DRNI – MEP Placement

Version 01

Stephen Haddock

September 19, 2011

Part 1: DRNI interworking with VLAN based (non-Traffic-Engineered) topologies in the network:

Example E-LAN (mp2mp) service

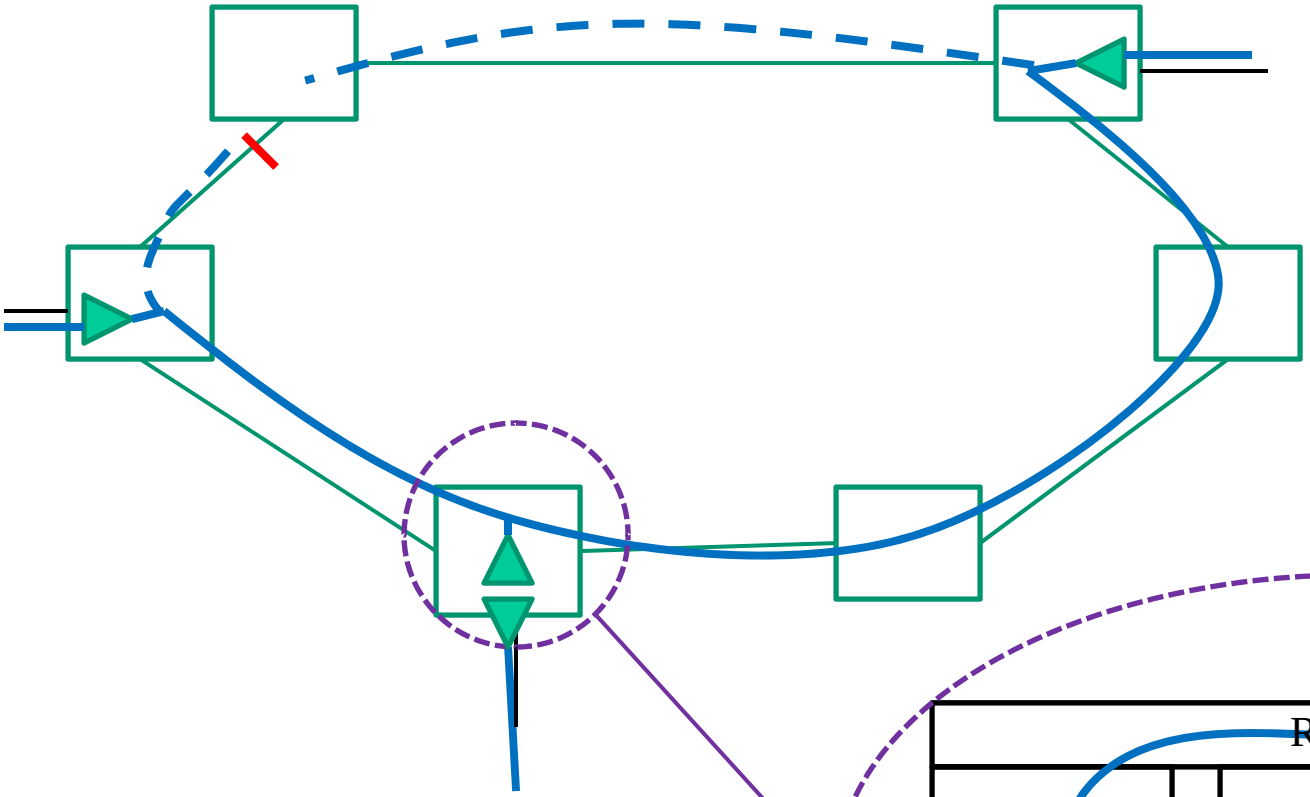

- S-VLAN with 3 points of attachment.
- Solid line is the pruned VLAN path.
- MEP at each point of attachment monitors the VLAN path on the active topology.

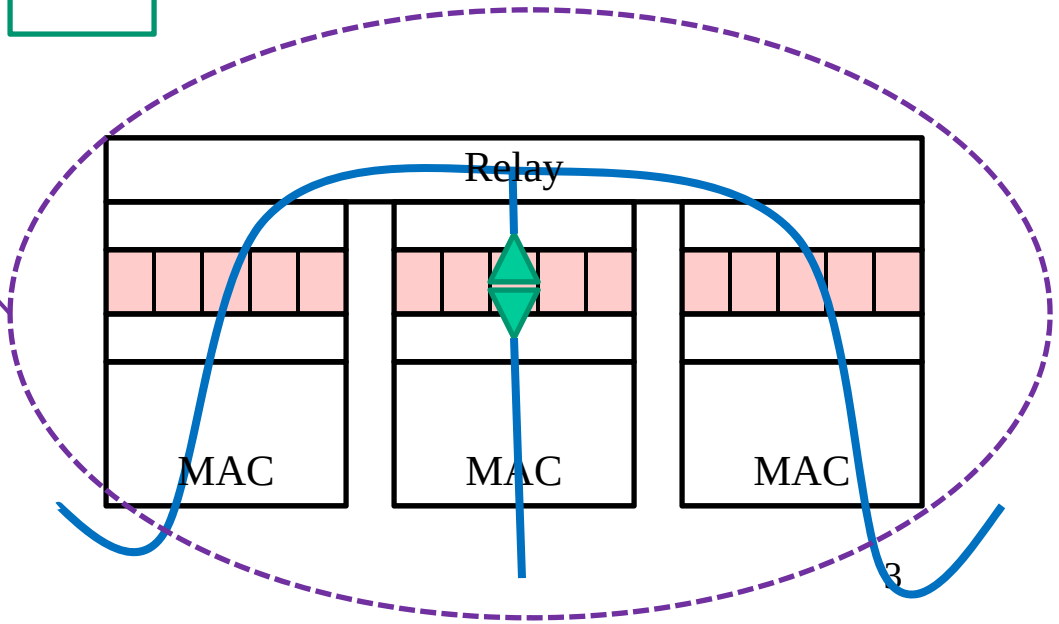

- Bridge model of a Bridge containing one of the points of attachment.

Bridge Model with Link Aggregation

- Bridge model without Link Aggregation.

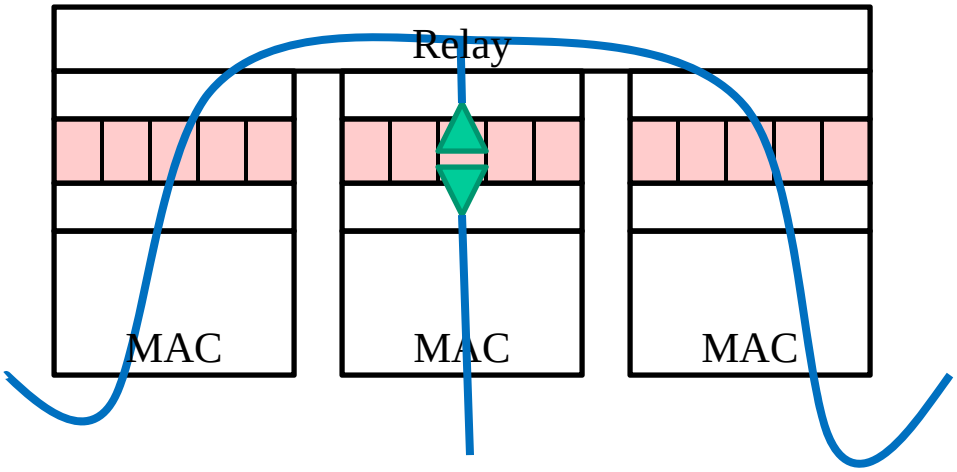

- Bridge model with Link Aggregation.
- The Aggregated Link appears to be a single logical link to a single Bridge Port.
- No change to the MEP placement in the model, though the implementation may need to distribute the MEP functionality across multiple line cards.

