ISIS for ISPs

AfNOG 2011 AR-E Workshop

Configuring ISIS in Cisco IOS

- Starting ISIS in Cisco IOS

  ```
  router isis [name]
  ```

  - Where the optional ***name*** can specify the name of the ISIS process
- ISIS name is unique to the router
  - Gives possibility of running multiple instances of ISIS on one router
  - ISIS process name is not passed between routers in an AS
  - **Some ISPs configure the ISIS name to be the same as their BGP Autonomous System Number**

    e.g. `router isis as64510`

Configuring ISIS in Cisco IOS

- Once ISIS started, other ISP required configuration under the ISIS process includes:
  - Capture adjacency changes in the system log

    ```
    log-adjacency-changes
    ```
  - Set metric-style to wide – modern & scalable

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

    ```
    is-type level-2-only
    ```
  - Set NET address

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

Adding interfaces to ISIS

- To activate ISIS on an interface:

```
interface HSSI 4/0
 ip router isis isp-bb
```

- To disable ISIS on an interface:

```
router isis isp-bb
 passive-interface GigabitEthernet 0/0
```

  - Disables CLNS on that interface
  - Puts the interface subnet address into the LSDB
  - No need for “redistribute connected”
- No ISIS configuration for an interface
  - No CLNS run on interface, no interface subnet in the LSDB

Adding interfaces to ISIS

- Scaling ISIS: **`passive-interface default`**
  - Disables ISIS processing on all interfaces apart from those marked as no-passive
  - Places all IP addresses of all connected interfaces into ISIS
    - **NB. This is NOT the same behaviour as for OSPF**
  - There must be at least one non-passive interface:

```
router isis isp-bb
 passive-interface default
 no passive-interface GigabitEthernet 0/0
interface GigabitEthernet 0/0
 ip router isis isp-bb
```

Configuration for A&B

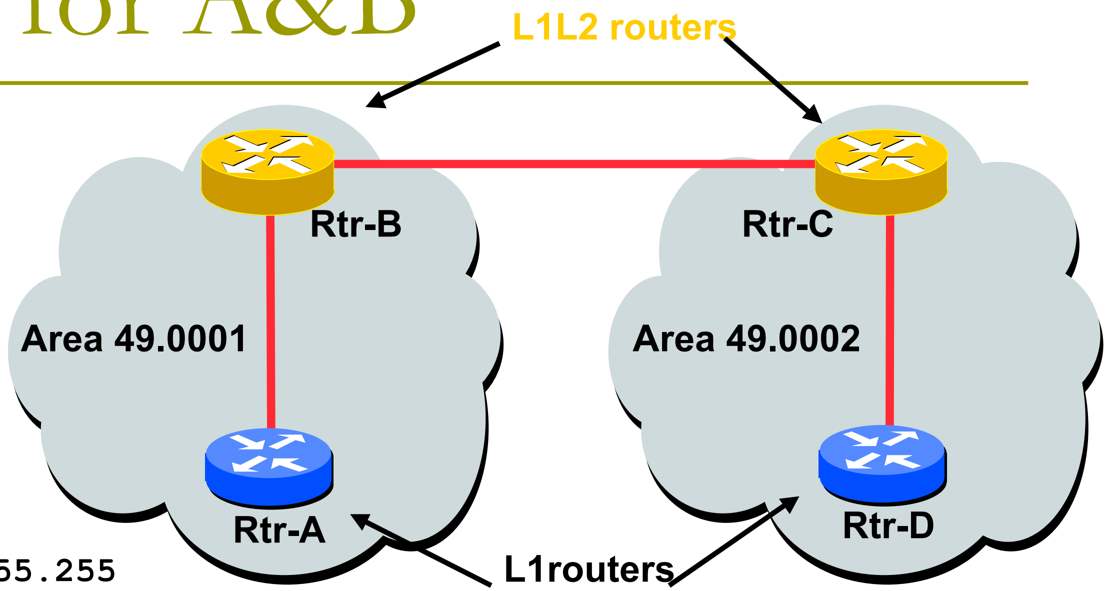

Router-B

```
interface Loopback0
 ip address 192.168.1.1 255.255.255.255
!
interface POS2/0/0
 ip address 192.168.222.1 255.255.255.0
 ip router isis
 isis circuit-type level-2
!
Interface FastEthernet4/0/0
 ip address 192.168.120.10 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
router isis
 passive-interface Loopback0
 net 49.0001.1921.6800.1001.00
```

Router-A

```
interface Loopback0
 ip address 192.168.1.5 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.120.5 255.255.255.0
 ip router isis
!
router isis
 is-type level-1
 passive-interface Loopback0
 net 49.0001.1921.6800.1005.00
```

