

BGP-4

Border Gateway Protocol 4 (BGP-4) Primer

Diarmuid Ó Briain



Last updated: 18 April 2017



© 2017 by C²S Consulting

Policies made available under the Creative Commons Attribution-NonCommercial 4.0 license (International): <https://creativecommons.org/licenses/by-nc/4.0>

Table of Contents

1. Introduction.....	4
2. BGP-4 Finite State Machine (FSM).....	4
2.1 Testing the BGP-4 FSM state.....	5
3. BGP-4 for IPv6.....	7
4. BGP-4 Messages.....	7
4.1 BGP-4 Update message.....	8
5. Application of BGP-4.....	9
6. Configuration of the Border Router.....	9
6.1 Interface configuration.....	10
6.2 BGP-4 configuration.....	11
6.3 OSPFv2 configuration.....	12
6.4 OSPFv3 configuration.....	12
6.5 Testing the network.....	12
7. Router Configurations.....	13
7.1 BR-A1.....	13
7.2 CR-A2.....	14
7.3 CR-A3.....	15
7.4 BR-B1.....	16
7.5 CR-B2.....	17
7.6 CR-B3.....	18
8. Host configurations.....	19
8.1 Host A7.....	19
8.2 Host A8.....	19
8.3 Host A9.....	19
8.4 Host A10.....	19

Illustration Index

Illustration 1: BGP-4 Simplified Finite State Machine.....	4
Illustration 2: BGP-4 test network.....	9

1. Introduction

Exterior Gateway Protocol (EGP) provide routing information between Autonomous System (AS). An AS is a set of routers under a single technical administration, using an Interior Gateway Protocol (IGP), like Open Shortest Path First (OSPF) and common metrics to determine how to route packets within the AS.

Each Internet Service Provider (ISP) manages the Internet Protocol (IP) space for its customers and aggregates routing information on their behalf to other ISPs and Internet eXchange Points (IXP) in neighbouring ASes. The first such protocol was Border Gateway Protocol (BGP) defined in RFC1105 back in June 1989. This protocol has evolved over the years to RFC4271 which defines BGP version 4 (BGP-4) as the standard inter-AS routing protocol today.

2. BGP-4 Finite State Machine (FSM)

BGP-4 is essentially a distance vector routing protocol that functions as a speaking system to exchange network reachability information between ASes. These ASes are linked by inter-AS routing protocol like BGP-4 to determine how to route packets to other ASes. BGP-4 uses Transmission Control Protocol (TCP) on port 179 as its transport protocol which eliminates the need to implement explicit update fragmentation, retransmission, acknowledgement, and sequencing. A TCP connection is formed between two systems. They exchange messages to open and confirm the connection parameters.

