WAN WANOP appliance is not involved in the connection. **Bandwidth Savings (%)** column shows a bar graph of how much WAN bandwidth was saved by the transaction, whether through caching or compression. While the aim of caching and compression is to increase speed and usability and not reduce bandwidth usage, speed and usability increases are often related to bandwidth reduction. That is, a 90% bandwidth savings implies a 10x increase in speed.

Monitoring > Optimization > Connections > Accelerated Connections

Accelerated Connections

Unaccelerated Connections

Action

Details	Initiator	Responder	Duration	Idle	Bytes Transferred↑	Compression Ratio/Type	Bandwidth Savings (%)
	172.16.0.50 : 56501	192.229.163.33 : 80	0m 45s	0m 21s	504.95 KB	169.8 to 1 (Disk)	<div><div></div></div> 95.8
	172.16.0.193 : 1060	77.234.41.64 : 80	2h 52m 51s	2m 8s	393.43 KB	1.3 to 1 (Disk)	<div><div></div></div> 15.6
	172.16.0.58 : 55987	104.20.12.86 : 80	18m 23s	0m 5s	327.75 KB	N/A (None)	<div><div></div></div> 0
	172.16.0.50 : 56074	192.229.163.33 : 80	1m 10s	0m 22s	289.83 KB	91.2 to 1 (Disk)	<div><div></div></div> 95.2
	172.16.0.50 : 56092	216.58.216.130 : 80	1m 8s	0m 6s	241.33 KB	90.4 to 1 (Disk)	<div><div></div></div> 94.9
	172.16.0.50 : 56558	31.13.76.100 : 80	0m 42s	0m 3s	156.73 KB	2.8 to 1 (Disk)	<div><div></div></div> 60.6
	172.16.0.50 : 56335	216.58.216.130 : 80	1m 2s	0m 2s	96.65 KB	85.8 to 1 (Disk)	<div><div></div></div> 95.4
	172.16.0.50 : 56559	31.13.76.100 : 80	0m 42s	0m 6s	86.77 KB	2.9 to 1 (Disk)	<div><div></div></div> 62.7

Manage video caching sources

March 12, 2021

You can manage your video sources either globally, by configuring global settings, or individually, by changing status of a video source.

Configure global settings

Global settings enable you to configure the feature at the appliance level. Irrespective of the video sources you have added, these settings are applicable to the entire video caching feature on the appliance. You can:

- Configure the maximum size of the cached objects
- Configure a DNS suffix
- Configure Caching Ports
- Update the video caching policy file


You can configure a maximum size for cached objects. An object larger than this limit is not cached. By default, the maximum caching object size is 100 MB.

For the URLs that do not contain complete domain names and require domain name suffixes to be added to the hostname of the video server, appending a default domain name is necessary for eliciting a response from the server. For example, when you access the http://training/CitrixSD-WANWANOP_VideoCaching.mp4 video, the appliance might be expected to translate the URL to http://training.example.com/CitrixSD-WANWANOP_VideoCaching.mp4. In this case, you must specify example.com as the domain name suffix.

The video caching feature requires a port number for the HTTP video server. The default is port 80. If your HTTP video server uses a port other than this well-known HTTP port, you must add the port number to the list of caching ports.

To configure global settings for video caching:

1. Navigate to **Configuration > Video Caching > Set Global Parameters**.



Set Global Parameters

Max Caching Object Size*

200 MB ▼

DNS Suffix

example.com

Caching Ports

80,89,6789

OK Close

2. In the **MaxCaching Object Size** field, set the maximum size for cached objects.
Select a value from the available limits. An object larger than this limit is not cached.
3. In the **DNS Suffix** field, enter a domain name to append to URLs that do not contain complete domain names and require domain name suffixes to be added to the host name of the video server.
4. In the **Caching Ports** field, type the HTTP video server's port to add it to the list of caching ports. Optionally, add multiple port numbers separated by commas.
5. Click **OK**.

The appliance uses 10% of the allocated disk space for management purposes. When the disk usage reaches 90% of the allocated disk space, that is an indication of the disk being full. To cache more video objects, the appliance removes the least used objects from the video cache. You need not clear the cache unless the cache serves stale video objects.

To clear video cache, navigate to **Configuration > Video Caching** and click **Clear Video Cache**.

WAN insight

March 16, 2021


The Citrix SD-WAN WANOP appliances optimize the delivery of a large number of applications through the WAN, by improving the efficiency of data flow across the network between the datacenter and the

branch sites. WAN Insight analytics enable administrators to easily monitor the accelerated and un-accelerated WAN traffic that flows between the datacenter and branch WAN optimization appliances. WAN Insight provides visibility into clients, applications and branches on the network, to help troubleshoot network issues effectively. Live and historical reports enable you to proactively address issues, if any.

Enabling analytics on the datacenter WAN optimization appliance enables the Citrix Application Delivery Management (ADM) to collect data and provide reports and statistics for the datacenter and the branch WAN optimization appliances.

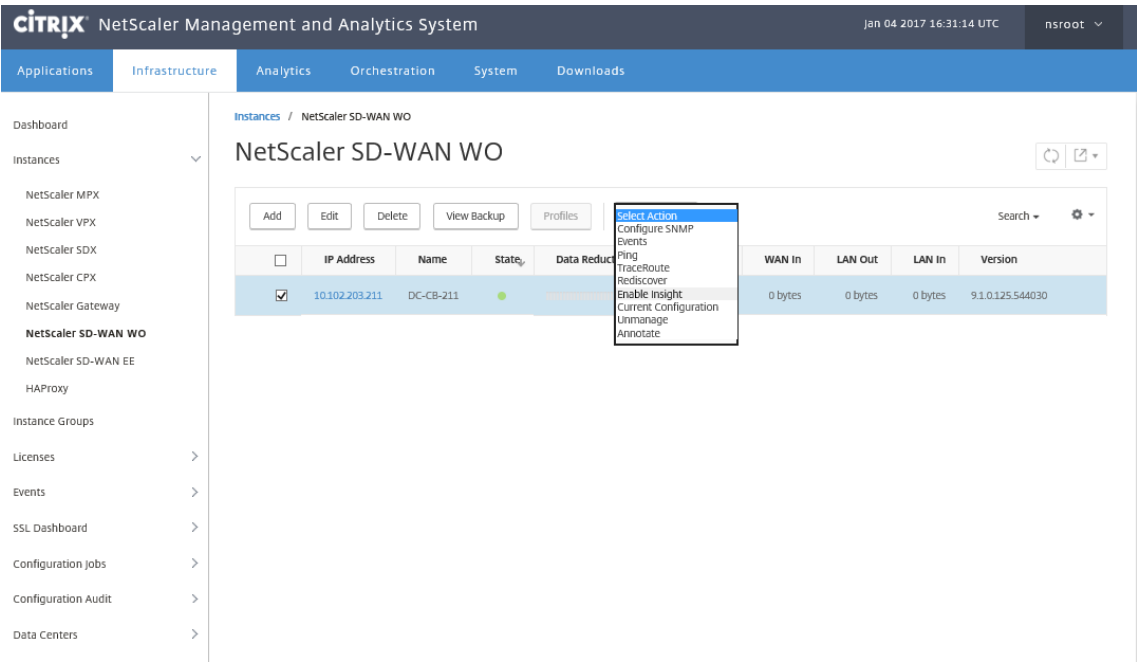
Note

For information on adding an instance, see [Add instances to Citrix ADM](#).

**To enable analytics on the WAN optimization appliance:**

1. In a web browser, type the IP address of the Citrix ADM (for example, <http://192.168.100.1>).
2. In the **User Name** and **Password** fields, enter the administrator credentials.

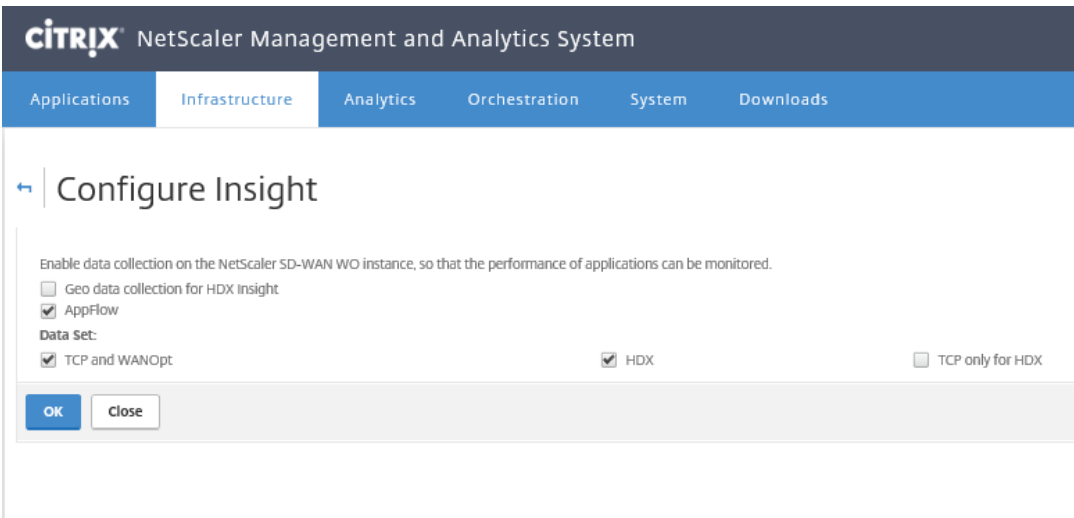
3. Navigate to **Infrastructure > Instances > Citrix SD-WAN WO**, and select the datacenter WAN optimization appliance.



4. From the **Action** drop-down, select **Enable Insight**.

5. Select the following parameters as required:

- **Geo data collection for HDX Insight:** Shares client IP address with the Google Geo API.
- **AppFlow:** Starts collecting data from WAN optimization instances.
- **TCP and WANOpt:** Provides TCP and WANOpt Insight reports.
- **HDX:** Provides HDX Insight reports.
- **TCP only for HDX:** Provides TCP only for HDX Insight reports.



6. Click **OK**.

To view WAN Insight reports:


1. In a web browser, type the IP address of the Citrix ADM (for example, <http://192.168.100.1>).
2. In the **User Name** and **Password** fields, enter the administrator credentials.
3. Navigate to **Analytics > WAN Insight**.

Note

The WAN Insight option is visible only after you add an SD-WAN WO instance to Citrix ADM.

You can view the following reports:

- **Applications** - Displays the usage and performance statistics of all the applications for the selected duration.
- **Branches** - Displays the usage and performance statistics of all the WAN optimization branch appliances.
- **Clients** - Displays the usage and performance statistics of all the clients accessing the WAN optimization appliances, in each branch.



