Hierarchical IS-IS

More levels, please

- At scale, what happens to IS-IS network architecture?
  - A pod in a data center becomes an L1 area.
  - The entire data center becomes an L2 area.
  - Now we need Level 3 for the WAN. How?

Levels 3 thru 8

- IS-IS encoding already reserves bits for more levels
- Circuit type (ISO 10589, section 9.5):
  - 1 - Level 1
  - 2 - Level 2
  - 3 - Level 1 & 2
  - 6 reserved bits

New bits

- Bit values:
  - 4 - Level 3
  - 8 - Level 4
  - 16 - Level 5
  - 32 - Level 6
  - 64 - Level 7
  - 128 - Level 8
- Set bits MUST be contiguous

New LAN Hello PDU

- Existing:
  - LAN L1 IIH PDU
  - LAN L2 IIH PDU
- Add:
  - LAN HELLO PDU (same format, just separate for backward compatibility)
  - Covers L3 - L8
  - If only some levels are supported, this applies only to the common levels.

New LSPs

- Existing:
  - L1 LSP
  - L2 LSP
- Add
  - L3 LSP, L4 LSP, …, L8 LSP
  - Inherits everything from L2, translated to the target level

New CSNP, PSNP

- Existing: L1 CSNP, L2 CSNP, L1 PSNP, L2 PSNP
- Add:
  - L3 CSNP, L4 CSNP, …, L8 CSNP
  - L3 PSNP, L4 PSNP, …, L8 PSNP

Summary

- Enable level 3 through 8 for IS-IS.
- Each level multiplies overall scale, so this gives extreme scalability (10^{24}) for little effort.
- Request WG adoption.

