Improve Throughput on Catalyst 8000V in Azure

Contents

Introduction

This document explains how to improve performance of Cisco Catalyst 8000Vs deployed in Azure.

Catalyst 8000V Throughput Improvement in Azure

With Cisco Cloud OnRamp for Multicloud, users can deploy Cisco Catalyst 8000V virtual routers in NVA in Azure directly with SD-WAN Manager (UI or API).

Cloud OnRamp automation allows users to seamlessly create and discover virtual WAN, virtual hubs, and build connections to virtual networks in Azure.

Once Cisco Catalyst 8000Vs are deployed in Azure, the virtual appliances can be monitored and managed from SD-WAN Manager.

This document explains how to improve performance in Azure from three perspectives:

- installation of HSEC license;
- throughput limitations on TCP 12346 port in Azure;
- auto-negotiated speed on transport interface.

Installation of HSEC License

Devices that use Smart Licensing Using Policy, and that must support an encrypted traffic throughput of 250 Mbps or greater, require an HSEC license.

This is a requirement of US export control regulation. You can use Cisco SD-WAN Manager to install HSEC licenses.

Cisco SD-WAN Manager contacts Cisco Smart Software Manager (SSM), which provides a smart license authorization code (SLAC) to load onto a device.

Loading the SLAC on a device enables an HSEC license.

Refer to Managing HSEC Licenses in Cisco Catalyst SD-WAN for details on installing and management of the licenses.

Throughput Limitations on TCP 12346 Port in Azure

Currently, the automation deploys C8000V with one transport interface (GigabitEtherent1) and one service interface (GigabitEtherent2).

Due to Azure inbound limitations on SD-WAN port TCP 12346, the throughput can be limited on per-transport interface basis as traffic enters Azure infrastructure.

The inbound limit of 200K PPS is enforced by Azure infrastructure and thus users cannot achieve more than ~1Gbps per C8000V NVA instance (an example assumption: a packet size of 600B, calculation: 600B * 8 = 4800bits; 4800b * 200Kpps = 960 Mbps).

> **Note**: Azure can increase the inbound limit to 400K PPS on per case (ticket) basis. Customers need to contact Azure directly and request the increase.

To overcome this limitation, Cisco collaborated with Azure to allow SD-WAN branches to build multiple SD-WAN tunnels to each NVA instance.

To make this configuration change, the administrator must follow these steps:

1. In SD-WAN Manager, deploy cloud gateway with C8000V in Azure using Cloud OnRamp automation.

2. In Azure portal, change IP settings for NVA in virtual hub.
3. In SD-WAN Manager, create and push a new config group utilizing the settings from cloud portal.

