



DRNI IB-BEBs

Rev. 4

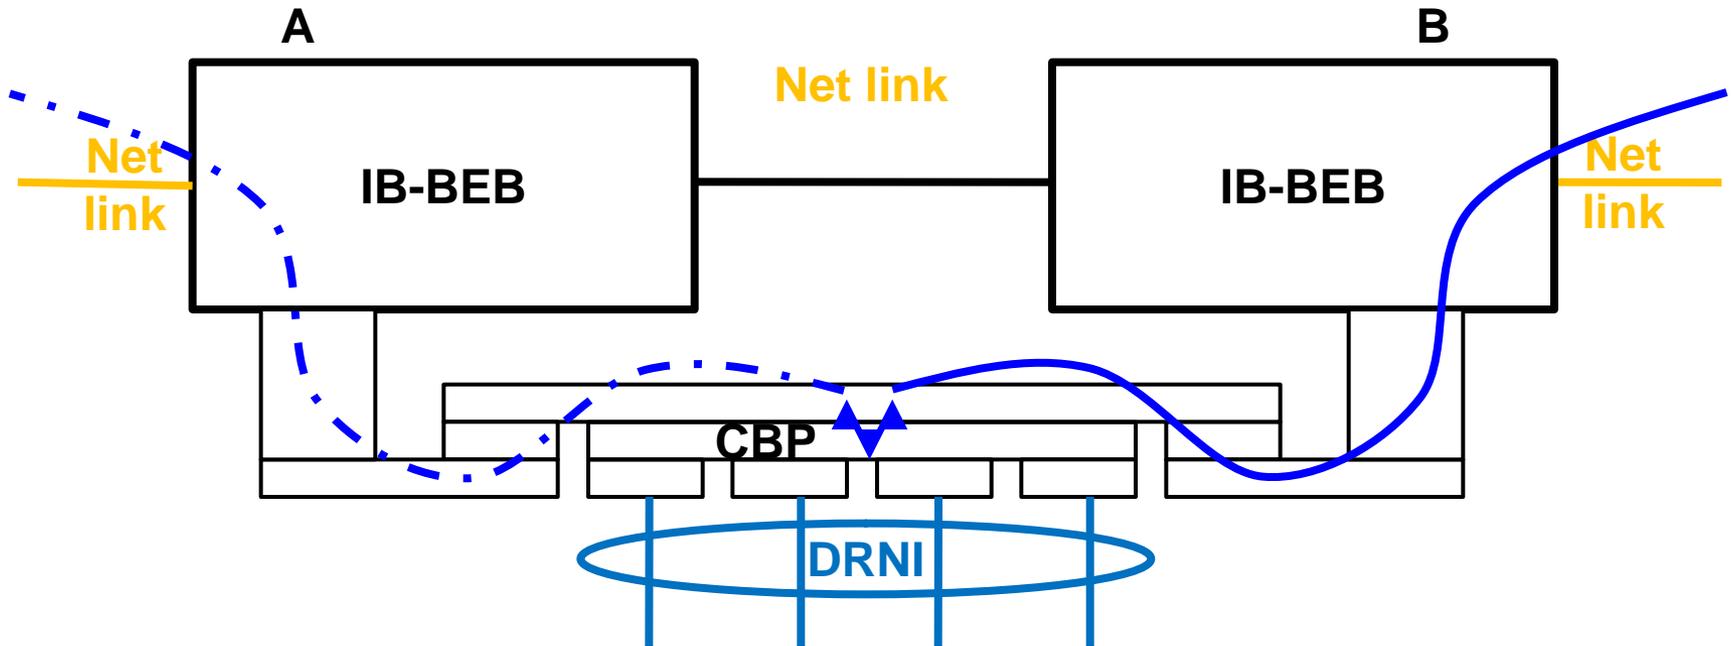
Norman Finn

nfinn@cisco.com

What is the IB-BEB DRNI problem?

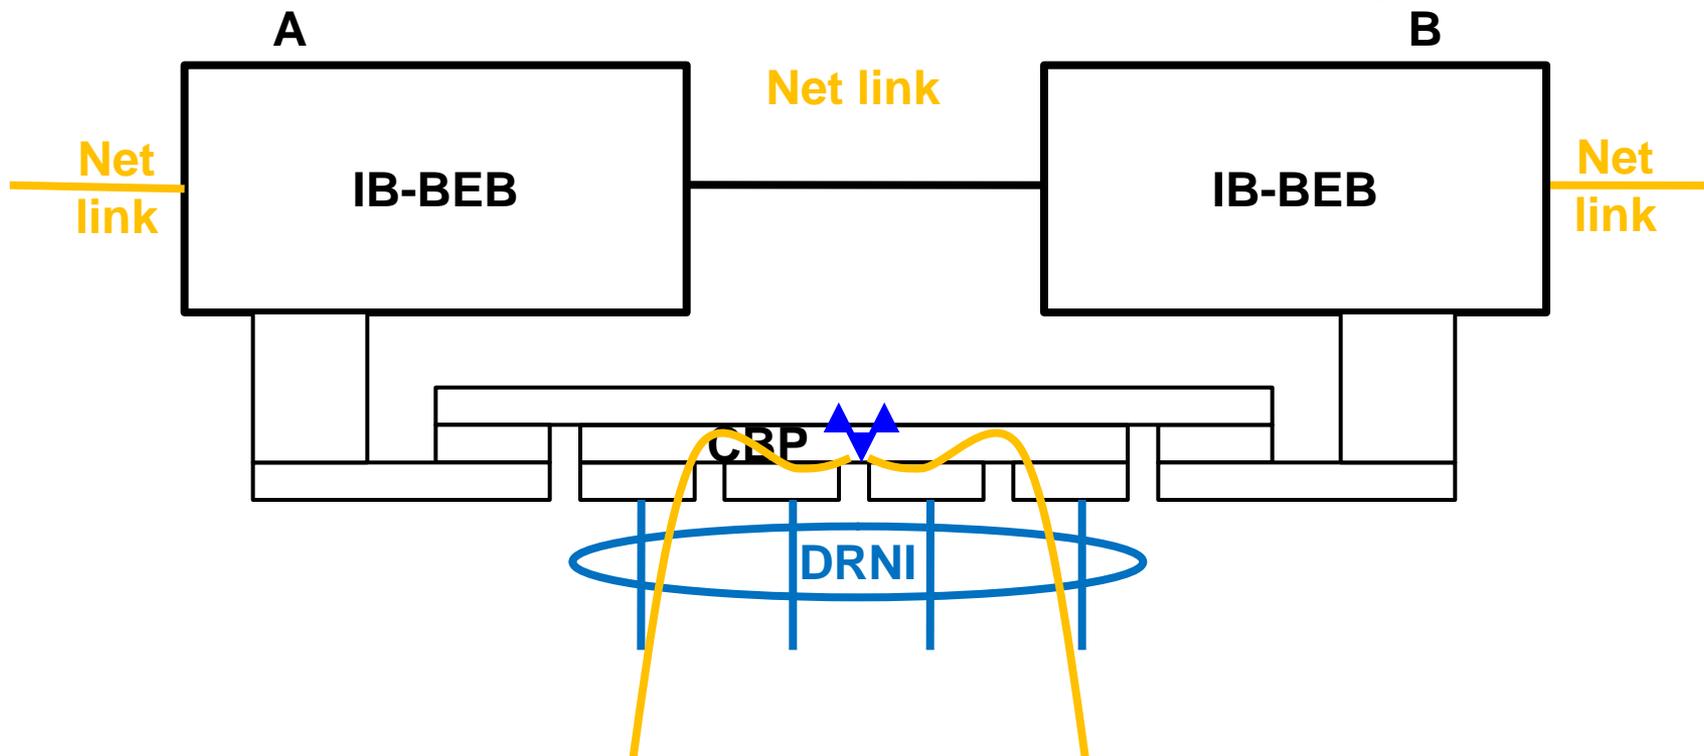
- In protection switching, all of the alternate network paths need to terminate in a CBP in the virtual switch at which the DRNI Link Aggregation Group terminates.
- You need all alternate network paths to split/merge at one CBP in one physical Node in the Portal so that that one Node can observe the states of all of the paths and make a rapid selection.
- The above means that **the Inter-DAS Link must be B-VLAN tagged.**
- But, the S-VLANs bundled in an I-SID may be spread across the Link Aggregation Group links, so to get them all to the PIP connected to that CBP, **the Inter-DAS Link must be S-VLAN tagged.**

