Introduction to IS-IS

ISP Workshops

Last updated 26th May 2018

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

4

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 ~400 routers) can happily use 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

Designated
IS

Designated
IS

12
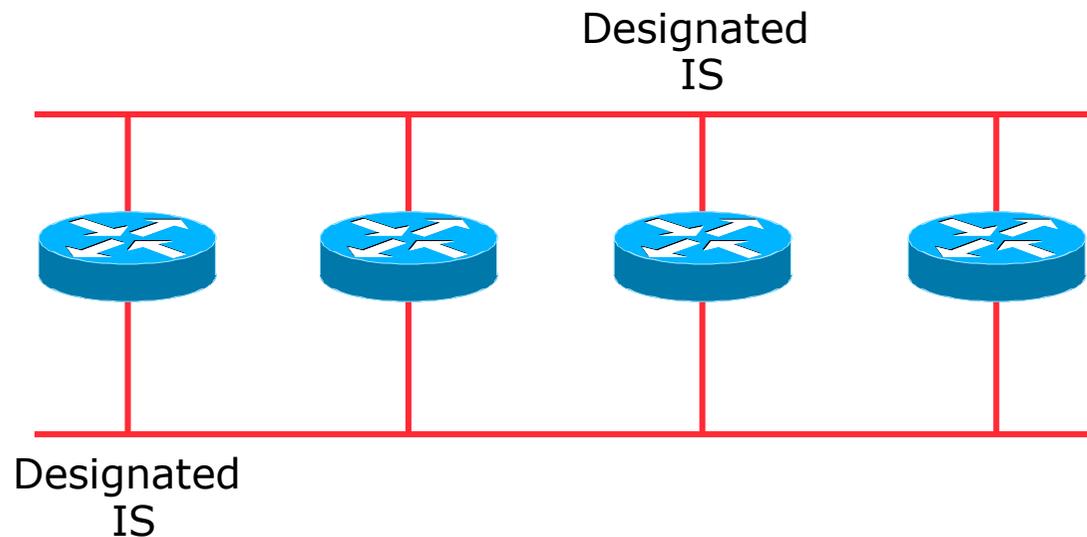
Selecting the Designated Router

- ❑ 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

e0:f8:47:1d:93:3c   e0:f8:47:1d:81:32

R1 ← DIS   R2

e0:f8:47:1d:93:30
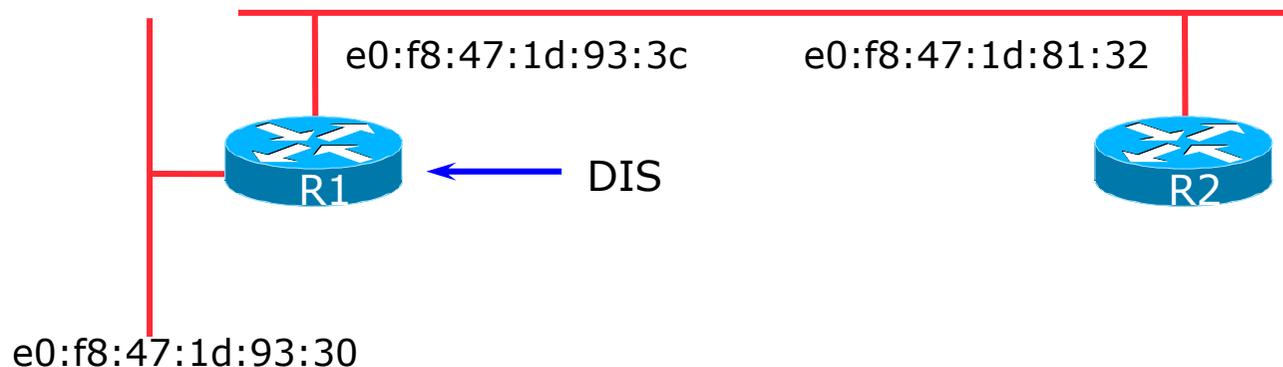
Adjacencies: Examples

- 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
```

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

15

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 in 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

19

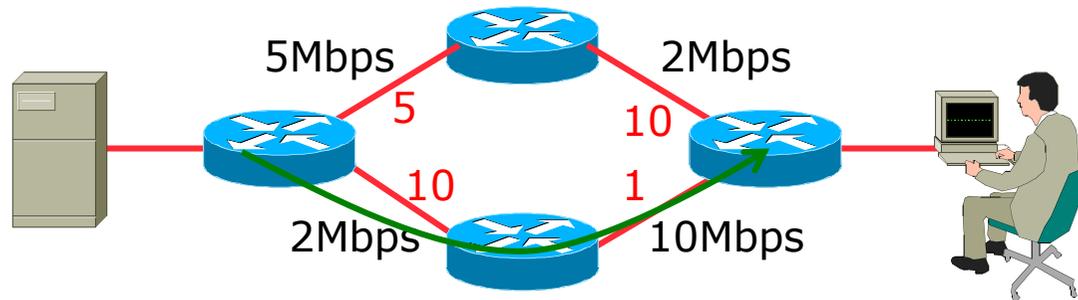# IS-IS interface costs

- 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
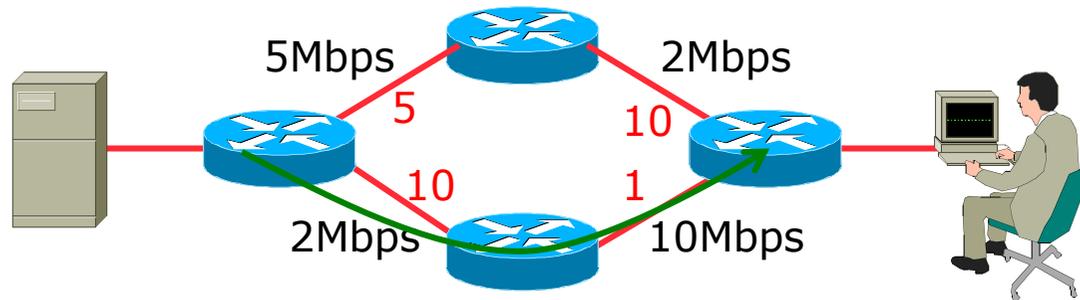- IS-IS will load balance over paths with equal total cost to the same destination

