How to Select a Particular Site to Be a Preferred Regional Internet Breakout?

Contents

Introduction

This document describes how to configure SD-WAN fabric in order to configure particular branch vEdge as preferred regional Internet breakout with the help of Direct Internet Access (DIA) and centralized data policy. This solution could be useful in case, for example, when a regional site uses some centralized service like Zscaler® and should be used as a preferred Internet exit point. Such deployment requires Generic Routing Encapsulation (GRE) or Internet Protocol Security (IPSec) tunnels to be configured from a transport VPN and data flow is different from the regular DIA solution, where traffic reaches Internet directly.

Prerequisites

Requirements

Cisco recommends that you have knowledge of this topic:

- Basic understanding of SD-WAN Policy Framework.

Components Used

The information in this document is based on these software and hardware versions:

- vEdge Routers
- vSmart Controller with 18.3.5 software version.

Background Information

Service VPN traffic from vEdge2, that should reach the Internet, is forwarded to another branch vEdge1, using data plane tunnels. vEdge1 is the router where DIA configured for local Internet breakout.

Network Diagram

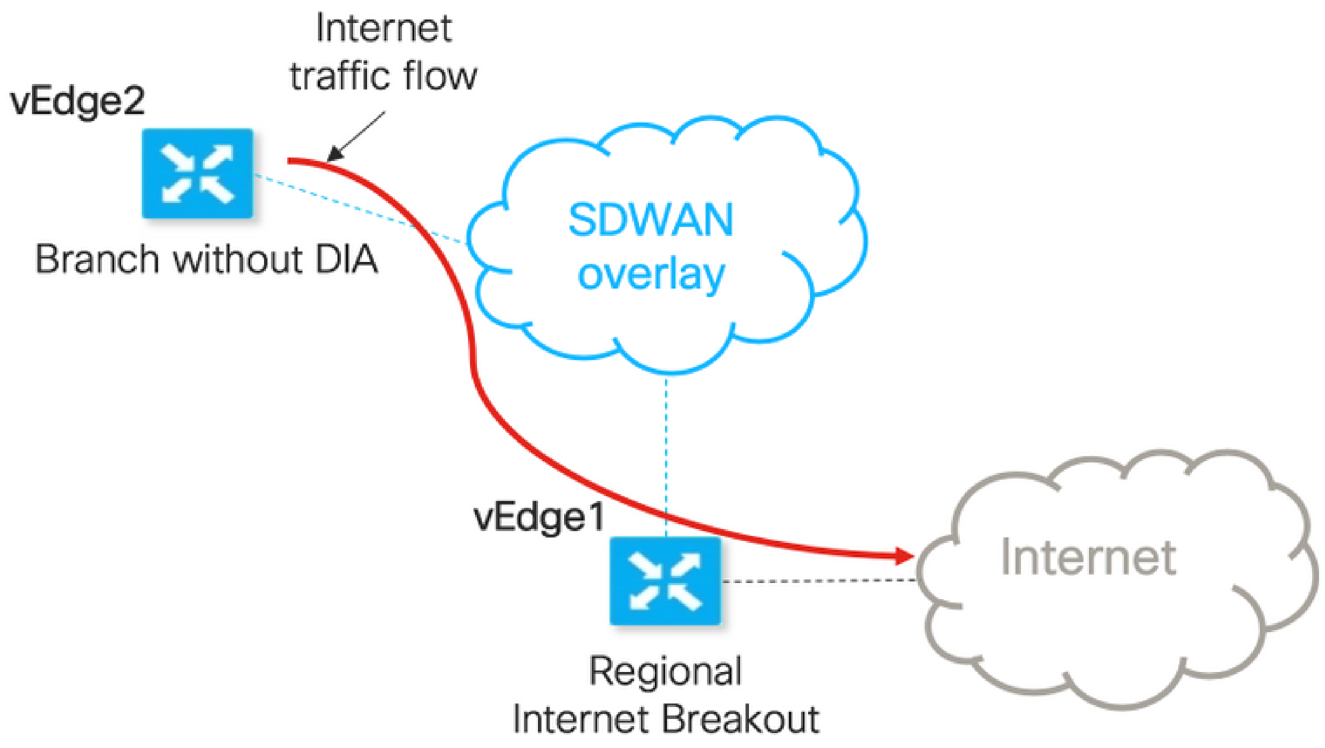

Host name	vEdge1	vEdge2
Host role	Branch device that has DIA (regional Internet breakout)	Branch device that has no DIA configured
VPN 0		
Transport Locations (TLOC) 1	biz-internet, ip: 192.168.110.6/24	biz-internet, ip: 192.168.110.5/24
Transport Locations (TLOC) 2	public-internet, ip: 192.168.109.4/24	public-internet, ip: 192.168.109.5
Service VPN 40	Interface ge0/1, ip: 192.168.40.4/24	Interface ge0/2, ip: 192.168.50.5

