S-38.3192 Verkkopalvelujen tuotanto

S-38.3192 Network Service Provisioning

Lecture 3: IS-IS / OSPF

IS-IS versions

- **IS-IS**
  - ISO 10589
  - OSI network layer routing protocol
  - Works only with
    - CLNS / CLNP
      - OSI Network protocol for connectionless services
- **Integrated IS-IS**
  - RFC 1195
  - Integration of IP addresses into OSI CLNP reference model
  - Three operation modes:
    - Pure CLNS
    - Pure IP
    - Mixed CLNS/IP
  - Operation is always based on CLNS model

IS-IS

- Linkstate protocol like OSPF
  - Dijkstra SPF
- **Operates directly on top of link layer**
  - Parallel to IP
    - Pros: Independent of IP addressing
    - Con: Difficulties with ATM null encapsulation (TCP ACK falls within 48 bytes with null encap otherwise 56 bytes -> two cells)
- Uses SNPA:s in L2 addressing
  - Sub Network Point of Attachment (MAC, VC, DLCI)
- Supports
  - Point-to-point links
  - Broadcast links
    - Uses multicast MAC addresses in communication

Addressing

- IS-IS uses OSI addressing at L3
  - Network Service Access Point (NSAP)
    - Also used in ATM, CMIP …
  - Conceptually max 20 byte long address format
    - Different structuring depending on sponsoring organisation and usage

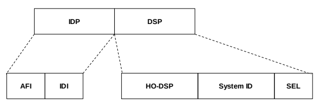

IDP: Initial Domain Part
DSP: Domain Specific Part

AFI: Address Format Identifier (represents coding of address octects)
IDI: Initial Domain Identifier (represents address authority)
HO-DSP: High Order Domain Specific Part (Domain specifig structure of address)
System ID: Unique System Identifier
SEL: Selector (process multiplexing tag)

NSAP Addresses

Address Domain	Authority	AFI
X.121 (X.25)	ITU TSB	36/37
ISO Data Country Code (DCC)	National Standard Bodies	38/39
E.164 (ISDN)	ITU TSB	44/45
ISO 6523 International Code Designator (ICD)	British Standards Institute	46/47
Local		48/49

IS-IS NSAP Addressing

- Special version of NSAP address – **Network Entity Title (NET)**
  - Area ID defines the L2 or L1 area the router belongs to
  - System ID is unique identifier of system within the area
    - Same functionality as OSPF router-ID
  - Selector is internal process ID with IP routers this is always 00

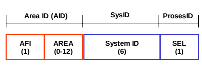

Addressing

- Administration of **SysID is similar to** administration of **Router-ID**
  - It should be unique within domain
    - Duplicates cause problems for SPF
      - SPF is executed over sysIDs not over prefixes
  - It should be easily understandable (helps troubleshooting)
  - One convention is to use IP address of the loopback interface as the source for this information (like in OSPF)
- Loopback IP address: 10.100.100.4
  - Hex encoded zeroes at front: 0000.0A64.6404
  - Hex encoded zeroes at end: 0A64.6404.0000
  - Positional decimal coding: 0101.0010.0004
  - Direct decimal coding: 1010.0100.4000

Addressing

- **AreaID (AID) is coded in the frontpart of the NET**
  - First octet is AFI but with NET it does not have general meaning
    - NET addresses are not visible outside the domain
    - Good practise to use local AFIs 0x49
  - The rest is actual area identifier
    - Upto 12 bytes (if present)
    - No special area identifiers
      - OSPF: backbone – 0 or 0.0.0.0
      - IS-IS: L2 can be which ever area ID

Area

- IS-IS area is determined by router NET
  - Router can have multiple NETs
    - Resulting multiple parallel adjacencies
  - Area boundary is the logical boundary defined by the adjacencies
    - There either is or not adjacency over particular link
- Two types of areas
  - L2 e.g. Backbone
  - L1 e.g. Non-backbone
- Three types of routers
  - L1 router
  - L2 router
  - L1/L2 router

