Configure BGP Multipath in Secure Firewall Threat Defense

Contents

Introduction

This document describes how to configure Border Gateway Protocol (BGP) Multipath in Cisco Secure Firewall Threat Defense.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- BGP configuration on Cisco Secure Firewall Threat Defense (FTD)
- General BGP
- Cisco Secure Firewall Management Center (FMC)

Components Used

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

- Cisco FTD version 7.6
- Cisco FMC version 7.6

Disclaimer: The networks and IP addresses referenced in this document are not associated with any individual users, groups, or organizations. This configuration has been created exclusively for use in a lab environment.

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.

Background Information

This document describes how to configure BGP Multipath Load sharing in Firepower Threat Defence when a route to the same destination is received from different routers in the same AS (For example, same ISP).

BGP Multipath allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix. Traffic to the destination prefix is then shared across all installed paths.

These paths are added to the routing table alongside the best path for load-sharing purposes. BGP Multipath does not influence the process of selecting the best path. For instance, a FTD still selects one path as the best based on the algorithm and advertises this best path to its BGP peers.

To qualify as candidates for multipath, paths to the same destination must match the best path in these characteristics:

- Weight
- Local preference
- AS-PATH length
- Origin code
- Multi Exit Discriminator (MED)

One of the following:

- Neighboring AS or sub-AS (prior to BGP Multipath addition)
- AS-PATH (following BGP Multipath addition)

Certain BGP Multipath features impose further requirements on multipath candidates:

- The path must originate from an external or confederation-external neighbor (eBGP).
- The IGP metric to the BGP next hop must match the IGP metric of the best path.

For internal BGP (iBGP) multipath candidates, these additional requirements apply:

- The path must be learned from an internal neighbor (iBGP).
- The IGP metric to the BGP next hop must match the best paths IGP metric, unless the router is set for unequal-cost iBGP multipath.

BGP inserts up to *n* most recently received paths from multipath candidates into the IP routing table, where *n* is the number of routes to install to the routing table, as specified when you configure BGP Multipath. The range of values for *n* are from 1-8 for FTD. The default value, when multipath is disabled, is 1.

