ERICSSON

P802.1Qca D0.4 Use Cases and Operation

János Farkas

November 12, 2013

Outline

- Presentation objectives
- Disclaimer
- Highlights
- Use cases
- VIDs, MACs, and paths
- Examples
- Summary
- Backup

Presentation Objectives

- Explore the operation of explicit path establishment as described in P802.1Qca D0.4 through examples

Disclaimer

- The operation presented here is not the final standard!
- This is only as the view is after the first Task Group ballot
- There are open items and items under debate

- Some use cases investigated here, but the application possibilities are broader than that.

Highlights

- 802.1Qca is an extension to IS-IS
- It is control plane
- Main goal: path establishment
- Path establishment does not require hardware changes!
- Paths are symmetric
- Forwarding can be made unidirectional by
  - VID or
  - VID + MAC
- The algorithm the PCE uses for path computation is not specified

Use Case 1 Centralized PCE

- Main target: Data Center network
  - Main goal: traffic steering
  - This does not preclude the different case when a time-sensitive network has a single central PCE, which is also supported
- PCE brings up adjacency with an SPT Bridge
  - PCE resides in an end station directly connected to an SPT Bridge
  - PCE most likely resides in a VM in a DC
  - PCE has a link state database LSDB
- Single central PCE
  - A single control entity for the explicit paths
  - → No conflicts are possible

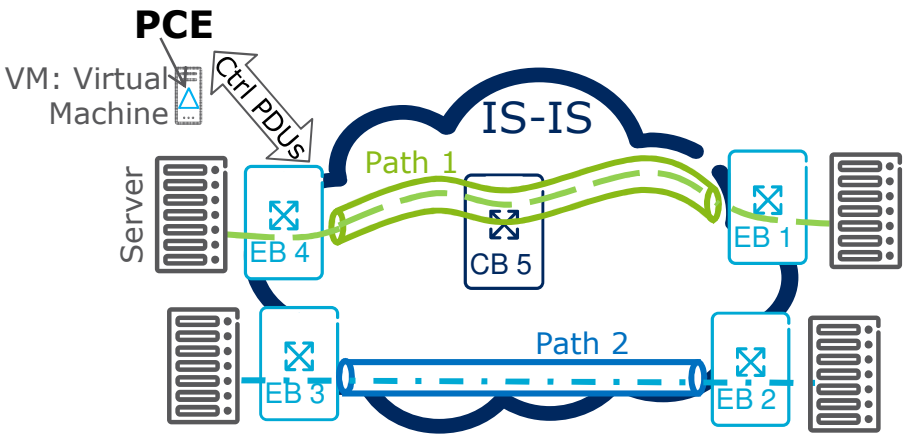

Use Case 2
Multiple PCEs

- Main target: time-sensitive network
  - Main goals: paths for streams, and paths for seamless redundancy
  - There may be other applications relying on multiple PCEs
- PCE is resident of an SPT Bridge
  - For instance, each Edge Bridge (EB) may has its own PCE
  - Each PCE has its link state database LSDB
  - PCE can get MSRPDUs as well

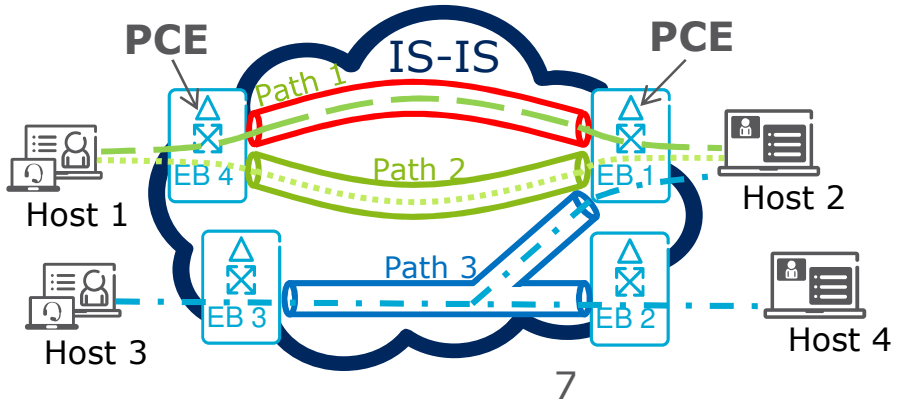

TSN Use Case
SRP (802.1Qcc) is also there

- One or more paths can be requested from a PCE (by Listener/Talker)
- PCE determines the explicit path(s) if needed
- Path(s) get installed and VIDs, MACs configured
- SRP performs reservation on either the shortest path or on the explicit path

- A lot of details still have to be nailed down!

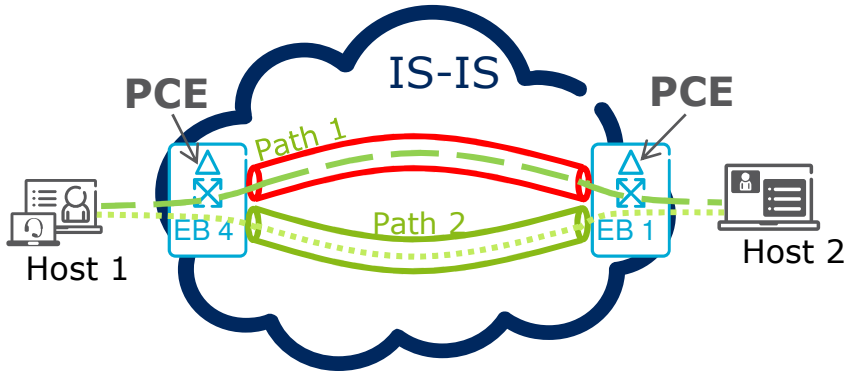

VIDs, MACs, and Paths

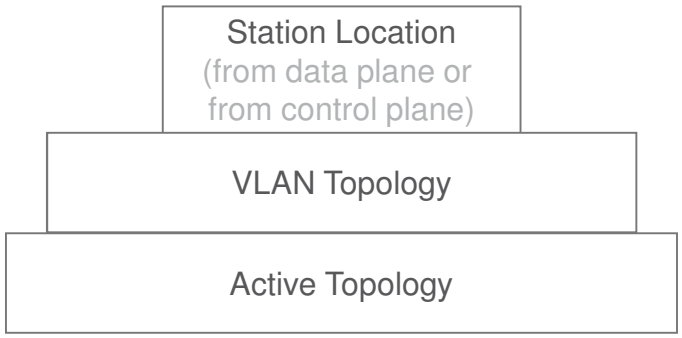

based on Figure 7-1

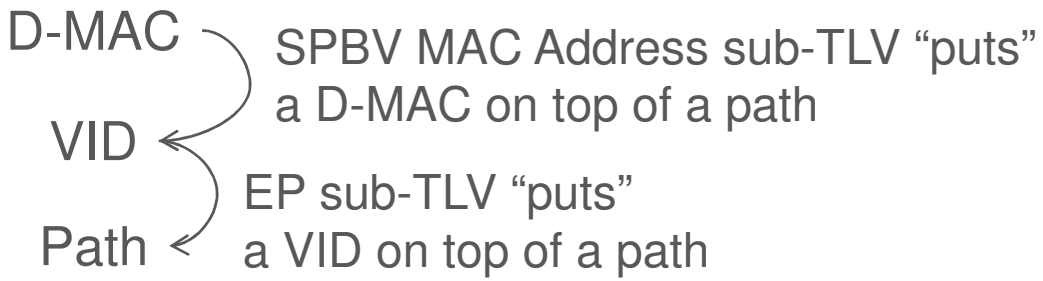

- VID + D-MAC together determine the egress port, i.e. the path that a frame takes
- **VIDs can be reused ➔ a VID does not identify an EP**
  - see details on VID reuse in http://www.ieee802.org/1/files/public/docs2013/ca-nbragg-routing-structures-1113-v02.pdf, and Example 2 below

Example 1
Two Disjoint Paths

- Goal: explore how two disjoint paths: path1 and path2 can be set-up by ISIS-PCR
- This is Use Case 2; each bridge has its own PCE in this example
  - The PCE of E is responsible for path1 and path2
- DA1 the is the only DA in this example (control plane DA)
- Out of scope for now:
  - How TA of T1 got to L1
  - How L1 requests the paths from E
  - How reservation is done
  - How frame duplication/elimination is done

