OSPF for IPv6

ISP Workshops

Last updated 17th May 2014

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged and this notice remains in place
- Bug fixes and improvements are welcomed
  - Please email *workshop (at) bgp4all.com*

Philip Smith

Recap: OSPFv2

- April 1998 was the most recent revision (RFC 2328)
- OSPF uses a 2-level hierarchical model
- SPF calculation is performed independently for each area
- Typically faster convergence than DVRPs
- Relatively low, steady state bandwidth requirements

OSPFv3 overview

- OSPF for IPv6
- Based on OSPFv2, with enhancements
- Distributes IPv6 unicast prefixes
- Runs directly over IPv6
- Ships-in-the-night with OSPFv2
- OSPFv3 does **not** carry IPv4 prefixes
  - RFC5838 proposes an extension which adds address family support

OSPFv3 / OSPFv2 Similarities

- Basic packet types
  - Hello, DBD, LSR, LSU, LSA
- Mechanisms for neighbor discovery and adjacency formation
- Interface types
  - P2P, P2MP, Broadcast, NBMA, Virtual
- LSA flooding and aging
- Nearly identical LSA types

OSPFv3 / OSPFv2 Differences

OSPFv3 runs on a Link instead of per IP Subnet

- A link by definition is a medium over which two nodes can communicate at link layer
- In IPv6 multiple IP subnet can be assigned to a link and two nodes in different subnet can communicate at link layer therefore OSPFv3 is running per link instead of per IP subnet
- An Interface connect to a link and multiple interface can be connected to a link

OSPFv3 / OSPFv2 Differences

Support of Multiple Instance per Link

- New field (instance) in OSPF packet header allow running multiple instance per link
- Instance ID should match before packet being accepted
- Useful for traffic separation, multiple areas per link and address families (RFC5838)

OSPFv3 / OSPFv2 Differences

Address Semantic Change in LSA

- Router and Network LSA carry only topology information
- Router LSA can be split across multiple LSAs; Link State ID in LSA header is a fragment ID
- Intra area prefix are carried in a new LSA payload called intra-area-prefix-LSAs
- Prefix are carried in payload of inter-area and external LSA

OSPFv3 / OSPFv2 Differences

Generalisation of Flooding Scope

- In OSPFv3 there are three flooding scope for LSAs (link-local scope, area scope, AS scope) and they are coded in LS type explicitly
- In OSPFv2 initially only area and AS wide flooding was defined; later opaque LSAs introduced link local scope as well

OSPFv3 / OSPFv2 Differences

Explicit Handling of Unknown LSA

- The handling of unknown LSA is coded via U-bit in LS type
- When U bit is set, the LSA is flooded with the corresponding flooding scope, as if it was understood
- When U bit is clear, the LSA is flooded with link local scope
- In v2 unknown LSA were discarded

OSPFv3 / OSPFv2 Differences

Authentication is Removed from OSPF

- Authentication in OSPFv3 has been removed
- OSPFv3 relies now on the IPv6 authentication header since OSPFv3 run over IPv6
- Autype and Authentication field in the OSPF packet header therefore have been suppressed

OSPFv3 / OSPFv2 Differences

OSPF Packet format has been changed

- The mask field has been removed from Hello packet
- IPv6 prefix are only present in payload of Link State update packet

OSPFv3 / OSPFv2 Differences

Two New LSAs Have Been Introduced

- Link-LSA has a link local flooding scope and has three purposes:
  - The router link local address
  - List all IPv6 prefixes attached to the link
  - Assert a collection of option bit for the Router-LSA
- Intra-area-prefix-LSA
  - Used to advertise router's IPv6 address within the area

Inter-Area Prefix LSA

- Describes the destination outside the area but still in the AS
- Summary is created for one area, which is flooded out in all other areas
- Originated by an ABR
- Only intra-area routes are advertised into the backbone
- Link State ID simply serves to distinguish inter-area-prefix-LSAs originated by the same router
- Link-local addresses must never be advertised in inter-area-prefix-LSAs

LSA Types

	LSA Function Code	LSA Type
Router-LSA	1	0x2001
Network-LSA	2	0x2002
Inter-Area-Prefix-LSA	3	0x2003
Inter-Area-Router-LSA	4	0x2004
AS-External-LSA	5	0x4005
Deprecated	6	0x2006
NSSA-LSA	7	0x2007
Link-LSA	8	0x0008
Intra-Area-Prefix-LSA	9	0x2009

