Area Abstraction

Better areas for IS-IS
draft-li-area-abstraction-00

Trend: Turning the router inside out

- Traditional multi-chassis router design:
  - Packet forwarding engines at the edges
  - Clos or Benes fabric in the middle
  - Single system abstraction
- What happens if you put the fabric on the outside?
  - Many smaller systems for doing packet forwarding
  - Clos or leaf-spine topology
  - Need abstraction of the result
- IGP abstraction mechanism: the area

Review: IS-IS areas

- Level 1 areas abstracted into Level 2
- But, for transit, Level 1 topology must also be Level 2 topology
- If most of the Level 1 topology is used for transit, this provides no benefit
- Result: Level 2 scale problem. Areas not particularly useful.

Requirements

- A stronger Level 1 abstraction
  - Level 1 area looks like a Level 2 node
  - Internal topology NOT advertised at all
  - All external connectivity represented

Proposed Architecture

- Represent L1 area as a single L2 pseudo-node
  - L1 area elects an Area Leader
  - Area Leader picks a pseudo-node ID
  - On external links, border routers generate IIH's using pseudo-node ID.
  - Internally, border routers create tunneled L2 adjacencies with Area Leader.

- Area Leader creates pseudo-node LSP listing external adjacencies from border router LSPs.
- Only the pseudo-node LSP flooded on external links
- Other L2 area-originated LSPs NOT flooded externally
- L2 transit LSPs flooded normally
- Result: L1 area looks like a single L2 pseudo-node with full external connectivity

Internals

- L1 area border routers must provide forwarding for L2
  - ABRs are full fledged L2 routers
  - Logical connectivity is via the Area Leader, but that's suboptimal for forwarding
  - Compute 'shortcuts' from entry ABR to exit ABR for direct tunneling across L1
  - Can use SR (or even GRE) for forwarding

Area Pseudo-node TLV

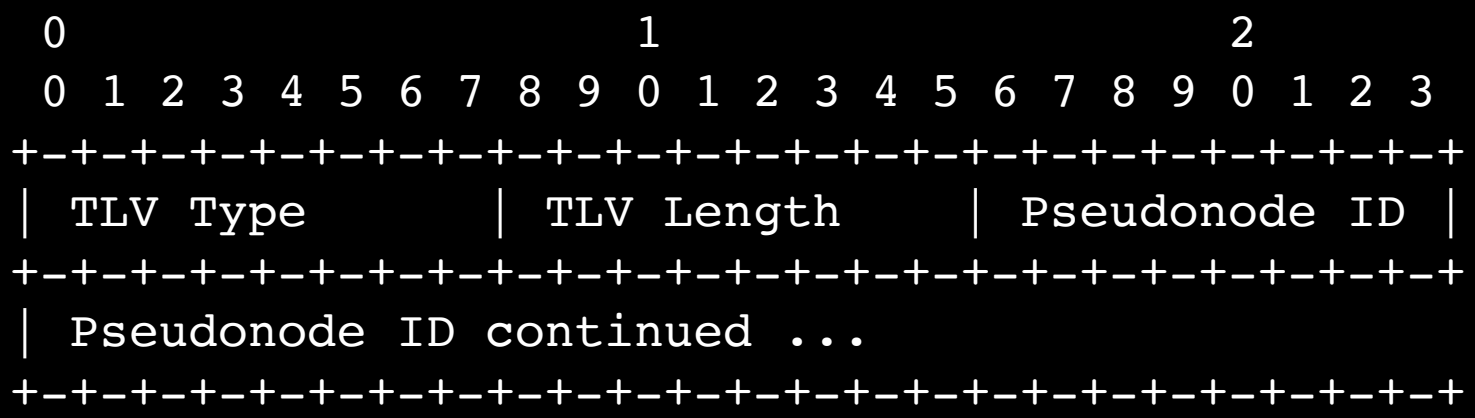

- TLV Type: TBD
- TLV Length: 2 + (system ID length + 1)
- Pseudonode ID: Pseudonode ID for the area