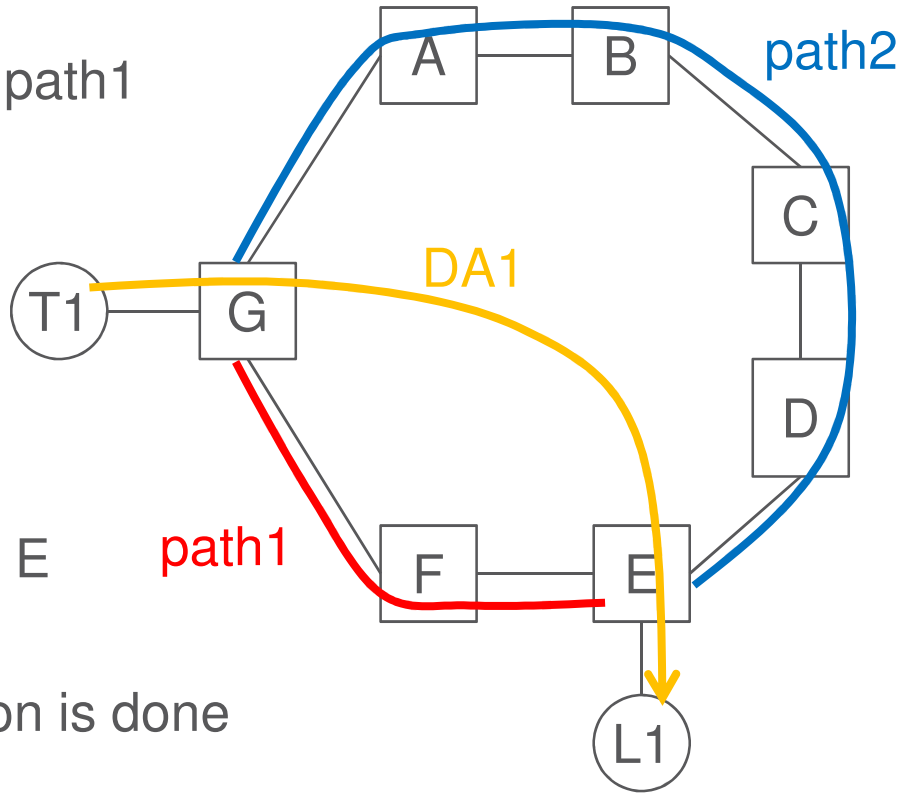

Example 1
VID – MSTI

- Two different VIDs are used
  - Two egress ports are associated with DA1 in G's FDB
  - Path1 → VID1
  - Path2 → VID2
- VID → MSTID configuration
  - We only want to have DA1 along these paths
    - No data plane MAC learning
    - MAC only populated by IS-IS
  - VID1 → SPBM MSTID (0xFFC)
  - VID2 → SPBM MSTID(0xFFC)
  - **Therefore, neither VID1 or VID2 are SPVIDs!**

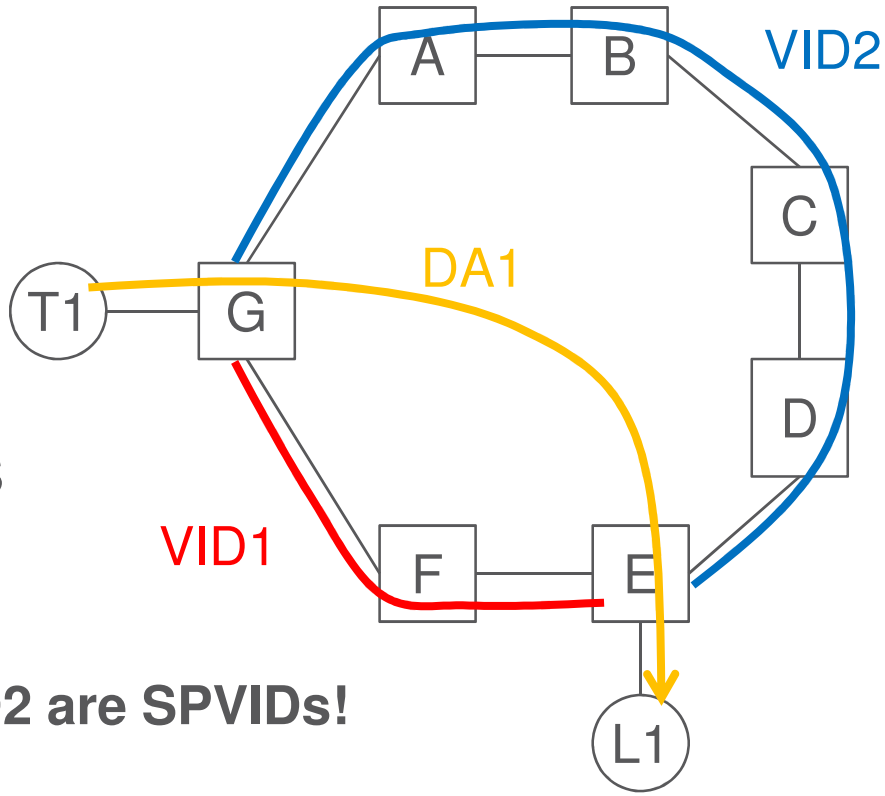

Example 1
ECT Algorithm

- Explicit path operation is selected by allocating the VLAN's Base VID to an EP ECT Algorithm (Table 45-1)
- SPB Base VLAN-Identifiers sub-TLV
  - Static paths

Type = 6
Length = 12
ECT-Algorithm = **SE ECT Algorithm (00-80-C2-20)**
Base VID = **VID1**
U =1
M = 1
ECT-Algorithm = **SE ECT Algorithm (00-80-C2-20)**
Base VID = **VID2**
U =1
M = 1

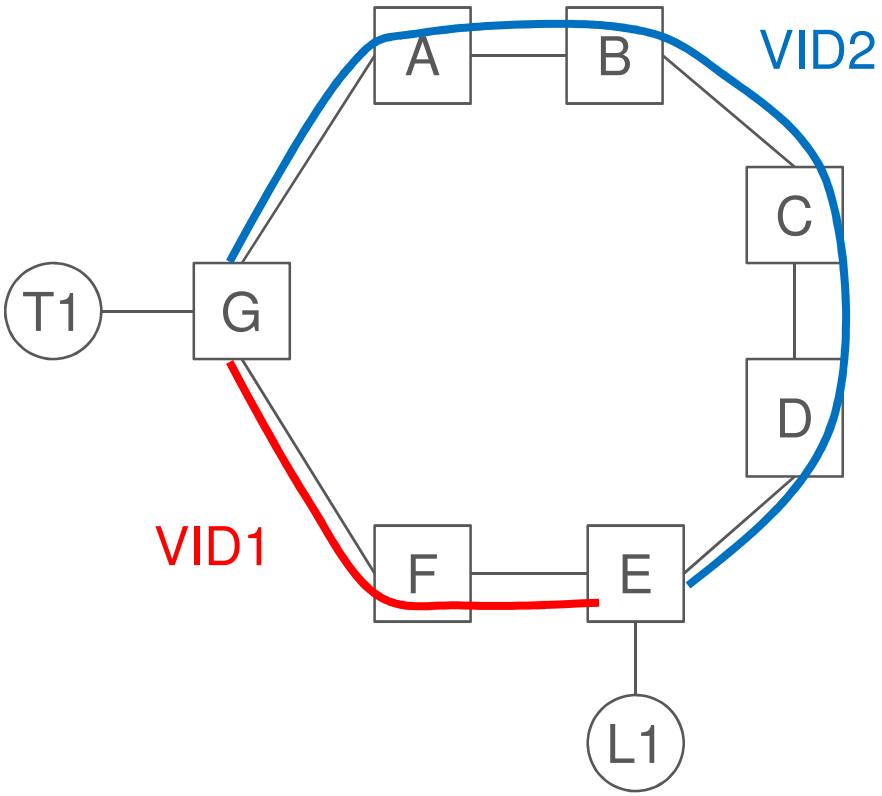

Example 1
MAC DA

- DA1 the is the only DA in this example
- MAC → VID association is only used in this example
- ➔ SPBV MAC Address sub-TLV is used
  - **Neither VID1 or VID2 are SPVIDs!** see VID → MSTID allocations on page 11
- DA1: G Transmits (T flag), E Receives (R flag)

