IS-IS

Scalable Infrastructure
Workshop

AfNOG 2010

Why IS-IS?

- Link State IGP Protocol
- Over CLNP (ISO protocol) not IP, so harder to attack
- Very simple to configure
- All the power of OSPF and you can get as many knobs if you want, but don't
- Almost no one uses IS-IS, only the world's largest ISPs (and clueful enterprises)
- IPv6 Support is natural
- And it does not lock you in to a vendor

Converting to/from IS-IS

- IGPs are 'ships in the night' i.e. they are quite independent
- You want to convert from OSPF to IS-IS?
- Leave OSPF configuration as is
- Configure IS-IS
- Carefully inspect IS-IS database etc
- Turn off OSPF
- AOL did it without dropping a packet, see http://nanog.org/meetings/nanog29/abstracts.php?pt=Njg2Jm5hbm9nMjk=&nm=nanog29

Simple (& real) IS-IS Configuration

```
router isis RGnet
 net 47.0042.0001.0000.0000.0004.0002.1981.8015.0000.00
 is-type level-2
 log-adjacency-changes
 metric-style wide
 passive-interface Loopback 0

interface Loopback 0
 ip address 198.180.150.252 255.255.255.255
 ipv6 address 2001:418:8006::252/121

interface GigabitEthernet 0/1
 ip address 198.180.150.121 255.255.255.128
 ipv6 address 2001:418:8006::121/121
 isis circuit-type
 ip router isis RGnet
 isis metric 1 level-2
 isis circuit-type level-2
```

IS-IS Standards History

- ISO 10589 specifies OSI IS-IS routing protocol for CLNS traffic
  - Tag/Length/Value (TLV) options to enhance the protocol
  - A Link State protocol with a 2 level hierarchical architecture.
- RFC 1195 added IP support
  - I/IS-IS runs on top of the Data Link Layer
  - Requires CLNP to be configured
- Internet Draft defines how to add IPv6 address family support to IS-IS

  www.ietf.org/internet-drafts/draft-ietf-isis-ipv6-07.txt
- Internet Draft introduces Multi-Topology concept for IS-IS

  www.ietf.org/internet-drafts/draft-ietf-isis-wg-multi-topology-12.txt

Very Large Scale IS-IS Design

- When you have over 200+ routers

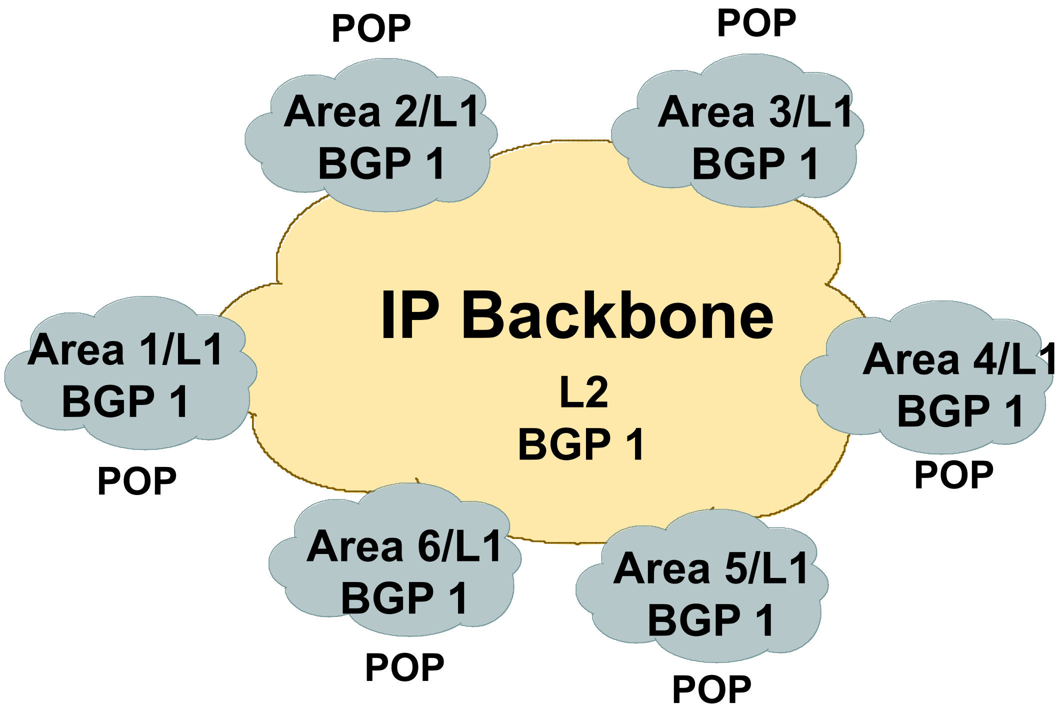

IS-IS Levels

- IS-IS has a 2 layer hierarchy
  - Level-2 (the backbone)
  - Level-1 (the areas)
- A router can be
  - Level-1 (L1) router
  - Level-2 (L2) router
  - Level-1-2 (L1L2) router

