

Introduction to ISIS



AfNOG 2011 SI-E Workshop

IS-IS Standards History

- ISO 10589 specifies OSI IS-IS routing protocol for 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
 - Requires CLNP 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
 - (Required for an incremental roll-out of IPv6 on existing IPv4 infrastructure)

ISIS Levels

- ISIS 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


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

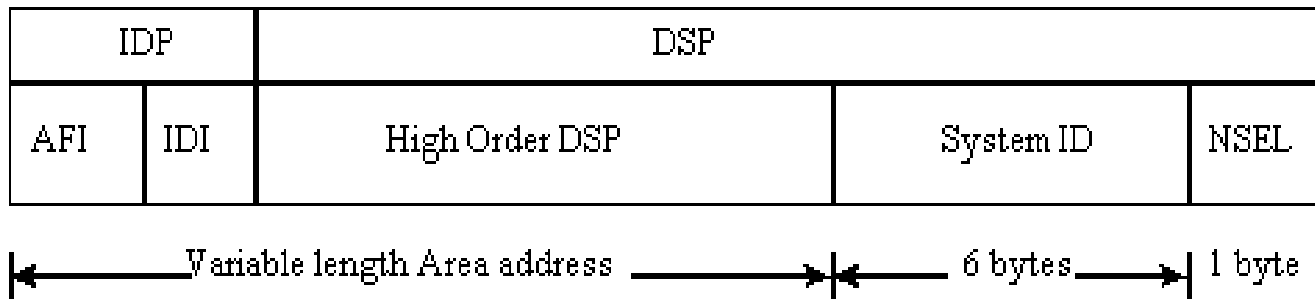
Backbone & Areas

- ISIS does not have a backbone area as such (like OSPF)
- Instead the backbone is the contiguous collection of Level-2 capable routers
- ISIS area borders are on links, 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 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

An Addressing Example



Addressing Common Practices


- 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

Addressing & Design Practices

- ISPs usually only use one area
 - 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
 - Note that Cisco IOS default is L1L2
 - Set L2 under ISIS generic configuration (can also be done per interface)

Adjacencies

- Hello Protocol Data Units (PDUs) are exchanged between routers to form adjacencies



- Area addresses are exchanged in IIH PDUs
 - Intermediate-System to Intermediate System Hello PDUs
 - (PDU is ISIS 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
 - level-1 LSP(s) and
 - level-2 LSP(s)

The ISIS LSP

- LSPs have a Fixed Header and TLV coded contents
- The LSP header contains
 - LSP-id
 - Remaining Lifetime
 - Type of LSP (level-1, level-2)
 - Attached bit
 - Sequence number
 - Checksum
 - Overload bit
- The LSP contents are coded as TLV (Type, Length, Value)
 - Area addresses
 - IS neighbours
 - Authentication Information

Link State Database Content

- Each router maintains a separate LSDB for level-1 and level-2 LSPs
- The LSDB contains:
 - 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
- All routers get all LSPs
- Each LSP has a sequence number
- There are 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

- Each LAN has a Designated Router (DIS)
- The DIS has two tasks
 - Conducting the flooding over the LAN
 - Creating and updating a special LSP describing the LAN topology (Pseudonode LSP)
- 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 for the other
 - Tie break is by the highest MAC address

Flooding on a LAN

- ❑ 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 (and may ask specific re-transmissions with PSNPs)

Complete Sequence Number PDU

- ❑ Describes all LSPs in your LSDB (in range)
- ❑ 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

- 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

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 network is so large requiring the use of areas, use summary-addresses
- ❑ >400 routers in the backbone is quite doable


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
 - 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
 - Future implementation of level-1 areas will be easier
 - Backbone continuity is ensured from start

Typical ISP Design



Introduction to ISIS



AfNOG 2011 AR-E Workshop