Traffic Engineering for 802.1 unicast/multicast networks

Nigel Bragg, Paul Unbehagen

Background

- Both AVB and Interworking have seen the need for
  - Bandwidth reservation
  - Engineered paths
- SPB not only learns topology and shares location of services, but shares link specific information
- SPB currently only controls default paths
- ISIS has existing mechanisms to distribute TE information
- Priority Flow Control and Congestion Notification are complementary, but orthogonal to path placement

“Concern”

We are unsure whether we have fully understood the requirements of the environment AVB is targeting :

- so we’ll replay our understanding of the headlines :
  - explicitly engineered **p2p** and **p2mp** trees,
  - with (restorable) backup
  - and a protection-switched (OAM-driven) mechanism to move between trees on an all-or-nothing basis (i.e. **no** requirement for protection-switched local repair, which will be handled by restoration)
  - lossless **manual** switchover between trees.
  - … what else ??
- How similar are these to Interworking?

Question

- What needs to be added to SPB/IS-IS?

Answer

- TE with SPB using underlying ISIS TLVs
- Protection switching mechanism to control flow to and from engineered paths
  - using the principle inherited from PBB-TE
- AVB specific information?
- Flow awareness?

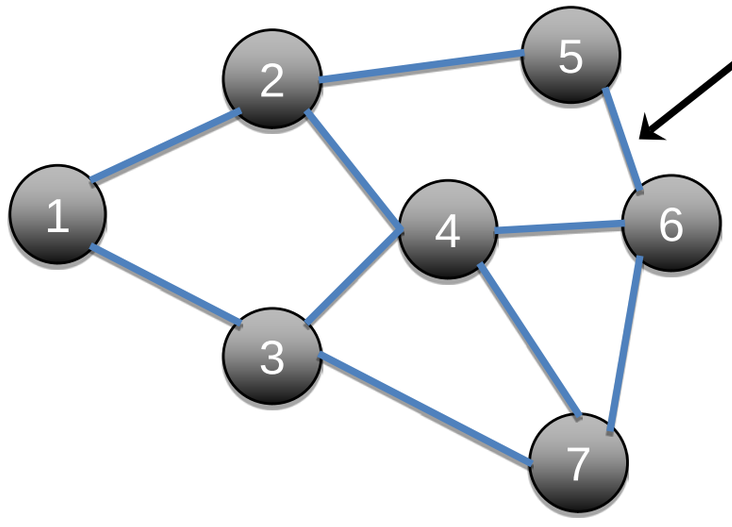

- Sub-TLV 3: Administrative Group (color, resource class)
- Sub-TLV 9: Maximum Link Bandwidth
- Sub-TLV 10: Maximum Reservable Link Bandwidth
- Sub-TLV 11: Unreserved Bandwidth
- Sub-TLV 29: SPB Link Metric
- *Potential new TLV?*

SPB Protocol Modes

- SPB has two data plane modes; SPBv and SPBm
- Both use IS-IS as a base control plane to manage their own data planes
- ISIS learns and shares everything through its Link State DataBase (LSDB)

