

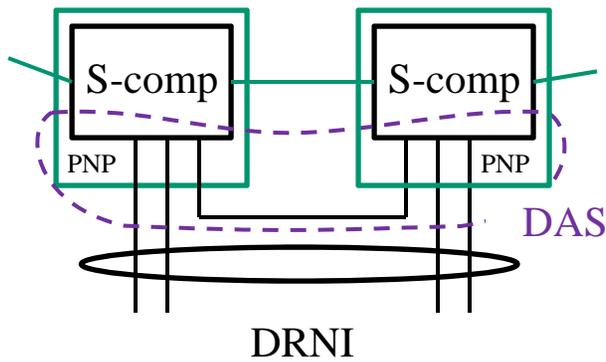
Multiple DRNI Support

Version 01

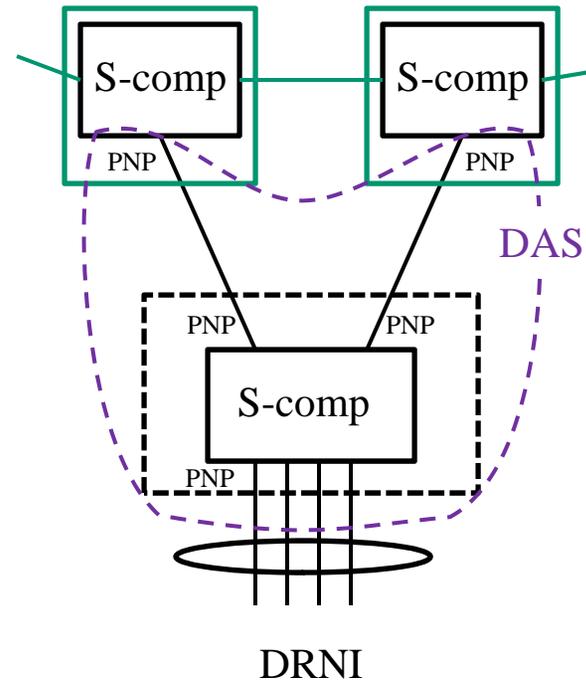
Stephen Haddock

October 20, 2011

Physical and Logical Topology of S-tagged DRNI

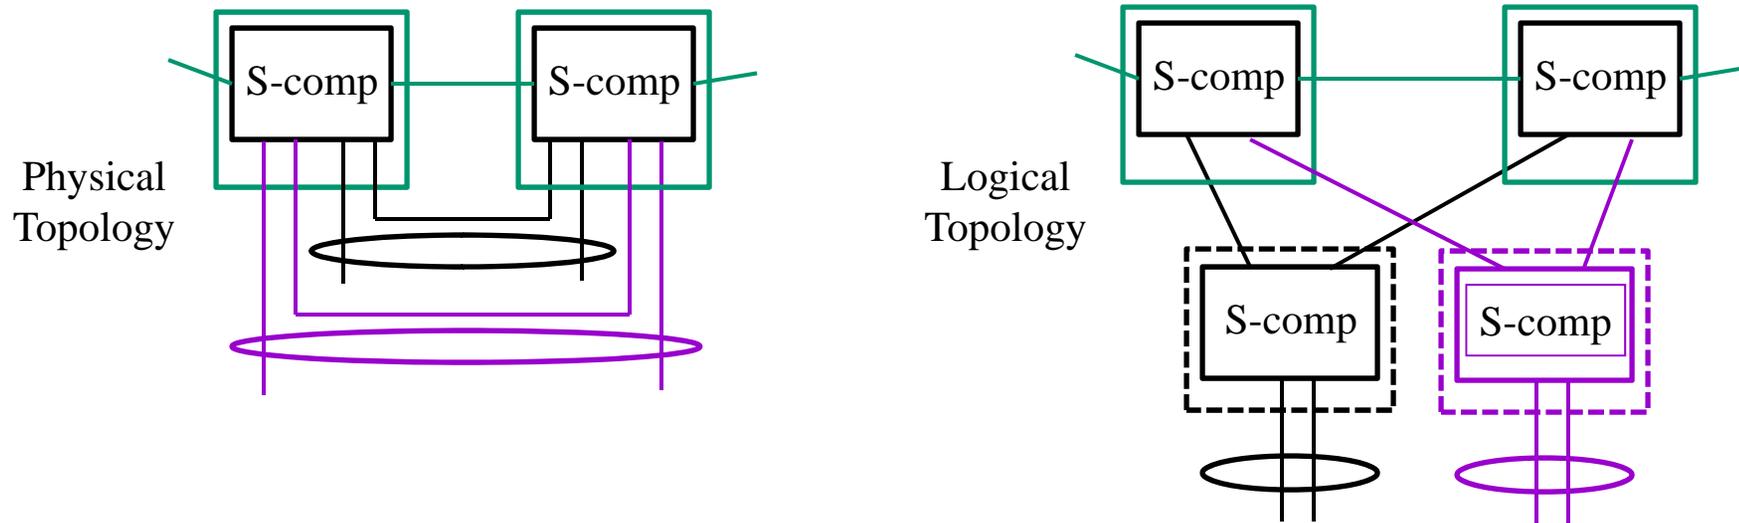


Physical
Topology



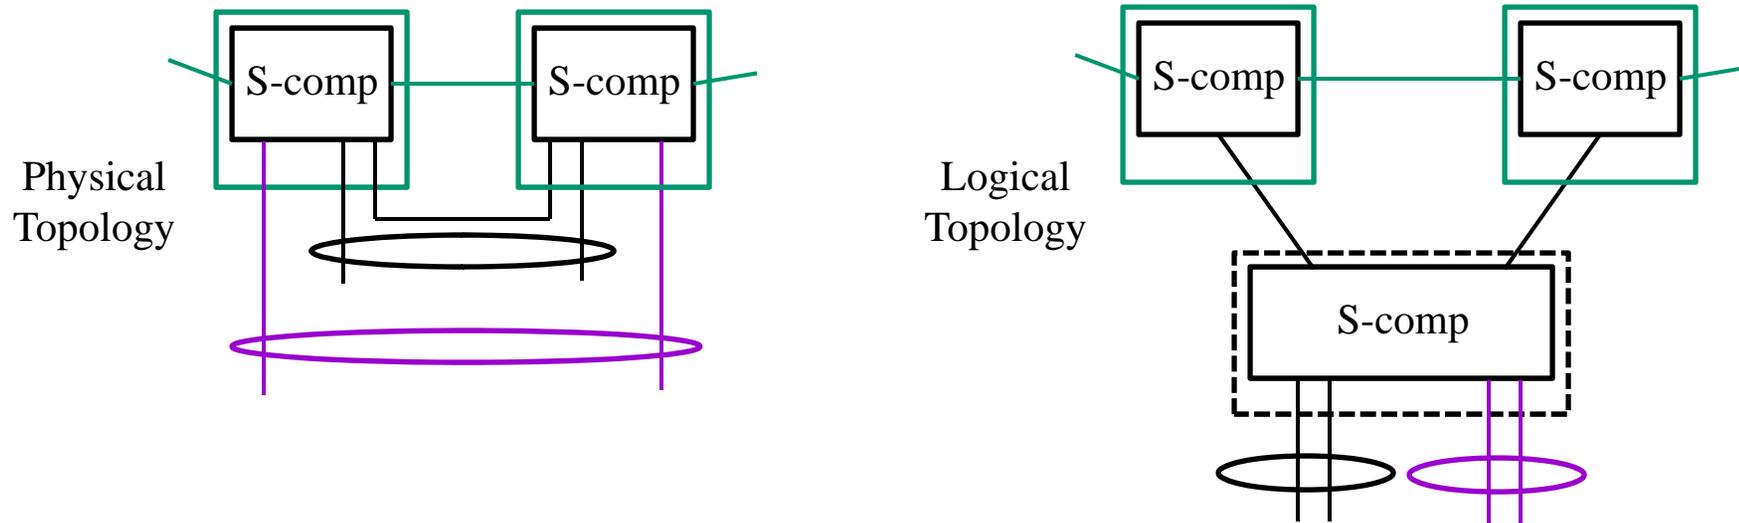
Logical
Topology

Possible Models of Multiple DRNIs



- **Option 1: Create logical component per DRNI**
 - Advantages:
 - Keeps all DRNI aggregation layers separate.
 - Disadvantages:
 - Requires Intra-DAS Link per DRNI.
 - Separate physical dedicated Intra-DAS Links, or use encapsulation to create separate virtual Intra-DAS Links on a single physical link.
 - Requires separate instances of ‘Distributed Higher Layers’ per DRNI.

Possible Models of Multiple DRNIs

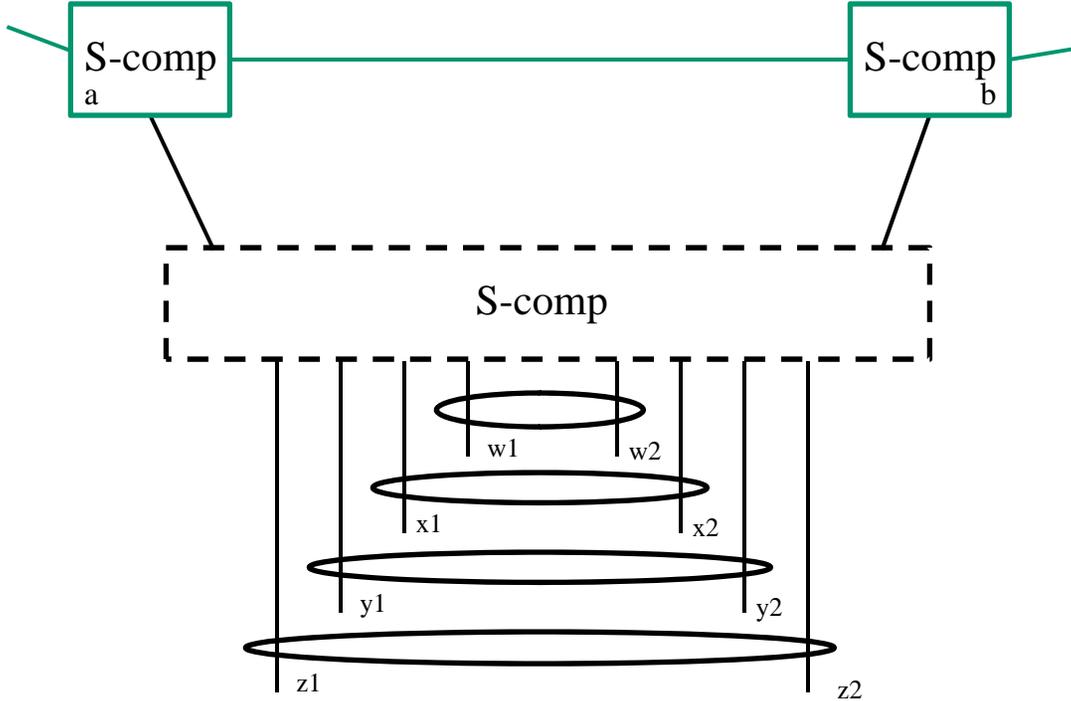


- **Option 2: Single logical component for multiple DRNIs**
 - Disadvantages
 - Potentially intertwines some aspects of aggregation layers for different DRNIs.
 - Advantages
 - Same Intra-DAS Link supports all DRNIs.
 - One (or possibly two) physical links used for the Intra-DAS Links of all DRNIs without requiring encapsulation.
 - Single instance of ‘Distributed Higher Layers’ for all DRNIs.