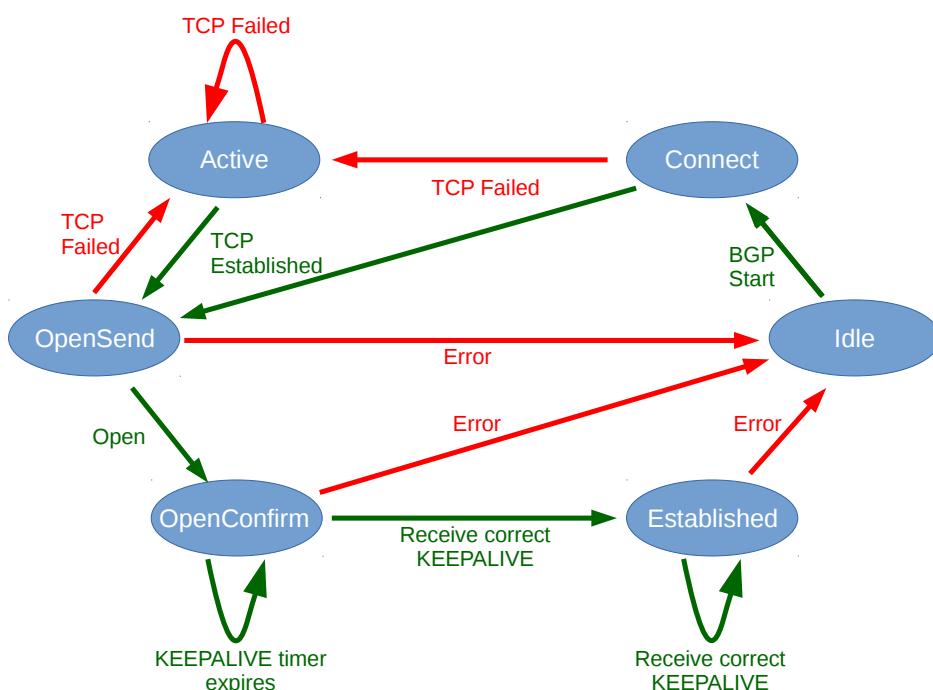


Illustration 1: BGP-4 Simplified Finite State Machine

BGP-4 forms a neighbour relationship with peers. As demonstrated in Illustration 1 the relationship is managed by a Finite State Machine (FSM) that traverses the following states:

State	Meaning
Idle	BGP-4 refuses all incoming BGP connections
Connect	BGP-4 is waiting for the TCP connection to be completed
Active	BGP-4 is trying to acquire a peer by initiating a TCP connection
OpenSent	BGP-4 waits for an OPEN message from its peer
OpenConfirm	BGP-4 waits for a KEEPALIVE or NOTIFICATION message from its peer
Established	BGP-4 can exchange UPDATE, NOTIFICATION and KEEPALIVE messages with its peer

2.1 Testing the BGP-4 FSM state

Review the state of the BGP-4 neighbours via the router command line.

```
BR-B1# show ip bgp summary
BGP router identifier 220.22.22.1, local AS number 2222
RIB entries 7, using 504 bytes of memory
Peers 2, using 5064 bytes of memory

Neighbor      V   AS MsgRcvd MsgSent     TblVer  InQ OutQ Up/Down  State/PfxRcd
195.5.5.1      4 1111     459     461          0    0    0 07:16:39      2
2005:5::       4 1111     463     468          0    0    0 07:16:38      2

Total number of neighbors 2

BR-B1# show ipv6 bgp summary
BGP router identifier 220.22.22.1, local AS number 2222
RIB entries 7, using 504 bytes of memory
Peers 2, using 5064 bytes of memory

Neighbor      V   AS MsgRcvd MsgSent     TblVer  InQ OutQ Up/Down  State/PfxRcd
2005:5::       4 1111     463     468          0    0    0 07:16:55      2

Total number of neighbors 1
```

Once the Internet Protocol (IP) address of a neighbour is identified a much more detailed review of state can be made using the following command.

```
BR-B1# sh ipv6 bgp neighbors 195.5.5.1
BGP neighbor is 195.5.5.1, remote AS 1111, local AS 2222, external link
  BGP version 4, remote router ID 210.11.11.1
  BGP state = Established, up for 07:23:09
  Last read 00:00:06, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    4 Byte AS: advertised and received
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    Inq depth is 0
    Outq depth is 0
      Sent          Rcvd
    Opens:           5            3
    Notifications:  0            3
    Updates:        3            3
    Keepalives:     460          457
    Route Refresh:  0            0
    Capability:    0            0
    Total:          468          466
  Minimum time between advertisement runs is 30 seconds

  For address family: IPv4 Unicast
    Community attribute sent to this neighbor(both)
    2 accepted prefixes

    Connections established 3; dropped 2
    Last reset 07:23:22, due to BGP Notification received
    Local host: 195.5.5.2, Local port: 37261
    Foreign host: 195.5.5.1, Foreign port: 179
    Nexthop: 195.5.5.2
    Nexthop global: 2005:5::1
    Nexthop local: fe80::200:ff:feaa:16
    BGP connection: non shared network
    Read thread: on  Write thread: off
```

Should the state show as “Active” then it is likely there is a problem establishing a TCP connection to the remote peer. Test this as follows by using telnet to try and open a TCP connection to the port. A fail will show like this:

```
BR-B1# telnet 195.5.5.1 179
Trying 195.5.5.1...
telnet: Unable to connect to remote host: Connection refused
```

A pass will show like this:

```
BR-B1# telnet 195.5.5.1 179
Trying 195.5.5.1...
Connected to 195.5.5.1.
Escape character is '^]'.
```

3. BGP-4 for IPv6

RFC4760 “Multiprotocol Extensions for BGP-4” and RFC7606 “Revised Error Handling for BGP-4 UPDATE Messages” define additions to BGP-4 that have allowed it to operate for both IPv4 and IPv6 network layer protocols. To enable BGP-4 to support routing for both protocols it requires the ability to associate a particular IP protocol with the NEXT-HOP information and the ability to associate with associated Network Layer Reachability Information (NLRI). To identify individual IP protocols associated with NEXT-HOP information the concept of an *Address Family* to differentiate was developed.