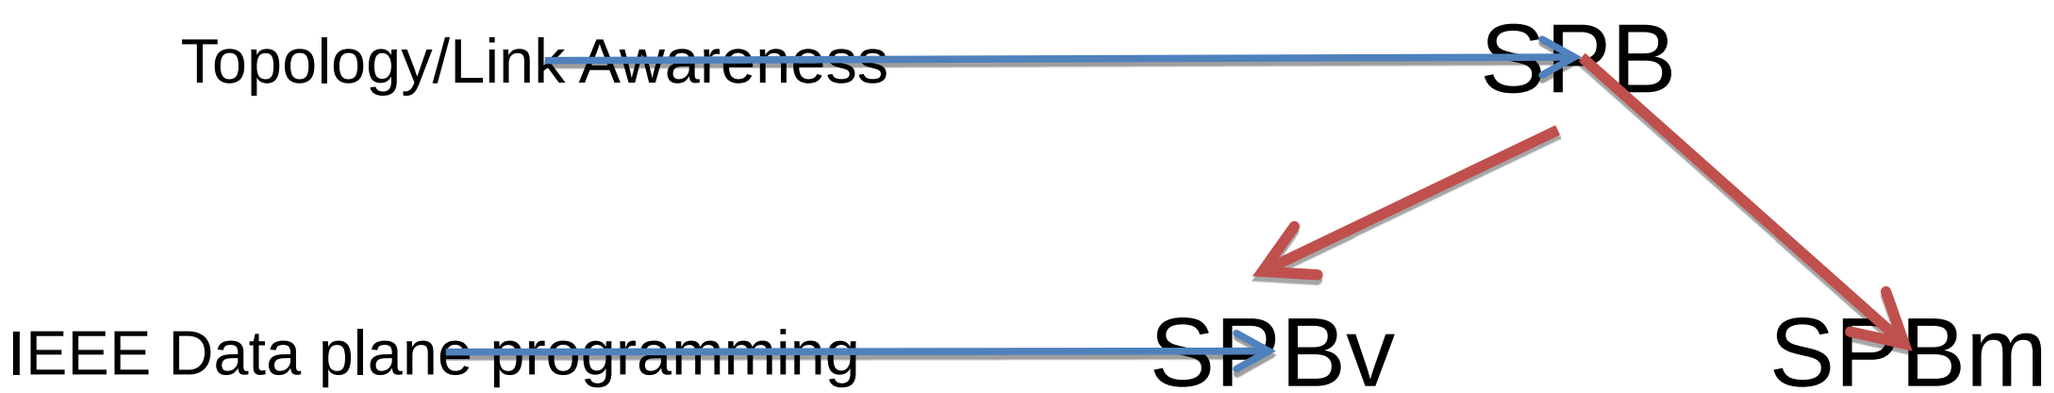

Existing TE in IS-IS

- SPB adjacency information is at the same level as the IS-IS TLV's for TE
  - so what does this mean ?

We can just reuse the IS-IS mechanisms to flood link occupancy information

How do we install path state ?

We assume that engineered paths are computed by a management function. Then :

- Current IS-IS-TE uses a signalling protocol to seize resources and install path state, or
- At modest scale, we could just flood the path way-points in IS-IS using a new SPB TLV

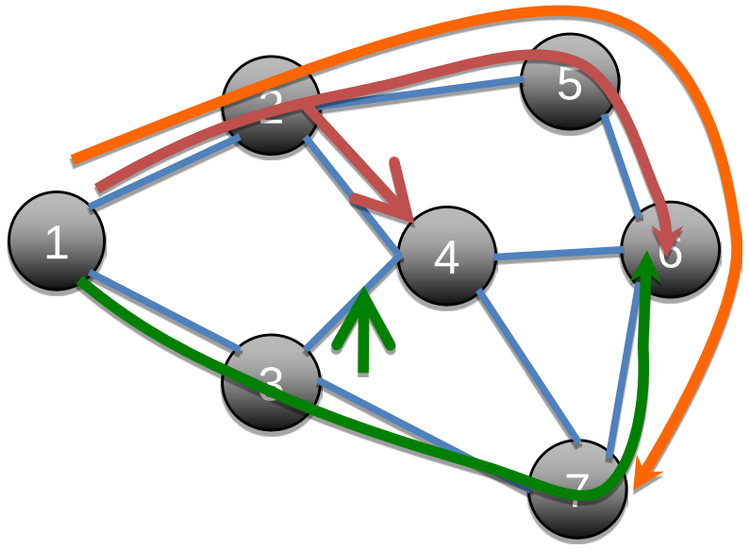

Path 1 (1->4, 6) Default SPT

Path 2 (1->4, 6) Explicit routed (diverse)

Path 3 (1-7) Explicit Routed (non-SPF)

But, what about simple bandwidth reservation with no alt path computation?

How do we do the protection switch ?

- Actually, the switch-over is easy :
  - Swap VLAN Id at head-end to invoke a second forwarding plane.
  - all receivers accept off either VID
- How do we trigger the head-end to switch ?
  - data plane OAM ? – one CCM session per recipient
  - rely on local detection and IS-IS “link down”

What’s already available in IS-IS

- ISIS has per link TE awareness built in
  - SPB adjacencies use the same base TLV
  - Implementation chooses what to add in the adj messages
- Should we reuse what’s there or define new TLV?

Current ISIS-TE TLV

- The Extended IS Reachability TLV
  - Sub-TLV 3: Administrative Group (color, resource class)
  - Sub-TLV 6: IPv4 Interface Address (n/a)
  - Sub-TLV 8: IPv4 Neighbor Address (n/a)
  - Sub-TLV 9: Maximum Link Bandwidth
  - Sub-TLV 10: Maximum Reservable Link Bandwidth
  - Sub-TLV 11: Unreserved Bandwidth
  - Sub-TLV 18: Traffic Engineering Default Metric
  - *Sub-TLV 29: SPB Link metric (24-bit unsigned number)*
- The Extended IP Reachability TLV
  - Most likely not applicable