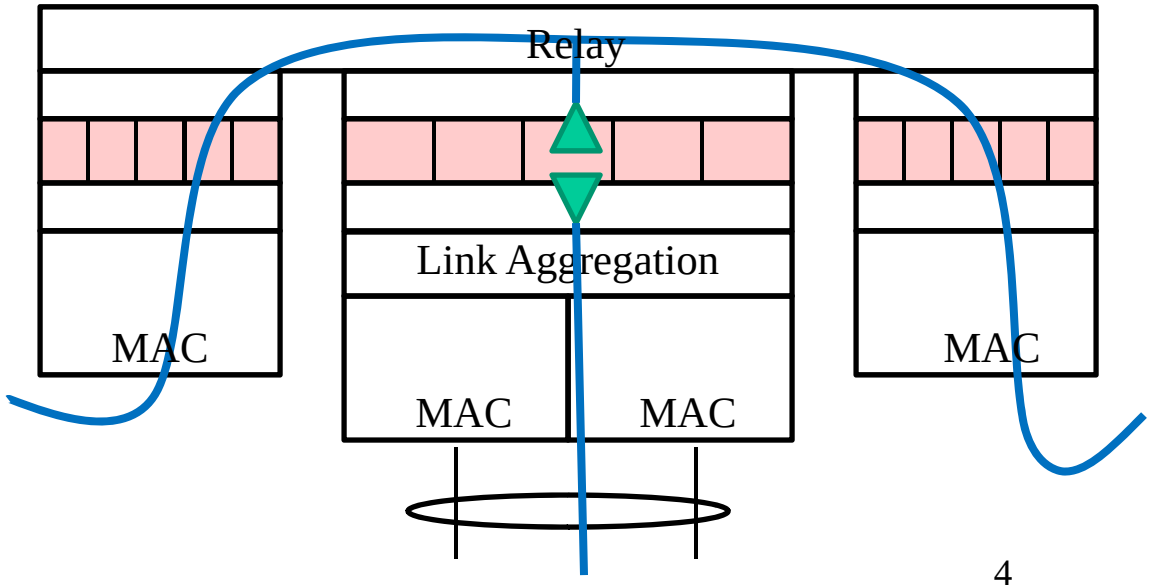

Example with DRNI

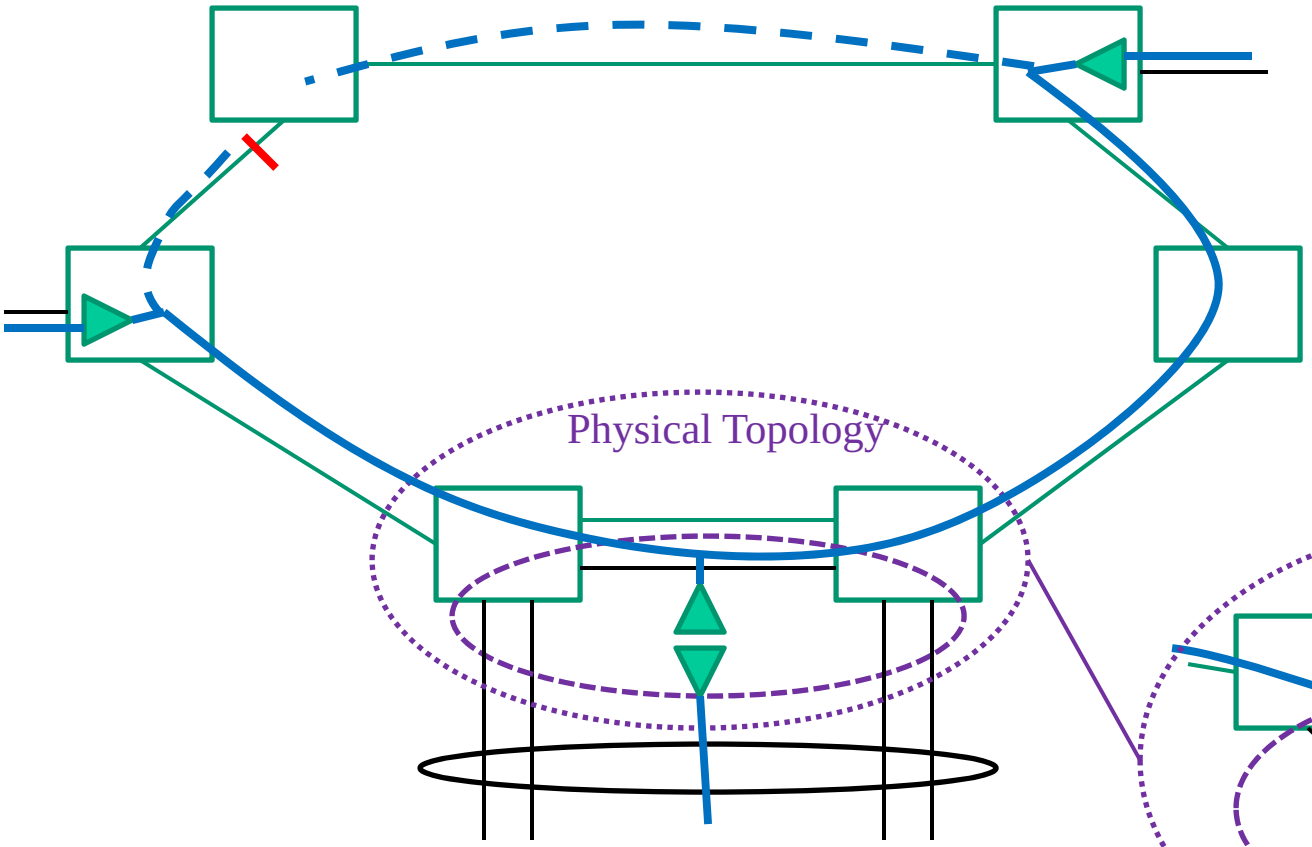

- DRNI created between two bridges.
- The DRNI still appears to be a single point of attachment (a single logical link with a single Bridge Port on a logical bridge) with a single MEP.
- Where is this MEP?

