Configure Static NAT for TLOC Extension for Interoperability with Symmetric NAT

Contents

Introduction

This document describes configuring static NAT on a TLOC Extension Router using NAT Overload to work with peers behind Symmetric NAT.

Recommendations

Cisco recommends that you have knowledge of these topics:

- Cisco Catalyst Software-Defined Wide Area Network (SD-WAN)
- Network Address Translation (NAT)
- TLOC Extension

Components Used

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

- C8000V version 17.15.1a

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Problem

The Cisco Catalyst SD-WAN Design Guide highlights that certain types of Network Address Translation (NAT) can impact the formation of Control Connections and BFD Tunnels.

The two types of NAT which do not work together are **Port/Address Restricted NAT and Symmetric NAT.** These NAT types require that sessions be initiated from the internal network to allow traffic on each port. This means external traffic cannot initiate a connection to the internal network without a prior request from the inside.

Sites behind a symmetric NAT frequently experience difficulties establishing BFD sessions with peer sites. This is particularly challenging when peering with a site using TLOC Extension behind NAT Overload (also known as Port/Address Restricted NAT).

Topology

Conditions

1. S30_Edge1 is behind a Symmetric NAT

2. S20_Edge2 is behind the TLOC Extension where S20_Edge1 is using NAT Overload (PAT) to NAT the flows from Edge2.

This results in the BFD hellos getting dropped on the Symmetric NAT device and the S20_Edge1 due to no session is present for the unknown port from the peer.

The S20_Edge1 device shows Implicit ACL Drop for these hellos due to they do not match any session in the NAT table.