Configuration for C&D

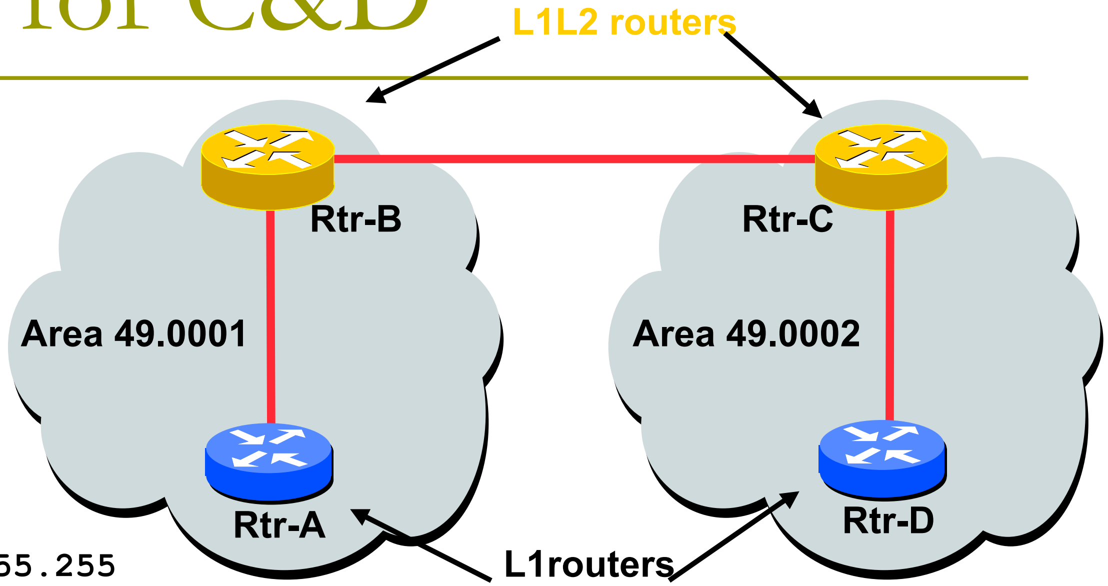

Router-C

```
interface Loopback0
 ip address 192.168.2.2 255.255.255.255
!
interface POS1/0/0
 ip address 192.168.222.2 255.255.255.0
 ip router isis
 isis circuit-type level-2
!
interface Fddi3/0
 ip address 192.168.111.2 255.255.255.0
 ip router isis
 isis circuit-type level-1
!
router isis
 passive-interface Loopback0
 net 49.0002.1921.6800.2002.00
```

Router-D

```
interface Loopback0
 ip address 192.168.2.4 255.255.255.255
!
interface Fddi6/0
 ip address 192.168.111.4 255.255.255.0
 ip router isis
!
router isis
 is-type level-1
 passive-interface Loopback0
 net 49.0002.1921.6800.2004.00
```

Status Commands in ISIS

- **show clns**
  - Shows the global CLNS status as seen on the router, e.g.

```
Rtr-B>show clns
Global CLNS Information:
  2 Interfaces Enabled for CLNS
  NET: 49.0001.1921.6800.1001.00
  Configuration Timer: 60, Default Holding Timer: 300,
   Packet Lifetime 64
  ERPDU's requested on locally generated packets
  Intermediate system operation enabled (forwarding
   allowed)
  IS-IS level-1-2 Router:
    Routing for Area: 49.0001
```

Status Commands in ISIS

- show clns neighbors
  - Shows the neighbour adjacencies as seen by the router:

```
Rtr-B> show clns neighbors
System Id      SNPA           Interface State Holdtime Type Protocol
1921.6800.2002 *PPP*            PO2/0/0    Up        29   L2     IS-IS
1921.6800.1005 00e0.1492.2c00 Fa4/0/0    Up         9   L1     IS-IS
```

  - More recent IOSes replace system ID with router hostname – much easier troubleshooting

Status Commands in ISIS

- show clns interface
  - Shows the CLNS status on a router interface:

```
Rtr-B>  show clns interface POS2/0/0
POS2/0/0 is up, line protocol is up
  Checksums enabled, MTU 4470, Encapsulation PPP
  ERPDUs enabled, min. interval 10 msec.
  RDPDUs enabled, min. interval 100 msec., Addr Mask enabled
 Congestion Experienced bit set at 4 packets
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 47 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x100
    Level-1 Metric: 10, Priority: 64, Circuit ID: 1921.6800.2002.00
    Number of active level-1 adjacencies: 0
    Level-2 Metric: 10, Priority: 64, Circuit ID: 1921.6800.1001.00
    Number of active level-2 adjacencies: 1
    Next IS-IS Hello in 2 seconds
```

