BSc in Telecommunications Engineering

TEL3214

Computer Communication Networks

Lecture 07
Routing – IPv6 with OSPFv3

Eng Diarmuid O'Briain, CEng, CISSP

Department of Electrical and Computer Engineering,
College of Engineering, Design, Art and Technology,
Makerere University

Table of Contents

Illustration Index

This page is intentionally blank

1. Open Shortest Path First version 3 (OSPFv3)

Open Shortest Path First version 3 (OSPFv3) is defined in RFC5340. The fundamental mechanisms of OSPFv2 for IPv4 like flooding, Designated Router (DR) election, area support, Short Path First (SPF) calculations, etc. are identical for OSPFv3. Some change is necessary, either due to changes in protocol semantics between IPv4 and IPv6, or simply to handle the increased address size of IPv6.

Changes between OSPFv2, and OSPFv3 include:

- Addressing semantics have been removed from OSPF packets and the basic Link State Advertisements (LSAs).
- New LSAs have been created to carry IPv6 addresses and prefixes.
- OSPF now runs on a per-link basis rather than on a per-IP-subnet basis.
- Flooding scope for LSAs has been generalised.
- Authentication has been removed from the OSPFv3 protocol and instead relies on IPv6's underlying Authentication Header (AH) and Encapsulating Security Payload (ESP).

Even with larger IPv6 addresses, most packets in OSPFv3 are almost as compact as those in OSPFv2.

1.1 Removal of address semantics

IPv6 addresses are not present in OSPFv3 packets, except in LSA payloads carried by the Link State Update packets. Router-LSAs and network-LSAs no longer contain network addresses, but simply express topology information.

OSPF Router IDs, Area IDs, and LSA Link State IDs remain at the IPv4 size of 32 bits.

Neighboring routers are now always identified by Router ID. In OSPFv4 they had been identified by an IPv4 address on broadcast, Non-Broadcast Multi-Access (NBMA), and point-to-multipoint links.

1.2 Authentication Changes

In OSPFv3, authentication has been removed from the OSPFv3 protocol. The "AuType" and "Authentication" fields have been removed from the OSPFv3 packet header, and all authentication-related fields have been removed from the OSPFv3 area and interface data structures.

When running over IPv6, OSPFv3 relies on the IPv6 AH and ESP headers to ensure integrity and authentication/confidentiality of routing exchanges.

2. Configuring IPv6

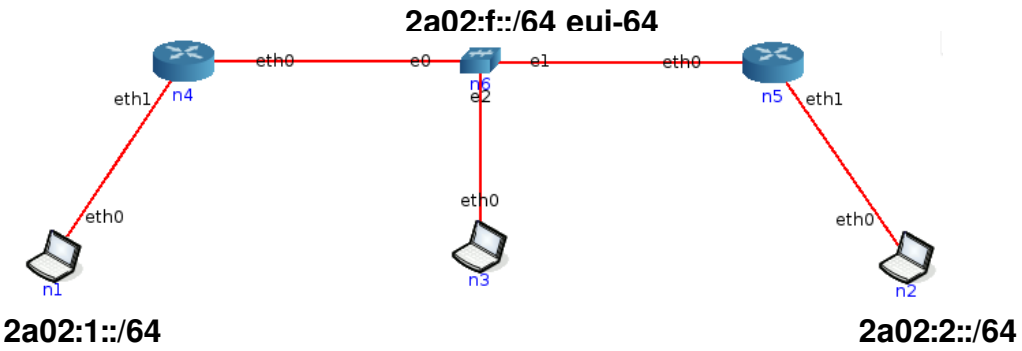

Illustration 1: IPv6 routed network

2.1 Configuring Router n4

Some of the configuration is identical to the IPv4 and some have subtle differences that need to be watched for. Follow along with the explanations for this router.

```
n4# conf t
n4(config)# hostname RTR_n4
```

IPv6 forwarding must be specifically enabled, IPv4 is enabled by default.

```
RTR_n4(config)# ipv6 forwarding
```

