

Introduction to ISIS



SI-E Workshop
AfNOG 2013 - Zambia

IS-IS Standards History

- ISO 10589 specification that defines IS-IS as an OSI routing protocol for CLNS traffic
 - A Link State protocol with a 2 level hierarchical architecture
 - With Type/Length/Value (TLV) options for protocol enhancements
- The RFC 1195 added Support for IP
 - Thus called Integrated IS-IS
 - I/IS-IS runs on top of the Data Link Layer or rather L2
 - Requires CLNP (Connectionless Network Protocol) to be configured
- RFC5308 adds IPv6 address family support to IS-IS
- RFC5120 defines Multi-Topology concept for IS-IS
 - Permits IPv4 and IPv6 topologies which are not identical

ISIS Levels

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

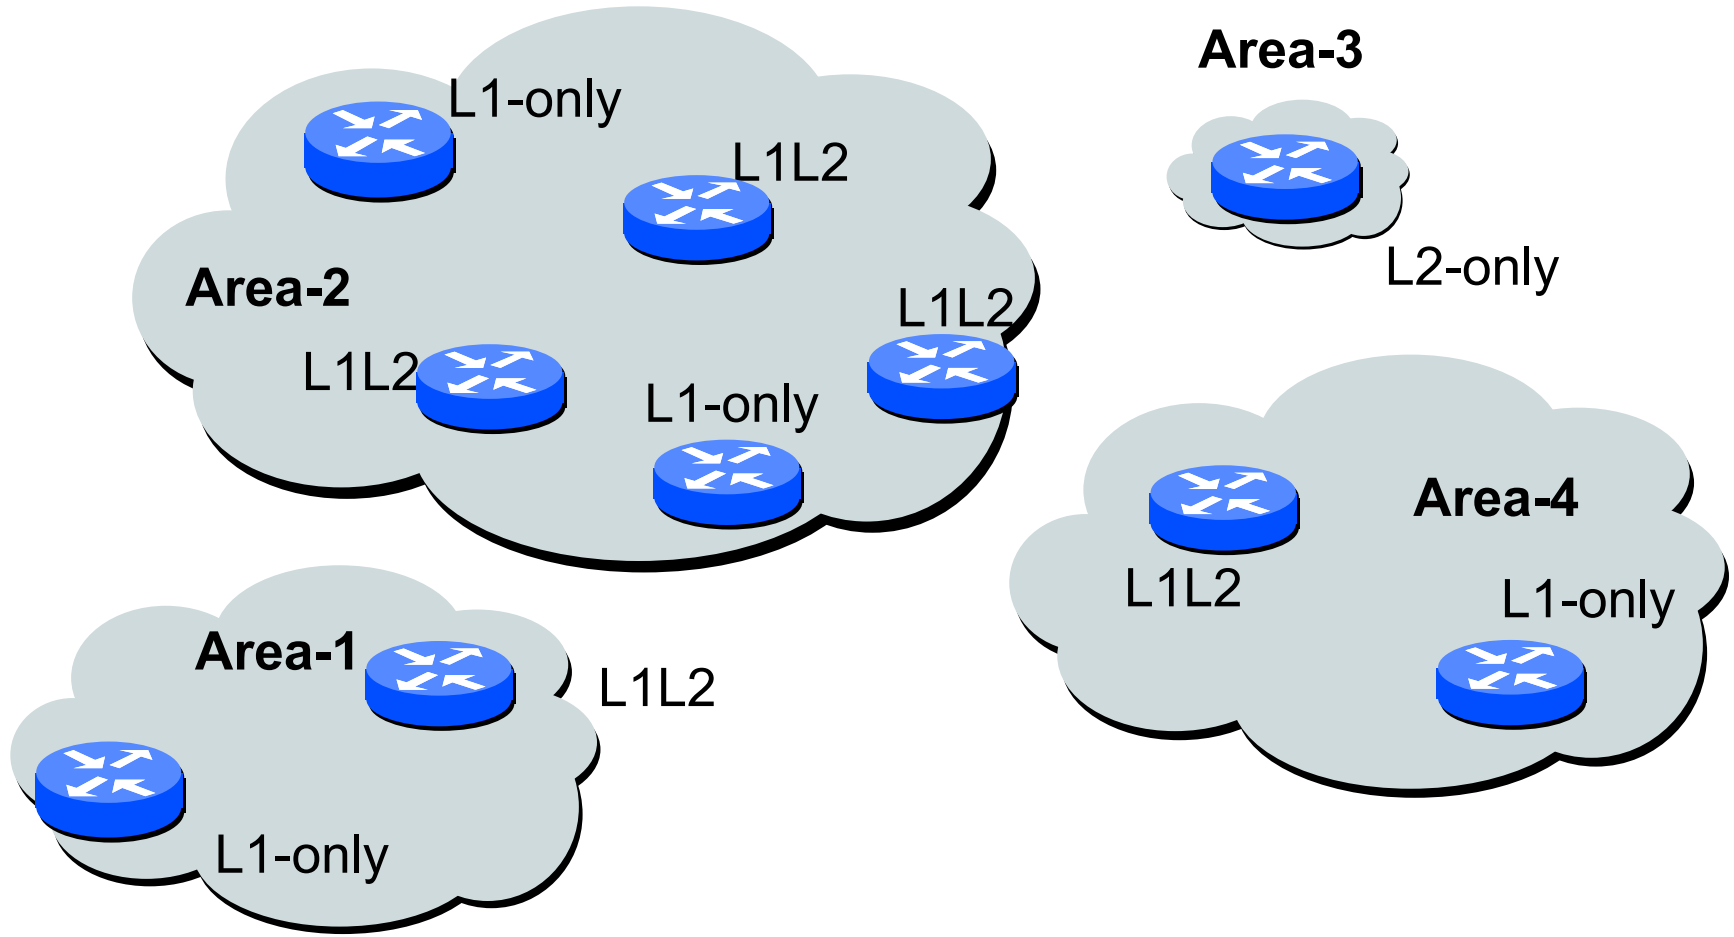
ISIS Levels

- Level-1 router
 - Has neighbours only on the same area
 - Has a level-1 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

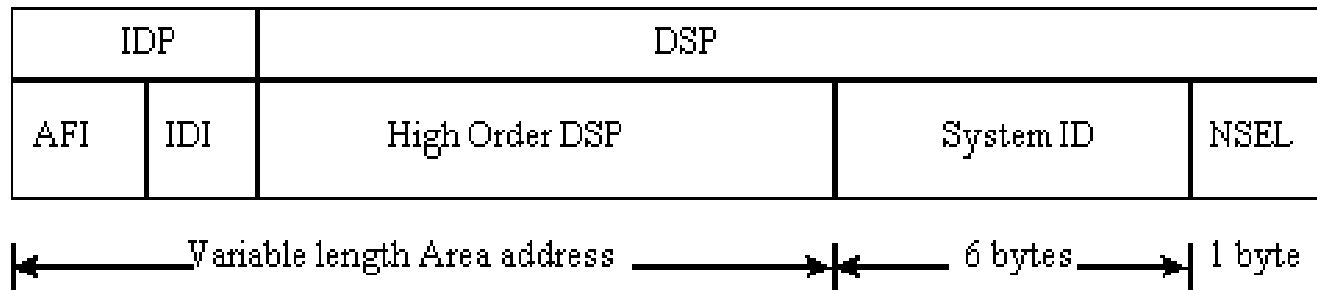
Backbone & Areas

- ❑ ISIS does not have a backbone area as such (like OSPF typical area 0)
- ❑ Instead the backbone is the contiguous collection of Level-2 capable routers
- ❑ ISIS area borders are on the wire or rather links and not routers
- ❑ Each router is identified with a unique **Network Entity Title (NET)**
 - NET is a Network Service Access Point (NSAP) where the n-selector is 0
 - (Compare with each router having a unique Router-ID with IP routing protocols)

