Configure ASA Virtual Tunnel Interfaces in dual ISP Scenario

Contents

Introduction

This document describes how to configure VTI (Virtual Tunnel Intrfaces) between two ASAs (Adaptive Security Appliances) with use of IKEv2 (Internet Key Exchange version 2) protocol to provide secure connectivity between two branches. Both of the branches have two ISP links for high availablility and load balancing purposes. Border Gateway Protocol (BGP) neighborship is established over the tunnels in order to exchange internal routing information.
This feature is introduced in ASA version 9.8(1). ASA VTI implementation is compatible with VTI implementation available on IOS routers.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- BGP protocol

Components Used

The information in this document is based on ASAv firewalls running 9.8(1)6 software version.

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.

Differences between VTI and Crypto Map

- Crypto map is an output feature of the interface. In order to send the traffic through crypto map based tunnel, the traffic needs to be routed to the internet facing interface (traditionally called outside interface) and must be matched against crypto ACL. On the other hand, VTI is a logical interface. Tunnel to every VPN peer is represented by a different VTI. If the routing points towards VTI, the packet will be encrypted and sent to the corresponding peer.

- VTI eliminates the need to use crypto access lists and Network Address Translation (NAT) exemption rules.

- Crypto map Access Control List (ACL) does not allow for overlapping entries. VTI is a route based VPN and regular routing rules apply for the VPN traffic, which simplifies configuration and processes to troubleshoot.

- Crypto map automatically prevents traffic between sites to be sent in cleartext if tunnel is down. VTI does not automatically protect against it. Null routes need to be added to ensure equal functionality.

Configure

Network Diagram

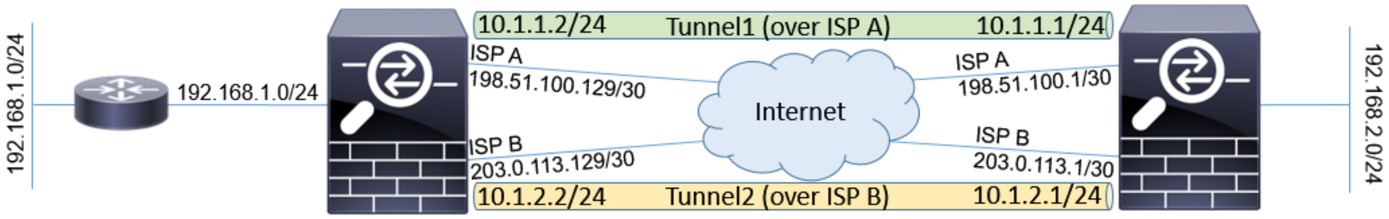

Configurations

> **Note**: This example is not suitable for the scenario where the ASA is a member of independed autonomous system and has BGP peerings with ISP networks. It covers the topology where ASA has two independent ISP links with public addresses from different autonomous systems. In such case, ISP may deploy anti-spoofing protection that verifies if the received packets are not sourced from public IP that belongs to another ISP. In this configuration, proper measures are taken to prevent this.

1. Common encryption and authentication parameters. Information about recommended cryptographic parameters can be found at: https://www.cisco.com/c/en/us/about/security-center/next-generation-cryptography.html

   On both ASAs:

```
crypto ikev2 policy 10
encryption aes-256
```

```
integrity sha256
group 24
prf sha256
lifetime seconds 86400
!
crypto ipsec ikev2 ipsec-proposal PROP
protocol esp encryption aes-256
protocol esp integrity sha-256
```

2. Configure the IPsec profile. One of the sides has to be initiator and one needs to be a responder of the IKEv2 negotiation:

   ASA left:

   ```
   crypto ipsec profile PROF
   set ikev2 ipsec-proposal PROP
   set pfs group24
   responder-only
   ```

   ASA right:

   ```
   crypto ipsec profile PROF
   set ikev2 ipsec-proposal PROP
   set pfs group24
   ```

3. Enable IKEv2 protocol on both ISP interfaces.

   Both ASAs:

   ```
   crypto ikev2 enable ispa
   crypto ikev2 enable ispb
   ```

4. Configure the Pre-shared key to mutually authenticate the ASAs:

   ASA left:

   ```
   tunnel-group 198.51.100.1 type ipsec-l2l
   tunnel-group 198.51.100.1 ipsec-attributes
   ikev2 remote-authentication pre-shared-key *****
   ikev2 local-authentication pre-shared-key *****
   !
   tunnel-group 203.0.113.1 type ipsec-l2l
   tunnel-group 203.0.113.1 ipsec-attributes
   ikev2 remote-authentication pre-shared-key *****
   ikev2 local-authentication pre-shared-key *****
   ```

   ASA right:

   ```
   tunnel-group 198.51.100.129 type ipsec-l2l
   tunnel-group 198.51.100.129 ipsec-attributes
   ikev2 remote-authentication pre-shared-key *****
   ikev2 local-authentication pre-shared-key *****
   !
   tunnel-group 203.0.113.129 type ipsec-l2l
   tunnel-group 203.0.113.129 ipsec-attributes
   ikev2 remote-authentication pre-shared-key *****
   ikev2 local-authentication pre-shared-key *****
   ```

5. Configure the ISP interfaces:

   ASA left:
   ```
   interface GigabitEthernet0/1
   nameif ispa
   security-level 0
   ip address 198.51.100.129 255.255.255.252
   !
   interface GigabitEthernet0/2
   nameif ispb
   security-level 0
   ip address 203.0.113.129 255.255.255.252
   !
   ```
   ASA right:
   ```
   interface GigabitEthernet0/1
   nameif ispa
   security-level 0
   ip address 198.51.100.1 255.255.255.252
   !
   interface GigabitEthernet0/2
   nameif ispb
   security-level 0
   ip address 203.0.113.1 255.255.255.252
   !
   ```
6. The primary link is ISP A interface. ISP B is secondary. The primary link availability is tracked with use of ICMP ping request to a host in the internet, in this example the ASAs use each other ISP A interface as ping destination:

   ASA left:
   ```
   sla monitor 1
   type echo protocol ipIcmpEcho 198.51.100.1 interface ispa
   !
   sla monitor schedule 1 life forever start-time now
   !
   track 1 rtr 1 reachability
   !
   route ispa 0.0.0.0 0.0.0.0 198.51.100.130 1 track 1
   route ispb 0.0.0.0 0.0.0.0 203.0.113.130 10
   ```
   ASA right:
   ```
   sla monitor 1
   type echo protocol ipIcmpEcho 198.51.100.129 interface ispa
   !
   sla monitor schedule 1 life forever start-time now
   !
   track 1 rtr 1 reachability
   !
   route ispa 0.0.0.0 0.0.0.0 198.51.100.2 1 track 1
   route ispb 0.0.0.0 0.0.0.0 203.0.113.2 10
   ```
7. The primary VTI is always established over the ISP A. Secondary VTI is established over ISP B. Static routes towards tunnel destination are needed. This ensures that the encrypted packets leave from the correct physical interface to avoid ISP anti-spoofing drops:

   ASA left:
   ```
   route ispa 198.51.100.1 255.255.255.255 198.51.100.130 1
   route ispb 203.0.113.1 255.255.255.255 203.0.113.130 1
   ```
   ASA right:
   ```
   route ispa 198.51.100.129 255.255.255.255 198.51.100.2 1
   route ispb 203.0.113.129 255.255.255.255 203.0.113.2 1
   ```

8. VTI configuration:

   ASA left:

```
interface Tunnel1
nameif tuna
ip address 10.1.1.2 255.255.255.0
tunnel source interface ispa
tunnel destination 198.51.100.1
tunnel mode ipsec ipv4
tunnel protection ipsec profile PROF
!
interface Tunnel2
nameif tunb
ip address 10.1.2.2 255.255.255.0
tunnel source interface ispb
tunnel destination 203.0.113.1
tunnel mode ipsec ipv4
tunnel protection ipsec profile PROF
```