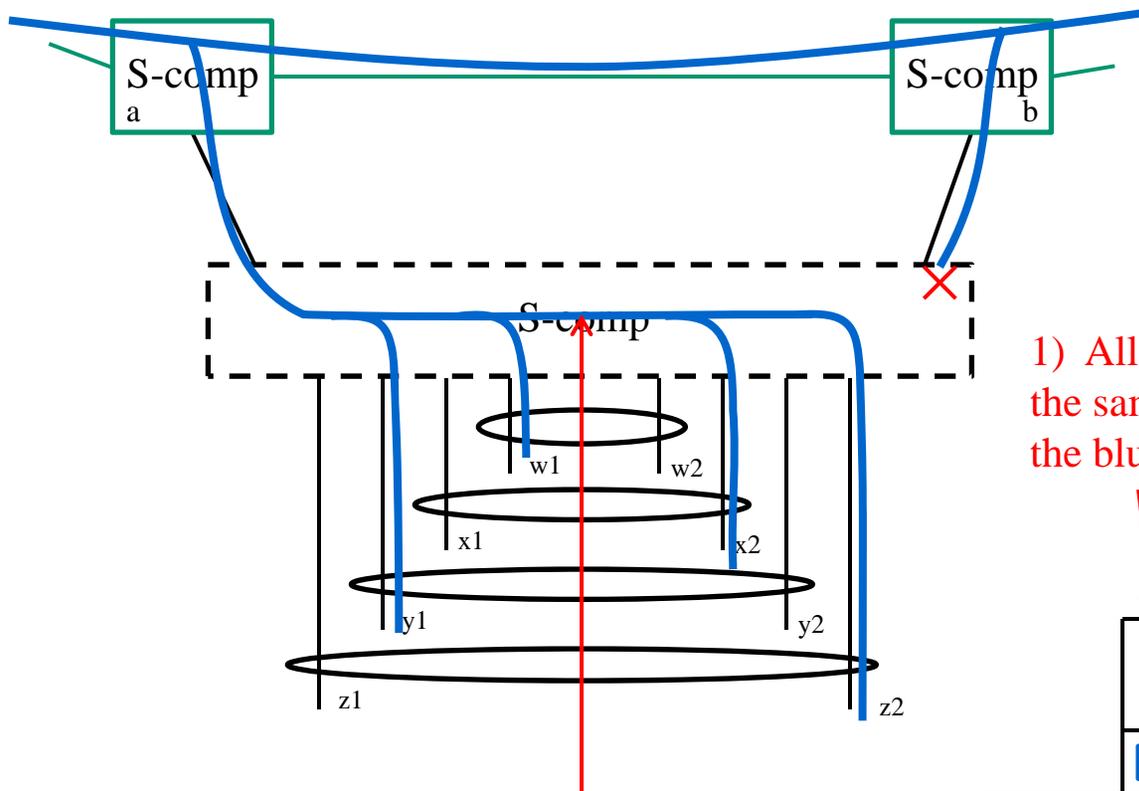
Should we consider multiple DRNIs?

- Norm has suggested that it shouldn't be in the first version:
From final slide of 'new-nfinn-why-LACP-for-NNI-0111-v01':
 - The links connecting the physical nodes of a Portal (intra-Portal links) are very simple if they support a single RNI.
 - These links get much more complex, both in the data plane and in the control plane, if a single intra-Portal link supports multiple RNIs.
 - Therefore, standardization of support by a single intra-Portal link of multiple RNIs should be deferred.
- I think it is worth some consideration:
 - May not be as complex as feared.
 - Even if inclusion in the standard is deferred, the current standard should try to construct a DRNI model that allows multiple DRNI support to be included in the future.

Single Logical Component with 4 DRNIs



Single Logical Component with 4 DRNIs



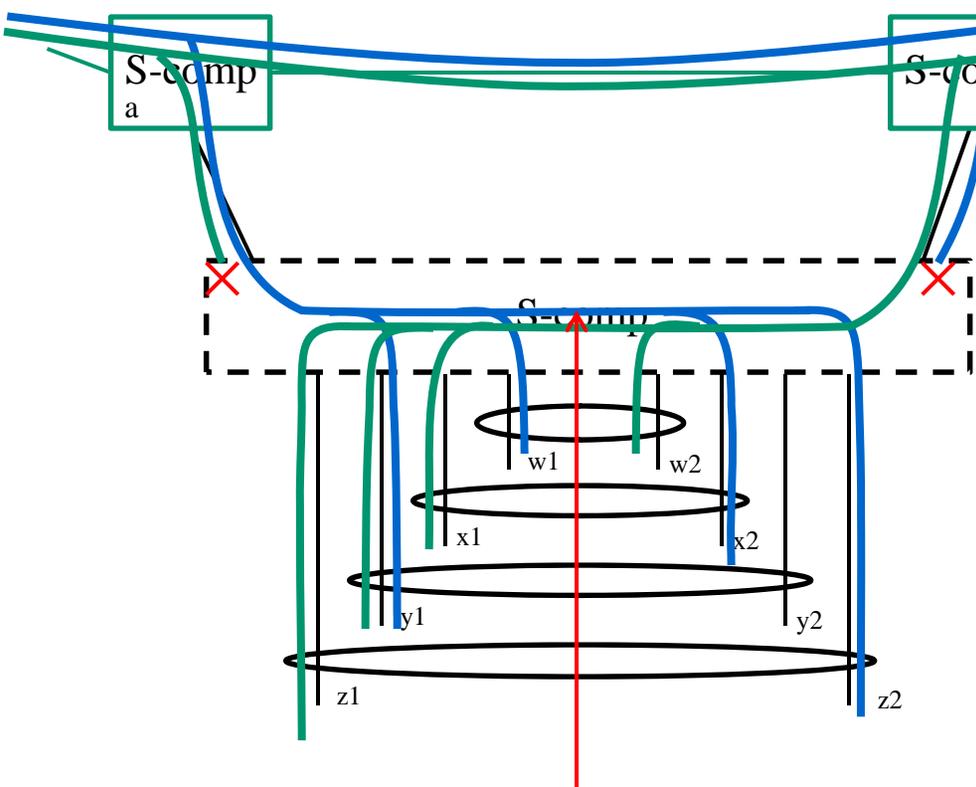
For this example the Link selector is the same as the Gateway selector (the VID), though this isn't always the case.

1) All DRNIs select the same Gateway for the blue VLAN

		Gateway VID	Link SID
		■	■
D	w	a	w1
R	x	a	x2
N	y	a	y1
I	z	a	z2

2) When split logical component between two physical devices, all blue VLAN frames on Intra-DAS link are either moving toward or away from the Gateway. This is sufficient to determine how to forward each frame.

Single Logical Component with 4 DRNIs



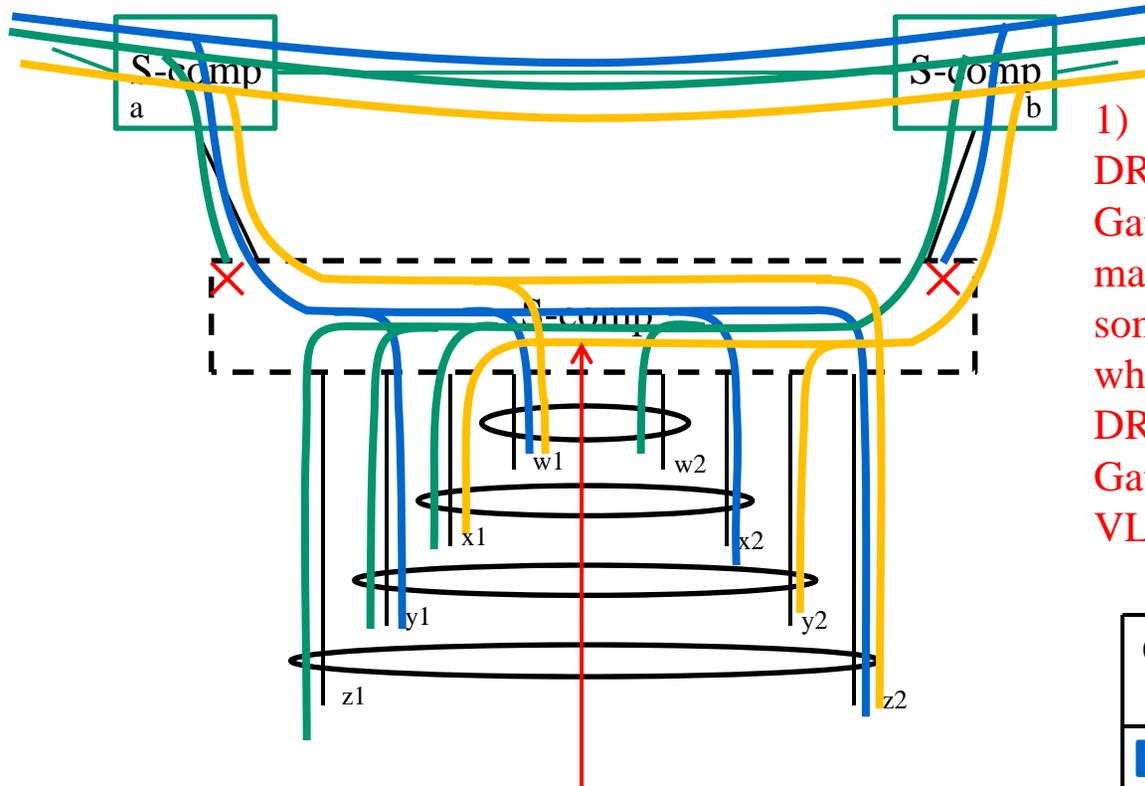
1) All DRNIs select the other Gateway for the green VLAN



		Gateway VID		Link SID	
		Blue	Green	Blue	Green
D	w	a	b	w1	w2
	x	a	b	x2	x1
R	y	a	b	y1	y1
	z	a	b	z2	z1

2) When split logical component between two physical devices, all frames on Intra-DAS link are either moving toward or away from their respective Gateways. This, together with the VID, is still sufficient to determine how to forward each frame.