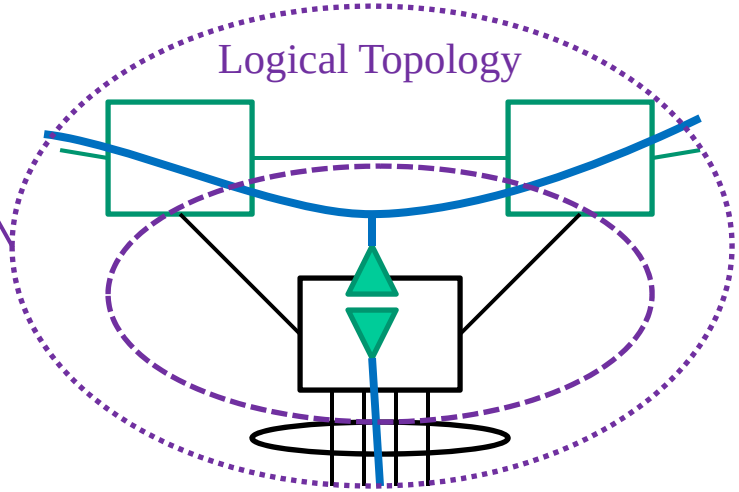

Logical Bridge Model of DRNI

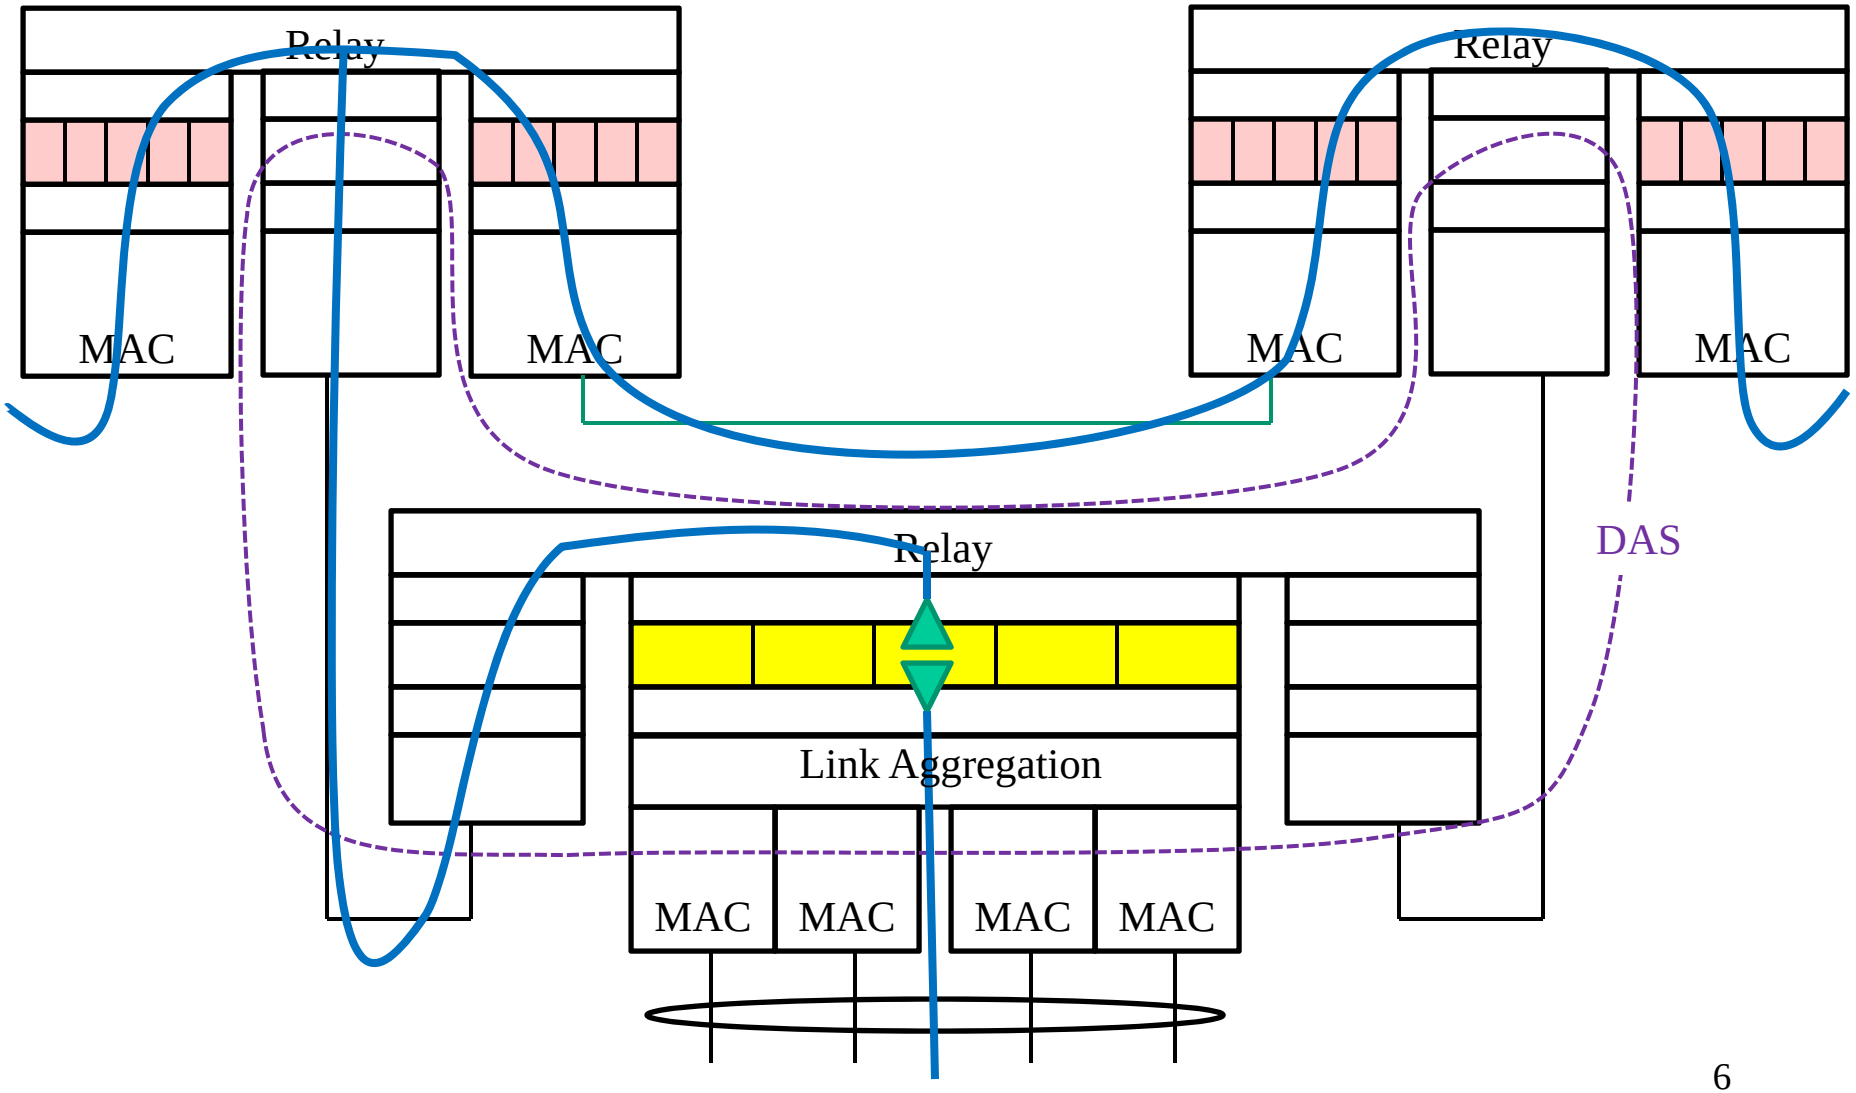

Bridge Model of Distributed Aggregation Sublayer

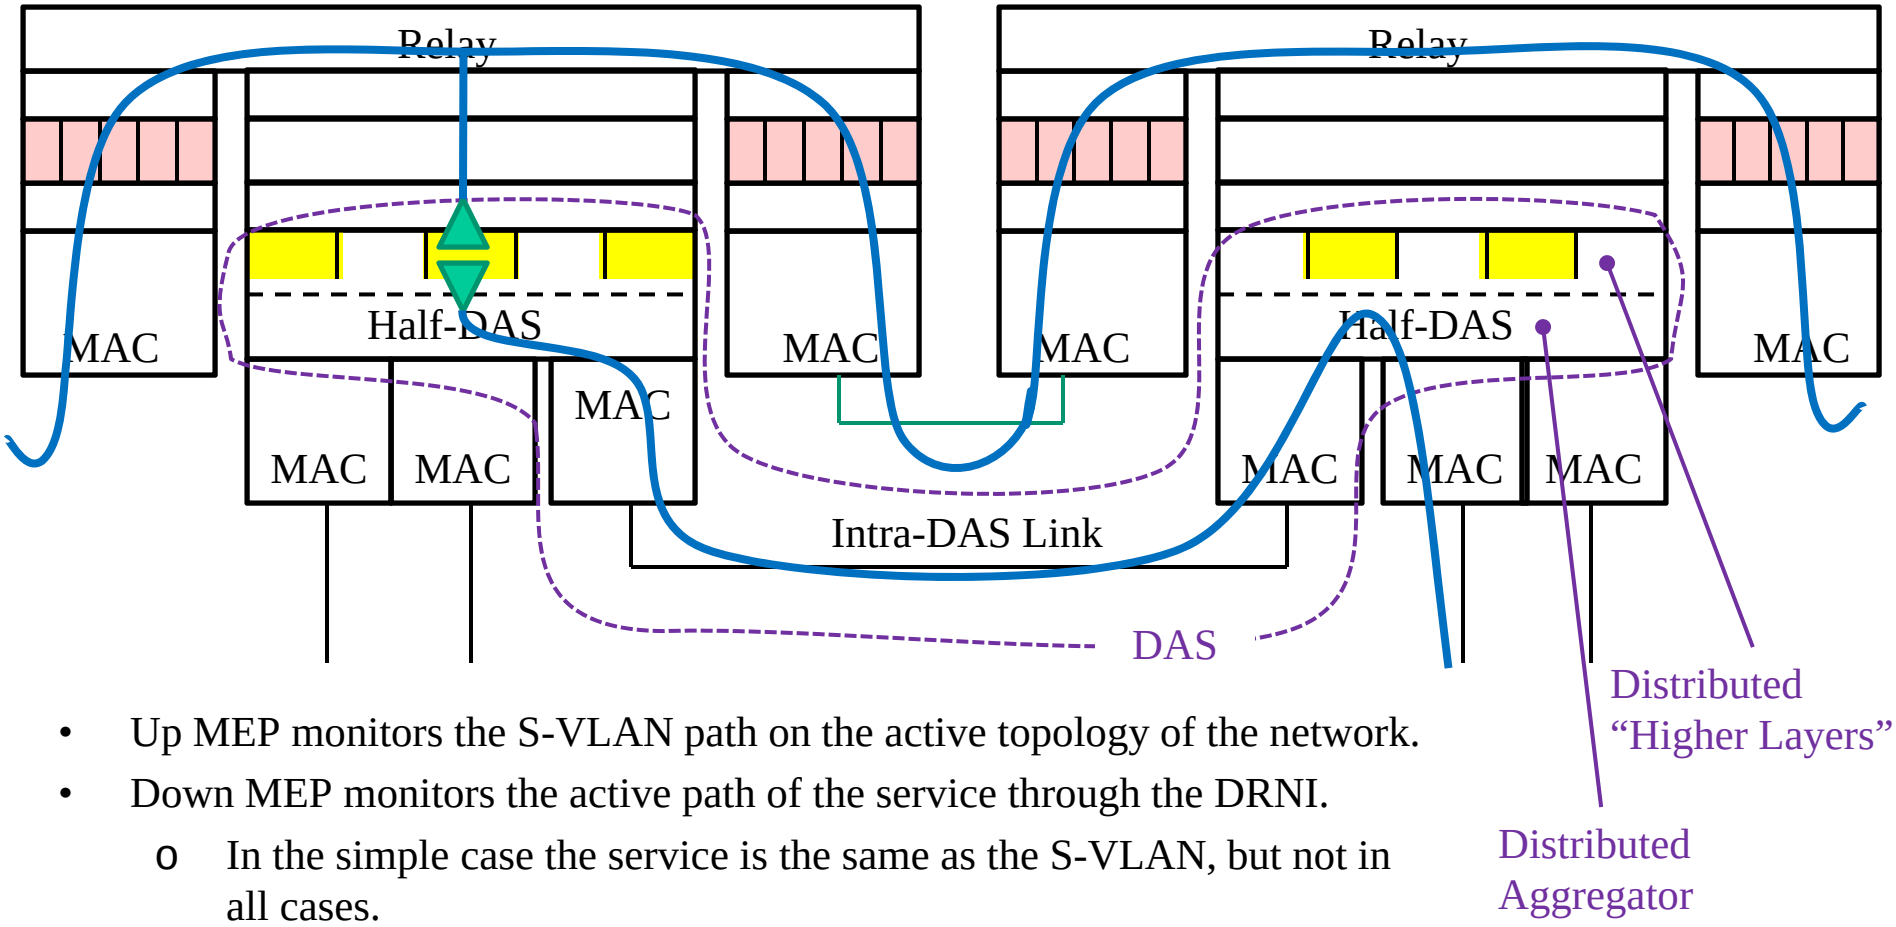

- Up MEP monitors the S-VLAN path on the active topology of the network.
- Down MEP monitors the active path of the service through the DRNI.
  - In the simple case the service is the same as the S-VLAN, but not in all cases.
- May also have Port-based Down MEPs on each physical link of the DRNI (not shown).

Conceptual Model vs. Implementation - 1

- Simple case: S-VID used for both gateway and link selection.
  - MEPs could be either at gateway or at link.
  - The standard needs to specify to assure two different implementations can create a DRNI between them.
  - I believe the MEPs should be at the gateway, because this will work not only for the simple case, but for the exceptions.
- Exception #1: Backwards compatibility objective
  - Need to interoperate with a legacy implementation that cannot do link selection based on a service identifier.
  - Means frames for the service may be received on multiple links.
  - MEP functionality (at least performance counters, and perhaps CFM message reception on Down MEP) may need to be distributed across multiple links and therefore possibly multiple line cards.
  - How this is done will be very implementation dependent, but locating the MEPs at the gateway contains the problem to a single bridge, so becomes the same problem that already needs to be solved for traditional LAG.
- Exception #2: DRNI on Backbone Edge Bridges

Are MEPs “close enough” to the network “edge”?

- The gateway is the appropriate place for Up MEPs monitoring the flow through the network and Down MEPs monitoring the flow across the DRNI .
  - Loop prevention protocols (RSTP/MSTP, G.8032, SPB, …) are responsible for providing an active path through the network from Gateway to Gateway. The MEPs monitor the active path. A failure or topology change in the network may result in a change to the active path, however there are no there are no other MEPs at the same level that monitor potential alternate paths through the network nor do the MEPs at the Gateway play a role in restoration of the active path following a failure or topology change.
    - Note this is fundamentally different than edge-to-edge Protection Switching.
  - The DRNI protocols are responsible for providing an active path through the DRNI from Gateway to Gateway. The MEPs monitor the active path. A failure or change in link selection may result in a change to the active path, however there are no service level MEPs at the Gateway that monitor potential alternate paths across the DRNI nor do the MEPs at the Gateway play a role in the restoration of the active path following a failure or change in link selection.