L1, L2, and L1L2 Routers



NSAP and Addressing

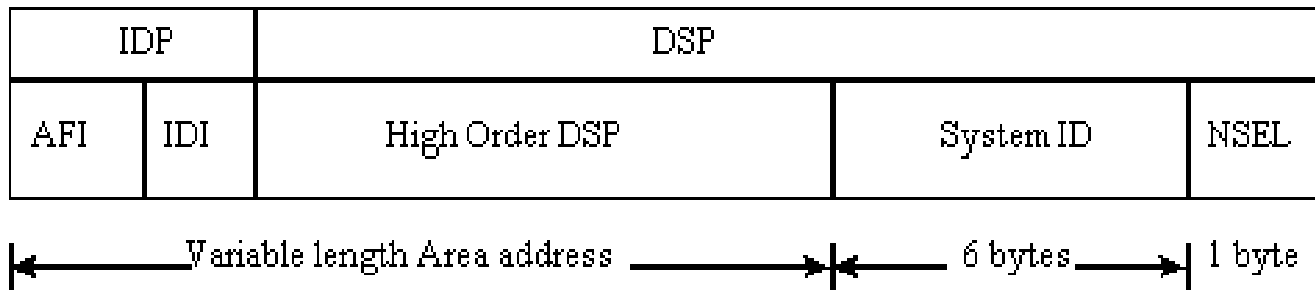


- NSAP: Network Service Access Point
 - Total length between 8 and 20 bytes
 - Area Address: variable length field (up to 13 bytes)
 - System ID: defines either an ES or IS in an area.
 - NSEL: N-selector. identifies a network service
- NET: The address of the network entity itself