Interface ***eth0*** forms one side of the link between the routers. There is no need to assign a specific IP address so a prefix is given as an address and the router will form its own global scope address using its EUI-64 Medium Access Control (MAC) address. As no host will require a prefix from this interface RA messages can be suppressed. As this is effectively a Point to Point (P2P) link Open Shortest Path First version 3 (OSPFv3) can be informed that as in a P2P scenario an election process is unnecessary and OSPFv3 control messages can be reduced.

```
RTR_n4(config-if)# int eth0
RTR_n4(config-if)# ipv6 address 2a02:f::/64
RTR_n4(config-if)# ipv6 nd suppress-ra
RTR_n4(config-if)# ipv6 ospf6 network point-to-point
```

For this router the ***eth1*** interfaces is the edge network, the computers will expect to get IPv6 prefix from the router so they can carry out the SLAAC process to make a global scope address. The first command specifies a specific address for the router interface. The second command is a prefix for the Network Discovery (ND) process with a valid lifetime and preferred lifetime in seconds. The reachable-time is the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred. It is possible to replace these numbers with *infinite*. As the hosts on this link want to be given a prefix Router Advertisement (RA) messages must not be suppressed.

```
RTR_n4(config)# int eth1
RTR_n4(config-if)# ipv6 address 2a02:1::1/64
RTR_n4(config-if)# ipv6 nd prefix 2a02:1::/64 86400 86400
RTR_n4(config-if)# ipv6 nd reachable-time 5000
RTR_n4(config-if)# no ipv6 nd suppress-ra
```

From an OSPFv3 perspective this is a broadcast network and the interface is set to passive, meaning that no adjacency will be formed on this interface.

```
RTR_n4(config-if)# ipv6 ospf6 network broadcast
RTR_n4(config-if)# ipv6 ospf6 passive
```

Interface ***lo*** is configured with an IPv4 and an IPv6 address. In truth the IPv6 address is unnecessary in this configuration as it is so simple. The IPv4 address is used as the ID for OSPFv3 which uses the 4 octet notation.

```
RTR_n4(config-if)# int lo
RTR_n4(config-if)# ip address 10.0.0.1/32
RTR_n4(config-if)# ipv6 address 2a00:0::1/128
```

2.2 Configuring Router n5

Configure Router n5 is a similar fashion to R4. Note the abbreviations from Router n4.

```
n5# conf t
n5(config)# hostname RTR_n5
RTR_n5(config)# ipv6 forwarding

RTR_n5(config)# int eth0
RTR_n5(config-if)# ipv6 a 2a02:f::/64
RTR_n5(config-if)# ipv6 nd s
RTR_n5(config-if)# ipv6 ospf6 n p

RTR_n5(config-if)# int eth1
RTR_n5(config-if)# ipv6 a 2a02:2::1/64
RTR_n5(config-if)# ipv6 nd p 2a02:2::/64 86400 86400
RTR_n5(config-if)# ipv6 nd re 5000
RTR_n5(config-if)# ipv6 ospf6 n b
RTR_n5(config-if)# ipv6 ospf6 pa
RTR_n5(config-if)# no ipv6 nd s

RTR_n5(config-if)# int lo
RTR_n5(config-if)# ip a 10.0.0.2/32
RTR_n5(config-if)# ipv6 a 2a00:0::2/128
```

2.3 Testing the ND SLAAC process

Check the IPv6 addresses on Hosts n1 and n2. They should have built a global scope IPv6 address from assigned prefixes from the routers. See that this has indeed happened.

Host n1

```
root@n1:/tmp/pycore.52704/n1.conf# ip -6 addr list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
73: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qlen 1000
    inet6 2a02:1::200:ff:feaa:0/64 scope global mngtmpaddr dynamic
       valid_lft 86055sec preferred_lft 86055sec
    inet6 fe80::200:ff:feaa:0/64 scope link
       valid_lft forever preferred_lft forever
```

Host n2

```
root@n2:/tmp/pycore.52704/n2.conf# ip -6 addr list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
84: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qlen 1000
    inet6 2a02:2::200:ff:feaa:5/64 scope global mngtmpaddr dynamic
       valid_lft 86088sec preferred_lft 86088sec
    inet6 fe80::200:ff:feaa:5/64 scope link
       valid_lft forever preferred_lft forever
```

2.4 OSPFv3

OSP

2.5 OSPFv3 configuration

So to the routing. For this OSPFv3 is employed. The configuration is different that that for OSPFv2. In this case assign a router ID and specify the participating interfaces and to which area they belong. In this case they all belong to the backbone area.

