Track Tunnels Health Status when Connected to the Internet

Contents

Introduction
Background Information
Prerequisites
Requirements
Components Used
Configure
Network Diagram
Track Interface Status
Configurations
Verify
Troubleshoot

Introduction

This document describes how to track transport tunnels' health status in VPN 0. In Releases 17.2.2 and later, on Network Address Translation (NAT) enabled transport interfaces are used for local internet exit. You can track the status of the internet connection with the help of these. If the internet becomes unavailable, traffic is automatically redirected to the non-NATed tunnel on the transport interface.

Background Information

In order to provide users at a local site with direct, secure access to Internet resources, such as websites, you can configure the vEdge router to function as a NAT device, that performs both address and port translation (NAPT). When you enable NAT, it allows traffic exiting from a vEdge router to pass directly to the Internet rather than being backhauled to a co-location facility that provides NAT services for Internet access. If you use NAT in this way on a vEdge router, you can eliminate traffic "tromboning" and allow for efficient routes, that have shorter distances, between users at the local site and the network-based applications that they use.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

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.

Configure

Network Diagram

vEdge1 router here acts as a NAT device. The vEdge router splits its traffic into two flows, which you can think of as two separate tunnels. One traffic flow, shown in green, remains within the overlay network and travels between the two routers in the usual fashion, on the secure IPsec tunnels that form the overlay network. The second traffic stream, shown in grey, is redirected through the vEdge router's NAT device and then out of the overlay network to a public network.

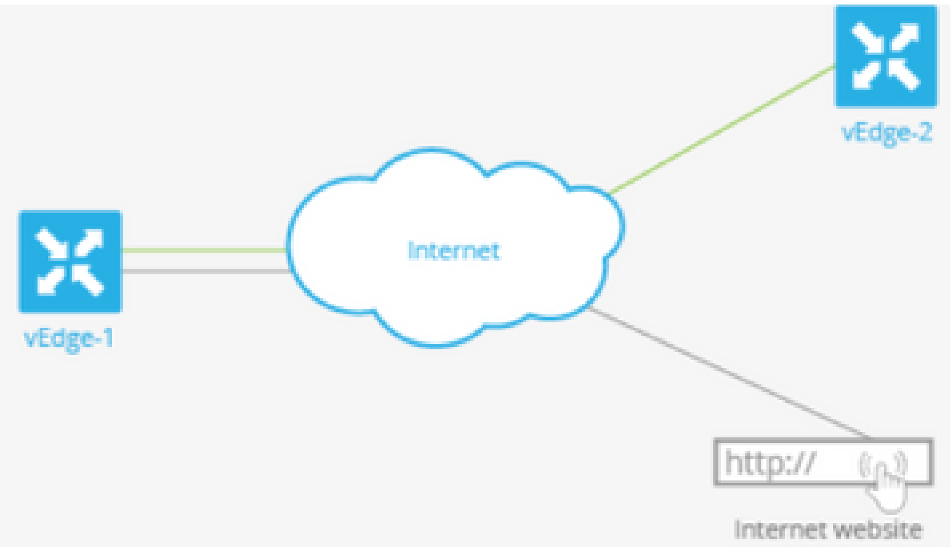

This image explains how the NAT functionality on the vEdge router splits traffic into two flows (or two tunnels) so that some of it remains within the overlay network and some go directly to the Internet or other public networks.

Here, the vEdge router has two interfaces:

- Interface ge0/1 faces the local site and is in VPN 1. Its IP address is 10.1.12.0/24.
- Interface ge0/0 faces the transport cloud and is in VPN 0 (the transport VPN). Its IP address is 192.23.100.0/24, and it uses the default OMP port number, 12346, for overlay network tunnels.

In order to configure the vEdge router to act as a NAT device so that some traffic from the router can go directly to a public network, you do three things:

- Enable NAT in the transport VPN (VPN 0) on the WAN-transport–facing interface, which here is ge0/0. All traffic exiting from the vEdge router, going either to other overlay network sites or to a public network, passes through this interface.
- To direct data traffic from other VPNs to exit from the vEdge router directly to a public network, enable NAT in those VPNs or ensure that those VPNs have a route to VPN 0.

When NAT is enabled, all traffic that passes through VPN 0 is NATed. This includes both the data traffic from VPN 1 that is destined for a public network and all control traffic, including the traffic required to establish and maintain DTLS control plane tunnels between the vEdge router and the vSmart controller and between the router and the vBond orchestrator.