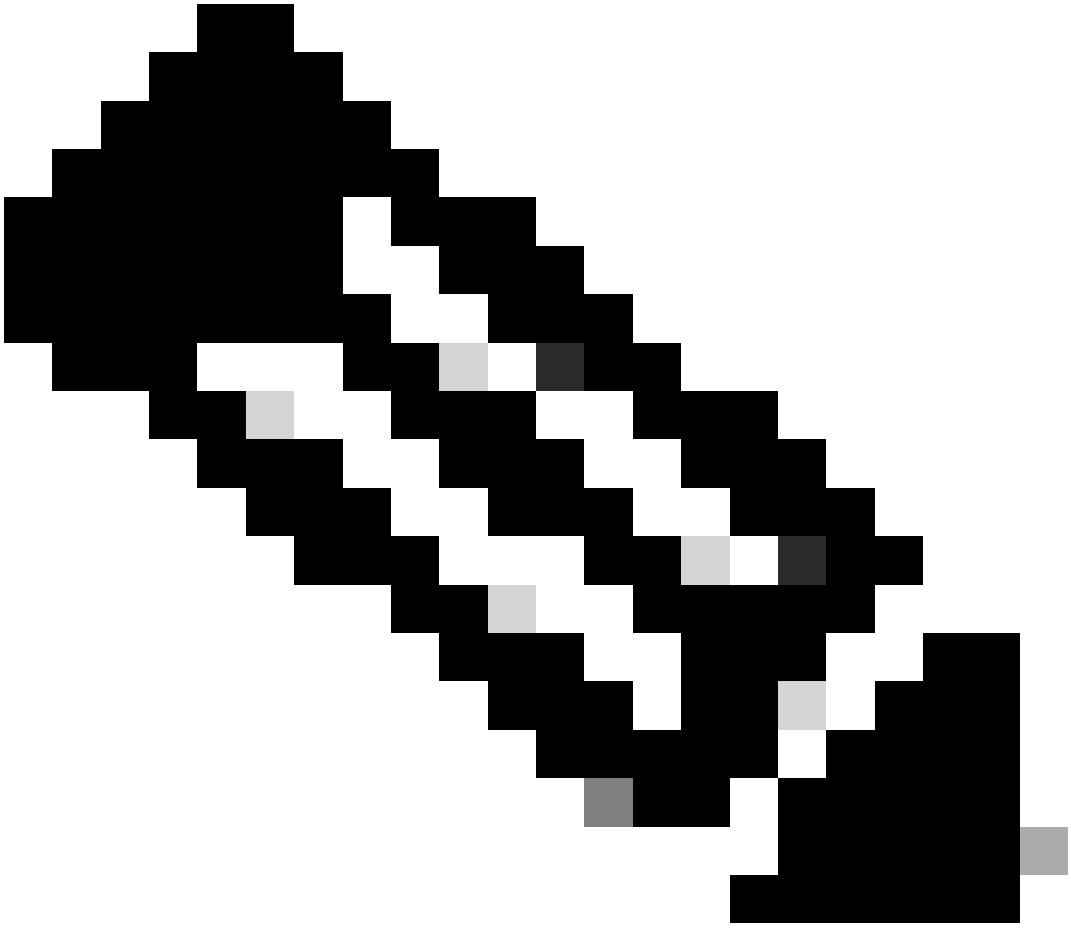

Note: The equivalent next-hop-self is performed on the best path that is selected among eBGP multipaths before it is forwarded to internal peers.

Configure

Diagram

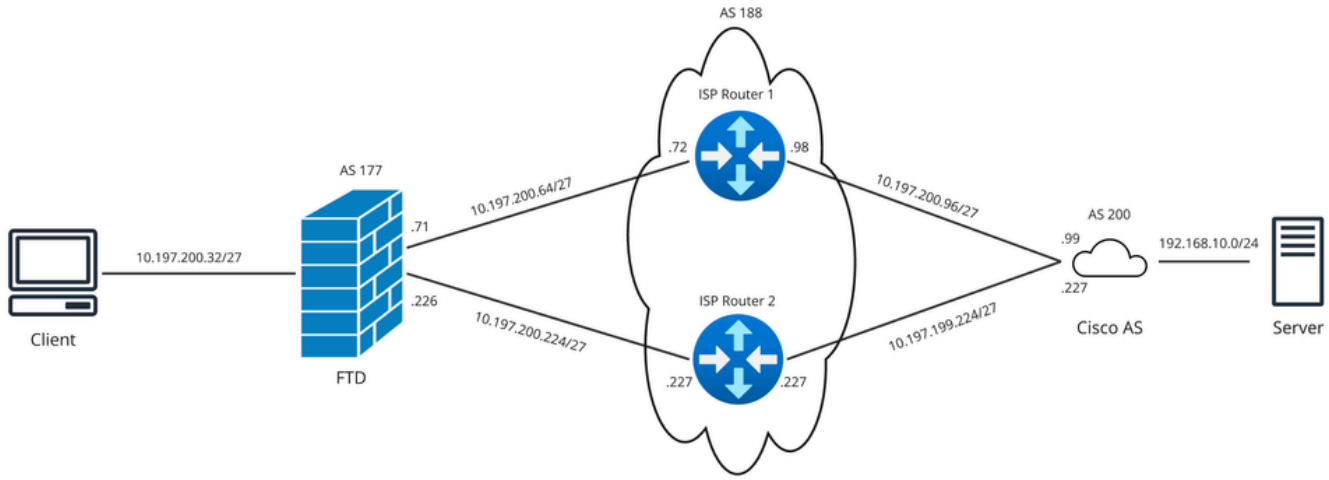

Network Diagram

BGP Configuration

Navigate to **Devices > Device management > Edit Device > Routing > BGP > IPv4** to configure BGP after enabling it.