4. BGP-4 Messages

BGP-4 messages are sent over TCP connections via port 179. A message is processed only after it is entirely received. The maximum message size is 4096 octets. After a TCP connection is established, the first message sent is an OPEN message which, if acceptable is responded with a KEEPALIVE message confirming receipt of the OPEN message.

Message	Function
OPEN	After a TCP connection is established, the first message sent is an OPEN message which is responded with a KEEPALIVE message confirming the OPEN is sent back
UPDATE	Exchange routing information between BGP-4 peers
NOTIFICATION	Sent when an error condition is detected and the BGP-4 connection is closed immediately after it is sent
KEEPALIVE	Determine peer are reachability, these are exchanged often enough not to cause the Hold Timer to expire

4.1 BGP-4 Update message

The BGP-4 UPDATE Message is used to transfer routing information between BGP-4 peers. Consider the following extract from a BGP-4 UPDATE which carried NLRI information for two IPv4 networks 220.0.9.0/24 and 220.0.9.0/24 with the AS_PATH and NEXT_HOP IP address defined.

```
Border Gateway Protocol - UPDATE Message
Marker: ffffffffffffffffffffff
Length: 56
Type: UPDATE Message (2)
Withdrawn Routes Length: 0
Total Path Attribute Length: 25
Path attributes
  Path Attribute - ORIGIN: IGP
  Path Attribute - AS_PATH: 2222 1111
  Path Attribute - NEXT_HOP: 195.5.5.2
Network Layer Reachability Information (NLRI)
  210.0.7.0/24
    NLRI prefix length: 24
    NLRI prefix: 210.0.7.0
  210.0.8.0/24
    NLRI prefix length: 24
    NLRI prefix: 210.0.8.0
```

The following extract was in the same BGP-4 UPDATE packet, it carries NLRI information for the IPv6 Prefixes 2201:9::/64 and 2201:10::/64, again the AS_PATH and NEXT_HOP IPv6 addresses are defined. This allows the routing table to be updated at the peer.

```
Border Gateway Protocol - UPDATE Message
Marker: ffffffffffffffffffffff
Length: 103
Type: UPDATE Message (2)
Withdrawn Routes Length: 0
Total Path Attribute Length: 80
Path attributes
  Path Attribute - MP_REACH_NLRI
    Address family identifier (AFI): IPv6 (2)
    Subsequent address family identifier (SAFI): Unicast (1)
    Next hop network address (32 bytes)
      Next Hop: 2005:5::1
      Next Hop: fe80::200:ff:fea:16
    Number of Subnetwork points of attachment (SNPA): 0
    Network layer reachability information (18 bytes)
      2201:9::/64
        MP Reach NLRI prefix length: 64
        MP Reach NLRI IPv6 prefix: 2201:9::
      2201:10::/64
        MP Reach NLRI prefix length: 64
        MP Reach NLRI IPv6 prefix: 2201:10::
  Path Attribute - ORIGIN: IGP
  Path Attribute - AS_PATH: 2222
  Path Attribute - MULTI_EXIT_DISC: 0
```

5. Application of BGP-4

Each ISP has an assigned AS number from their local Regional Internet Registry (RIR) that is globally unique. IP and IPv6 address blocks are also assigned to ISPs by their respective RIR and the ISP as a Local Internet Registry (LIR) will re-assign parts of the assigned blocks to organisations and individuals. To advertise these blocks it is necessary for the ISP to connect a Border Router (BR) acting as a BGP-4 speaker to peer BR in other ASes. At this BR point the ISP can aggregate the routes they wish advertised to the peer.