Router n4

```
RTR_n4(config-if)# router ospf6
RTR_n4(config-router)# router-id 10.0.0.1
RTR_n4(config-router)# interface eth0 area 0.0.0.0
RTR_n4(config-router)# interface eth1 area 0.0.0.0
```

Router n5

```
RTR_n5(config-if)# router ospf6
RTR_n5(config-router)# router-id 10.0.0.2
RTR_n5(config-router)# int eth0 a 0.0.0.0
RTR_n5(config-router)# int eth1 a 0.0.0.0
```

2.6 Reviewing the configurations

2.6.1 Router n4

```
RTR_n4# show ipv6 ospf6 neighbor

Neighbor ID     Pri    DeadTime  State/IfState         Duration I/F[State]
10.0.0.2          1    00:00:37   Full/DROther         00:52:20 eth0[PointToPoint]



RTR_n4# show ipv6 ospf6 database

        Area Scoped Link State Database (Area 0.0.0.0)

Type LSId           AdvRouter       Age   SeqNum                       Payload
Rtr  0.0.0.0        0.0.0.0        1702 8000000b             10.0.0.2/0.0.0.79
Rtr  0.0.0.0        10.0.0.1       1368 80000003             10.0.0.2/0.0.0.79
Rtr  0.0.0.0        10.0.0.2       1339 8000000b             10.0.0.1/0.0.0.77
INP  0.0.0.0        0.0.0.0        1708 80000005                   2a02:1::/64
INP  0.0.0.0        0.0.0.0        1708 80000005                   2a02:f::/64
INP  0.0.0.0        10.0.0.1       1371 80000002                   2a02:1::/64
INP  0.0.0.0        10.0.0.1       1371 80000002                   2a02:f::/64
INP  0.0.0.0        10.0.0.2        161 80000004                   2a02:2::/64
INP  0.0.0.0        10.0.0.2        161 80000004                   2a02:f::/64

        I/F Scoped Link State Database (I/F eth0 in Area 0.0.0.0)

Type LSId           AdvRouter       Age   SeqNum                       Payload
Lnk  0.0.0.75       0.0.0.0         871 80000003           fe80::200:ff:feaa:1

        I/F Scoped Link State Database (I/F eth1 in Area 0.0.0.0)

Type LSId           AdvRouter       Age   SeqNum                       Payload
Lnk  0.0.0.77       0.0.0.0         867 80000003           fe80::200:ff:feaa:2
Lnk  0.0.0.79       10.0.0.2        162 80000005           fe80::200:ff:feaa:3

        AS Scoped Link State Database

Type LSId           AdvRouter       Age   SeqNum                       Payload
```

```
RTR_n4# show ipv6 ospf6 route
*N IA 2a02:1::/64                   ::                          eth0 00:53:37
*N IA 2a02:2::/64                   fe80::200:ff:feaa:3         eth1 00:53:33
*N IA 2a02:f::/64                   ::                          eth1 00:53:37
 N IA 2a02:f::/64                   fe80::200:ff:feaa:3         eth1 00:53:33


RTR_n4# sh ipv6 route
Codes: K - kernel route, C - connected, S - static, R - RIPng,
       O - OSPFv6, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* ::1/128 is directly connected, lo
C>* 2a00::1/128 is directly connected, lo
O   2a02:1::/64 [110/10] is directly connected, eth0, 00:54:05
C>* 2a02:1::/64 is directly connected, eth0
O>* 2a02:2::/64 [110/20] via fe80::200:ff:feaa:3, eth1, 00:54:02
O   2a02:f::/64 [110/10] is directly connected, eth1, 00:54:05
C>* 2a02:f::/64 is directly connected, eth1
C * fe80::/64 is directly connected, eth1
C>* fe80::/64 is directly connected, eth0
```