Area Abstraction

An IS-IS network

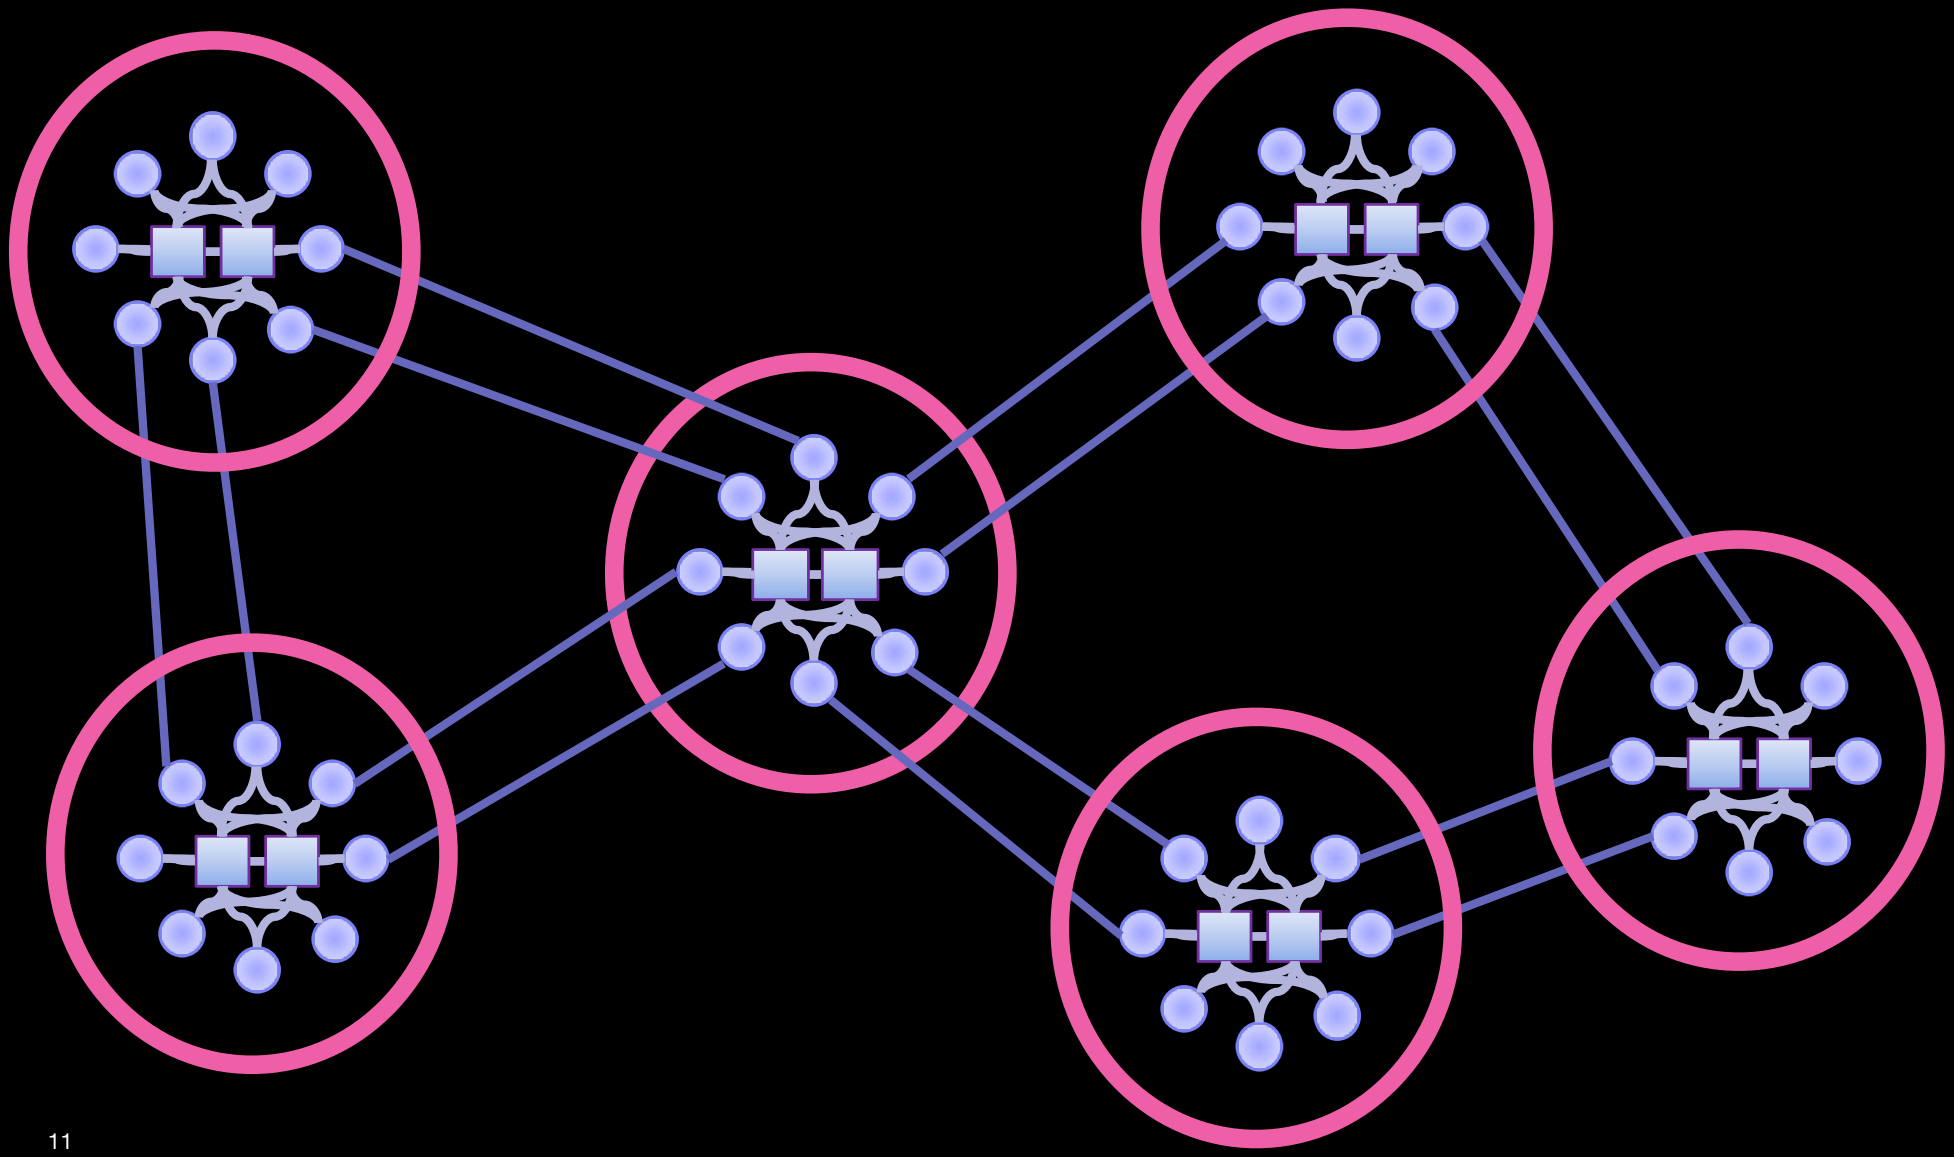

An IS-IS network