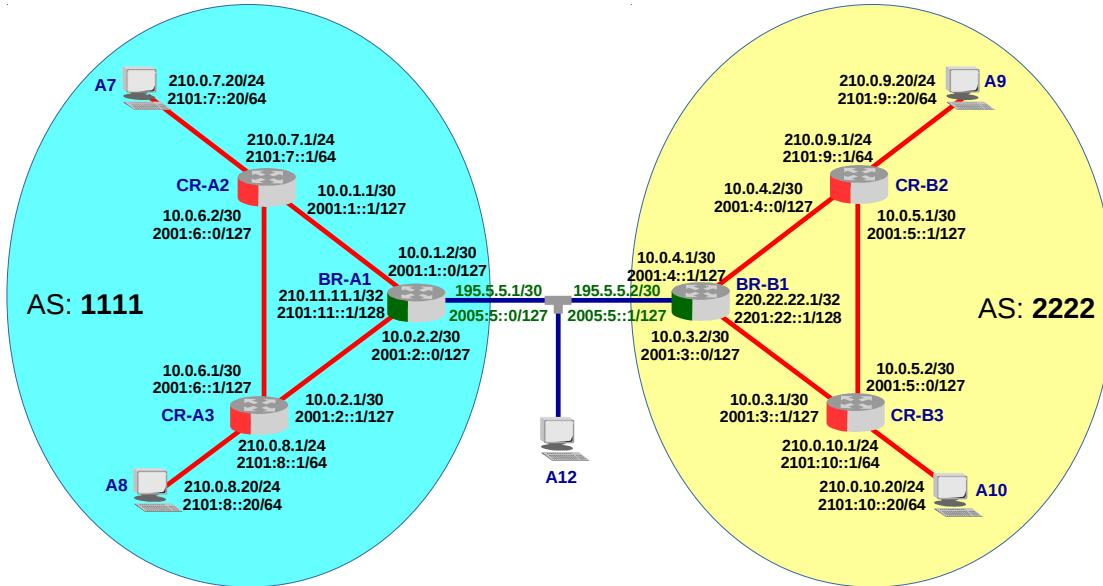


Illustration 2: BGP-4 test network

Considering the networks in Illustration 2. Each of the two ASes share internal routing information via OSPF. The ASes are inter-connected via the BRs BR-A1 in AS 1111 and BR-B1 in AS 2222. Both ASes wish to limit the IP subnets advertised to the other AS to the public IP networks 210.0.X.0/24, 2101:X::/64 and 220.0.X.0/24, 2201:X::/64. Without going into the detail of the configuration on the four Core Routers (CR) it is suffice to say that all interfaces are routed by OSPFv2 and OSPFv3.

6. Configuration of the Border Router

Using BR-A1 for the explanation. The configuration for BR-B1 is quite similar.

6.1 Interface configuration

The first part of the configuration simply assigns IP addresses to the interfaces and creates a loopback interface and assigns it IP addresses. As IPv6 addresses are assigned manually for simplicity there is no need for Router Advertisement (RA) messages so they are suppressed.

```
BR-A1# configure terminal
BR-A1(config)# interface lo
BR-A1(config-if)# ip address 210.11.11.1/32
BR-A1(config-if)# ipv6 address 2101:11::1/128

BR-A1(config)# interface eth0
BR-A1(config-if)# ip address 195.5.5.1/30
BR-A1(config-if)# ipv6 address 2005:5::1/127
BR-A1(config-if)# ipv6 nd suppress-ra

BR-A1(config)# interface eth1
BR-A1(config-if)# ip address 10.0.1.2/30
BR-A1(config-if)# ipv6 address 2001:1::1/127
BR-A1(config-if)# ipv6 nd suppress-ra
BR-A1(config-if)# ipv6 ospf6 network broadcast

BR-A1(config)# interface eth2
BR-A1(config-if)# ip address 10.0.2.2/30
BR-A1(config-if)# ipv6 address 2001:2::2/127
BR-A1(config-if)# ipv6 nd suppress-ra
BR-A1(config-if)# ipv6 ospf6 network broadcast
```