Configurations

Solution 1: Centralized Data-Policy usage to Change Next-Hop.

vEdge2 has data plane tunnel established with vEdge1 and other sites (Full-mesh style connectivity)

vEdge1 has DIA configured with **ip route 0.0.0.0/0 vpn 0**.

vSmart centralized data-policy configuration:

```
policy
 data-policy DIA_vE1
```

```
  vpn-list VPN_40
   sequence 5
    match
     destination-data-prefix-list ENTERPRISE_IPs
    !
    action accept
    !
   !
   sequence 10
    action accept
     set
      next-hop 192.168.40.4
     !
    !
   !
   default-action accept
  !
 !
!
lists
  vpn-list VPN_40
   vpn 40
  !
  data-prefix-list ENTERPRISE_IPs
   ip-prefix 10.0.0.0/8
   ip-prefix 172.16.0.0/12    ip-prefix 192.168.0.0/16 ! apply-policy site-list SITE2 data-policy DIA_vE1 from-service
```

vEdge2 - doesn't require any special configuration.

Here you can find steps to perform verification if a policy was applied properly.

1. Check that policy is absent from vEdge2:

```
vedge2# show policy from-vsmart
% No entries found.
```

2. Check Forwarding Information Base (FIB) programming. It should show route absence (Blackhole) for the destination on the Internet:

```
vedge2# show policy service-path vpn 40 interface ge0/2 source-ip 192.168.50.5 dest-ip 173.37.145.84 protocol 1 all
Number of possible next hops: 1
Next Hop: Blackhole
```

3. Apply vSmart data-policy under **apply-policy** section of vSmart configuration or activate in vManage GUI.

4. Check that vEdge2 successfully received data-policy from vSmart:

```
vedge2# show policy from-vsmart
from-vsmart data-policy DIA_vE1
 direction from-service
 vpn-list VPN_40
  sequence 5
   match
    destination-data-prefix-list ENTERPRISE_IPs
   action accept
  sequence 10
   action accept
```

```
    set
     next-hop 192.168.40.4
  default-action accept
from-vsmart lists vpn-list VPN_40
 vpn 40
from-vsmart lists data-prefix-list ENTERPRISE_IPs
 ip-prefix 10.0.0.0/8
 ip-prefix 172.16.0.0/12
 ip-prefix 192.168.0.0/16
```

5. Check Forwarding Information Base (FIB) programming, that shows possible routes for the destination on the Internet:

```
vedge2# show policy service-path vpn 40 interface ge0/2 source-ip 192.168.50.5 dest-ip
173.37.145.84 protocol 1 all
Number of possible next hops: 4
Next Hop: IPsec
  Source: 192.168.110.5 12366 Destination: 192.168.110.6 12346 Color: biz-internet
Next Hop: IPsec
  Source: 192.168.109.5 12366 Destination: 192.168.110.6 12346 Color: public-internet
Next Hop: IPsec
  Source: 192.168.110.5 12366 Destination: 192.168.109.4 12346 Color: biz-internet
Next Hop: IPsec
  Source: 192.168.109.5 12366 Destination: 192.168.109.4 12346 Color: public-internet
```

6. Confirm reachability to the destination on the Internet:

```
vedge2# ping vpn 40 173.37.145.84
Ping in VPN 40
PING 173.37.145.84 (173.37.145.84) 56(84) bytes of data.
64 bytes from 173.37.145.84: icmp_seq=1 ttl=63 time=0.392 ms
64 bytes from 173.37.145.84: icmp_seq=3 ttl=63 time=0.346 ms
^C
--- 173.37.145.84 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 2000ms
rtt min/avg/max/mdev = 0.345/0.361/0.392/0.021 ms
```

Here you can find vEdge1 configuration steps.

1. Activate Network Address Translation (NAT) on the transport interface, where DIA should be used:

```
vpn 0
 !
 interface ge0/0
  description "DIA interface"
  ip address 192.168.109.4/24
 nat <<<<==== NAT activated for a local DIA !
```

2. Add static route **ip route 0.0.0.0/0 vpn 0** in a Service VPN to activate DIA:

```
vpn 40
 interface ge0/4
  ip address 192.168.40.4/24
  no shutdown
 !
 ip route 0.0.0.0/0 vpn 0     <<<<==== Static route for DIA !
```

3. Check if RIB contains NAT route:

```
vedge1# show ip route vpn 40 | include nat
40 0.0.0.0/0 nat - ge0/0 - 0 - - - F,S
```