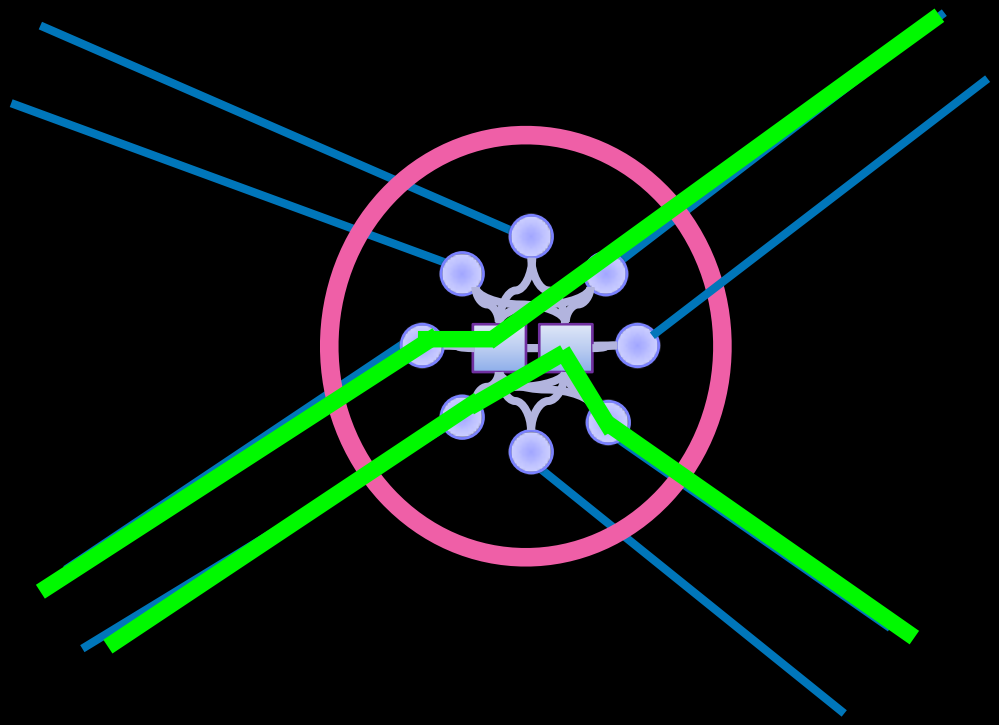

An IS-IS network

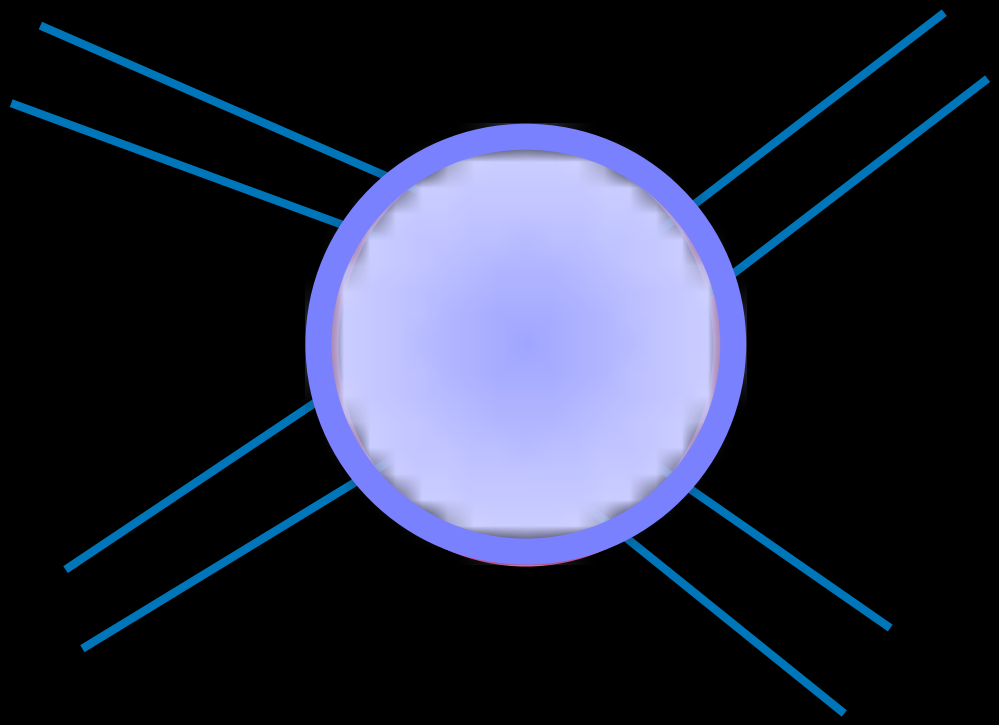

Requirements

1. Send out level 2 hello packets (IIHs). Same system ID on all level 2 interfaces. We'll need a special system ID for this.