Single Logical Component with 4 DRNIs

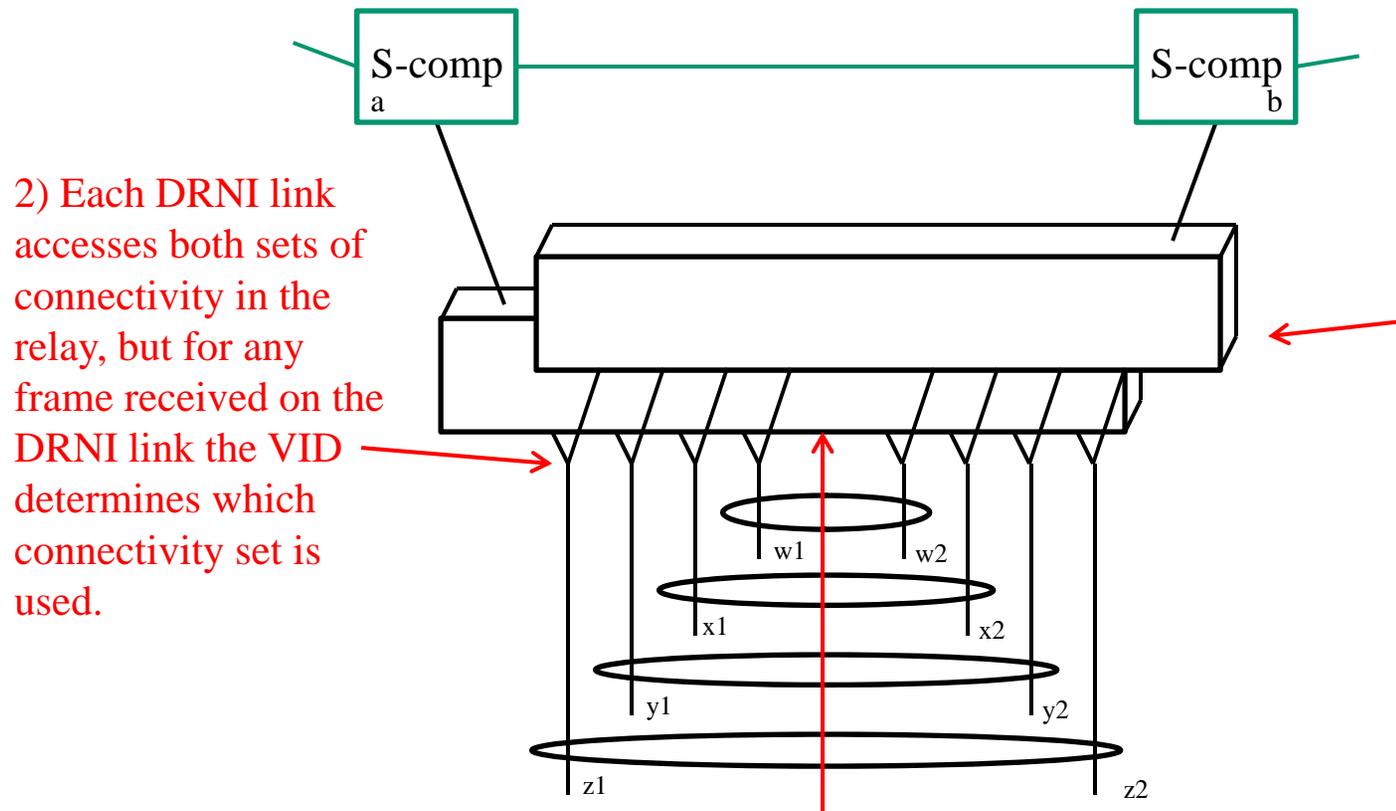


1) The constraint that all DRNIs select the same Gateway for a given VLAN may be overly restrictive for some networks. Consider what happens if different DRNIs select different Gateways for the orange VLAN.

2) When split logical component between two physical devices, if all orange VLAN frames use the same Intra-DAS link there is not enough information available on the receiving end to determine how to forward each frame. Need to know whether frame going from b to a is going toward Gateway a or away from Gateway b.

		Gateway VID			Link SID		
		Blue	Green	Yellow	Blue	Green	Yellow
D	w	a	b	a	w1	w2	w1
R	x	a	b	b	x2	x1	x1
N	y	a	b	b	y1	y1	y2
I	z	a	b	a	z2	z1	z2

Single Logical Component with 4 DRNIs

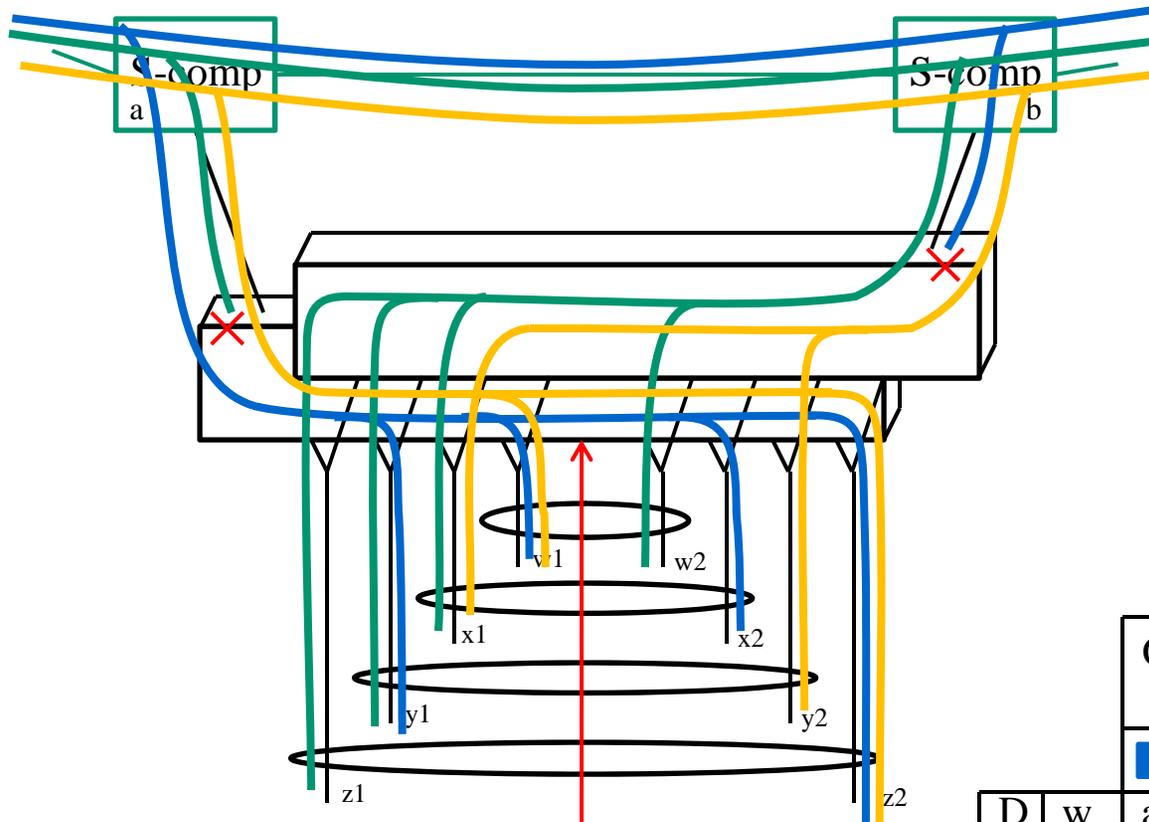


2) Each DRNI link accesses both sets of connectivity in the relay, but for any frame received on the DRNI link the VID determines which connectivity set is used.

1) Observe that in the filtering database, the connectivity of frames using a for the Gateway never overlaps the connectivity of frames using b for the Gateway. Conceptually the gateway selection divides the relay into two disjoint sets of connectivity.

3) When split logical component between two physical devices, use two Intra-DAS links. Now on each link frames are either moving toward or away from the Gateway. This, together with the VID, is sufficient to determine how to forward each frame.