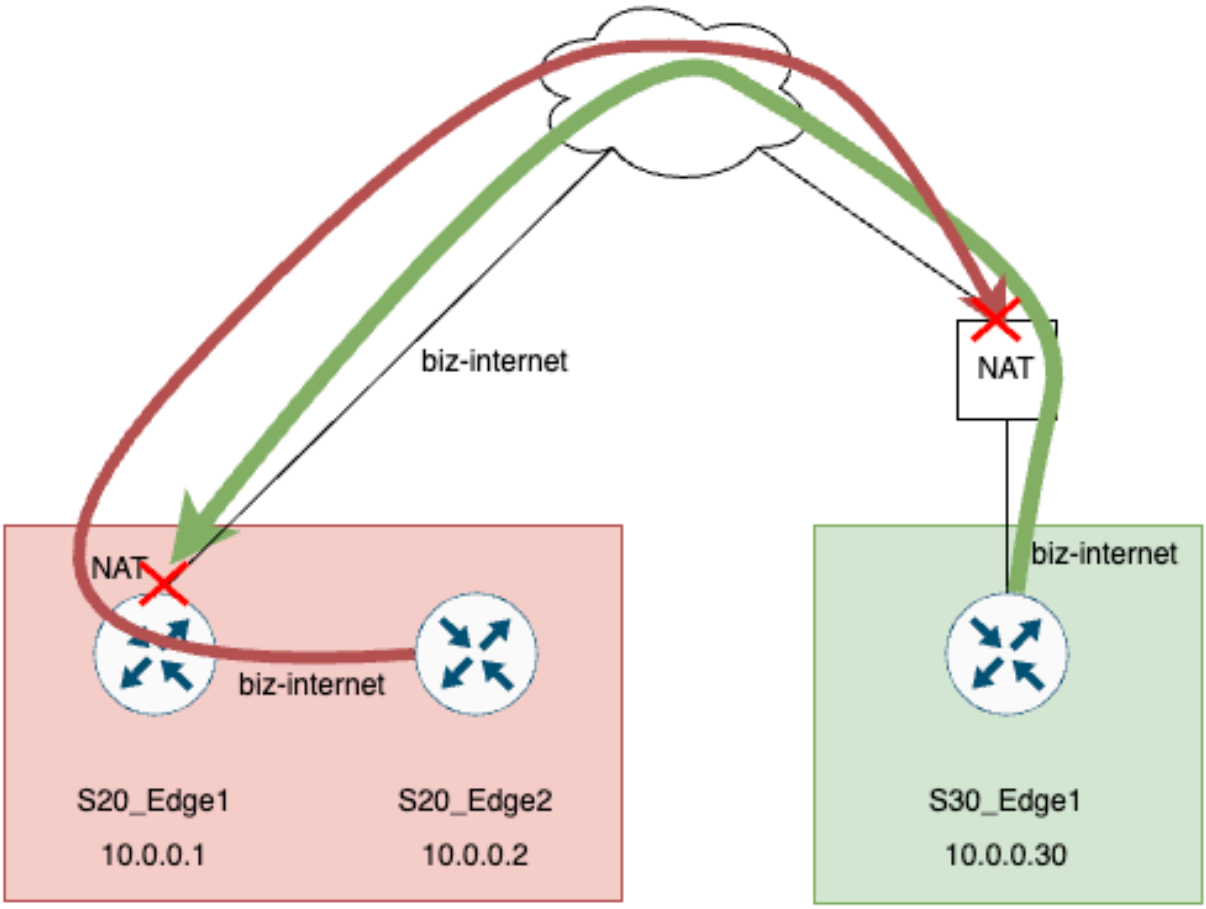

Identify the Issue

Step 1. Check BFD Sessions

From the show **sdwan bfd sessions output** on S30_Edge1, it is seen that the BFD session to S20_Edge2, 10.0.0.2 is down.

```
S30_Edge1#show sdwan bfd sessions
                                       SOURCE TLOC      REMOTE TLOC
SYSTEM IP        SITE ID     STATE        COLOR            COLOR            SOURCE IP
-----------------------------------------------------------------------------------------------------------
10.0.0.2         20          down         biz-internet     biz-internet     192.168.30.2
10.0.0.1         20          up           biz-internet     biz-internet     192.168.30.2
```

Step 2. Check the NAT Type

At the bottom of the output, the NAT Type A is seen on S30_Edge1. This indicates Symmetric NAT. Also note the public IP 172.16.1.34 and port 31048.

```
S30_Edge1# show sdwan control local-properties
```

```
<SNIP>
site-id                         30
domain-id                       1
protocol                        dtls
tls-port                        0
system-ip                       10.0.0.30
<SNIP>

 NAT TYPE: E -- indicates End-point independent mapping
           A -- indicates Address-port dependent mapping
           N -- indicates Not learned
           Note: Requires minimum two vbonds to learn the NAT type

                          PUBLIC          PUBLIC PRIVATE         PRIVATE
INTERFACE                 IPv4            PORT   IPv4            IPv6

-------------------------------------------------------------------------------------------------------------
GigabitEthernet1              172.16.1.34     31048  192.168.30.2    ::
```

Step 3. Check the NAT configuration

From the topology it is known that S20_Edge2 is behind the TLOC Extension. At this point we can check for the PAT configuration on the S20_Edge1.

NAT overload configuration is already present on S20_Edge1

```
S20_Edge1#sh run int gi1
interface GigabitEthernet1
 description biz-internet
 ip dhcp client default-router distance 1
 ip address 192.168.20.2 255.255.255.0
 no ip redirects
 ip nat outside
 load-interval 30
 negotiation auto
 arp timeout 1200
end
```

```
S20_Edge1#sh run | i nat
<SNIP>
ip nat inside source list nat-dia-vpn-hop-access-list interface GigabitEthernet1 overload
```

Step 4. Check the public IP and port

Check **show sdwan control local properties** output on S20_Edge2 to see the public IP and port 172.16.1.18 and port 5063

```
S20_Edge2#show sdwan control local-properties
<SNIP>
site-id                         20
domain-id                       1
```

```
protocol                          dtls
tls-port                          0
system-ip                         10.0.0.2
<SNIP>

 NAT TYPE: E -- indicates End-point independent mapping
           A -- indicates Address-port dependent mapping
           N -- indicates Not learned
           Note: Requires minimum two vbonds to learn the NAT type

                          PUBLIC          PUBLIC PRIVATE         PRIVATE
INTERFACE                 IPv4            PORT   IPv4            IPv6

-------------------------------------------------------------------------------------------------------------
GigabitEthernet2.100          172.16.1.18     5063   192.168.100.2   ::
```

Step 5. Check the NAT translations

Now check the NAT translations on the S20_Edge1 device. There is only a NAT session to the advertised IP and port for S30_Edge1, IP 172.16.1.34 and port 31048. Considering what we know about symmetric NAT, this is not be the case. There must be at least one different port than 31048 (not a standard SD-WAN port like 12346), if not a different IP AND port combination.

```
 S20_Edge1#sh ip nat translations
Pro  Inside global         Inside local          Outside local         Outside global
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.1.69:12346     172.16.1.69:12346
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.0.102:12446    172.16.0.102:12446
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.1.50:12346     172.16.1.50:12346
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.0.202:12346    172.16.0.202:12346
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.1.82:12346     172.16.1.82:12346
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.1.34:31048     172.16.1.34:31048
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.0.201:12346    172.16.0.201:12346
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.0.101:12446    172.16.0.101:12446
udp  192.168.20.2:5063     192.168.100.2:12346   172.16.1.98:12346     172.16.1.98:12346
```