IS-IS Levels

- Level-1 router
  - Has neighbours only on the same area
  - Has a level-1 Link State Data Base (LSDB) with all routing information for the area
- Level-2 router
  - May have neighbours in the same or other areas
  - Has a Level-2 LSDB with all routing information about inter-area
- Level-1-2 router
  - May have neighbours on any area.
  - Has two separate LSDBs: level-1 LSDB & level-2 LSDB

Backbone & Areas

- IS-IS does not have a backbone area as such (like OSPF)
- Instead the backbone is the contiguous collection of Level-2 capable routers
- IS-IS area borders are on links, not routers
- Each router is identified with Network Entity Title (NET)
  - NET is an NSAP where the n-selector is 0

L1, L2, and L1L2 Routers

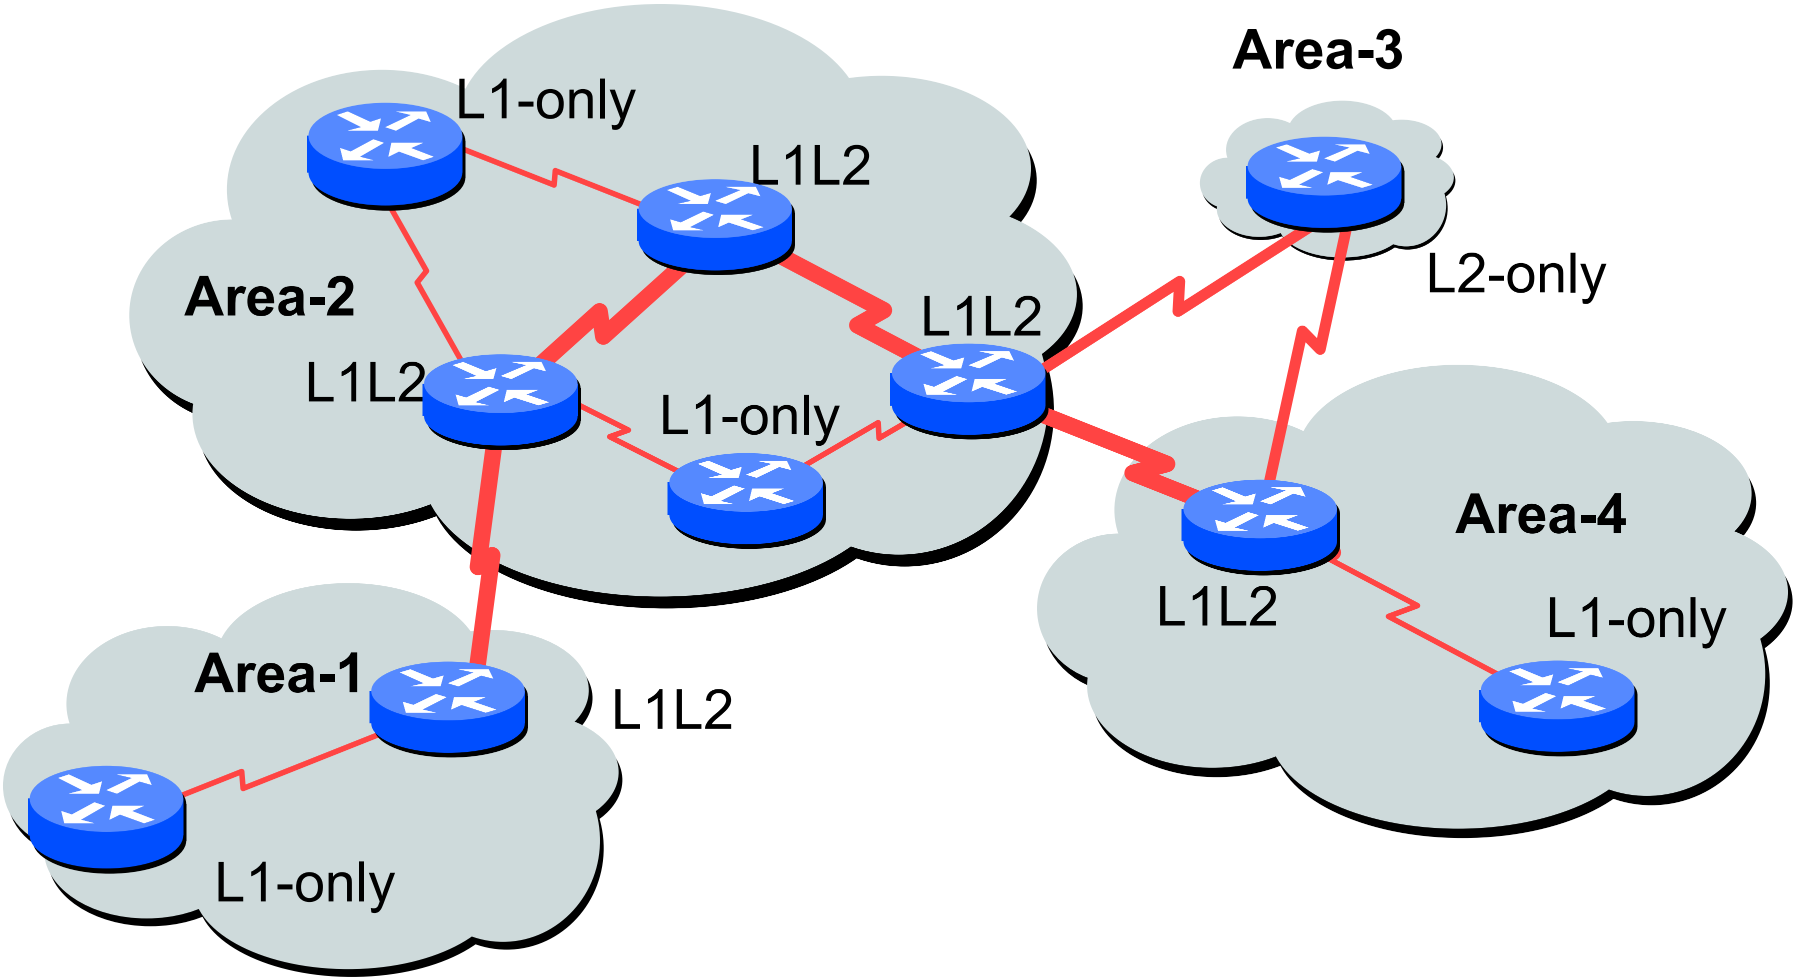

NSAP and Addressing

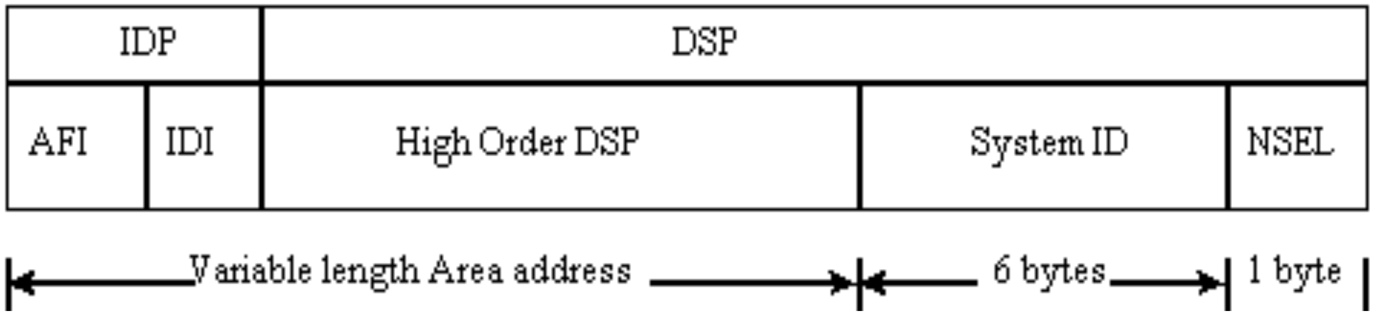

- NSAP: Network Service Access Point
  - Total length between 8 and 20 bytes
  - Area Address: variable length field (up to 13 bytes)
  - System ID: defines an ES or IS in an area.
  - NSEL: N-selector. identifies a network service user (transport entity or the IS network entity itself)
- NET: the address of the network entity itself

Addressing Common Practices

- ISPs typically choose NSAP addresses thus:
  - First 8 bits – pick a number
  - 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