2.6.2 Router n5

```
RTR_n5# show ipv6 ospf6 neighbor
Neighbor ID     Pri    DeadTime  State/IfState         Duration I/F[State]
10.0.0.1          1    00:00:36   Full/DROther        00:55:55 eth0[PointToPoint]
```

```
RTR_n5# show ipv6 ospf6 database

        Area Scoped Link State Database (Area 0.0.0.0)

Type LSId            AdvRouter       Age   SeqNum                       Payload
Rtr  0.0.0.0         0.0.0.0         108 8000000c             10.0.0.2/0.0.0.79
Rtr  0.0.0.0         10.0.0.1       1574 80000003             10.0.0.2/0.0.0.79
Rtr  0.0.0.0         10.0.0.2       1543 8000000b             10.0.0.1/0.0.0.77
INP  0.0.0.0         0.0.0.0         114 80000006                   2a02:1::/64
INP  0.0.0.0         0.0.0.0         114 80000006                   2a02:f::/64
INP  0.0.0.0         10.0.0.1       1577 80000002                   2a02:1::/64
INP  0.0.0.0         10.0.0.1       1577 80000002                   2a02:f::/64
INP  0.0.0.0         10.0.0.2        365 80000004                   2a02:2::/64
INP  0.0.0.0         10.0.0.2        365 80000004                   2a02:f::/64

        I/F Scoped Link State Database (I/F eth0 in Area 0.0.0.0)

Type LSId            AdvRouter       Age   SeqNum                       Payload
Lnk  0.0.0.77        0.0.0.0        1073 80000003           fe80::200:ff:feaa:2
Lnk  0.0.0.79        10.0.0.2        366 80000005           fe80::200:ff:feaa:3

        I/F Scoped Link State Database (I/F eth1 in Area 0.0.0.0)

Type LSId            AdvRouter       Age   SeqNum                       Payload
Lnk  0.0.0.82        0.0.0.0         623 80000004           fe80::200:ff:feaa:4
Lnk  0.0.0.82        10.0.0.2        365 80000005           fe80::200:ff:feaa:4

        AS Scoped Link State Database

Type LSId            AdvRouter       Age   SeqNum                       Payload


RTR_n5# show ipv6 ospf6 route
*N IA 2a02:1::/64                   ::                             0 00:56:35
*N IA 2a02:2::/64                   ::                          eth1 01:06:33
*N IA 2a02:f::/64                   ::                          eth0 01:06:33
 N IA 2a02:f::/64                   ::                             0 00:56:35


RTR_n5# sh ipv6 route
Codes: K - kernel route, C - connected, S - static, R - RIPng,
       O - OSPFv6, I - IS-IS, B - BGP, A - Babel,
       > - selected route, * - FIB route

C>* ::1/128 is directly connected, lo
C>* 2a00::2/128 is directly connected, lo
O>* 2a02:1::/64 [110/20] via fe80::200:ff:feaa:2, eth0, 01:02:28
O   2a02:2::/64 [110/10] is directly connected, eth1, 01:06:52
C>* 2a02:2::/64 is directly connected, eth1
O   2a02:f::/64 [110/10] is directly connected, eth0, 01:06:52
C>* 2a02:f::/64 is directly connected, eth0
C * fe80::/64 is directly connected, eth1
C>* fe80::/64 is directly connected, eth0
```

2.7 Reviewing traffic on the wire at Hub n6

ICMPv6 Neighbour Solicitation for 2a02:f::

```
Frame: 78 bytes on wire (624 bits)on interface 0
Ethernet II, Src: 00:00:00_aa:00:02, Dst: 33:33:ff:00:00:00
Internet Protocol Version 6, Src: :: (::), Dst: ff02::1:ff00:0
Internet Control Message Protocol v6
    Type: Neighbor Solicitation (135)
    Code: 0
    Checksum: 0x5097 [correct]
    Reserved: 00000000
    Target Address: 2a02:f::
```