The following metrics are displayed:

Metric	Description
Active Accelerated Connections	Number of active WAN connections that are accelerated.
Active Unaccelerated Connections	Number of active WAN connections that are not accelerated.
WAN Latency	Delay, in milliseconds, that the user experiences while interacting with an

application. |

| Compression Ratio | Ratio of data compression between the branch office and datacenter appliances for the selected duration. |

| Packets Sent | Number of packets that the WAN optimization appliance has sent over the network for the selected duration. |

| Packets Received | Number of packets that the WAN optimization appliance has received from the network for the selected duration. |

| Bytes Sent over WAN | Number of bytes that the Citrix WAN optimization appliance has sent over the WAN for the selected duration. |

| Bytes Received over WAN | Number of bytes that the WAN optimization appliance received from the WAN for the selected duration. |

| LAN RTO | Number of times the WAN optimization appliance has timed out retransmission to the LAN for the selected duration. |

| WAN RTO | Number of times the WAN optimization appliance has timed out retransmission to the WAN for the selected duration. |

| Retransmit Packets (LAN) | Number of packets the WAN optimization appliance has retransmitted to the LAN network for the selected duration. |

| Retransmit Packets (WAN) | Number of packets the WAN optimization appliance has retransmitted to the WAN network for the selected duration. |


Asymmetric routing

March 12, 2021

In Citrix SD-WAN WANOP network, asymmetric routing occurs when packets flowing from client to server or server to client for the same TCP connection do not pass through one or both the client-side and server-side WANOP appliances. The following cases of asymmetry are observed.


Complete asymmetry:

Complete asymmetry occurs when packets flow from a client to the server through both the client-side and server-side Citrix SD-WAN WANOP appliances. However, on the return path from server to client the packets take a different route bypassing both the Citrix SD-WAN WANOP Appliances.




Server-side asymmetry:

Server-side asymmetry occurs when packets flow from a client to the server through both the client-side and server-side Citrix SD-WAN WANOP appliances. However, on the return path the packets bypass the server-side Citrix SD-WAN WANOP appliance but traverses the client-side Citrix SD-WAN WANOP appliance.



Client-side asymmetry:

Client-side asymmetry occurs when packets flow from a client to the server through both the client-side and server-side Citrix SD-WAN WANOP appliances. However, on the return path the packets traverse the server-side Citrix SD-WAN WANOP appliance but bypass client-side Citrix SD-WAN WANOP appliance.



Handle asymmetry in Citrix SD-WAN WANOP network

In Citrix SD-WAN WANOP network, when complete asymmetry occurs the TCP connection is reset. To avoid TCP connection break and to continue sending unaccelerated traffic, an asymmetric connection list is introduced in SD-WAN WANOP 10.1. This feature is disabled by default; you can enable this feature on both the client-side and server-side SD-WAN WANOP appliances.

On detecting an asymmetric connection for the first time, the TCP connection between client and server is reset and an entry of the tuple is made in the asymmetric connection list. The tuple consists of the client IP address and server IP address. Subsequent connections from the tuple pass through unaccelerated. The connection tuple remains in the asymmetric connection list for a default time-out period of four hours or until symmetry is detected. The unaccelerated pass-through is effective until the time-out occurs or until the appliance dynamically detects that the asymmetry is no longer present.

When client-side asymmetry or server-side asymmetry is detected, the TCP connection is retained and the packets pass through the Citrix SD-WAN WANOP appliance unaccelerated, by default.

To enable asymmetric connection list on Citrix SD-WAN WANOP appliances:

1. Access the WANOP CLI command prompt (WANOP Accelerator/Broker IP).
2. Log in with the following credentials:

```
1  **Login as:** *cli*****
2
3  **Login**: **** *admin*****
4
5  **Password**: **** *nsroot*****
```

Note

The default password for admin is *nsroot*. If you have changed the password, use the right one.

3. Type the following command and hit enter.

```
1  *Set parameter AssymmetricConnectionList.Enable on*
```

Note

You can configure the time-out period as per your network requirement, using the *AssymmetricConnectionList.AutoFlushDuration* command.

There are multiple parameters available with asymmetry list that can be fine-tuned, on-demand, based on your network environment. For more information, contact Citrix Customer Support.

Citrix SD-WAN WANOP client plug-in


March 12, 2021

The Citrix 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 appliances. It is ideal for remote workers. It connects to a Citrix WANOP appliance at the other end of the link (typically in the data center or cloud).

The principles of WANOP Client Plug-in operation are generally the same as those of a WANOP 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 appliances at the other ends of the links. If you customize the MSI file with the DNS or IP addresses of the WANOP 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 all currently supported releases of Citrix Workspace App for Windows and Citrix Receiver for Windows.

Hardware and software requirements

March 12, 2021

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:

- WANOP VPX
- WANOP 1000
- WANOP 2000
- WANOP 3000
- WANOP 4000
- WANOP 4100
- WANOP 5000
- WANOP 5100

How WANOP plug-in works

March 12, 2021

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

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 Citrix 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, Citrix SD-WAN WANOP A2 does not need to be licensed for plug-in acceleration, because Citrix SD-WAN WANOP 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 Citrix SD-WAN WANOP B. For this to work, Citrix SD-WAN WANOP 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



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 WANOP Plug-in looks up the destination address and sees that it matches a subnet accelerated by the appliance. It attaches WANOP options to the TCP header of the SYN packet. No addresses are changed.

Src: 10.0.0.50, Dst: 10.200.0.10

3. The appliance notes the SYN options and recognizes that this is an accelerable connection. It strips the options from the packet and allows it to pass through to the server. No addresses are changed.

Src: 10.0.0.50, 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: 10.0.0.50

5. The appliance tags the SYN-ACK packet with a TCP header option that shows that acceleration will take place.

Src: 10.200.0.10, Dst: 10.0.0.50

6. The WANOP Plug-in receives the SYN-ACK packet. The options in the packet headers indicate that the connection is accelerated. The Plug-in strips the options and passes the SYN-ACK packet to the application. The connection is now fully open and accelerated.

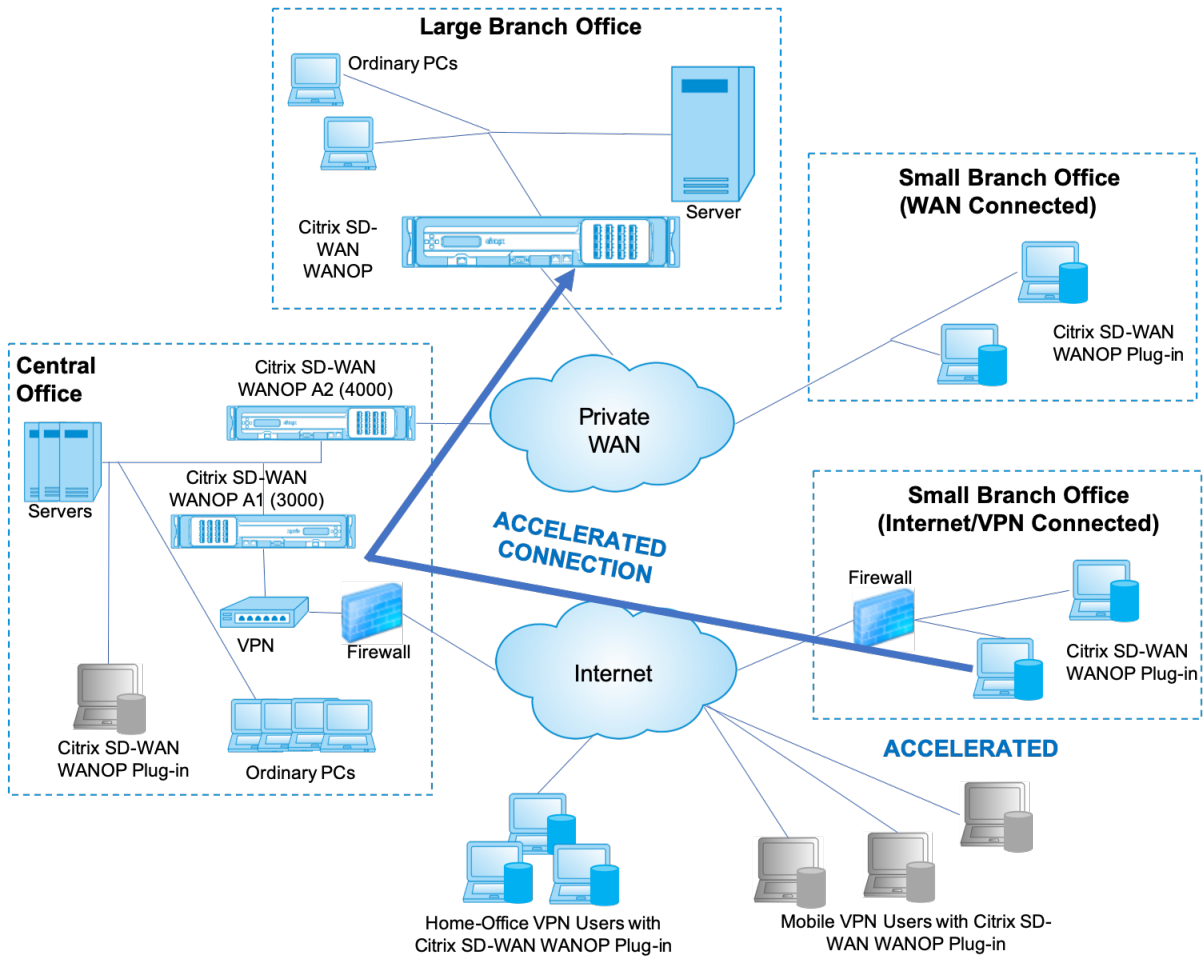
Redirector mode

Redirector mode works differently from transparent mode in the following ways:

- The WANOP Client 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. Citrix SD-WAN WANOP Plug-in looks up the destination 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 destination address from the TCP options field), and giving itself as the source.

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 layered 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 packet. CIFS acceleration will perform speculative read-ahead and write-behind operations. Also, if packets are dropped between appliance and the Reporter plug-in, the retransmission is handled by the appliance, not the server, using advanced recovery algorithms.

How the plug-in selects an appliance

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.

Deploy appliances for use with plug-ins

March 12, 2021

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.

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.

Customize plug-in's MSI file

March 12, 2021

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 its 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.