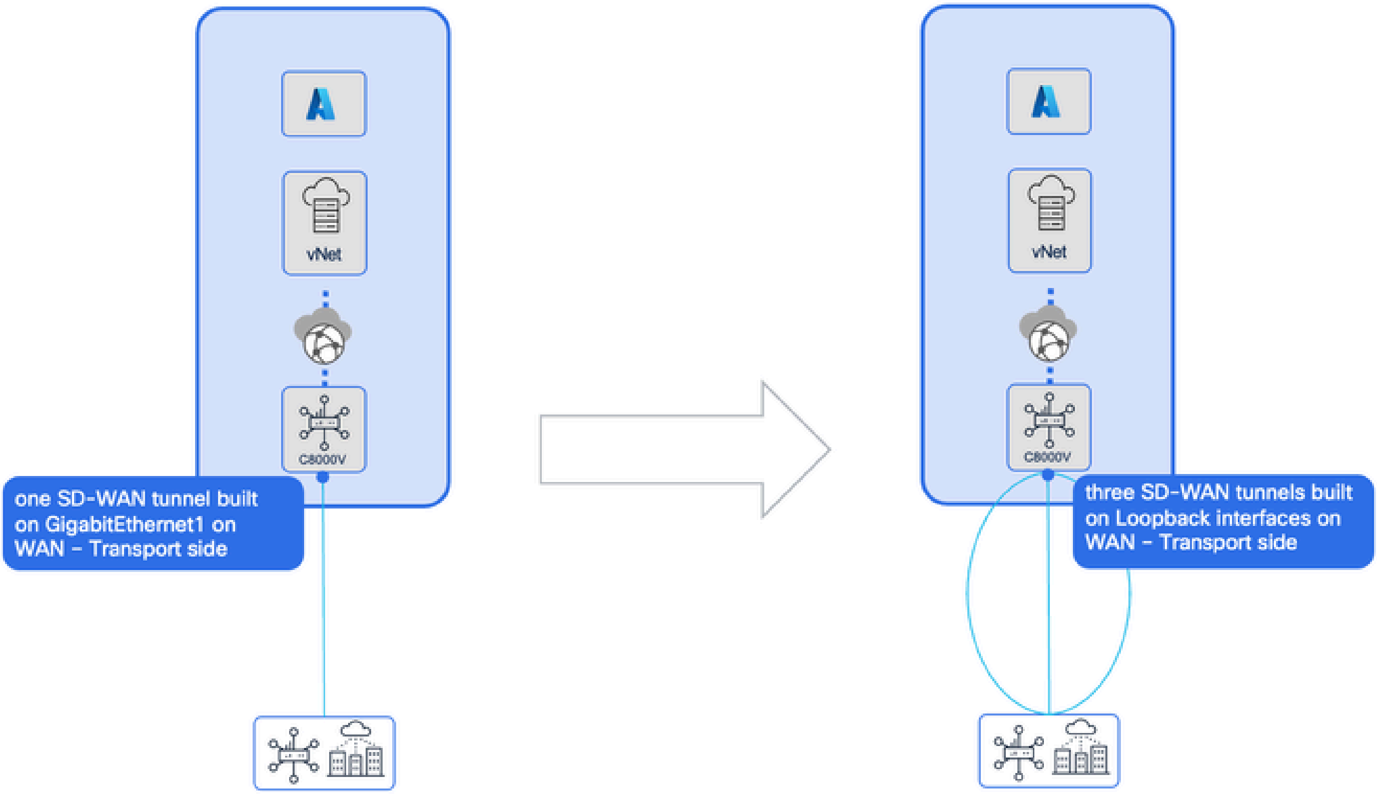
vNet
C8000V
one SD-WAN tunnel built on GigabitEthernet1 on WAN – Transport side
vNet
C8000V
three SD-WAN tunnels built on Loopback interfaces on WAN – Transport side

Step 1:

Deploy Cisco Catalyst 8000V in Azure using the procedure found here at this [Youtube channel]() or [Release Notes]().

Step 2:
For changing the IP settings, navigate to **Azure porta > Virtual WANs > chosen virtual WAN > virtual hub > NVA** in virtual hub.