Step 6. Check FIA trace

Run a FIA trace just to check that packets are getting dropped on S20_Edge1. Keep in mind that the IP does not have to be the same as the advertised one, but in this case for simplicity, it is.

```
S20_Edge1#debug platform condition ipv4 172.16.1.34/32 both
S20_Edge1#debug platform condition start
S20_Edge1#debug platform packet packet 1024 fia
S20_Edge1#debug platform packet packet 1024 fia-trace
S20_Edge1#show platform packet summary
Pkt   Input                    Output                   State  Reason
0     Gi2.100                  Gi1                      FWD
1     internal0/0/recycle:0    Gi1                      FWD
2     Gi2.100                  Gi1                      FWD
3     internal0/0/recycle:0    Gi1                      FWD
4     Gi2.100                  Gi1                      FWD
5     internal0/0/recycle:0    Gi1                      FWD
```

```
6     Gi2.100                  Gi1                       FWD
7     internal0/0/recycle:0    Gi1                       FWD
8     Gi1                      Gi1                       DROP   479 (SdwanImplicitAclDrop)
```

Check packet 8 to see if this is the suspected packet.

```
S20_Edge1#show platform packet packet 8
Packet: 8           CBUG ID: 482
Summary
  Input     : GigabitEthernet1
  Output    : GigabitEthernet1
  State     : DROP 479 (SdwanImplicitAclDrop)
  Timestamp
    Start   : 6120860350139 ns (04/18/2025 02:35:03.873687 UTC)
    Stop    : 6120860374021 ns (04/18/2025 02:35:03.873710 UTC)
Path Trace
  Feature: IPV4(Input)
    Input       : GigabitEthernet1
    Output      : <unknown>
    Source      : 172.16.1.34
    Destination : 192.168.20.2
    Protocol    : 17 (UDP)
      SrcPort   : 3618
      DstPort   : 12346
<SNIP>
```

This does seem to be the packet from S30_Edge1.

Checking back on the NAT table in step 6, we can see there is no session for this packet. That is the reason for the drop.

Step 7. Check BFD counters

BFD packets from S20_Edge2 are not be seen at S30_Edge1 due to they are dropped outside the device, on the NAT device. The BFD Tx/Rx counters can be checked via **show sdwan tunnel statistics** command.

```
S30_Edge1#show sdwan tunnel statistics
tunnel stats ipsec 192.168.30.2 172.16.1.18 12346 12347
 system-ip            10.0.0.2
 local-color          biz-internet
 remote-color         biz-internet
 tunnel-mtu           1438
 tx_pkts              10
 tx_octets            1060
 rx_pkts              0   <<<<<<<<<<<<<
 rx_octets            0
 tcp-mss-adjust       1358
 ipv6_tx_pkts         0
 ipv6_tx_octets       0
 ipv6_rx_pkts         0
 ipv6_rx_octets       0
 tx_ipv4_mcast_pkts   0
 tx_ipv4_mcast_octets 0
```

```
rx_ipv4_mcast_pkts   0
rx_ipv4_mcast_octets 0
tx-ipv6-mcast-pkts   0
tx-ipv6-mcast-octets 0
rx-ipv6-mcast-pkts   0
rx-ipv6-mcast-octets 0
```

Solution

To solve this, a static NAT can be configured on top of the NAT Overload (PAT) on S20_Edge1 to NAT all Control and BFD Packets to a single IP/Port combination.

1. First, it is be necessary to disable port-hopping on this color, or system-wide on S20_Edge2.

A port-offset is also be added as a best practice for S20_Edge2 so S20_Edge1 and S10_Edge2 do not use the same source port for control connections or BFD tunnels.