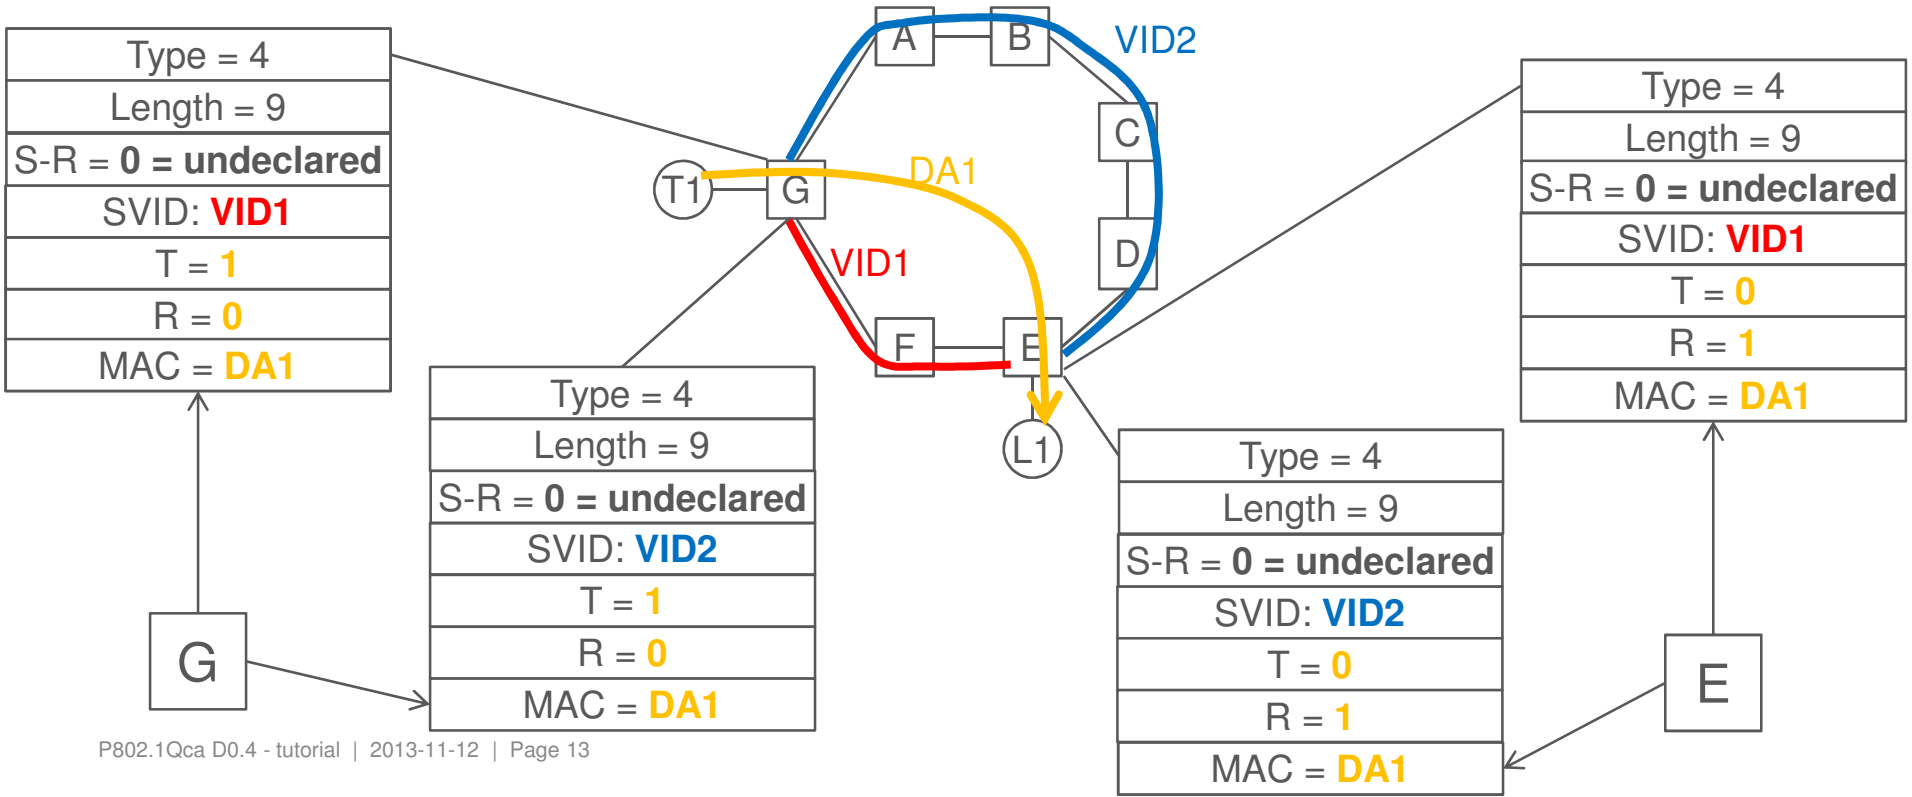

Example 1 Path1

- Fully specified path, i.e. no loose hop
- Path is specified by the EP Hop sub-TLVs
- VID is on top of the path

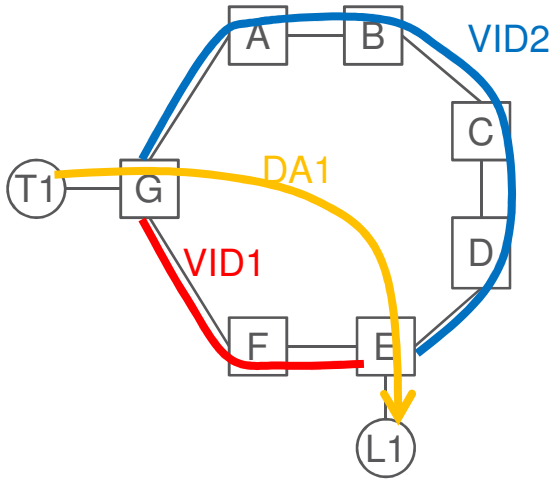

Explicit Path sub-TLV of path2 in an LSP of E:

Type = TBD3
Length = 31
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID1**
EP Hop sub-TLV for G
EP Hop sub-TLV for F
EP Hop sub-TLV for E

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **E**

see more detailed picture on page 33

Example 1 Path2

- Fully specified path, i.e. no loose hop

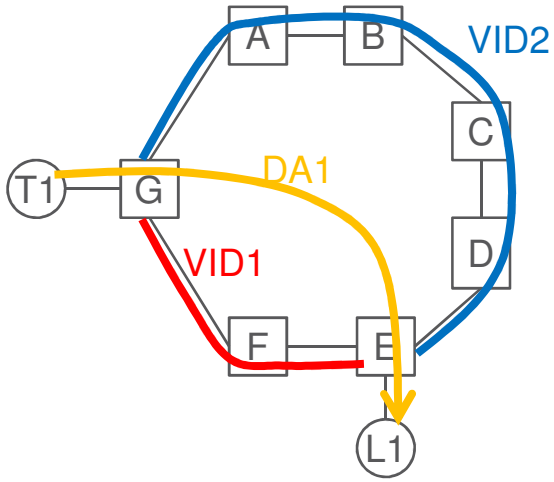

Explicit Path sub-TLV of path2 in an LSP of E:

Type = TBD3
Length = 58
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID2**
EP Hop sub-TLV for G
EP Hop sub-TLV for A
EP Hop sub-TLV for B
EP Hop sub-TLV for C
EP Hop sub-TLV for D
EP Hop sub-TLV for E

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **E**

see more detailed picture on page 34

Example 1
FDB Entries

- T/R flags are specified for DA1, see page 13

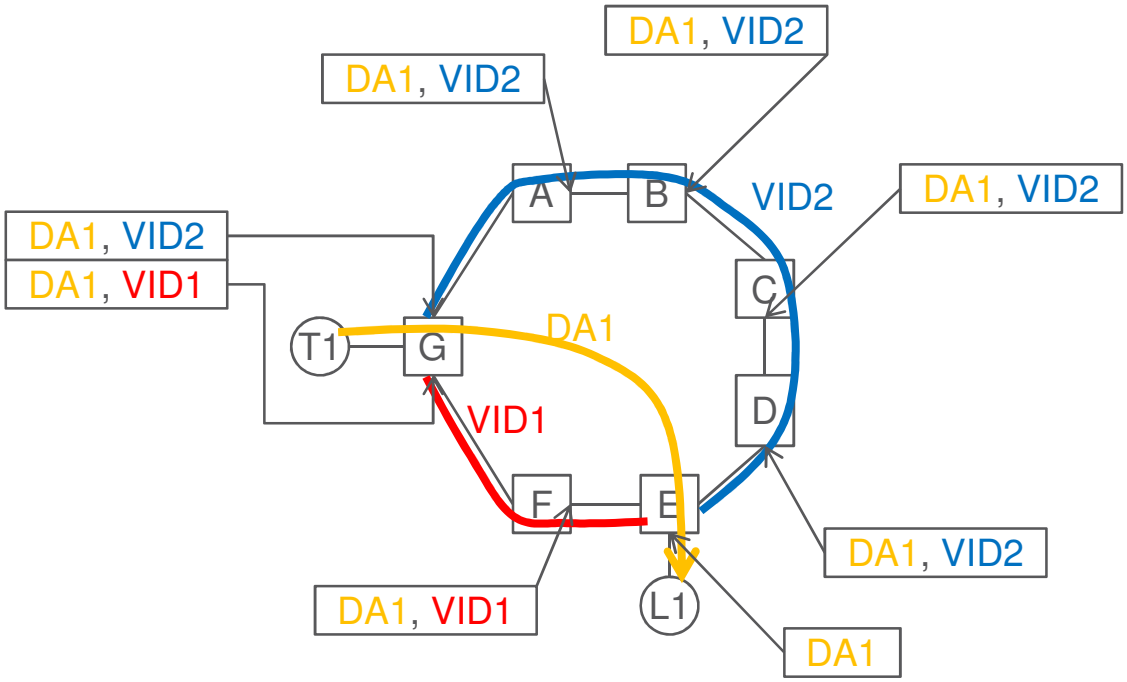

Example 2
Two More Disjoint Paths

- DA2 is the new DA (control plane DA)
- VID1 and VID2 are aimed to be reused
  - It can be done because of the traffic pattern, i.e. we only have DA1 and DA2 FDB entries for VID1 and VID2
  - Path3 → VID1
  - Path4 → VID2
- Two egress ports are associated with DA2 in B's FDB

