

Introduction to IS-IS

ISP Workshops



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)

Last updated 30th January 2021

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

IS-IS

- Intermediate System to Intermediate System
- ISO 10589 specifies OSI IS-IS routing protocol for ConnectionLess-mode Network Services (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

IS-IS

- Known as a Link State Routing Protocol
 - The other link state routing protocol is OSPF
 - Each node in the network computes the map of connectivity through the network
 - Both use 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
- The other type of Routing Protocol is Distance Vector
 - Like Cisco's EIGRP or RIP
 - Each node shares its view of the routing table with other nodes

IS-IS

- Routers with IS-IS enabled on them look for neighbouring routers also running IS-IS
 - Hello Protocol Data Units (PDUs) are exchanged
 - The “Hello” packet includes the list of known neighbours, and details such as “hello interval” and “router dead interval”
 - Hello interval – how often the router will send Hellos
 - Router dead interval – how long to wait before deciding router has disappeared
 - The values of “hello interval” and “router dead interval” do not need to match on both neighbours (unlike for OSPF)
 - When a neighbouring router responds with matching details, a neighbour relationship is formed

IS-IS Neighbour Relationships

- A relationship is formed between neighbouring routers for the purpose of exchanging routing information
 - This is called an **ADJACENCY**

IS-IS Adjacencies

- Once an adjacency is formed, neighbours share their link state information
 - Information goes in a **Link State PDU** (LSP)
 - LSPs are flooded to all neighbours
- New information received from neighbours is used to compute a new view of the network
- On a link failure
 - New LSPs are flooded
 - The routers recompute the routing table

IS-IS across a network

- ❑ All routers across the network form neighbour relationships with their directly attached neighbours
- ❑ Each router computes the routing table
- ❑ Once each router has the same view of the network, the network has **converged**
- ❑ The IGP design for a network is crucially important to ensure **scalability** and **rapid convergence**
- ❑ Generally: **the fewer the prefixes, the faster the convergence**

IS-IS Levels

- IS-IS has a 2 layer hierarchy
 - Level-2 (the backbone)
 - Level-1 (the edge)
- A router can be
 - Level-1 (L1) router
 - Level-2 (L2) router
 - Level-1-2 (L1L2) router
- Most small to medium networks (up to ~500 routers) are happily using just Level-2

IS-IS

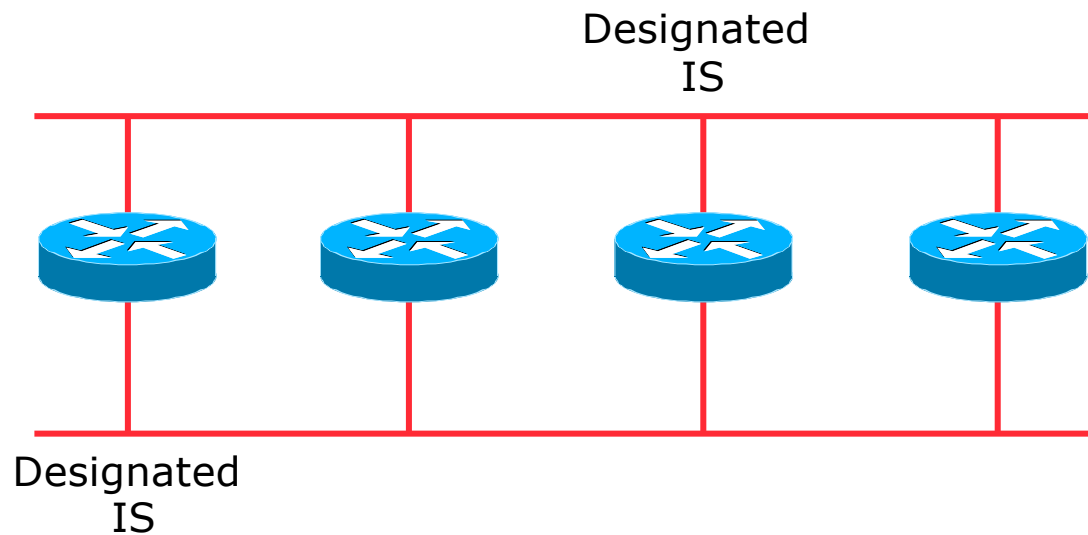
- IS-IS is multiprotocol
 - Integrated IS-IS carries CLNS and IPv4 address families
 - RFC5308 adds IPv6 address family support
 - RFC5120 adds multi-topology support
- IS-IS extended to carry IPv6 prefixes
 - Either sharing topology with IPv4
 - When IPv4 and IPv6 topologies are identical
 - Or using “multi-topology”, independent of IPv4
 - Allows incremental rollout of IPv6