Intra-DAS Link needs B-VLAN tag



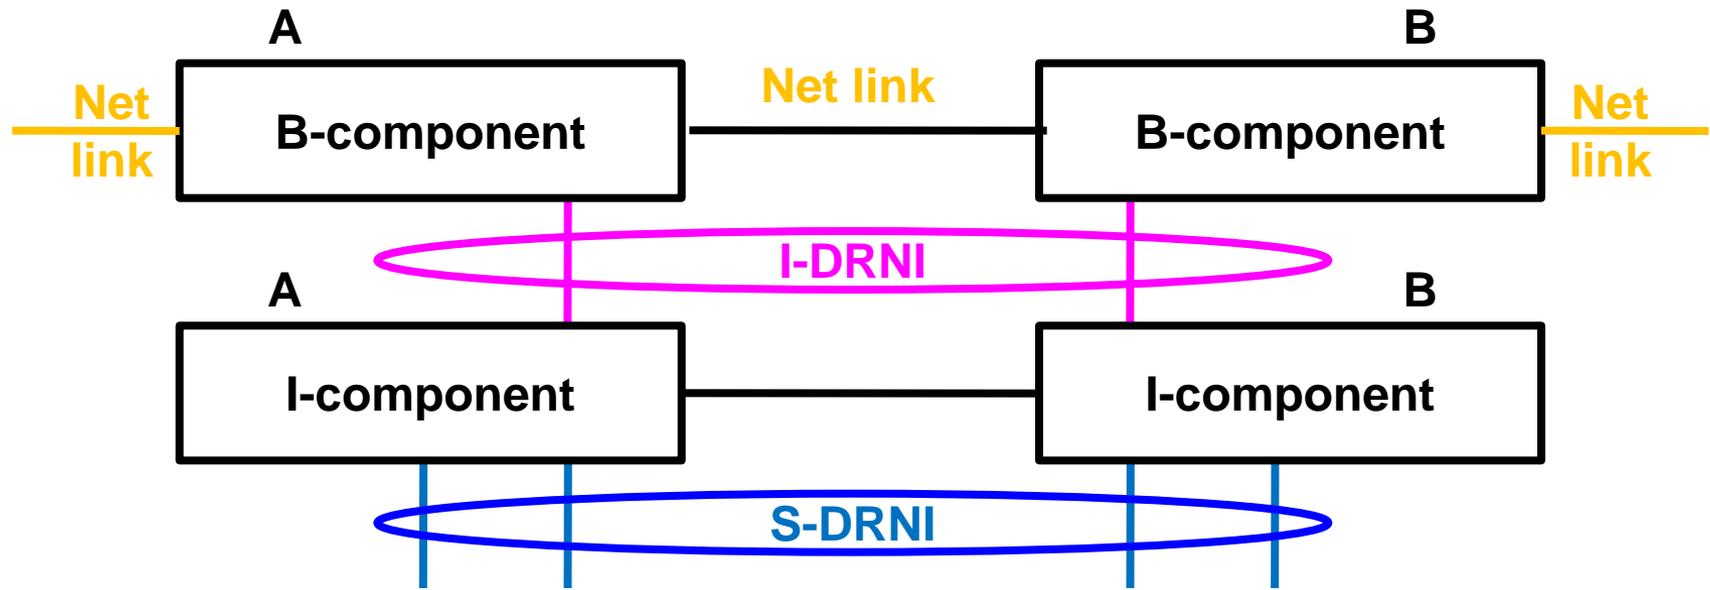
- To have a single CBP at which all network paths terminate, the Intra-DAS Link needs to be **tagged with a B-VLAN**.