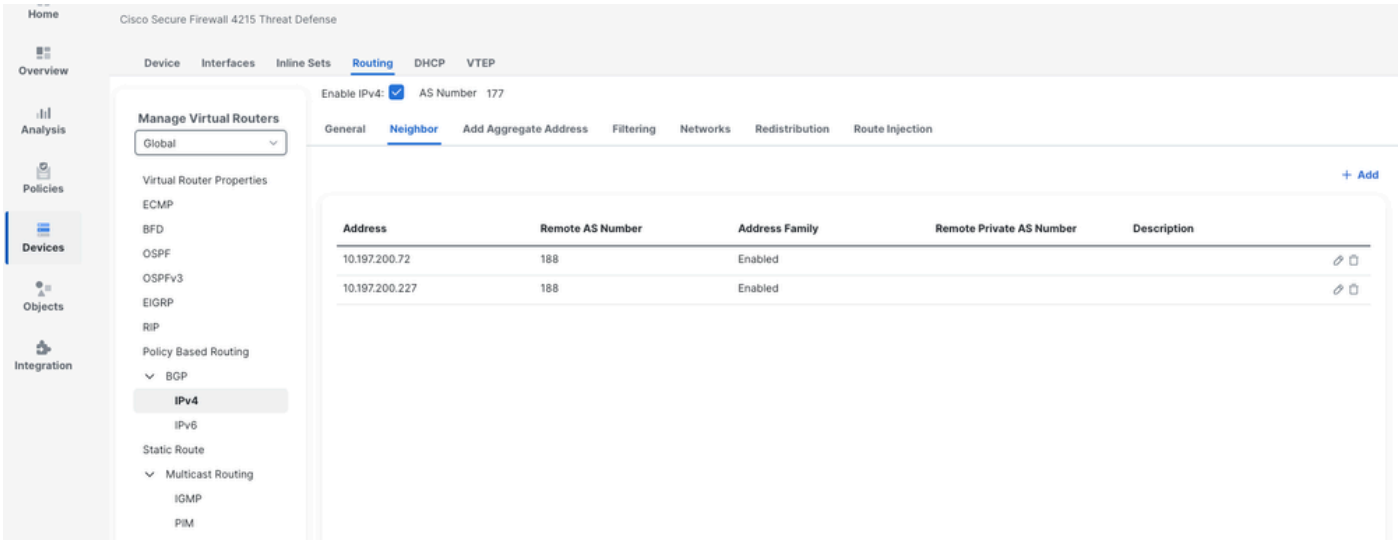

BGP Configuration

BGP configuration from LINA:

```
router bgp 177
 bgp log-neighbor-changes
 bgp router-id 1.1.x.x
 bgp router-id vrf auto-assign
 address-family ipv4 unicast
  neighbor 10.197.200.72 remote-as 188
  neighbor 10.197.200.72 transport path-mtu-discovery disable
  neighbor 10.197.200.72 activate
  neighbor 10.197.200.227 remote-as 188
  neighbor 10.197.200.227 transport path-mtu-discovery disable
  neighbor 10.197.200.227 activate
  no auto-summary
  no synchronization
```

```
 exit-address-family
!
```

Two BGP neighbors from same AS:

```
ftd1# show bgp summary
BGP router identifier 1.1.x.x, local AS number 177
BGP table version is 9, main routing table version 9
2 network entries using 400 bytes of memory
4 path entries using 320 bytes of memory
1/1 BGP path/bestpath attribute entries using 208 bytes of memory
1 BGP AS-PATH entries using 40 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 968 total bytes of memory
BGP activity 4/2 prefixes, 10/6 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.197.200.72   4          188 67      66             9    0    0 01:10:15  1
10.197.200.227  4          188 60      60             9    0    0 01:00:56  1
```

Notice that there are two valid routes received for same destination network, one from each neighbor.

```
ftd1# show bgp 192.168.10.0 255.255.255.0
BGP routing table entry for 192.168.10.0/24, version 9
Paths: (2 available, best #2, table default)
  Advertised to update-groups: 3
  188 200
    10.197.200.227 from 10.197.200.227 (3.3.x.x)
      Origin incomplete, localpref 100, valid, external
  188 200
    10.197.200.72 from 10.197.200.72 (2.2.x.x)
      Origin incomplete, localpref 100, valid, external, best
```

You can see that the best path selected and installed in routing table is the one which is received from neighbor 10.197.200.72.

```
ftd1# show route 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, 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
SI - Static InterVRF
Gateway of last resort is not set

B 192.168.10.0 255.255.255.0 [20/0] via 10.197.200.72, 00:01:55
```

BGP Multipath Configuration

Configure BGP multi path under **Devices > Device management > Edit Device > Routing > BGP > IPv4 > Edit Forward Packets Over Multiple Paths**.

