

Aggregate Route Indirect Next-Hop Option

In This Chapter

This section provides information about aggregate route indirect next-hop option configurations.

Topics in this section include:

- [Applicability on page 1726](#)
- [Overview on page 1727](#)
- [Configuration on page 1728](#)
- [Conclusion on page 1735](#)

Applicability

This section is applicable to all of the 7950 XRS series, 7750 SR series (SR-7, SR-12, SR-c4 and SR-c12), 7710 SR, as well as 7450 ESS (ESS-6, ESS-7 and ESS-12) series in mixed mode. This configuration is supported by all IOM/IMMs and supported on all chassis types as long as the protocol is supported.

The configuration was tested in release 11.0R1.

Overview

The 7x50s have for many releases supported the ability to configure IPv4 and IPv6 aggregate routes. A configured aggregate route that has the best preference for the prefix is added to the routing table (activated) when it has at least one contributing route and removed when there are no longer any more contributing routes. A contributing route is any route installed in the forwarding table that is a more-specific match of the aggregate. (10.16.12.0/24 is a contributing route to the aggregate route 10.16.12.0/22, but for this same aggregate 10.16.12.0/22 and 10.0.0.0/8 are not contributing routes).




Figure 263: Aggregate Routes

In Figure 263, Router A could choose to advertise all the four routes or one aggregate route. By aggregating the four routes, fewer updates are sent on the link between routers A and B, router B needs to maintain a smaller routing table resulting in better convergence and router B saves on computational resources by evaluating fewer entries in its routing table.

Different network operators have different requirements for how to forward a packet that matches an aggregate route but not any of the more-specific routes in the forwarding table that activated it. In general, there are three different options:

1. The packet can be forwarded according to the next-most specific route, ignoring the aggregate route. This can lead to routing loops in some topologies.
2. The packet can be discarded.
3. The packet can be forwarded towards an indirect next-hop address that is configured by the operator. The indirect next-hop could be the address of a threat management server that analyzes the packets it receives for security threats. This option requires the aggregate route to be installed in the forwarding table with a resolved next-hop interface determined from a route lookup of the indirect next-hop address.

Configuration

The test topology is shown in [Figure 264](#).

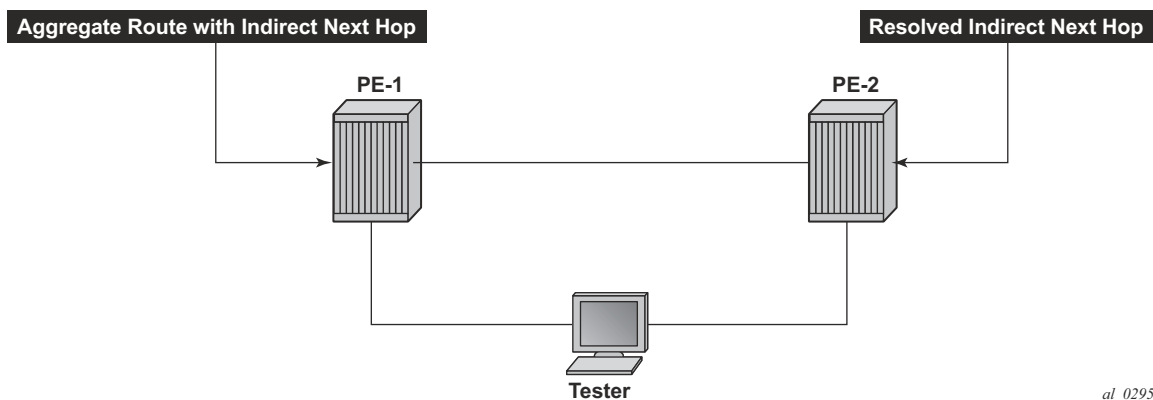


Figure 264: Test topology

This feature adds a new **indirect** keyword and associated IP address parameter to the existing aggregate command in these configuration contexts:

```
config>router
config>service>vprn
```