4. Confirm that DIA works and we can see Internet Control Message Protocol (ICMP) session to 173.37.145.84 from vEdge2 in NAT translations

```
vedge1# show ip nat filter | tab

                            PRIVATE                       PRIVATE  PRIVATE
PUBLIC  PUBLIC
NAT  NAT                    SOURCE         PRIVATE DEST   SOURCE   DEST     PUBLIC SOURCE
PUBLIC DEST    SOURCE  DEST    FILTER       IDLE        OUTBOUND  OUTBOUND   INBOUND  INBOUND
VPN  IFNAME  VPN  PROTOCOL  ADDRESS        ADDRESS        PORT     PORT     ADDRESS
ADDRESS        PORT    PORT    STATE        TIMEOUT     PACKETS   OCTETS     PACKETS  OCTETS
DIRECTION
------------------------------------------------------------------------------------------------
------------------------------------------------------------------------------------------------
------
0 ge0/0 40 icmp 192.168.50.5 173.37.145.84 9269 9269 192.168.109.4 173.37.145.84 9269 9269
established 0:00:00:02 10 840 10 980 -
```

Note: This solution doesn't allow us to organize redundancy or load sharing with different regional exits usage.

Doesn't work with IOS-XE routers

Solution 2: Inject Required GRE\IPSec\NAT Default Route to OMP.

As of now, there is no possibility to get the default route, pointing to GRE\IPSec tunnel on vEdge1, to be advertised through OMP to vEdge2 (redistribute nat route OMP protocol). Please note that behavior may change in future software versions.

Our goal is to create a regular static default route (**IP route 0.0.0.0/0 <next-hop IP addr>**) that could be originated by vEdge2 (device preferred for DIA) and further propagated via OMP.

To achieve this, dummy VPN is created on vEdge1 and a physical port loop is performed with cable. Loop is created between ports assigned to dummy VPN and port in the desired VPN which requires static default route. Also, you can create a loop with just one physical interface that is attached to the switch with dummy VLAN and two sub-interfaces assigned to corresponding VPNs s on the picture below:

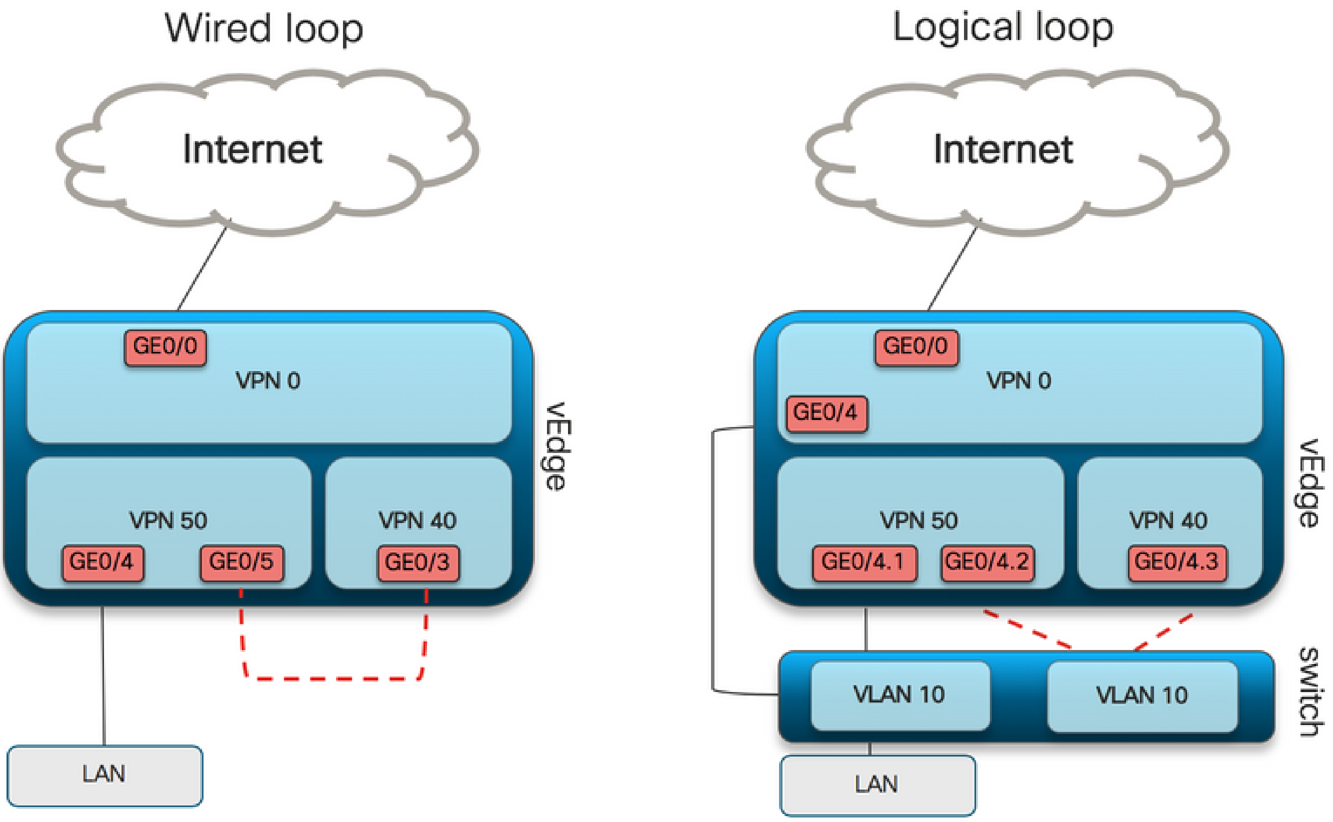

Here you can find vEdge1 configuration example.

1. Create a dummy VPN:

```
vpn 50
 interface ge0/3
description DIA_for_region ip address 192.168.111.2/30 no shutdown ! ip route 0.0.0.0/0 vpn 0
<<<<==== NAT activated for a local DIA
 ip route 10.0.0.0/8 192.168.111.1 <<<<==== Reverse routes, pointing to loop interface GE0/3
ip route 172.16.0.0/12 192.168.111.1
ip route 192.168.0.0/16 192.168.111.1 !
```

2. Check FIB that DIA route, pointing to the NAT interface, was successfully added to the routing table:

```
vedge1# show ip route vpn 50 | i nat
50      0.0.0.0/0 nat - ge0/0 - 0 - - - F,S
```

3. Service VPN used for production purposes, where regular default route is configured (which OMP will be able to advertise):

```
vpn 40
 interface ge0/4
  description CORPORATE_LAN
  ip address 192.168.40.4/24
  no shutdown
 !
 interface ge0/5
description LOOP_for_DIA ip address 192.168.111.1/30 no shutdown ! ip route 0.0.0.0/0
192.168.111.2 <<<<==== Default route, pointing to loop interface GE0/5 omp advertise connected
advertise static ! !
```

4. Check the RIB for the presence of default route pointing to the loop interface:

```
vedge1# show ip route vpn 40 | include 0.0.0.0
40 0.0.0.0/0 static - ge0/5 192.168.111.2 - - - - F,S
```

5. Check that vEdge1 advertised default route via OMP:

```
vedge1# show omp routes detail | exclude not\ set

---------------------------------------------------
omp route entries for vpn 40 route 0.0.0.0/0 <<<<==== Default route OMP entry ------------------
--------------------------------- RECEIVED FROM: peer 0.0.0.0 <<<<==== OMP route is locally
originated path-id 37 label 1002 status C,Red,R Attributes: originator 192.168.30.4 type
installed tloc 192.168.30.4, public-internet, ipsec overlay-id 1 site-id 13 origin-proto static
origin-metric 0 ADVERTISED TO: peer 192.168.30.3 Attributes: originator 192.168.30.4 label 1002
path-id 37 tloc 192.168.30.4, public-internet, ipsec site-id 13 overlay-id 1 origin-proto static
origin-metric 0
```

6. vEdge2 doesn't require any configuration, the default route is received via OMP, which points to vEdge1

```
vedge2# show ip route vpn 40 | include 0.0.0.0
40 0.0.0.0/0 omp - - - - 192.168.30.4 public-internet ipsec F,S
```

7. Confirm reachability to 173.37.145.84:

```
vedge2# ping vpn 40 173.37.145.84
Ping in VPN 40
PING 173.37.145.84 (173.37.145.84) 56(84) bytes of data.
64 bytes from 173.37.145.84: icmp_seq=2 ttl=62 time=0.518 ms
64 bytes from 173.37.145.84: icmp_seq=5 ttl=62 time=0.604 ms
^C
--- 192.168.109.5 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 3999ms
rtt min/avg/max/mdev = 0.518/0.563/0.604/0.032 ms
```

> **Note:** This solution allows you to organize redundancy or load sharing with different regional exits usage.
>
> Doesn't work with IOS-XE routers

Solution 3: Inject Default Route to OMP when Centralized Data-Policy used for DIA.

When centralized data-policy is used for local DIA, the possible way to inject the default route, it points to a regional device with DIA which is the usage of this static default route: **ip route 0.0.0.0/0 Null0**.

Due to internal packet flow, traffic that arrives from branches reach DIA thanks to data-policy, and never reach route to Null0. As you can see here, next-hop lookup happens only after a policy deployment.

