Migrating from OSPF to IS-IS

Philip Smith
philip@apnic.net

MyNOG 3

25th – 29th November 2013

Kuala Lumpur

Presentation slides

- Available at:
  - http://thyme.apnic.net/ftp/seminars/MyNOG3-OSPF-to-ISIS-migration.pdf
  - And on the MyNOG 3 website
- Feel free to ask questions any time

Introduction

- With the advent of IPv6 and dual stack networks, more ISPs expressing interest to migrate to IS-IS
  - This is not as difficult as it sounds
- Presentation describes:
  - The differences between OSPF and IS-IS
  - The migration process
    - Based on several successful migrations
    - Uses Cisco IOS CLI as examples

IS-IS

- Intermediate System to Intermediate System
- ISO 10589 specifies OSI IS-IS routing protocol for CLNS traffic
  - A Link State protocol with a 2 level hierarchical architecture
  - Type/Length/Value (TLV) options to enhance the protocol
- RFC 1195 added IP support
  - Integrated IS-IS
  - I/IS-IS runs on top of the Data Link Layer

OSPF

- Open Shortest Path First
- Open:
  - Meaning an Open Standard
  - Developed by IETF (OSPF Working Group) for IP – RFC1247
  - Current standard is OSPFv2 (RFC2328)
- Shortest Path First:
  - Edsger Dijkstra's algorithm for producing shortest path tree through a graph
    - Dijkstra, E. W. (1959). "A note on two problems in connexion with graphs". *Numerische Mathematik* **1**: 269–271

IS-IS & OSPF: Similarities

- Both are Interior Gateway Protocols (IGP)
  - They distribute internal reachability information between routers belonging to a single Autonomous System (AS)
- With support for:
  - IPv4 and IPv6
  - Authentication
  - Multi-path
  - Unnumbered links

IS-IS and OSPF Terminology

OSPF	IS-IS
Host	End System (ES)
Router	Intermediate System (IS)
Link	Circuit
Packet	Protocol Data Unit (PDU)
Designated router (DR)	Designated IS (DIS)
Backup DR (BDR)	N/A (no BDIS is used)
Link-State Advertisement (LSA)	Link-State PDU (LSP)
Hello packet	IIH PDU
Database Description (DBD)	Complete sequence number PDU (CSNP)

IS-IS and OSPF Terminology (Cont.)

OSPF	IS-IS
Area	Sub domain (area)
Non-backbone area	Level-1 area
Backbone area	Level-2 Sub domain (backbone)
Area Border Router (ABR)	L1L2 router
Autonomous System Boundary Router (ASBR)	Any IS

Transport

- OSPF uses IP Protocol 89 as transport

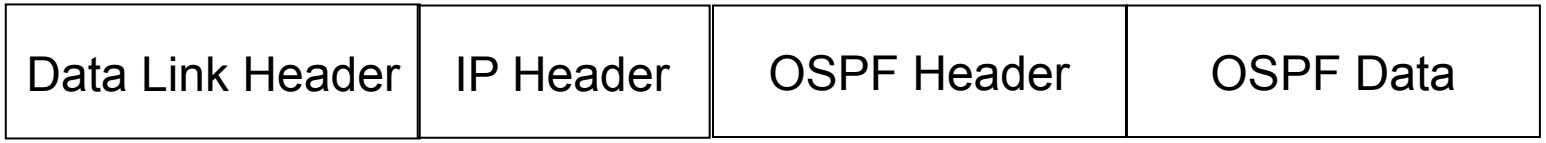

- IS-IS is directly encapsulated in Layer 2

For Service Providers

- Which IGP should an ISP choose?
  - Both OSPF and IS-IS use Dijkstra SPF algorithm
  - Exhibit same convergence properties
  - IS-IS less widely implemented on router platforms
  - IS-IS runs on data link layer, OSPF runs on IP layer
- Why do we keep discussing the merits of each IGP?