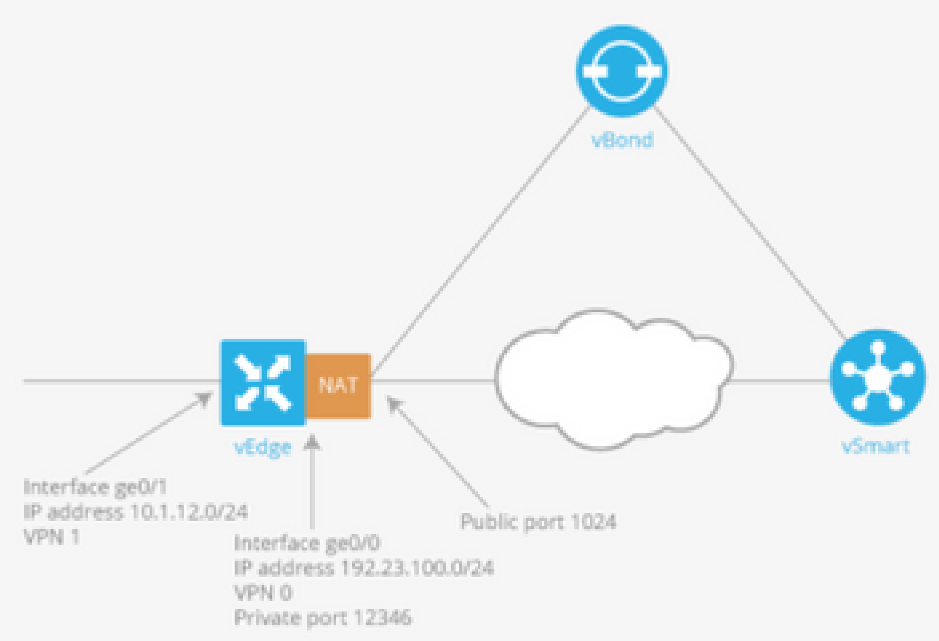

Track Interface Status

Tracking the interface status is useful when you enable NAT on a transport interface in VPN 0 to allow data traffic from the router to exit directly to the internet rather than having to first go to a router in a data center. In this situation, enabling NAT on the transport interface splits the TLOC between the local router and the data center into two, with one going to the remote router and the other going to the internet.

When you enable transport tunnel tracking, the software periodically probes the path to the internet to determine whether it is up. If the software detects that this path is down, it withdraws the route to the internet destination, and traffic destined to the internet is then routed through the data center router. When the software detects that the path to the internet is again functioning, the route to the internet is reinstalled.

Configurations

1. Configure **tracker** under the **system** block.

endpoint-dns-name<*dns-name*> is the DNS name of the endpoint of the tunnel interface. This is the destination on the internet to which the router sends probes to determine the status of the transport interface.

```
system
 tracker tracker
  endpoint-dns-name google.com
 !
!
```

2. Configure **nat** and **tracker** on the transport interface.

```
vpn 0
 interface ge0/0
  ip address 192.0.2.70/24
  nat
  !
 tracker tracker
  tunnel-interface
!
!
```

3. Direct traffic to existing locally via VPN 0.

```
vpn 1
 ip route 0.0.0.0/0 vpn 0
!
```

Verify

Use this section to confirm that your configuration works properly.

1. The check default route is in VPN 0.

```
vEdge# show ip route vpn 0
Codes Proto-sub-type:
  IA -> ospf-intra-area, IE -> ospf-inter-area,
  E1 -> ospf-external1, E2 -> ospf-external2,
  N1 -> ospf-nssa-external1, N2 -> ospf-nssa-external2,
  e -> bgp-external, i -> bgp-internal
Codes Status flags:
  F -> fib, S -> selected, I -> inactive,
  B -> blackhole, R -> recursive

                                            PROTOCOL  NEXTHOP     NEXTHOP          NEXTHOP

VPN    PREFIX              PROTOCOL         SUB TYPE  IF NAME     ADDR             VPN      TLOC
 IP          COLOR            ENCAP  STATUS
------------------------------------------------------------------------------------------------
---------------------------------------------
0      0.0.0.0/0           static           -         ge0/0       192.0.2.1       -        -
            -                 -      F,S
0      192.0.2.255/32      connected        -         system      -                -        -
            -                 -      F,S
0      192.0.2.70/24       connected        -         ge0/0       -                -        -
            -                 -      F,S
```

