Configure Google Cloud Interconnect as a Transport with Cisco SD-WAN in a Click

Contents

[Introduction](#)
[Background Information](#)
[Problem](#)
[Solution](#)
[Design Overview](#)
[Solution Details](#)
[Step 1. Preparation](#)
[Step 2. Create Cisco Cloud Gateway with Cloud onRamp for Multicloud Workflow](#)
[Step 3. In GCP Console Add a Partner Interconnect Connection](#)
[Step 4. Use Cloud onRamp Interconnect in Cisco vManage to Create the DC Connection](#)
[Step 5. Configure DC Router to Establish Tunnels over Internet and over GCP Cloud Interconnect](#)
[Verify](#)
[DC Megaport SD-WAN Router Configuration](#)

Introduction

This document describes how to use Google Cloud Interconnect as Software-defined Wide Area Network (SD-WAN) transport.

Background Information

Enterprise customers with workloads on Google Cloud Platform (GCP) use Cloud Interconnect for Data Center or Hub connectivity. At the same time, public internet connection is also very common in Data Center and is used as an underlay for SD-WAN connectivity with other locations. This article describes how GCP Cloud Interconnect can be used as an underlay for Cisco SD-WAN.

It is very similar to that describes the same solution for AWS.

Key benefit of using GCP Cloud Interconnect as just another transport for Cisco SD-WAN is the ability to use SD-WAN policies over all transports including GCP Cloud Interconnect. Customers can create SD-WAN application-aware policies and route critical applications over GCP Cloud Interconnect and reroute via public internet in case of SLA violations.

Problem

GCP Cloud Interconnect does not provide native SD-WAN capabilities. Typical questions from Enterprise SD-WAN customers are:

- "Can I use GCP Cloud Interconnect as an underlay for Cisco SD-WAN"?
- "How can I interconnect GCP Cloud Interconnect and Cisco SD-WAN"?

- "How can I create a resilient, secure and scalable solution"?

Solution

Design Overview

The key design point is the connection of the Data Center via GCP Cloud Interconnect to Cisco SD-Routers created by Cloud onRamp for Multicloud provisioning as shown in the image.

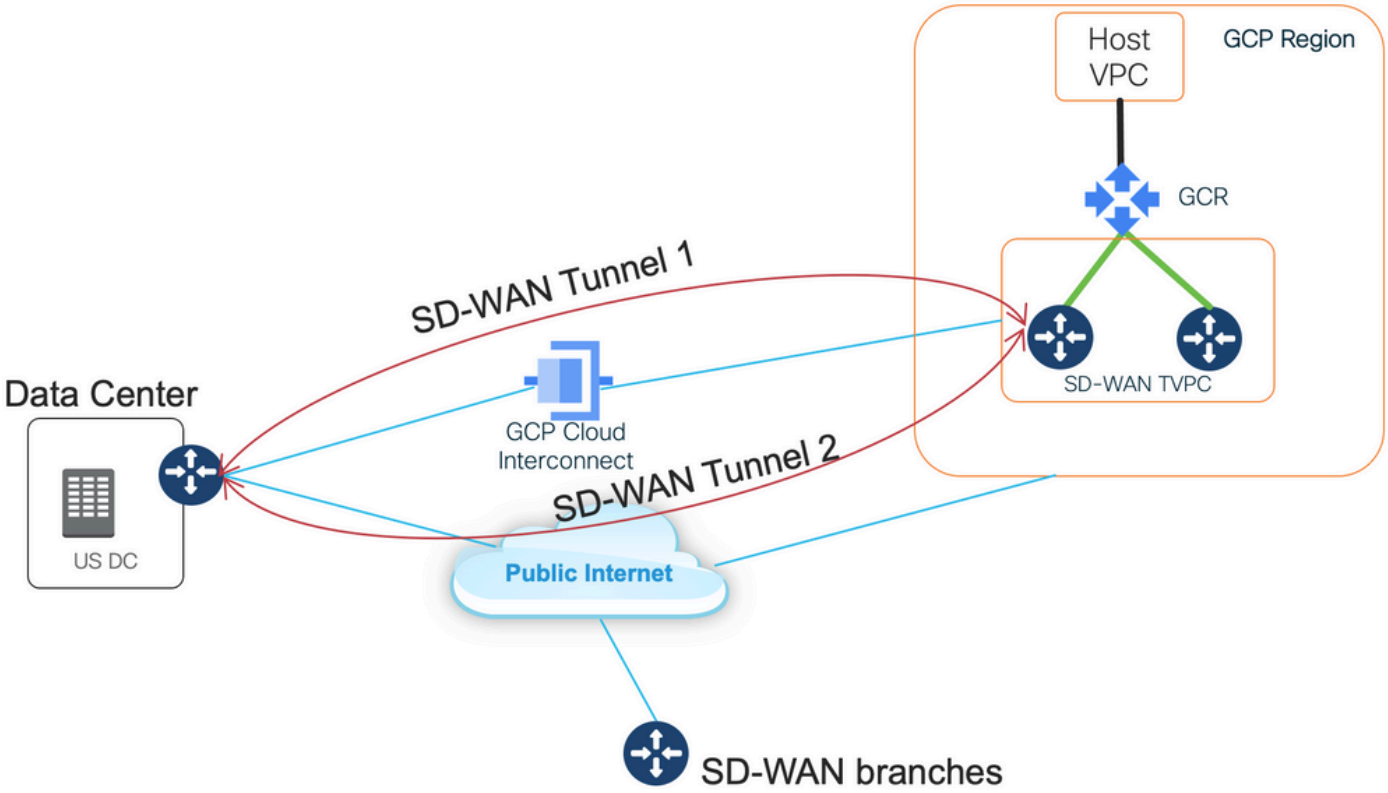

The benefits of this solution are:

- Fully Automatic: Cisco Cloud onRamp for Multicloud automation can be used to deploy SD-WAN transit VPC with two SD-WAN routers. Host VPCs can be discovered as a part of Cloud onRamp and mapped to SD-WAN networks with one click.
- Full SD-WAN over GCP Cloud Interconnect: GCP Cloud Interconnect is just another SD-WAN transport. All SD-WAN features like application-aware policies, encryption, etc. can be natively used on the SD-WAN tunnel over GCP Cloud Interconnect.