Intra-DAS Link needs S-VLAN tag



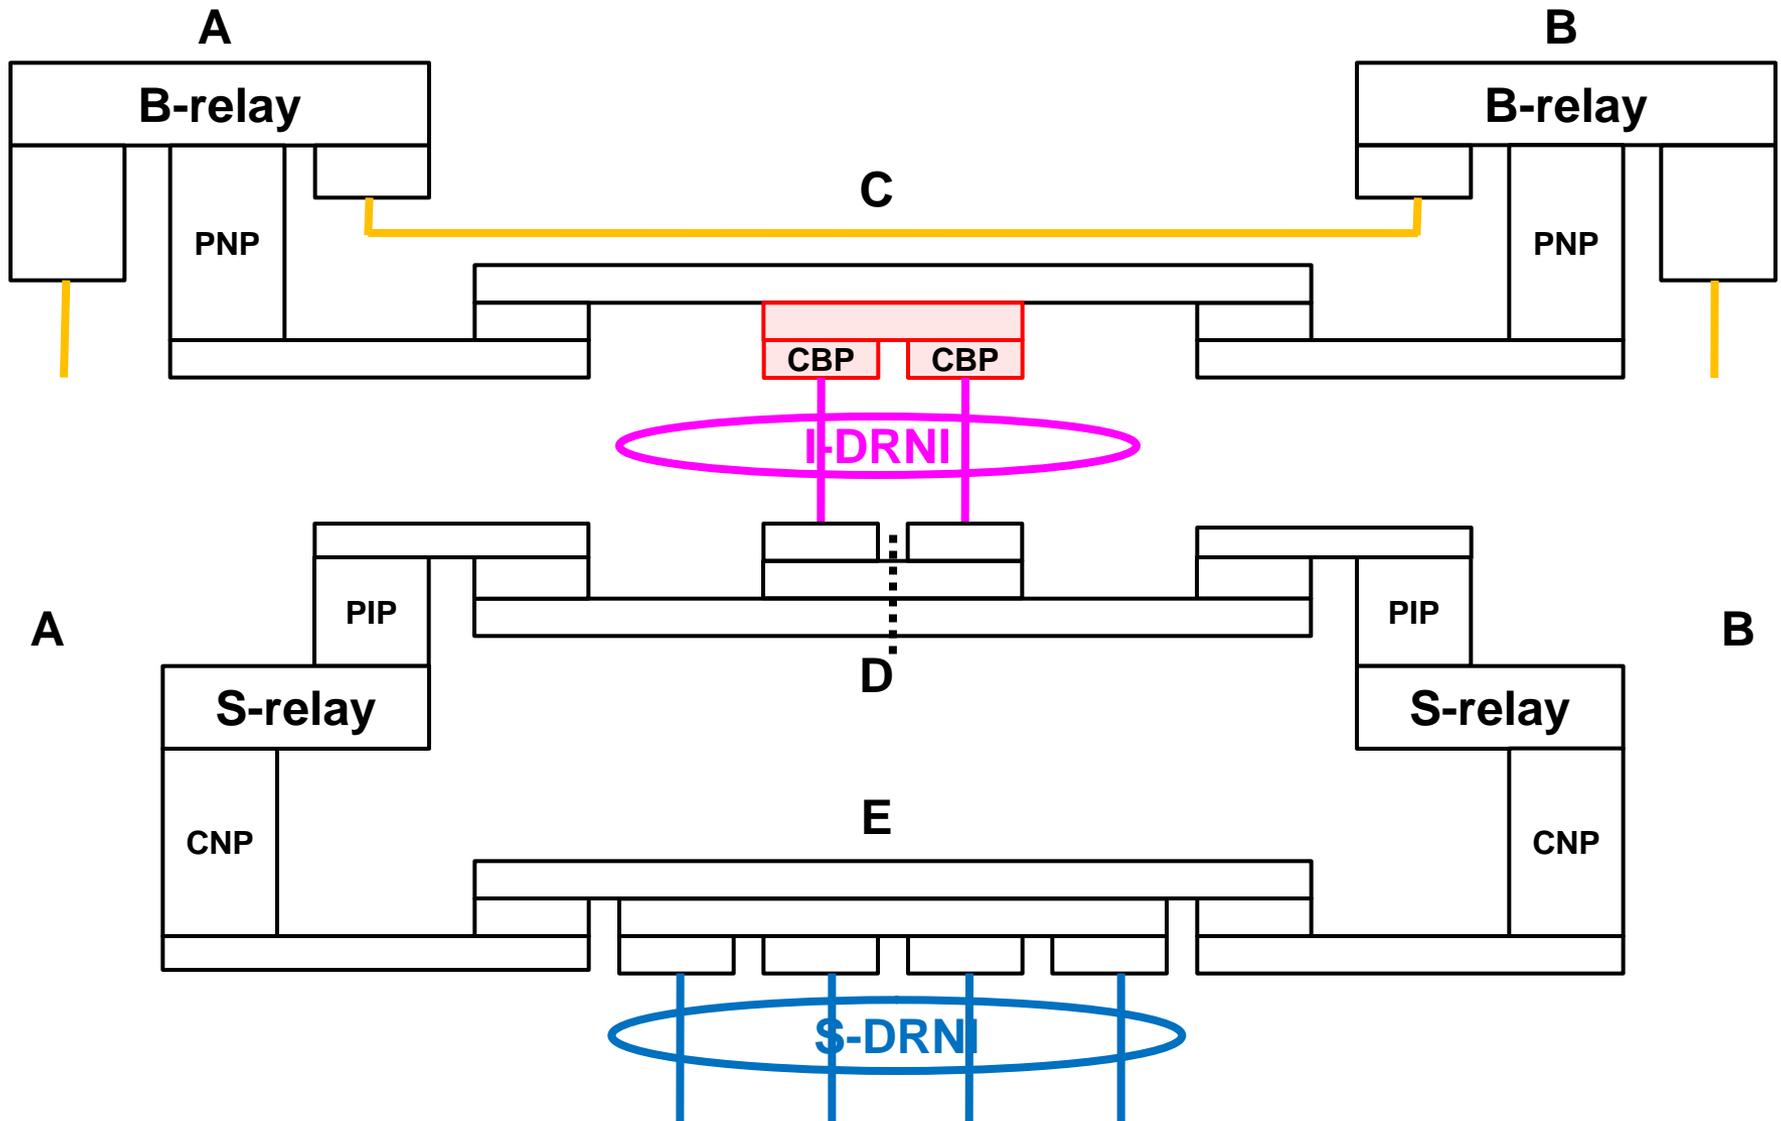
- To distribute the S-tagged VLANs collected into a single I-tagged service to their respective Link Aggregation Group links, the Intra-DAS Link needs to be **tagged with an S-VLAN**.

IB-BEB Component view



- An additional DRNI between the B-components and the I-components provides an extra Intra-DAS Link so that one can be S-tagged and the other B-tagged.