6.2 BGP-4 configuration

The next step involves the creation of a BGP-4 instance for the AS, in this case 1111. The *router-id* is set to the IPv4 loopback address. The two IPv4 public networks to be advertised are specified as well as the neighbour IP addresses. Both the IPv4 and IPv6 addresses are specified here though BGP-4 will carry routing information for both protocols over a single protocol. Finally an instruction to redistribute static routes with a metric of 30. This will be explained further in a moment.

```
BR-A1(config-if)# router bgp 1111
BR-A1(config-router)# bgp router-id 210.11.11.1
BR-A1(config-router)# network 210.0.7.0/24
BR-A1(config-router)# network 210.0.8.0/24
BR-A1(config-router)# neighbor 195.5.5.2 remote-as 2222
BR-A1(config-router)# neighbor 2005:5::1 remote-as 2222
BR-A1(config-router)# redistribute static metric 30
```

For IPv6 it is necessary to configure the IPv6 *address-family*. Specify the two public IPv6 prefixes to be advertised and activate the neighbour already defined to also carry routing information for IPv6. Again redistribute IPv6 static routes with a metric of 30.

```
BR-A1(config-router)# address-family ipv6
BR-A1(config-router-af)# network 2101:7::/64
BR-A1(config-router-af)# network 2101:8::/64
BR-A1(config-router-af)# neighbor 2005:5::1 activate
BR-A1(config-router-af)# redistribute static metric 30
BR-B1(config-router-af)# exit-address-family
BR-A1(config-router)# exit
```

The last step defines static routes for the public networks. Null0 is a virtual interface that allows the routes to be added to the Router Information Base (RIB) with no specific gateway interface. These are advertised through the redistribution of static routes into BGP-4 while the OSPF routes in the RIB are not. When packets arrive at the router it can direct them due to the OPSF routing information in the RIB.

```
BR-A1(config)# ip route 210.0.7.0/24 Null0 250
BR-A1(config)# ip route 210.0.8.0/24 Null0 250
BR-A1(config)# ipv6 route 2101:7::/64 null0 250
BR-A1(config)# ipv6 route 2101:8::/64 null0 250
```

6.3 OSPFv2 configuration

The internal AS IPv4 traffic is routed using OSPFv2 as the IGP. The loopback IPv4 address is used as the router-id. All the connected networks are advertised but while the network on eth0, 195.5.5.0 is advertised the interface is set to passive so it will not try and establish a neighbour state with the BR on the other AS.

```
BR-A1(config)# router ospf
BR-A1(config-router)# ospf router-id 210.11.11.1
BR-A1(config-router)# redistribute bgp metric 30
BR-B1(config-router)# passive-interface eth0
BR-B1(config-router)# network 195.5.5.0/30 area 0.0.0.0
BR-A1(config-router)# network 10.0.1.0/30 area 0.0.0.0
BR-A1(config-router)# network 10.0.2.0/30 area 0.0.0.0
BR-A1(config-router)# network 210.11.11.1/32 area 0.0.0.0
```

6.4 OSPFv3 configuration

This configuration is similar to that for OSPFv2 however the interfaces are added to the default area and not the networks. Another interesting difference is the lack of a passive-interface option. Therefore the interface eth0 has been left out of the default area.

```
BR-A1(config-ospf6)# router ospf6
BR-A1(config-ospf6)# router-id 210.11.11.1
BR-A1(config-ospf6)# redistribute bgp
BR-A1(config-ospf6)# interface lo area 0
BR-A1(config-ospf6)# interface eth1 area 0
BR-A1(config-ospf6)# interface eth2 area 0
```

6.5 Testing the network

Grabbing a trace from the test node A12. This file can be analysed in Wireshark.

```
A12# tshark -i eth0 -w myBGPtrace.pcapng -F pcap
Capturing on 'eth0'
16 ^C

A12# ls
myBGPtrace.pcapng
```

Ping to the public networks on the opposite AS.