Please note, that the scalability of this solution goes along with C8000V performance on GCP. Please refer to SalesConnect for details on C8000v performance on GCP.

Solution Details

The key point to understand this solution is SD-WAN Colors. Please note, that GCP SD-WAN routers will have **private color private2** for the internet connectivity as well as connectivity via Interconnect, SD-WAN tunnels will be formed over the Internet using public IP addresses as well as SD-WAN tunnels will be established (using the same interface) over the Interconnect circuits using private IP addresses to a DC/Site. This means, that the Data Center router (biz-internet color) will establish a connection to GCP SD-WAN routers (private2 color) via the Internet with public IP addresses and via Its Private colour over Private IP.

Generic Information about SD-WAN Colors:

Transport Locators (TLOCs) refer to the WAN transport (VPN 0) interfaces by which SD-WAN routers connect to the underlay network. Each TLOC is uniquely identified through a combination of the system IP address of the SD-WAN router, the color of the WAN interface, and the transport encapsulation (GRE or IPsec). The Cisco Overlay Management Protocol (OMP) is used to distribute TLOCs (also known as TLOC routes), SD-WAN overlay prefixes (also known as OMP routes), and other information between SD-WAN routers. It is through TLOC routes that SD-WAN routers know how to reach each other and establish IPsec VPN tunnels with each other.

SD-WAN routers and/or controllers (vManage, vSmart, or vBond) may sit behind Network Address Translation (NAT) devices within the network. When an SD-WAN router authenticates to a vBond controller, the vBond controller will learn both the private IP address/port number and the public IP address/port number settings of the SD-WAN router during the exchange. vBond controllers act as Session Traversal Utilities for NAT (STUN) servers, allowing SD-WAN routers to discover mapped and/or translated IP addresses and port numbers of their WAN transport interfaces.

On SD-WAN routers every WAN transport is associated with a public and private IP address pair. The private IP address is considered to be the pre-NAT address. This is IP address assigned to the WAN interface of the SD-WAN router. Although this is considered to be the private IP address, this IP address can be either part of the publicly routable IP address space or part of the IETF RFC 1918 non-publicly routable IP address space. The public IP address is considered to be the post-NAT address. This is detected by the vBond server when the SD-WAN router initially communicates and authenticates with the vBond server. The public IP address can also be either part of the publicly routable IP address space or part of the IETF RFC 1918 non-publicly routable IP address space. In the absence of NAT, both the public and private IP addresses of the SD-WAN transport interface are the same.

TLOC colors are statically defined keywords used to identify individual WAN transports on each SD-WAN router. Each WAN transport on a given SD-WAN router must have a unique color. Colors are also used to identify an individual WAN transport as being either public or private. The colors metro-ethernet, Mpls, and private1, private2, private3, private4, private5, and private6 are considered private colors. They are intended for use in private networks or places where there is no NAT. The colors are 3g, biz-internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, public-internet, red, and silver are considered public colors. They are intended to be used in public networks or in places with public IP addressing of the WAN transport interfaces, either natively or through NAT.

Color dictates the use of either private or public IP addresses when communicating through the control and data planes. When two SD-WAN routers attempt to communicate with each other, both using WAN transport interfaces with private colors, each side will attempt to connect to the remote router’s private IP address. If one or both sides are using public colors, then each side will attempt to connect to the remote router’s public IP address. An exception to this is when the Site IDs of two devices are the same. When the Site IDs are the same, but the colors are public, the private IP addresses will be used for communication. This may occur for SD-WAN routers attempting to communicate to a vManage or vSmart controller located within the same site. Note that SD-WAN routers do not, by default, establish IPsec VPN tunnels between each other when they have the same Site IDs.

Here is the output from the Data Center router, which shows two tunnels via Internet (color biz-internet) and two tunnels via GCP Cloud Interconnect (color private1) to two SD-WAN routers. Refer to the full DC router configuration in the attachment for more details.

```
MP-IC-US-R1#sh sdwan bfd sessions
SOURCE TLOC REMOTE TLOC DST PUBLIC DST PUBLIC DETECT TX
SYSTEM IP SITE ID STATE COLOR COLOR SOURCE IP IP PORT ENCAP MULTIPLIER INTERVAL(msec UPTIME
TRANSITIONS
-----------------------------------------------------------------------------------------------
-----------------------------------------------------------------------------------------------
------------------
35.35.35.2 35 up biz-internet private2 162.43.150.15 35.212.162.72 12347 ipsec 7 1000 10
4:02:55:32 0
35.35.35.1 35 up biz-internet private2 162.43.150.15 35.212.232.51 12347 ipsec 7 1000 10
4:02:55:32 0
35.35.35.1 35 up private1 private2 192.168.9.9 10.35.0.2 12347 ipsec 7 1000 10 0:00:00:16 0
35.35.35.2 35 up private1 private2 192.168.9.9 10.35.0.3 12347 ipsec 7 1000 10 0:00:00:16 0
...
MP-IC-US-R1#
```

This image illustrates topology details with IP addresses and SD-WAN colors, that are used to verify the solution.