NEW LSAs

Configuring OSPFv3 in Cisco IOS

- Similar to OSPFv2
  - Prefixing existing Interface and Exec mode commands with "**`ipv6`**"
- Interfaces configured directly
  - Replaces **`network`** command
  - (Also available in OSPFv2 from IOS 12.4)
- "Native" IPv6 router mode
  - Not a sub-mode of **`router ospf`**

Configuring OSPFv3

- Setting up the OSPFv3 process:

```
[no] ipv6 router ospf <process ID>
```

- Applying the OSPFv3 process to an interface:

```
interface <router-int-name>
 [no] ipv6 ospf <process ID> area <area ID>
```

- Configuring summarisation:

```
ipv6 router ospf <process ID>
 [no] area <area ID> range <prefix>/<length>
```

OSPFv3 exec mode commands

- Exec mode commands:

```
show ipv6 ospf [<process ID>]
clear ipv6 ospf [<process ID>]
```

- Showing new LSA:

```
show ipv6 ospf [<process ID>] database link
show ipv6 ospf [<process ID>] database prefix
```

OSPFv3 Authentication

- Configuring authentication per area:
  - SPI value has to be unique per area:

```
ipv6 router ospf <process ID>
 area 0 authentication ipsec spi 256 md5 <passwd>
```

- Disabling authentication on a specific link when area authentication is activated:

```
interface fastethernet 0/0
 ipv6 ospf authentication null
```

- Configuring authentication per interface:
  - SPI value has to be unique per link:

```
interface fastethernet 0/0
 ipv6 ospf authentication ipsec spi 256 md5 <passwd>
```

OSPFv3 Debug Commands

- Adjacency is not appearing

```
[no] debug ipv6 ospf adj
[no] debug ipv6 ospf hello
```

- SPF is running constantly

```
[no] debug ipv6 ospf spf
[no] debug ipv6 ospf flooding
[no] debug ipv6 ospf events
[no] debug ipv6 ospf lsa-generation
[no] debug ipv6 ospf database-timer
```

- General purpose

```
[no] debug ipv6 ospf packets
[no] debug ipv6 ospf retransmission
[no] debug ipv6 ospf tree
```

OSPFv3 Configuration Example

```
Router1#
interface POS1/1
 ipv6 address 2001:db8:FFFF:1::1/64
 ipv6 ospf 100 area 0
!
interface POS2/0
 ipv6 address 2001:db8:1:1::2/64
 ipv6 ospf 100 area 1
!
 ipv6 router ospf 100
  log-adjacency-changes
!

Router2#
interface POS3/0
 ipv6 address 2001:db8:1:1::1/64
 ipv6 ospf 100 area 1
!
ipv6 router ospf 100
 log-adjacency-changes
```

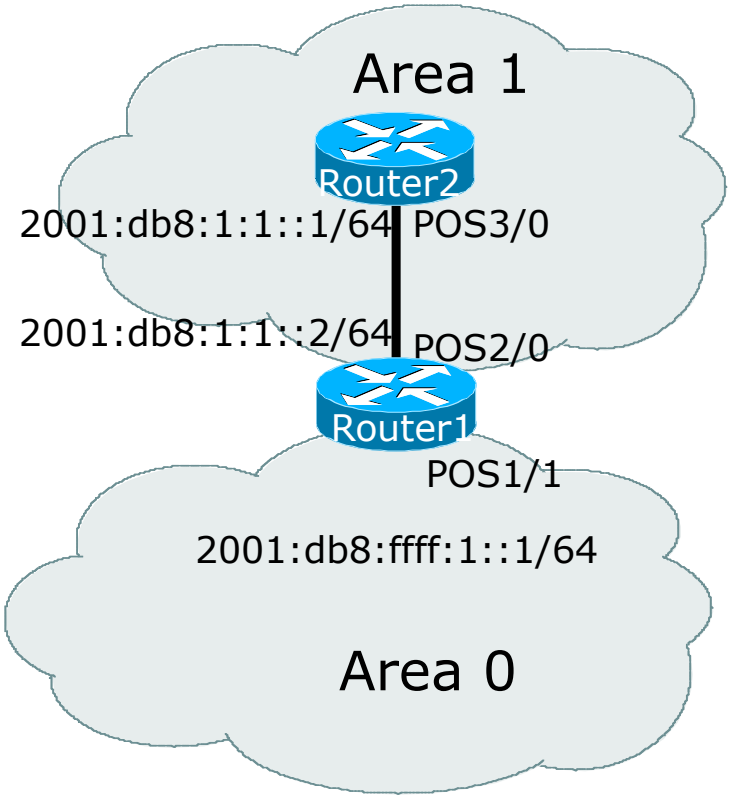

OSPFv3 Interface Status

```
Router2#sh ipv6 ospf int pos 3/0
POS3/0 is up, line protocol is up
  Link Local Address FE80::290:86FF:FE5D:A000, Interface ID 7
  Area 1, Process ID 100, Instance ID 0, Router ID 10.1.1.4
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 3, maximum is 3
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.1.1.3
  Suppress hello for 0 neighbor(s)
```