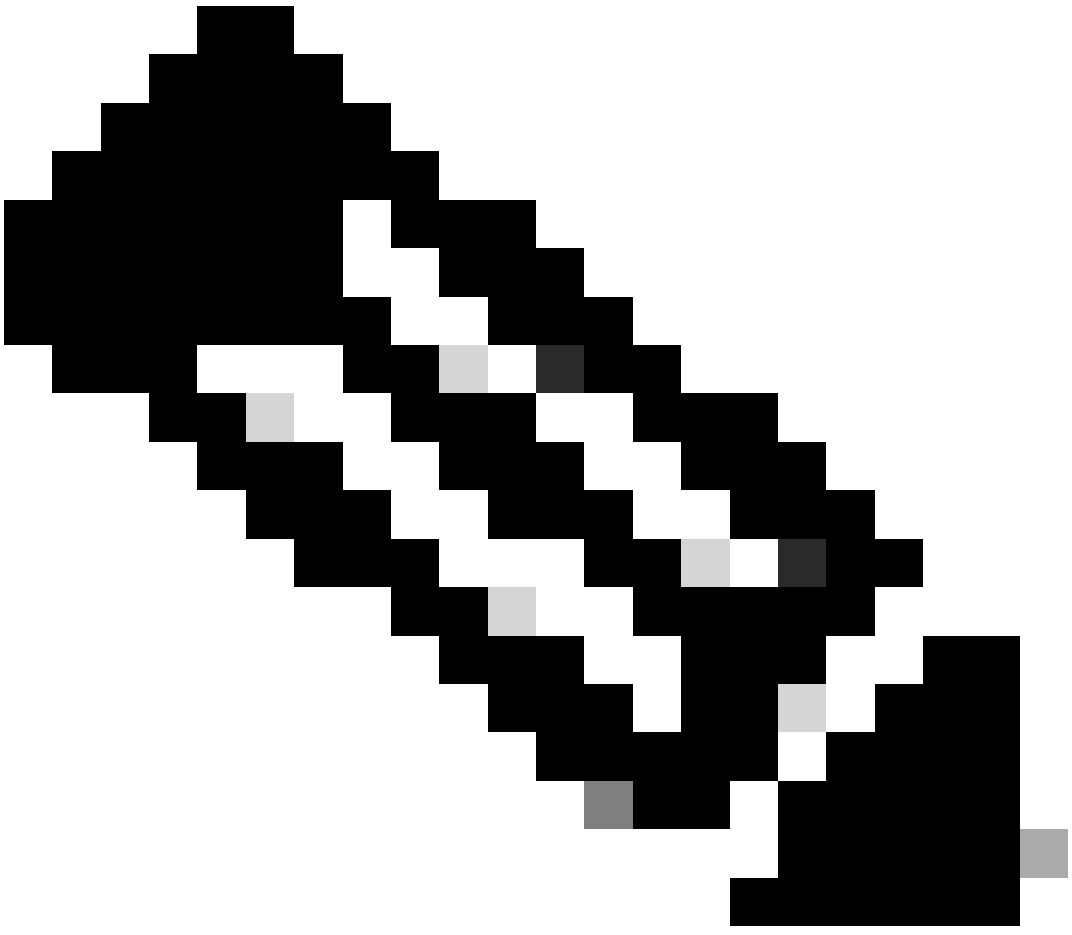

Note: This configuration can be performed through the router CLI or through a vManage CLI Add-On template.

```
S20_Edge2#config-t
S20_Edge2(config)# system
S20_Edge2(config-system)# no port-hop
S20_Edge2(config-system)# port-offset 1
S20_Edge2(config-system)# commit
```