OSPFv3 Hello packet.

```
Frame: 94 bytes on wire (752 bits) on interface 0
Ethernet II, Src: 00:00:00_aa:00:02, Dst: 33:33:00:00:00:05
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::5
Open Shortest Path First
    OSPF Header
        Version: 3
        Message Type: Hello Packet (1)
        Packet Length: 40
        Source OSPF Router: 10.0.0.1
        Area ID: 0.0.0.0 (Backbone)
        Checksum: 0xe78b [correct]
        Instance ID: IPv6 unicast AF (0)
        Reserved: 00
    OSPF Hello Packet
        Interface ID: 77
        Router Priority: 1
        Options: 0x000013 (R, E, V6)
        Hello Interval [sec]: 10
        Router Dead Interval [sec]: 40
        Designated Router: 0.0.0.0
        Backup Designated Router: 0.0.0.0
        Active Neighbor: 10.0.0.2
```

OSPFv3 Database Description

```
Frame: 82 bytes on wire (656 bits) on interface 0
Ethernet II, Src: 00:00:00_aa:00:02, Dst: 33:33:00:00:00:05
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::5
Open Shortest Path First
    OSPF Header
        Version: 3
        Message Type: DB Description (2)
        Packet Length: 28
        Source OSPF Router: 10.0.0.1 (10.0.0.1)
        Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
        Checksum: 0xb8e2 [correct]
        Instance ID: IPv6 unicast AF (0)
        Reserved: 00
    OSPF DB Description
    LSA Header
        LS Age: 450 seconds
        Do Not Age: False
        LS Type: Link-LSA (0x0008)
        Link State ID: 0.0.0.77
        Advertising Router: 0.0.0.0 (0.0.0.0)
        LS Sequence Number: 0x80000004
        LS Checksum: 0xab84
        Length: 56
    LSA Header
        LS Age: 1543 seconds
        Do Not Age: False
        LS Type: Link-LSA (0x0008)
        Link State ID: 0.0.0.79
        Advertising Router: 10.0.0.2 (10.0.0.2)
        LS Sequence Number: 0x80000005
        LS Checksum: 0x45da
        Length: 56

     ~~~~~ more ~~~~
```

OSPFv3 Link State Request.

```
Frame 12: 106 bytes on wire (848 bits) on interface 0
Ethernet II, Src: 00:00:00_aa:00:03, Dst: 33:33:00:00:00:05
Internet Protocol Version 6, Src: fe80::200:ff:feaa:3, Dst: ff02::5
Open Shortest Path First
    OSPF Header
        Version: 3
        Message Type: LS Request (3)
        Packet Length: 52
        Source OSPF Router: 10.0.0.2
        Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
        Checksum: 0x94a1 [correct]
        Instance ID: IPv6 unicast AF (0)
        Reserved: 00
    Link State Request
        Reserved: 0
        LS Type: Link-LSA (0x0008)
        Link State ID: 0.0.0.77
        Advertising Router: 10.0.0.1
    Link State Request
        Reserved: 0
        LS Type: Router-LSA (0x2001)
        Link State ID: 0.0.0.0
        Advertising Router: 10.0.0.1
    Link State Request
        Reserved: 0
        LS Type: Intra-Area-Prefix-LSA (0x2009)
        Link State ID: 0.0.0.0
        Advertising Router: 10.0.0.1
```