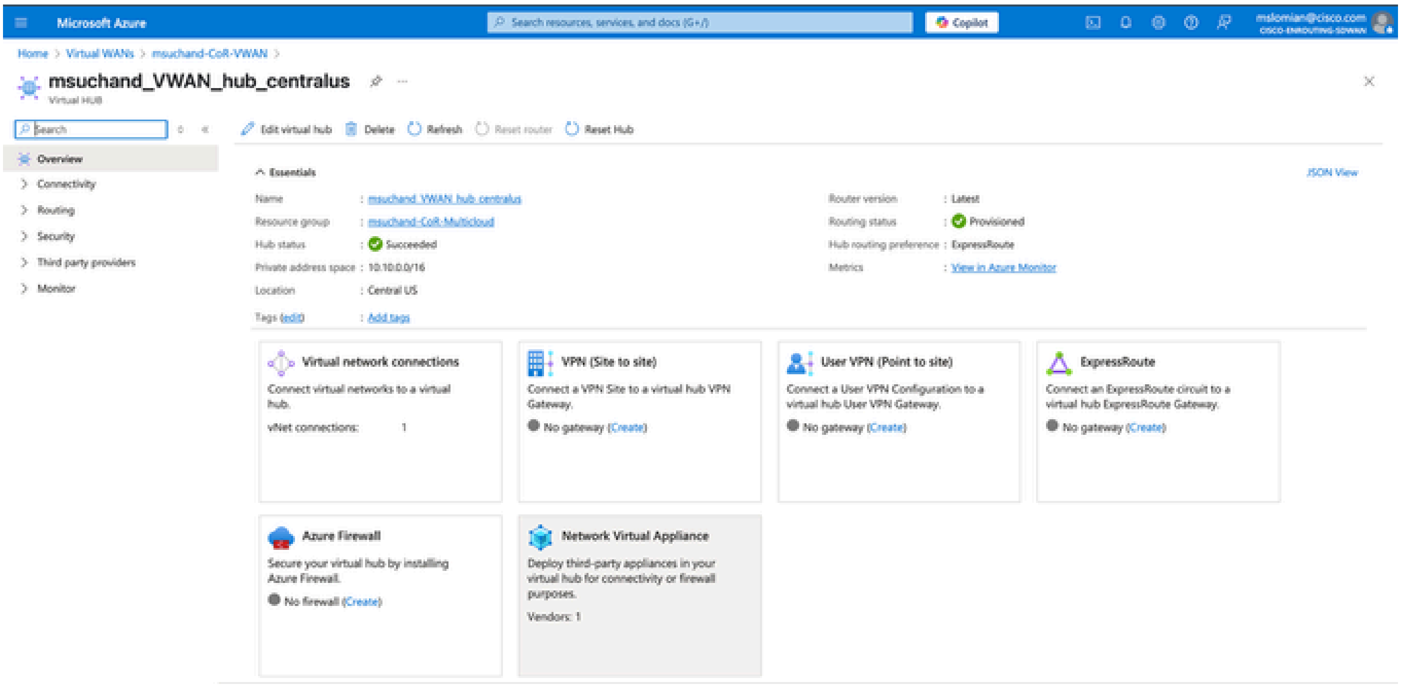

Name : msuchand_VWAN_hub_centralus
Resource group : msuchand-CoR-Multicloud
Hub status : Succeeded
Private address space : 10.10.0.0/16
Location : Central US
Tags (edit) : Add tags

Router version : Latest
Routing status : Provisioned
Hub routing preference : ExpressRoute
Metrics : View in Azure Monitor

Virtual network connections
Connect virtual networks to a virtual hub.
vNet connections: 1

VPN (Site to site)
Connect a VPN Site to a virtual hub VPN Gateway.
● No gateway (Create)

User VPN (Point to site)
Connect a User VPN Configuration to a virtual hub User VPN Gateway.
● No gateway (Create)

ExpressRoute
Connect an ExpressRoute circuit to a virtual hub ExpressRoute Gateway.
● No gateway (Create)

Azure Firewall
Secure your virtual hub by installing Azure Firewall.
● No firewall (Create)

Network Virtual Appliance
Deploy third-party appliances in your virtual hub for connectivity or firewall purposes.
Vendors: 1

In the virtual hub view on NVA, navigate to **Third party providers > Manage Configurations**.

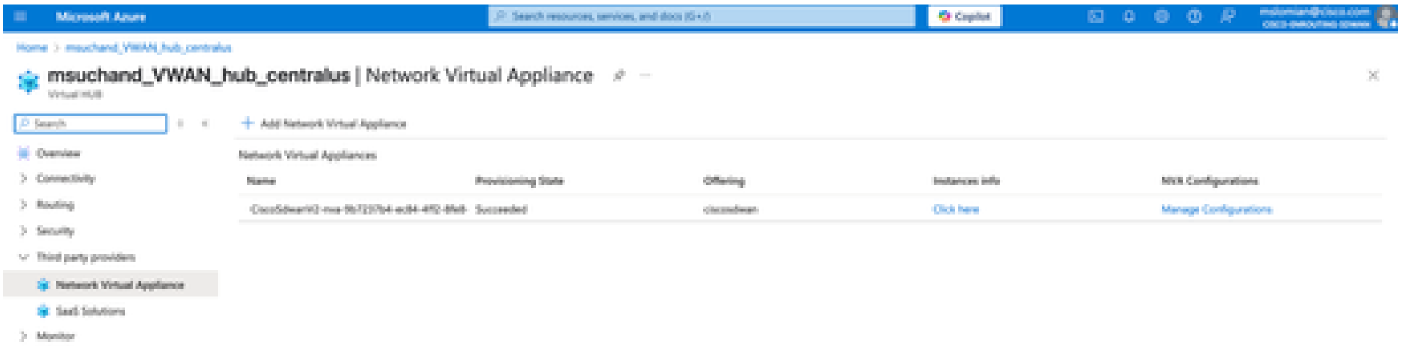
Microsoft Azure
Search resources, services, and docs (G+/)
Copilot
Home > msuchand_VWAN_hub_centralus
Search
Overview
Connectivity
Routing
Security
Third party providers
Network Virtual Appliance
SaaS Solutions
Monitor
Add Network Virtual Appliance