Example **47.0001.aaaa.bbbb.cccc.00** Where,

- Area Address = **47.0001**
- SysID = **aaaa.bbbb.cccc**
- Nsel = **00**

NSAP and Addressing

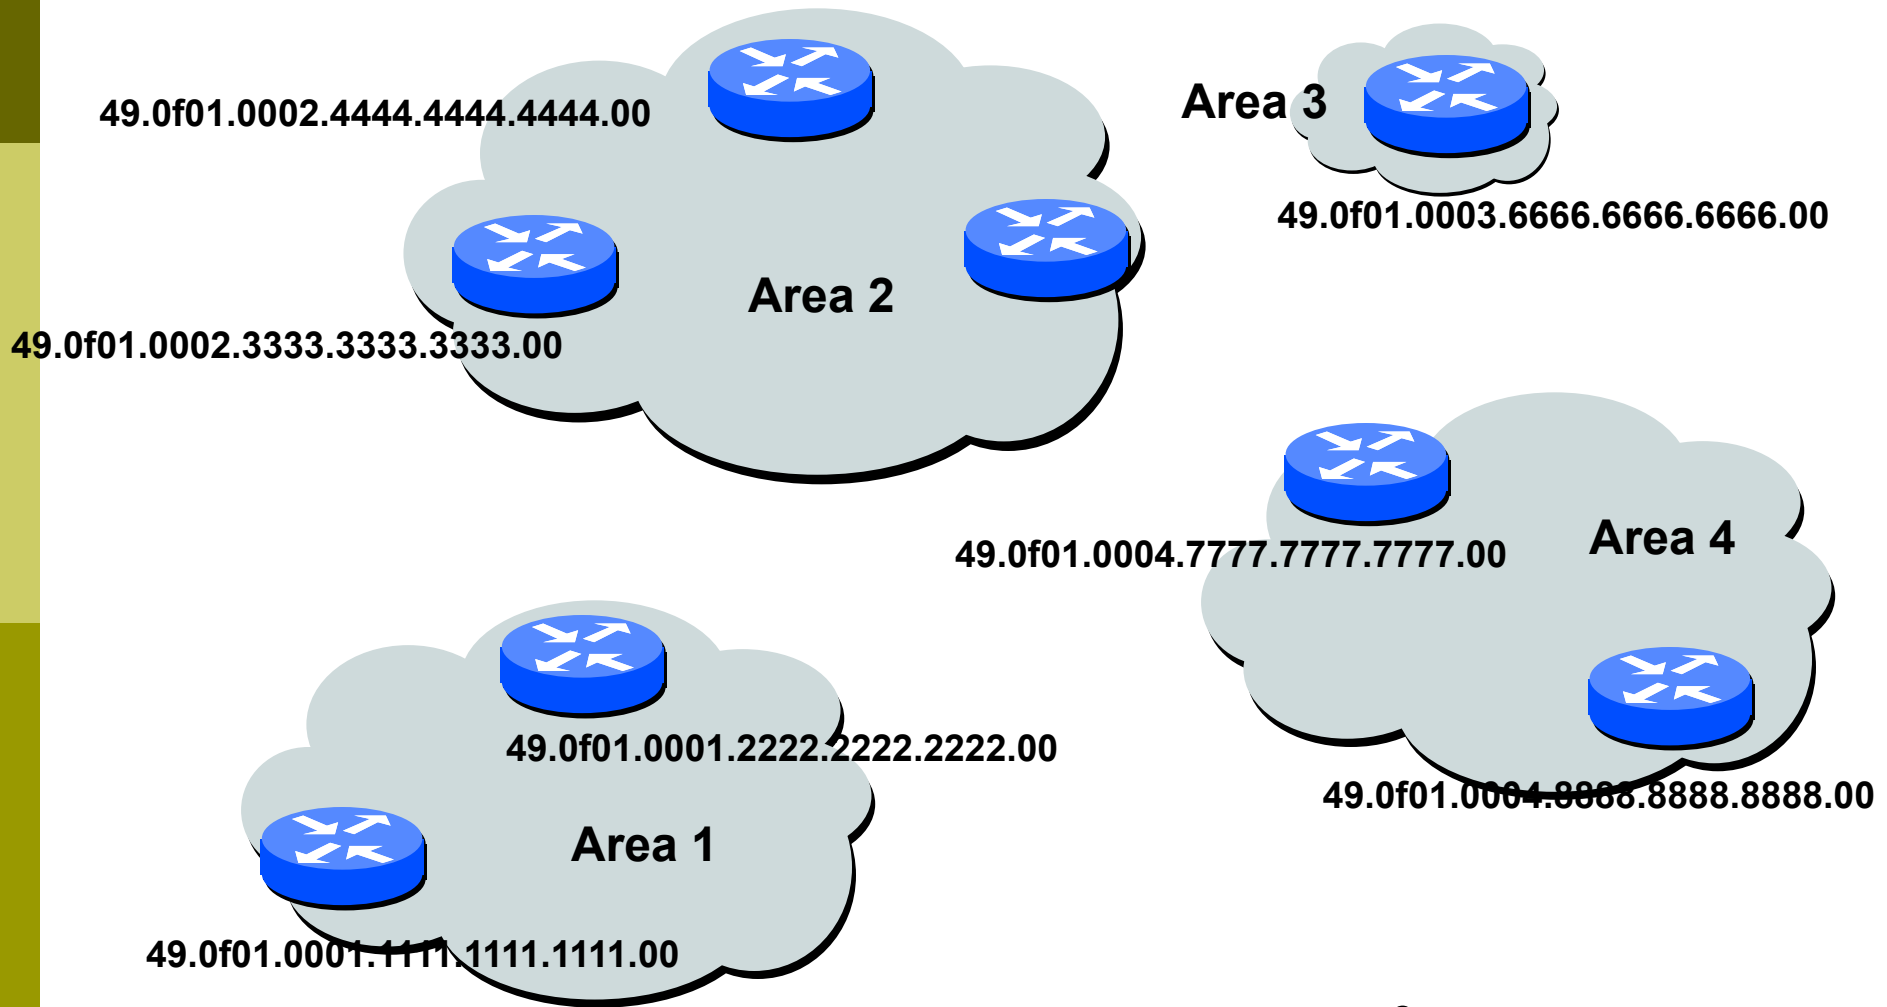


- NSAP: Network Service Access Point
 - Total length between 8 and 20 bytes
 - Area Address: variable length field (up to 13 bytes)
 - System ID: defines either an ES or IS in an area.
 - NSEL: N-selector. identifies a network service
- NET: The address of the network entity itself

Example **47.0001.aaaa.bbbb.cccc.00** Where,

- Area Address = **47.0001**
- SysID = **aaaa.bbbb.cccc**
- Nsel = **00**

Typical NSAP Addressing



Addressing Common Practices


- ISP's typically choose NSAP addresses thus:
 - First 8 bits – pick a number (usually 49)
 - Next 16 bits – area
 - Next 48 bits – router loopback address (BCP)
 - Final 8 bits – zero
- Example:
 - NSAP: 49.0001.1921.6800.1001.00
 - Router: 192.168.1.1 (loopback) in Area 1

Addressing & Design Practices

- ISPs typically use one area (eg.49.0001)
 - Multiple areas only come into consideration once the network is several hundred routers big
- NET begins with 49
 - “Private” address range
- All routers are in L2 only (Core Network)
 - Note: Cisco IOS defaults to L1L2
 - Set L2 under ISIS router configuration (can also be done per interface)

Adjacencies – Hello PDU (IIS)

- Hello Protocol Data Units (PDUs) are exchanged between routers.
- Typically to establish and maintain adjacencies between IS's.



