

Migrating from OSPF to IS-IS



ISP Workshops

Introduction

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

Motivation

- “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

- ❑ **next-hop-self** for iBGP
 - 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 **network** statements
 - Only Loopback and internal point-to-point links should remain
 - (For Cisco IOS 12.4 onwards and IOS-XR ensure that OSPF is only activated on internal and loopback interfaces – same for OSPFv3 configuration)

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
```


Configuration Example: IOS-XR

```
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 ISP
  area 0
    interface Loopback0
      passive enable
    !
    interface fastethernet 0/0
    !
    interface fastethernet 0/1
```

IPv6 configuration

- ❑ If IPv6 has already been deployed
 - OSPFv3 configuration also needs to be tidied up
- ❑ For IOS:
 - `router ospf 100` configuration should look identical to the `ipv6 router ospf 100` configuration
- ❑ For IOS-XR:
 - `router ospf ISP` configuration should look identical to the `router ospfv3 ISP` configuration
- ❑ 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:
 - 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 (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 address 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 address 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 level-2
  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



Configuration Example: IOS-XR

```
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 enable
interface fastethernet 0/1
  ip address 172.16.0.5 255.255.255.252
  ipv6 enable
...
router isis ISP
  set-overload-bit on-startup wait-for-bgp
  is-type level-2-only
  net 49.0001.1720.1600.1001.00
  address-family ipv4 unicast
    metric-style wide
  address-family ipv6 unicast
    metric-style wide
(next slide)
```


Configuration Example: IOS-XR (cont)

```
router isis ISP
!
interface Loopback0
  passive
  address-family ipv4 unicast
    metric 1 level 2
  address-family ipv6 unicast
    metric 1 level 2
!
interface fastethernet 0/0
  address-family ipv4 unicast
    metric 20 level 2
  address-family ipv6 unicast
    metric 20 level 2
!
interface fastethernet 0/1
  address-family ipv4 unicast
    metric 20 level 2
  address-family ipv6 unicast
    metric 20 level 2
!
```

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
- ❑ 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
- ❑ 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
- ❑ 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 does not remove these when the routing process is removed
- ❑ IOS-XR
 - `no router ospf ISP`
 - `no router ospfv3 ISP`
 - Performs a clean removal

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 in 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 EIGRP to IS-IS
 - Follow the same procedures described here
 - EIGRP's administrative distance is either 90 or 170, depending on prefix origin ⇒ set IS-IS admin distance appropriately

Migrating from OSPF to IS-IS



ISP Workshops