2. Participate in level 2 flooding.

3. Generate a level 2 link state PDU with all of the level 2 adjacencies. Nothing else is advertised into level 2.

4. Forward level 2 transit traffic.

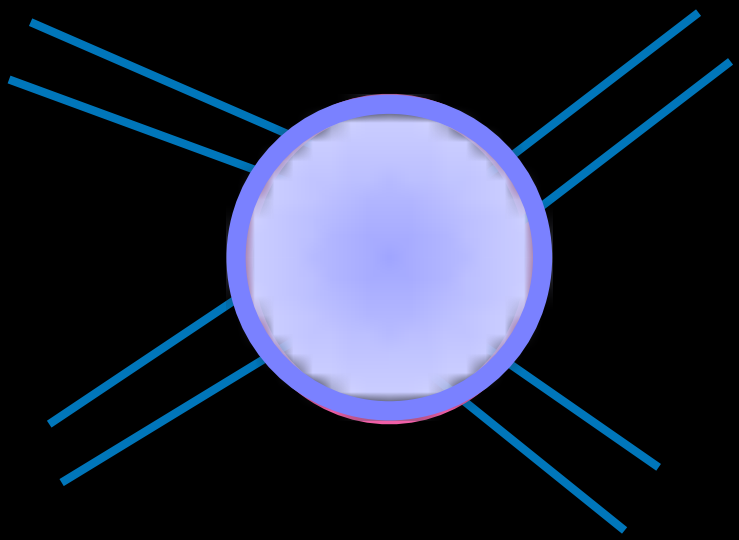

Area Proxy System ID TLV

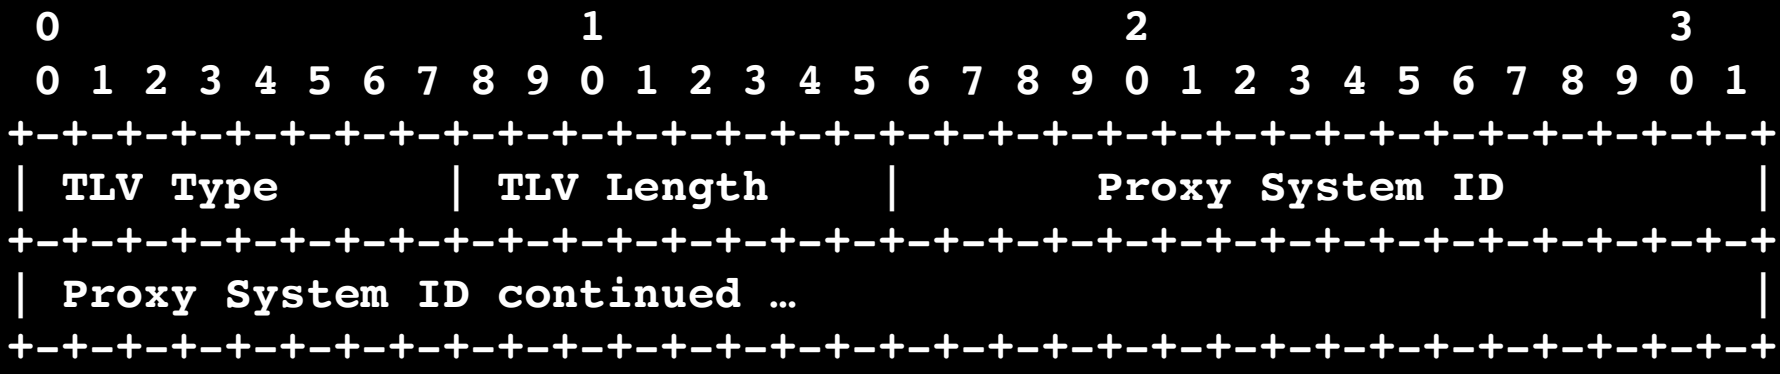

- Area Leader advertises a System ID to be used as a proxy for entire area.
- Proxy ID used by edge level 2 systems to generate level 2 hellos, using the Proxy ID as the source.

L2 Flooding

- Need to provide L2 flooding across the L1 area.
- Tunnel from L1L2 edge routers back to Area Leader. Run L2 on top of the tunnel. This creates the L2 flooding topology.

L2 Proxy LSP

- Area Leader generates L2 Proxy LSP.
- Examines L2 LSDB for external neighbors. These were advertised as adjacencies by the L1L2 edge routers.
- External adjacencies are then advertised as neighbors to the proxy (e.g. A, B, C, …, H).
- All other L2 LSPs from the area itself are NOT flooded outside of the area.
- All other normal area ‘leaking’ happens into the proxy LSP.

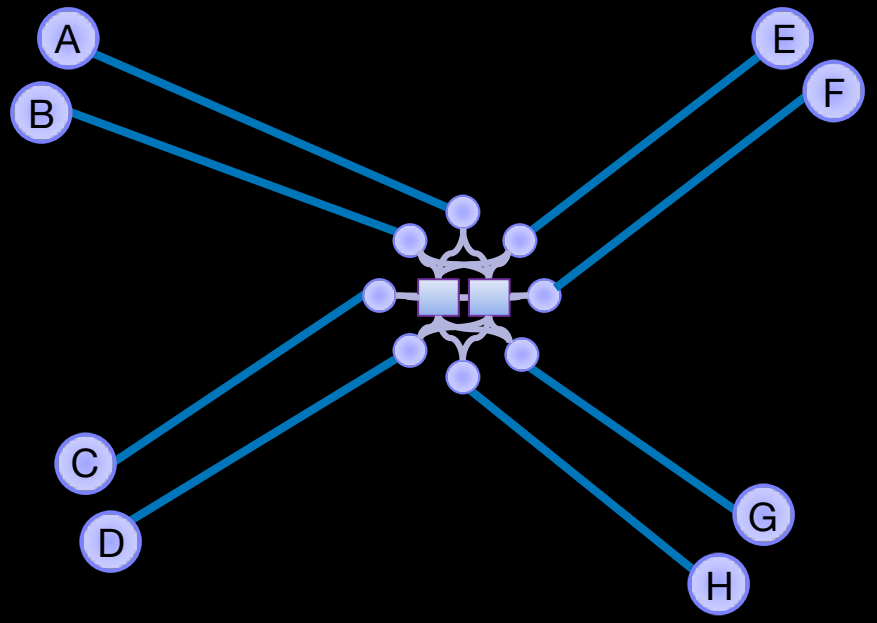

L2 Forwarding

- Traffic arrives at L1L2 edge router.
- Transit traffic would match an L2 prefix.
- As part of the L2 SPF computation, the entry edge router uses the area exit router as a next hop and tunnels (SR, RSVP-TE, GRE) the traffic to the exit.
- L2 SPF can only use intra-area metrics as an inter-area tiebreaker.