Logical Topology of DRNI with IB-BEBs

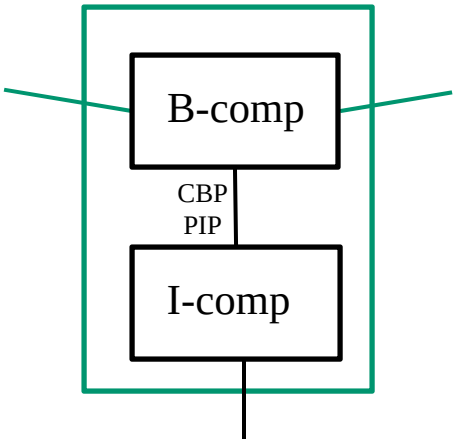

Single Link

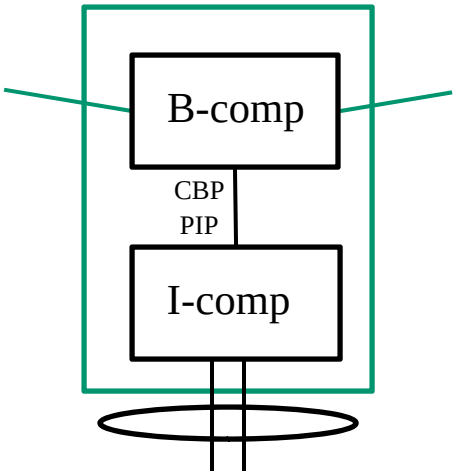

Link Aggregation

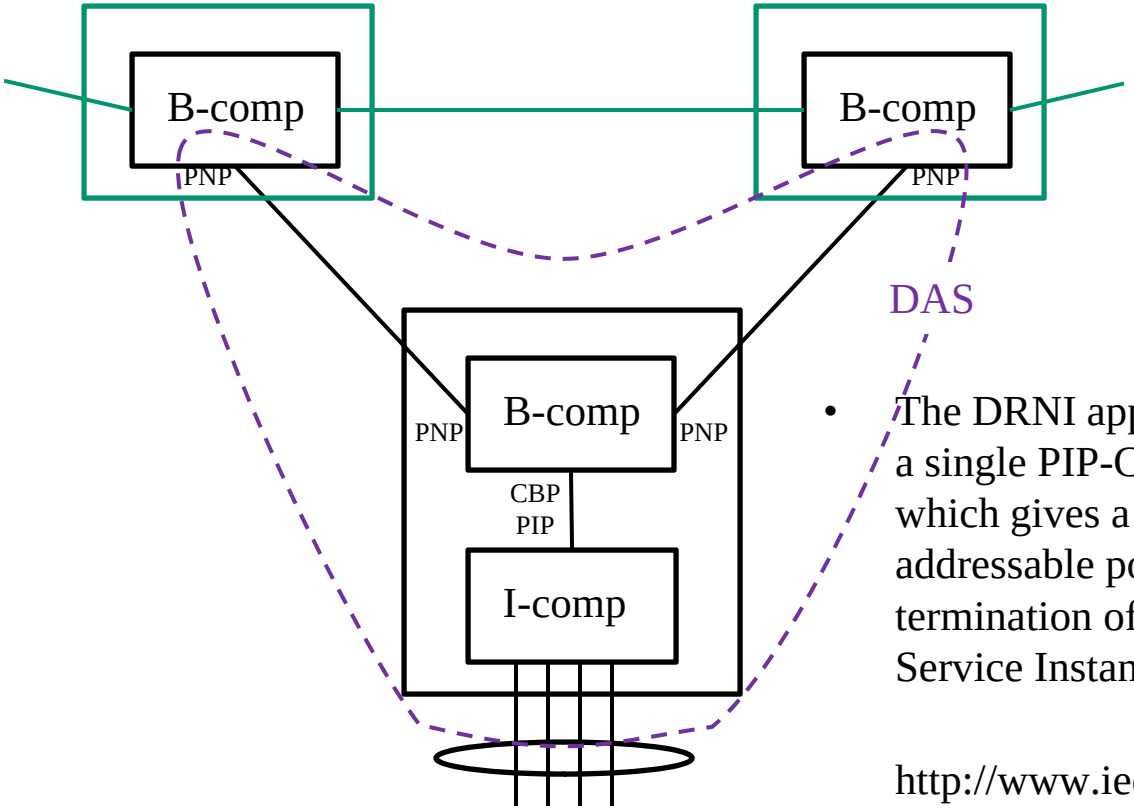

DRNI

- The DRNI appears to have a single PIP-CBP pair which gives a single addressable point for the termination of Backbone Service Instances. See:

http://www.ieee802.org/1/files/public/docs2010/new-haddock-resilient-network-interconnect-addressing-1110-v1.pdf

MEP Placement in DRNI with IB-BEBs

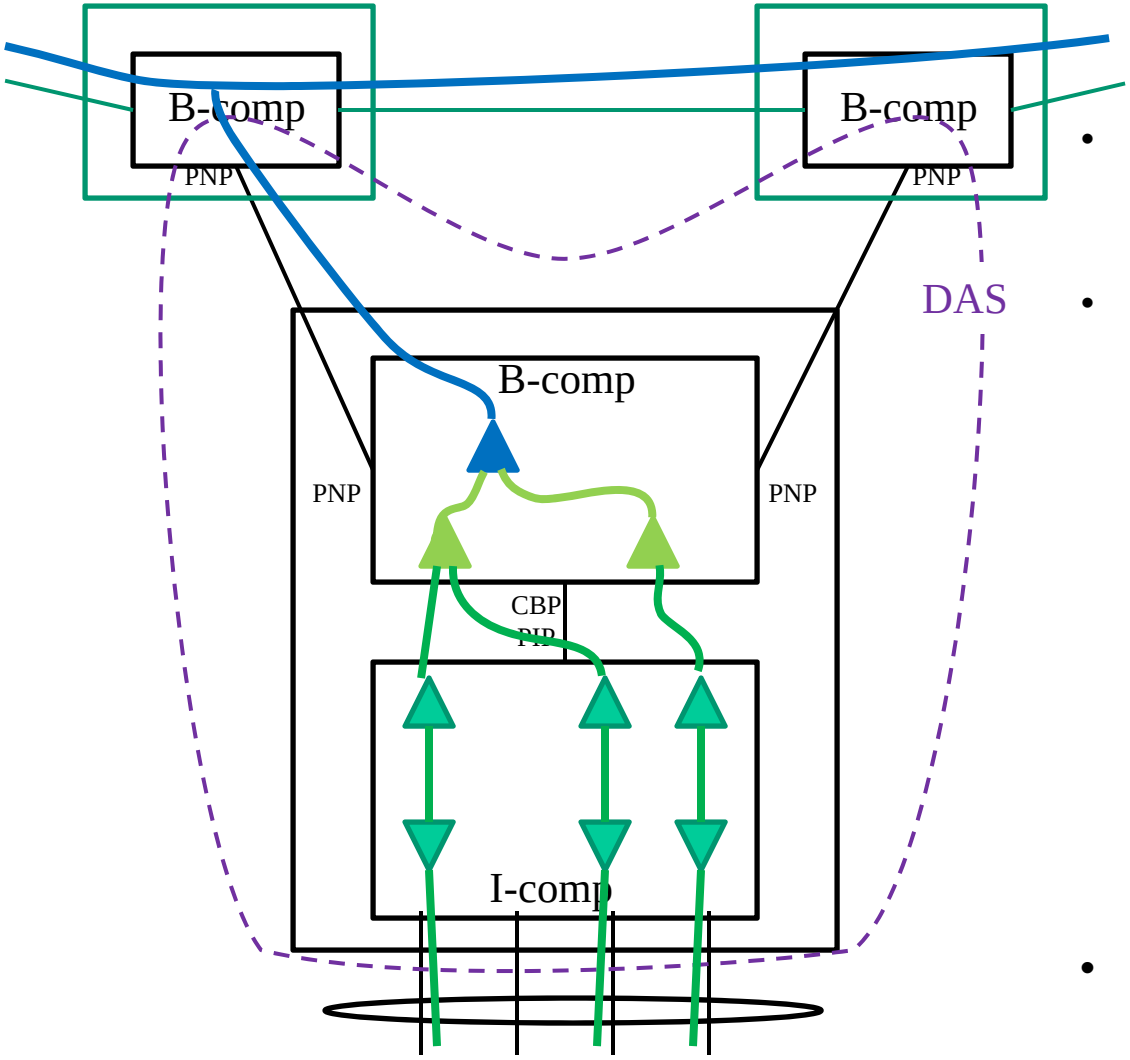

- Example shows three S-VLANs mapping to two Backbone Service Instances mapping to one B-VLAN
- MEP placements:
  1. One MEP at logical CBP monitoring B-VLAN.
  2. Two MEPs at logical CBP (or at logical PIP) monitoring BSIs.
  3. Three MEPs at logical PIP monitoring S-VLANs through the network.
  4. Three MEPs at logical Aggregated Link monitoring S-VLANs across the DRNI
- Note DRNI link selection is based on S-VID, while gateway selection is based on B-VID.