An Addressing Example

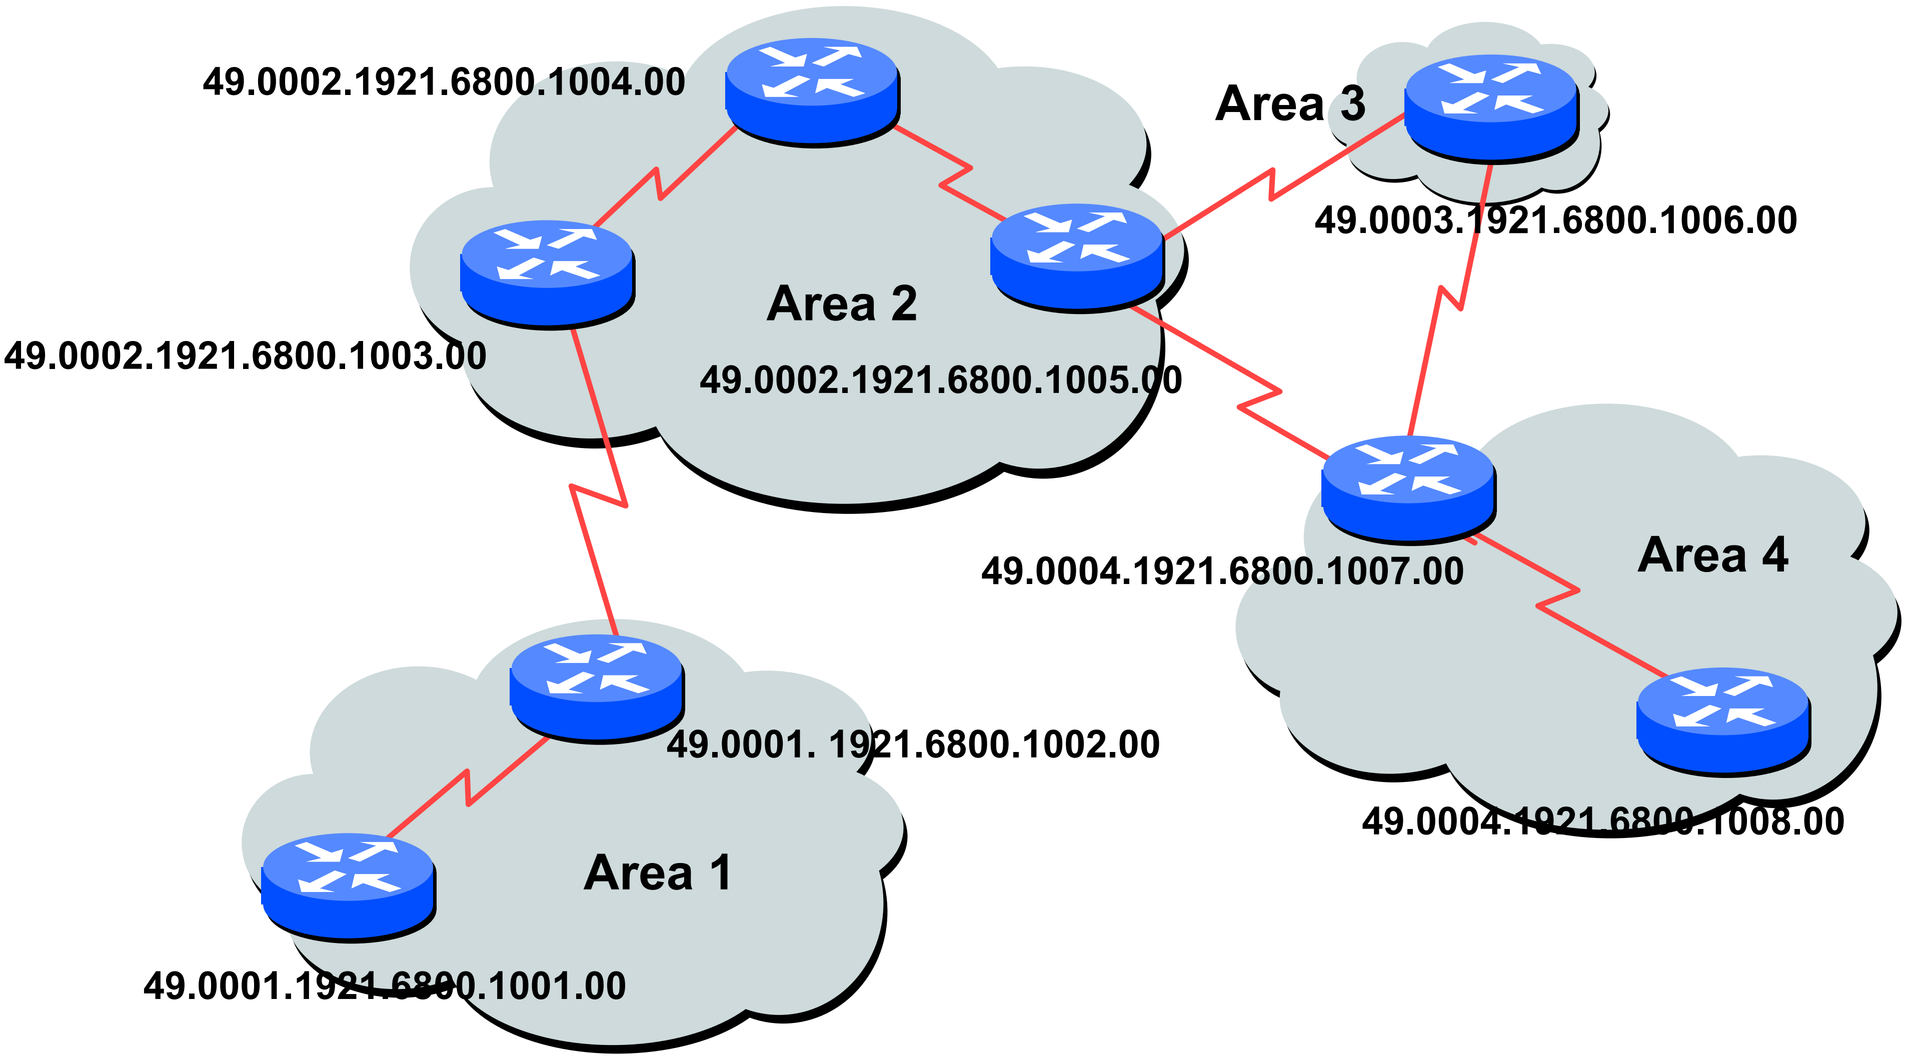

Adjacencies

- Hello PDU IIHs are exchanged between routers to form adjacencies

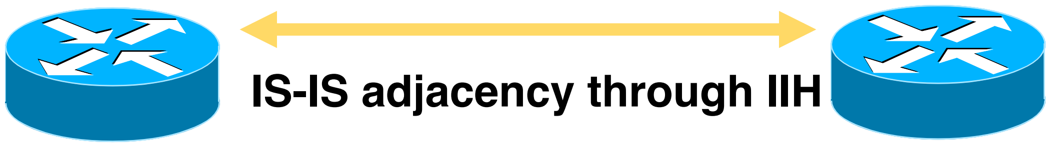

- Area addresses are exchanged in IIH PDUs

Link State PDU (LSP)