The aggregate route configuration commands are shown below.

```
config
  router
    aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as
      -number: ip-address] [community comm-id] [black-hole | indirect ip-address]

config
  service
    vprn
    aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as
      -number: ip-address] [community comm-id] [black-hole | indirect ip-address]
```

Parameters:

- **indirect** — This indicates that the aggregate route has an indirect address. Note from this syntax that the indirect option is mutually exclusive with the black-hole option. To change the next-hop type of an aggregate route (for example from black-hole to indirect) the route must be deleted and then re-added with the new next-hop type (however other configuration attributes can generally be changed dynamically).

- `<ip-address>` — Installing an aggregate route with an indirect next-hop is supported for both IPv4 and IPv6 prefixes, however if the aggregate prefix is IPv6 the indirect next-hop must be an IPv6 address and if the aggregate prefix is IPv4 the indirect next-hop must be an IPv4 address.

An indirect next-hop address of an aggregate route may be resolved by any of the following route types:

- Direct/local route
- Static route with regular next-hop, black-hole next-hop or an indirect next-hop
- OSPFv2 or RIP IPv4 route (applicable only to IPv4 aggregate routes)
- LDP shortcut route (applicable only to IPv4 aggregate routes)
- OSPFv2 or IS-IS shortcut route (IPv4 route with an LDP/RSVP or RSVP tunnel next-hop) (applicable only to IPv4 aggregate routes)
- OSPFv3 or IS-IS route
- BGP route resolved by an IGP route
- BGP route resolved by a BGP route
- BGP labeled-IPv4 route resolved by an LDP or RSVP tunnel (applicable only to IPv4 aggregate routes)
- 6PE route resolved by an LDP tunnel or static route with black-hole next-hop (applicable only to IPv6 aggregate routes)
- BGP-VPN route resolved by a BGP labeled-IPv4 route, LDP tunnel, or RSVP tunnel (applicable only to aggregate routes configured in a VPRN context)

If an indirect next-hop is not resolved, the aggregate route will show up as black-hole.

Step 1. Configure the aggregate route.

Command: `aggregate`

Syntax `aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-number: ip-address] [community comm-id] [black-hole | indirect ip-address]`
`no aggregate ip-prefix/ip-prefix-length`

Context `config>router`
`config>service>vprn`

Description Use this command to automatically install an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.

The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.

By default aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.

The **no** form of the command removes the aggregate.

Default No aggregate routes are defined.

Parameters *ip-prefix/ip-prefix-length* — The destination address of the aggregate route.

Values:

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 — 32
ipv6-prefix	x:x:x:x:x:x:x x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D
ipv6-prefix-length	0 — 128

summary-only — This optional parameter suppresses the advertisement of more specific component routes for the aggregate. To remove the summary-only option, enter the same aggregate command without the summary-only parameter.

as-set — This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

aggregator as-number:ip-address — This optional parameter adds the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

community comm-id — This configuration option associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

Aggregate Route Indirect Next-Hop Option

Values: *comm-id* *asn:comm-val | well-known-comm*
asn 0 — 65535
comm-val 0 — 65535
well-known-comm no-advertise, no-export, no-export-subconfed

black-hole — This optional parameter installs the aggregate route, when activated, in the FIB with a black-hole next-hop. Packets matching an aggregate route with a black-hole next hop are discarded.

indirect *ip-address* — This configuration option specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to *ip-address*.