- IS-IS area addresses are also exchanged in this IIH PDUs.
- A PDU is an IS-IS equivalent of a packet

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
 - Independent level-1 LSP(s) and
 - Independent level-2 LSP(s)

Link State Database Content

- Each IS maintains a separate LSDB for either level-1 or level-2 LSPs
- The LSDB contains:
 - LSP headers and contents
 - **SRM** = Send Routing Message
 - **SSN** = Send Sequence Number
 - **SRM bits**: set per interface when a router has to flood an LSP through that interface
 - **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
- All IS's get all LSPs
- Each LSP has a sequence number
- There are 2 kinds of flooding:
 - Flooding on a point to point link and
 - Flooding on a LAN

Flooding on a p2p link

- Once the adjacency is established either routers send CSNP packet.
- And in case of any missing LSP's, if not present in the received CSNP both routers would send a request!!!!
- This is done through a PSNP packet request
- PSNP (Partial Sequence Number PDU)
- CSNP (Complete Sequence Number PDU)

Flooding on a LAN

- Each LAN has a Designated Router (DIS)
- The DIS has two tasks
 - Conducting LSP flooding over the LAN
 - Creating and updating a special LSP describing the LAN topology (Pseudo-node LSP)
- DIS election is based on priority
 - Best practice is to select two routers and give them higher priority
 - Thus, in case of any failure one provides deterministic backup for the other
 - DIS Tie breaker is router with the highest MAC address

Flooding on a LAN Cont...

- ❑ DIS conducts the flooding over the LAN
- ❑ DIS multicasts CSNP every 10 seconds
- ❑ All routers on the LAN check the CSNP against their own LSDB.
- ❑ In case of any missing content within the LSP, the IS may request for specific re-transmissions of upto date LSP's via a PSNP request

Complete Sequence Number PDU

- Used to distribute a routers complete link-state database
- If the LSDB is large, multiple CSNPs are sent
- Used on 2 occasions:
 - Periodic multicast by DIS (every 10 seconds) to synchronise the LSDB over LAN subnets
 - On p2p links when link comes up

Partial Sequence Number PDUs

- Typically exchanged on p2p links, PSNP are used to ack and request link-state info
- 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

Network Design Issues

- ❑ As in all IP network designs, the key issue is the addressing layout
- ❑ ISIS supports a large number of routers in a single area
- ❑ When network is so large requiring the use of areas, employ summary-addresses
- ❑ >400 routers in the backbone is quite doable ... according to Philip Smith :-)