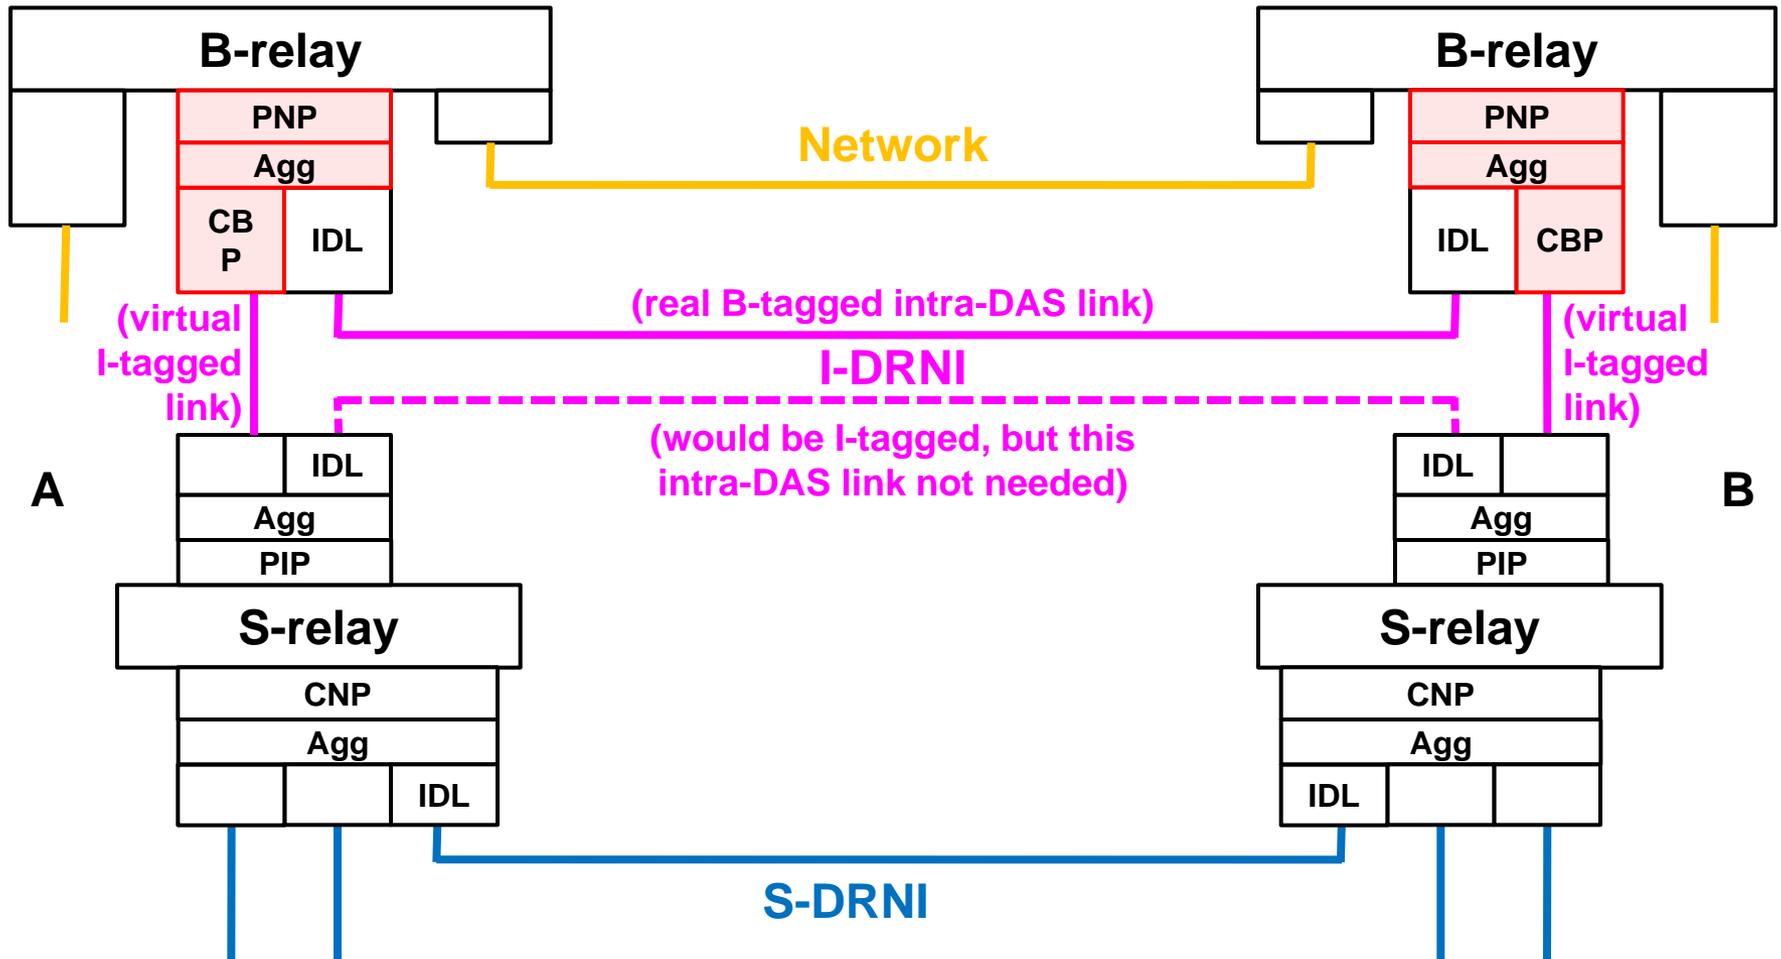
IB-BEB DRNI: Logical view



I-DRNI

- I-DRNI is an I-tagged interface, but the distributed relay in virtual bridge (**C**) operates on B-tags, so the intra-DAS link of virtual bridge (**C**) is B-tagged..
- Note the position of the CBP – it is **below** Link Aggregation! Virtual relay C operates on B-tags, and virtual relay D is a no-op, because it aggregates only one link on each side, and does not need an Intra-DAS Link.
- This position is OK, because there is exactly one “physical” link in the I-DRNI on each Node, so the inversion does not change the number or identity of Bridge Ports in the B-components.