- Each router creates an LSP and floods it to neighbours
- A level-1 router will create level-1 LSP(s)
- A level-2 router will create level-2 LSP(s)
- A level-1-2 router will create
  - level-1 LSP(s) and
  - level-2 LSP(s)

LSP Header

- LSPs have
  - Fixed header
  - Type-Length-Value (TLV) coded contents
- The LSP header contains
  - LSP-id
  - Sequence number
  - Remaining Lifetime
  - Checksum
  - Type of LSP (level-1, level-2)
  - Attached bit
  - Overload bit

LSP Contents

- The LSP contents are coded as TLV (Type, Length, Value)
  - Area addresses
  - IS neighbors
  - Authentication Info

LSDB content

- Each router maintains a separate Link State Database (LSDB) for level-1 and level-2 LSPs
- LSP headers and contents
- SRM bits: set per interface when router has to flood this LSP
- SSN bits: set per interface when router has to send a PSNP for this LSP

Flooding of LSPs

- New LSPs are flooded to all neighbors
- It is necessary that all routers get all LSPs
- Each LSP has a sequence number
- 2 kinds of flooding
  - Flooding on a p2p link
  - Flooding on LAN

Flooding on a p2p link

- Once the adjacency is established both routers send CSNP packet
- Missing LSPs are sent by both routers if not present in the received CSNP
- Missing LSPs may be requested through PSNP

Flooding on a LAN

- There’s a Designated IS-IS Router (DIS)
- DIS election is based on priority
  - Best practice is to select two routers and give them higher priority – then in case of failure one provides deterministic backup to the other
- Tie break is by the highest MAC address
- DIS has two tasks
  - Conducting the flooding over the LAN
  - Creating and updating a special LSP describing the LAN topology (Pseudonode LSP)
- Pseudo-node represents LAN (created by the DIS)

Flooding on a LAN

- DIS conducts the flooding over the LAN
- DIS multicasts CSNP every 10 seconds
- All routers in the LAN check the CSNP against their own LSDB (and may ask specific re-transmissions with PSNPs)

Complete Sequence Number PDU

- Describes all LSPs in your LSDB (in range)
- If LSDB is large, multiple CSNPs are sent
- Used at 2 occasions
  - Periodic multicast by DIS (every 10 seconds) to synchronise LSDB over LAN subnets
  - On p2p links when link comes up

Partial Sequence Number PDUs

- PSNPs Exchanged on p2p links (ACKs)
- Two functions
  - Acknowledge receipt of an LSP
  - Request transmission of latest LSP
- PSNPs describe LSPs by its header
  - LSP identifier
  - Sequence number
  - Remaining lifetime
  - LSP checksum

Configuration

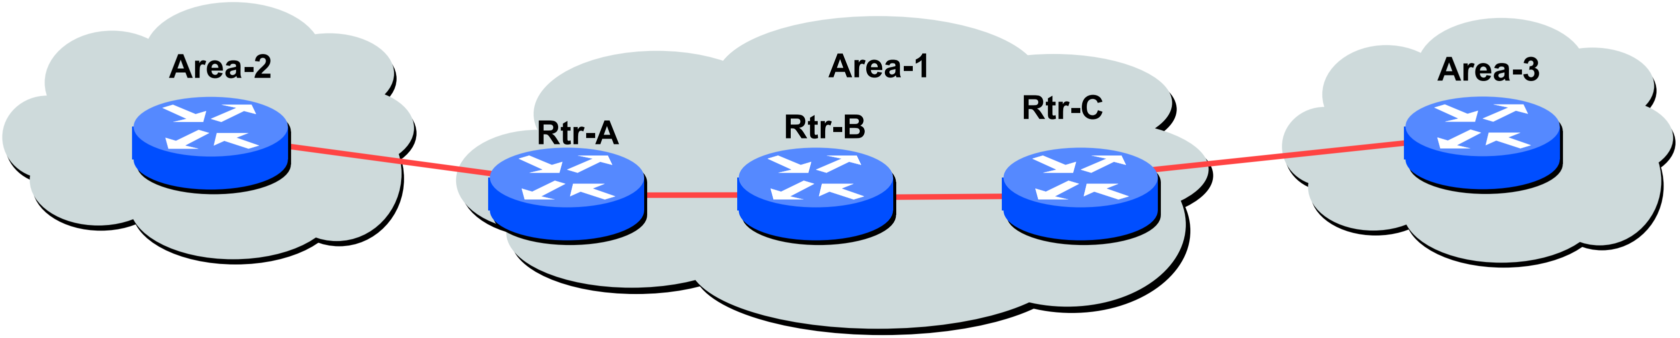

- L1, L2, L1-L2
  - By default Cisco routers will be L1L2 routers
  - Routers can be manually configured to behave as
    - Level-1 only, Level-2 only, Level-1-2
    - This is what most ISPs and enterprises/campuses do
  - Configuration can be done per interface or at the router level