2. Tracker Status should be 'UP' in show interface VPN 0.

```
vEdge# show interface ge0/0

                                      IF      IF      IF
                        TCP
                AF                    ADMIN   OPER    TRACKER  ENCAP
          SPEED          MSS                  RX        TX
VPN  INTERFACE  TYPE  IP ADDRESS      STATUS  STATUS  STATUS   TYPE   PORT TYPE  MTU   HWADDR
          MBPS   DUPLEX  ADJUST  UPTIME       PACKETS   PACKETS
------------------------------------------------------------------------------------------------
------------------------------------------------------------------
0    ge0/0      ipv4  192.0.2.70/24  Up      Up      Up       null   transport  1500
12:b7:c4:d5:0c:50  1000   full    1420    19:17:56:35  21198589  24842078
```

3. Look for 'NAT' route entry in the RIB.

```
vEdge# show ip routes nat
Codes Proto-sub-type:
  IA -> ospf-intra-area, IE -> ospf-inter-area,
  E1 -> ospf-external1, E2 -> ospf-external2,
  N1 -> ospf-nssa-external1, N2 -> ospf-nssa-external2,
  e -> bgp-external, i -> bgp-internal
Codes Status flags:
  F -> fib, S -> selected, I -> inactive,
  B -> blackhole, R -> recursive

                                            PROTOCOL  NEXTHOP     NEXTHOP          NEXTHOP

VPN    PREFIX              PROTOCOL         SUB TYPE  IF NAME     ADDR             VPN      TLOC
IP          COLOR            ENCAP  STATUS
------------------------------------------------------------------------------------------------
---------------------------------------------
1      0.0.0.0/0           nat              -         ge0/0       -                0        -
            -                -      F,S
```

4. Cross-check that the default-route from the service-side points to the Transport interface with NAT on.

```
vEdge# show ip route vpn 1 0.0.0.0
Codes Proto-sub-type:
  IA -> ospf-intra-area, IE -> ospf-inter-area,
  E1 -> ospf-external1, E2 -> ospf-external2,
  N1 -> ospf-nssa-external1, N2 -> ospf-nssa-external2,
  e -> bgp-external, i -> bgp-internal
Codes Status flags:
  F -> fib, S -> selected, I -> inactive,
  B -> blackhole, R -> recursive

                                      PROTOCOL  NEXTHOP     NEXTHOP          NEXTHOP

VPN    PREFIX              PROTOCOL   SUB TYPE  IF NAME     ADDR             VPN     TLOC IP
      COLOR            ENCAP  STATUS
------------------------------------------------------------------------------------------------
------------------------------
1      0.0.0.0/0           nat        -         ge0/0       -                0        -
      -                -      F,S
```

Troubleshoot

Use this section in order to confirm that your configuration works properly.

1. Ensure that the endpoint-ip or endpoint-dns-name is something on the Internet that can respond to HTTP requests. Also, verify that the endpoint IP address is not the same as the Transport interface. In the case, "Tracker Status" will show as "Down".

```
vEdge# show interface ge0/0

                                      IF      IF      IF
                        TCP
                AF                    ADMIN   OPER    TRACKER  ENCAP
          SPEED          MSS                  RX         TX
VPN  INTERFACE  TYPE  IP ADDRESS      STATUS  STATUS  STATUS   TYPE   PORT TYPE  MTU   HWADDR
          MBPS   DUPLEX  ADJUST  UPTIME       PACKETS    PACKETS
------------------------------------------------------------------------------------------------
------------------------------------------------------------------
0    ge0/0      ipv4  192.0.2.70/24  Up      Up      Down     null   transport  1500
```

```
12:b7:c4:d5:0c:50  1000   full    1420    19:18:24:12  21219358  24866312
```

2. Here is an example that can be used in order to verify that packets go out to the Internet. For example, 8.8.8.8 is Google DNS. Packets from VPN 1 are sourced.

```
vEdge# ping vpn 1 8.8.8.8
Ping in VPN 1
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=51 time=0.473 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=51 time=0.617 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=51 time=0.475 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=51 time=0.505 ms
64 bytes from 8.8.8.8: icmp_seq=5 ttl=51 time=0.477 ms
--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 3999ms
rtt min/avg/max/mdev = 0.473/0.509/0.617/0.058 ms
```

Verify the NAT translational filters. You will see that the NAT filter is built for Internet Control Message Protocol (ICMP).

```
vEdge# show ip nat filter

                            PRIVATE                       PRIVATE  PRIVATE  PUBLIC
        PUBLIC  PUBLIC

NAT  NAT                    SOURCE        PRIVATE DEST    SOURCE   DEST     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   1    icmp      192.0.0.70  8.8.8.8         13067    13067    192.0.2.70  8.8.8.8
      13067   13067   established  0:00:00:02  5         510       5        490      -
```