Status Commands in ISIS

- **`show clns protocol`**
  - Displays the status of the CLNS protocol on the router:

```
Rtr-B> show clns protocol
IS-IS Router: <Null Tag>
  System Id: 1921.6800.1001.00  IS-Type: level-1-2
  Manual area address(es):
        49.0001
  Routing for area address(es):
        49.0001
  Interfaces supported by IS-IS:
        FastEthernet4/0/0 - IP
        POS2/0/0 - IP
  Redistributing:
    static
  Distance: 110
```

Other status commands

- **`show clns traffic`**
  - Shows CLNS traffic statistics and activity for the network
- **`show isis database`**
  - Shows the ISIS link state database
  - i.e. the “routing table”

Network Design Issues

- As in all IP network designs, the key issue is the addressing lay-out
- ISIS supports a large number of routers in a single area
- When using areas, use summary-addresses
- >400 routers in the backbone is quite doable

Network Design Issues

- Possible link cost
  - Default on all interface is 10
  - (Compare with OSPF which set cost according to link bandwidth)
  - Manually configured according to routing strategy
- Summary address cost
  - Equal to the best more specific cost
  - Plus cost to reach neighbor of best specific
- Backbone has to be contiguous
  - Ensure continuity by redundancy
- Area partitioning
  - Design so that backbone can NOT be partitioned

Scaling Issues

- Areas vs. single area
  - Use areas where
    - sub-optimal routing is not an issue
    - areas with one single exit point
- Start with L2-only everywhere is a good choice
  - Future implementation of level-1 areas will be easier
  - Backbone continuity is ensured from start

Narrow to Wide Metrics Transition

- When migrating from narrow to wide metrics, care is required
  - Narrow and wide metrics are NOT compatible with each other
  - Migration is a two stage process, using the “transition” keyword
- Networks using narrow metrics should first configure transition metrics across all routers:

```
router isis isp
 metric-style transition
```

- Once the whole network is changed to transition support, the metric style can be changed to wide:

```
router isis isp
 metric-style wide
```

ISP best common practices

- NET address construction
  - Area and loopback address
- L2
  - Configure globally in ISIS process
  - L1-L2 and L1 can be used later for scaling large networks
- Wide metrics
  - Narrow metrics are too limiting
- Deploying IPv6 in addition to IPv4
  - **Recommendation:** use single-topology (IOS default) unless an incremental roll-out of IPv6 is intended
  - Multi-topology is useful for an incremental roll-out, and if there should be future differences in topology between IPv4 and IPv6

Advanced ISIS Configuration

(for experts only)

Developed by Mark Tinka &
Philip Smith

Purging the RIB on link failure

- For routing protocols that are capable of responding to link failures, IOS allows such routing protocols to quickly and more efficiently delete associated routes from the RIB when a link, and the interface is removed from the routing table
- Without this command, the "less efficient" RIB process is used to delete the associated next-hop routes of the failed interface, by default
  - If this process has to work through a very large routing table, it can use up a number of CPU cycles and potentially increase convergence time.

```
ip routing protocol purge interface
```

ISIS neighbour authentication

- Create key chains to be used for HMAC-MD5 authentication for both Level-1 and Level-2

```
key chain isis-level1
 key 1
   key-string xxxxx
key chain isis-level2
 key 1
   key-string xxxxx
```

Setting up Loopback Interface

- Create the Loopback interface/Router-ID
  - It will NOT have IS-IS running on it because it is not a transit interface
  - Disabling IS-IS on it, while announcing the IP prefixes into IS-IS, allows the IS-IS domain to scale because LSP/Hello packets are not unnecessarily generated for the Loopback interface
  - An IS-IS metric will NOT be set, which will default the Loopback interface's metric to zero (0).

```
interface loopback0
 ip address 192.168.0.1 255.255.255.255
 ipv6 address 2001:db8:192:168:0:1/128
```

Level-1 Interface Configuration

- Configure addresses and enable ISIS for IPv4 and IPv6

```
interface gigabitethernet0/1
 ip address 192.168.1.1 255.255.255.192
 ipv6 address 2001:db8:192:168:1:1/112
 !
 ip router isis 1
 ipv6 router isis 1
```

- Ensure this interfaces runs at Level-1

```
 isis circuit-type level-1
```

Level-1 Interface: Metrics & Auth

- Set the costs for IPv4 and IPv6

```
interface gigabitethernet0/1
 isis metric 400 level-1
 isis ipv6 metric 400 level-1
```