For Service Providers

- Biggest ISPs tend to use IS-IS – why?
- Looking back to the early 1990s:
  - Cisco implementation of IS-IS was much more stable and reliable than OSPF implementation – ISPs naturally preferred IS-IS
  - Main IS-IS implementations more tuneable than equivalent OSPF implementations – because biggest ISPs using IS-IS put more pressure on Cisco to implement “knobs”

For Service Providers

- Moving forward a decade
  - Early Cisco OSPF implementation was substantially rewritten
    - Now competitive with IS-IS in features and performance
  - Router vendors wishing a slice of the core market needed an IS-IS implementation as solid and as flexible as that from Cisco
    - Those with IS-IS & OSPF support tend to ensure they exhibit performance and feature parity

How to choose an IGP?

- OSPF
  - Rigid area design – all networks must have area 0 core, with sub-areas distributed around
  - Suits ISPs with central high speed core linking regional PoPs

How to choose an IGP?

- IS-IS
  - Relaxed two level design – L2 routers must be linked through the backbone
  - Suits ISPs with “stringy” networks, diverse infrastructure, etc, not fitting central core model of OSPF
  - More flexible than OSPF, but easier to make mistakes too

Considerations

- "Security"
  - IS-IS runs on link layer
  - Not possible to "attack" the IGP using IP as with OSPF
- Not dependent on IP addressing
  - IS-IS's NSAP addressing scheme avoids dependencies on IP as with OSPF
- "Reliability"
  - IS-IS has long been used by the majority of the world's biggest ISPs
  - Belief that equipment vendors pay more attention to IS-IS reliability, scalability, and features

More considerations

- Migration to IPv6
  - Adding IPv6 means OSPFv2 and OSPFv3 in network
    - Two protocols, two sets of identical configuration
  - IS-IS simply requires the addition of the IPv6 address-family
    - Most networks operate single topology for IPv4 and IPv6
  - Is this why there is now RFC5838 describing support of multiple address families in OSPFv3?
    - Vendor support?

Migration Plan

1. Verify OSPF configuration and operation
2. Deploy IS-IS over entire backbone
3. Set OSPF admin distance to be higher than IS-IS
4. Check for remnants in OSPF
5. Remove OSPF from entire backbone
6. Confirm IGP operation

Verify OSPF Configuration

- Set BGP next hop to be local router
  - No external point-to-point links need to be carried on OSPF
  - If external point-to-point links are required (for monitoring), carry in iBGP tagged with specific community visible to monitoring system only
- Remove surplus OSPF configuration
  - Only Loopback and internal point-to-point links should be carried in OSPF
  - (Loopback needed for iBGP etc)

Configuration Example: IOS <12.4

```
interface loopback 0
 ip addr 172.16.1.1 255.255.255.255
interface fastethernet 0/0
 ip address 172.16.0.1 255.255.255.252
interface fastethernet 0/1
 ip address 172.16.0.5 255.255.255.252
…
router ospf 100
 max-metric router-lsa on-startup wait-for-bgp
 passive-interface default
 no passive-interface fastethernet 0/0
 no passive-interface fastethernet 0/1
 network 172.16.0.0 mask 0.0.0.3.area 0
 network 172.16.0.4 mask 0.0.0.3 area 0
 network 172.16.1.1 mask 0.0.0.0 area 0
```

Configuration Example: IOS 12.4

```
interface loopback 0
 ip addr 172.16.1.1 255.255.255.255
 ip ospf 100 area 0
interface fastethernet 0/0
 ip address 172.16.0.1 255.255.255.252
 ip ospf 100 area 0
interface fastethernet 0/1
 ip address 172.16.0.5 255.255.255.252
 ip ospf 100 area 0
…
router ospf 100
 max-metric router-lsa on-startup wait-for-bgp
 passive-interface default
 no passive-interface fastethernet 0/0
 no passive-interface fastethernet 0/1
```

IPv6 configuration

- If IPv6 has already been deployed
  - OSPFv3 configuration also needs to be tidied up
- In Cisco IOS:
  - **`router ospf 100`** configuration should look identical to the **`ipv6 router ospf 100`** configuration
  - If not, fix it
- Check that the IPv4 adjacencies match the IPv6 adjacencies

Verifying OSPF operation

- Verifying operation is important after clean up
  - iBGP peers all stable
  - Next hop values are all valid
  - Check OSPF routing table
- If OSPFv3 deployed for IPv6, compare with OSPFv2
  - As well as adjacencies, compare **routing** table entries