Links in IS-IS

- Two types of links in IS-IS:
 - Point-to-point link
 - Only one other router on the link, forming a point-to-point adjacency
 - Multi-access network (e.g. ethernet)
 - Potential for many other routers on the network, with several other adjacencies
- IS-IS in multi-access networks has optimisations to aid scaling
 - One router is elected to originate the LSPs for the whole multi-access network
 - Called “**Designated Intermediate System**”
 - Other routers on the multi-access network form adjacencies with the DIS

Designated IS

- There is ONE designated router per multi-access network
 - Generates network link advertisements
 - Assists in database synchronization
 - Scales IS-IS for multi-access (ethernet) networks




Selecting the Designated Router

- Configured priority (per interface)
 - Configure high priority on the router to be the DIS

```
interface gigabitethernet0/1
isis priority 127 level-2
```

- Else priority determined by highest MAC address
 - Best practice is to set two routers to be highest priority – then in case of failure of the DIS there is deterministic fall back to the other



Adjacencies: Cisco IOS Examples

- To find CLNS adjacency state, use:

```
show clns neighbor
```

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
Router2	Fa0/0	ca01.9798.0008	Up	23	L2	M-ISIS
Router3	Se1/0	*HDLC*	Up	26	L2	M-ISIS

- To find IS-IS adjacency state, use:

```
show isis neighbor
```

System Id	Type	Interface	IP Address	State	Holdtime	Circuit Id
Router2	L2	Fa0/0	10.10.15.2	UP	24	Router2.01
Router3	L2	Se1/0	10.10.15.6	UP	27	00

Adjacencies: JunOS Examples

- To find IS-IS adjacency state, use:

```
philip@R1> show isis adjacency
```

Interface	System	L State	Hold (secs)	SNPA
xe-1/0/0.0	R2	2 Up	8	3c:61:4:26:40:a5
xe-1/1/0.0	R3	2 Up	9	3c:8a:b0:cf:72:95

IS-IS NSAP Address

- ❑ IP based routing protocols have a router-id to uniquely identify a router
- ❑ In IS-IS, the IS (router) is identified by a Network Entity Title (NET)
 - Can be from 64 to 160 bits long
 - The NET is the address of a Network Service Access Point (NSAP), identifying an instance of IS-IS running on the IS
- ❑ ISPs typically choose NSAP addresses thus:
 - First 8 bits – pick a number (usually 49)
 - Next 16 bits – area
 - Next 48 bits – router loopback address
 - Final 8 bits – zero
- ❑ Example:
 - NSAP: 49.0001.1921.6800.1001.00
 - Router: 192.168.1.1 (loopback) in Area 1

IS-IS NSAP Address (Alternative)

- A simpler alternative, assuming a well documented ISP design
 - First 8 bits – pick a number (usually 49)
 - Next 16 bits – area
 - Next 16 bits – PoP identifier
 - Next 16 bits – Router identifier
 - Final 8 bits – zero
- Example:
 - NSAP: 49.0001.0009.0003.00
 - Router: #3 in PoP 9 in Area 1

IS-IS on Cisco IOS

❑ Starting IS-IS in Cisco's IOS

```
router isis as42
```

- Where “as42” is the process ID

❑ IS-IS process ID is local to the router

- Allows the possibility of running multiple instances of IS-IS on one router
- The process ID is not passed between routers
- Some ISPs configure the process ID to be the same as their BGP Autonomous System Number