IS-IS Metric Calculation
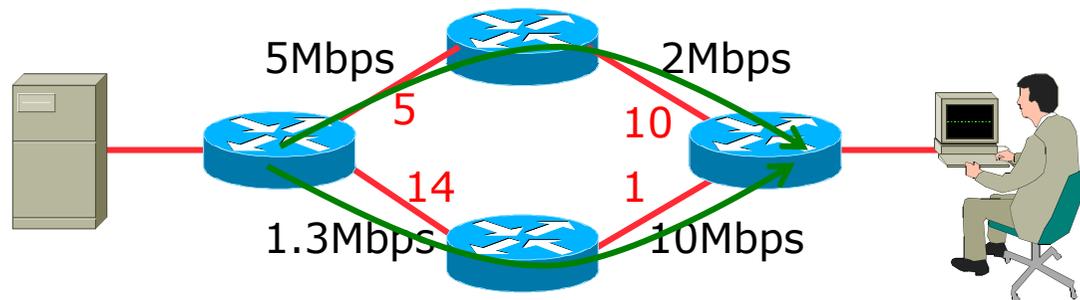
□ Best path/lowest cost = 11

IS-IS Metric Calculation

- Best path/lowest cost = 11

5Mbps   2Mbps

5   10

10   1

2Mbps   10Mbps

- Equal cost paths = 15

5Mbps   2Mbps

5   10

14   1

1.3Mbps   10Mbps

IS-IS Neighbour Authentication

- 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

- 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
!
```

Originating a Default Route (IPv4)

- Always originating 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 (IPv6)

- Always originating 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
```

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

    ```
    interface Gigabit 4/0
      isis network 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

- 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
```

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
  - >400 routers in running in L2 is entirely feasible

Introduction to IS-IS

ISP Workshops