- Enable HMAC-MD5 for level-1

```
 isis authentication mode md5 level-1
```

- Associate the key-chain defined earlier

```
 isis authentication key-chain isis-level1 level-1
```

Level-1 Interface: DIS and BFD

- Set this IS (router) to be the DIS in this Level-1 area
  - A DIS of 126 (higher than the default of 64) configured on another IS in this area sets it up as the backup DIS

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

- Enable BFD for fast failure detection
  - BFD helps reduce the convergence times of IS-IS because link failures will be signalled much quicker

```
interface gigabitethernet0/1
 bfd interval 250 min_rx 250 multiplier 3
```

Level-2 interface

- This interface is used for a trunk link to another PoP forming part of your network-wide backbone
  - As such it will be a Level-2 interface, making this router a Level-1/Level-2 IS.
  - Metric and authentication are all configured for Level-2

```
interface gigabitethernet0/2
 ip address 192.168.2.1 255.255.255.252
 ipv6 address 2001:db8:192:168:2:1:/126
 ip router isis 1
 ipv6 router isis 1
 isis circuit-type level-2-only
 isis metric 400 level-2
 isis ipv6 metric 400 level-2
 isis authentication mode md5 level-2
 isis authentication key-chain isis-level2 level-2
```

Level 2 interface: more details

- To make this IS-IS BCP more interesting, we will assume this trunk link is a broadcast multi-access link, i.e., Ethernet.
- As this is an Ethernet interface, IS-IS will attempt to elect a DIS when it forms an adjacency
  - Because it is running as a point-to-point WAN link, with only 2 IS's on the wire, configuring IS-IS to operate in "point-to-point mode" scales the protocol by reducing the link failure detection times
  - Point-to-point mode improves convergence times on Ethernet networks because it:
    - Prevents the election of a DIS on the wire,
    - Prevents the flooding process from using CSNP's for database synchronization
    - Simplifies the SPF computations and reduces the IS's memory footprint due to a smaller topology database.

```
int gi0/2
 isis network point-to-point
```

ISIS Process

- Configure parameters specific to the IS-IS routing protocol
  - This covers both IPv4 and IPv6, as IS-IS supports both IP protocols in the same implementation

```
router isis as100
```

ISIS Process

- Create an NET for the ISIS process:
  - This is made up of:
    - a private AFI (49)
    - an area part
    - a System ID (taken from the padded Loopback interface IP address), and
    - an N-SEL of zero (0)

```
net 49.0001.1921.6800.0001.00
```

ISIS Process

- Under the ISIS process enable HMAC-MD5 authentication:

```
authentication mode md5
authentication key-chain isis-level1 level-1
authentication key-chain isis-level2 level-2
```

Advanced Configuration

- Enable iSPF (incremental SPF).
  - This, in the long run, reduces CPU demand because SPF calculations are run only on the affected changes in the SPT.
  - As this is a Level-1/Level-2 router, enable iSPF at both levels 60 seconds after the command has been entered into the configuration.
  - Note that IOS only supports iSPF for IPv4.

```
ispf level-1-2 60
```

Advanced Configuration

- Enable wide/extended metric support for IS-IS.
  - IOS, by default, supports narrow metrics, which means you can define cost values between 1-63. This is not scalable.
  - To solve this problem, enable wide metrics, which allows you to define cost values between 1-16777214.

```
metric-style wide
```

Advanced Configuration

- Increase ISIS default metric
  - Default value is 10
  - All interfaces in both L1 and L2 have this value
  - Not useful if configured value is “accidentally” removed - a low priority interface could end up taking full load by mistake
  - Configure a “very large” value as default

```
metric 100000
```

Advanced Configuration

- Disable IIH padding
  - On high speed links, it may strain huge buffers
  - On low speed links, it may waste bandwidth and affect other time sensitive applications, e.g., voice.
  - Disabling IIH padding is safe because IOS will still pad the first 5 IIH's to the full MTU to aid in the discovery of MTU mismatches

```
no hello padding
```

Advanced Configuration

- Allow the Loopback interface IP address to be carried within IS-IS, while preventing it from being considered in the flooding process.

```
passive-interface Loopback0
```

- Log changes in the state of the adjacencies.

```
log-adjacency-changes
```

Advanced Configuration

- Tell the IS to ignore LSP's with an incorrect data-link checksum, rather than purge them
  - Purging LSP's with a bad checksum causes the initiating IS to regenerate that LSP, which could overload the IS if perpetuated in a cycle
  - So rather than purge them, ignore them.

```
ignore-lsp-errors
```