Configuration for A&B

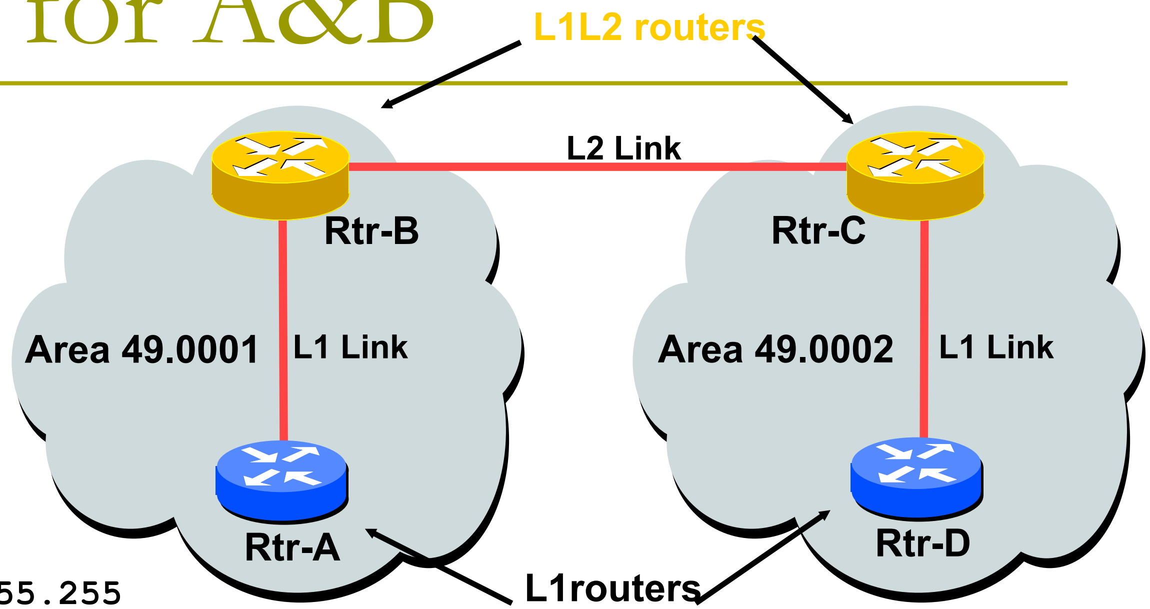

Router-B

```
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
!
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

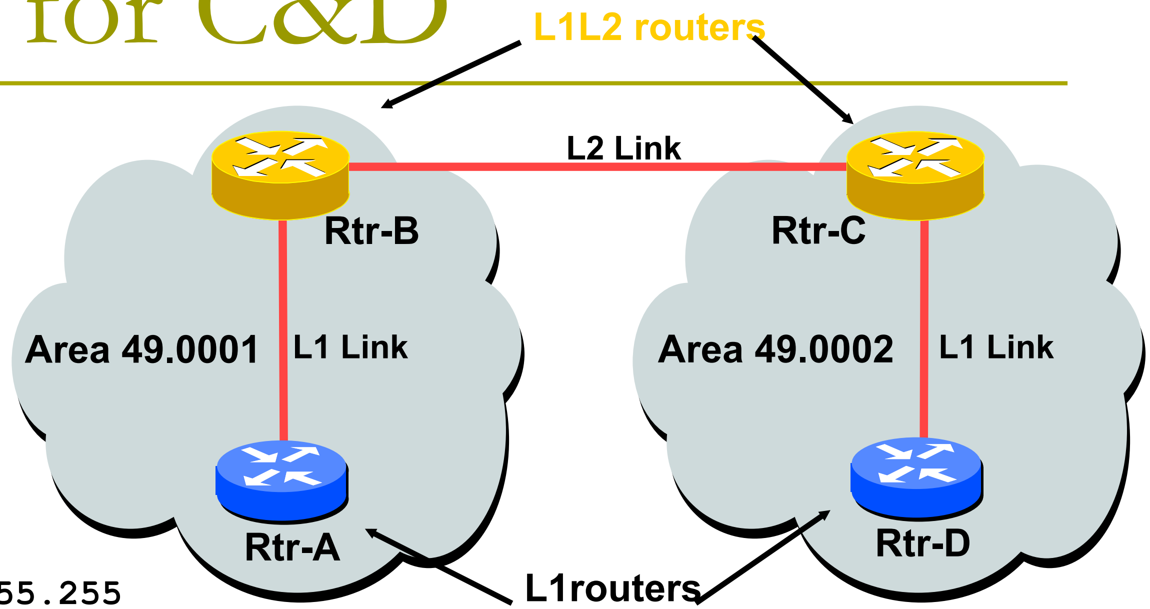

Router-C

```
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
```

Adding interfaces to IS-IS

- To activate IS-IS on an interface:
  - **`interface FastEthernet 4/0`**
  - **` ip route isis isp-bb`**
  - **` isis circuit-type level-2`**
- To disable IS-IS 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 IS-IS configuration on an interface
  - No CLNS run on interface, no interface subnet in the LSDB

Adding interfaces to IS-IS

- Scaling IS-IS: **`passive-interface default`**
  - Disables IS-IS processing on all interfaces apart from those marked as no-passive
  - Places all IP addresses of all connected interfaces into IS-IS
  - 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
 isis metric 1 level-2
```

Status Commands in IS-IS

- 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 IS-IS

- 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 – ease of troubleshooting

Status Commands in IS-IS

- 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 IS-IS

- 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 IS-IS 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
- IS-IS 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
    - so trading efficiency for very very large scale
    - areas have only single exit points
  - Start with L2-only everywhere is a good choice
  - Future implementation of level-1 areas will be easier
  - Backbone continuity is ensured from start