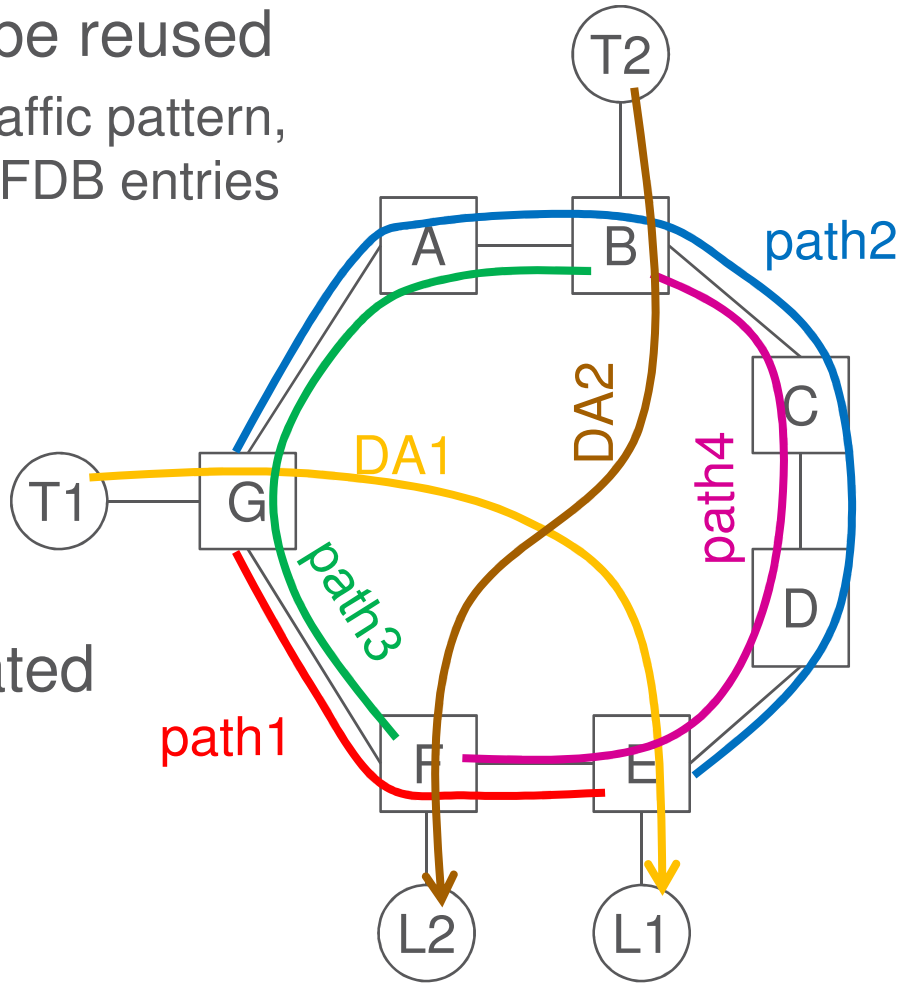

Example 2
VID – MSTI

- VID – MSTI allocation is the same as in case of Example 1
- VID → MSTID configuration
  - We only want to have DA1 along these paths
    - No data plane MAC learning
    - MAC only populated by IS-IS
  - VID1 → SPBM MSTID (0xFFC)
  - VID2 → SPBM MSTID(0xFFC)

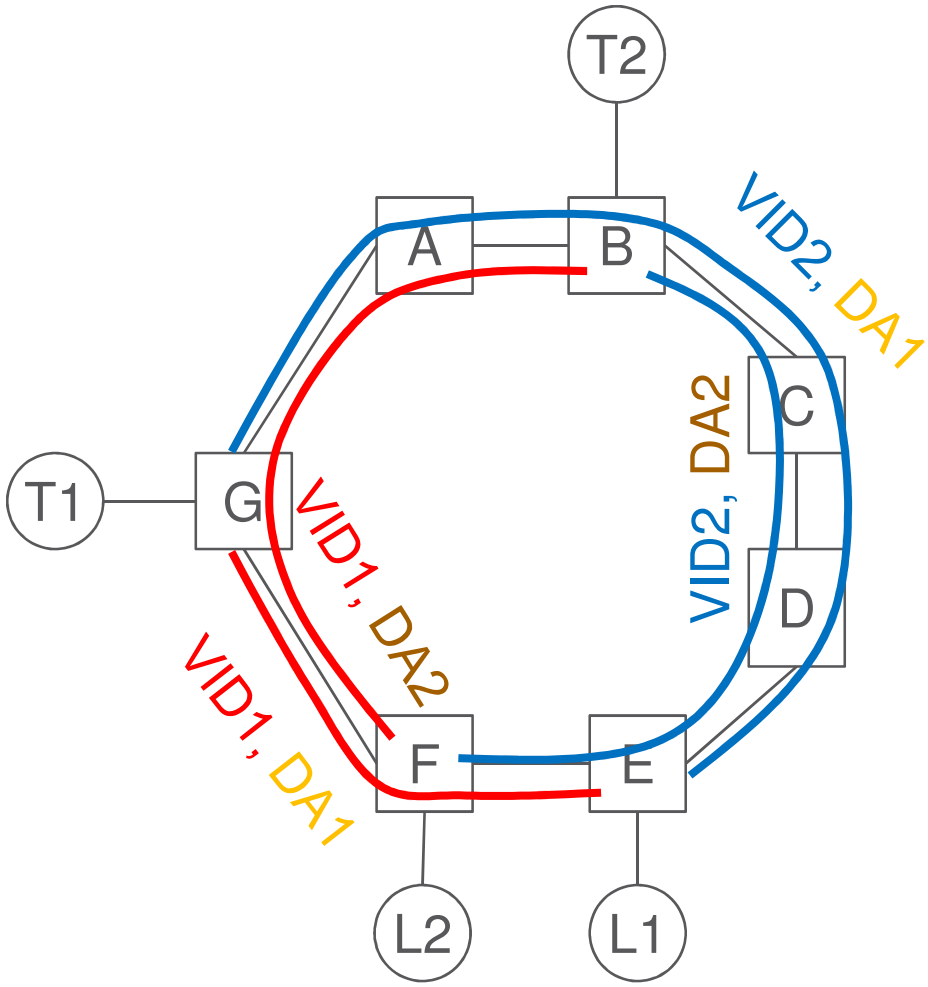

Example 2
ECT Algorithm

- Base VID → ECT Algorithm is the same as in case of Example 1
- SPB Base VLAN-Identifiers sub-TLV
  - Static paths

Type = 6
Length = 12
ECT-Algorithm = **SE ECT Algorithm (00-80-C2-20)**
Base VID = **VID1**
U =1
M = 1
ECT-Algorithm = **SE ECT Algorithm (00-80-C2-20)**
Base VID = **VID2**
U =1
M = 1

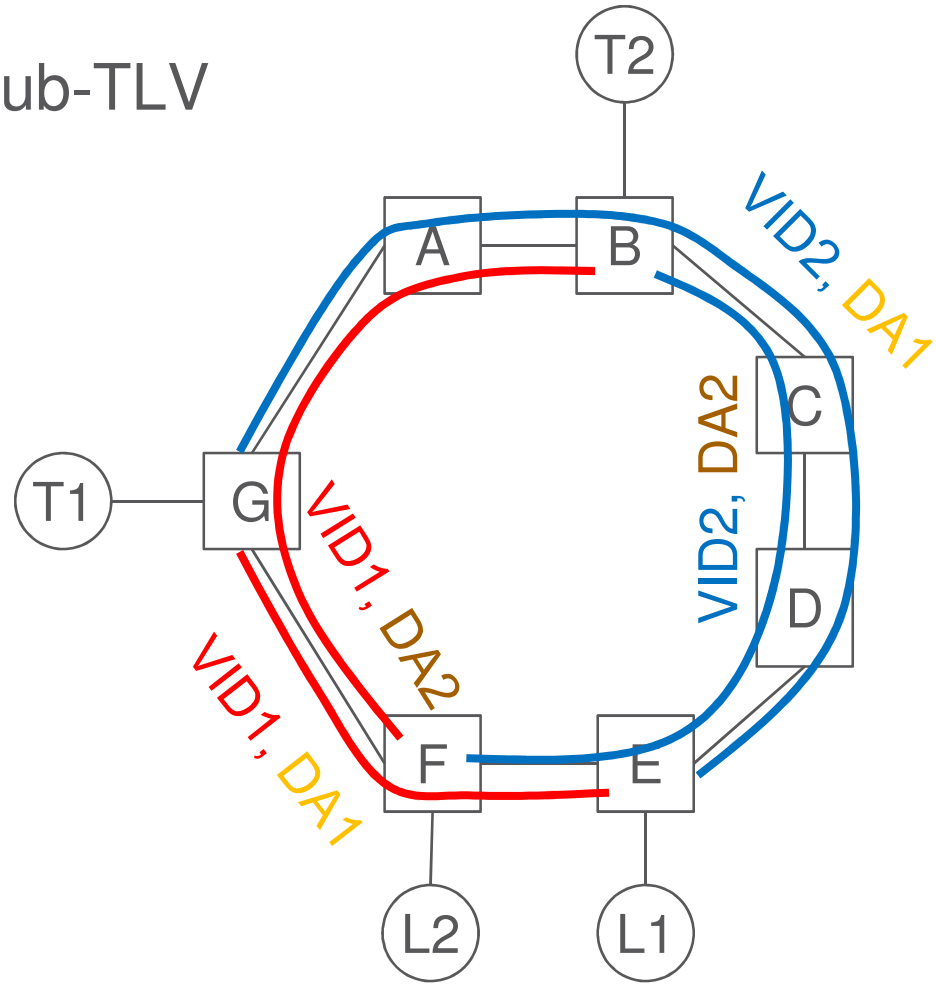

Example 2
MAC DA

- SPBV MAC Address sub-TLV is used
- DA2: G Transmits (T flag), E Receives (R flag)