Single Logical Component with 4 DRNIs

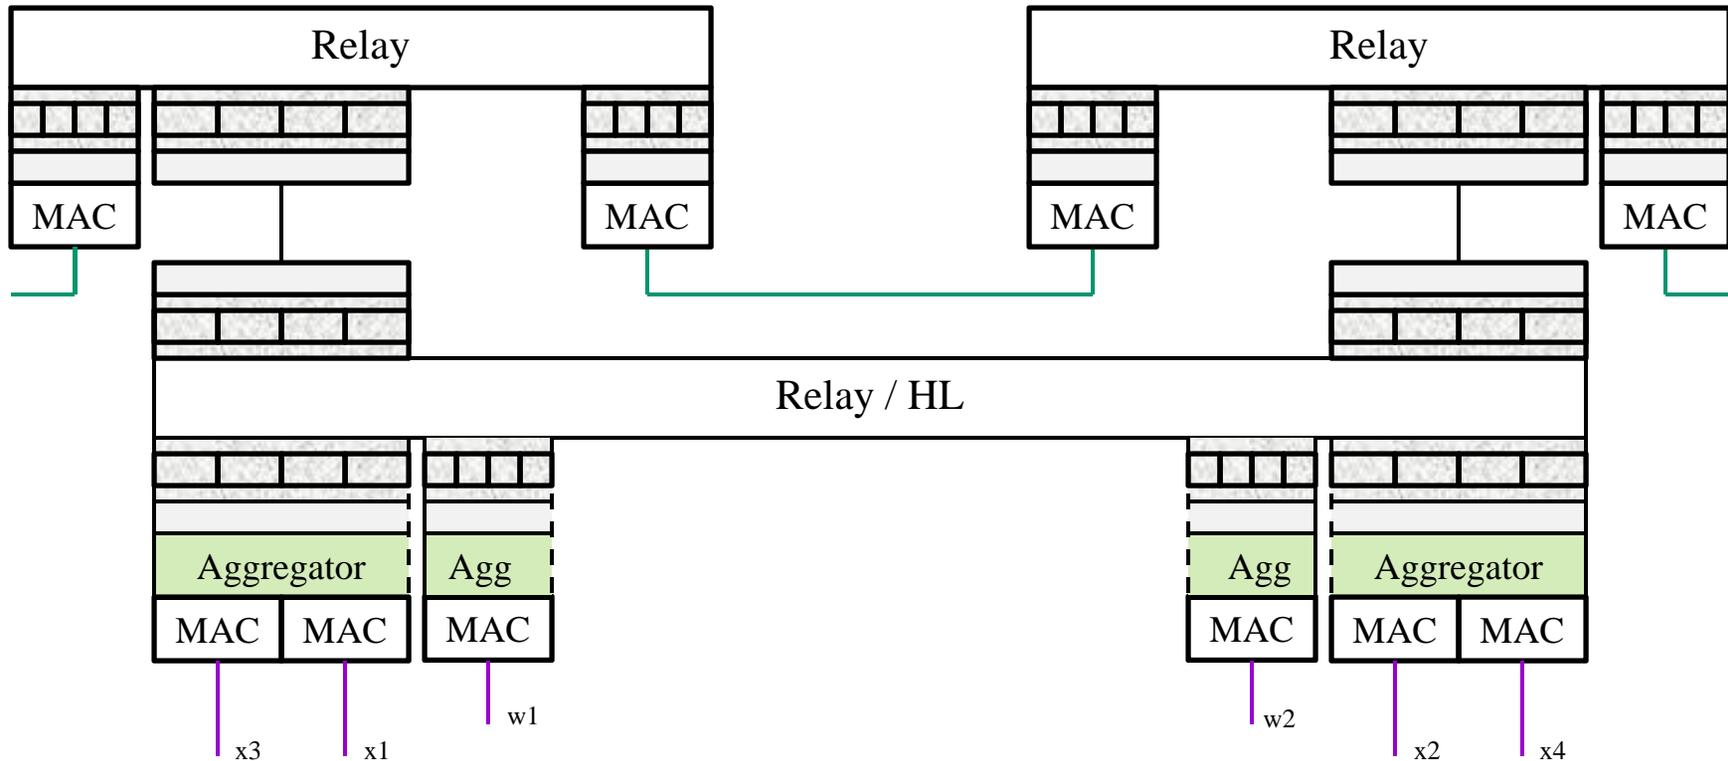


3) When split logical component between two physical devices, use two Intra-DAS links. Now on each link frames are either moving toward or away from the Gateway. This, together with the VID, is sufficient to determine how to forward each frame.

		Gateway VID			Link SID		
		Blue	Green	Yellow	Blue	Green	Yellow
D	w	a	b	a	w1	w2	w1
R	x	a	b	b	x2	x1	x1
N	y	a	b	b	y1	y1	y2
I	z	a	b	a	z2	z1	z2

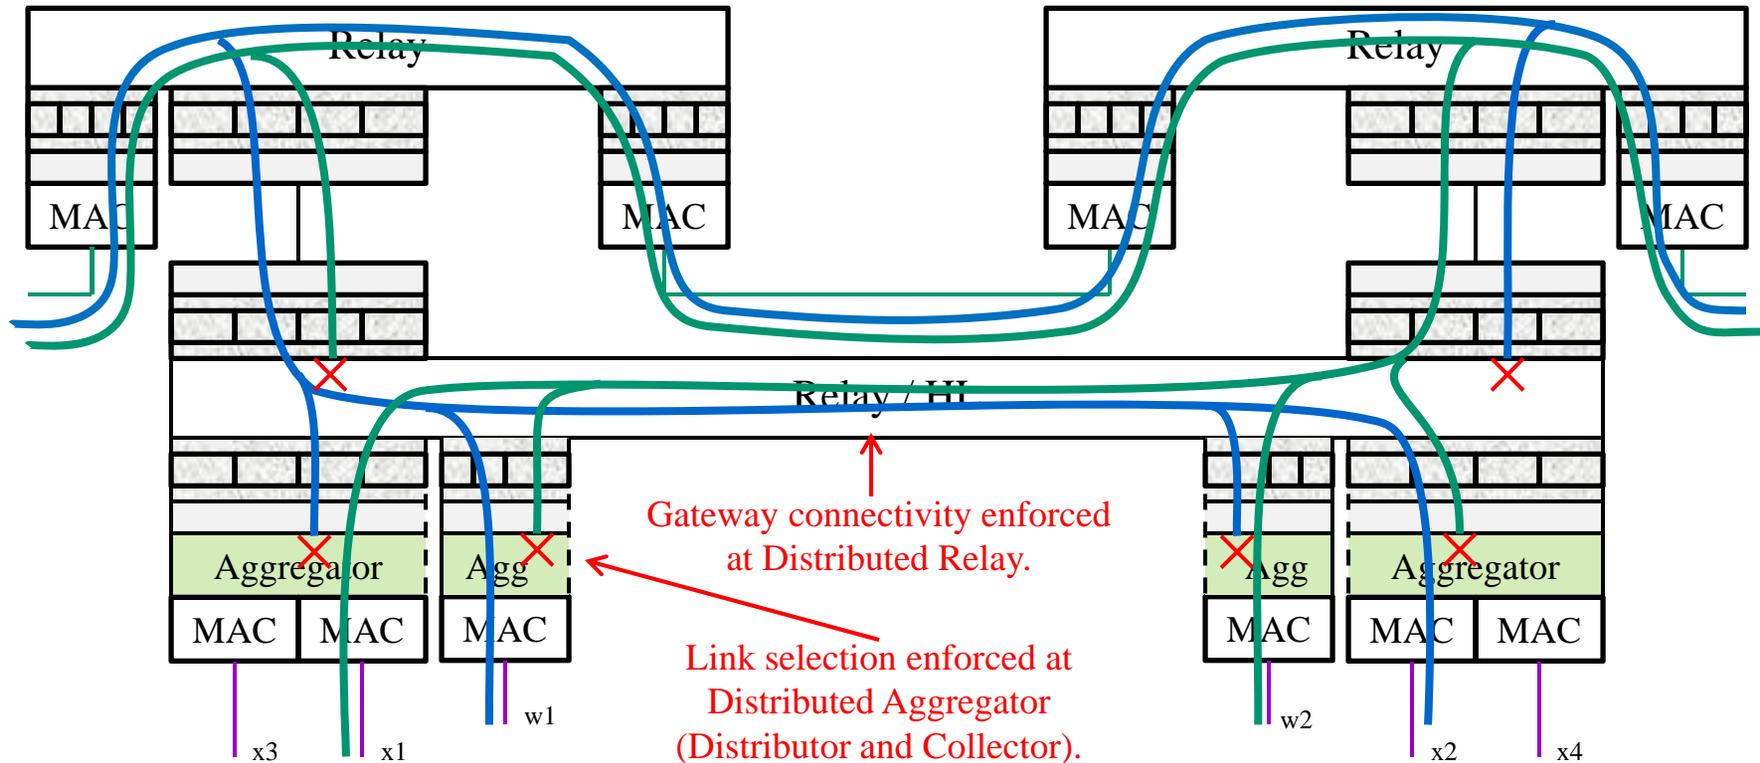
Logical Component Model of two S-tagged DRNIs

With constraint that each VLAN has the same Gateway for all DRNIs.