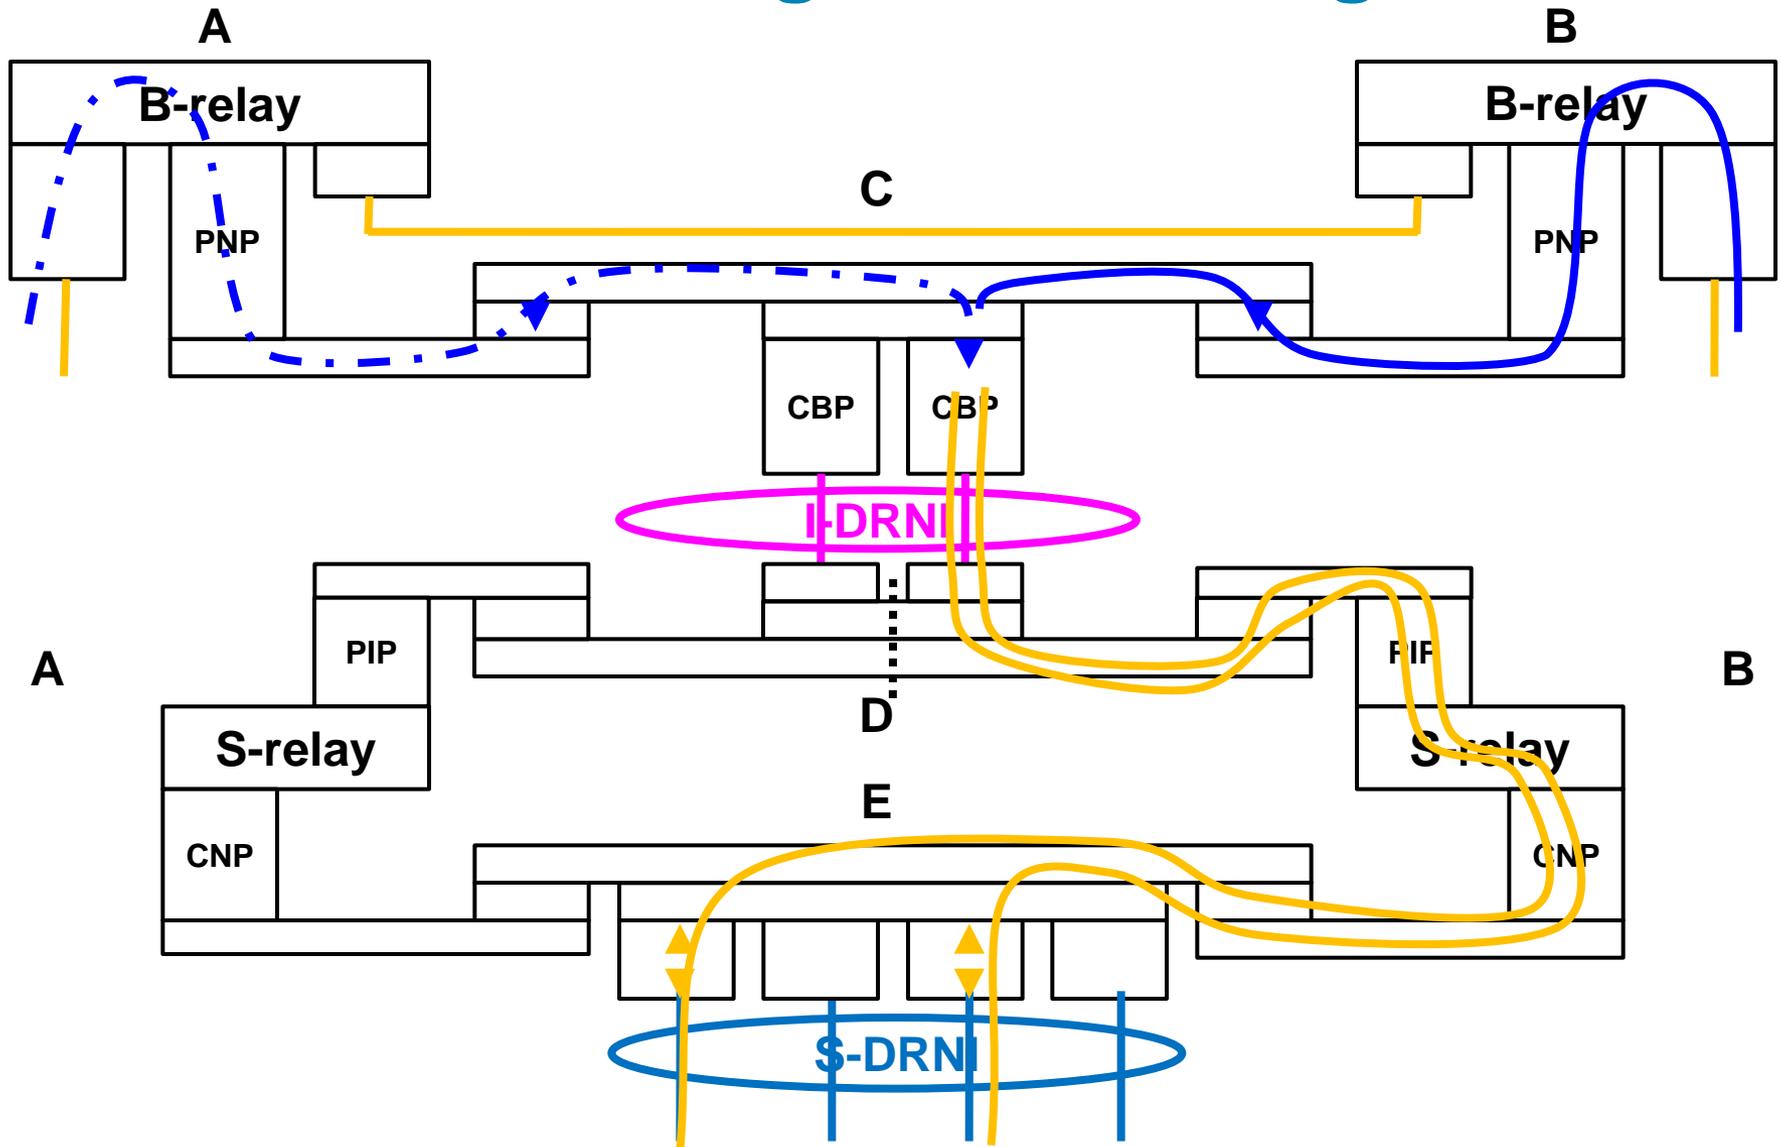
IB-BEB DRNI: Bridge view



I-DRNI

- If you want to have multiple CBPs on one S-DRNI, then you must have a separate I-DRNI (with a separate IDL) per CBP.