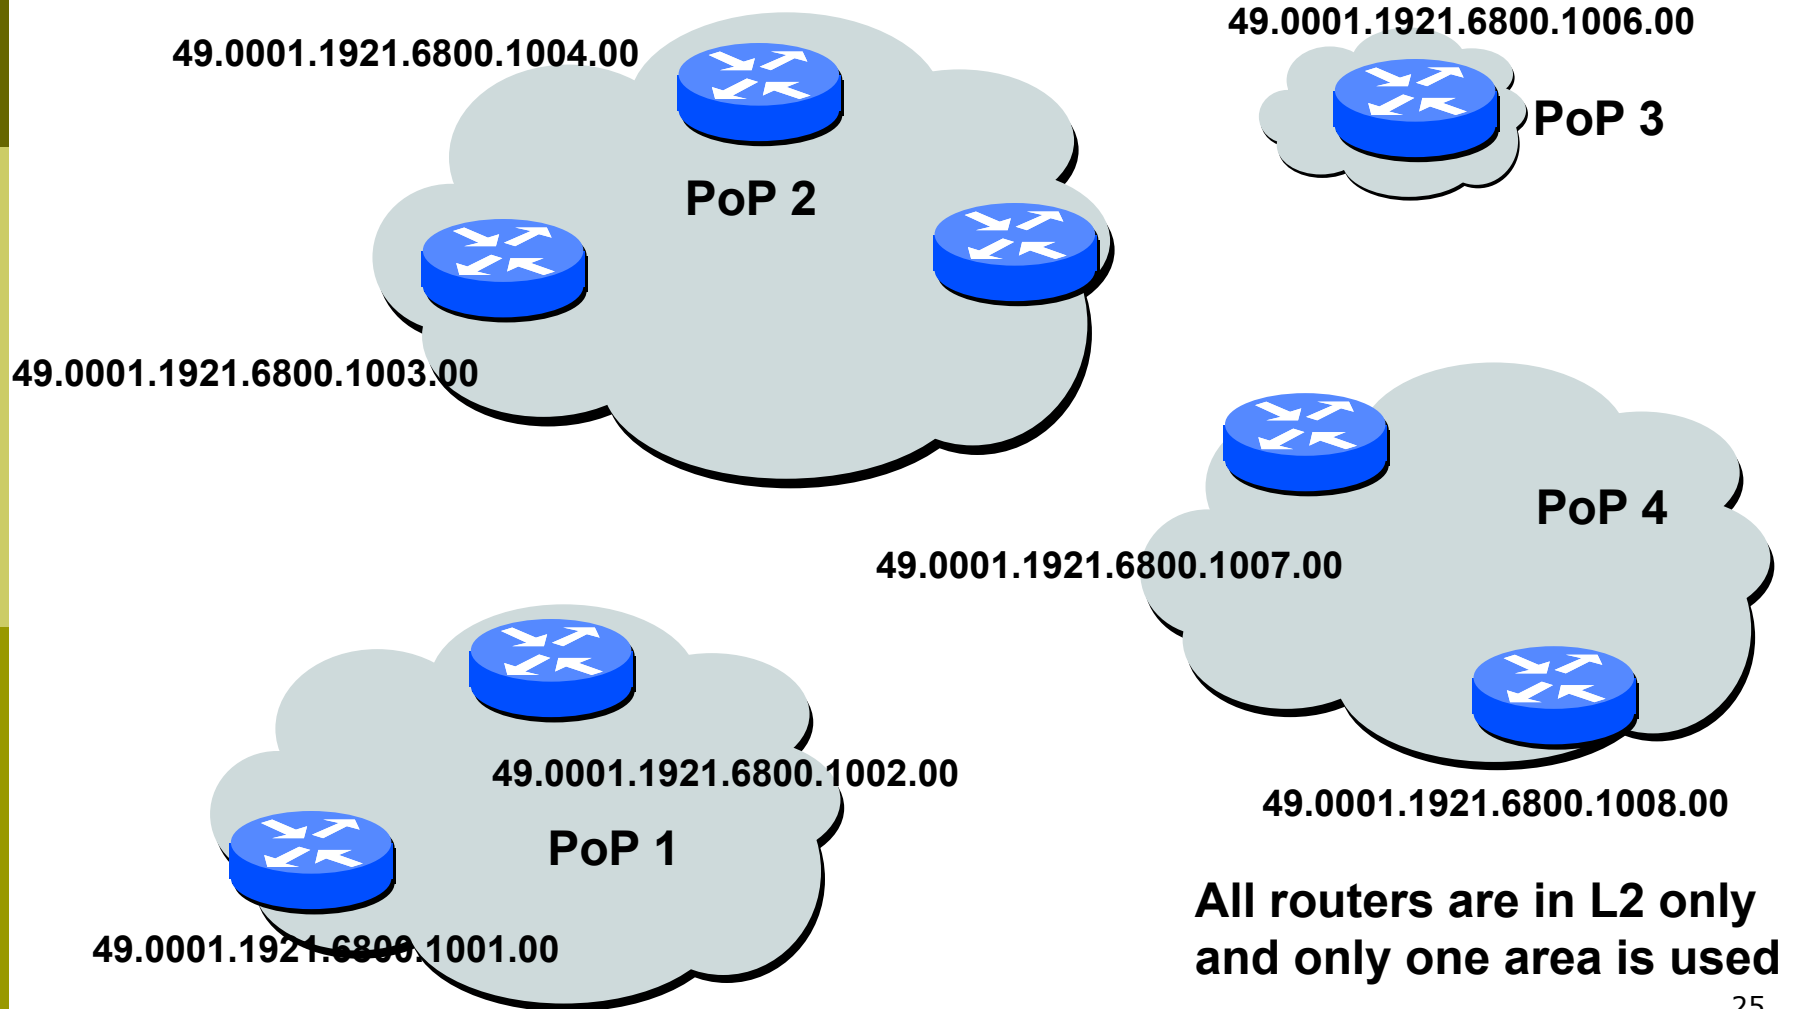
Network Design Issues

- Link cost
 - Default on all interfaces is 10
 - (Compare with OSPF which sets cost according to link bandwidth)
 - Manually configured according to routing strategy
- Summary address cost
 - Equal to the best more specific cost
- Backbone has to be contiguous
 - Ensures continuity through redundancy
- Area partitioning
 - Design in a way 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
 - Thus future implementation of any level-1 areas would become easier
 - Backbone continuity is ensured from start

Typical ISP Design



Questions?