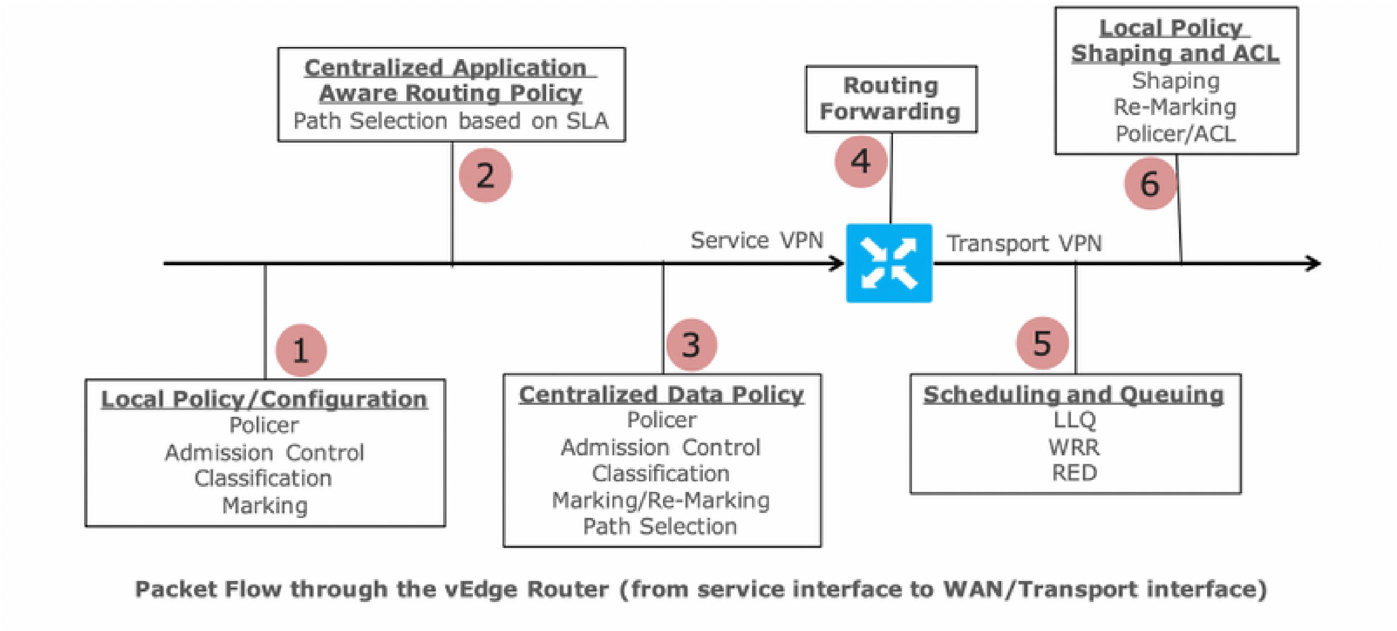

Packet Flow through the vEdge Router (from service interface to WAN/Transport interface)

vEdge2 has data plane tunnel established with vEdge1 and other sites (Full-mesh style connectivity). It doesn't require any special configuration.

vEdge1 has DIA configured with centralized data-policy.

Here you can find vEdge1 configuration steps.

1. Activate Network Address Translation (NAT) on the transport interface, where DIA should be used:

```
vpn 0
 !
 interface ge0/0
  description "DIA interface"
  ip address 192.168.109.4/24
 nat <<<<==== NAT activated for a local DIA !
```

2. Add static route **ip route 0.0.0.0/0 null0** in a Service VPN to advertise default to branches:

```
vpn 40
 interface ge0/4
  ip address 192.168.40.4/24
  no shutdown
 !
 ip route 0.0.0.0/0 null0     <<<<==== Static route to null0 that will be advertised to branches
via OMP !
```

3. Check if RIB contains default route:

```
vedge1# show ip route vpn 40 | include 0.0.0.0
40 0.0.0.0/0 static - - - 0 - - - B,F,S
```

4. Check that vEdge1 advertised default route via OMP:

```
vedge1# show omp routes detail | exclude not\ set

---------------------------------------------------
```

```
omp route entries for vpn 40 route 0.0.0.0/0 <<<<==== Default route OMP entry ------------------
--------------------------------- RECEIVED FROM: peer 0.0.0.0 <<<<==== OMP route is locally
originated path-id 37 label 1002 status C,Red,R Attributes: originator 192.168.30.4 type
installed tloc 192.168.30.4, public-internet, ipsec overlay-id 1 site-id 13 origin-proto static
origin-metric 0 ADVERTISED TO: peer 192.168.30.3 Attributes: originator 192.168.30.4 label 1002
path-id 37 tloc 192.168.30.4, public-internet, ipsec site-id 13 overlay-id 1 origin-proto static
origin-metric 0
```

5. Check that policy is absent on vEdge1 and that DIA isn't enabled:

```
vedge1# show policy from-vsmart
% No entries found.
```