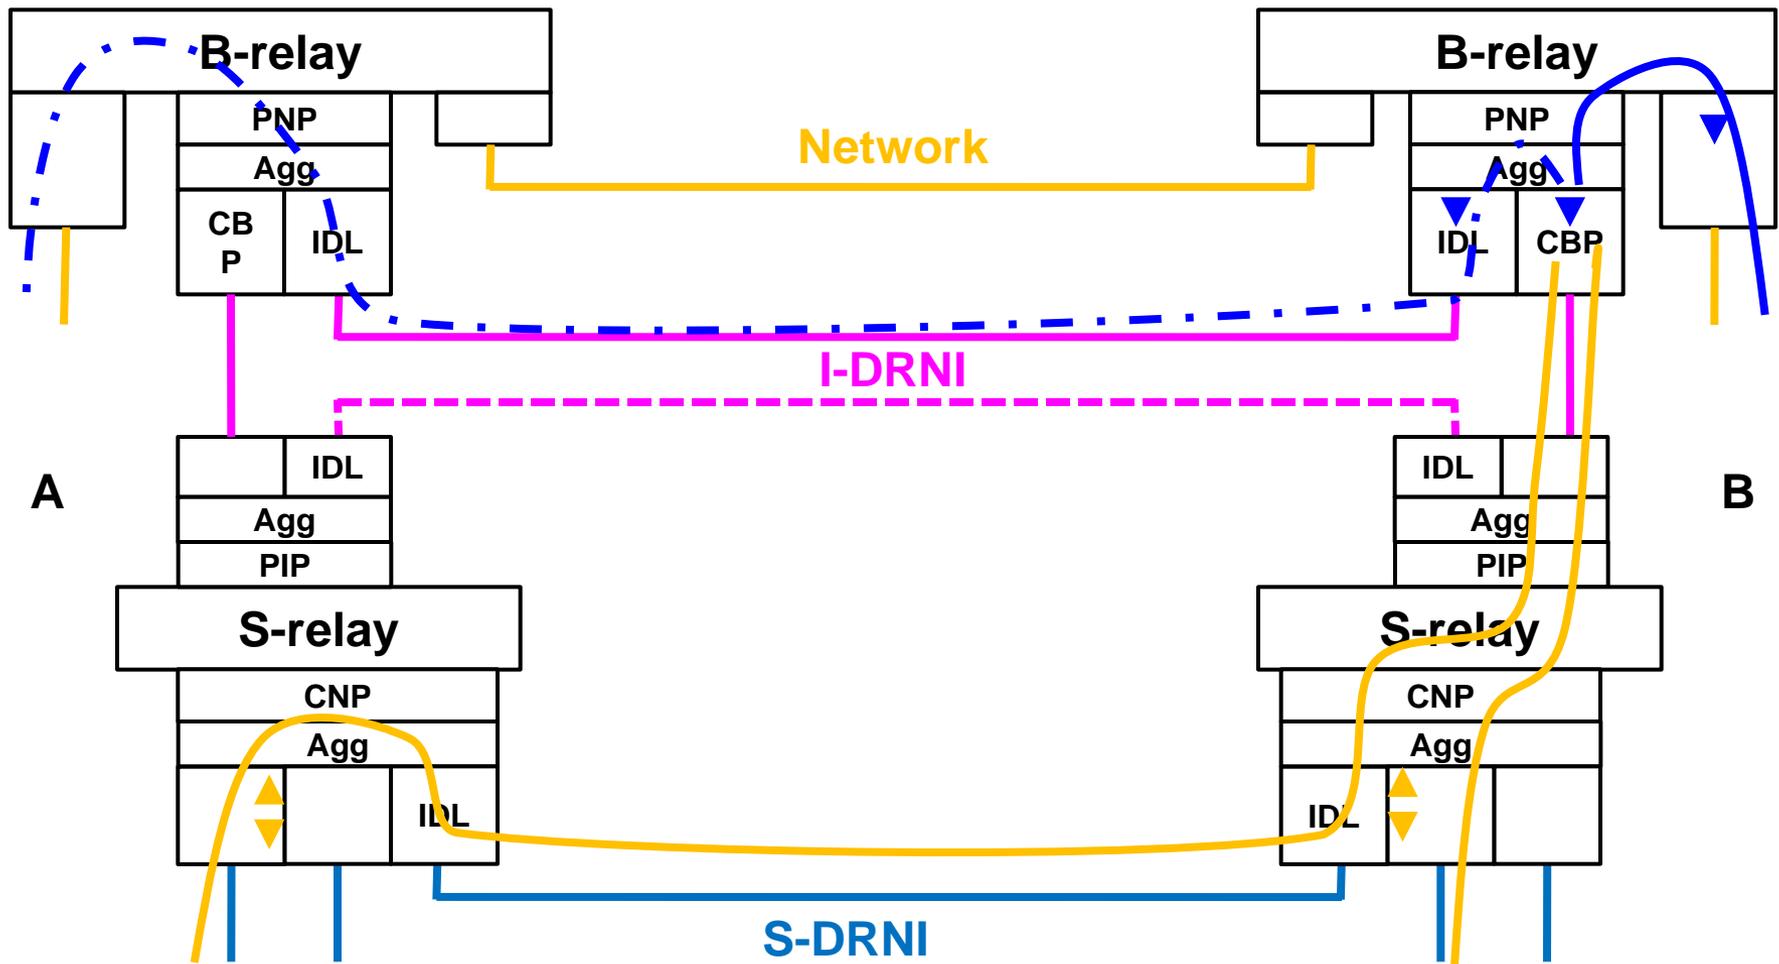
IB-BEB DRNI: Logical view, Seg. Prot.