msuchand_VWAN_hub_centralus | Network Virtual Appliance

Virtual HUB

Network Virtual Appliances

Name	Provisioning State	Offering	Instances info	NVA Configurations
CiscoSdwanV2-nva-9b7237b4-ec84-4ff2-8fe8-	Succeeded	ciscosdwan	Click here	Manage Configurations

In the NVA configuration, navigate to **Interface IP Configurations** and **Add Configurations**. It can take up to 30 minutes to assign IP addresses,

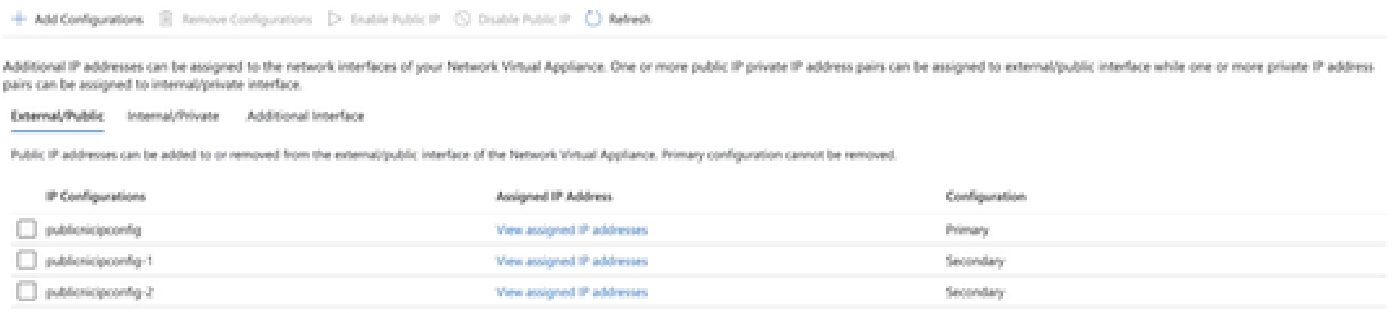

Additional IP addresses can be assigned to the network interfaces of your Network Virtual Appliance. One or more public IP private IP address pairs can be assigned to external/public interface while one or more private IP address pairs can be assigned to internal/private interface.

Public IP addresses can be added to or removed from the external/public interface of the Network Virtual Appliance. Primary configuration cannot be removed.

IP Configurations	Assigned IP Address	Configuration
□ publicnicipconfig	View assigned IP addresses	Primary
□ publicnicipconfig-1	View assigned IP addresses	Secondary
□ publicnicipconfig-2	View assigned IP addresses	Secondary

Step 3:

Once addresses are assigned, make a note of them and go to SD-WAN Manager. All C8000Vs need this config update.