6. Check Forwarding Information Base (FIB) programming. It should show route absence (Blackhole) for the destination on the Internet as DIA isn't enabled:

```
vedge1# show policy service-path vpn 40 interface ge0/2 source-ip 192.168.40.4 dest-ip
173.37.145.84 protocol 1 all
Number of possible next hops: 1
Next Hop: Blackhole
```

vSmart centralized data-policy configuration for DIA:

```
policy
 data-policy DIA_vE1
  vpn-list VPN_40
   sequence 5
    match
     destination-data-prefix-list ENTERPRISE_IPs
    action accept
   sequence 10
    action accept
     nat-use vpn0             <<<<==== NAT reference for a DIA default-action accept lists
vpn-list VPN_40    vpn 40   data-prefix-list ENTERPRISE_IPs    ip-prefix 10.0.0.0/8    ip-prefix
172.16.0.0/12    ip-prefix 192.168.0.0/16
site-list SITE1
site-id 1001 apply-policy site-list SITE1 <<<<==== policy applied to vEdge1 data-policy DIA_vE1
from-service
```

Apply vSmart data-policy under **apply-policy** section of vSmart configuration or activate in vManage GUI.

7. Check that vEdge1 successfully received data-policy from vSmart:

```
vedge1# show policy from-vsmart
from-vsmart data-policy DIA_vE1
 direction from-service
 vpn-list VPN_40
  sequence 5
   match
    destination-data-prefix-list ENTERPRISE_IPs
   action accept
  sequence 10
   action accept
nat-use vpn0 default-action accept from-vsmart lists vpn-list VPN_40 vpn 40 from-vsmart lists
data-prefix-list ENTERPRISE_IPs ip-prefix 10.0.0.0/8 ip-prefix 172.16.0.0/12    ip-prefix
192.168.0.0/16
```

8. Check Forwarding Information Base (FIB) programming, that shows possible routes for the destination on the Internet:

```
vedge1# show policy service-path vpn 40 interface ge0/2 source-ip 192.168.40.4 dest-ip
173.37.145.84 protocol 1 all
Number of possible next hops: 1
Next Hop: Remote
Remote IP:173.37.145.84, Interface ge0/0 Index: 4
```

9. Confirm reachability to the destination on the Internet:

```
vedge1# ping vpn 40 173.37.145.84
Ping in VPN 40
PING 173.37.145.84 (173.37.145.84) 56(84) bytes of data.
64 bytes from 173.37.145.84: icmp_seq=1 ttl=63 time=0.192 ms
64 bytes from 173.37.145.84: icmp_seq=3 ttl=63 time=0.246 ms
64 bytes from 173.37.145.84: icmp_seq=3 ttl=63 time=0.236 ms ^C --- 173.37.145.84 ping
statistics --- 3 packets transmitted, 3 received, 0% packet loss, time 2000ms rtt
min/avg/max/mdev = 0.245/0.221/0.192/0.021 ms
```

vEdge2 verification steps:

1. Confirm that default route was successfully received and installed in RIB:

```
vEdge2# sh ip route vpn 40 | include 0.0.0.0
40     0.0.0.0/0          omp              -         -           -                -
192.168.30.4   biz-internet     ipsec  F,S
40 0.0.0.0/0 omp - - - - 192.168.30.4 public-internet ipsec F,S
```

2. Check Forwarding Information Base (FIB) programming, that shows possible routes for the destination on the Internet:

```
vedge2# show policy service-path vpn 40 interface ge0/2 source-ip 192.168.50.5 dest-ip
173.37.145.84 protocol 1 all
Number of possible next hops: 4
Next Hop: IPsec
  Source: 192.168.110.5 12366 Destination: 192.168.110.6 12346 Color: biz-internet
Next Hop: IPsec
  Source: 192.168.109.5 12366 Destination: 192.168.110.6 12346 Color: public-internet
Next Hop: IPsec
  Source: 192.168.110.5 12366 Destination: 192.168.109.4 12346 Color: biz-internet
Next Hop: IPsec
  Source: 192.168.109.5 12366 Destination: 192.168.109.4 12346 Color: public-internet
```

3. Confirm reachability to the destination on the Internet:

```
vedge2# ping vpn 40 173.37.145.84
Ping in VPN 40
PING 173.37.145.84 (173.37.145.84) 56(84) bytes of data.
64 bytes from 173.37.145.84: icmp_seq=1 ttl=63 time=0.382 ms
64 bytes from 173.37.145.84: icmp_seq=1 ttl=63 time=0.392 ms 64 bytes from 173.37.145.84:
icmp_seq=3 ttl=63 time=0.346 ms ^C --- 173.37.145.84 ping statistics --- 3 packets transmitted,
3 received, 0% packet loss, time 2000ms rtt min/avg/max/mdev = 0.392/0.361/0.346/0.023 ms
```