IS-IS on Cisco IOS

- ❑ Cisco IOS default is for all routers to be L1L2
 - This is suboptimal – all routers need to be L2 only
- ❑ Once IS-IS is started, other required configuration under the IS-IS process includes:

```
log-adjacency-changes
```

- Capture adjacency changes in the system log

```
metric-style wide
```

- Set metric-style to wide

```
is-type level-2-only
```

- Set IS type to level 2 only (router-wide configuration)

```
net 49.0001.<loopback>.00
```

- Set NET address

Adding interfaces to IS-IS

- To activate IS-IS on an interface:

```
interface Gigabit 4/0  
ip router isis as42
```

- Puts interface subnet address into the LSDB
 - Enables CLNS on that interface
- To disable IS-IS on an interface:
- ```
router isis as42
passive-interface Gigabit 2/0
```
- Disables CLNS on that interface
  - Puts the interface subnet address into the LSDB
- No IS-IS configuration for an interface
    - No CLNS run on interface, no interface subnet in the LSDB

# IS-IS interface costs

---

- All interfaces have a default metric of 10
  - Fine for a uniform network, but most backbones have different link capacities between routers & PoPs
  - Many operators set default metric to 100000
- Many operators develop their own interface metric strategy
  - ```
isis metric 100 level-2
```
 - Sets interface metric to 100
 - Care needed as the sum of metrics determines the best path through the network
- IS-IS chooses lowest cost path through a network
- IS-IS will load balance over paths with equal total cost to the same destination

IS-IS on JunOS (1)

- The NET is set on the Loopback interface:

```
set interfaces lo0 unit 0 family iso address 49.0001.<loopback>.00
```

- IS-IS is all configured within "protocols"

- Set wide metrics only:

```
set protocols isis level 2 wide-metrics-only
```

- Set level-2 only:

```
set protocols isis level 1 disable
```

- Set multi-topology:

```
set protocols isis topologies ipv6-unicast
```

IS-IS on JunOS (2)

- All interfaces participating in IS-IS need to have address family ISO added:

```
set interfaces xe-0/1/0 unit 0 family iso
set interfaces xe-0/1/3 unit 0 family iso
```


- Remaining IS-IS configuration

- Use “passive” for interfaces where no adjacencies are expected but link address will be distributed by IS-IS:

```
set protocols isis interface lo0.0 passive
set protocols isis interface xe-0/1/0.0 point-to-point
set protocols isis interface xe-0/1/0.0 level 1 disable
set protocols isis interface xe-0/1/0.0 level 2 metric 20
set protocols isis interface xe-0/1/0.0 level 2 ipv6-unicast-metric 20
set protocols isis interface xe-0/1/3.0 passive
```


IS-IS Metric Calculation

- Best path/lowest cost = 11




IS-IS Metric Calculation

- Best path/lowest cost = 11



- Equal cost paths = 15



IS-IS Neighbour Authentication on Cisco IOS

- Neighbour authentication is highly recommended
 - Prevents unauthorised routers from forming neighbour relationships and potentially compromising the network
- Create a suitable key-chain

```
key chain isis-as42
  key 1
    key-string <password>
!
```

- There can be up to 255 different keys in each key chain

IS-IS Neighbour Authentication on Cisco IOS

- Apply key-chain per interface:

```
interface Gigabit 4/0
  isis authentication mode md5 level-2
  isis authentication key-chain isis-as42 level-2
!
```

- Apply key-chain to IS-IS process (all interfaces):

```
router isis as42
  authentication mode md5 level-2
  authentication key-chain isis-as42 level-2
!
```

IS-IS Neighbour Authentication on JunOS

- Neighbour authentication is highly recommended
 - Prevents unauthorised routers from forming neighbour relationships and potentially compromising the network

- Create and apply key to all processes:

```
set protocols isis level 2 authentication-key "$9$RLzhfW4iG.mTwZ5Fd9OC7-V"  
set protocols isis level 2 authentication-type md5
```

- Cisco IOS interoperability:

```
set protocols isis loose-authentication-check
```

Originating a Default Route (Cisco IOS IPv4)