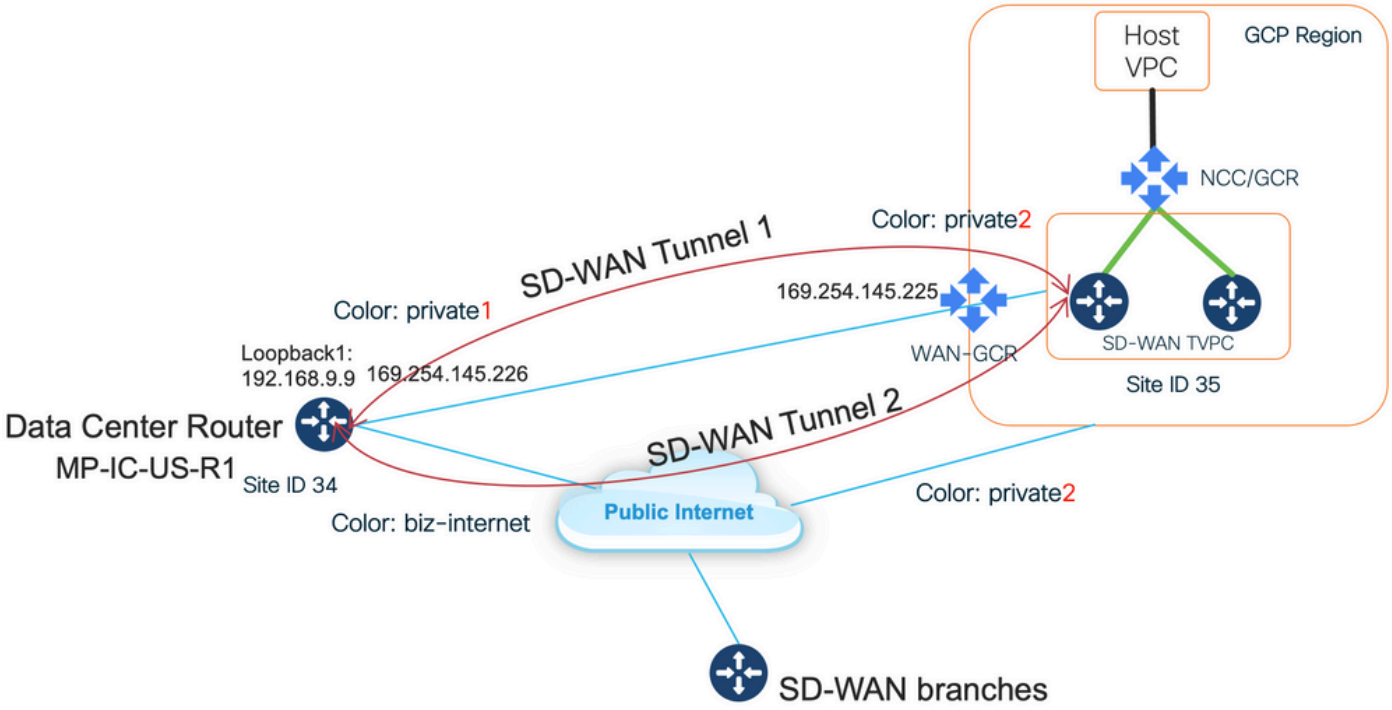

Software used:

- SD-WAN Controllers running CCO Version 20.7.1.1
- Data Center Router simulated with C8000v running 17.06.01a provisioned via vManage Cloud onRamp for Interconnect with Megaport
- Two SD-WAN routers in GCP: C8000v running 17.06.01a provisioned via vManage Cloud onRamp for Multicloud

Step 1. Preparation

Ensure that Cisco vManage has a working GCP account defined and Cloud onRamp Global Settings are configured properly.

Please also define an Interconnect Partner Account in vManage as well. In this blog Megaport is used as Interconnect partner, so you can define an appropriate account and global settings.

Step 2. Create Cisco Cloud Gateway with Cloud onRamp for Multicloud Workflow

This is a straightforward process: select two SD-WAN devices, attach the default GCP template, deploy. Please refer to Cloud onRamp for Multicloud documentation for details.

Step 3. In GCP Console Add a Partner Interconnect Connection

Use GCP step-by-step configuration workflow (**Hybrid Connectivity > Interconnect**) to create a Partner Interconnect connection with a selected partner, in the case of this blog - with Megaport as shown in the image.