R1 area type	R2 area type	AIDs	Adjacency
L1	L1	Same	L1
L1	L1	Different	-
L2	L2	Same	L2
L2	L2	Different	L2
L1	L2	Same	-
L1	L2	Different	-
L1	L1/L2	Same	L1
L1	L1/L2	Different	-
L2	L1/L2	Same	L2
L2	L1/L2	Different	L2
L1/L2	L1/L2	Same	L1 + L2
L1/L2	L1/L2	Different	L2

Broadcast support

- **Problem of scaling the number of adjacencies**
  - But is this really a problem
    - With OSPF: YES
    - With IS-IS: NO
  - Broadcast networks are just not that nice in SPF calculations due to ghost nodes
    - Make broadcast links point-to-point when it is used as such
  - Adjacency is formed between each router which fulfills area requirements on the same link
    - Exponential increase of adjacencies over the number of nodes sharing the link and area
  - Pseudonode emulation
    - Link is emulated as a node with zero cost
    - Implemented as designated router (DR) / Designated Intermediate System (DIS)
      - Reflector of routing information
      - Adjacencies only to DR / DIS

Broadcast support

- Pseudonode emulation
  - Lowers the amount of link state traffic
    - Saves processing
    - Saves link capacity
  - Adds ghost into SPF calculation
    - Complicates ECMP

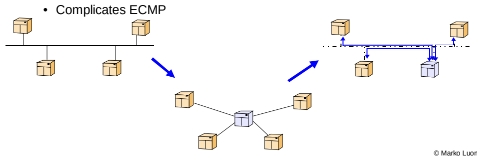

Broadcast support

- **IS-IS**
  - Only one intermediate system acts as a pseudo-node
  - L2 multicast addresses used
    - 0180.c200.0014: AllL1ISs
    - 0180.c200.0015: AllL2ISs
  - Everybody hears everybody
    - Fullmesh of adjacencies
- **OSPF**
  - Two routers act as pseudonodes
    - Designated Router
    - Backup Designated Router
  - L3 multicast addresses used
    - 224.0.0.5: AllSPFRouters
      - From DR to clients
    - 224.0.0.6: AllDRouters
      - From clients to DRs

IS-IS and OSPF

- **IS-IS**
  - Operates on top of L2-interface
    - NLPID coded
  - Link state operation
  - Dijkstra SPF
  - Two level hierarchy
    - Level-1
    - Level -2
  - Area border on links
  - Pseudonode emulation on LANs
    - Designated Router (DIS)
  - Designed for OSI CLNP
    - Support for IP added later on
- **OSPF**
  - Operates on top of IP
    - Protocol 89
  - Link state operation
  - Dijkstra SPF
  - Two level hierarchy
    - Backbone
    - Other area
  - Area border on nodes
    - Split nodes
  - Pseudonode emulation on LANs
    - Designated Router (DR)
    - Backup Designated Router (BDR)
  - Designed for IP

IS-IS vs OSPF

- **OSPF**
  - Packet format is variable
    - Positional fields with 32-bit alignment
  - Link State Database (LSDB) presented as LSAs
    - Content mixture of positional coding and TLV-coding
  - Unknown LSA types are discarded (not flooded)
- **IS-IS**
  - Positional header with TLV-coded messages
    - No particular alignment
  - Link State Database (LSDB) presented as LSPs
    - Content TLV coded
  - Unknown messages are ignored (flooded)
  - Actual information about the network is in TLVs
    - Easily extensible to new features

OSPF LSDB

- Built from the Link State Advertisements (LSA)
- Separate protocol elements and state machines handle the synchronization
  - **LS Update**
    - LSAs grouped into LSUpdates during flooding
    - LSUpdates are built individually at each hop based on accepted LSAs
  - **LS Acknowledge**
    - Each LSA has to be acknowledged otherwise they are retransmitted
  - **LS Request**
    - Missing or outdated information can be updated from the neighbor
  - **Database description**
    - LSDB structure as LS headers