   ASA right:

```
interface Tunnel1
nameif tuna
ip address 10.1.1.1 255.255.255.0
tunnel source interface ispa
tunnel destination 198.51.100.129
tunnel mode ipsec ipv4
tunnel protection ipsec profile PROF
!
interface Tunnel2
nameif tunb
ip address 10.1.2.1 255.255.255.0
tunnel source interface ispb
tunnel destination 203.0.113.129
tunnel mode ipsec ipv4
tunnel protection ipsec profile PROF
```

9. BGP configuration. The tunnel associated with ISP A is a primary. Prefixes advertised over the tunnel formed over ISP B have lower local-prefernce which makes them less preferred by the routing table:

   ASA left:

```
route-map BACKUP permit 10
set local-preference 80
!
router bgp 65000
bgp log-neighbor-changes
address-family ipv4 unicast
neighbor 10.1.1.1 remote-as 65000
neighbor 10.1.1.1 activate
neighbor 10.1.1.1 next-hop-self
neighbor 10.1.2.1 remote-as 65000
neighbor 10.1.2.1 activate
neighbor 10.1.2.1 next-hop-self
neighbor 10.1.2.1 route-map BACKUP out
network 192.168.1.0
no auto-summary
no synchronization
exit-address-family
```

   ASA right:

```
route-map BACKUP permit 10
```

```
set local-preference 80
!
router bgp 65000
bgp log-neighbor-changes
address-family ipv4 unicast
neighbor 10.1.1.2 remote-as 65000
neighbor 10.1.1.2 activate
neighbor 10.1.1.2 next-hop-self
neighbor 10.1.2.2 remote-as 65000
neighbor 10.1.2.2 activate
neighbor 10.1.2.2 next-hop-self
neighbor 10.1.2.2 route-map BACKUP out
network 192.168.2.0
no auto-summary
no synchronization
exit-address-family
```

10. (Optional) In order to advertise additional network behind left ASA that is not directly connected to it, static route redistribution can be configured:

    ASA left:

    ```
    route inside 192.168.10.0 255.255.255.0 192.168.1.100 1
    !
    prefix-list REDISTRIBUTE_LOCAL seq 10 permit 192.168.10.0/24
    !
    route-map REDISTRIBUTE_LOCAL permit 10
    match ip address prefix-list REDISTRIBUTE_LOCAL
    !
    router bgp 65000
    address-family ipv4 unicast
    redistribute static route-map REDISTRIBUTE_LOCAL
    ```

11. (Optional) The traffic can be load balanced between the tunnels based on the packet destination. In this example, route towards 192.168.10.0/24 network is preferred over backup tunnel (ISP B tunnel)

    ASA left:

    ```
    route-map BACKUP permit 5
    match ip address prefix-list REDISTRIBUTE_LOCAL
    set local-preference 200
    !
    route-map BACKUP permit 10
    set local-preference 80
    ```

12. To prevent traffic between sites from being sent in cleartext to the internet if tunnels are down, Null routes need to be added. All RFC1918 addresses were added for simplicity:

    Both ASAs:

    ```
    route Null0 10.0.0.0 255.0.0.0 250
    route Null0 172.16.0.0 255.240.0.0 250
    route Null0 192.168.0.0 255.255.0.0 250
    ```

13. (Optional) By default, the ASA BGP process sends keepalives once per 60 seconds. If the keepalive response is not received from the peer for 180 seconds, it is declared dead. In order to speed up the detection neigbor failure, you can configure BGP timers. In this example, the keepalives are sent every 10 seconds and neighbor is declared down after 30 seconds.

```
router bgp 65000
address-family ipv4 unicast
neighbor 10.1.1.2 timers 10 30
neighbor 10.1.2.2 timers 10 30
exit-address-family
```