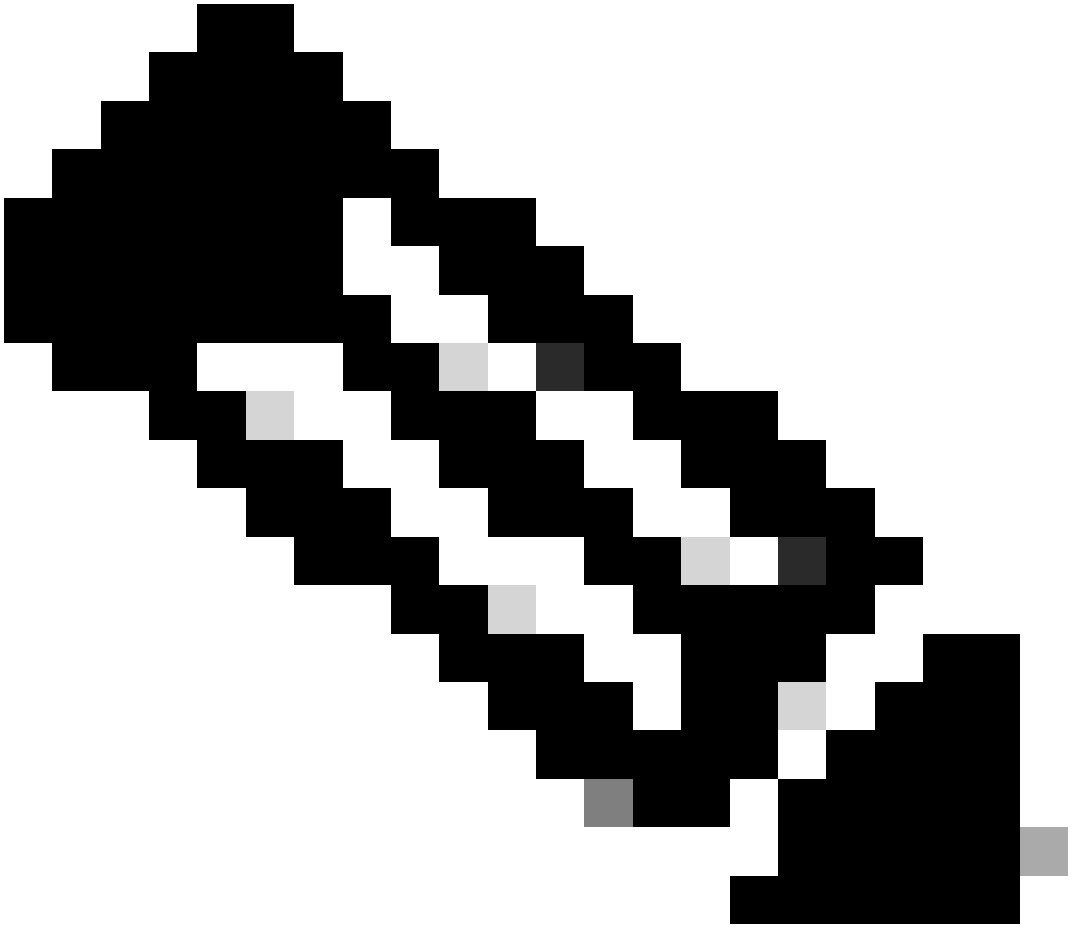

Note: Ensure that the S20_Edge2 is using the base port 12347 after this configuration by checking **show sdwan control local-properties.** If it is not using the base port, use the command **clear sdwan control port-index** to reset the port back to the base port. This prevents the port from changing if it were running on a higher port and then reboots later. The clear command resets control connections and bfd tunnels.

---

2. Configure the Static NAT on S20_Edge1.

```
S20_Edge1#config-t
S20_Edge1(config)# ip nat inside source static udp 192.168.100.2 12347 192.168.20.2 12347 egress-interf
S20_Edge1(config)# commit
```

3. Clear the NAT Translations on S20_Edge1.

```
S20_Edge1#clear ip nat translation *
```

Verification

1. Check the BFD Sessions on one of the peers.

```
S30_Edge1#show sdwan bfd sessions
                                      SOURCE TLOC      REMOTE TLOC
SYSTEM IP        SITE ID     STATE       COLOR            COLOR            SOURCE IP
-------------------------------------------------------------------------------------------------------
10.0.0.2         20          up          biz-internet     biz-internet     192.168.30.2
```

2. Check the NAT sessions on S20_Edge1.

```
S20_Edge1#sh ip nat translations
Pro  Inside global         Inside local          Outside local         Outside global
udp  192.168.20.2:12347    192.168.100.2:12347   ---                   ---
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.0.202:12346    172.16.0.202:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.1.50:12346     172.16.1.50:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.0.102:12446    172.16.0.102:12446
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.1.34:50890     172.16.1.34:50890
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.1.69:12346     172.16.1.69:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.1.98:12346     172.16.1.98:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.0.101:12446    172.16.0.101:12446
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.0.201:12346    172.16.0.201:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.1.82:12346     172.16.1.82:12346
udp  192.168.20.2:12347    192.168.100.2:12347   172.16.0.1:13046      172.16.0.1:13046
Total number of translations: 11
```

Now it is seen that all the control connections and BFD Tunnels are NAT to the configured IP and port, 192.168.20.2:12347. Also the connection to 172.16.1.34 is to a completely different port than advertised to vSmart by S30_Edge1. See port 50890.

3. Notice in the **show sdwan control local properties** output from S30_Edge1 that the advertised IP and port are 172.16.1.34 and port 60506.

```
S30_Edge1#show sdwan control local-properties
<SNIP>
site-id                         30
domain-id                       1
protocol                        dtls
```

```
tls-port                        0
system-ip                       10.0.0.30
<SNIP>

 NAT TYPE: E -- indicates End-point independent mapping
           A -- indicates Address-port dependent mapping
           N -- indicates Not learned
           Note: Requires minimum two vbonds to learn the NAT type

                         PUBLIC          PUBLIC PRIVATE         PRIVATE
INTERFACE                IPv4            PORT   IPv4            IPv6

-------------------------------------------------------------------------------------------------------------
GigabitEthernet1              172.16.1.34     60506  192.168.30.2    ::
```

References

[Cisco Catalyst SD-WAN Design Guide]