Summary

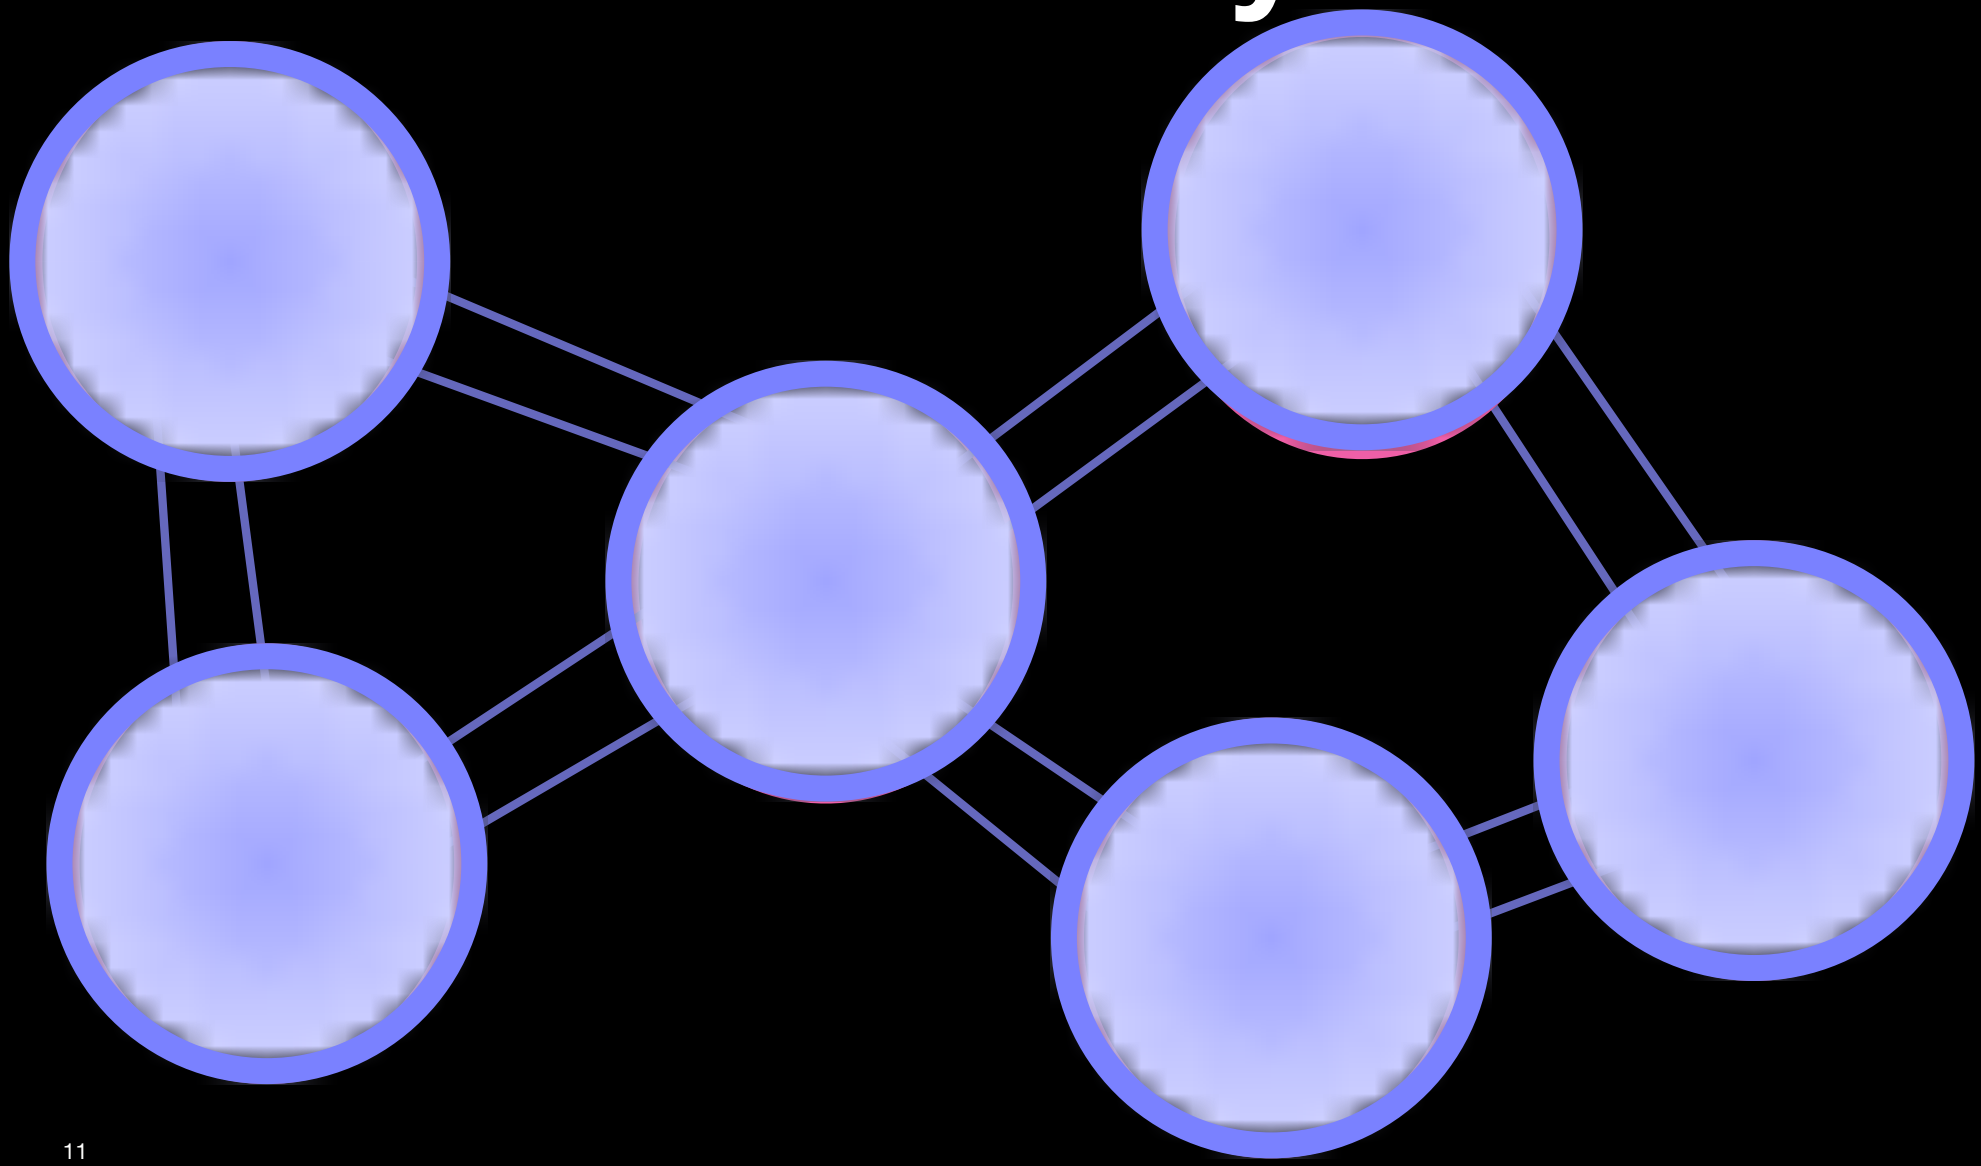

Relation to Topology-Transparent Zones

- Prior (independent) work: draft-chen-isis-ttz-05
- Similar goals

Primary differences

- Zone boundaries don’t coincide with area boundaries.
- All zone edge routers add an LSP to the L2 LSDB.
  - Each LSP has links to all other zone edge routers.
- Smooth migration mechanism for zones

- Combine drafts?