Verify

Verify if IKEv2 tunnel is up:

```
ASA-right(config)# show crypto ikev2 sa

IKEv2 SAs:

Session-id:32538, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local Remote Status Role
836052177 198.51.100.1/500 198.51.100.129/500 READY INITIATOR
Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:24, Auth sign: PSK, Auth verify: PSK
Life/Active Time: 86400/7 sec
Child sa: local selector 0.0.0.0/0 - 255.255.255.255/65535
remote selector 0.0.0.0/0 - 255.255.255.255/65535
ESP spi in/out: 0xc6623962/0x5c4a3bce

IKEv2 SAs:

Session-id:1711, Status:UP-ACTIVE, IKE count:1, CHILD count:1

Tunnel-id Local Remote Status Role
832833529 203.0.113.1/500 203.0.113.129/500 READY INITIATOR
Encr: AES-CBC, keysize: 256, Hash: SHA256, DH Grp:24, Auth sign: PSK, Auth verify: PSK
Life/Active Time: 86400/29 sec
Child sa: local selector 0.0.0.0/0 - 255.255.255.255/65535
remote selector 0.0.0.0/0 - 255.255.255.255/65535
ESP spi in/out: 0x2e3715af/0xc20e22b4
```

Verify BGP neighborship status:

```
ASA-right(config)# show bgp summary
BGP router identifier 203.0.113.1, local AS number 65000
BGP table version is 29, main routing table version 29
3 network entries using 600 bytes of memory
5 path entries using 400 bytes of memory
5/3 BGP path/bestpath attribute entries using 1040 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 2040 total bytes of memory
BGP activity 25/22 prefixes, 69/64 paths, scan interval 60 secs

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
10.1.1.2 4 65000 6 5 29 0 0 00:00:51 2
10.1.2.2 4 65000 7 6 29 0 0 00:01:20 2
```

Verify routes received from BGP. Routes marked with ">" are installed in the routing table:

```
ASA-right(config)# show bgp
```

```
BGP table version is 29, local router ID is 203.0.113.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale, m multipath
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
*>i192.168.1.0 10.1.1.2 0 100 0 i
* i 10.1.2.2 0 80 0 i
*> 192.168.2.0 0.0.0.0 0 32768 i
* i192.168.10.0 10.1.1.2 0 100 0 ?
*>i 10.1.2.2 0 200 0 ?
```

Verify routing table:

```
ASA-right(config)# show route

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, V - VPN
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, + - replicated route
Gateway of last resort is 198.51.100.2 to network 0.0.0.0

S* 0.0.0.0 0.0.0.0 [1/0] via 198.51.100.2, ispa
S 10.0.0.0 255.0.0.0 is directly connected, Null0
C 10.1.1.0 255.255.255.0 is directly connected, tuna
L 10.1.1.1 255.255.255.255 is directly connected, tuna
C 10.1.2.0 255.255.255.0 is directly connected, tunb
L 10.1.2.1 255.255.255.255 is directly connected, tunb
S 172.16.0.0 255.240.0.0 is directly connected, Null0
S 192.168.0.0 255.255.0.0 is directly connected, Null0
B 192.168.1.0 255.255.255.0 [200/0] via 10.1.1.2, 00:02:06
C 192.168.2.0 255.255.255.0 is directly connected, inside
L 192.168.2.1 255.255.255.255 is directly connected, inside
B 192.168.10.0 255.255.255.0 [200/0] via 10.1.2.2, 00:02:35
C 198.51.100.0 255.255.255.252 is directly connected, ispa
L 198.51.100.1 255.255.255.255 is directly connected, ispa
S 198.51.100.129 255.255.255.255 [1/0] via 198.51.100.2, ispa
C 203.0.113.0 255.255.255.252 is directly connected, ispb
L 203.0.113.1 255.255.255.255 is directly connected, ispb
S 203.0.113.129 255.255.255.255 [1/0] via 203.0.113.2, ispb
```

Troubleshoot

Debugs used to troubleshoot IKEv2 protocol:

```
debug crypto ikev2 protocol 4
debug crypto ikev2 platform 4
```

For more information about troubleshooting IKEv2 protocol:
https://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/115935-asa-ikev2-debugs.html

For more information about troubleshooting BGP protocol:

https://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/118050-config-bgp-00.html#anc37

Related Information

- **BGP route selection rules: https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html**
- **ASA BGP configuration guide: https://www.cisco.com/c/en/us/support/docs/security/asa-5500-x-series-next-generation-firewalls/118050-config-bgp-00.html**
- **Technical Support & Documentation - Cisco Systems**