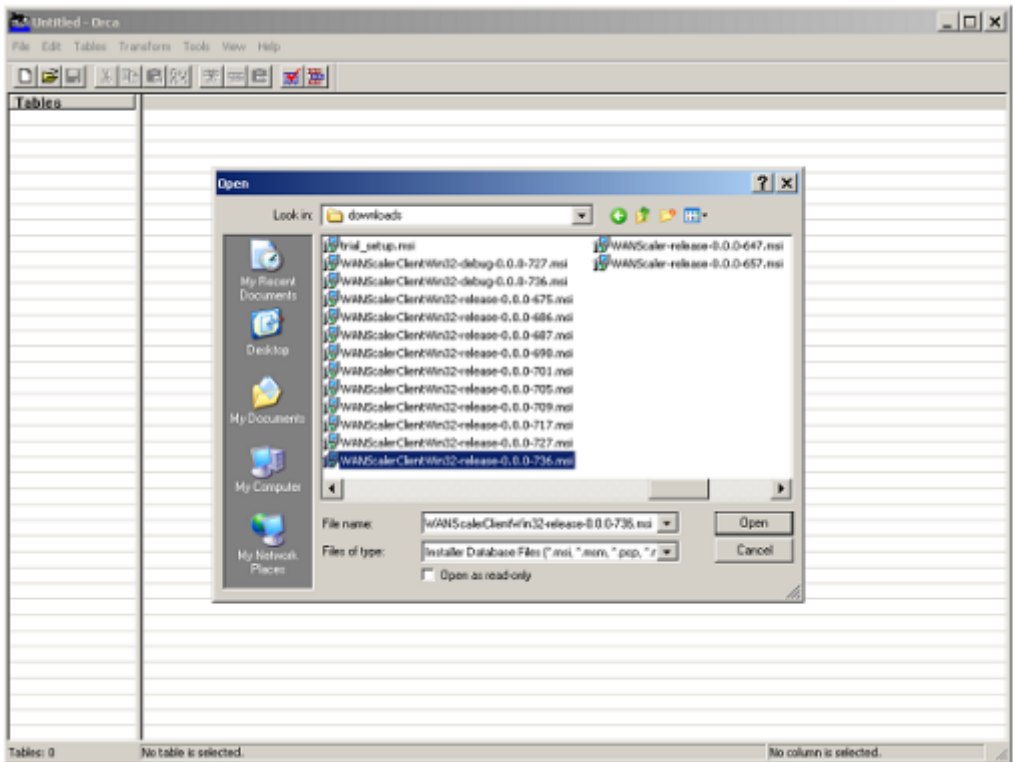
Install the Orca MSI editor:

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,

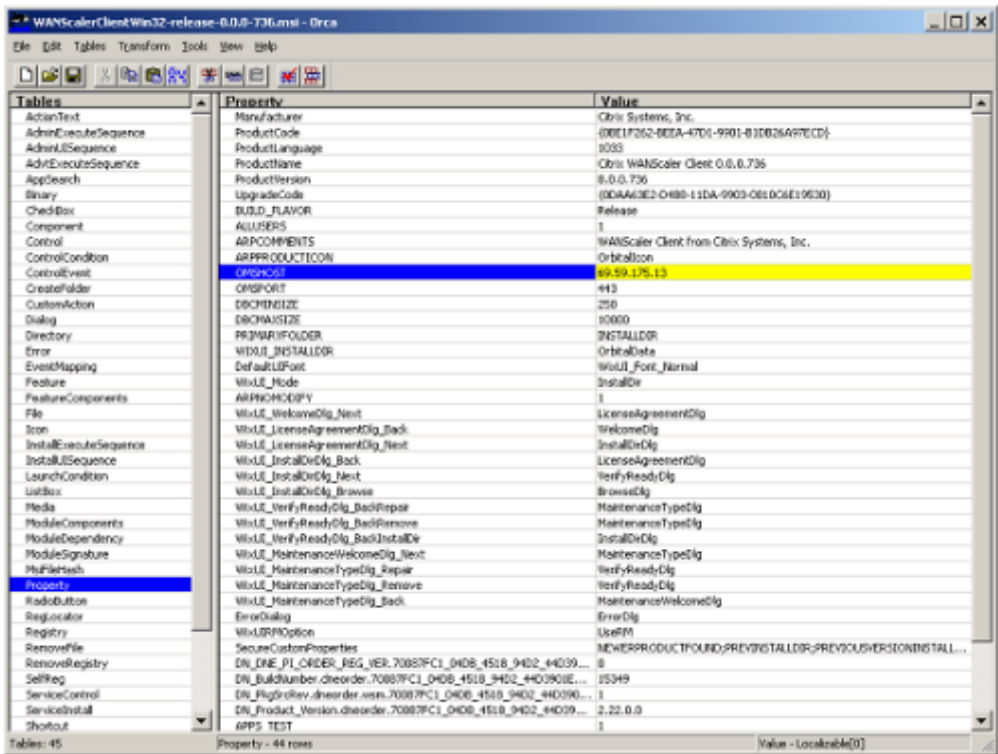
type the new value, and press **Enter**.

For more information, see the [table](#)

- a) 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**.

For more information, see the [table](#)

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.

Deploy plug-ins on Windows

March 12, 2021

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.”

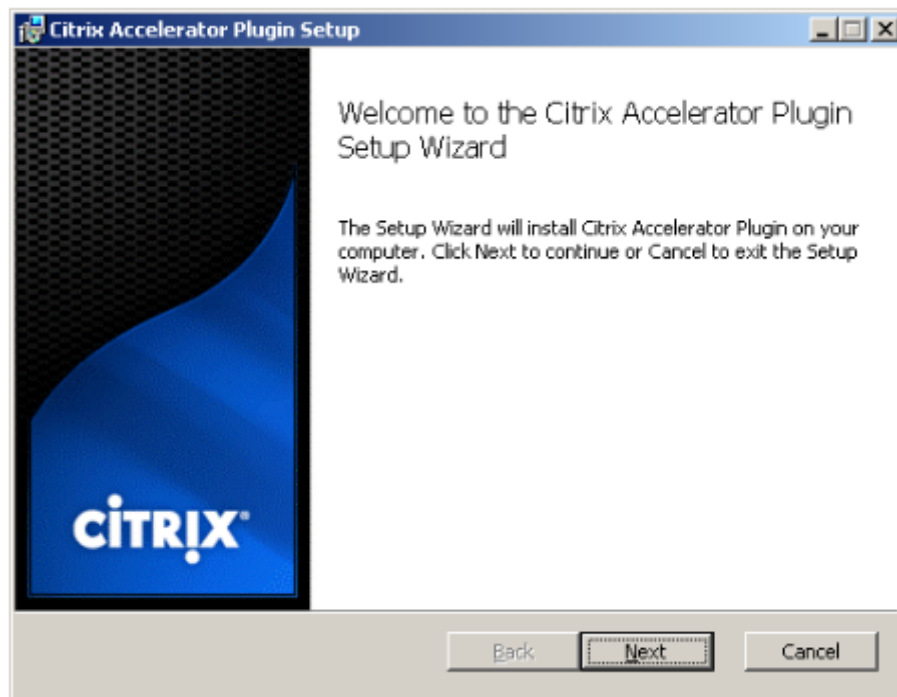
The only user configuration needed by the plug-in is the list of appliance addresses. This list can consist of a comma-separated list of IP or DNS address. The two forms can be mixed. You can customize the distribution file so that the list points to your appliance by default. Once installed, operation 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

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:

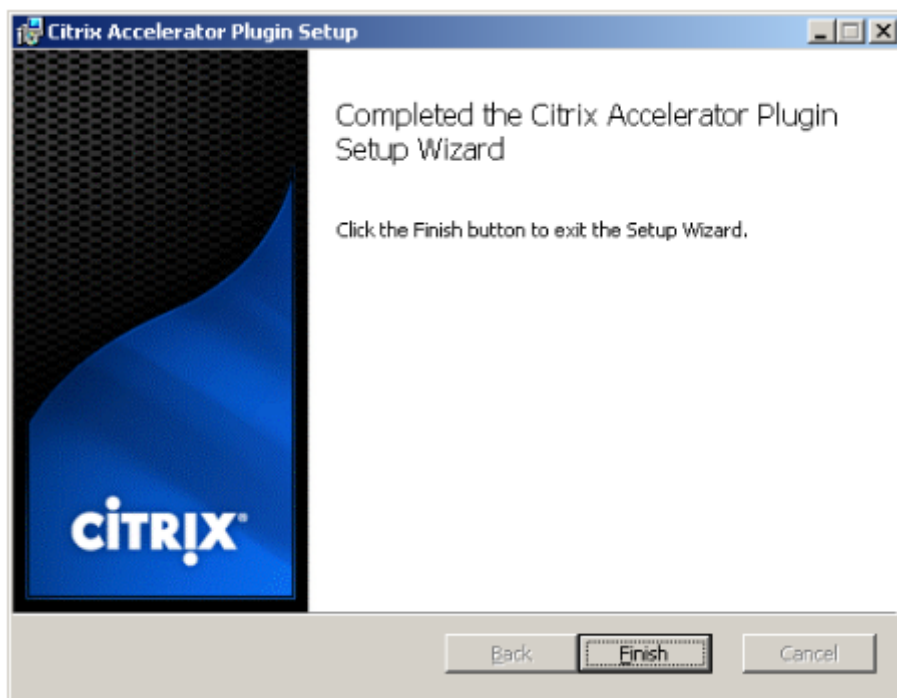


The steps below are for an interactive installation. A silent installation can be performed with the command:

```
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 appli-

ance. 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.

Troubleshoot plug-ins

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.

Citrix SD-WAN WANOP plug-in GUI

March 12, 2021

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

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

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 Citrix Virtual Apps and Desktops 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.

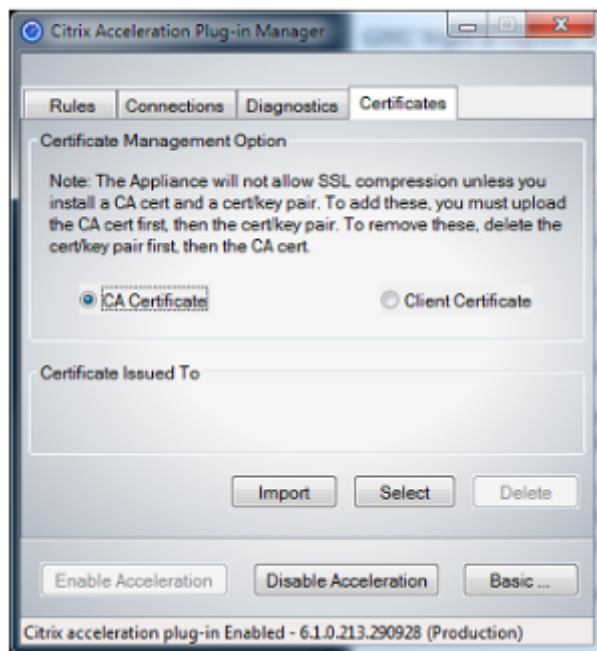
When you stop tracing, a pop-up window shows the trace files. Send them to your Citrix representative by the means he or she recommends.

- **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.

Update Citrix SD-WAN WANOP plug-in

March 12, 2021

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.

Uninstall the WANOP client 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**.

Restart the system to finish uninstalling the client.

Citrix Virtual Apps and Desktops acceleration

March 12, 2021

Note

In this discussion, *Virtual Apps* refers to the ICA and CGP protocol streams. Therefore, what is said about Virtual Apps applies also to Virtual Desktops.

Virtual Apps/Virtual Desktops (ICA/CGP) acceleration has three components:

- **Compression**—The appliance cooperates with Virtual Apps clients and servers to compress Virtual Apps data streams for interactive data (keyboard/mouse/display/audio) and batch data (printing and file transfers). This interaction takes place transparently and requires no configuration of the appliance. A small amount of configuration, described below, is required on older Virtual Apps servers (release 4.x).
- **Multistream ICA**—In addition to compression, Citrix SD-WAN WANOP appliances support the new Multistream ICA protocol, in which up to four connections are used for the different ICA priorities, instead of multiplexing all priorities over the same connection. This approach gives interactive tasks greater responsiveness, especially when combined with the appliance's traffic shaping.
- **Traffic shaping**—The Citrix SD-WAN WANOP traffic shaper uses the priority bits in the Virtual Apps data protocols to modulate the connection's priority in real time, matching the bandwidth share of each connection to what the connection is transmitting at the moment.

Note

Multistream ICA is disabled by default. It can be enabled on the Features page. Multistream ICA and AutoQoS requires Session Reliability to be enabled.

To optimize ICA connections for Citrix Virtual Apps and Desktops release 7.0 and later, Citrix SD-WAN WANOP appliance supports Citrix Receiver for Chrome release 1.4 and later, and Citrix Receiver for HTML5 release 1.4 and later.

HDX Transport protocol from UDP/EDT to TCP –In certain network conditions, UDP/EDT cannot be used as the optimized protocol to deliver HDX traffic. You can change the protocol to TCP so that WANOP can provide:

- Compression/DDup benefits
- Visibility (local reports and HDX Insight)

WANOP can block EDT traffic and force the session to TCP. During session initiation Citrix Receiver starts session on both TCP and EDT. If EDT session is not established then TCP session is used. WANOP GUI provides an option to force the session on TCP protocol on the features page.

Configure Virtual Apps acceleration

March 12, 2021

Virtual Apps acceleration applies to both the ICA and CGP protocols within Virtual Apps. The Citrix SD-WAN WANOP appliances, Virtual Apps servers, and Virtual Apps clients provide cooperative acceleration of Virtual Apps connections, providing substantial speedup compared to Virtual Apps alone. This cooperation requires up-to-date versions of all three components.

Virtual Apps compression dynamically switches between memory based compression for interactive channels (such as mouse, keyboard, and screen data) and disk based compression for bulk tasks (such as file transfers and print jobs). Compression ratios increase as compression history fills, increasing the amount of data that can be matched against new data. Virtual Apps compression provides several times as much data reduction as does unassisted Virtual Apps, often exceeding 50:1 on repetitive bulk transfers such as printing or saving successive versions of the same document.

Virtual Apps compression achieves high link utilization without congestion, by preventing users from interfering with each other.

To enable Virtual Apps acceleration

1. Check the ICA service class policy. On the Configuration: Service Classes page, the ICA service class should show disk in the Acceleration column and ICA Priorities in the Traffic Shaping column. If not, edit the service class definition.
2. Update Virtual Apps 4.x servers and clients. (Not necessary on Virtual Apps 5.0 or later). Use Presentation Server 4.5 with Hotfix Rollup Pack PSE450W2K3R03 (Beta) or later. This release

includes the following server and client software, both of which must be installed for Virtual Apps compression:

- a) Server package PSE450R03W2K3WS.msp or later.
 - b) Client version 11.0.0.5357 or later.
3. Update Virtual Desktops servers and clients to release 4.0 or later.
 4. Verify Virtual Apps server registry settings. (Not necessary on Virtual Apps 5.0 or later.) On the Virtual Apps servers, verify the following settings and correct or create them as necessary:

```
pre codeblock HKLM\System\CurrentControlSet\Control\Citrix\
WanScaler\EnableForSecureIca = 1 HKLM\System\CurrentControlSet
\Control\Citrix\WanScaler\EnableWanScalerOptimization = 1 HKLM\
System\CurrentControlSet\Control\Citrix\WanScaler\UchBehavior = 2
<!--NeedCopy-->
```

These are all DWORD values.

5. Open and use Virtual Apps connections, between updated Virtual Apps clients and servers, that pass through the updated Citrix SD-WAN WANOP. By default, these sessions use CGP. For ICA, on the client, under Citrix Program Neighborhood, clear the Custom ICA Connections check box. Then, right-click a connection icon, navigate to **Properties > Options**, and click **Enable Session Reliability** check box. Multi-Stream ICA and AutoQoS requires Session Reliability to be enabled.
6. Verify acceleration.

After you start Virtual Apps sessions over the accelerated link, accelerated ICA connections should appear on the appliance's Monitoring: Connections page. A compression ratio of greater than 1:1 indicates that compression is taking place.

Optimize Citrix Receiver for HTML5

March 12, 2021

Application that must serve dynamic content work on HTML5 WebSockets. Citrix Receiver for Chrome and Citrix Receiver for HTML5 are such applications that support HTML5 WebSockets. These applications have simplified access to Virtual Desktops as these can be integrated with most recent Web browsers that support HTML5 WebSockets.

Note

You do not need to make any changes to the appliance configuration to use this feature.

How a Citrix SD-WAN WANOP appliance optimizes Citrix Receiver for HTML5

In a typical branch office and datacenter setup, shared resources like Virtual Desktop Agent (VDA) are installed on a Citrix Hypervisor server in the datacenter. Clients from the branch offices access these shared resources over the network by using Citrix Receiver.

In a typical branch office and datacenter setup, shared resources like Virtual Desktop Agent (VDA) are installed on a Citrix Hypervisor server in the datacenter. Clients from the branch offices access these shared resources over the network by using Citrix Receiver.

Being HTML compliant, VDA uses a WebSocket listener that runs on port 8008. When accessing an application, the client initiates a TCP connection at port 8008, and uses it to send an HTTP request to the server to upgrade the connection and use the WebSocket protocol. After the client negotiates the WebSocket connection with VDA, Independent Computing Architecture (ICA) negotiations begin and the client and the server use ICA over HTML5 to exchange data. For more information about the sequence of messages exchanged between the client and server, see [Messages Exchanged Between the Client and the Server](#).

After connections are established between the clients and the server, the Citrix SD-WAN WANOP appliance starts optimizing the connections by speeding up the traffic over the network, and accelerating Web page and other applications using Citrix Receiver for HTML5. The functionality of optimizing the Citrix Receiver for HTML5 connections is similar to HTTP Acceleration.

Note

- For more information about HTML5, see [How HTML5 Works](#).
- For more information about Citrix Receiver for HTML5, see <http://support.citrix.com/proddocs/topic/receiver-for-html5.html>.
- For more information about the system requirements of Receiver for HTML5, see <http://support.citrix.com/proddocs/topic/receiver-html5-14/receiver-html5-sys-req.html>.

Configure a Citrix SD-WAN WANOP appliance to optimize Citrix Receiver for HTML5

Optimization of Citrix Receiver for HTML5 connections is a zero configuration feature. You do not have to make any configuration changes to the appliance. Upgrading the Citrix SD-WAN WANOP software to release CB 7.3.1 or later creates the alt-http application classifier on the appliance and maps this application classifier to port 8008, which is the default for Virtual Desktops. As soon as you upgrade the software the appliance, it is ready to optimize native Chrome connections that use Citrix Receiver for HTML5.

If you are using SSL encryption for connections over Citrix Receiver for HTML5, connections use ICA over SSL. To enable ICA over SSL acceleration with Citrix Receiver for HTML5, you need to configure standard SSL acceleration, which includes the appropriate destination IP address in the service class

and SSL profile mapping. If you are planning to deploy the appliance in ICA proxy mode, you must map the StoreFront VIP address to StoreFront certificates. Similarly, if you plan to deploy the appliance in any end-to-end SSL encryption deployment mode, you must map the VDA IP address to VDA certificates.

Warning

Make sure that you do not change the port number of the alt-http application to any other port number. If you delete this application classifier or need to make any changes to it, you must add the port 8008 to the HTTP application classifier.

Verify Citrix Receiver for HTML5 connections

To verify that the appliance is optimizing Citrix Receiver for HTML5 connections, you can check to see if connections are listed in Citrix (ICA/CGP) and ICA Advanced monitoring pages. Existence of HTML5 connections in the monitoring pages is an indication that the appliance is optimizing the Citrix Receiver for HTML5 connections.

To verify Citrix Receiver for HTML5 connection on a Citrix SD-WAN WANOP appliance:

1. Navigate to the **Monitoring > Optimization > Citrix (ICA/CGP)** page.
2. On the **ICA Connections** tab, verify that the HTML5 connections are listed. An HTML5 connection is shown with HTML as a prefix in the Client Computer Name column, as shown in the following screen shot:

Published Application or Desktop	Client Computer Name	Client IP Address	Server IP Address	Protocol	Duration	Transferred Bytes #	Acceleration Status	Encryption
Word 2013_1	HTML-2922-1550	14.141.5.5	10.102.255.210	ICA over SSL	11h 45m 17s	1.19 MB	●	Basic (XOR)
SC Ar26 Win 2008 R2 RDS	HTML-1184-5111	14.141.5.5	10.102.255.210	ICA over SSL	4m 7s	196.88 KB	●	Basic (XOR)

3. Navigate to the **Monitoring > Optimization > ICA Advanced** page.
4. In the **Conn Info** tab, scroll down to the ICA Client and Server Information section. Entries for HTML5 connections have Citrix HTML5 client in the Product ID column, as shown in the following screen shot:

Dashboard

Monitoring

Configuration

Notifications (0)

Optimization

Citrix (ICA/CGP)

Connections

Compression

Filesystem (CIFS/SMB)

LAN vs WAN

Links Usage

Outlook (MAP)

Service Classes

Top Applications

Traffic Shaping

Usage Graph

ICA Advanced

Appliance Performance

Partners & Plugins

Monitoring > Optimization > ICA Advanced

Show Acceleration Status and Diagnostics: [ALL Connections Toggle](#)

Acceleration Status and Diagnostics

Conn ID	Connection Status	Session Status	Diagnostics	Remedy
116	●	●	OK	None
113	●	●	OK	None

Connection Attributes

Conn ID	Protocol	Stream	ICA Priority	Encryption	CB Pair Compression	CB Conn Compression Algorithm	CB Side	Client CB Compression	Server CB Compression	Acceleration Partner Type
116	ICA over SSL	Single	mixed	Basic (XOR)	on	DBC	Server	Disk	Disk	Appliance
113	ICA over SSL	Single	mixed	Basic (XOR)	on	DBC	Server	Disk	Disk	Appliance

ICA Client and Server Information

Client Info							Server Info				
Conn ID	Stream	Initial Program	Name	Version	Product ID	Directory	Launcher	Farm Name	Name	User Name	Domain
116	Single	SC Ar26 Win 2008 R2 RDS	HTML-1184-5111	1.4.0.5018	Citrix HTML5 client	none	ReceiverWeb		SC-RDS-AR26-02	sanjays	citrite
113	Single	Word 2013_1	HTML-2922-1550	1.5	Citrix HTML5 client	none	ReceiverWeb		CH-RDS-AR26-05	thavamanir	citrite

ICA Session Information


Session ID	Attributes	Session CIDR	Launched App	Multi Stream
------------	------------	--------------	--------------	--------------

Deployment modes

March 12, 2021

In a typical Citrix SD-WAN WANOP deployment, the Citrix SD-WAN WANOP appliances are paired across branch offices and datacenter. You install shared resources, such as VDA, in the datacenter. Clients from various branch offices access datacenter resources by using Citrix Receiver, as shown in the following figure.


A typical Citrix SD-WAN WANOP deployment topology



Clients install a Citrix Receiver software product, such as Citrix Receiver for HTML5, on their local computers and use it to access resources in the datacenter. Connections through the pair of Citrix SD-WAN WANOP appliances are optimized.

Understand messages exchanged between the client and the server

As with any type of network connection, a client using Citrix Receiver for HTML5 exchanges various messages with the server. The following figure shows a typical flow of messages between the client and server when a connection is established between them.



As shown in the above figure, the following sequence of messages is exchanged between the client and server when a client from a branch office wants to access datacenter server resources:

1. Client uses Citrix Receiver for HTML5 to send a TCP connection request to VDA on port 8008.
2. After establishing the TCP connection, the client sends a WebSocket upgrade request to VDA.
3. VDA responds to the upgrade request and switches to the WebSocket protocol.
4. Client and VDA negotiate WebSocket authorization.
5. Client sends a WebSocket connection request to VDA.
6. VDA responds to the WebSocket connection request.
7. VDA initiates ICA negotiation with the client.

8. After ICA negotiation, VDA starts transmitting ICA data.
9. VDA sends packet termination message.
10. Client responds with the packet termination message.

Note

The above example lists the sample messages exchanged for ICA over WebSocket. If you are using ICA over Common Gateway Protocol (CGP), the client and server negotiate CGP instead of WebSocket. However, for ICA over TCP, the client and server negotiate ICA.

Depending on the components you have deployed on the network, the connection is terminated at different points. The preceding figure represents a topology that does not have any additional components deployed on the network. As a result, the client communicates directly with VDA at port 8008. However, if you have installed a gateway, such as Citrix Gateway, at the datacenter, the connection is established with the gateway and it proxies VDA. Until the gateway negotiates the WebSocket authorization, there is no communication with VDA. After the gateway has negotiated WebSocket authorization, it opens a connection with VDA. Thereafter, the gateway acts as a middleman and passes messages from the client to VDA and vice versa.

Similarly, if a VPN tunnel is created between a Citrix gateway plugin installed on the client and Citrix Gateway installed at the datacenter, the gateway transparently forwards all client messages, immediately upon establishing a TCP connection, to VDA, and vice versa.

Note