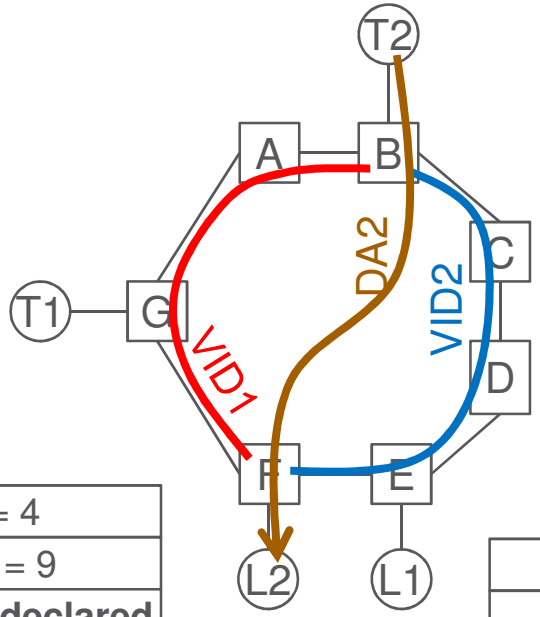

Type = 4
Length = 9
S-R = **0 = undeclared**
SVID: **VID1**
T = **1**
R = **0**
MAC = **DA2**

B

Type = 4
Length = 9
S-R = **0 = undeclared**
SVID: **VID2**
T = **1**
R = **0**
MAC = **DA2**

Type = 4
Length = 9
S-R = **0 = undeclared**
SVID: **VID2**
T = **0**
R = **1**
MAC = **DA2**

F

Type = 4
Length = 9
S-R = **0 = undeclared**
SVID: **VID1**
T = **0**
R = **1**
MAC = **DA2**

Example 2 Path3

- Fully specified path, i.e. no loose hop

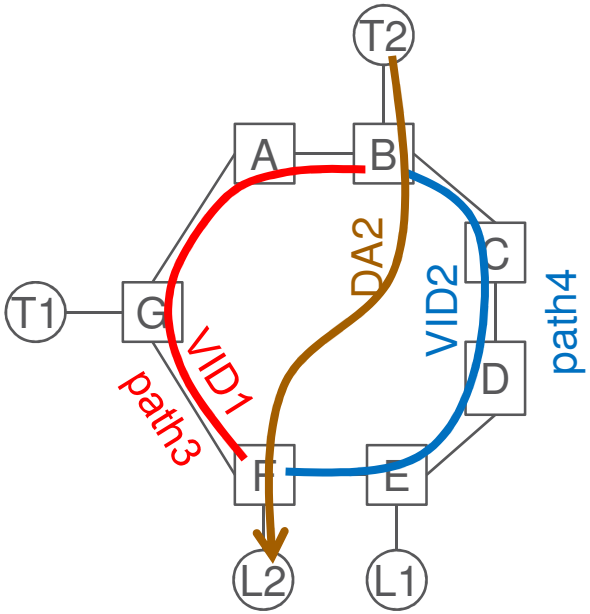

Explicit Path sub-TLV of path3 in an LSP of F:

Type = TBD3
Length = 40
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID1**
EP Hop sub-TLV for B
EP Hop sub-TLV for A
EP Hop sub-TLV for G
EP Hop sub-TLV for F

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **G**

see more detailed picture on page 35

Example 2 Path4

- Fully specified path, i.e. no loose hop

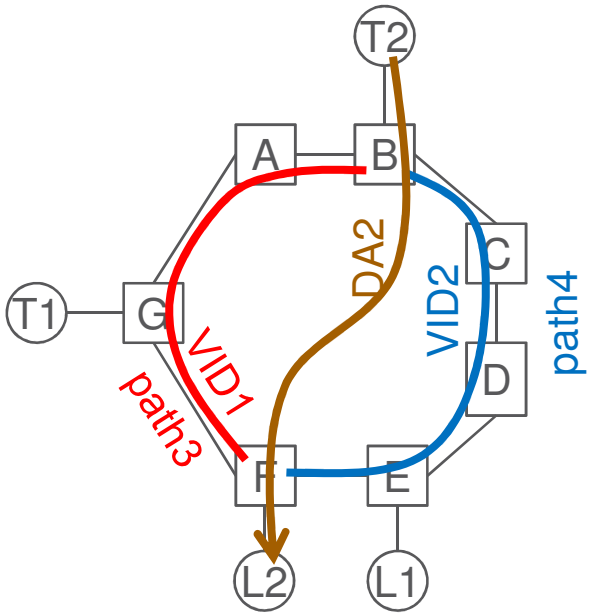

Explicit Path sub-TLV of path4 in an LSP of F:

Type = TBD3
Length = 49
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID2**
EP Hop sub-TLV for B
EP Hop sub-TLV for C
EP Hop sub-TLV for D
EP Hop sub-TLV for E
EP Hop sub-TLV for F

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **D**

see more detailed picture on page 36

Example 2
FDB Entries

- T/R flags are specified for DA2, see page 20

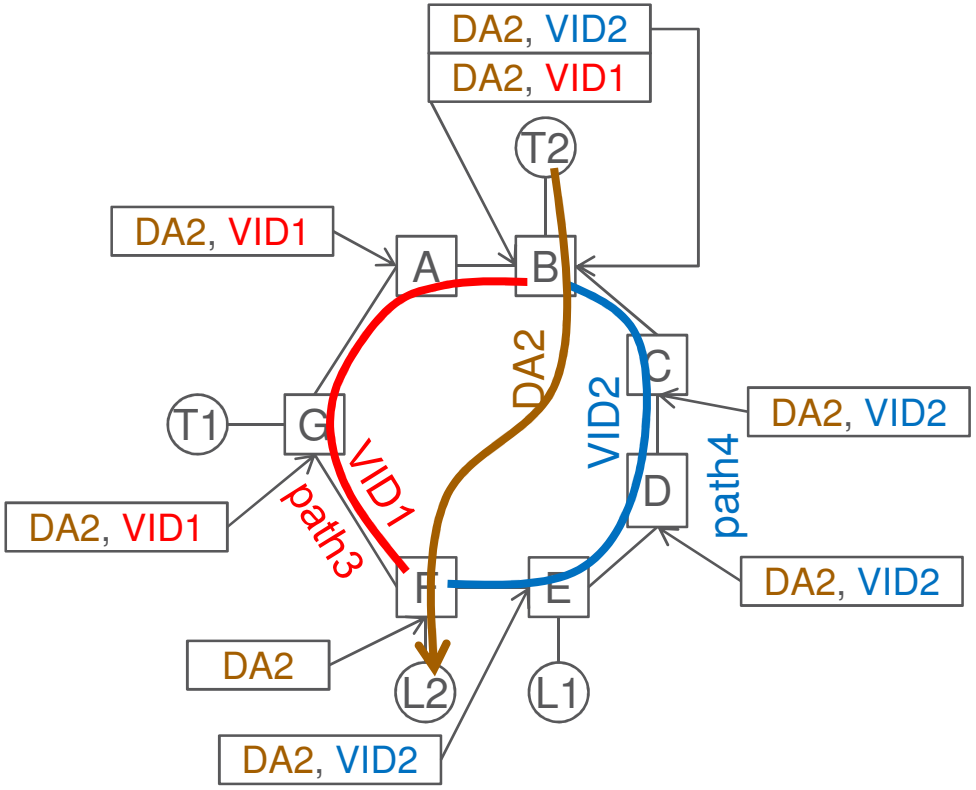

Example 3
Reservation on path2

- Reservation is performed by SRP
- L1 sends Listener Ready (LR)
- LR is propagated along path2
- Reservations are performed along path2

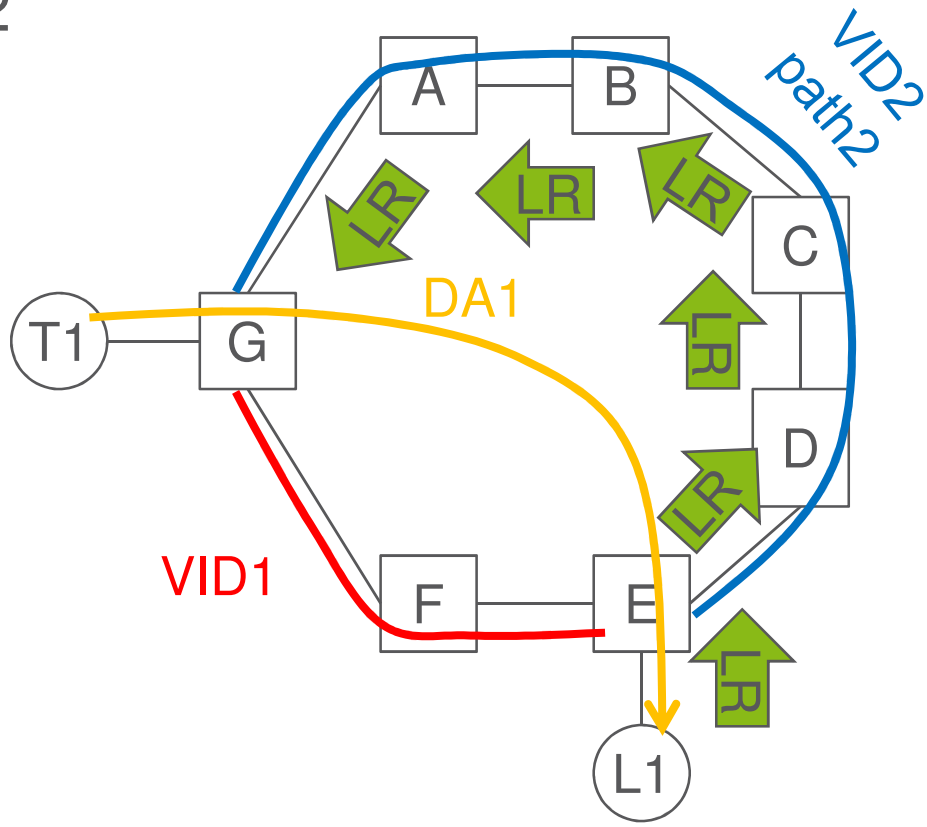

Example 4
Traffic Steering in a PBBN

- Centralized, single PCE
- PCE is connected to B
- Shortest path between A and E is overloaded
- Traffic steering path needed for I-SID5

