

Migrating from OSPF to ISIS



Philip Smith
APRICOT 2012

New Delhi

21st February – 2nd March 2012

Presentation slides

- Available at:
 - <http://thyme.apnic.net/ftp/seminars/APRICOT2012-OSPF-to-ISIS-migration.pdf>
 - And on the APRICOT 2012 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 ISIS
 - This is not as difficult as it sounds
- Presentation describes:
 - The differences between OSPF and ISIS
 - The migration process
 - Based on several successful migrations
 - Uses Cisco IOS CLI as examples

Comparing ISIS and OSPF

- ❑ Both are Link State Routing Protocols using the Dijkstra SPF Algorithm
- ❑ So what's the difference then?
- ❑ And why do ISP engineers end up arguing so much about which is superior?

What Is IS-IS ?

- ❑ Intermediate System to Intermediate System
- ❑ An “IS” is ISO terminology for a router
- ❑ IS-IS was originally designed for use as a dynamic routing protocol for ISO CLNP, defined in the ISO 10589 standard
- ❑ Later adapted to carry IP prefixes in addition to CLNP (known as Integrated or Dual IS-IS) as described in RFC 1195
- ❑ Predominantly used in ISP environment

IS-IS Timeline

- 1978ish “New” Arpanet Algorithm
 - Eric Rosen et al
- 1986 to 90 Decnet Phase V
 - Radia Perlman, Mike Shand
- 1987 ISO 10589 (IS-IS)
 - Dave Oran
- 1990 RFC 1195 (Integrated IS-IS)
 - Ross Callon, Chris Gunner
- 1990 to present: All sorts of enhancements
 - Everyone contributed!
- 2008 RFC5308 adds IPv6 support
 - And RFC5120 adds Multi-Topology Routing support

What Is OSPF ?

- ❑ Open Shortest Path First
- ❑ Link State Protocol using the Shortest Path First algorithm (Dijkstra) to calculate loop-free routes
- ❑ Used purely within the TCP/IP environment
- ❑ Designed to respond quickly to topology changes but using minimal protocol traffic
- ❑ Used in both Enterprise and ISP Environment

OSPF Timeline

- ❑ Development began in 1987 by IETF
- ❑ OSPFv1 published in 1989 with RFC 1131
- ❑ OSPFv2 published in 1991 with RFC 1247
- ❑ Further enhancements to OSPFv2 in 1994 with RFC 1583 and in 1997 with RFC 2178
- ❑ Last revision was in 1998 with RFC 2328 to fix minor problems
- ❑ All above OSPF RFCs authored by John Moy
- ❑ RFC2740 introduced OSPFv3 (for IPv6) in 1999, replaced by RFC5340 in 2008

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

- ❑ Host
- ❑ Router
- ❑ Link
- ❑ Packet
- ❑ Designated router (DR)
- ❑ Backup DR (BDR)
- ❑ Link-State Advertisement (LSA)
- ❑ Hello packet
- ❑ Database Description (DBD)

ISIS

- ❑ End System (ES)
- ❑ Intermediate System (IS)
- ❑ Circuit
- ❑ Protocol Data Unit (PDU)
- ❑ Designated IS (DIS)
- ❑ N/A (no BDIS is used)
- ❑ Link-State PDU (LSP)

- ❑ IIH PDU
- ❑ Complete sequence number PDU (CSNP)

IS-IS and OSPF Terminology

(Cont.)

OSPF

- ❑ Area
- ❑ Non-backbone area
- ❑ Backbone area

- ❑ Area Border Router (ABR)
- ❑ Autonomous System Boundary Router (ASBR)


ISIS

- ❑ Sub domain (area)
- ❑ Level-1 area
- ❑ Level-2 Sub domain (backbone)
- ❑ L1L2 router

- ❑ 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 ISIS use Dijkstra SPF algorithm
 - Exhibit same convergence properties
 - ISIS less widely implemented on router platforms
 - ISIS runs on data link layer, OSPF runs on IP layer

For Service Providers

- Biggest ISPs tend to use ISIS – why?
- Looking back to the early 1990s:
 - Cisco implementation of ISIS was much more stable and reliable than OSPF implementation – ISPs naturally preferred ISIS
 - Main ISIS implementations more tuneable than equivalent OSPF implementations – because biggest ISPs using ISIS 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 ISIS in features and performance
 - Router vendors wishing a slice of the core market needed an ISIS implementation as solid and as flexible as that from Cisco
 - Those with ISIS & 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 network linking regional PoPs
- Teaches good routing protocol design practices

How to choose an IGP?

□ ISIS

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

More considerations

- Migration to IPv6
 - Adding IPv6 means OSPFv2 and OSPFv3 in network
 - Two protocols, two sets of identical configuration
 - ISIS 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 ISIS over entire backbone
3. Set OSPF admin distance to be higher than ISIS
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 ISIS over entire backbone

- Ten years ago ISPs were experimenting with partial IPv6 deployments before extending over entire backbone
 - Fears about router code stability
 - Uncertainty about need to deploy IPv6 (given lack of “market demand” and continued abundance of IPv4 addresses)

Deploy ISIS over entire backbone

- Today, IPv6 deployment is fundamentally important to ensure continued network and Internet growth
 - Which means that ISPs will deploy dual-stack
 - And every device running an IPv4 IGP will also require to run an IPv6 IGP
 - ⇒ Single congruent topology (no multi-topology ISIS)

Deploy ISIS over entire backbone

- ISIS deployment (Cisco IOS):
 - Leave distance at default of 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)
 - Deploy both IPv4 and IPv6 at the same time
 - Passive interface configuration means ISIS 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)
```

Both IPv4 and IPv6
configurations

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
    set-overload-bit on-startup wait-for-bgp
  exit-address-family
!
```

Both IPv4 and IPv6 configurations



Set OSPF Admin Distance High

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

```
router ospf 100
  distance 120
!
ipv6 router ospf 100
  distance 120
```
- ❑ All ISIS paths learned by the router now take priority over the OSPF paths
 - For both IPv4 and IPv6

OSPF remnants

- As ISIS 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 ISIS
 - 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 ISIS**

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

- ❑ ISIS 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 ISIS
- ❑ Verify customer and external access
- ❑ Task complete

Conclusion

- Migration from OSPFv2 and OSPFv3 to ISIS 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 ISIS to OSPF
 - Use the reverse of the described process
 - But why would anyone?
- Migrating from EIGRP to ISIS
 - Follow the same procedures described here
 - EIGRP's administrative distance is either 90 or 170, depending on prefix origin ⇒ set ISIS admin distance appropriately