4. Confirm that DIA works and we can see Internet Control Message Protocol (ICMP) session to 173.37.145.84 from vEdge2 in NAT translations

```
vedge1# show ip nat filter | tab

                            PRIVATE                           PRIVATE  PRIVATE
PUBLIC  PUBLIC
```

```
NAT   NAT                    SOURCE          PRIVATE DEST    SOURCE    DEST      PUBLIC SOURCE
PUBLIC DEST     SOURCE  DEST     FILTER        IDLE         OUTBOUND  OUTBOUND   INBOUND  INBOUND
VPN  IFNAME  VPN  PROTOCOL  ADDRESS          ADDRESS         PORT      PORT      ADDRESS
ADDRESS         PORT     PORT     STATE         TIMEOUT      PACKETS   OCTETS      PACKETS  OCTETS
DIRECTION
-----------------------------------------------------------------------------------------------
-----------------------------------------------------------------------------------------------
------
0 ge0/0 40 icmp 192.168.50.5 173.37.145.84 9175 9175 192.168.109.4 173.37.145.84 9175 9175
established 0:00:00:04 18 1440 18 1580 -
```

Note: This solution allows to organize redundancy or load sharing with different regional exits usage.

Doesn't work with IOS-XE routers

Solution 4: Inject Default Route to OMP when Local DIA Used.

This solution can be used for both IOS-XE and Viptela OS based SD-WAN routers.

In brief, in this solution, a default route for DIA (0.0.0.0/0 Null0) is split into two subnetworks 0.0.0.0/1 and 128.0.0.0/1 pointing to Null0. This step is done to avoid overlapping of a default route that should be advertised to branches and default route, used for local DIA. In IOS-XE routes used for DIA have Administrative Distance (AD) equal to 6, while AD of static default is 1. The benefit of the solution is the ability to use redundancy schema when Regional DIA is configured in two different locations.

1. Activate NAT on a transport interface

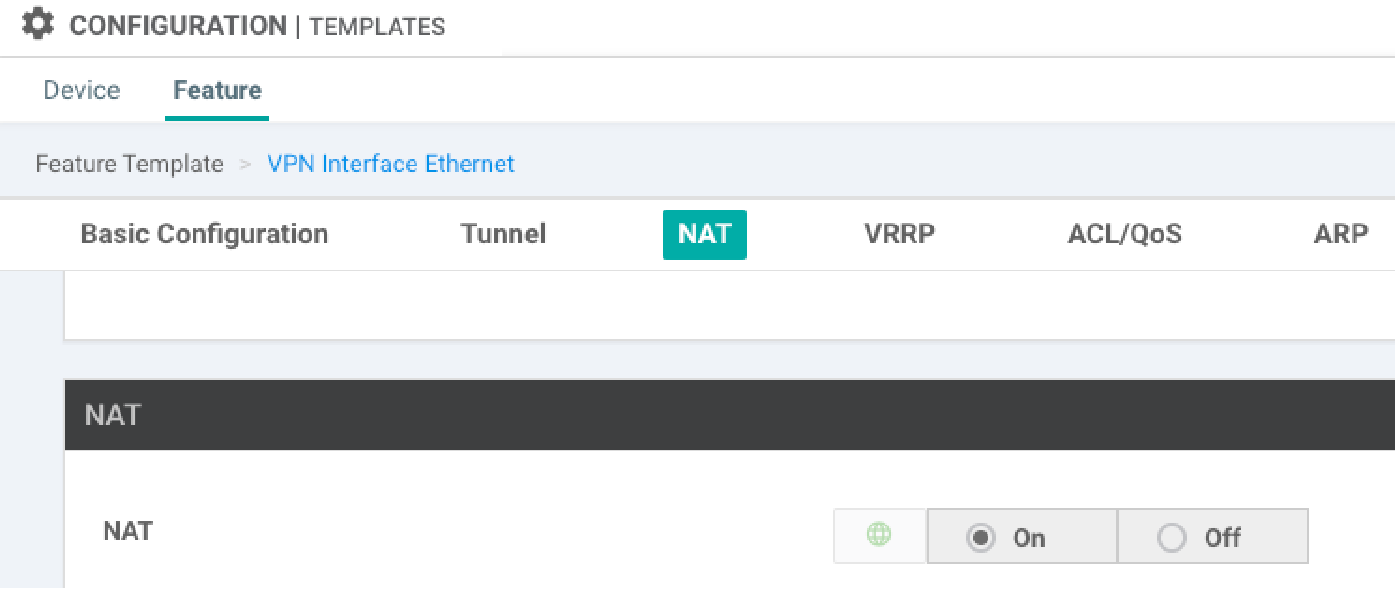
CONFIGURATION | TEMPLATES
Device
Feature
Feature Template > VPN Interface Ethernet
Basic Configuration
Tunnel
NAT
VRRP
ACL/QoS
ARP
NAT