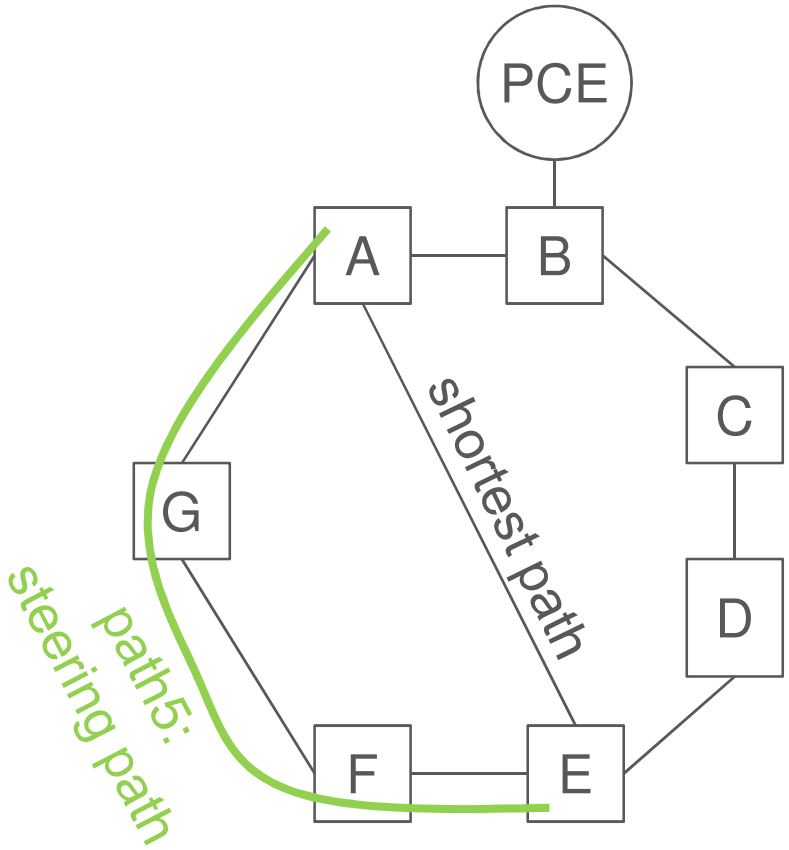

Example 4
VID – MSTI

- Path5 → B-VID5

- VID → MSTID configuration
  - VID5 → SPBM MSTID (0xFFC)
    - No data plane MAC learning
    - MAC only populated by IS-IS

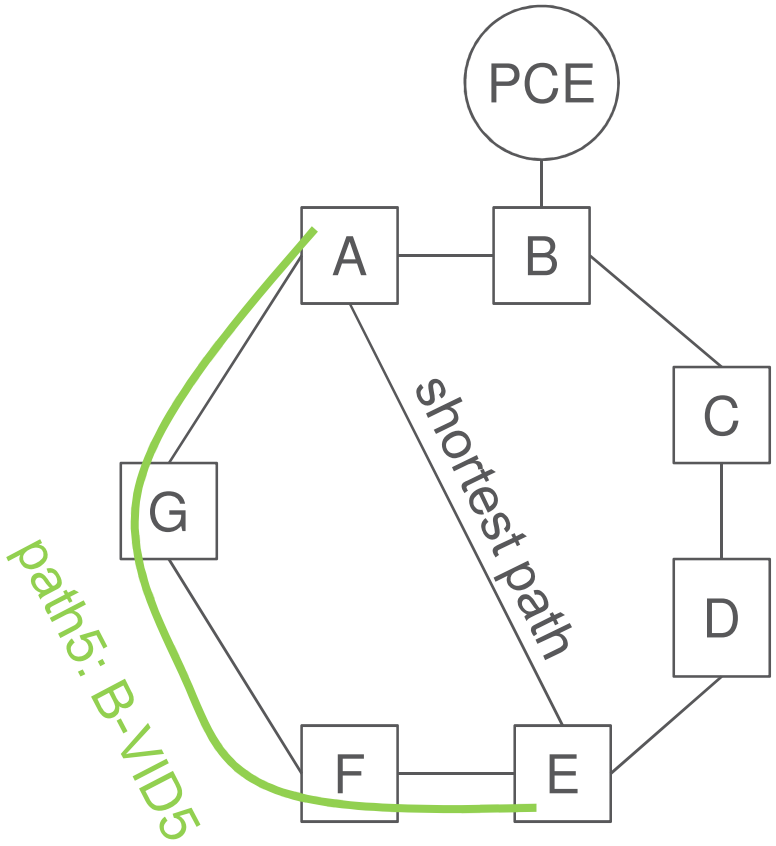

Example 4
ECT Algorithm

- SPB Base VLAN-Identifiers sub-TLV
  - Static path

Type = 6
Length = 12
ECT-Algorithm = **SE ECT Algorithm (00-80-C2-20)**
Base VID = **B-VID5**
U =1
M = 1

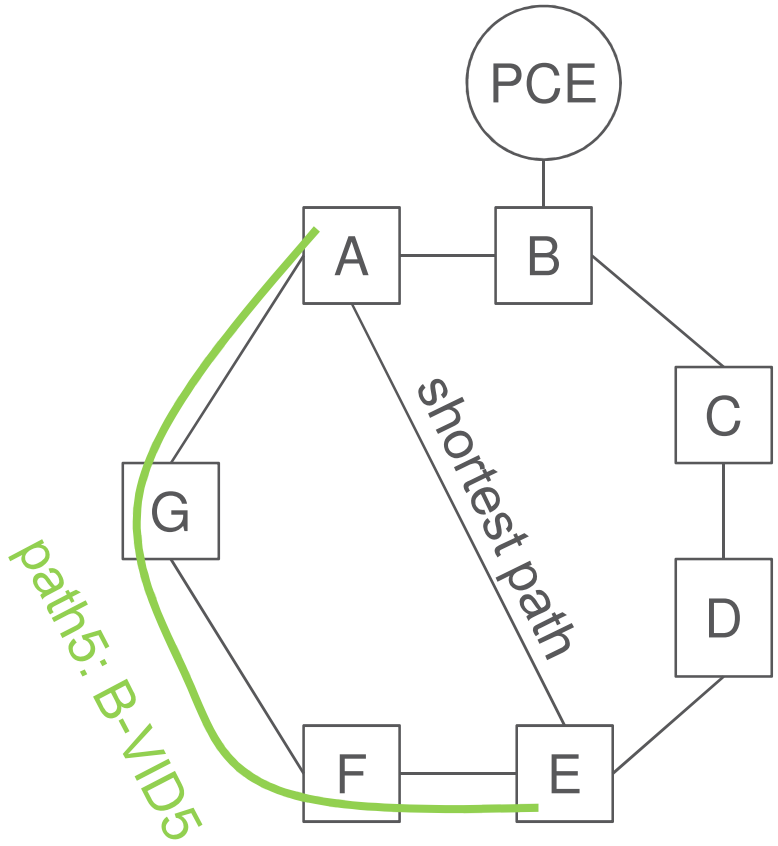

Example 4
MAC DA

- SPBM Service Identifier and Unicast Address sub-TLV provides the association of a B-MAC and I-SID to a B-VID
- Both A and E can Transmit and Receive on I-SID5
  - A: B-MAC-A
  - E: B-MAC-E