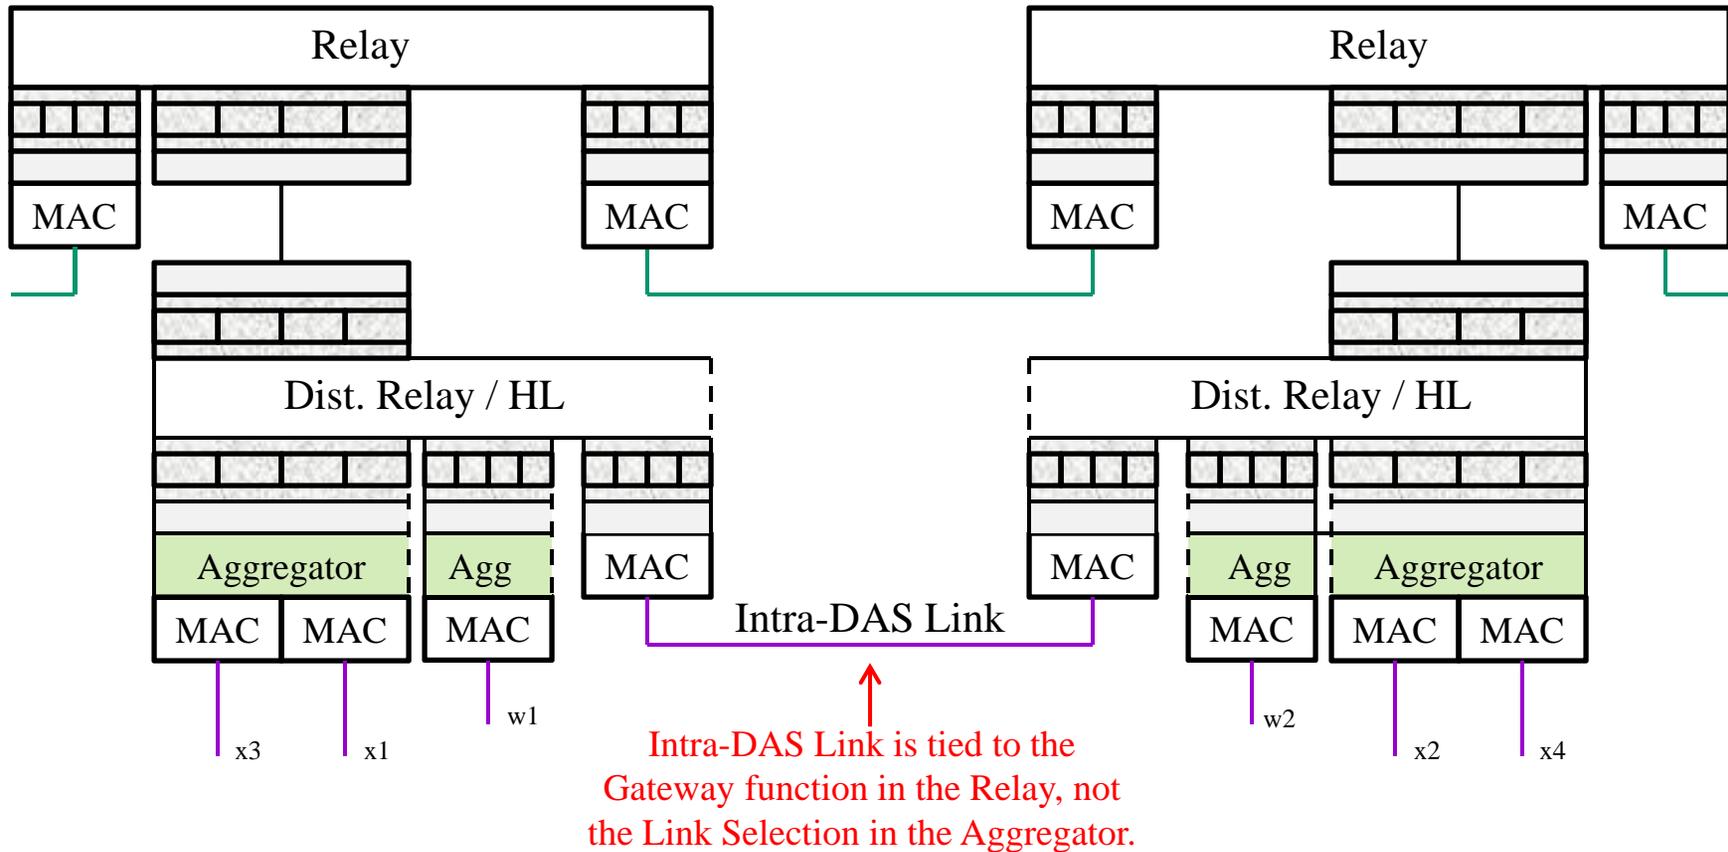
Logical Component Model of two S-tagged DRNIs

With constraint that each VLAN has the same Gateway for all DRNIs.



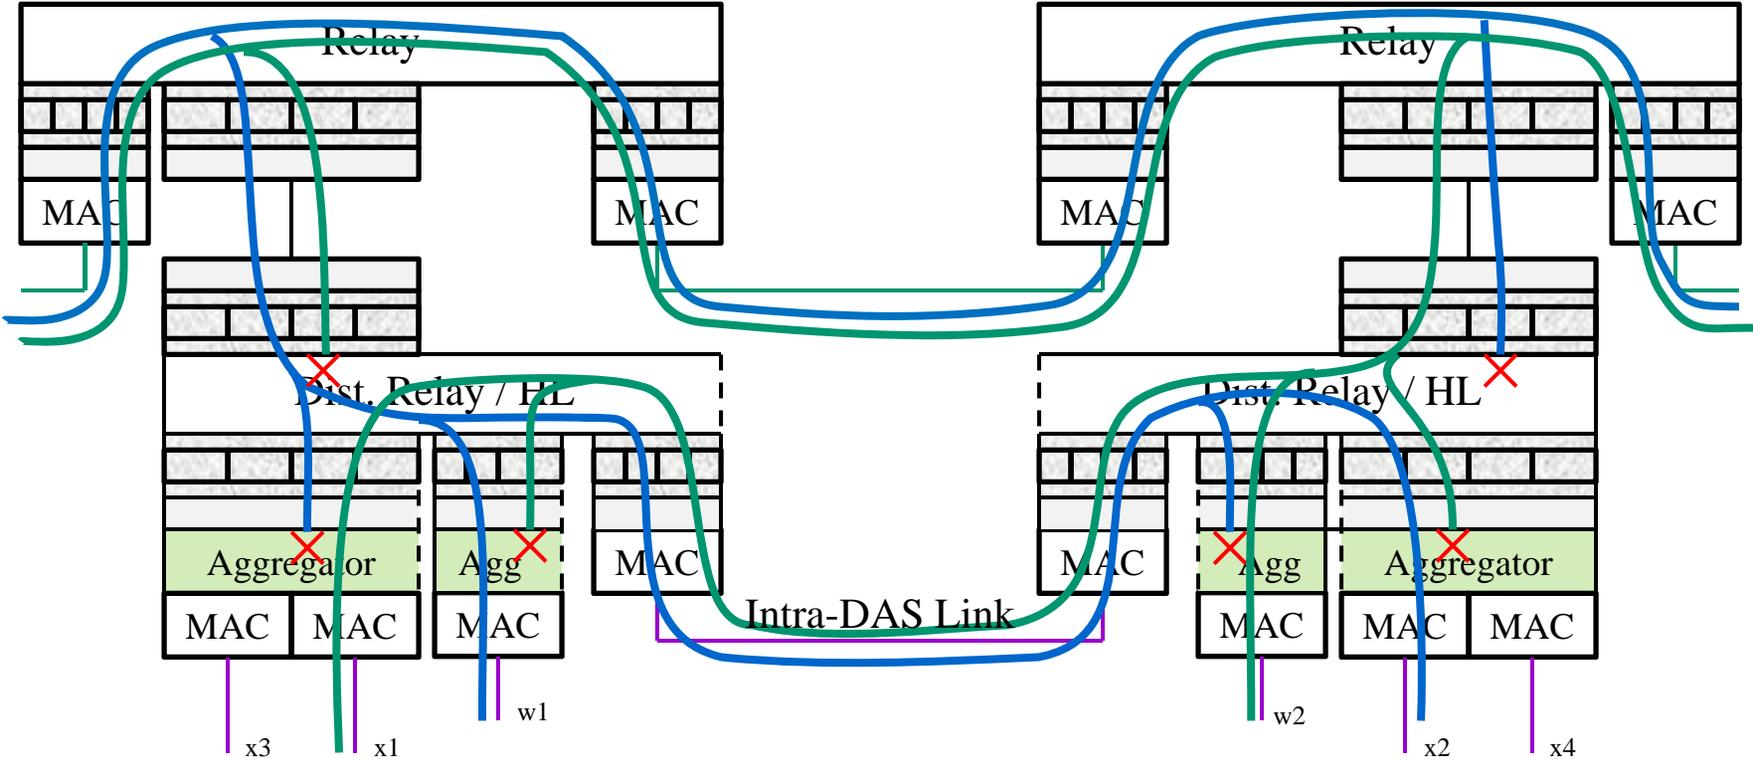
Distributed Component Model of two S-tagged DRNIs

With constraint that each VLAN has the same Gateway for all DRNIs.



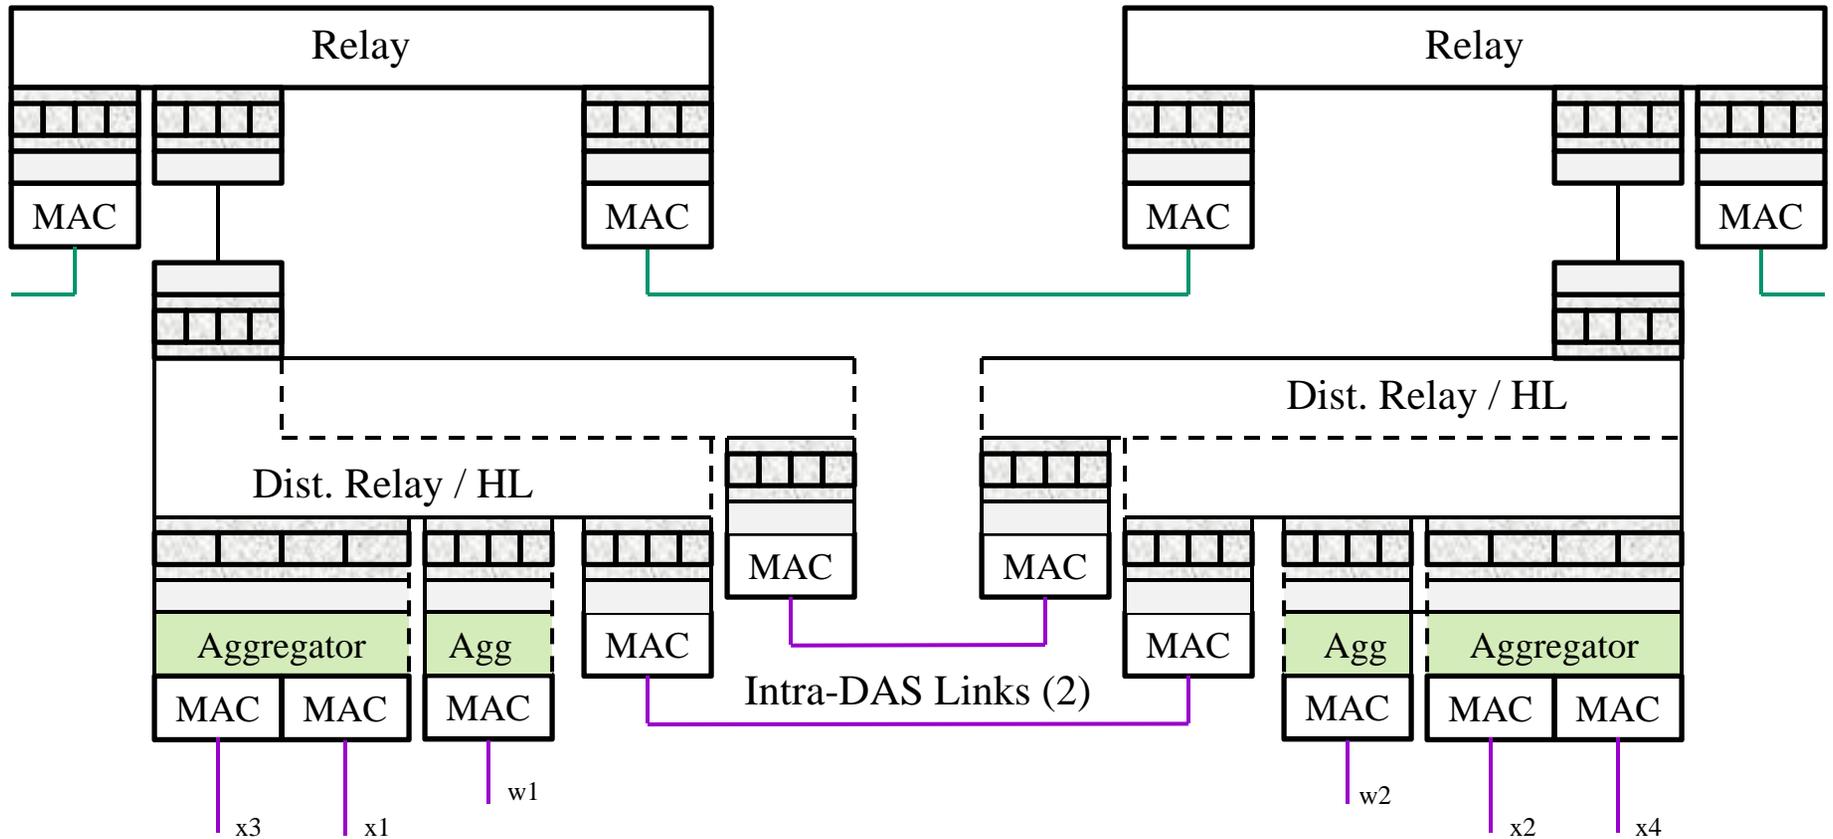
Distributed Component Model of two S-tagged DRNIs