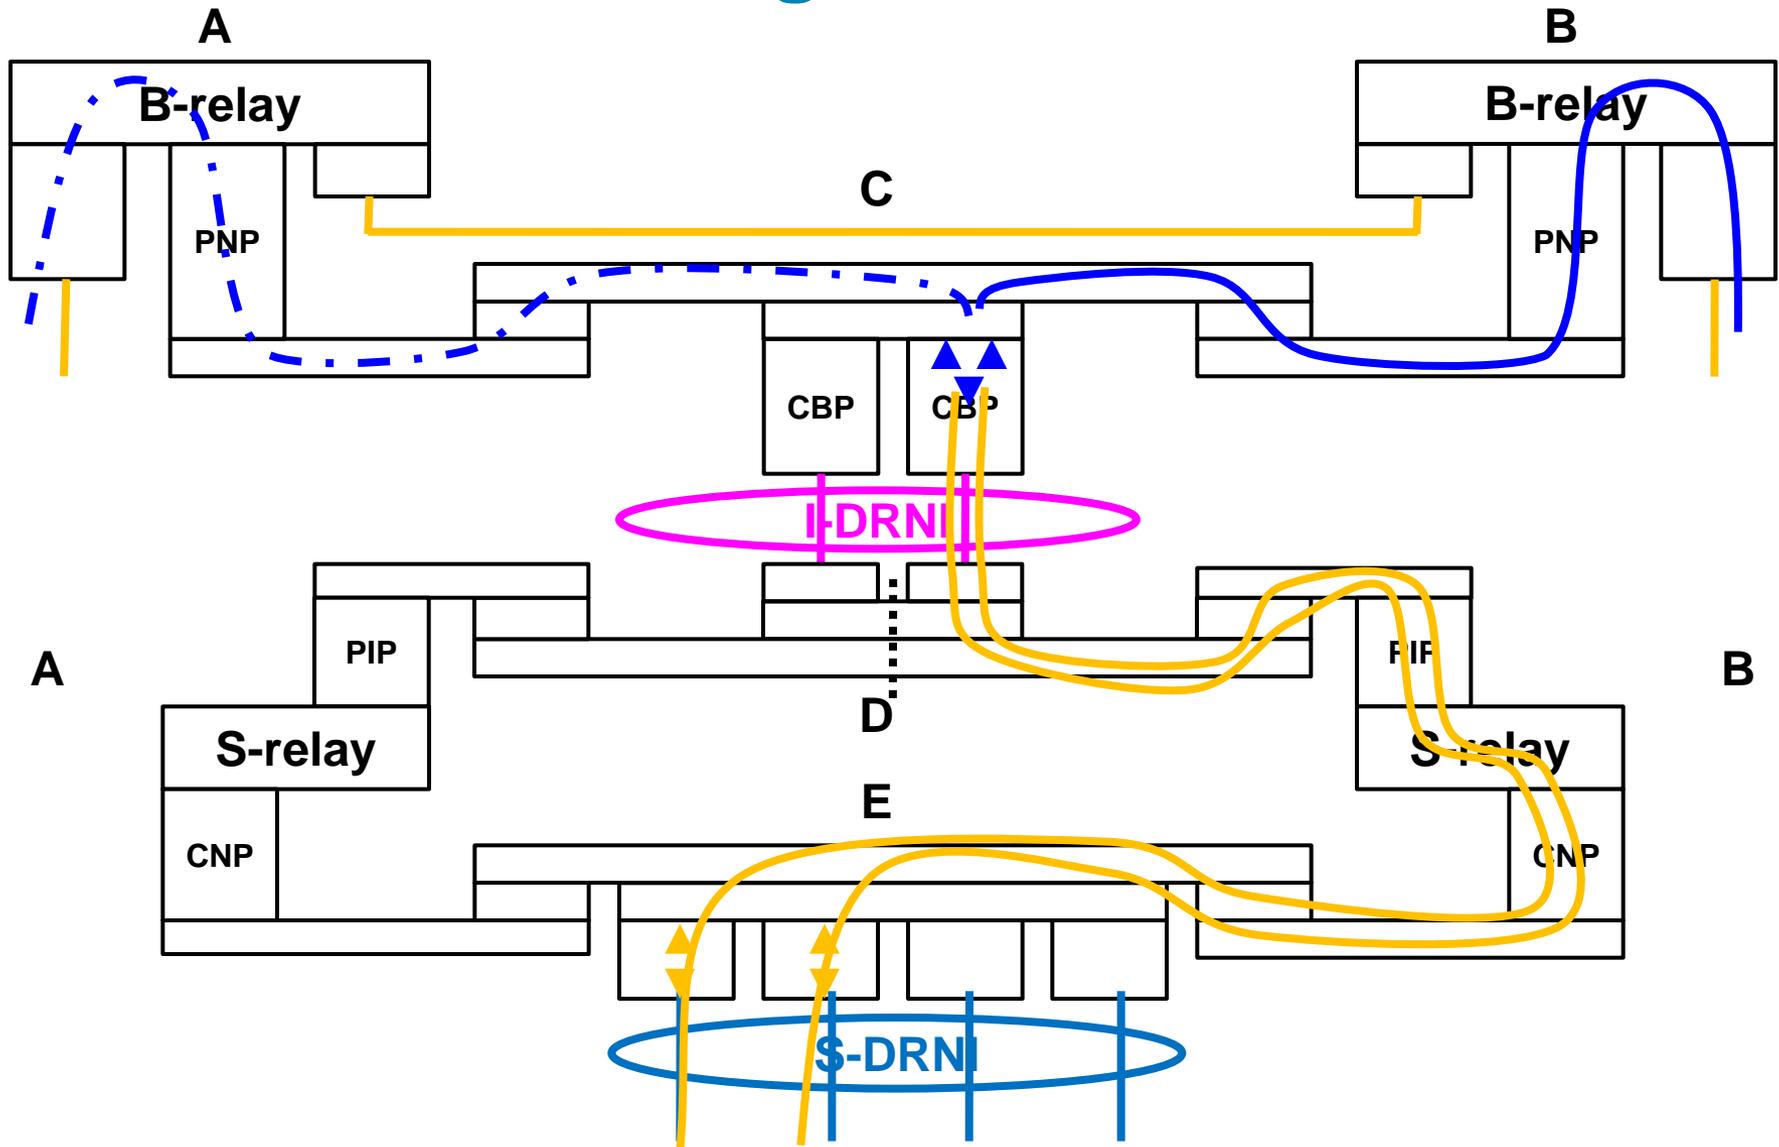
S-DRNI

- Note that the per-S-VLAN MEPs/MIPs have been placed at the very edge of the network. This satisfies Maarten's requirement.
- Placing the MEPs on the physical ports of the Link Aggregation Group requires a pair of S-VLAN splitters/muxes to serve an array of MEPs and MIPs on a “physical” port, below Link Aggregation.
- We could also place the S-VLAN MEPs in the CBP.

IB-BEB DRNI: Bridge view: Seg. Prot.