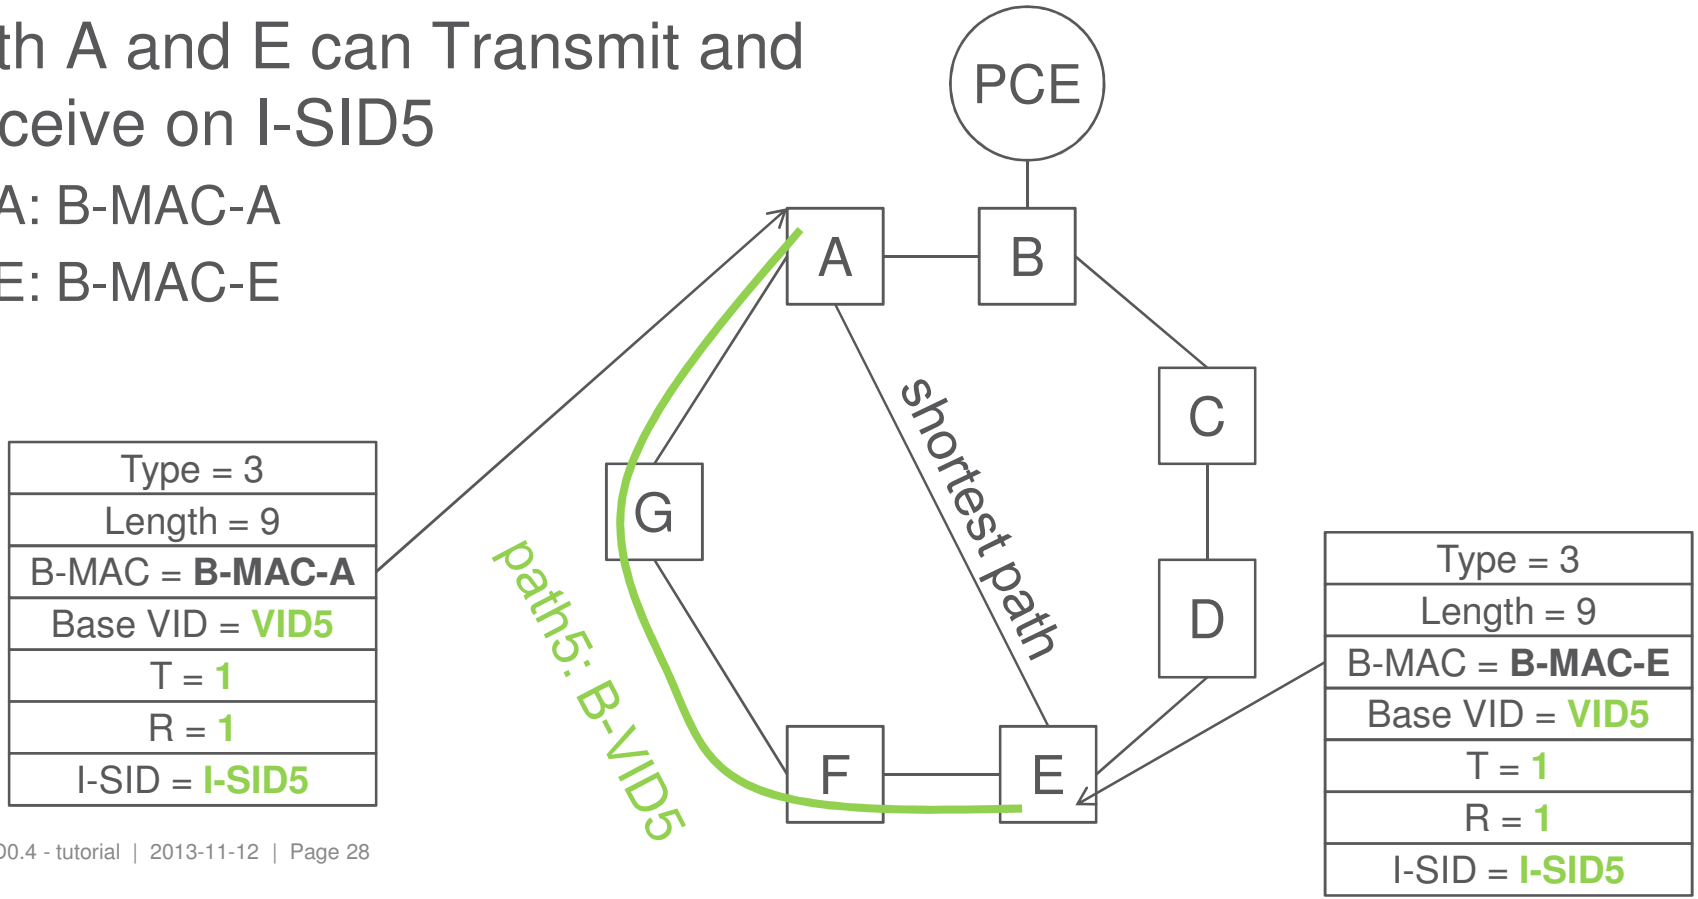

Example 4 Path5

- Fully specified path, i.e. no loose hop

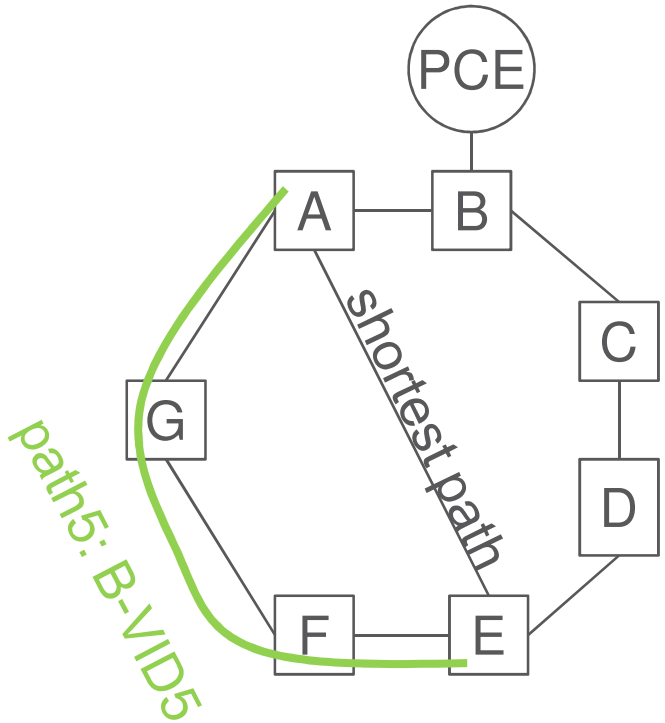

Explicit Path sub-TLV of path3 in an LSP of F:

Type = TBD3
Length = 40
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **B-VID5**
EP Hop sub-TLV for A
EP Hop sub-TLV for G
EP Hop sub-TLV for F
EP Hop sub-TLV for E

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **B**

see more detailed picture in Backup

Example 4
FDB Entries

- Bidirectional, symmetric forwarding is set up

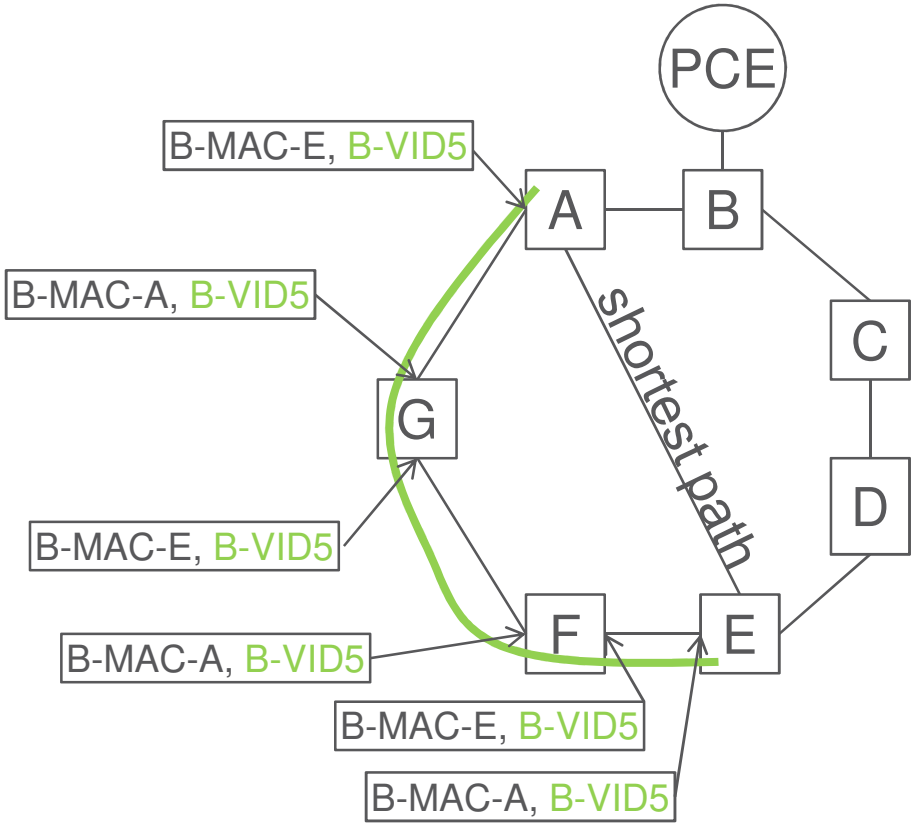

Summary

- 802.1Qca provides the tools for explicit path establishment
  - More features are provided than shown by the above examples
- We have all the sub-TLVs there in D0.4 for path set-up
- It is only a software upgrade
- VIDs can be reused

- P802.1Qca is in progress, just having second Task Group ballot
  - There is work to be done, e.g. interworking with SRP (802.1Qcc)