With constraint that each VLAN has the same Gateway for all DRNIs.



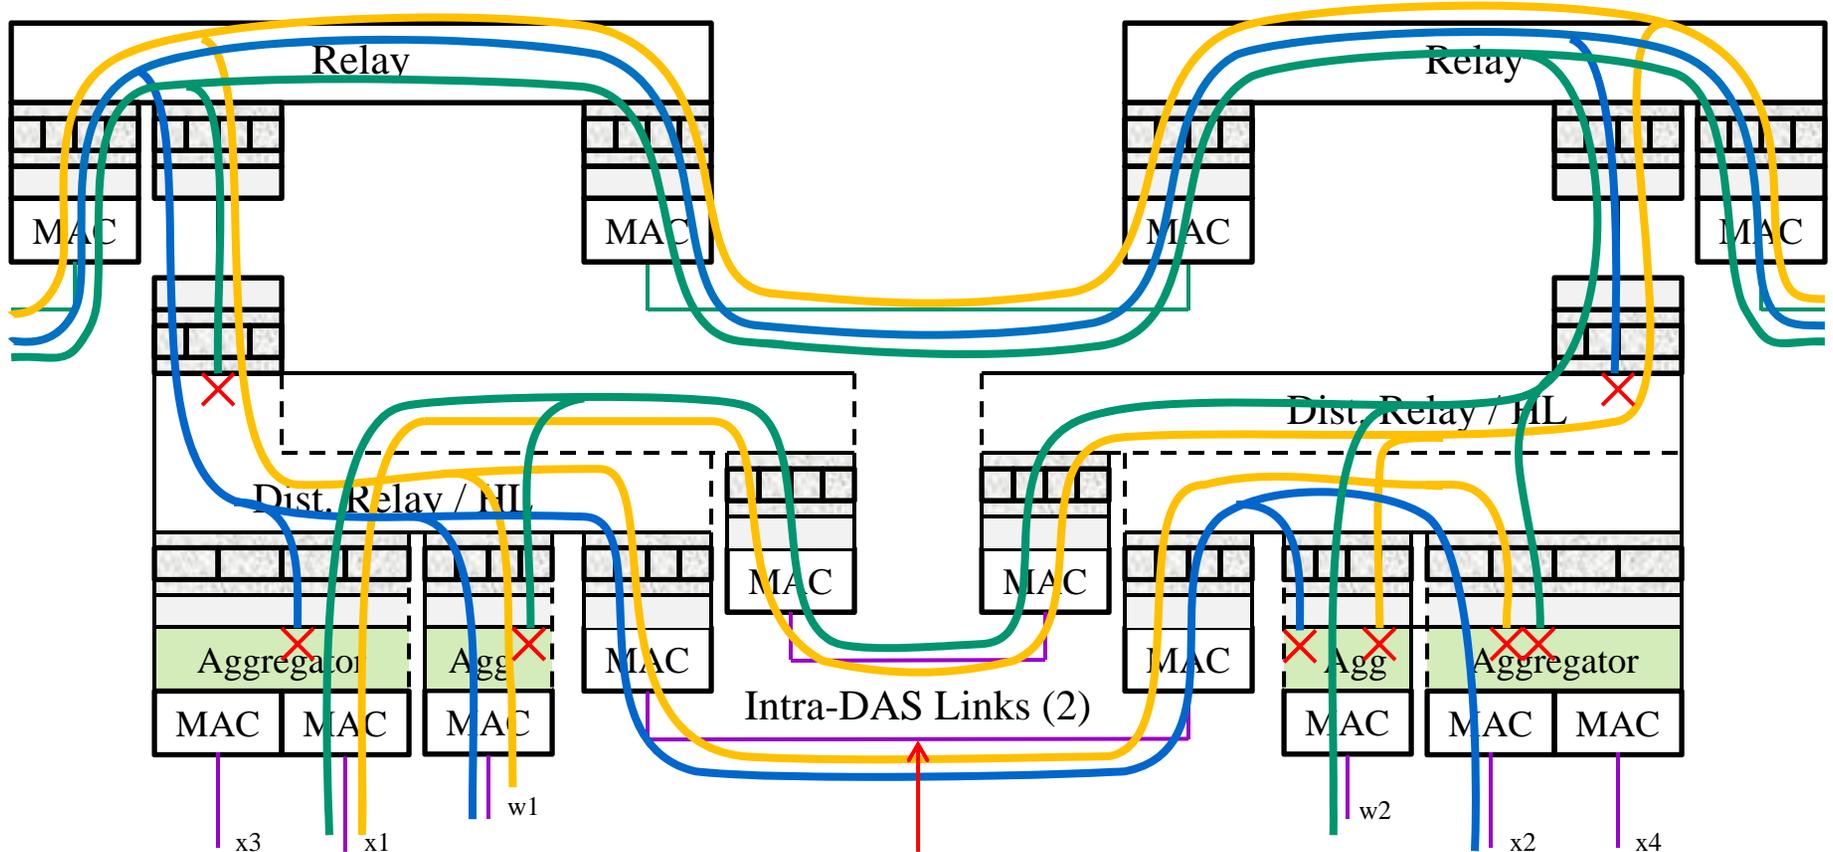
Distributed Component Model of two S-tagged DRNIs

No constraints on Gateway selection for each DRNI.



Distributed Component Model of two S-tagged DRNIs

No constraints on Gateway selection for each DRNI.