Advanced Configuration

- Reduce the amount of control traffic, conserving CPU usage for generation and refreshing of LSP's.
  - Do this by increasing the LSP lifetime to its limits.

```
max-lsp-lifetime 65535
```

- Reduce the frequency of periodic LSP flooding of the topology, which reduces link utilization
  - This is safe because there other mechanisms to guard against persistence of corrupted LSP's in the LSDB.

```
lsp-refresh-interval 65000
```

Advanced Configuration

- Customize IS-IS throttling of SPF calculations.
  - Good for when you also use BFD for IS-IS.
  - These are recommended values for fast convergence.

```
spf-interval 5 1 20
```

- Customize IS-IS throttling of LSP generation.
  - These are recommended values for fast convergence.

```
lsp-gen-interval 5 1 20
```

Advanced Configuration

- Customize IS-IS throttling of PRC calculations.
  - PRC calculates routes without performing a full SFP calculation.
  - This is done when a change is signaled by another IS, but without a corresponding change in the basic network topology, e.g., the need to reinstall a route in the IS-IS RIB.
  - These are recommended values for fast convergence.

```
prc-interval 5 1 20
```

Advanced Configuration

- Enable IS-IS fast-flooding of LSP's.
  - This tells the IS to always flood the LSP that triggered an SPF before the router actually runs the SPF computation.
  - This command used to be 'ip fast-convergence' and has since been replaced from IOS 12.3(7)T.
  - Below, we shall tell the IS to flood the first 10 LSP's which invoke the SPF before the SPF computation is started

```
fast-flood 10
```

Advanced Configuration

- Enable IS-IS IETF Graceful Restart.
  - This ensures an IS going through a control plane switchover continues to forward traffic as if nothing happened
  - Software and platform support is limited, so check whether your particular platform/code supports this
  - Also, deploy only if it's necessary.

```
nsf ietf
```

Advanced Configuration

- Enable BFD support for IS-IS.
  - With BFD running on the interface, a failure of the link would signal IS-IS immediately
  - IS-IS will then converge accordingly.

```
bfd all-interfaces
```

Advanced Configuration

- Tell IS-IS to ignore the attached bit
  - The Attached bit is set when an L1/L2 IS learns L1 routes from other L1 routers in the same area
  - The Attached bit causes the installation of an IS-IS-learned default route in the IS-IS RIB on L1 routers in the same area, as well as in the forwarding table if IS-IS is the best routing protocol from which the default route was learned – this can lead to suboptimal routing.

```
ignore-attached-bit
```

Advanced Configuration

- Wait until iBGP is running before providing transit path

```
set-overload-bit on-startup wait-for-bgp
```

  - Avoids blackholing traffic on router restart
  - Causes ISIS to announce its prefixes with highest possible metric until iBGP is up and running
  - When iBGP is running, ISIS metrics return to normal, make the path valid

Advanced Configuration

- Enable the IPv6 address family for in IS-IS.

```
address-family ipv6
```

- Enable multi-topology support for IPv6 in IS-IS.
  - Multi-topology support allows the IPv4 network topology to be independent of that of IPv6

```
multi-topology
```

Advanced Configuration

- Things to consider on routers operating as Level-1-only IS's:
  - IS-IS BCP techniques under the IS-IS routing process
  - In addition to the interface, tell the IS-IS routing process to operate in a Level-1 area only

```
router isis 1
 is-type level-1
```

Advanced Configuration

- Things to consider on routers operating as Level-1 and Level-2 IS’s:
  - To prevent sub-optimal routing of traffic from L1 IS's in one area to L1 IS's in another area, configure and enable Route Leaking on L1/L2 routers that form the backbone connectivity between two or more different areas
  - Route Leaking permits L1/L2 routers to install L1 routes learned from one area into L1 IS’s routing/forwarding tables in another area
  - This allows for reachability between L1 routers located behind L1/L2 routers in different areas

```
router isis 1
 redistribute isis ip level-2 into level-1 route-map FOO
!
ip prefix-list foo permit 0.0.0.0/0 le 32
!
route-map FOO permit 10
 match ip address prefix-list foo
```

Advanced Configuration

- Doing the same for IPv6:

```
router isis 1
 address-family ipv6
  redistribute isis level-2 into level-1 route-map FOO6
!
ip prefix-list foo6 permit ::/0 le 128
!
route-map FOO6 permit 10
 match ipv6 address prefix-list foo6
!
```

Advanced Configuration

- Summary
  - Best practice recommendations are commonly implemented on many ISP backbones
  - Ensures efficient and scalable operation of ISIS

ISIS for ISPs

AfNOG 2011 AR-E Workshop