ERICSSON

Backup Slides

Example 1 Path1

- Fully specified path, i.e. no loose hop

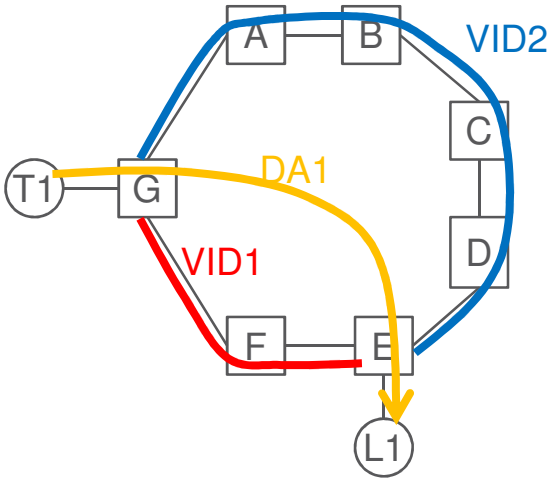

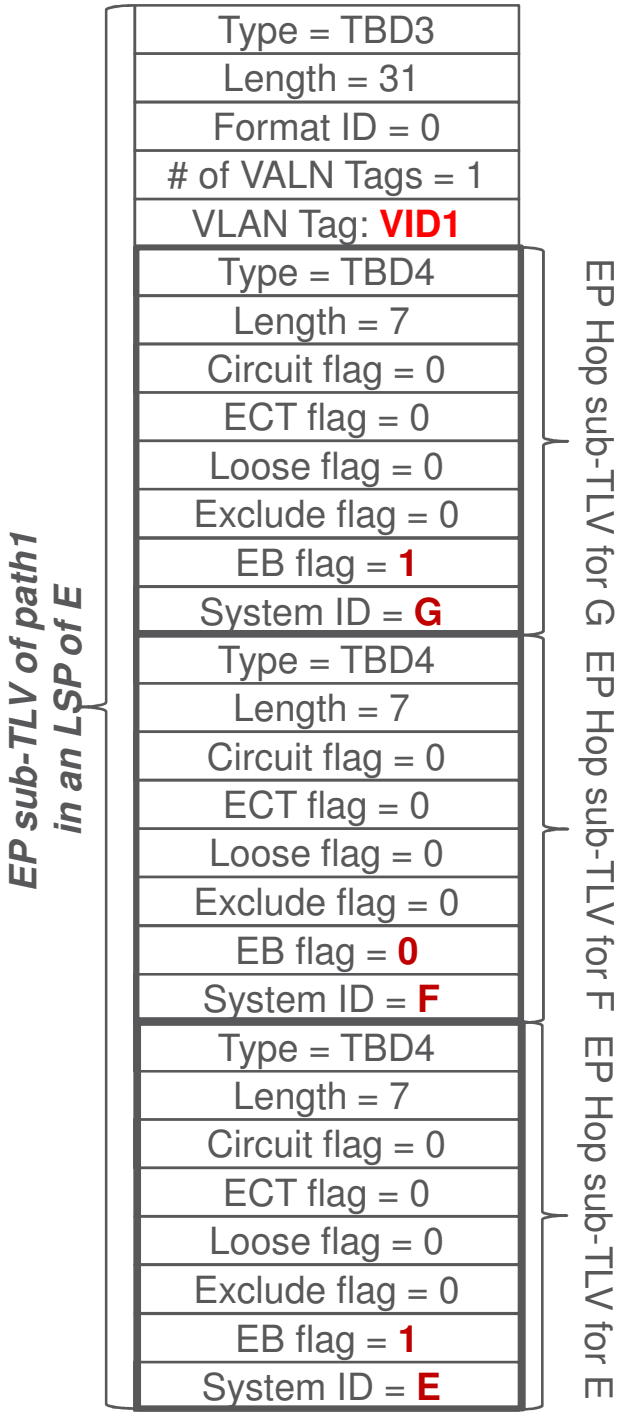

Example 1 Path2

- Fully specified path, i.e. no loose hop
- As each bridge knows the topology, 802.1Qca D0.4 allows arbitrary order of hops, which is under debate

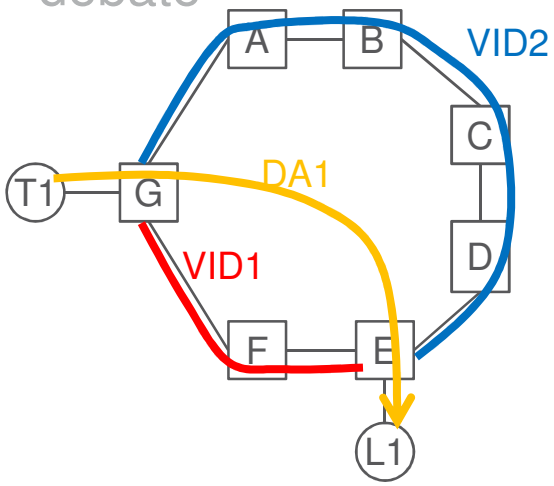

Explicit Path sub-TLV of path2 in an LSP of E:

Type = TBD3
Length = 58
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID2**
EP Hop sub-TLV for C
EP Hop sub-TLV for G
EP Hop sub-TLV for A
EP Hop sub-TLV for E
EP Hop sub-TLV for B
EP Hop sub-TLV for D

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **C**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **G**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **A**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **E**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **D**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **B**

Example 2 Path3

- Fully specified path, i.e. no loose hop

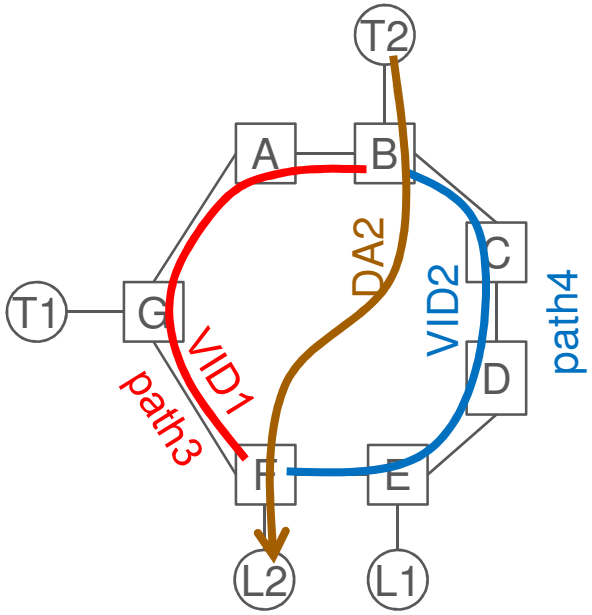

Explicit Path sub-TLV of path3 in an LSP of F:

Type = TBD3
Length = 40
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID1**
EP Hop sub-TLV for A
EP Hop sub-TLV for B
EP Hop sub-TLV for F
EP Hop sub-TLV for G

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **B**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **F**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = A

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **G**

Example 2 Path4

- Fully specified path, i.e. no loose hop

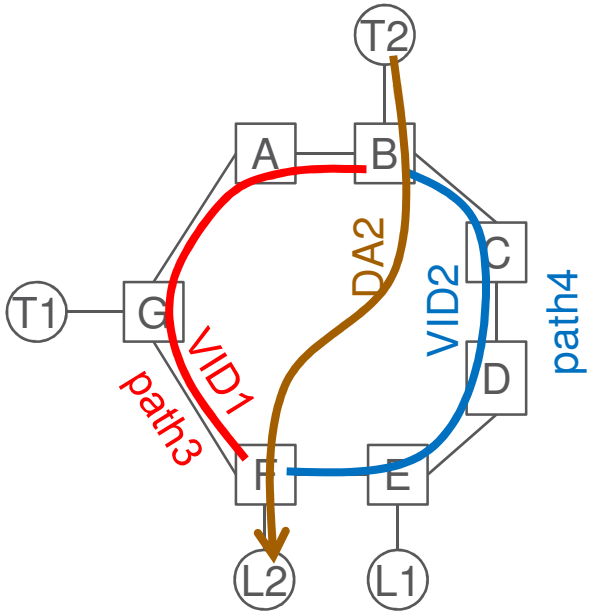

Explicit Path sub-TLV of path4 in an LSP of F:

Type = TBD3
Length = 49
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **VID2**
EP Hop sub-TLV for C
EP Hop sub-TLV for F
EP Hop sub-TLV for D
EP Hop sub-TLV for B
EP Hop sub-TLV for E

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **F**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **D**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **C**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **E**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **B**

Example 4 Path5

- Fully specified path, i.e. no loose hop

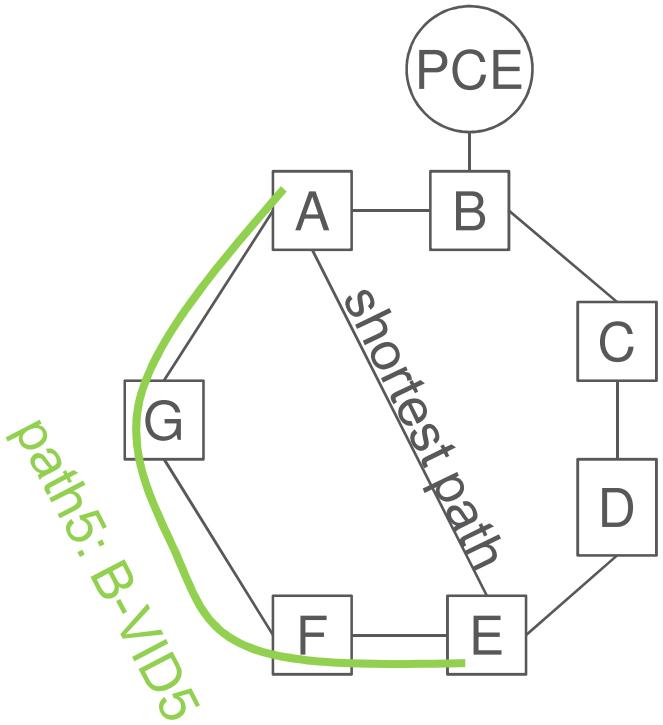

Explicit Path sub-TLV of path3 in an LSP of F:

Type = TBD3
Length = 40
Format ID = 0
# of VALN Tags = 1
VLAN Tag: **B-VID5**
EP Hop sub-TLV for A
EP Hop sub-TLV for E
EP Hop sub-TLV for F
EP Hop sub-TLV for G

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = **E**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **F**

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **1**
System ID = A

Type = TBD4
Length = 7
Circuit flag = 0
ECT flag = 0
Loose flag = 0
Exclude flag = 0
EB flag = **0**
System ID = **G**