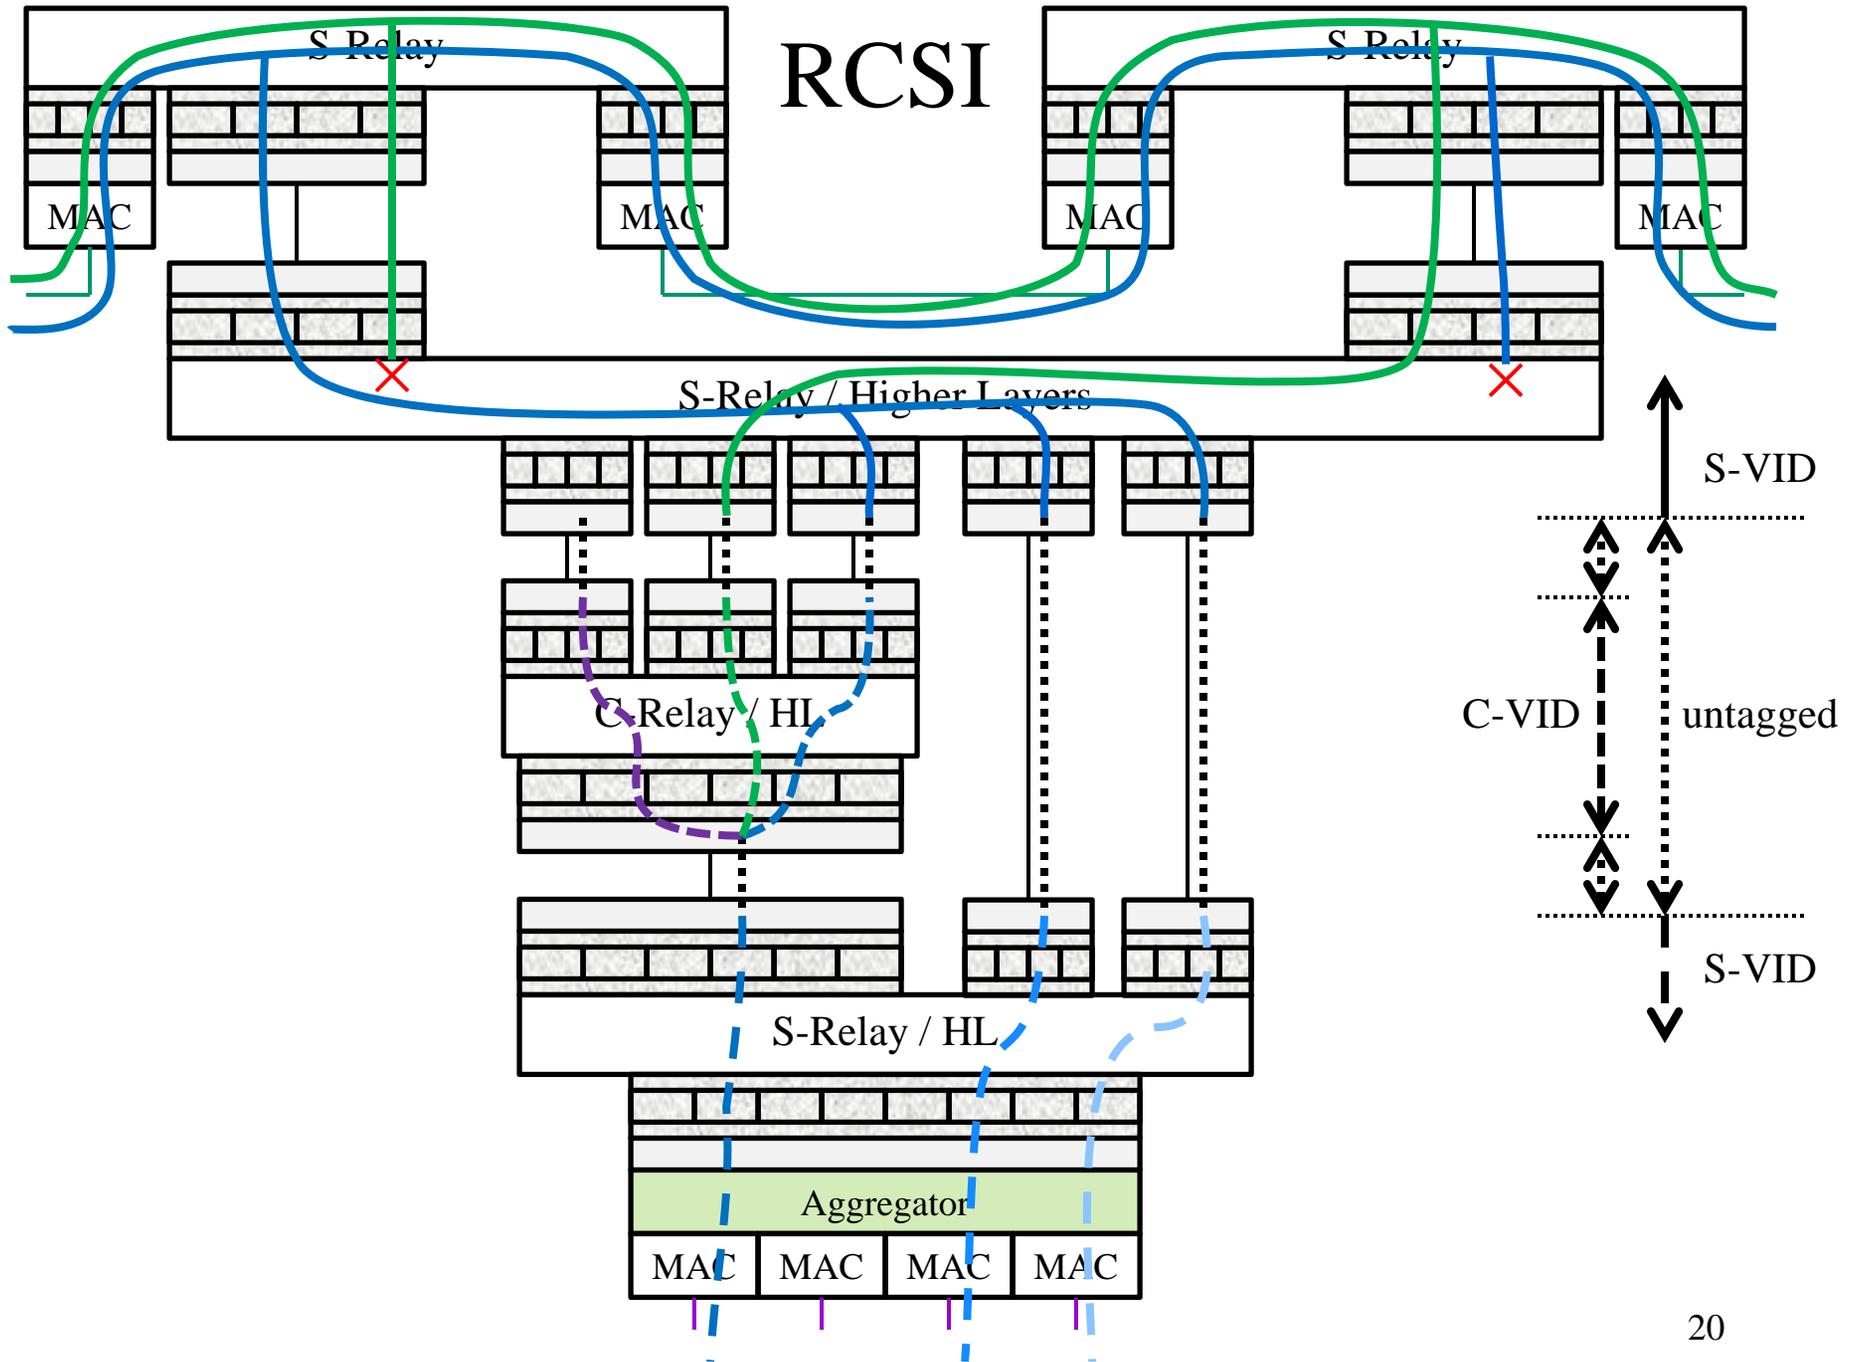


The two links provide resiliency for the Intra-DAS Link. In absence of a failure there are no constraints on the Gateway selection for any DRNI. If one link fails then connectivity is still maintained but may require a change of Gateway for some VLANs on some DRNIs.