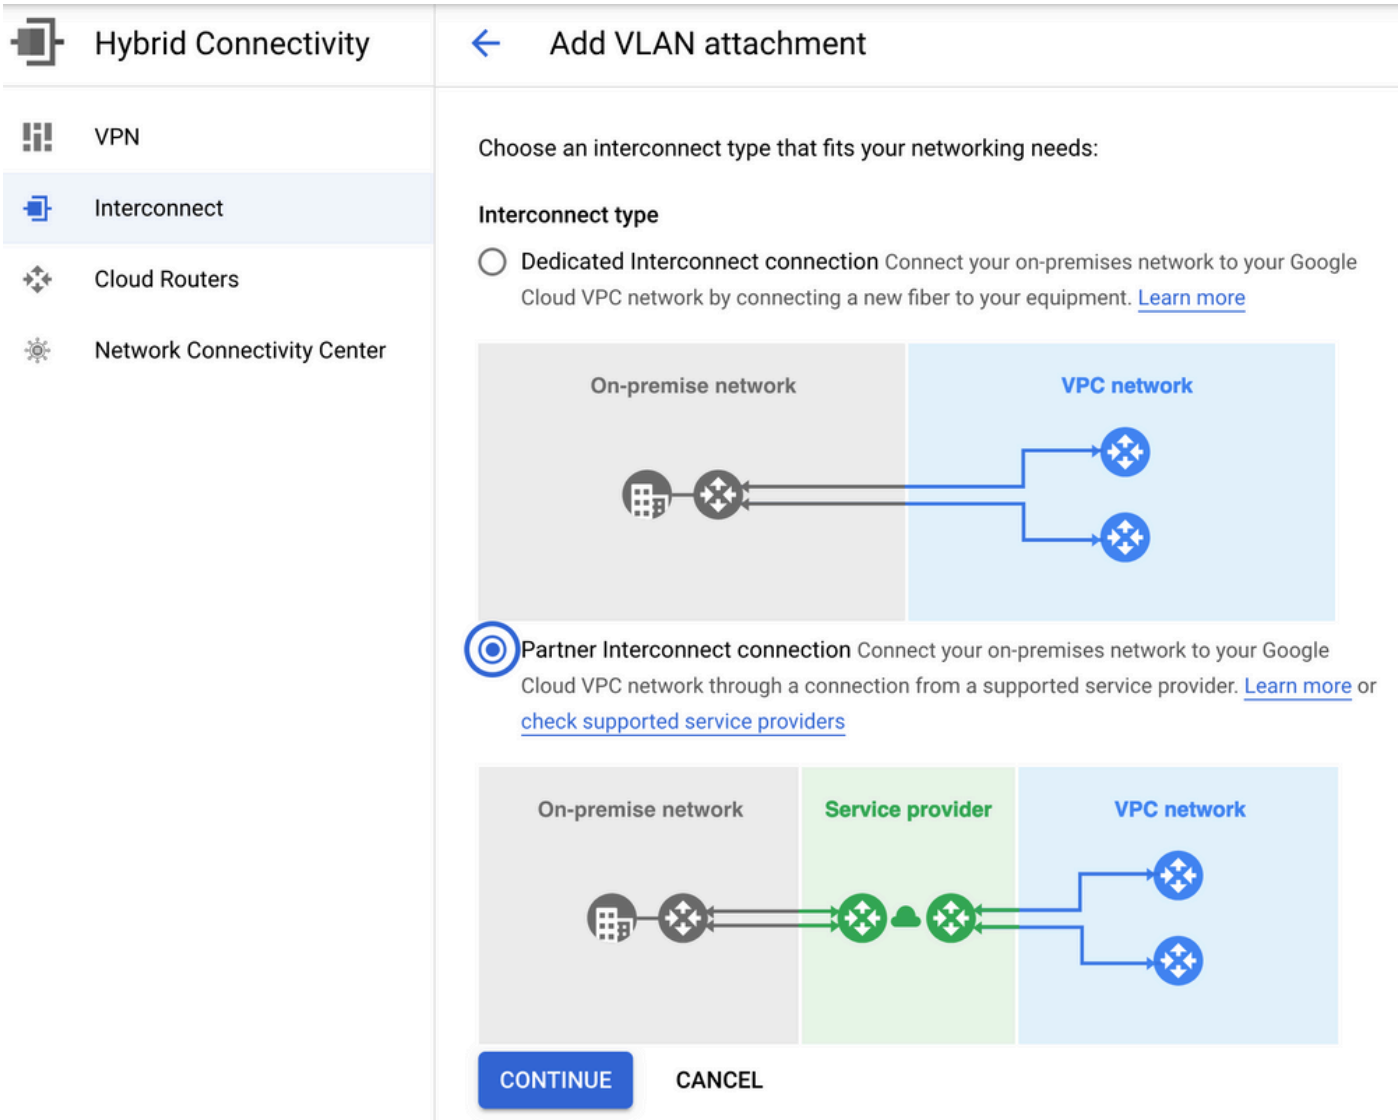

Please select the option **I ALREADY HAVE A SERVICE PROVIDER**.

For ease of demonstration, **Create a single VLAN** option is used without redundancy.

Select the correct network name, which was previously created by Cloud onRamp for Multicloud workflow. Under VLAN section, you can create a new GCR router and define a name for the VLAN, which later will be shown in the Cloud onRamp Interconnect section.

This image reflects all the points that are mentioned.