To optimize a connection that requires end-to-end SSL encryption, a TCP connection is established at port 443 on VDA.

Supported deployment modes

When configuring a Citrix SD-WAN WANOP appliance for optimizing Citrix Receiver for HTML5, you can consider any of the following deployment modes, depending on your network requirements. To optimize Citrix Receiver for HTML5 connections, Citrix SD-WAN WANOP appliances support the following deployment modes:

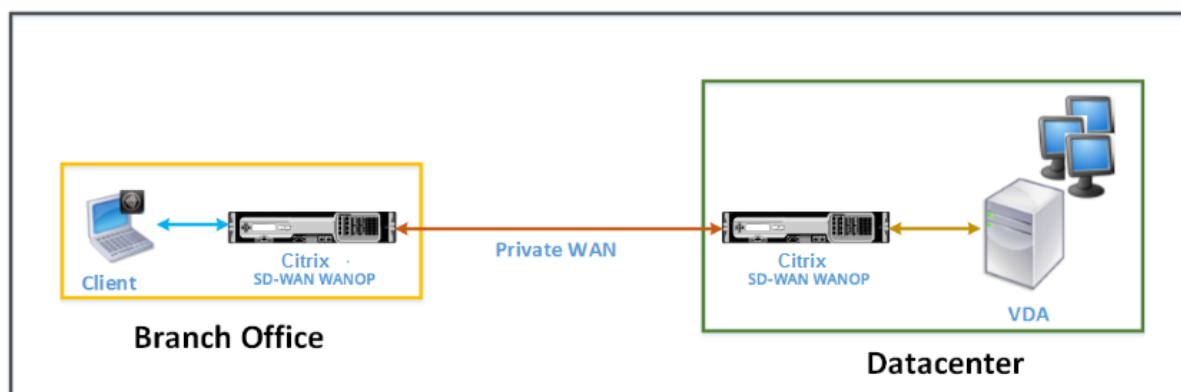
- Direct Access
- Direct Access with End-to-End SSL Encryption
- ICA Proxy Mode
- ICA Proxy Mode with End-to-End SSL Encryption
- Full Virtual Private Network (VPN) Mode

- Full Virtual Private Network (VPN) Mode with End-to-End SSL Encryption

Direct access:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in the direct access mode.

Citrix SD-WAN WANOP appliances deployed in direct access mode



In the direct access mode, a pair of Citrix SD-WAN WANOP appliances is installed across a branch office and the datacenter in inline mode. A client accesses VDA resources through Citrix Receiver for HTML5 over the private WAN. Connections from the client to the VDA resources is secured by using encryption at the ICA level. Messages exchanged between the client and VDA are explained in Understanding Messages Exchanged Between the Client and the Server.

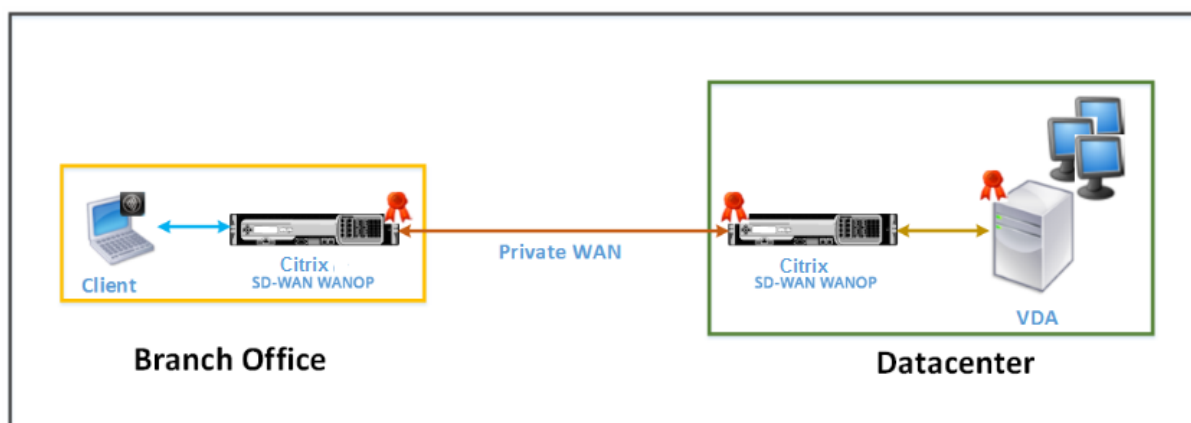
The Citrix SD-WAN WANOP appliances installed between the client and VDA datacenter optimize the Citrix Receiver for HTML5 connections established between them.

A direct access deployment is suitable for a corporate intranet on which clients connect without using Citrix Gateway or any other firewall. You deploy a set up with direct access when Citrix SD-WAN WANOP appliances are deployed in the inline mode and a client from a private WAN connects to the VDA resources.

Direct access with end-to-end SSL encryption:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in the direct access mode secured with end-to-end SSL encryption.

Citrix SD-WAN WANOP appliances deployed in direct access mode secured with end-to-end SSL encryption



The direct access with end-to-end SSL encryption mode is similar to the Direct Access mode, with the difference that the connection between the client and VDA resources is secured by SSL encryption and uses port 443 instead of port 8008 for the connection.

In this deployment, communication between a pair of Citrix SD-WAN WANOP appliances is secured by making the two appliance secured partners. This deployment is suitable for a corporate network where connections between the client and VDA resources are secured by SSL encryption.


Note

You must configure appropriate certificates on the appliances to create secure partners. For more information about secure partnering, see [Secure Peering](#).

ICA proxy mode:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in ICA proxy mode.

Citrix SD-WAN WANOP appliances deployed in ICA proxy mode



In the ICA proxy mode, a pair of Citrix SD-WAN WANOP appliances is installed across the branch office and a datacenter in inline mode. In addition, you install Citrix Gateway, which proxies VDA, at the datacenter. A client accesses VDA resources through Citrix Receiver for HTML5 over the public WAN. Because the gateway proxies the VDA, two connections are established: an SSL connection between

the client and Citrix Gateway and an ICA secured connection between Citrix Gateway and VDA. The Citrix Gateway establishes a connection with VDA resources on behalf of the client. Connections from the gateway to the VDA resources is secured by encryption at the ICA level.

Messages exchanged between the client and VDA are explained in Understanding Messages Exchanged Between the Client and the Server. However, in this case the connection is terminated at Citrix Gateway. The gateway proxies VDA and opens a connection to VDA only after the gateway has negotiated WebSocket authorization. The gateway then transparently passes messages from client to VDA and vice versa.

If you expect users to access VDA resources from a public WAN, you can consider deploying the ICA Proxy mode set up.


Note

You must configure appropriate certificates on the appliances to create secure partners. For more information about secure partnering, see [Secure Peering](#).

ICA proxy mode with end-to-end SSL encryption:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in ICA proxy mode secured with end-to-end SSL encryption.

Citrix SD-WAN WANOP appliances deployed in ICA proxy mode secured with end-to-end SSL encryption



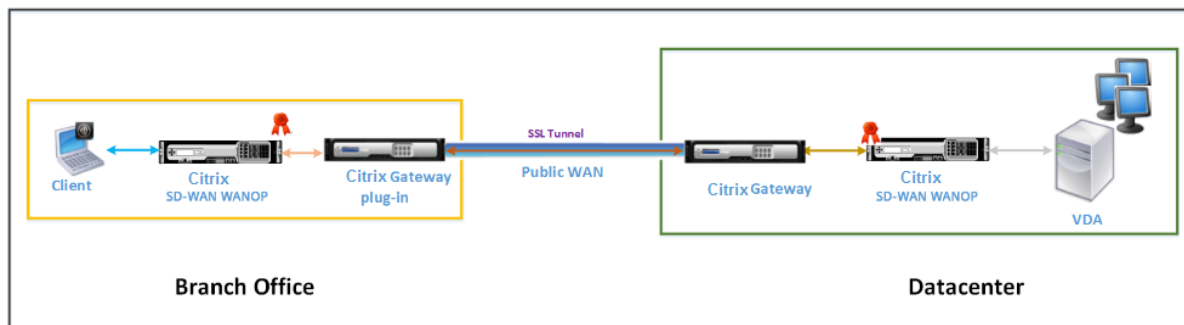
ICA Proxy mode with end-to-end SSL encryption mode is similar to ordinary ICA Proxy mode, with the difference that the connection between the Citrix Gateway and VDA is secured by SSL encryption instead of using an ICA secured connection. In this scenario, you must install appropriate certificates on the Citrix SD-WAN WANOP appliance and VDA. The connection between the Citrix Gateway and VDA uses port 443 instead of port 8008, as in case of ordinary ICA Proxy mode.

This deployment is suitable for a network where you must secure end-to-end communication between clients and VDA, including the connection between Citrix Gateway and VDA.

Full virtual private network (VPN) mode:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in the full Virtual Private Network (VPN) mode.

Citrix SD-WAN WANOP appliances deployed in VPN mode



In full VPN mode, a pair of Citrix SD-WAN WANOP appliances is installed across a branch office and the datacenter in inline mode. In addition to Citrix receiver for HTML5, you install the Citrix Gateway plugin on the client and Citrix Gateway interfacing external network at the datacenter. The Citrix Gateway plugin on the client and Citrix Gateway on the datacenter create an SSL tunnel or VPN over the network when they establish a connection. As a result, the client has a direct secure access to the VDA resources, with transparent connection through the Citrix SD-WAN WANOP appliance. When the client connection is terminated at Citrix Gateway, the gateway opens a transparent connection to port 8008 on VDA.

Messages exchanged between the client and VDA are explained in the Understanding Messages Exchanged Between the Client and the Server section. However, in this case the connection is terminated at Citrix Gateway. The gateway proxies VDA and opens a transparent connection to VDA at port 8008, and transparently passes all messages from client to VDA and vice versa.


The Citrix SD-WAN WANOP plug-in enables the client to access resources regardless of the location of the client. When you expect clients to need access to the VDA resources from locations other than their desktops, you can deploy the setup in full virtual Private Network (VPN) mode.

This deployment is suitable for organizations expecting their employees to access resources when they are traveling.

Full virtual private network (VPN) mode with end-to-end SSL encryption:

The following figure shows the deployment topology of Citrix Receiver for HTML5 installed on the client in the full VPN mode secured with end-to-end SSL encryption.

Citrix SD-WAN WANOP appliances deployed in VPN mode secured with end-to-end SSL encryption



Full Virtual Private Network (VPN) mode with end to end SSL encryption deployment is similar to ordinary full VPN mode, with the difference that the communication between Citrix Gateway and VDA is secured by SSL encryption and uses port 443 instead of port 8008.

This deployment is suitable for organizations that need end-to-end SSL encryption for resources accessed by the employees who are traveling.

Adaptive transport interoperability

March 12, 2021

Adaptive transport is a data transport mechanism for Citrix Virtual Apps and Desktops. It is faster, can scale, improves application interactivity, and is more interactive on challenging long-haul WAN and internet connections. Adaptive transport maintains high server scalability and efficient use of bandwidth. By using adaptive transport, ICA virtual channels automatically respond to changing network conditions. They intelligently switch the underlying protocol between the Citrix protocol called Enlightened Data Transport (EDT) and TCP to deliver the best performance. By default, adaptive transport is enabled, and EDT is used when possible, with fallback to TCP.

Citrix SD-WAN WANOP offers cross-session tokenized compression (data deduplication), including URL-based video caching. It provides significant bandwidth reduction if two or more people at the office location watch the same client-fetched video, or transfer or print significant portions of the same file or document. Furthermore, by running the processes for ICA data reduction and print job compression on the branch office appliance, WANOP offers VDA server CPU offload and enables higher Citrix Virtual Apps and Desktops server scalability.

When TCP is used as the data transport protocol, Citrix SD-WAN WANOP supports the optimization as described above. When using Citrix SD-WAN WANOP on network connections, choose TCP and disable EDT. By using TCP flow control and congestion control, Citrix SD-WAN WANOP ensures the equivalent interactivity to EDT at high latency and moderate packet loss.

For information on configuring adaptive transport on Citrix Virtual Apps and Desktops, see [Adaptive transport](#).

Citrix Hypervisor 6.5 upgrade

March 12, 2021

Important

To upgrade to Citrix Hypervisor version 6.5, the appliances must be running Citrix SD-WAN WANOP software release 9.0.x or later.

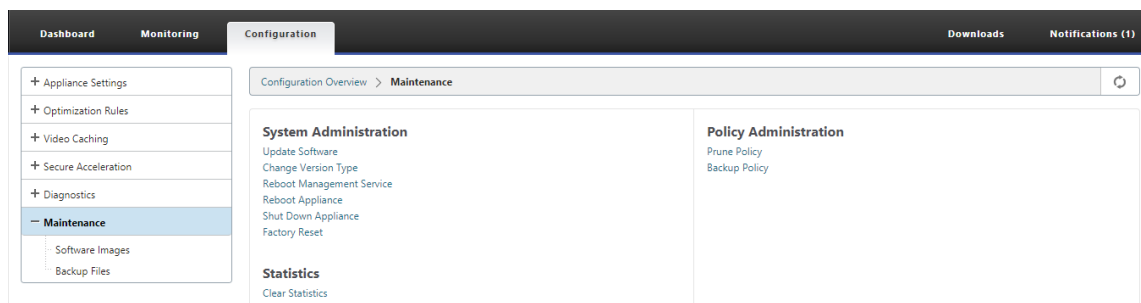
Note

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 Citrix Hypervisor 6.5

To upgrade to Citrix Hypervisor 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 Citrix SD-WAN WANOP GUI, go to **Configuration > Maintenance > Update Software**. Download the *ns-sdw-wo-<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.

Maintenance

March 12, 2021

Use the **Maintenance** page to perform maintenance activities such as upgrading/downgrading system software, backing up and restoring configurations and clearing statistics.



Upgrade/Downgrade

Upgrade system software


There is a different Citrix SD-WAN software package for each appliance model. You need to download the appropriate SD-WAN WANOP software package for an appliance you want to include in a network and save it in your local drive.

The appliance software is upgraded by means of patch files that you obtain from Citrix.

NOTE:

If the appliances are running older software release version, you need to upgrade to the latest software release version first.

To upgrade system software, go to **Configuration > Maintenance**. Select **Upgrade System Software** under **Upgrade/Downgrade**. select the patch file, and upload it to the appliance.




The patch file will be examined by the appliance. Only a valid patch file can upgrade the system to a different release from the one currently in use.