A7# fping 220.0.9.20 220.0.9.20 is alive	A10# fping 210.0.7.20 210.0.7.20 is alive
A7# fping6 2201:9::20 2201:9::20 is alive	A10# fping6 2101:7::20 2101:7::20 is alive
A7# fping 220.0.10.20 220.0.10.20 is alive	A10# fping 210.0.8.20 210.0.8.20 is alive
A7# fping6 2201:10::20 2201:10::20 is alive	A10# fping6 2101:8::20 2101:8::20 is alive

7. Router Configurations

7.1 BR-A1

```
BR-A1# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
    ip address 195.5.5.1/30
    ipv6 address 2005:5::1/127
    ipv6 nd suppress-ra
!
interface eth1
    ip address 10.0.1.2/30
    ipv6 address 2001:1::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth2
    ip address 10.0.2.2/30
    ipv6 address 2001:2::2/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface lo
    ip address 210.11.11.1/32
    ipv6 address 2101:11:11::1/128
!
router bgp 1111
    bgp router-id 210.11.11.1
    network 210.0.7.0/24
    network 210.0.8.0/24
    redistribute static metric 30
    neighbor 195.5.5.2 remote-as 2222
    neighbor 2005:5::1 remote-as 2222
!
    address-family ipv6
        network 2101:7::/64
        network 2101:8::/64
        redistribute static metric 30
        neighbor 2005:5::1 activate
    exit-address-family
!
router ospf
    ospf router-id 210.11.11.1
    redistribute bgp metric 30
    passive-interface eth0
    network 10.0.1.0/30 area 0.0.0.0
    network 10.0.2.0/30 area 0.0.0.0
    network 195.5.5.0/30 area 0.0.0.0
    network 210.11.11.1/32 area 0.0.0.0
!
```

7.2 CR-A2

```
CR-A2# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
    ip address 10.0.1.1/30
    ipv6 address 2001:1::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth1
    ip address 10.0.6.2/30
    ipv6 address 2001:6::/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth2
    ip address 210.0.7.1/24
    ipv6 address 2101:7::1/64
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface lo
    ip address 210.11.11.2/32
    ipv6 address 2101:11::2/128
!
router ospf
    ospf router-id 210.11.11.2
    network 10.0.1.0/30 area 0.0.0.0
    network 10.0.6.0/30 area 0.0.0.0
    network 210.0.7.0/24 area 0.0.0.0
    network 210.11.11.2/32 area 0.0.0.0
!
router ospf6
    router-id 210.11.11.2
    interface lo area 0.0.0.0
    interface eth0 area 0.0.0.0
    interface eth1 area 0.0.0.0
    interface eth2 area 0.0.0.0
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

7.3 CR-A3

```
CR-A3# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
    ip address 10.0.2.1/30
    ipv6 address 2001:2::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth1
    ip address 10.0.6.1/30
    ipv6 address 2001:6::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth2
    ip address 210.0.8.1/24
    ipv6 address 2101:8::1/64
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface lo
    ip address 210.11.11.3/32
    ipv6 address 2101:11::3/128
!
router ospf
    ospf router-id 210.11.11.3
    network 10.0.2.0/30 area 0.0.0.0
    network 10.0.6.0/30 area 0.0.0.0
    network 210.0.8.0/24 area 0.0.0.0
    network 210.11.11.3/32 area 0.0.0.0
!
router ospf6
    router-id 210.11.11.3
    interface lo area 0.0.0.0
    interface eth0 area 0.0.0.0
    interface eth1 area 0.0.0.0
    interface eth2 area 0.0.0.0
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

7.4 BR-B1

```

BR-B1# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
  ip address 195.5.5.2/30
  ipv6 address 2005:5::1/127
  ipv6 nd suppress-ra
!
interface eth1
  ip address 10.0.3.2/30
  ipv6 address 2001:3::1/127
  ipv6 nd suppress-ra
  ipv6 ospf6 network broadcast
!
interface eth2
  ip address 10.0.4.1/30
  ipv6 address 2001:4::1/127
  ipv6 nd suppress-ra
  ipv6 ospf6 network broadcast
!
interface lo
  ip address 220.22.22.1/32
  ipv6 address 2201:22::1/128
!
router bgp 2222
  bgp router-id 220.22.22.1
  network 220.0.9.0/24
  network 220.0.10.0/24
  redistribute static metric 30
  neighbor 195.5.5.1 remote-as 1111
  neighbor 2005:5:: remote-as 1111
!
  address-family ipv6
    network 2201:9::/64
    network 2201:10::/64
    redistribute static metric 30
    neighbor 2005:5:: activate
  exit-address-family
!
router ospf
  ospf router-id 220.22.22.1
  redistribute bgp metric 30
  passive-interface eth0
  network 10.0.3.0/30 area 0.0.0.0
  network 10.0.4.0/30 area 0.0.0.0
  network 195.5.5.0/30 area 0.0.0.0
  network 220.22.22.1/32 area 0.0.0.0
!
```