Values *ipv4-prefix* a.b.c.d
ipv6-prefix x:x:x:x:x:x:x
x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

```
A:PE-1>config#router aggregate 10.10.10.0/24 community 64496:64497 indirect 192.168.11.11
```

```
A:PE-1# show router aggregate
```

```
=====
```

```
Aggregates (Router: Base)
```

```
=====
```

Prefix	Aggr IP-Address	Aggr AS
Summary	AS Set	State
NextHop	Community	NextHopType
10.10.10.0/24	0.0.0.0	0
False	False	Inactive
192.168.11.11	64496:64497	Indirect

```
-----
```

```
No. of Aggregates: 1
```

```
=====
```

Step 2. Configure the contributing routes to activate the aggregate route.

```
*A:PE-1# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]
  Next Hop[Interface Name]
Type      Proto    Age      Pref
Metric
-----
10.10.10.0/24
  Black Hole      Remote Aggr    00h00m20s 130
  0
10.10.10.3/32
  10.18.0.89      Remote Static  00h00m24s 5
  1
10.10.10.4/32
  10.18.0.89      Remote Static  00h00m24s 5
  1
10.10.10.5/32
  10.18.0.89      Remote Static  00h00m24s 5
  1
10.10.10.8/32
  10.18.0.89      Remote Static  00h00m24s 5
  1
10.12.0.0/24
  to_PE-2         Local  Local    00h00m24s 0
  0
10.18.0.0/24
  to_Routers      Local  Local    00h00m24s 0
  0
10.19.0.0/24
  to_Tester       Local  Local    00h00m24s 0
  0
10.20.1.1/32
  system          Local  Local    00h00m24s 0
  0
-----
No. of Routes: 9
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
```

The aggregate route is now active:

```
A:PE-1# show router aggregate
=====
Aggregates (Router: Base)
=====
Prefix
  Summary
  NextHop
Aggr IP-Address
  AS Set
  Community
Aggr AS
  State
  NextHopType
-----
10.10.10.0/24
  False
  192.168.11.11
  0.0.0.0
  False
  64496:64497
  0
  Active
  Indirect
-----
No. of Aggregates: 1
=====
```


Step 3. Configure the resolving route.

```
A:PE-2# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type      Proto    Age      Pref
Metric
-----
10.10.10.0/24
  10.29.0.99      Remote  Static   00h00m14s  5
                                     1
10.12.0.0/24
  to_PE-1        Local   Local    00h00m14s  0
                                     0
10.20.1.2/32
  system         Local   Local    00h00m14s  0
                                     0
10.29.0.0/24
  to_Tester      Local   Local    00h00m14s  0
                                     0
192.168.11.0/24
  10.12.0.2      Remote  Static   00h00m14s  5
                                     1
-----
No. of Routes: 5
Flags: L = LFA nexthop available    B = BGP backup route available
      n = Number of times nexthop is repeated
=====
```

The aggregate route now has the resolved next-hop.

```
*A:PE-2# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]
Next Hop[Interface Name]
Type      Proto    Age      Pref
Metric
-----
10.10.10.0/24
  10.12.0.2      Remote  Aggr     00h00m20s  130
                                     0
10.10.10.3/32
  10.18.0.89     Remote  Static   00h00m24s  5
                                     1
10.10.10.4/32
  10.18.0.89     Remote  Static   00h00m24s  5
                                     1
10.10.10.5/32
  10.18.0.89     Remote  Static   00h00m24s  5
                                     1
10.10.10.8/32
  10.18.0.89     Remote  Static   00h00m24s  5
                                     1
10.12.0.0/24
  to_PE-2        Local   Local    00h00m24s  0
                                     0
10.18.0.0/24
  to_Routers     Local   Local    00h00m24s  0
                                     0
10.19.0.0/24
  to_Tester      Local   Local    00h00m24s  0
                                     0
10.20.1.1/32
  system         Local   Local    00h00m24s  0
                                     0
192.168.11.0/24
  10.12.0.2      Remote  Static   00h00m24s  5
                                     1
-----
No. of Routes: 10
```

Configuration

Flags: L = LFA nexthop available B = BGP backup route available
 n = Number of times nexthop is repeated

=====

Conclusion

Aggregate routes offer several advantages, the key being reduction in the routing table size and overcoming routing loops, among other things. Aggregate routes with indirect next hop option helps in faster network convergence by decreasing the number of route table changes. This example shows how to configure aggregate routes with indirect next hop option.

Conclusion