NAT ● On ○ Off

2. In a feature template for a service VPN, where DIA should be used add the following static IPv4 routes:

- 0.0.0.0/1 and 128.0.0.0/1 pointing to VPN. These routes are used for DIA

- 0.0.0.0/0 pointing to Null 0. This route is used for advertising via OMP to branches (similar as in Solution 3)

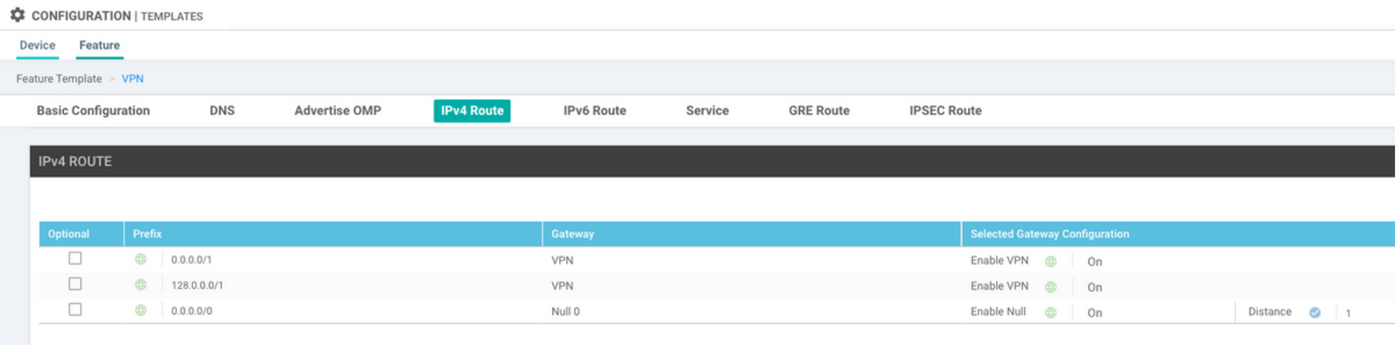

Optional	Prefix	Gateway	Selected Gateway Configuration			
□	0.0.0.0/1	VPN	Enable VPN	On		
□	128.0.0.0/1	VPN	Enable VPN	On		
□	0.0.0.0/0	Null 0	Enable Null	On	Distance	1

3. Check that routes were successfully added to RIB :

```
cedge1#show ip route vrf 40

Routing Table: 40
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, o - ODR, P -
periodic downloaded static route, H - NHRP, l - LISP
       a - application route, + - replicated route, % - next hop override, p - overrides from
PfR

Gateway of last resort is 0.0.0.0 to network 0.0.0.0

S*     0.0.0.0/0 is directly connected, Null0                     <<<<==== Static route to null0
that will be advertised to branches via OMP n Nd 0.0.0.0/1 [6/0], 00:08:23, Null0 <<<<==== DIA
route n Nd 128.0.0.0/1 [6/0], 00:08:23, Null0 <<<<==== DIA route 192.40.1.0/32 is subnetted, 1
subnets m 192.40.1.1 [251/0] via 192.168.30.207, 3d01h 192.40.2.0/32 is subnetted, 1 subnets m
192.40.2.1 [251/0] via 192.168.30.208, 3d01h
```

4. Check that DIA works well locally:

```
cedge1#ping vrf 40 173.37.145.84
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 173.37.145.84, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```

5. Check that default route successfully advertised to a branch and installed in RIB

```
cedge3#show ip route vrf 40

Routing Table: 40
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, o - ODR, P -
periodic downloaded static route, H - NHRP, l - LISP
       a - application route, + - replicated route, % - next hop override, p - overrides from
PfR

Gateway of last resort is 192.168.30.204 to network 0.0.0.0
```

```
m*    0.0.0.0/0 [251/0] via 192.168.30.204, 00:02:45    <<<<==== Default route that advertised
via OMP 192.40.1.0/32 is subnetted, 1 subnets m 192.40.11.1 [251/0] via 192.168.30.204, 00:02:45
192.40.13.0/32 is subnetted, 1 subnets C 192.40.13.1 is directly connected, Loopback40
```

6. Check that DIA works well locally:

```
cedge3#ping vrf 40 173.37.145.84
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 173.37.145.84, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/10 ms
```

7. Check on Regional DIA router successful NAT translation.

```
cedge1#show ip nat translations
Pro  Inside global         Inside local          Outside local         Outside global
icmp 192.168.109.204:1     192.40.13.1:1         173.37.145.84:1       173.37.145.84:1
Total number of translations: 1
```

Note: This solution allows to organize redundancy or load sharing with different regional exits usage.

Note: CSCvr72329 - enhancement request "NAT route redistribution to OMP"

Related Information

- **Centralized Data Policy**
- **Configuring Centralized Data Policy**
- **Centralized Data Policy Configuration Examples**
- **OMP Routing Protocol**
- **Configuring OMP**