OSPFv3 Neighbour Status

```
Router2#sh ipv6 ospf neighbor detail
 Neighbor 10.1.1.3
    In the area 1 via interface POS3/0
    Neighbor: interface-id 8, link-local address FE80::2D0:FFFF:FE60:DFFF
    Neighbor priority is 1, State is FULL, 12 state changes
    Options is 0x630C34B9
    Dead timer due in 00:00:33
    Neighbor is up for 00:49:32
    Index 1/1/1, retransmission queue length 0, number of retransmission 1
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 2, maximum is 2
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

OSPFv3 entries in Routing Table

```
Router2#sh ipv6 route
IPv6 Routing Table - 5 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
OI  2001:db8:FFFF:1::/64 [110/2]
     via FE80::2D0:FFFF:FE60:DFFF, POS3/0
C   2001:db8:1:1::/64 [0/0]
     via ::, POS3/0
L   2001:db8:1:1::1/128 [0/0]
     via ::, POS3/0
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
```

OSPFv3 link troubleshooting

- Next router address in OSPFv3 is a link-local address

```
OI  2001:db8:FFFF:1::/64 [110/2]
     via FE80::2D0:FFFF:FE60:DFFF, POS3/0
```

- How to troubleshoot??
  - SSH to neighbouring router needs extended SSH command, for example:

```
ssh FE80::2D0:FFFF:FE60:DFFF /source-int POS3/0
```

  - Source interface has to be specified – a router with multiple interfaces has no idea which interface the remote link local address is attached to

Cisco IOS OSPFv3 Database Display

```
Router2# show ipv6 ospf database

OSPF Router with ID (3.3.3.3) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
0               1.1.1.1         2009        0x8000000A 0x2DB1   1
0               3.3.3.3         501         0x80000007 0xF3E6   1

                Net Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum
7               1.1.1.1         480         0x80000006 0x3BAD

                Inter Area Prefix Link States (Area 0)

ADV Router      Age         Seq#        Prefix
1.1.1.1         1761        0x80000005  2001:db8:2:2::/64
1.1.1.1         982         0x80000005  2001:db8:2:4::2/128

                Link (Type-8) Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Interface
11              3.3.3.3         245         0x80000006 0xF3DC   Lo0
7               1.1.1.1         236         0x80000008 0x68F    Fa2/0
7               3.3.3.3         501         0x80000008 0xE7BC   Fa2/0

                Intra Area Prefix Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Ref lstype
0               1.1.1.1         480         0x80000008 0xD670   0x2001
107             1.1.1.1         236         0x80000008 0xC05F   0x2002
0               3.3.3.3         245         0x80000006 0x3FF7   0x2001
```

Cisco IOS OSPFv3 Detailed LSA Display

```
show ipv6 ospf 1 database inter-area prefix

  LS age: 1714
  LS Type: Inter Area Prefix Links
  Link State ID: 0
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000006
  Checksum: 0x25A0
  Length: 36
  Metric: 1
  Prefix Address: 2001:db8:2:2::
  Prefix Length: 64, Options: None

 show ipv6 ospf 1 database link

  LS age: 283
  Options: (IPv6 Router, Transit Router, E-Bit, No Type 7-to-5, DC)
  LS Type: Link-LSA (Interface: Loopback0)
  Link State ID: 11 (Interface ID)
  Advertising Router: 3.3.3.3
  LS Seq Number: 80000007
  Checksum: 0xF1DD
  Length: 60
  Router Priority: 1
  Link Local Address: FE80::205:5FFF:FEAC:1808
  Number of Prefixes: 2
  Prefix Address: 2001:db8:1:3::
  Prefix Length: 64, Options: None
  Prefix Address: 2001:db8:1:3::
  Prefix Length: 64, Options: None
```

Conclusion

- Based on existing OSPFv2 implementation
- Similar CLI and functionality

OSPF for IPv6

ISP Workshops