An upgrade preserves license files and system settings. The upgraded unit requires no reconfiguration except for any new features that have been added with the new release.

Change release

The change release page displays the currently installed release. If you want to change the release version, click **Change Release** option and select the release from the drop-down list and click **Change**.



Change version type

The **Change Version Type** option allows you to select a debug version of the release. You can select the version type from the **Type** drop-down list and click **Change**. The following are the possible debug versions:


- Default
- Level 1
- Level 2
- Default MC
- Level 1 MC
- Level 2 MC

You need to perform this action as instructed by the Support team.

Restart system

Once a patch is installed, a pop-up message will ask if the appliance can be restarted. The patch will not be applied until the appliance is restarted. If you select not to restart the system immediately, a reminder will be placed at the top of each page.

Click **Restart System** to restart the SD-WAN WANOP appliance. This process takes several minutes.



Backup settings


You can back up the appliance configuration by saving it as a text file.

Click **Save Settings**, a text file is downloaded to your local drive. License files, SSH parameters, and the IP addresses on the Management IP page cannot be saved. The file is an ordinary text file, but should not be edited manually.

Restore settings

Once the file is saved, it can be restored to the same SD-WAN WANOP appliance.


The appliance maintains copies of older releases. **Restore Settings** option helps to restore configured settings. Licenses files, SSH parameters, and the IP addresses on the Management IP page are not copied back from the newer release to the older one. Instead, the appliance will revert to the settings that were in effect at the time the older release was upgraded.



The screenshot shows a dialog box titled "Upload Settings file" with a close button (X) in the top right corner. Below the title bar, there is a text input field containing the filename "CitrixNetScalerSD-WANforCitrixXer". To the right of the input field is a "Browse" button with a dropdown arrow. Below the input field and the "Browse" button is an "Upload" button.

Reset to factory defaults


Reset to Factory Defaults option allows resetting the settings. It sets all the parameters except IP addresses, bandwidth settings, and licenses to their factory defaults. Click **Reset to Factory Defaults**, a confirmation message appears. Click **Yes** if you want to restore the settings to factory defaults.



The screenshot shows a "Confirm" dialog box with a close button (X) in the top right corner. The main content area contains a warning icon (a question mark in a blue square) followed by the text: "This action will cause a loss of almost all configured settings. The interface adapters and IP address settings will not be reset." Below this text is the question: "Are you sure you want to restore the settings to their factory defaults?". At the bottom of the dialog box are two buttons: "Yes" (highlighted in blue) and "No".

Clear statistics

Clear Statistics page allows resetting the SD-WAN WANOP appliance's statistics. It also allows creating reports that start at the beginning of the desired sampling window. Select the statistic options you want to clear from the appliance and click **Clear**.



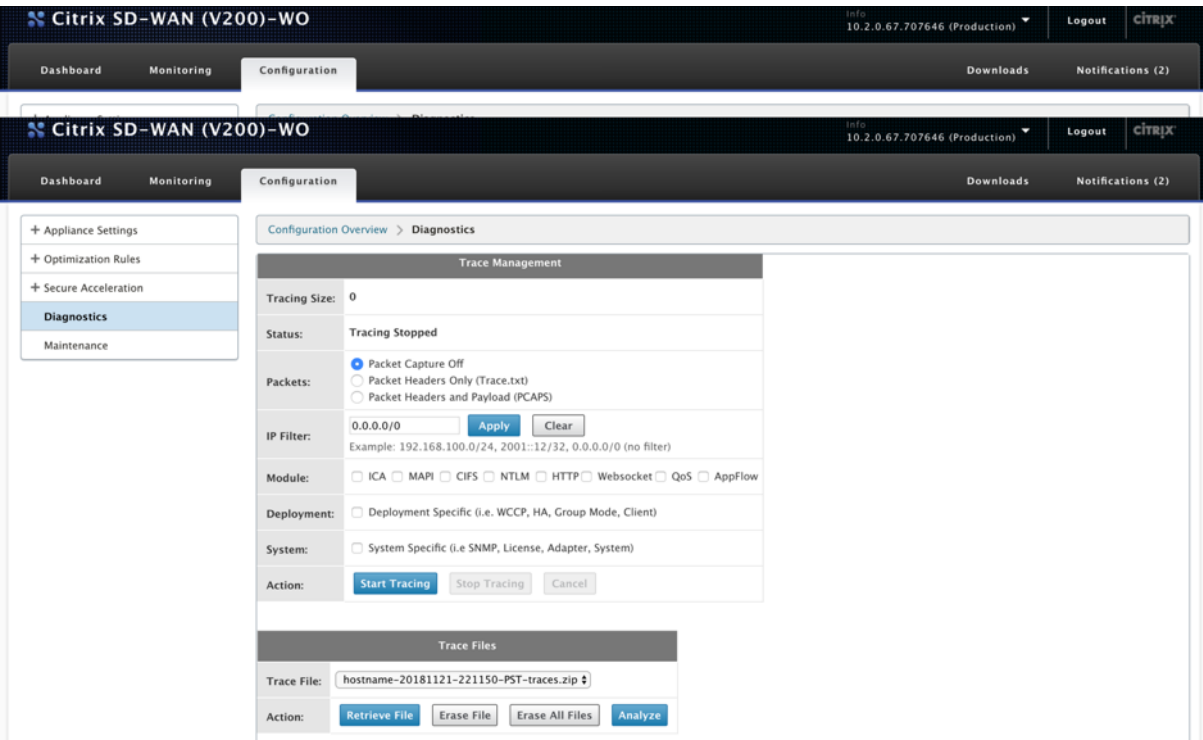
Diagnostics

March 12, 2021

This section provides diagnostic tools to identify network issues in your SD-WAN WANOP network and troubleshoot them. You can also obtain system log files, system information, and other necessary details that assist the Citrix SD-WAN Support team in diagnosing and resolving network issues.

Following are the Diagnostics tool available in SD-WAN WANOP:

- Tracing
- Packet Analyzer
- Bypass Card Test
- Retrieve Course
- Line Tester
- Ping
- Traceroute
- System Info
- Diagnostic Data




Tracing

The **Tracing** tool is used to watch the packets flowing over the SD-WAN WANOP network. It can open each packet and identify the protocol used, the IP address of the source and destination, and other payload information. This information is used by Citrix Support team to find the root cause of network issues.

You can choose to trace **Packet Headers Only** or **Packet Headers and Payload**. You can choose the module to trace and specify if the tracing should be deployment specific or system specific.

Click **Start Tracing**, the appliance begins to trace the packets. The results are packaged into a ZIP archive when you click **Stop Tracing**. This archive can be downloaded onto your computer, using the **Retrieve File** option. You can then forward these files to the Support team. The trace files also provide crash analysis data.


Click **Analyze** to view more information about the packets in **Packet Analyzer** tab.



You can view the time, source address, destination address, protocol, length, and payload information.

Bypass card test

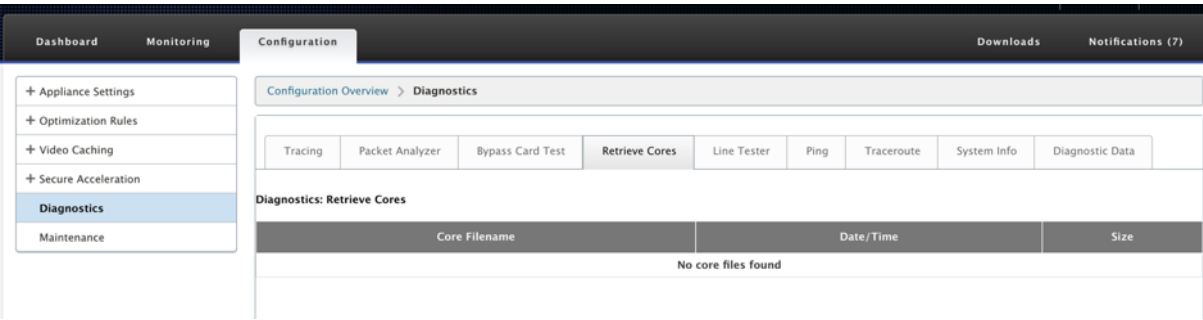
You can test the fail-to-wire functionality of Ethernet interface for an appliance deployment in inline (Fail-to-wire) mode. Enter the number of seconds for the appliance to stay in bypass mode and click **Start Test**. During this period, the appliance is bypassed. Normal operation will resume after that.



Retrieve cores


Core files are created when the SD-WAN WANOP appliance exits abnormally or crashes. The appliance restarts automatically after a crash. In case of persistent crashes, acceleration is disabled but the management interface remains active.

You can select and retrieve the required core files that were created during the appliance crash or when the appliance behaved abnormally. The retrieved files are saved in a ZIP archive. You can share this with the support team for further analysis.




Line tester

The **Line Test: SERVER** function starts an iperf server on the appliance, running in TCP mode. This option can be used to verify connectivity between WANOP appliances and troubleshooting network traffic. To run iperf tests, one system (an appliance or another host) must run iperf as a server, and another must connect to it as a client.




You can use the default **Line Tester Server** interface and port number. Click **Start Server** to start an iperf server on the appliance.




The **Line Test: CLIENT** function starts an iperf client on the unit, running in TCP mode. You can also specify the iperf server port number and the length of the test. When the test is complete, the connection speed will be reported. Click **Start Test** to see the WAN and LAN traffic result.

Test Results(COMPLETE)




Ping

Ping allows you to check connectivity of the network elements in your SD-WAN network. Enter the IP address of the network element and click **Run Ping** to see the result.



Traceroute

Traceroute allows you to record the route between your SD-WAN appliance and any other network element in your SD-WAN network or on the internet. It calculates and displays the amount of time each hop took.

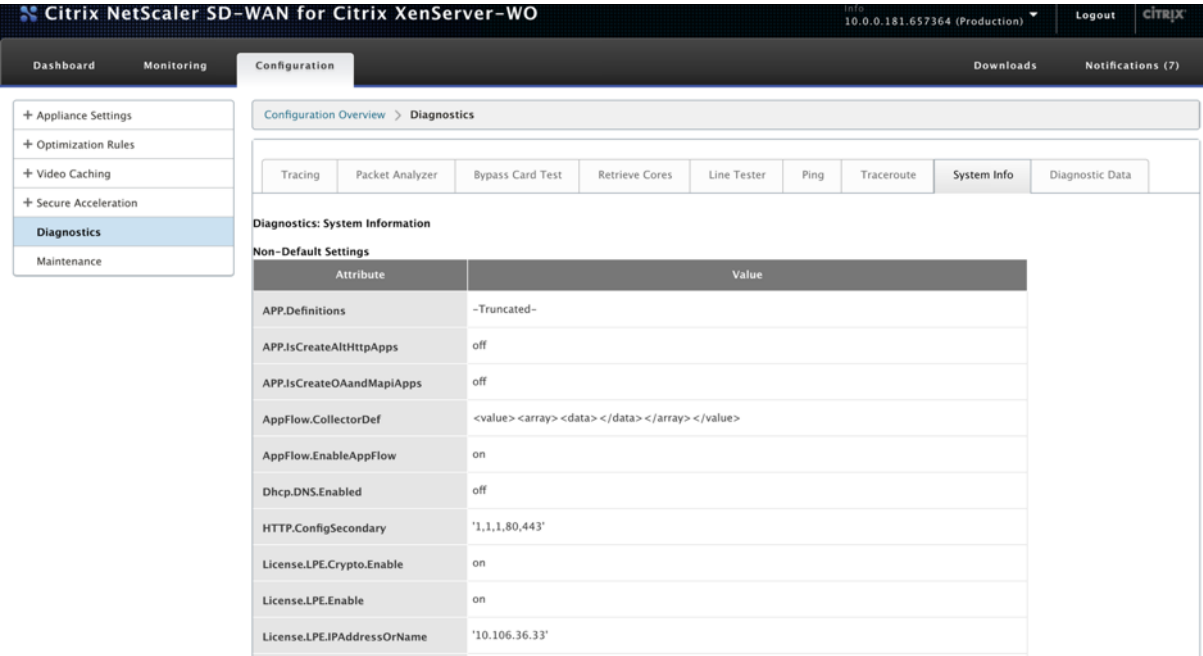


System info

The **System Info** lists all the parameters that are not set to their defaults. This information is read-only. It is used by Support when some kind of misconfiguration is suspected. When you report a problem,

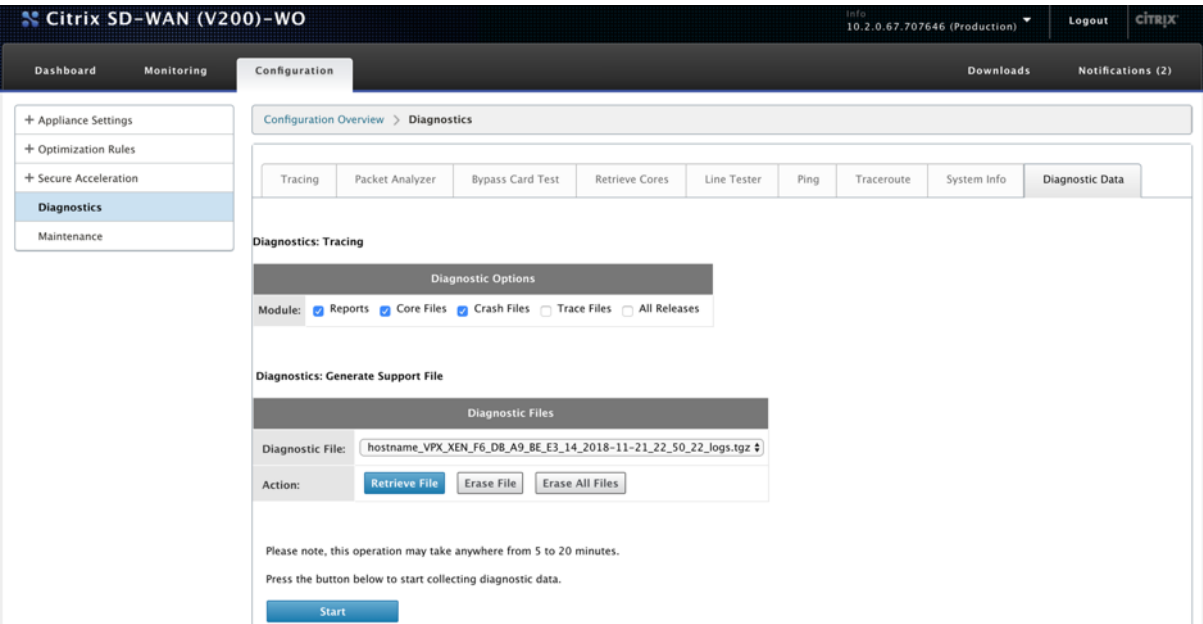
you may be asked to check one or more values on this page.

It provides **Non-Default Settings**, **Detailed Information For Adapter Primary**, **Detailed Information For Adapter apA.2**, and **Detailed Information For Adapter apA.1**.



Diagnostic Data

Diagnostic Data allows you to package diagnostic data for analysis by the Citrix Support team. Select the diagnostic files required and click **Start**. You can then, click **Retrieve File** to download the zip archive, and share it with Citrix Support.



Troubleshooting

March 12, 2021

The following topics provide a list of issues, the cause for the issue, and resolution steps for some Citrix SD-WAN WANOP features.

[CIFS and MAPI](#)

[Citrix SD-WAN WANOP plug-in](#)

[RPC over HTTPS](#)

[Video Caching](#)

[Citrix Virtual Apps and Desktops acceleration](#)

CIFS and MAPI

March 12, 2021

- **Issue:** A domain controller is removed from the network. However, the Citrix SD-WAN WANOP appliance is not able to leave the domain.

Cause: This is a known issue with the appliance.

Workaround: From the Windows Domain page, change the DNS to the one through which you can resolve the intended domain. Next, use the

Rejoin Domain option to make the Citrix SD-WAN WANOP appliance join that domain. Now try leaving from the domain.

- **Issue:** MAPI connections are not optimized and the following error message appears:
non-default setting in outlook is not supported

Cause: This is a known issue with release 6.2.3 and earlier releases.

Resolution: Upgrade the appliance to the latest release.

- **Issue:** The appliance optimized the MAPI connections. However, the monitoring pages display the number of send and received bytes as zero.

Cause: This is a known issue with the appliance.

Resolution: This is a benign issue and does not affect the functionality of the appliance. You can ignore it.

- **Issue:** Unable to establish secure peering between Citrix SD-WAN WANOP appliances.

Cause: Secure peering with the partner appliance is not properly configured.

Resolution: Do the following:

1. Verify that you have uploaded appropriate combination of CA and server certificates to the appliance.
 2. Navigate to the **Citrix SD-WAN WANOP > Configuration > SSL Settings > Secure Partners** page.
 3. In the **Partner Security** section, under **Certificate Verification**, select **None - allow all requests** option to make sure that certificate never expires.
 4. Verify that the appliance can establish secure peering with the partner appliance.
 5. Verify that the **Listen On** section has an entry for the IP address of the intended Citrix SD-WAN WANOP appliance.
- **Issue:** When connecting to an Exchange cluster, Outlook users with optimized connections are occasionally bypassed or prompted for logon credentials.

Cause: MAPI optimization requires that each node in the Exchange cluster be associated with the exchangeMDB service principal name (SPN). Over time, as you need more capacity, you add additional nodes to the cluster. However, sometimes, the configuration task might not be completed, leaving some nodes in cluster without SPN settings. This issue is most prevalent in Exchange clusters with Exchange Server 2003 or Exchange Server 2007.

Resolution: Do the following on each Exchange servers in the set up:

1. Access the domain controller.
2. Open the command prompt.
3. Run the following commands:

```
pre codeblock setspn -A exchangeMDB/Exchange1 Exchange1
setspn -A exchangeMDB/Exchange1.example.com Exchange1 <!--
NeedCopy-->
```

- **Issue:** When attempting to connect to Outlook, the Trying to connect message is displayed and then the connection is terminated.

Cause: The client-side Citrix SD-WAN WANOP appliance has blacklist entries that do not exist on the server-side appliance.

Resolution: Remove the blacklist entries from both appliances, or (recommended) upgrade the software of the appliances to release 6.2.5 or later.

- **Issue:** The appliance fails to join the domain even after passing the pre domain checks.

Cause: This is a known issue.

Resolution: Do the following:

1. Access the appliance by using an SSH utility.
2. Log on to the appliance by using the root credentials.
3. Run the following command:

```
/opt/likewise/bin/domainjoin-cli join \<Domain\\_Name\>  
administrator
```

- **Issue:** The LdapError error message appears when you add a delegate user to the Citrix SD-WAN WANOP appliance.

Resolution: Do one of the following:

- On the Citrix SD-WAN WANOP appliance's DNS server, verify that a reverse lookup zone is configured for every domain-controller IP address.
- Verify that the system clock of the client machine is synchronized with the system clock of the Active Directory server. When using Kerberos, these clocks must be synchronized.
- Update the delegate user on the Windows Domain page by providing the password for the delegate user once again.

- **Issue:** The Time skew error message appears when you add a delegate user to the Citrix SD-WAN WANOP appliance.

Resolution: Verify that the appliance is joined to the domain. If not, join the appliance to the domain. This synchronizes the appliance time with the domain-server time and resolves the issue.

- **Issue:** The Client is temporarily excluded for acceleration. Last Error (Kerberos error.) error message appears when you add a delegate user to the Citrix SD-WAN WANOP appliance.

Cause: The delegate user is configured for the **Use Kerberos only** authentication.

Resolution: Verify that, on the domain controller, the delegate user's authentication setting is **Use any authentication protocol**.

- **Issue:** The Delegate user not ready error message appears when you add a delegate user to the Citrix SD-WAN WANOP appliance.

Resolution: If the message appears only on the client-side appliance, ignore it. However, if the message is displayed on the server-side appliance, run the delegate user precheck tool, available on the **Windows Domain** page, and then configure the delegate user on the server-side appliance.

- **Issue:** The Last Error (The Server is not delegated for Kerberos authentication. Please add delegate user, check list for services and server allowed for delegation.) UR:4 error message appears when you add a delegate user to the Citrix SD-WAN WANOP appliance.

Resolution: Verify that the delegate user is correctly configured on the domain controller and that you have added appropriate services to the domain controller.

- **Issue:** The appliance is not able to join the domain.

Resolution: Run the domain precheck tool, available on the Windows Domain page, and resolve the issues, if any. If the domain precheck tool does not report any issues, contact Citrix Technical Support for further assistance in resolving the issue.

Citrix SD-WAN WANOP plug-in

March 12, 2021

- **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 Citrix 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.

RPC over HTTPS

March 12, 2021

- **Issue:** After upgrading the software of the appliance to release 7.3, the monitoring reports do not have a special category for RPC over HTTPS connections.

Cause: When you upgrade the appliance to release 7.3, the RPC over HTTPS applications do not belong to their own service class. As a result, all RPC over HTTPS connections are listed as TCP Other connections in the reports.

Resolution: To categorize these connections as RPC over HTTPS connections, create a service class for them applications.

- **Issue:** After creating a service class for RPC over HTTPS, all HTTP and HTTPS traffic is categorized as RPC over HTTP.

Cause: You have not added the destination IP address to the service class you have created for RPC over HTTPS applications.

Resolution: Modify the service class you have created for RPC over HTTPS applications, by adding the destination IP addresses of your servers.

Video caching

March 12, 2021

- **Issue:** After adding an entry to the list of prepopulation tasks, the entry is still in the Configured state.

Cause: A prepopulation task takes approximately one minute move to the Downloading state.

Resolution: Check the status of the entry after a minute or refresh the page to verify that the status changes to Downloading.

- **Issue:** After adding an entry to the list of prepopulation tasks, the status of the entry displays ERROR 403. However, the website works fine in a Web browser.

Cause: The IP address of the Citrix SD-WAN WANOP apA does not have access to the video server.

Resolution: To resolve this issue, verify and update the following:

- Access rules across the firewalls
- Source IP address based limitations in the httpd.conf file of the video server

Cause: The video server does not support the HEAD method.

Resolution: The video server must permit Citrix SD-WAN WANOP IP address for this method.

Cause: Directory listing for folders is not enabled on the video server.

Resolution: The video server must enable directory listing for the folders.

- **Issue:** After creating entries for prepopulation tasks, you cannot modify or delete entries.

Cause: You might have clicked **Start Now** for the entry.

Resolution: This is by design. You cannot modify or delete an entry after you have clicked **Start Now** for the entry and the entry is in the queued, starting, or downloading state. You can delete the entry only after the download is complete.

- **Issue:** After creating entries for propopulation tasks, video is not getting downloaded and cached. The status of the entry displays Failed to Download.

Cause: The prepopulation entry does not have absolute URL for the video.

Resolution: To resolve this issue, complete the following procedure:

1. Verify that the prepopulation entry has the actual URL of the video, such as [http://10.102.29.16/Citrix SD-WAN WANOP_demo.mp4](http://10.102.29.16/Citrix%20SD-WAN%20WANOP_demo.mp4), and not an HTML file. The Citrix SD-WAN WANOP appliance cannot search the content of the HTML file to find the video link.
2. Verify that the HTTP protocol is used to serve the video. You can verify this by using the View Source option of the web browser.
3. You can get the absolute URL of the video by using the Developer Tools option of the web browser.

Citrix Virtual Apps and Desktops acceleration

March 12, 2021

- **Issue:** After upgrading an appliance to release 7.3.1, the ICA connections are not categorized as Citrix Receiver for HTML5 connections in the ICA Monitoring pages.

Cause: Service class defined on the appliance is **HTTP (Private)** instead of Web (Private). When you upgrade an appliance to release 7.3.1, the **ALTHHTTP** application is not added to this service class. As a result, even though ICA connections over Citrix Receiver for HTML5 are optimized, these are not categorized as Citrix Receiver for HTML5 connections in the ICA Monitoring pages.

Resolution : To categorize ICA connections over Citrix Receiver for HTML5, complete the following procedure:

1. Navigate to the **Configuration > Optimization Rules > Service Classes** page.
2. Edit the **HTTP (Private)** service class.
3. Click **Add Rule**.
4. In Filter Rules, under Applications, click **Any**.
5. From the Applications list, select **ALTHHTTP**.
6. Click **Add**.
7. Click **Save**.
8. Make other changes to the filter rule, as required.
9. Click **Save**.



© 2024 Cloud Software Group, Inc. All rights reserved. Cloud Software Group, the Cloud Software Group logo, and other marks appearing herein are property of Cloud Software Group, Inc. and/or one or more of its subsidiaries, and may be registered with the U.S. Patent and Trademark Office and in other countries. All other marks are the property of their respective owner(s).