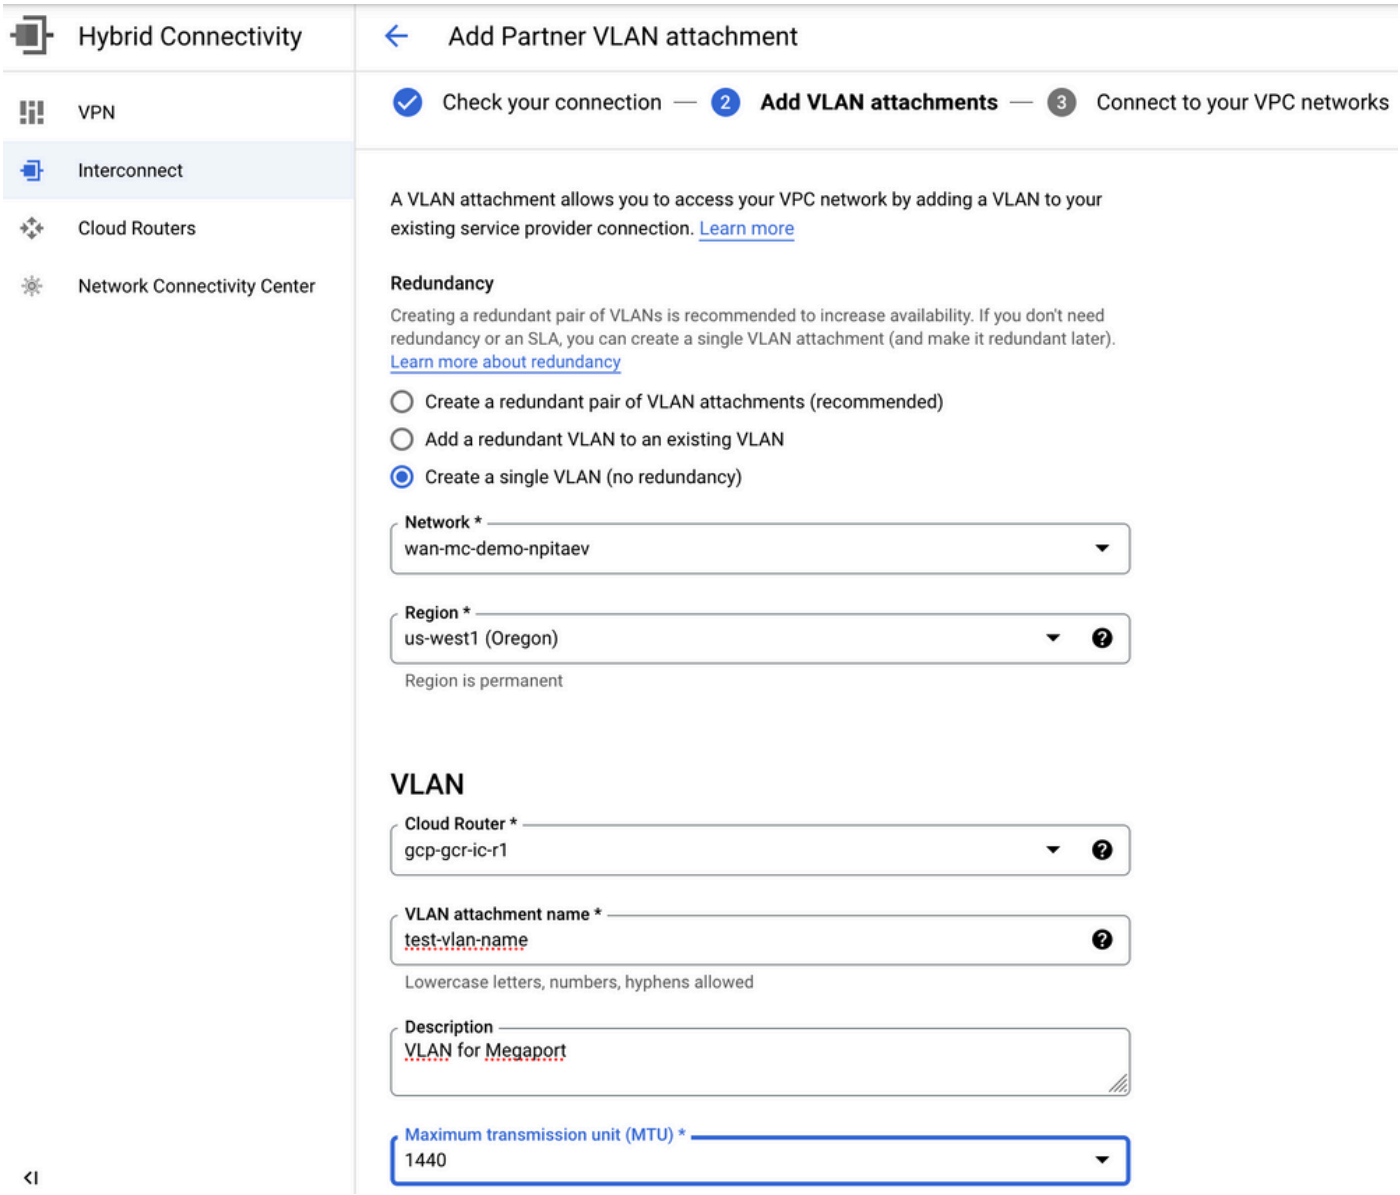

Basically, once Step 3. is completed, you can simply grab the BGP configuration and make the connectivity based on what the Interconnect provider has used. In this case, Megaport is used to test. However, you can use any sort of interconnect which can be via Megaport, Equinix, or an MSP.

Step 4. Use Cloud onRamp Interconnect in Cisco vManage to Create the DC Connection

Similar to the AWS Blog, use Cisco Cloud onRamp Interconnect workflow with Megaport to create a Data Center Router and use it for GCP Cloud Interconnect. Please note, that Megaport is used here just for testing purposes, if you already have a Data Center setup, there is no need to use Megaport.

In Cisco vManage select one free SD-WAN router, attach the default CoR Megaport template, and deploy it as Cisco Cloud Gateway in Megaport using CoR Interconnect workflow.

Once the Cisco SD-WAN router in Megaport will be active, use CoR Interconnect workflow to create a connection as shown in the image.

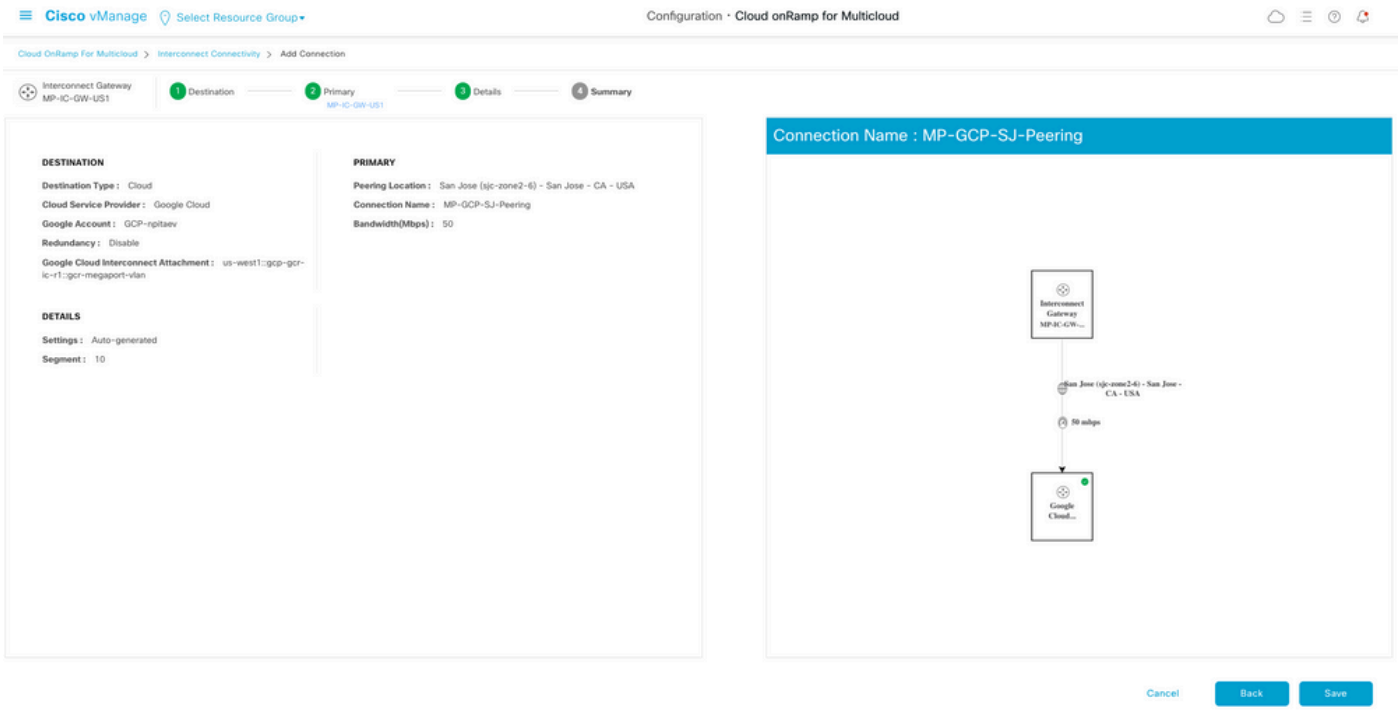

Step 5. Configure DC Router to Establish Tunnels over Internet and over GCP Cloud Interconnect

Bring SD-WAN Megaport Router into CLI mode and **move** the configuration from the service side to VPN0. Because GCP uses 169.254.x.y IP addresses, you can create Loopback1 Interface on the DC router and use it for SD-WAN communication over GCP Cloud Interconnect.