OSPF State Machine

- **Four phases**
  - Neighbour aquiry
    - Init + 2-way: Hello protocol
  - Database initialisation
    - ExStart + Exchange: DD
  - Database synhronization
    - Exchange + Loading: LS request, LS update, LS acknowledge
  - Flooding
    - Full: LS update, LS acknowledge

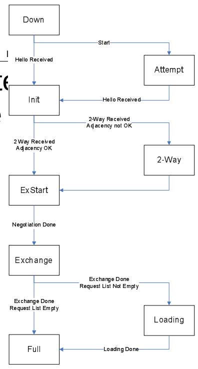

LSA types

- **OSPF contains separate LSA type for each different information element**
  - Tightly coded message structures
  - Optimized for 32bit processing
  - LSAs are valid only on certain points of network and areas
  - Unknown LSA types are rejected
    - Flooding is based only on accepted LSAs

Type	LSA Name	Description
1	Router	Link information
2	Network	DR adjacent RIDs
3	Network Summary	Prefixes from other area
4	ASBR Summary	Address of ASBR
5	AS External	External prefixes
6	Group membership	MOSPF groups
7	NSSA	Not So Stub Area Prefixes
8	External Attribute	BGP attributes
9	Opaque (link)	Traffic Engineering Hitless Restart Optional Capabilites
10	Opaque (area)	
11	Opaque (AS)	

IS-IS LSDB

- Built from Link State PDUs (LSP)
  - LSPs are basically TLV triplets
    - Type (Code) defines the format of value
    - Length defines the length of TLV
    - Informational element to be processed
  - LSP carries several TLV coded elements (default maximum 1492 bytes per LSP)
    - LSPs are flooded as is
      - Information is delivered intact from the originator to all of the routers in the area
      - If router does not support certain TLV it just ignores the TLV
        - » Gradual update of capabilities

IS-IS LSDB

- **Simple state machine**
  - Neighbour aquiry
    - Hello protocol
  - Database synhronization and Flooding
    - Complete Sequence Number PDUs
    - Partial Sequence number PDUs
    - Link State PDUs
- Synchronization is based on periodical descriptions of complete database

Some IP related TLVs

- Area Addresses (C=1)
  - List of all AIDs present at sender
- Neigbors (C=2)
  - List of senders neighbors (SysID)
  - Metrics to reach neighbors
- Protocols supported (C=129)
  - NLPIDs of supported protocols
    - IPv4 = 204
- IP Interface Addresses (C=132)
  - IP addresses of IS-IS interfaces on the sender
- IP Internal Reachability (C=128)
  - IP prefixes directly connected to sender and their metrics
- IP External Reachability (C=130)
  - IP prefixes external to IS-IS domain
- Extended IS Reachability (C=22)
  - Similar than neighbors but wide metrics
    - Narrow metrics = 6 bits
    - Wide metrics = 24 bits
- Extended IP Reachability (C=135)
  - Same as Internal and External but with wide metrics
  - Up/Down bit for controlled route leaking

Overload

- In large carrier networks restarting a router causes reload in
  - IGP
    - Fast convergence (few tens of seconds)
      - IS-IS faster than OSPF due to its simpler state machine
  - BGP
    - Slow convergence if default free (may take minutes to load all the prefixes)
- Overload bit in IS-IS makes router present in the network but not candidate for forwarding transit traffic
  - Timed clearing of overload bit is common feature in NOS
    - Causes new calculation of SPF when bit is cleared

OSPF vs IS-IS

- **OSPF**
  - Complex
  - Easier to manage
    - Default behavior is to do everything that adjacent is capable
  - Difficult to add new features
  - Good area control
  - Good for corporate environments
  - Dependens on 32bit IP addresses
    - Architectural dependency
      - OSPFv3 changes this
- **IS-IS**
  - Less complex
  - Difficult to manage advanced features
    - Requires routing policies
    - Default behavior is not to do anything but basic LS / SPF
  - Easier to add new features
  - Scalable on single area
  - Good for carrier environments
  - No architectural depencies to IP addresses