Deploy IS-IS over entire backbone

- ISPs will deploy IPv6 dual-stack across their infrastructure
  - Every device running an IPv4 IGP will also require to run an IPv6 IGP
- Single-topology IS-IS
  - IPv4 and IPv6 topology identical
  - Needs care as adjacent routers need to have both IPv4 and IPv6 on the link
- Multi-topology IS-IS
  - IPv4 and IPv6 topology could differ
  - More flexibility for operators doing incremental roll-outs of IPv6

Deploy IS-IS over entire backbone

- IS-IS deployment (Cisco IOS):
  - IS-IS protocol distance is 115 (higher than OSPF's 110)
  - Use wide metrics (required for IPv6 address family support)
  - Only using Level-2 IS (Cisco IOS default is L1L2)
  - Passive interface configuration means IS-IS is not run on the interface, but the address is announced in the IGP
- IPv6 addressing in backbone – choice of:
  - Global unicast addresses
  - Link local addressing/unnumbered interfaces

Configuration Example: IOS

```
interface loopback 0
 ip address 172.16.1.1 255.255.255.255
 ipv6 address 2001:db8::1/128
!
interface fastethernet 0/0
 ip address 172.16.0.1 255.255.255.252
 ipv6 unnumbered loopback 0
 ip router isis ISP
 isis metric 20 level-2
 ipv6 router isis ISP
 isis ipv6 metric 20 level-2
!
(next slide)
```

Configuration Example: IOS (cont)

```
interface fastethernet 0/1
 ip address 172.16.0.5 255.255.255.252
 ipv6 unnumbered loopback 0
 ip router isis ISP
 isis metric 20 level-2
 ipv6 router isis ISP
 isis ipv6 metric 20 level-2
!
router isis ISP
 net 49.0001.1720.1600.1001.00
 passive-interface Loopback 0
 is-type level-2-only
 metric-style wide
 set-overload-bit on-startup wait-for-bgp
!
 address-family ipv6
  multi-topology
  set-overload-bit on-startup wait-for-bgp
 exit-address-family
!
```

Both IPv4 and IPv6 configurations

Set OSPF Admin Distance High

- Once IS-IS is deployed over entire backbone, set OSPF's admin distance above that of IS-IS
  - For all routers across the backbone
- Cisco IOS example:

```
router ospf 100
 distance 120
!
ipv6 router ospf 100
 distance 120
```

- All IS-IS paths learned by the router now take priority over the OSPF paths
  - For both IPv4 and IPv6

OSPF remnants

- As IS-IS is now responsible for interior routing, if all the preparation work was completed, there should be no prefixes left in OSPF
  - If there are, check what they are, and what caused them
- Remnant prefixes could include:
  - Forgotten passive interfaces for IS-IS
  - Forgotten active adjacencies

OSPF remnants

- Check adjacencies across the backbone
  - Compare **`show ip ospf neigh`** with **`show isis neigh`**
  - There should be the same number of neighbours
  - If not, fix the problem
  - Don’t forget IPv6!
- End result of tidying up work should mean:
  - **No more prefixes left in OSPF**
  - **A successful deployment of IS-IS**

Remove OSPF

- OSPF can now be safely removed from the entire backbone
- Cisco IOS:

```
no router ospf 100
no ipv6 router ospf 100
```

  - Will also need to go to each interface and remove ospf metric, link type, and authentication configuration
    - IOS unfortunately does not remove these when the routing process is removed

Confirm IGP operation

- IS-IS should now be functioning normally
- Verify iBGP sessions
  - Should have been completely unaffected by the entire migration process
- Verify next hop values
  - Adjacencies should be known by IS-IS
- Verify customer and external access
- Task complete

Conclusion

- Migration from OSPFv2 and OSPFv3 to IS-IS is straightforward
  - With planning
  - With adherence to procedure developed during planning
- Can be carried out any time
  - (but planned maintenance slots strongly recommended)
- Now running single multi-address family IGP to support both IPv4 and IPv6

Footnote

- Migrating from IS-IS to OSPF
  - Use the reverse of the described process
  - But why would anyone?
- Migrating from other IGPs to IS-IS
  - Follow the same procedures described here