Here are the relevant parts of the DC router configuration.

```
interface Loopback1
no shutdown
ip address 192.168.9.9 255.255.255.255
!
!
interface Tunnel2
ip unnumbered Loopback1
tunnel source Loopback1
tunnel mode sdwan
!
!
interface GigabitEthernet1.215
encapsulation dot1Q 215
ip address 169.254.145.226 255.255.255.248
ip mtu 1440
!
!
router bgp 64513
bgp log-neighbor-changes
neighbor 169.254.145.225 remote-as 16550
neighbor 169.254.145.225 description MP-GCP-SJ-Peering
neighbor 169.254.145.225 ebgp-multihop 4
!
address-family ipv4
network 192.168.9.9 mask 255.255.255.255
neighbor 169.254.145.225 activate
neighbor 169.254.145.225 send-community both
exit-address-family
```

```
!
!
sdwan
interface Loopback1
tunnel-interface
encapsulation ipsec preference 100 weight 1
color private1
max-control-connections 0
allow-service all
!
```

Please refer to the full DC router configuration in the latter section of the document.

Verify

GCP Cloud Interconnect Status:

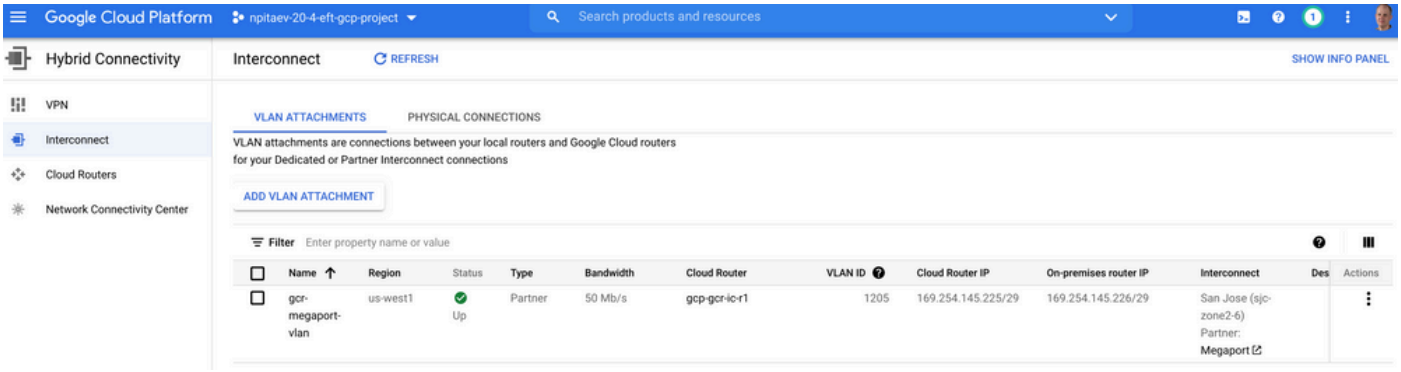

BGP connectivity between Data Center Router and WAN GCR implementing Cloud Interconnect:

```
MP-IC-US-R1#sh ip ro bgp
...
10.0.0.0/27 is subnetted, 1 subnets
B 10.35.0.0 [20/100] via 169.254.145.225, 01:25:26
MP-IC-US-R1#
```