Bridge Model of DRNI with IB-BEBs

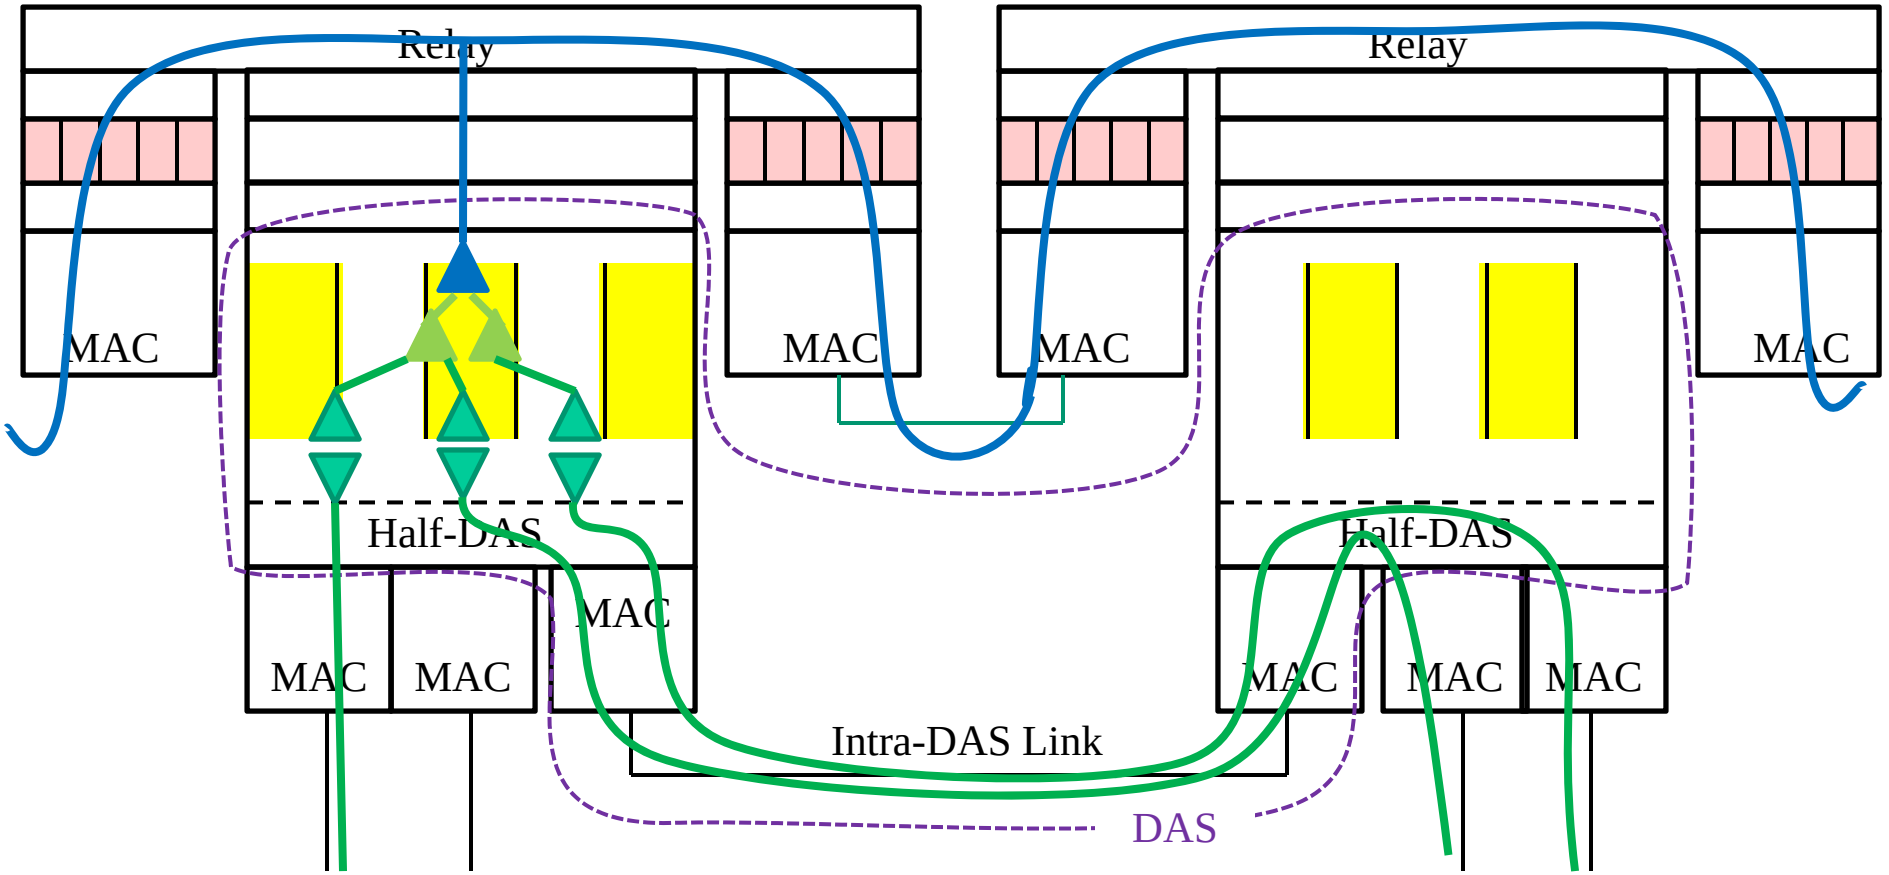

Conceptual Model vs. Implementation - 2

- Conceptually all MEPs are at the Gateway.
  - The S-VLAN up and down MEPs could be viewed as either at the selected Gateway or the selected Link, but putting them at the edge of the DRNI protocol domain means putting them at the Gateway.
  - The BSI level and B-VLAN level MEPs really need to be at the Gateway.
    - Puts them at the edge of the network and DRNI protocol domains.
    - The potential multiplexing of S-VLANs into BSIs and BSIs into B-VLANs mean there is at least a possibility that a single flow is distributed across multiple links, but conceptually putting the MEPs at the Gateway means the MEP functionality only needs to be distributed across links on one physical bridge (the DRNI links and the Intra-DAS link).
    - Allows maximum implementation flexibility for where MAC-in-MAC encapsulation occurs relative to switching frames across the Intra-DAS Link.
  - Allows the physical implementation of distributing the functionality of the active MEPs to be contained within one physical bridge. No need for the active MEPs to be distributed across physical bridges.

Part 2: DRNI interworking with Protection Switching in the network:

Example E-Line (p2p) service with Protection Switching

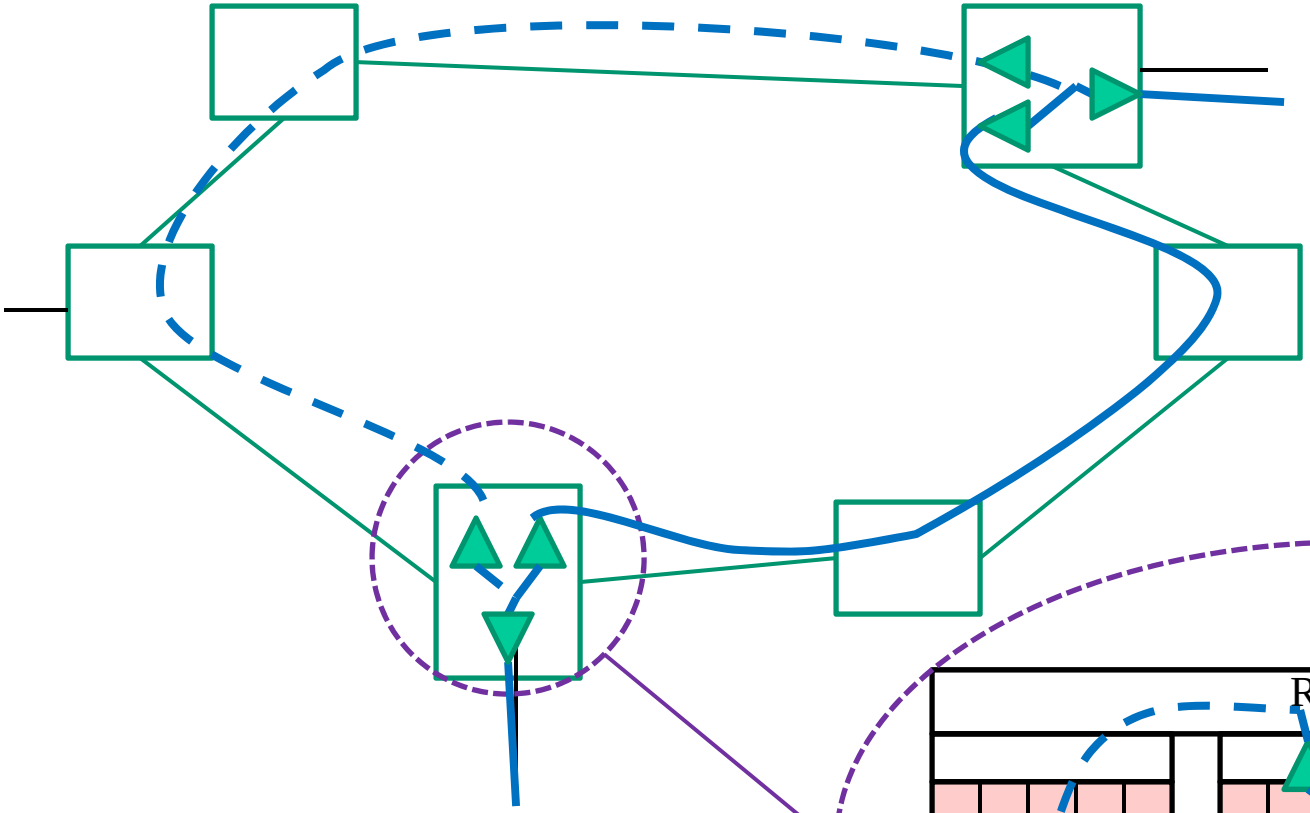

- Two points of attachment connected by two disjoint paths.
- Solid line is the working path; dashed line is the protection path.
- MEPs monitor each path.

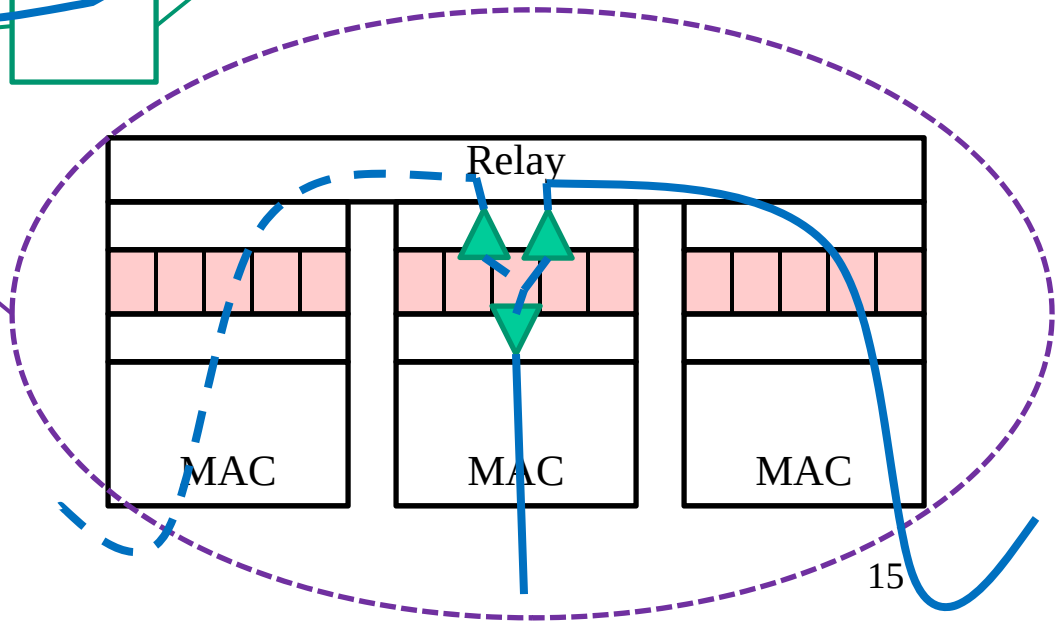

- Bridge model of a Bridge containing one of the points of attachment.