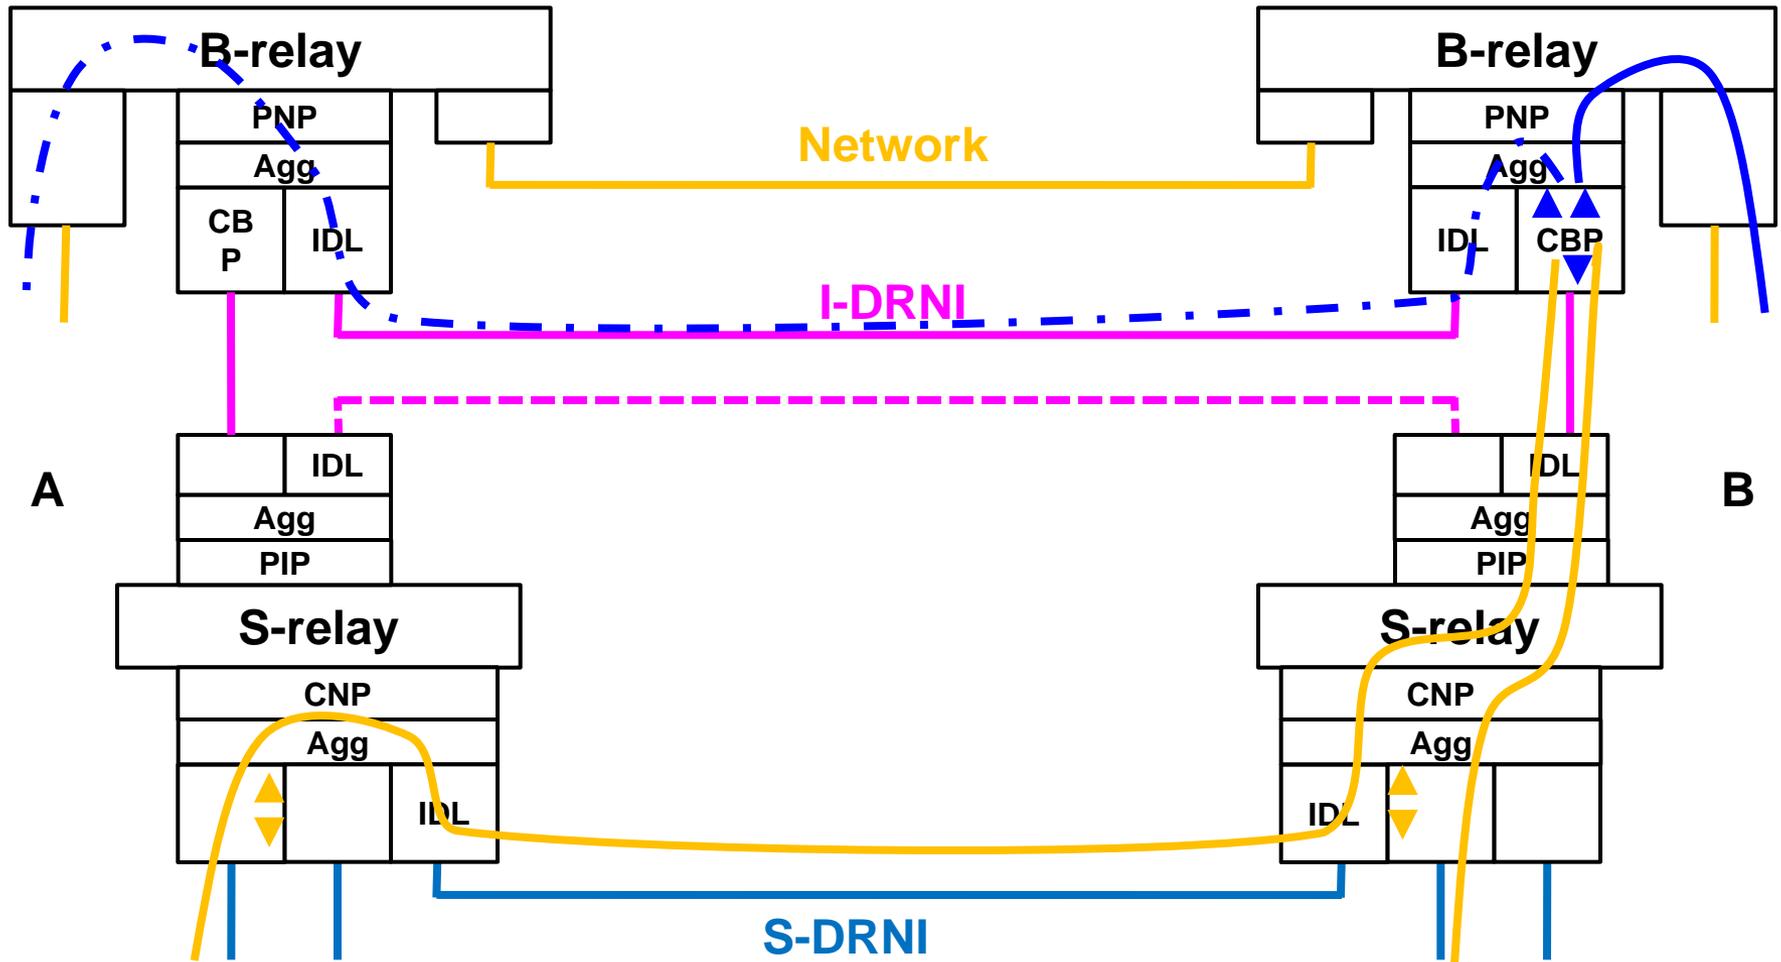
IB-BEB DRNI: Logical view, TESI Prot.



Virtual DRNI

- The virtual DRNI has only one vertical link per physical device.
- You do not need two intra-DAS links (encapsulations, really) for virtual bridges **C** and **D**. They are not needed for redundancy, since if virtual bridge **C**'s intra-DAS link fails, so does **D**'s. They are not needed for load sharing, because they are physically the same and belong to the same provider.
- The upper intra-DAS link in the virtual DRNI resolves mismatches between distribution by B-VLAN and distribution by service.

IB-BEB DRNI: Bridge view: TESI Prot.



Solves everyone's concerns

- One Node, the active one, has all the MEPs it needs to do protection switching.
- The two paths to the network do not make use of the network link(s), if any, between the Portal's Nodes.
- The S-DRNI's intra-DAS link is part of the network, not part of the DRNI, as far as the per-service MEPs/MIPs are concerned.

Gateway changes and shadow MEPs

- You can still change gateways. Assuming that the “shadow” CBP in the other Node has the same MAC address as the primary CBP, no learning event takes place when you switch gateways.
- The shadow CBP’s MEPs can have the same MAC addresses as the primary CBP’s MEPs, because the MEPs are silent as long as the Gateway is in the other Node.

Reasonable question:

- The I-DRNI's vertical links are virtual. Only one of the intra-DAS links are utilized, so the Frame Distribution and Frame Collection of the up-facing "Portal" of the I-DRNI are trivial (one input, one output). **Why bother to call it a DRNI?**
- Because the upper Portal of the I-DRNI does exactly what Frame Collectors and Frame Distributors do.
- And, because the I-DRNI's upper intra-DAS link definitely does not belong to the network, so it is difficult to make it a separate port on the B-component.

Remaining concerns

- This adds a second IDL, which increases the burden of encapsulation / tag-sharing.
- But, assuming that the Gateway is on the side where the primary path goes out the network link, and not the I-DRNI IDL, no data flows over the I-DRNI IDL; only CFM frames pass over the I-DRNI IDL.
- We have to make sure that the **backup Gateway** is known to be working and ready in case the primary Gateway fails. It is not yet clear to me what that entails.