DC Megaport SD-WAN Router Configuration

```
MP-IC-US-R1#sh sdwan bfd sessions
SOURCE TLOC REMOTE TLOC DST PUBLIC DST PUBLIC DETECT TX
SYSTEM IP SITE ID STATE COLOR COLOR SOURCE IP IP PORT ENCAP MULTIPLIER INTERVAL(msec UPTIME
TRANSITIONS
----------------------------------------------------------------------------------------------
----------------------------------------------------------------------------------------------
------------------
10.12.1.11 12 up biz-internet public-internet 162.43.150.15 13.55.49.253 12426 ipsec 7 1000 10
4:02:55:32 0
35.35.35.2 35 up biz-internet private2 162.43.150.15 35.212.162.72 12347 ipsec 7 1000 10
4:02:55:32 0
35.35.35.1 35 up biz-internet private2 162.43.150.15 35.212.232.51 12347 ipsec 7 1000 10
4:02:55:32 0
61.61.61.61 61 down biz-internet biz-internet 162.43.150.15 162.43.145.3 12427 ipsec 7 1000 NA 0
61.61.61.61 61 down biz-internet private1 162.43.150.15 198.18.0.5 12367 ipsec 7 1000 NA 0
35.35.35.1 35 up private1 private2 192.168.9.9 10.35.0.2 12347 ipsec 7 1000 10 0:00:00:16 0
35.35.35.2 35 up private1 private2 192.168.9.9 10.35.0.3 12347 ipsec 7 1000 10 0:00:00:16 0
10.12.1.11 12 down private1 public-internet 192.168.9.9 13.55.49.253 12426 ipsec 7 1000 NA 0
61.61.61.61 61 down private1 biz-internet 192.168.9.9 162.43.145.3 12427 ipsec 7 1000 NA 0
61.61.61.61 61 down private1 private1 192.168.9.9 198.18.0.5 12367 ipsec 7 1000 NA 0
```

```
MP-IC-US-R1#sh ip ro bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
H - NHRP, G - NHRP registered, g - NHRP registration summary
o - ODR, P - periodic downloaded static route, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
&- replicated local route overrides by connected

Gateway of last resort is 162.43.150.14 to network 0.0.0.0

10.0.0.0/27 is subnetted, 1 subnets
B 10.35.0.0 [20/100] via 169.254.145.225, 00:03:17
MP-IC-US-R1#
MP-IC-US-R1#sh sdwa
MP-IC-US-R1#sh sdwan runn
MP-IC-US-R1#sh sdwan running-config
system
location "55 South Market Street, San Jose, CA -95113, USA"
gps-location latitude 37.33413
gps-location longitude -121.8916
system-ip 34.34.34.1
overlay-id 1
site-id 34
port-offset 1
control-session-pps 300
admin-tech-on-failure
sp-organization-name MC-Demo-npitaev
organization-name MC-Demo-npitaev
port-hop
track-transport
track-default-gateway
console-baud-rate 19200
no on-demand enable
on-demand idle-timeout 10
vbond 54.188.241.123 port 12346
!
service tcp-keepalives-in
service tcp-keepalives-out
no service tcp-small-servers
no service udp-small-servers
hostname MP-IC-US-R1
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
vrf definition 10
rd 1:10
address-family ipv4
route-target export 64513:10
route-target import 64513:10
exit-address-family
!
address-family ipv6
exit-address-family
!
!
ip arp proxy disable
no ip finger
no ip rcmd rcp-enable
```

```
no ip rcmd rsh-enable
no ip dhcp use class
ip bootp server
no ip source-route
no ip http server
no ip http secure-server
ip nat settings central-policy
cdp run
interface GigabitEthernet1
no shutdown
arp timeout 1200
ip address dhcp client-id GigabitEthernet1
no ip redirects
ip dhcp client default-router distance 1
ip mtu 1500
load-interval 30
mtu 1500
negotiation auto
exit
interface GigabitEthernet1.215
no shutdown
encapsulation dot1Q 215
ip address 169.254.145.226 255.255.255.248
no ip redirects
ip mtu 1440
exit
interface Loopback1
no shutdown
ip address 192.168.9.9 255.255.255.255
exit
interface Tunnel1
no shutdown
ip unnumbered GigabitEthernet1
no ip redirects
ipv6 unnumbered GigabitEthernet1
no ipv6 redirects
tunnel source GigabitEthernet1
tunnel mode sdwan
exit
interface Tunnel2
no shutdown
ip unnumbered Loopback1
no ip redirects
ipv6 unnumbered Loopback1
no ipv6 redirects
tunnel source Loopback1
tunnel mode sdwan
exit
clock timezone UTC 0 0
logging persistent size 104857600 filesize 10485760
no logging monitor
logging buffered 512000
logging console
aaa authentication login default local
aaa authorization exec default local
aaa server radius dynamic-author
!
router bgp 64513
bgp log-neighbor-changes
neighbor 169.254.145.225 remote-as 16550
neighbor 169.254.145.225 description MP-GCP-SJ-Peering
neighbor 169.254.145.225 ebgp-multihop 4
address-family ipv4 unicast
neighbor 169.254.145.225 activate
```

```
neighbor 169.254.145.225 send-community both
network 192.168.9.9 mask 255.255.255.255
exit-address-family
!
timers bgp 60 180
!
snmp-server ifindex persist
line aux 0
stopbits 1
!
line con 0
speed 19200
stopbits 1
!
line vty 0 4
transport input ssh
!
line vty 5 80
transport input ssh
!
lldp run
nat64 translation timeout tcp 3600
nat64 translation timeout udp 300
sdwan
interface GigabitEthernet1
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
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 Loopback1
tunnel-interface
encapsulation ipsec preference 100 weight 1
color private1
max-control-connections 0
allow-service all
no allow-service bgp
allow-service dhcp
allow-service dns
allow-service icmp
no 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
appqoe
no tcpopt enable
no dreopt enable
!
omp
no shutdown
send-path-limit 4
ecmp-limit 4
graceful-restart
no as-dot-notation
timers
holdtime 60
advertisement-interval 1
graceful-restart-timer 43200
eor-timer 300
exit
address-family ipv4
advertise bgp
advertise connected
advertise static
!
address-family ipv6
advertise bgp
advertise connected
advertise static
!
!
!
licensing config enable false
licensing config privacy hostname false
licensing config privacy version false
licensing config utility utility-enable false
bfd color lte
hello-interval 1000
no pmtu-discovery
multiplier 1
!
bfd default-dscp 48
bfd app-route multiplier 2
bfd app-route poll-interval 123400
security
ipsec
rekey 86400
replay-window 512
!
!
sslproxy
no enable
rsa-key-modulus 2048
certificate-lifetime 730
eckey-type P256
ca-tp-label PROXY-SIGNING-CA
settings expired-certificate drop
settings untrusted-certificate drop
settings unknown-status drop
```

```
settings certificate-revocation-check none
settings unsupported-protocol-versions drop
settings unsupported-cipher-suites drop
settings failure-mode close
settings minimum-tls-ver TLSv1
dual-side optimization enable
!

MP-IC-US-R1#
MP-IC-US-R1#
MP-IC-US-R1#
MP-IC-US-R1#sh run
Building configuration...

Current configuration : 4628 bytes