Bridge Model with Link Aggregation

- Bridge model without Link Aggregation.

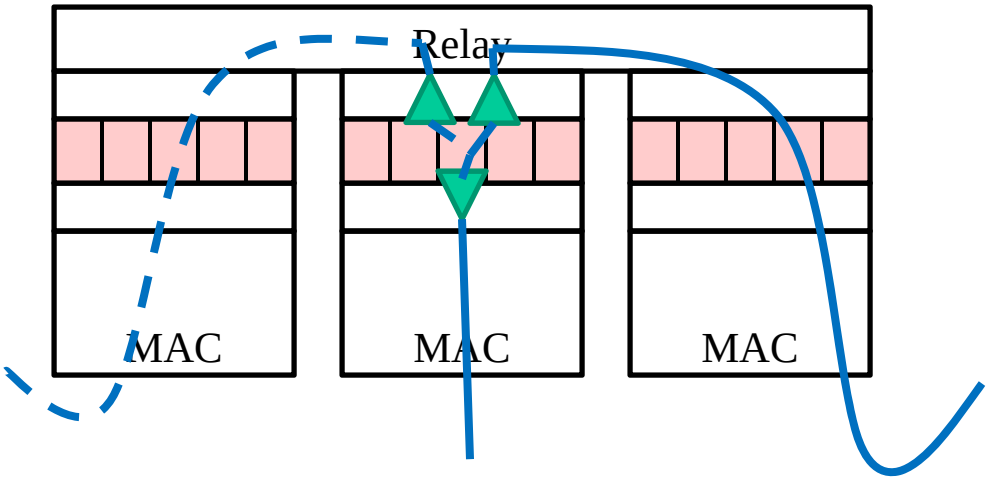

- Bridge model with Link Aggregation.
- The Aggregated Link appears to be a single logical link to a single Bridge Port.
- No change to the MEP placement in the model, though the implementation may need to distribute the MEP functionality.

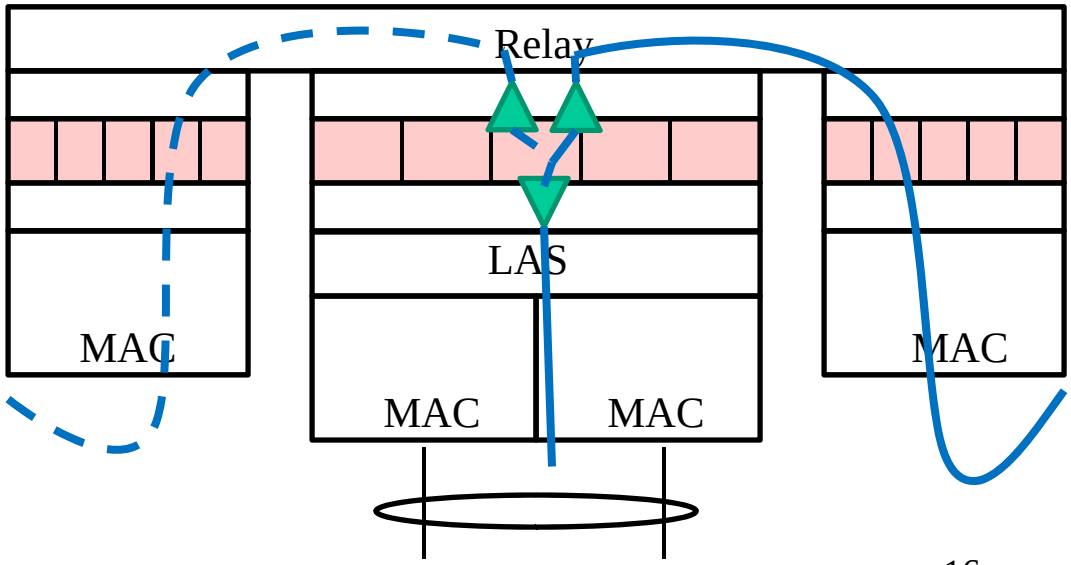

Example with DRNI

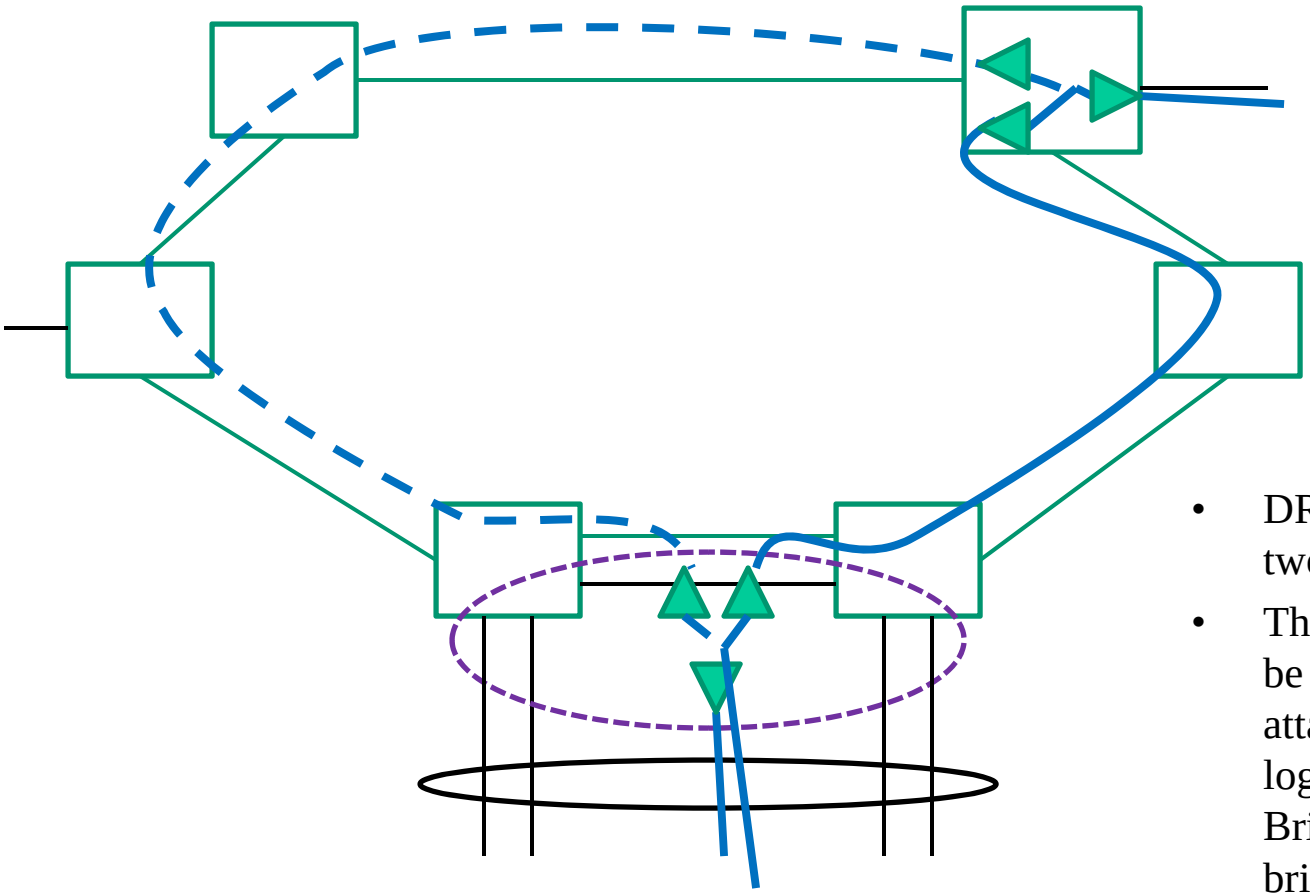

- DRNI created between two bridges.
- The DRNI still appears to be a single point of attachment (a single logical link with a single Bridge Port on a logical bridge).
- Where are these MEPs?