IS-IS for IPv6

IS-IS for IPv6

- 2 Tag/Length/Values added to introduce IPv6 routing
- IPv6 Reachability TLV (0xEC)
  - External bit
  - Equivalent to IP Internal/External Reachability TLV’s
- IPv6 Interface Address TLV (0xE8)
  - For Hello PDUs, must contain the Link-Local address
  - For LSP, must only contain the non-Link Local address
- IPv6 NLPID (0x8E) is advertised by IPv6 enabled routers

IOS IS-IS dual IP configuration

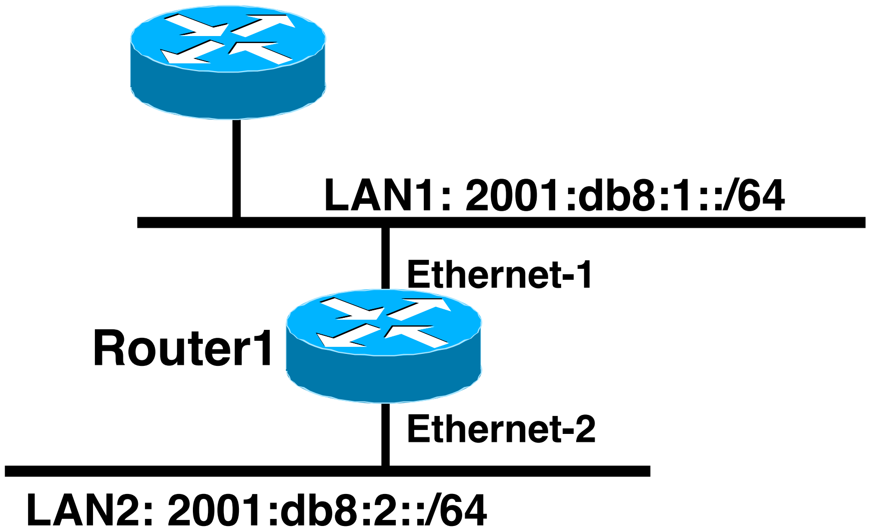

Dual IPv4/IPv6 configuration.
Redistributing both IPv6 static routes
and IPv4 static routes.

```
Router1#
interface ethernet-1
 ip address 10.1.1.1 255.255.255.0
 ipv6 address 2001:db8:1::1/64
 ip router isis
 ipv6 router isis

interface ethernet-2
 ip address 10.2.1.1 255.255.255.0
 ipv6 address 2001:db8:2::1/64
 ip router isis
 ipv6 router isis

router isis
 address-family ipv6
  redistribute static
  exit-address-family
 net 42.0001.0000.0000.072c.00
 redistribute static
```

IOS Configuration for IS-IS for IPv6 on IPv6 Tunnels over IPv4

```
On Router1:
interface Tunnel0
 no ip address
 ipv6 address 2001:db8:1::1/64
 ipv6 address FE80::10:7BC2:ACC9:10 link-local
 ipv6 router isis
 tunnel source 10.42.1.1
 tunnel destination 10.42.2.1
!
router isis
 net 42.0001.0000.0000.0001.00
```

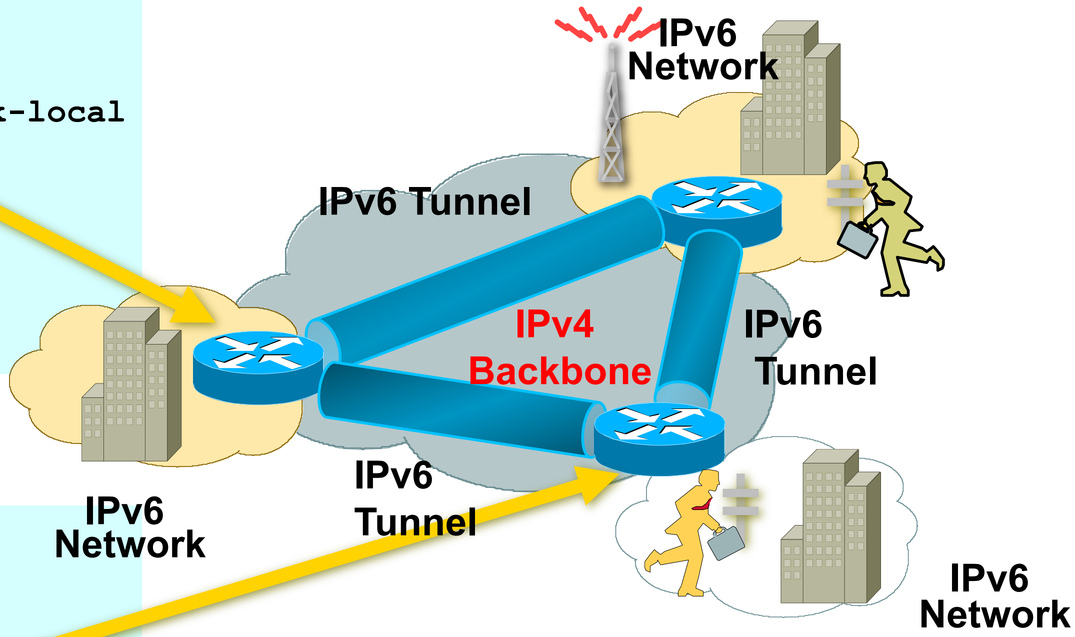


```
On Router2:
interface Tunnel0
 no ip address
 ipv6 address 2001:db8:1::2/64
 ipv6 address FE80::10:7BC2:B280:11 link-local
 ipv6 router isis
 tunnel source 10.42.2.1
 tunnel destination 10.42.1.1
!
router isis
 net 42.0001.0000.0000.0002.00
```