!
! Last configuration change at 19:42:11 UTC Tue Jan 25 2022 by admin
!
version 17.6
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps debug datetime msec
service timestamps log datetime msec
service password-encryption
! Call-home is enabled by Smart-Licensing.
service call-home
platform qfp utilization monitor load 80
no platform punt-keepalive disable-kernel-core
platform console virtual
!
hostname MP-IC-US-R1
!
boot-start-marker
boot-end-marker
!
!
vrf definition 10
rd 1:10
!
address-family ipv4
route-target export 64513:10
route-target import 64513:10
exit-address-family
!
address-family ipv6
exit-address-family
!
vrf definition 65528
!
address-family ipv4
exit-address-family
!
logging buffered 512000
logging persistent size 104857600 filesize 10485760
no logging monitor
!
aaa new-model
!
!
aaa authentication login default local
aaa authorization exec default local
!
!
!
```

```
!
!
aaa server radius dynamic-author
!
aaa session-id common
fhrp version vrrp v3
ip arp proxy disable
!
!
!
!
!
!
!
ip bootp server
no ip dhcp use class
!
!
!
no login on-success log
ipv6 unicast-routing
!
!
!
!
!
!
!
!
subscriber templating
!
!
!
!
!
!
!
!
multilink bundle-name authenticated
!
!
!
!
!
!
!
!
!
crypto pki trustpoint TP-self-signed-1238782368
enrollment selfsigned
subject-name cn=IOS-Self-Signed-Certificate-1238782368
revocation-check none
rsakeypair TP-self-signed-1238782368
!
crypto pki trustpoint SLA-TrustPoint
enrollment pkcs12
revocation-check crl
!
!
crypto pki certificate chain TP-self-signed-1238782368
crypto pki certificate chain SLA-TrustPoint
!
!
!
!
!
```

```
!
!
!
license udi pid C8000V sn 9SRWHHH66II
license boot level network-premier+dna-premier
diagnostic bootup level minimal
memory free low-watermark processor 202832
!
!
spanning-tree extend system-id
!
username admin privilege 15 secret 9
$9$3V6L3V6L2VUI2k$ysPnXOdg8RLj9KgMdmfHdSHkdaMmiHzGaUpcqH6pfTo
!
redundancy
!
!
!
!
no crypto ikev2 diagnose error
!
!
lldp run
cdp run
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback1
ip address 192.168.9.9 255.255.255.255
!
interface Loopback65528
vrf forwarding 65528
ip address 192.168.1.1 255.255.255.255
!
interface Tunnel1
ip unnumbered GigabitEthernet1
no ip redirects
ipv6 unnumbered GigabitEthernet1
no ipv6 redirects
tunnel source GigabitEthernet1
tunnel mode sdwan
!
interface Tunnel2
ip unnumbered Loopback1
no ip redirects
ipv6 unnumbered Loopback1
no ipv6 redirects
```

```
tunnel source Loopback1
tunnel mode sdwan
!
interface GigabitEthernet1
ip dhcp client default-router distance 1
ip address dhcp client-id GigabitEthernet1
no ip redirects
load-interval 30
negotiation auto
arp timeout 1200
!
interface GigabitEthernet1.215
encapsulation dot1Q 215
ip address 169.254.145.226 255.255.255.248
no ip redirects
ip mtu 1440
arp timeout 1200
!
router omp
!
router bgp 64513
bgp log-neighbor-changes
neighbor 169.254.145.225 remote-as 16550
neighbor 169.254.145.225 description MP-GCP-SJ-Peering
neighbor 169.254.145.225 ebgp-multihop 4
!
address-family ipv4
network 192.168.9.9 mask 255.255.255.255
neighbor 169.254.145.225 activate
neighbor 169.254.145.225 send-community both
exit-address-family
!
ip forward-protocol nd
no ip http server
no ip http secure-server
!
ip nat settings central-policy
ip nat route vrf 65528 0.0.0.0 0.0.0.0 global
no ip nat service H225
no ip nat service ras
no ip nat service rtsp udp
no ip nat service rtsp tcp
no ip nat service netbios-ns tcp
no ip nat service netbios-ns udp
no ip nat service netbios-ssn
no ip nat service netbios-dgm
no ip nat service ldap
no ip nat service sunrpc udp
no ip nat service sunrpc tcp
no ip nat service msrpc tcp
no ip nat service tftp
no ip nat service rcmd
no ip nat service pptp
no ip ftp passive
ip scp server enable
!
!
!
!
!
!
!
!
control-plane
```

```
!
!
mgcp behavior rsip-range tgcp-only
mgcp behavior comedia-role none
mgcp behavior comedia-check-media-src disable
mgcp behavior comedia-sdp-force disable
!
mgcp profile default
!
!
!
!
!
!
line con 0
stopbits 1
speed 19200
line aux 0
line vty 0 4
transport input ssh
line vty 5 80
transport input ssh
!
nat64 translation timeout udp 300
nat64 translation timeout tcp 3600
call-home
! If contact email address in call-home is configured as sch-smart-licensing@cisco.com
! the email address configured in Cisco Smart License Portal will be used as contact email
address to send SCH notifications.
contact-email-addr sch-smart-licensing@cisco.com
profile "CiscoTAC-1"
active
destination transport-method http
!
!
!
!
!
!
netconf-yang
netconf-yang feature candidate-datastore
end

MP-IC-US-R1#
MP-IC-US-R1#
MP-IC-US-R1#sh ver
Cisco IOS XE Software, Version 17.06.01a
Cisco IOS Software [Bengaluru], Virtual XE Software (X86_64_LINUX_IOSD-UNIVERSALK9-M), Version
17.6.1a, RELEASE SOFTWARE (fc2)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2021 by Cisco Systems, Inc.
Compiled Sat 21-Aug-21 03:20 by mcpre


Cisco IOS-XE software, Copyright (c) 2005-2021 by cisco Systems, Inc.
All rights reserved. Certain components of Cisco IOS-XE software are
licensed under the GNU General Public License ("GPL") Version 2.0. The
software code licensed under GPL Version 2.0 is free software that comes
with ABSOLUTELY NO WARRANTY. You can redistribute and/or modify such
GPL code under the terms of GPL Version 2.0. For more details, see the
documentation or "License Notice" file accompanying the IOS-XE software,
or the applicable URL provided on the flyer accompanying the IOS-XE
software.
```

ROM: IOS-XE ROMMON

MP-IC-US-R1 uptime is 4 days, 3 hours, 2 minutes
Uptime for this control processor is 4 days, 3 hours, 3 minutes
System returned to ROM by reload
System image file is "bootflash:packages.conf"
Last reload reason: factory-reset

Technology Package License Information:
Controller-managed

The current throughput level is 250000 kbps

Smart Licensing Status: Registration Not Applicable/Not Applicable

cisco C8000V (VXE) processor (revision VXE) with 2028465K/3075K bytes of memory.
Processor board ID 9SRWHHH66II
Router operating mode: Controller-Managed
1 Gigabit Ethernet interface
32768K bytes of non-volatile configuration memory.
3965112K bytes of physical memory.
11526144K bytes of virtual hard disk at bootflash:.

Configuration register is 0x2102

MP-IC-US-R1#