- How to always originate a default route into IS-IS:

```
router isis as42
  default-information originate
```

- Originate a default route into IS-IS only if a default route exists in the RIB from BGP:

```
ip access-list standard BGP-NH
  permit <ebgp neighbor address>
!
ip prefix-list DEFAULT-ROUTE permit 0.0.0.0/0
!
route-map DEFAULT-ORIG permit 10
  match ip address prefix-list DEFAULT-ROUTE
  match ip next-hop BGP-NH
!
router isis as42
  default-information originate route-map DEFAULT-ORIG
```

Originating a Default Route (Cisco IOS IPv6)

- How to always originate a default route into IS-IS:

```
router isis as42
  address-family ipv6
    default-information originate
```

- Originate a default route into IS-IS only if a default route exists in the RIB from BGP:

```
ipv6 access-list BGP-NHv6
  permit ipv6 host <ebgp neighbor link-local address> any
!
ipv6 prefix-list DEFAULT-v6ROUTE permit ::/0
!
route-map DEFAULT-ORIGv6 permit 10
  match ipv6 address prefix-list DEFAULT-v6ROUTE
  match ipv6 next-hop BGP-NHv6
!
router isis as42
  address-family ipv6
    default-information originate route-map DEFAULT-ORIGv6
```

Originating a Default Route (JunOS)

- How to originate a default route into IS-IS:
 - Set up Static default routes to Discard (or wherever needed):

```
set routing-options rib inet6.0 static route ::/0 discard
set routing-options static route 0.0.0.0/0 discard
```

- Create the (multi-protocol) policy:

```
set policy-options policy-statement ISIS-DEF-ORIG term IPv4 from family inet
set policy-options policy-statement ISIS-DEF-ORIG term IPv4 from protocol static
set policy-options policy-statement ISIS-DEF-ORIG term IPv4 from route-filter 0.0.0.0/0 exact
set policy-options policy-statement ISIS-DEF-ORIG term IPv4 then accept
set policy-options policy-statement ISIS-DEF-ORIG term IPv6 from family inet6
set policy-options policy-statement ISIS-DEF-ORIG term IPv6 from protocol static
set policy-options policy-statement ISIS-DEF-ORIG term IPv6 from route-filter ::/0 exact
set policy-options policy-statement ISIS-DEF-ORIG term IPv6 then accept
set policy-options policy-statement ISIS-DEF-ORIG term ALL from protocol isis
set policy-options policy-statement ISIS-DEF-ORIG term ALL then accept
```

- And then apply to IS-IS:

```
set protocols isis export ISIS-DEF-ORIG
```

IS-IS on Point-to-Point Ethernet

□ IS-IS on point-to-point ethernet:

- DIS election is not needed on a point-to-point link – so it is disabled, which is more efficient

- Cisco IOS:

```
interface Gigabit 4/0
  isis network point-to-point
```

- JunOS:

```
set protocols isis interface xe-0/1/0.0 point-to-point
```

- As DIS election is independent of IP, the above command is generic – there is no need for an IPv6 equivalent

Handling IPv6 in IS-IS for Cisco IOS

- To add IPv6 support in IS-IS:

```
interface Gigabit 4/0
  ipv6 router isis as42
```

- Topologies:

- For single topology, nothing else is required
- For multi-topology, include:

```
router isis as42
  address-family ipv6
  multi-topology
```

Handling IPv6 in IS-IS for JunOS

□ To add IPv6 support in IS-IS:

- Nothing needed if IPv6 configured on interface
- Remember to set metric appropriately:

```
set protocols isis interface xe-0/1/0.0 level 2 ipv6-unicast-metric 20
```

□ Topologies:

- For single topology, nothing else is required
- For multi-topology, include:

```
set protocols isis topologies ipv6-unicast
```

Conclusion

- IS-IS is a Link State Routing Protocol
- Quick and simple to get started
 - But has a myriad of options and features to cover almost all types of network topology
 - ISPs keep their IS-IS design **SIMPLE**
 - >500 routers running in L2 is entirely feasible

Introduction to IS-IS



ISP Workshops