It can be done by CLI Addon (it appends whatever is in templates / configuration profiles). Refer to this example configuration:

```
interface Loopback 1000
    ip address 10.0.0.244 255.255.255.255
    no shut
exit
interface Loopback 2000
    ip address 10.0.0.246 255.255.255.255
    no shut
exit
interface Loopback 3000
    ip address 10.0.0.247 255.255.255.255
    no shut
exit
interface GigabitEthernet1
    speed 10000
    no ip dhcp client default-router distance 1
    no ip address dhcp client-id GigabitEthernet1
    ip unnumbered Loopback1000
exit
interface GigabitEthernet2
    speed 10000
exit
```

```
ip route 0.0.0.0 0.0.0.0 10.0.0.241 → 10.0.0.241 IP is Loopback 1000 IP -3
ip route 10.0.0.241 255.255.255.255 GigabitEthernet1 → 10.0.0.241 IP is Loopback 1000 IP -3
interface Tunnel1
    no shutdown
    ip unnumbered Loopback1000
    ipv6 unnumbered Loopback1000
    tunnel source Loopback1000
    tunnel mode sdwan
interface Tunnel2
    no shutdown
    ip unnumbered Loopback2000
    ipv6 unnumbered Loopback2000
    tunnel source Loopback2000
    tunnel mode sdwan
interface Tunnel3
    no shutdown
    ip unnumbered Loopback3000
    ipv6 unnumbered Loopback3000
    tunnel source Loopback3000
    tunnel mode sdwan
sdwan
 interface Loopback1000
   tunnel-interface
    encapsulation ipsec weight 1
    no border
    color biz-internet
    no last-resort-circuit
    no low-bandwidth-link
    no vbond-as-stun-server
    vmanage-connection-preference 5
    port-hop
    carrier                       default
    nat-refresh-interval          5
    hello-interval                1000
    hello-tolerance               12
    no allow-service all
    no allow-service bgp
    allow-service dhcp
    allow-service dns
    allow-service icmp
    allow-service sshd
    no allow-service netconf
    no allow-service ntp
    no allow-service ospf
    no allow-service stun
    allow-service https
    no allow-service snmp
    no allow-service bfd
  exit
 exit
 interface Loopback2000
   tunnel-interface
    encapsulation ipsec weight 1
    no border
    color public-internet
    no last-resort-circuit
    no low-bandwidth-link
    no vbond-as-stun-server
    vmanage-connection-preference 4
    port-hop
    carrier                       default
    nat-refresh-interval          5
    hello-interval               1000
    hello-tolerance              12
    no allow-service all
    no allow-service bgp
    allow-service dhcp
    allow-service dns
    allow-service icmp
    allow-service sshd
    no allow-service netconf
    no allow-service ntp
    no allow-service ospf
    no allow-service stun
    allow-service https
    no allow-service snmp
    no allow-service bfd
  exit
 exit
 interface Loopback3000
   tunnel-interface
    encapsulation ipsec weight 1
    no border
    color custom1
    no last-resort-circuit
    no low-bandwidth-link
    no vbond-as-stun-server
    vmanage-connection-preference 3
    port-hop
    carrier                      default
    nat-refresh-interval         5
    hello-interval               1000
    hello-tolerance              12
    no allow-service all
    no allow-service bgp
    allow-service dhcp
    allow-service dns
    allow-service icmp
    allow-service sshd
    no allow-service netconf
    no allow-service ntp
    no allow-service ospf
    no allow-service stun
    allow-service https
    no allow-service snmp
    no allow-service bfd
  exit
 exit
 interface GigabitEthernet1
  no tunnel-interface
  exit
 exit
```

Auto-negotiated Speed on Transport Interface

Cisco transport interface on Cisco Catalyst 8000V, which used in default templates or auto-generated configuration groups (GigabitEthernet1), is configured with negotiate auto to make sure that the connection is established.

In order to get better performance (above 1Gb), it is recommended to set the speed on interfaces to 10Gb.

This is applicable also to the service interface (GigabitEthernet2). To verify the negotiated speed, run these commands:

```
azure-central-us-1#sh int gi1
GigabitEthernet1 is up, line protocol is up
  Hardware is vNIC, address is 000d.3a92.e2ff (bia 000d.3a92.e2ff)
  Internet address is 10.48.0.244/28
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,

…
azure-central-us-1#sh int gi2
GigabitEthernet2 is up, line protocol is up
  Hardware is vNIC, address is 000d.3a92.ea8a (bia 000d.3a92.ea8a)
  Internet address is 10.48.0.229/28
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
```

Note: Although this article is focused on C8000V deployment in Azure with Cloud OnRamp automation (NVA), the auto-negotiated speed is also applicable to Azure deployments in VNets, AWS, and Google deployments.

In order to change this, make changes in template (**Confugruation > Templates > Feature Template > Cisco VPN Interface Ethernet) / configuration group** (see the guide). Alternatively, administrators can edit this in CLI, if the device is CLI-managed.