Logical Bridge Model of DRNI with Infrastructure Protection Switching

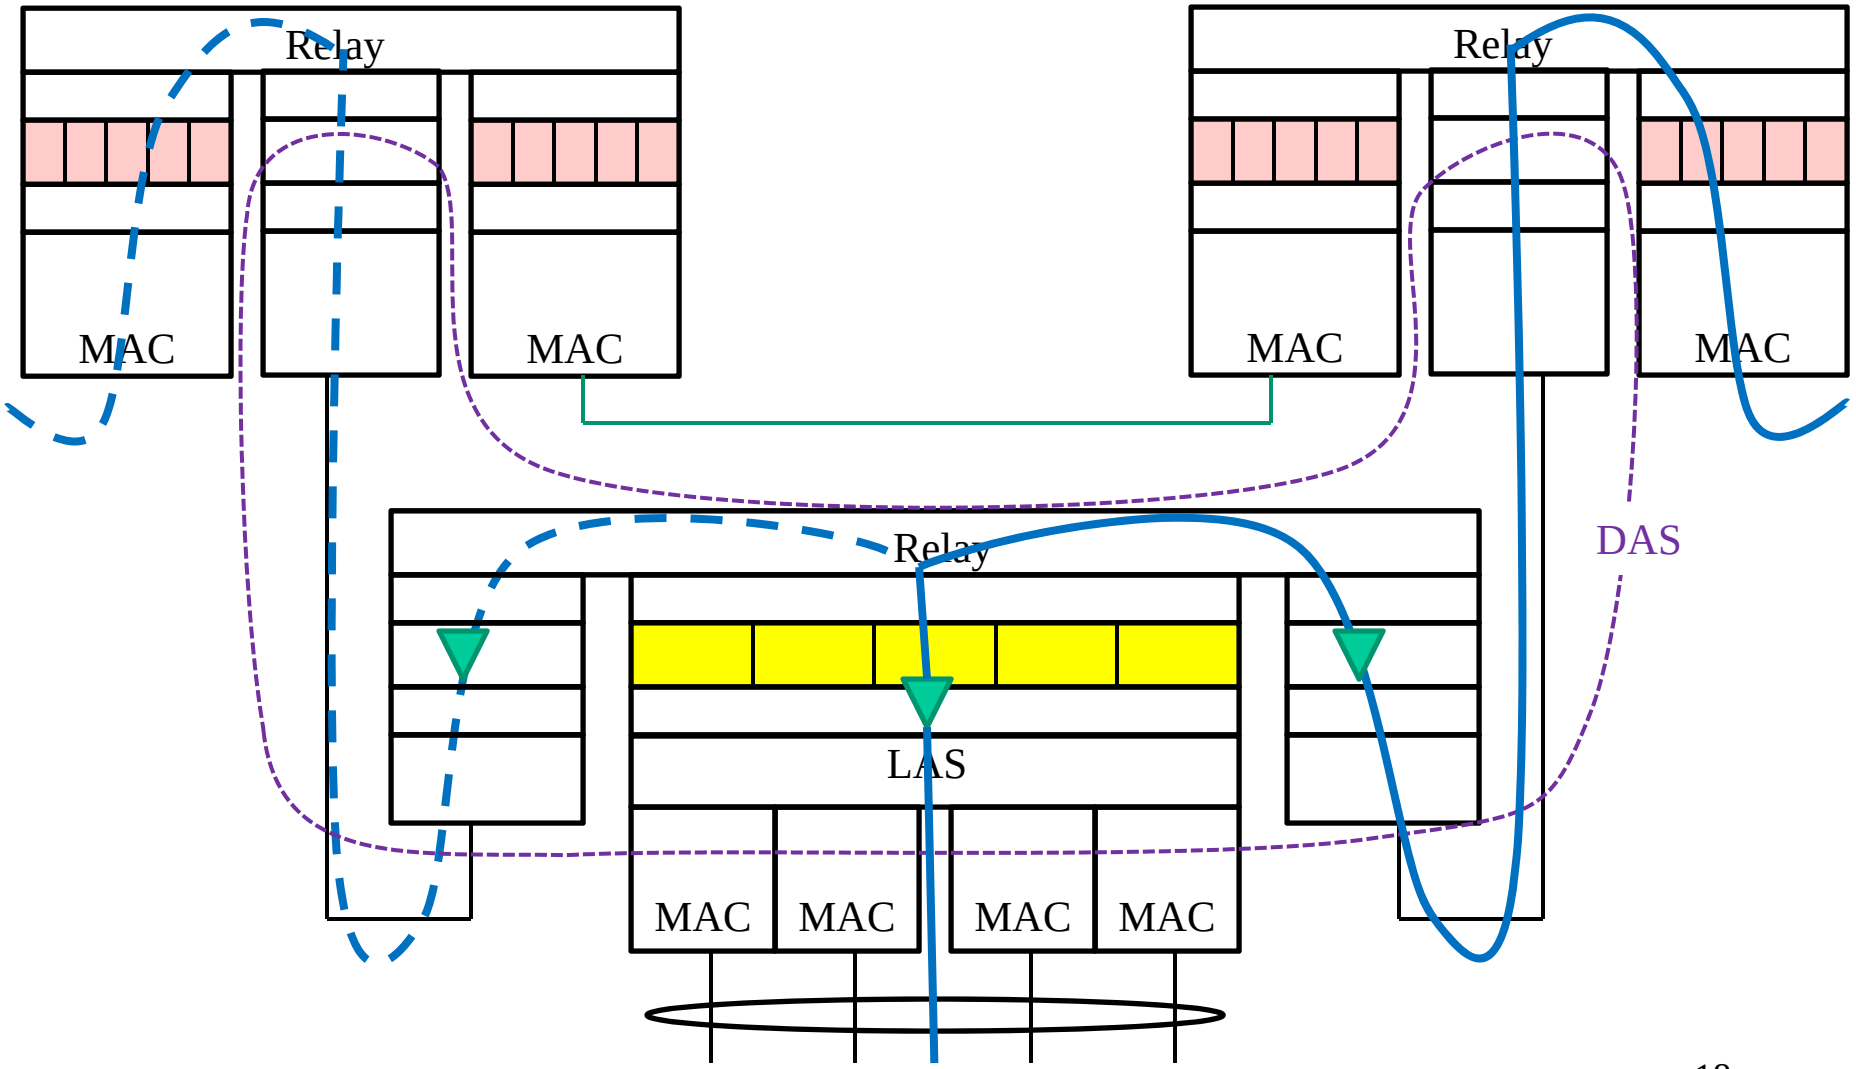

Bridge Model of DRNI with Infrastructure Segment Protection

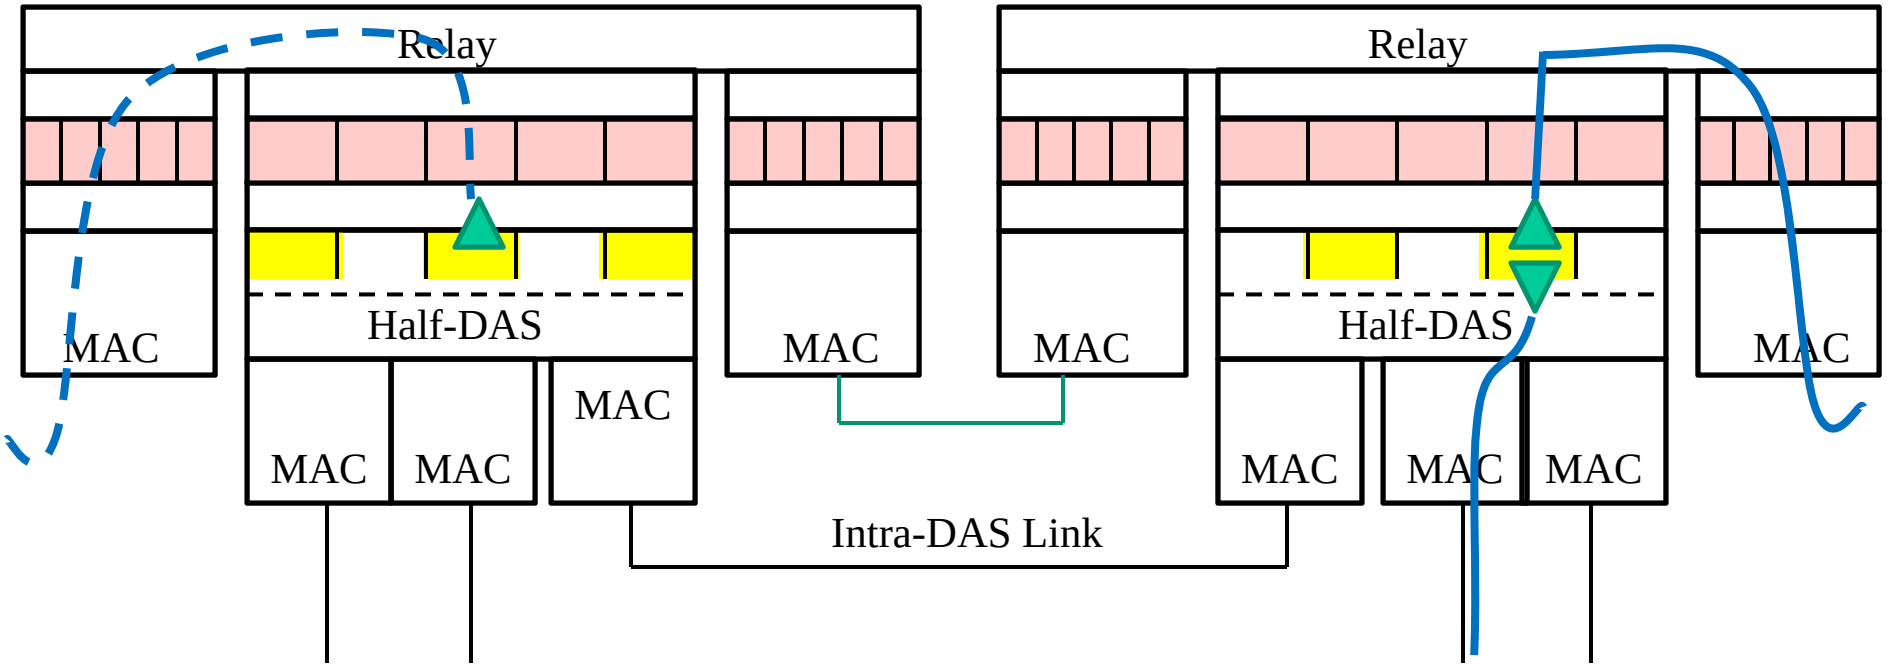

MEP Placement in DRNI with PBB-TE Protection Switching

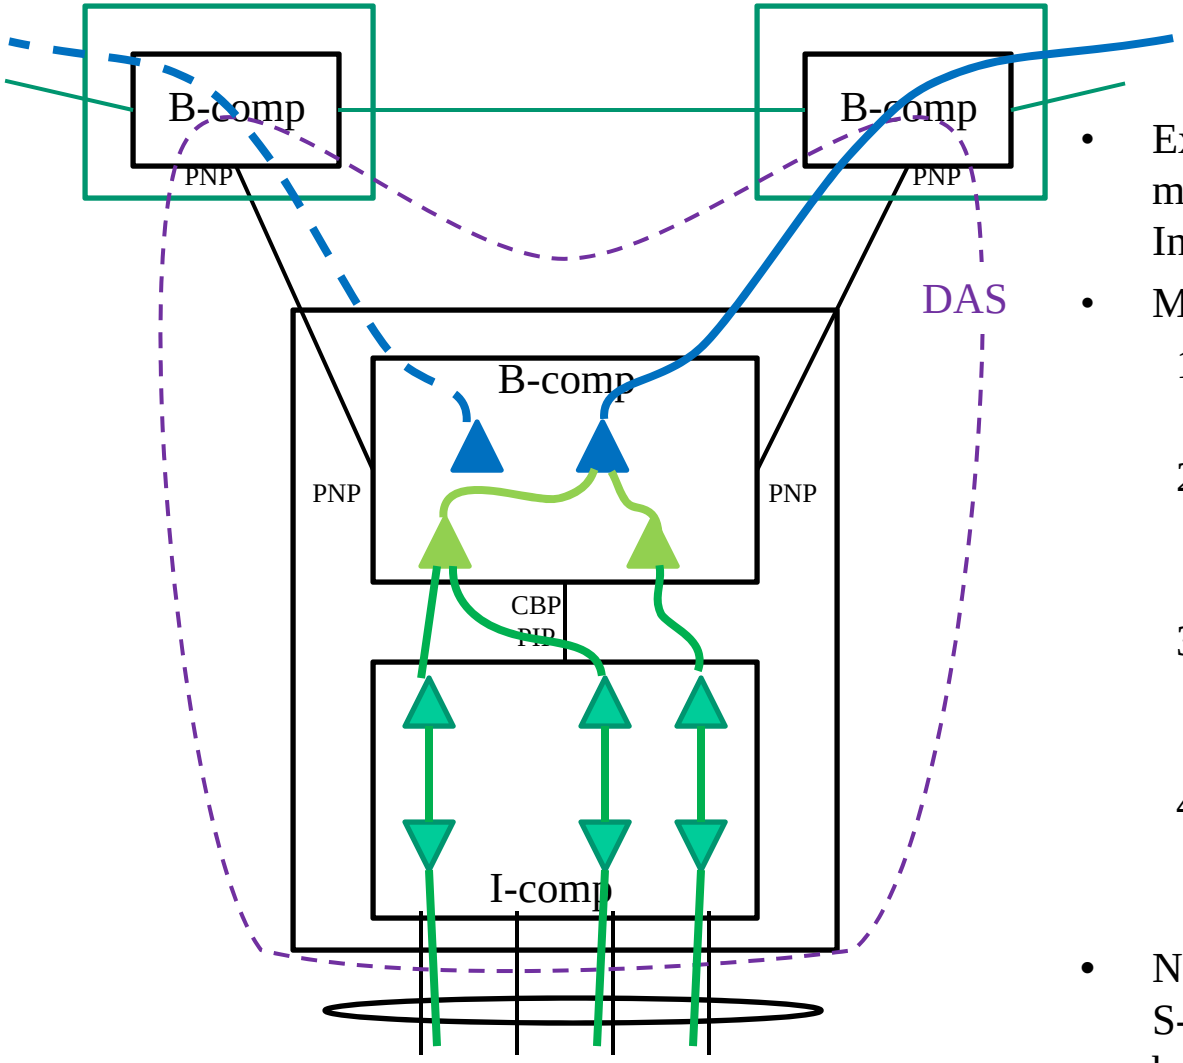

- Example shows three S-VLANs mapping to two Backbone Service Instances mapping to one B-VLAN
- MEP placements:
  1. Two MEPs at logical CBP monitoring TESIs.
  2. Two MEPs at logical CBP (or at logical PIP) monitoring BSIs.
  3. Three MEPs at logical PIP monitoring S-VLANs through the network.
  4. Three MEPs at logical Aggregated Link monitoring S-VLANs across the DRNI
- Note DRNI link selection is based on S-VID, while gateway selection is based on B-VID (or TESI).

Bridge Model of DRNI with PBB-TE Protection Switching

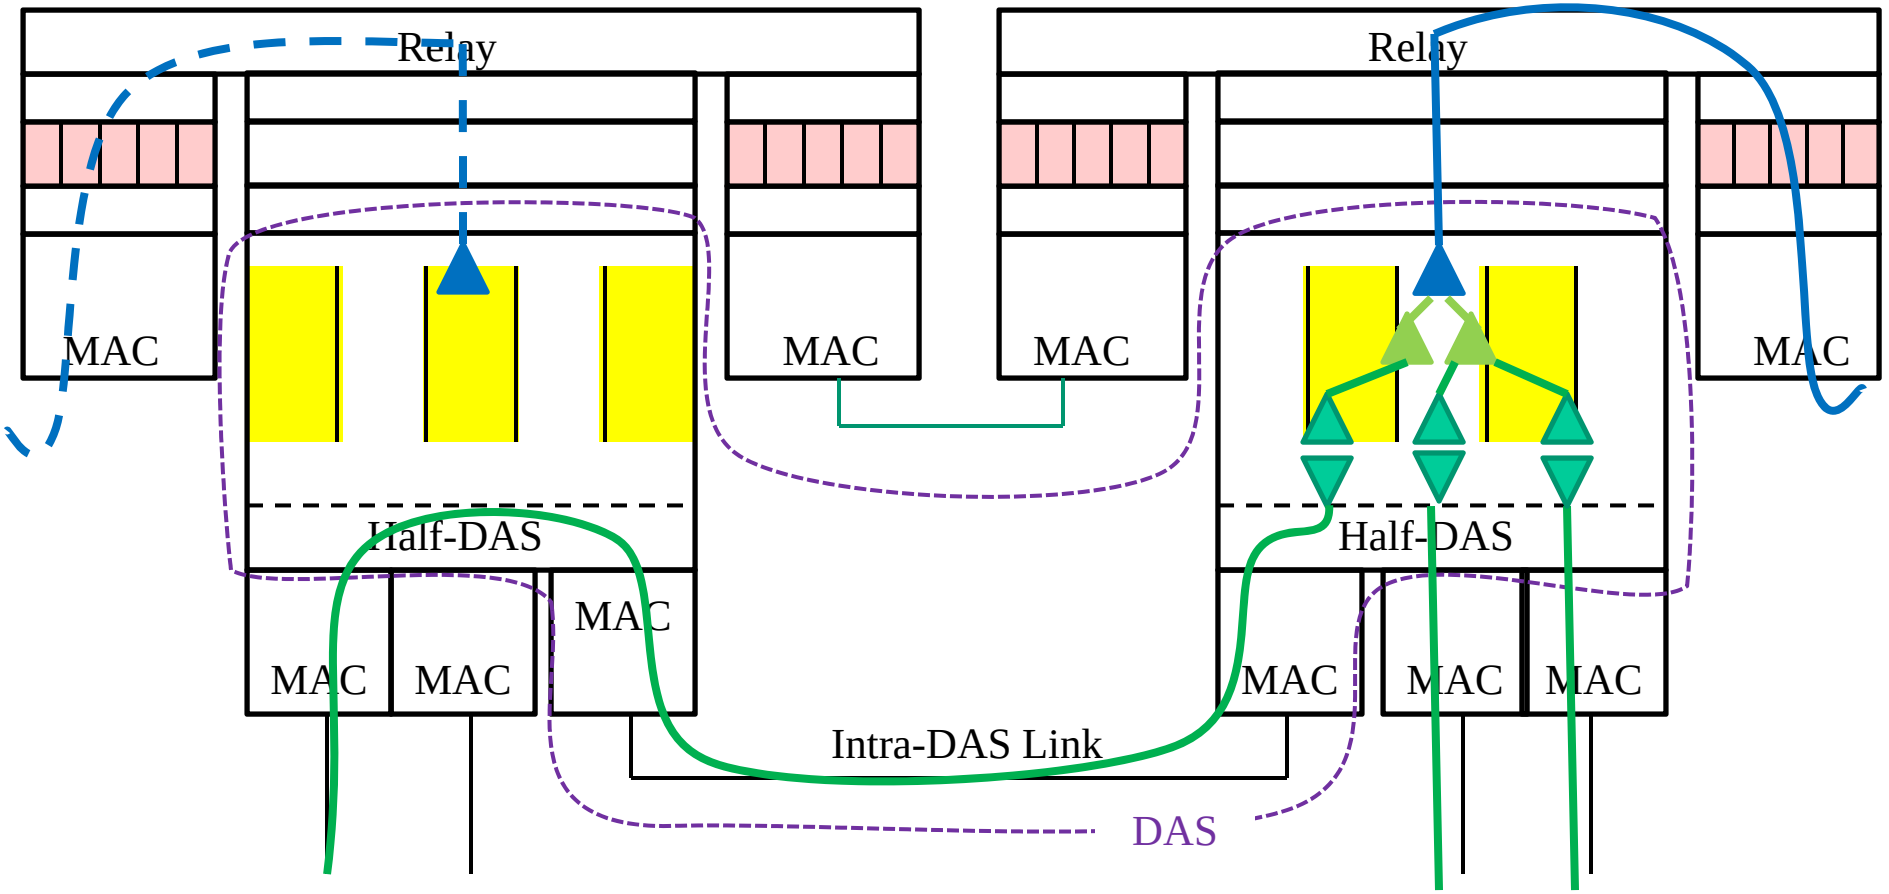

Conceptual Model vs. Implementation - 3

- Conceptually the MEPs should still be at the Gateway.
  - The Gateway is still the boundary between the domains of the protocols that control resiliency in the DRNI and in the network.
  - The network now uses Protection Switching for resiliency; the resiliency protocol in the DRNI is the same as in previous examples.
  - For Infrastructure Segment Protection, the two possible Gateways are logically the same as the two physical ports where the Working and Protection MEPs are located in a non-distributed interface.
  - For PBB-TE Protections Switching, the Working and Protection MEPs monitor TESIs, and as with PBB B-VLAN MEPs the only logical placement for them is at the Gateway.

Conclusions

- The architecturally consistent place to model the MEPs on an DRNI logical port is at the selected .
  - The Gateways are the boundary between the protocol domains for the resiliency protocols in the network and in the DRNI.
  - In some examples the MEPs could be modeled at either the selected Gateway or at the selected Link, however there are other examples where the MEP cannot be modeled as being at a single Link (i.e. packets flowing through the same MEP may come across different DRNI links).
- In some physical implementations the MEP functionality may need to be distributed across multiple line cards or port ASICs.
  - Modeling the MEPs at the Gateway means the MEP functionality may need to be distributed between the DRNI Links and Intra-DAS Link on a single physical bridge, but not between bridges.