7.5 CR-B2

```
CR-B2# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
    ip address 10.0.4.2/30
    ipv6 address 2001:4::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth1
    ip address 10.0.5.1/30
    ipv6 address 2001:5::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth2
    ip address 220.0.9.1/24
    ipv6 address 2201:9::1/64
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface lo
    ip address 220.22.22.2/32
    ipv6 address 2201:22::2/128
!
router ospf
    ospf router-id 220.22.22.2
    network 10.0.4.0/30 area 0.0.0.0
    network 10.0.5.0/30 area 0.0.0.0
    network 220.0.9.0/24 area 0.0.0.0
    network 220.22.22.2/32 area 0.0.0.0
!
router ospf6
    router-id 220.22.22.2
    interface lo area 0.0.0.0
    interface eth0 area 0.0.0.0
    interface eth1 area 0.0.0.0
    interface eth2 area 0.0.0.0
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

7.6 CR-B3

```
CR-B3# show run
Building configuration...

Current configuration:
!
!
service integrated-vtysh-config
!
debug ospf6 lsa unknown
!
interface eth0
    ip address 10.0.3.1/30
    ipv6 address 2001:3::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth1
    ip address 10.0.5.2/30
    ipv6 address 2001:5::1/127
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface eth2
    ip address 220.0.10.1/24
    ipv6 address 2201:10::1/64
    ipv6 nd suppress-ra
    ipv6 ospf6 network broadcast
!
interface lo
    ip address 220.22.22.3/32
    ipv6 address 2201:22::3/128
!
router ospf
    ospf router-id 220.22.22.3
    network 10.0.3.0/30 area 0.0.0.0
    network 10.0.5.0/30 area 0.0.0.0
    network 220.0.10.0/24 area 0.0.0.0
    network 220.22.22.3/32 area 0.0.0.0
!
router ospf6
    router-id 220.22.22.3
    interface eth0 area 0.0.0.0
    interface eth1 area 0.0.0.0
    interface eth2 area 0.0.0.0
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

8. Host configurations

8.1 Host A7

```
A7.conf# ip addr list dev eth0
2080: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:00:00:aa:00:0d brd ff:ff:ff:ff:ff:ff
        inet 210.0.7.20/24 scope global eth0
            valid_lft forever preferred_lft forever
        inet6 2101:7::20/64 scope global
            valid_lft forever preferred_lft forever
        inet6 fe80::200:ff:feaa:d/64 scope link
            valid_lft forever preferred_lft forever
```

8.2 Host A8

```
A8.conf# ip addr list dev eth0
2085: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:00:00:aa:00:0f brd ff:ff:ff:ff:ff:ff
        inet 210.0.8.20/24 scope global eth0
            valid_lft forever preferred_lft forever
        inet6 2101:8::20/64 scope global
            valid_lft forever preferred_lft forever
        inet6 fe80::200:ff:feaa:f/64 scope link
            valid_lft forever preferred_lft forever
```

8.3 Host A9

```
A9.conf# ip addr list dev eth0
2090: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:00:00:aa:00:11 brd ff:ff:ff:ff:ff:ff
        inet 220.0.9.20/24 scope global eth0
            valid_lft forever preferred_lft forever
        inet6 2201:9::20/64 scope global
            valid_lft forever preferred_lft forever
        inet6 fe80::200:ff:feaa:11/64 scope link
            valid_lft forever preferred_lft forever
```

8.4 Host A10

```
A10.conf# ip addr list dev eth0
2095: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 00:00:00:aa:00:13 brd ff:ff:ff:ff:ff:ff
        inet 220.0.10.20/24 scope global eth0
            valid_lft forever preferred_lft forever
        inet6 2201:10::20/64 scope global
            valid_lft forever preferred_lft forever
        inet6 fe80::200:ff:feaa:13/64 scope link
            valid_lft forever preferred_lft forever
```

This page is intentionally blank