IS-IS for IPv6 on an IPv6 Tunnel requires GRE Tunnel; it can't work with IPv6 configured tunnel as IS-IS runs directly over the data link layer

Multi-Topology IS-IS extensions

- IS-IS for IPv6 assumes that the IPv6 topology is the same as the IPv4 topology
  - Single SPF running, multiple address families
  - Some networks may be like this, but many others are not
- Multi-Topology IS-IS solves this problem
  - New TLV attributes introduced
  - New Multi-Topology ID #2 for IPv6 Routing Topology
  - Two topologies now maintained:
    - ISO/IPv4 Routing Topology (MT ID #0)
    - IPv6 Routing Topology (MT ID #2)

Multi-Topology IS-IS extensions

- New TLVs attributes for Multi-Topology extensions:
  - Multi-topology TLV: contains one or more multi-topology ID in which the router participates
  - MT Intermediate Systems TLV: this TLV appears as many times as the number of topologies a node supports
  - Multi-Topology Reachable IPv4 Prefixes TLV: this TLV appears as many times as the number of IPv4 announced by an IS for a given MT ID
  - Multi-Topology Reachable IPv6 Prefixes TLV: this TLV appears as many times as the number of IPv6 announced by an IS for a given MT ID

Multi-Topology IS-IS configuration example (IOS)

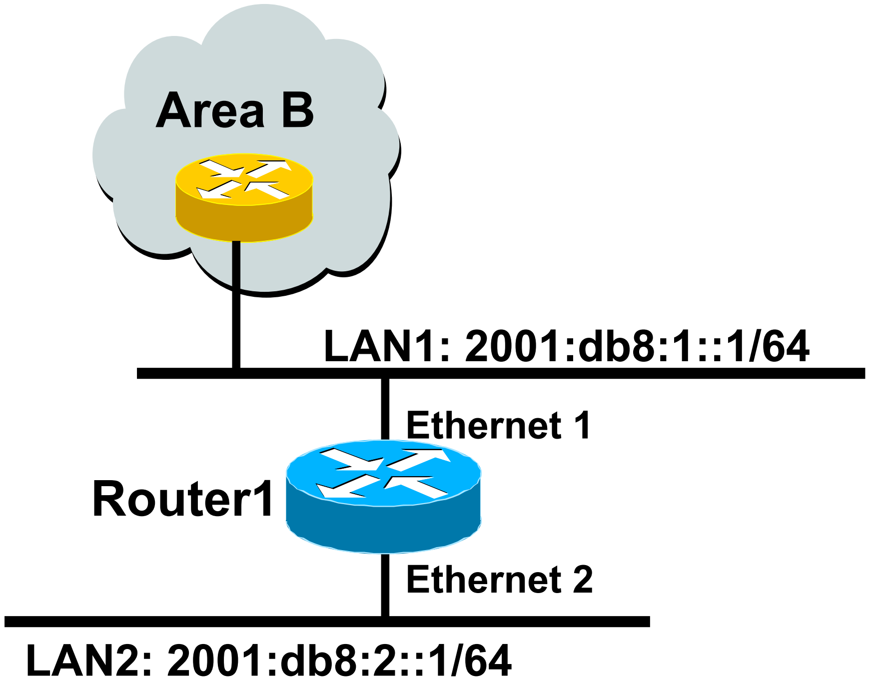

- The optional keyword **`transition`** may be used for transitioning existing IS-IS IPv6 single SPF mode to MT IS-IS
- Wide metric is mandated for Multi-Topology to work

```
Router1#
interface Ethernet 1
 ip address 10.1.1.1 255.255.255.0
 ipv6 address 2001:db8:1::1/64
 ip router isis
 ipv6 router isis
 isis ipv6 metric 20

interface Ethernet 2
 ip address 10.2.1.1 255.255.255.0
 ipv6 address 2001:db8:2::1/64
 ip router isis
 ipv6 router isis
 isis ipv6 metric 20

router isis
 net 42.0001.0000.0000.072c.00
 metric-style wide
 !
 address-family ipv6
  multi-topology
 exit-address-family
```

ISP common practices

- NSAP address construction
  - Area and loopback address
- L2
  - L1-L2 and L1 used later for scaling
- Wide metrics
  - Narrow metrics are too limiting
- Deploying IPv6 in addition to IPv4
  - Multi-topology is recommended – gives increased flexibility should there be future differences in topology

Summary

- You have learned about:
  - IS-IS for IPv4
  - L1, L2 and L1L2 routers
  - IS-IS areas
  - IS-IS configuration and status commands
  - IS-IS extensions for IPv6
  - ISP common practices