OSPFv3 Link State Update.

```
Frame: 210 bytes on wire (1680 bits) on interface 0
Ethernet II, Src: 00:00:00_aa:00:02, Dst: 33:33:00:00:00:05
Internet Protocol Version 6, Src: fe80::200:ff:feaa:2, Dst: ff02::5 (ff02::5)
Open Shortest Path First
    OSPF Header
        Version: 3
        Message Type: LS Update (4)
        Packet Length: 156
        Source OSPF Router: 10.0.0.1
        Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
        Checksum: 0x1ef2 [correct]
        Instance ID: IPv6 unicast AF (0)
        Reserved: 00
    LS Update Packet
        Number of LSAs: 3
        Link-LSA (Type: 0x0008)
            LS Age: 1 seconds
            Do Not Age: False
            LS Type: Link-LSA (0x0008)
            Link State ID: 0.0.0.77
            Advertising Router: 10.0.0.1
            LS Sequence Number: 0x80000002
            LS Checksum: 0x4fd7
            Length: 56
            Router Priority: 1
            Options: 0x000013 (R, E, V6)
            Link-local Interface Address: fe80::200:ff:feaa:2
            # prefixes: 1
            PrefixLength: 64
            PrefixOptions: 0x00
            Reserved: 0
            Address Prefix: 2a02:f::
        Router-LSA (Type: 0x2001)
            LS Age: 17 seconds
            Do Not Age: False
            LS Type: Router-LSA (0x2001)
            Link State ID: 0.0.0.0
            Advertising Router: 10.0.0.1
            LS Sequence Number: 0x80000005
            LS Checksum: 0xc55d
            Length: 24
            Flags: 0x00
            Options: 0x000013 (R, E, V6)
        Intra-Area-Prefix-LSA (Type: 0x2009)
            LS Age: 1 seconds
            Do Not Age: False
            LS Type: Intra-Area-Prefix-LSA (0x2009)
            Link State ID: 0.0.0.0
            Advertising Router: 10.0.0.1
            LS Sequence Number: 0x80000005
            LS Checksum: 0x756d
            Length: 56
            # prefixes: 2
            Referenced LS type 0x2001 (Router-LSA)
            Referenced Link State ID: 0.0.0.0
            Referenced Advertising Router: 10.0.0.1
            PrefixLength: 64
            PrefixOptions: 0x00
            Metric: 10
            Address Prefix: 2a02:1::
            PrefixLength: 64
            PrefixOptions: 0x00
            Metric: 10
            Address Prefix: 2a02:f::
```

2.8 Testing the IPv6 network

Using ping6 and traceroute6 to test from Host n1 to n2 demonstrated that the network is functional.

```
root@n1:/tmp/pycore.52704/n1.conf# ping6 -c1 2a02:2::200:ff:feaa:5
PING 2a02:2::200:ff:feaa:5(2a02:2::200:ff:feaa:5) 56 data bytes
64 bytes from 2a02:2::200:ff:feaa:5: icmp_seq=1 ttl=62 time=0.066 ms

--- 2a02:2::200:ff:feaa:5 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.066/0.066/0.066/0.000 ms


root@n1:/tmp/pycore.52704/n1.conf# traceroute6 2a02:2::200:ff:feaa:5
traceroute to 2a02:2::200:ff:feaa:5 (2a02:2::200:ff:feaa:5), 30 hops max, 80 byte
packets
 1  2a02:1::1 (2a02:1::1)  0.029 ms  0.007 ms  0.006 ms
 2  2a02:2::1 (2a02:2::1)  0.021 ms  0.011 ms  0.011 ms
 3  2a02:2::200:ff:feaa:5 (2a02:2::200:ff:feaa:5)  0.018 ms  0.012 ms  0.012 ms
```

2.9 Summary of the configuration in both routers

These lists can be copy and pasted into the respective vtysh terminals.

Router n4

```
conf t
hostname RTR_n4
ipv6 forwarding
int eth0
ipv6 address 2a02:1::1/64
ipv6 nd prefix 2a02:1::/64 86400 86400
ipv6 nd reachable-time 5000
ipv6 ospf6 network broadcast
ipv6 ospf6 passive
no ipv6 nd suppress-ra
int eth1
ipv6 address 2a02:f::/64
ipv6 nd suppress-ra
ipv6 ospf6 network point-to-point
int lo
ipv6 address 2a00:0::1/128
ip address 10.0.0.1/32
router ospf6
router-id 10.0.0.1
interface eth0 area 0.0.0.0
interface eth1 area 0.0.0.0
```

Router n5

```
conf t
hostname RTR_n5
ipv6 forwarding
int eth1
ipv6 address 2a02:2::1/64
ipv6 nd prefix 2a02:2::/64 86400 86400
ipv6 nd reachable-time 5000
ipv6 ospf6 network broadcast
ipv6 ospf6 passive
no ipv6 nd suppress-ra
int eth0
ipv6 address 2a02:f::/64
ipv6 nd suppress-ra
ipv6 ospf6 network point-to-point
int lo
ipv6 address 2a00:0::2/128
ip address 10.0.0.2/32
router ospf6
router-id 10.0.0.2
interface eth0 area 0.0.0.0
interface eth1 area 0.0.0.0
```