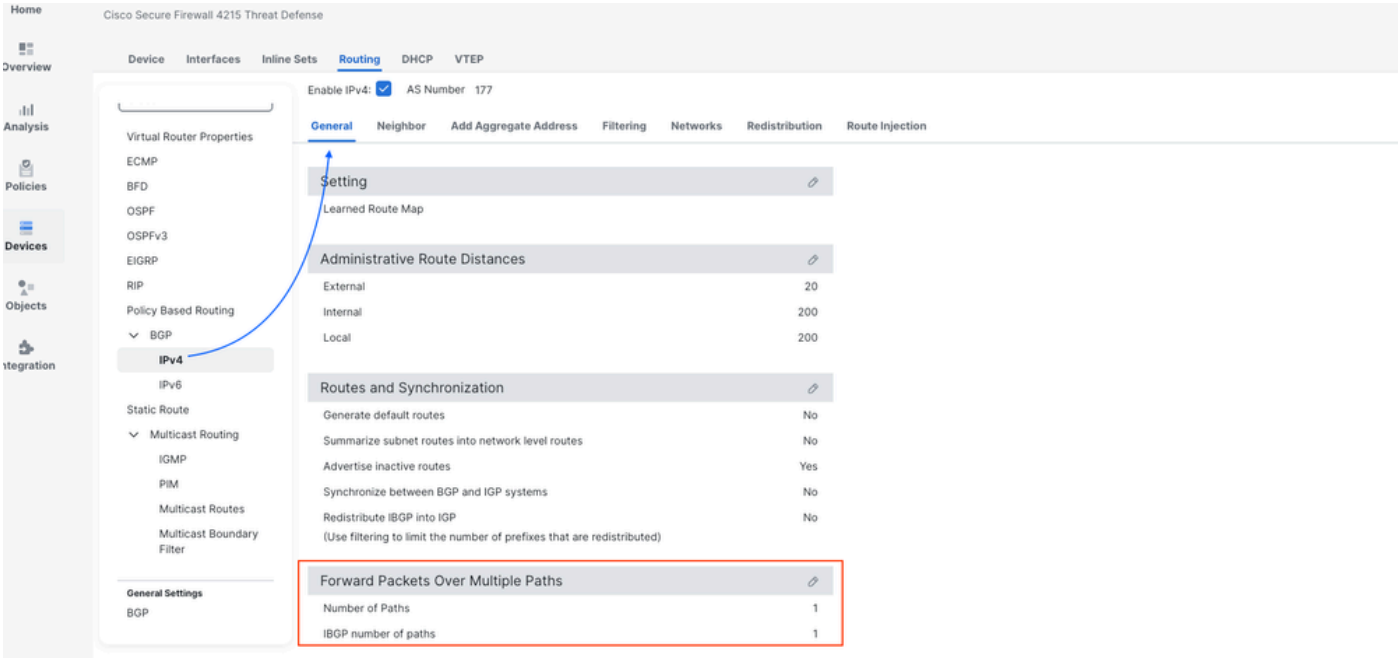

BGP Multipath in FMC

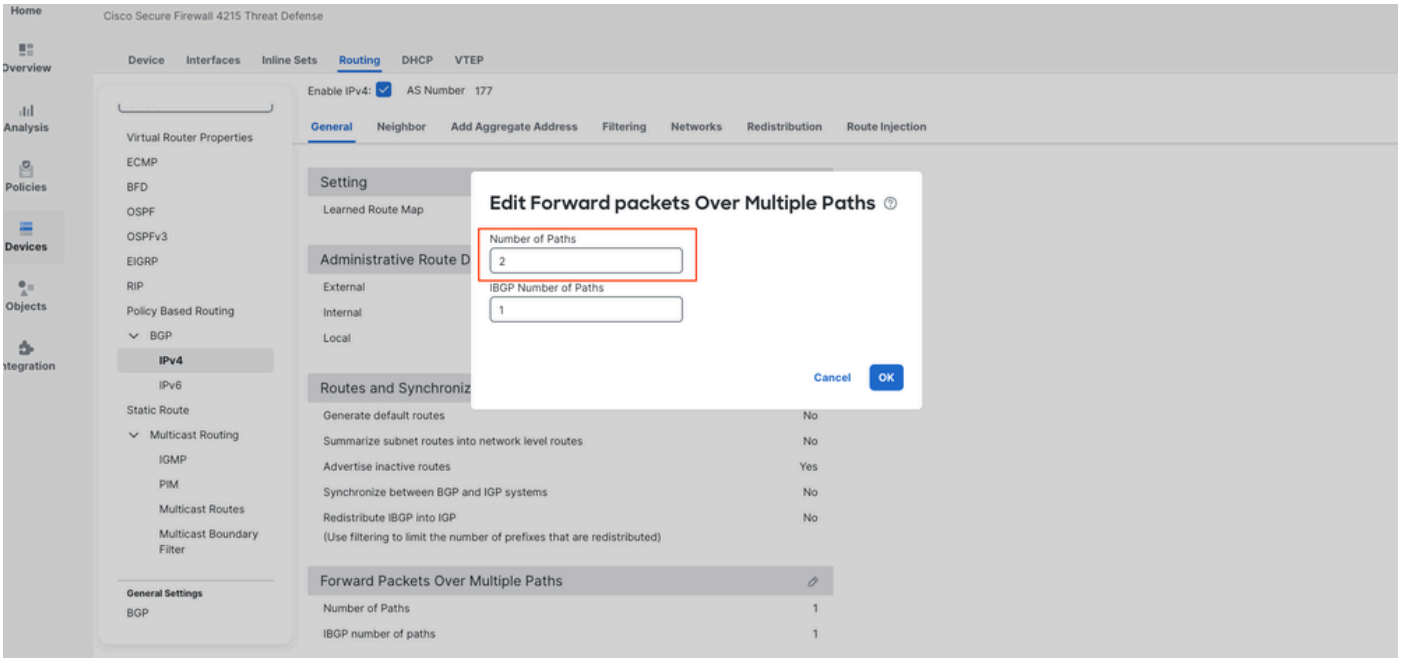

BGP Multipath in FMC

Save the changes and **Deploy**.

Verify

BGP configuration from LINA after enabling multipath:

<#root>

```
ftd1# sh run router
router bgp 177
bgp log-neighbor-changes
bgp router-id 1.1.x.x
bgp router-id vrf auto-assign
address-family ipv4 unicast
neighbor 10.197.200.72 remote-as 188
neighbor 10.197.200.72 transport path-mtu-discovery disable
neighbor 10.197.200.72 activate
neighbor 10.197.200.227 remote-as 188
neighbor 10.197.200.227 transport path-mtu-discovery disable
neighbor 10.197.200.227 activate
maximum-paths 2
no auto-summary
no synchronization
exit-address-family
```

Notice the *m* before one of the routes indicating multipath

<#root>

```
ftd1# show bgp

BGP table version is 11, local router ID is 1.1.x.x
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
*m 192.168.10.0 10.197.200.227 0 188 200 ?
*> 10.197.200.72 0 188 200 ?

ftd1# show bgp 192.168.10.0 255.255.255.0
BGP routing table entry for 192.168.10.0/24, version 11
Paths: (2 available, best #2, table default)
Multipath: eBGP
Advertised to update-groups: 3
188 200
10.197.200.227 from 10.197.200.227 (3.3.x.x)
Origin incomplete, localpref 100, valid, external,
multipath

188 200
10.197.200.72 from 10.197.200.72 (2.2.x.x)
Origin incomplete, localpref 100, valid, external,
multipath
, best
```

Notice that, now there are two routes to the same destination installed in the routing table after enabling

BGP multipath.

<#root>

```
ftd1# show route 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, 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
SI - Static InterVRF
Gateway of last resort is not set
```

B 192.168.10.0 255.255.255.0 [20/0] via 10.197.200.227, 00:01:17

[20/0] via 10.197.200.72, 00:01:17

Troubleshoot

1. Verify BGP Configuration:

Use **show bgp** command to check the BGP table and ensure that multiple paths are marked as multipat.

Confirm that **Number of Paths** is configured to allow multiple paths.

2. Check Path Attributes:

Ensure that the paths have equal BGP attributes required for multipath; such as weight, local preference, AS Path Length and so on.

3. Load Sharing Verification:

Use the **show route** command to verify that the paths are being used for load sharing. The output should show multiple paths for the same destination.

Q&A

1.Is the command **bgp bestpath as-path multipath-relax** supported in FTD via Flex config for load sharing?

No, an enhancement is already in place for this to be supported in FTD/ASA. Cisco bug ID CSCvw16654

Related information

Troubleshoot Common BGP Issues