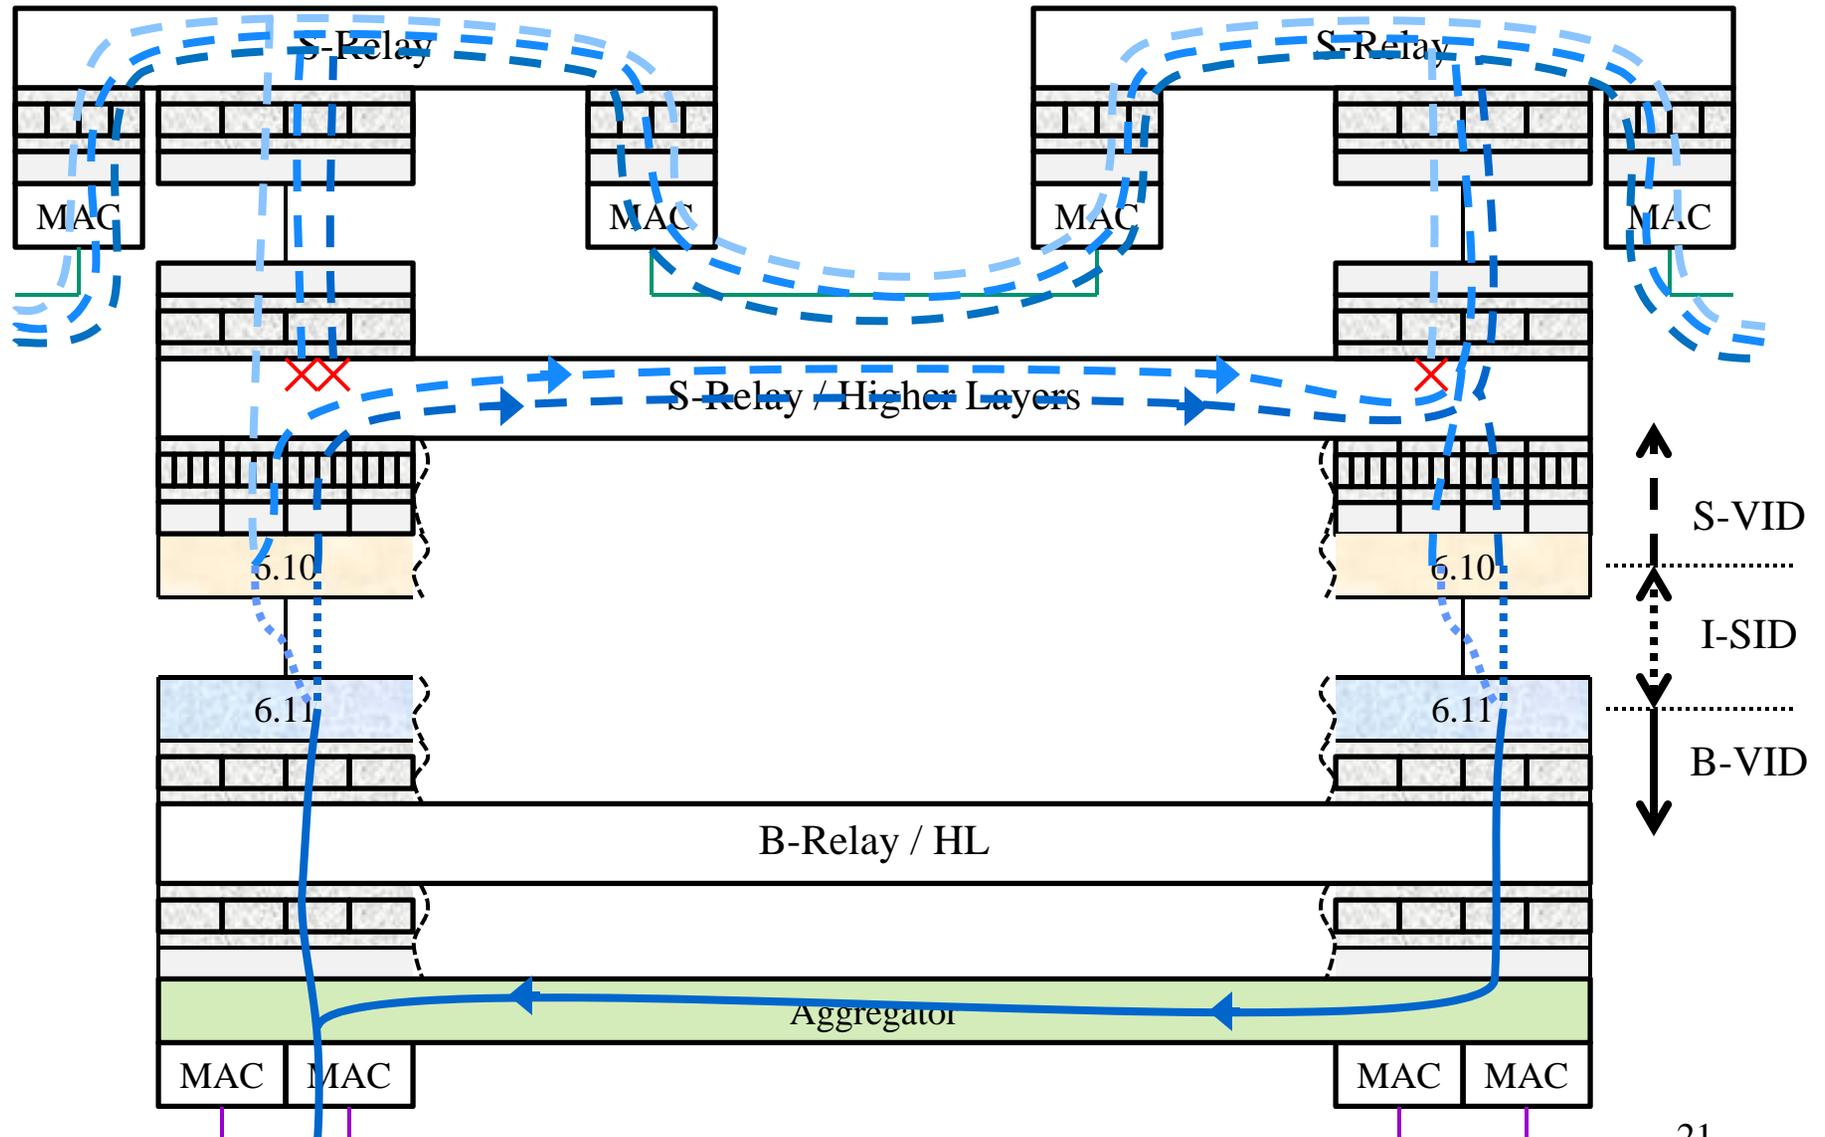
Implications on the DRNI model

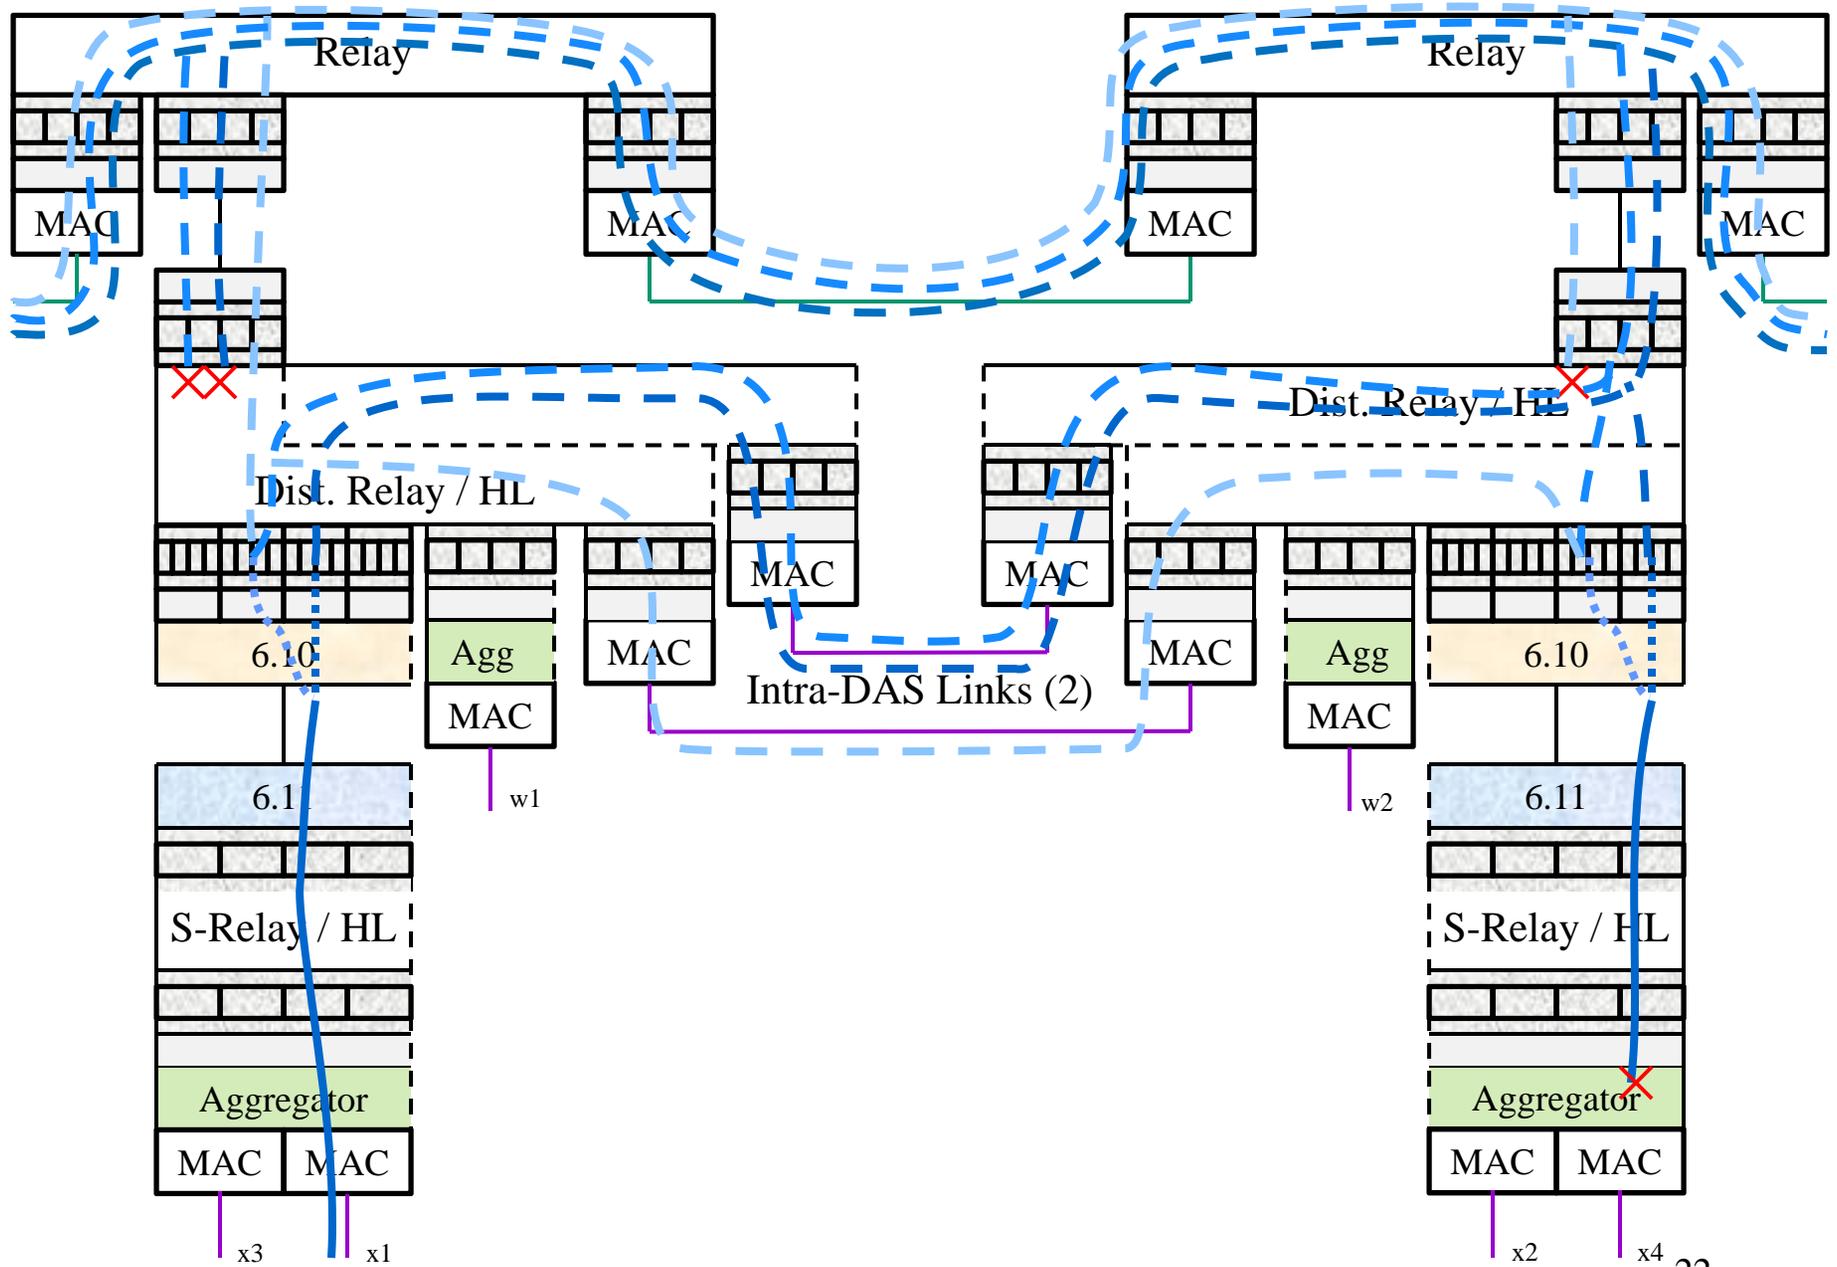
- The models supporting multiple DRNIs on a single logical component have two differences from previous models:
 1. Gateway and Distributed Relay functions:
 - In simple models of a single DRNI, the Distributed Relay is a simple multiplexer that enforces the Gateway selection by VID.
 - When supporting multiple DRNIs, the Distributed Relay is more than a simple multiplexer because it needs to be able to forward frames received on one logical Bridge Port to other Bridge Ports as well as to the Gateway.
 - Turns out a single DRNI on a Remote Customer Access Port (RCAP) of a PEB supporting RCSI also needs to be able to forward frames between logical Bridge Ports on the Distributed Relay.
 2. Intra-DAS Link is tied to the Gateway, not the Aggregator:
 - Some considerations even in the single DRNI model were driving to this conclusion anyway. It keeps all Gateway functionality in the Distributed Relay. If the Intra-DAS Link is tied to the Aggregator then the Gateway determination for received frames needs to be built into the Distributed Collector.

Backup Slides



Logical Component Model of 'B-tagged' DRNI w/ IB-BEB





Key to Interface Stacks in Diagrams

