Split-Horizon vs. VID filtering

Version 4

References

- This presentation is part of a response to an ITU-T to IEEE 802.1 liaison, available at: http://www.ieee802.org/1/files/public/docs2010/liaison-itut-oLS-213-0710.zip.
- This presentation is available at: http://www.ieee802.org/1/files/public/docs2010/liaison-nfinn-split-horizon-vid-filtering-0710-v04.pdf.

Rooted multipoint using VID filtering

Rooted multipoint service definition

- A “Rooted Multipoint” service has:

  One or more “root ports” that have connectivity to each other and have connectivity to all “leaf ports”.

  Zero or more “individual leaf ports” that each have connectivity with all root ports, but not with any other leaf port.

  Zero or more “leaf groups,” each consisting of two or more “group leaf ports” that have connectivity among themselves, and with all root ports, but not with any leaf port outside the leaf group.

- This definition is more general (leaf groups) than that supported by split horizon.
- Rooted multipoint has been fully supported by IEEE Std. 802.1Q since 2005.

Required VLAN IDs (VIDs)

- Each rooted multipoint service requires:

  One “root VID” *R*. (A Root port can support multiple tagged services.)

  One “individual VID” *I*, if there are any individual leaf ports, else VID *I* is not needed.

  One “group VID” V_g for each leaf group *g* (if any) in the service.

- For example, any number of individual leaves can be supported with two VIDs *R* and *I*. Three leaf groups plus any number of individual leaves can be supported with 5 VIDs *R*, *I*, V_1, V_2, and V_3.

- 4094 VIDs can support 2047 rooted multipoint services, or 3094 ordinary services plus 500 RMPs, which is more flexible than devoting a VID bit for 2047 RMP and 2047 ordinary, or sacrificing a priority or Drop Eligible bit.

Configuration: leaf ports see no tags, root ports are tagged (or untagged)

Configuration parameter	root port	group G leaf port	individual leaf port
Port VID Set: Bridge transmits only these VIDs.	*R*, *I*, all V_g	*R*, V_G	*R*
Untagged VID set: Bridge transmits these VIDs untagged.	None (or *R*, *I*, all V_g)	*R*, V_G	*R*
Filter VIDs on input: Bridge filters any input frame not in the Port VID Set.	Yes	No	No
Port VID: Bridge assigns this VID to any untagged frame received.	*R*	V_G	*I*
MVRP: Port attracts these VIDs via MVRP.	*R*, *I*, all V_g	*R*, V_G	*R*
Output VID translation: CE on Root port sees only VLAN *R*.	*I, all* $V_g \rightarrow R$	none	none
Admit tagged frames: Admit tagged frames received on port.	Yes	No	No

- All VIDs *R, I,* and all V_g use Shared Learning.
- A leaf port can support only one service, because it is untagged.

More advanced configurations

- The most common case is that one leaf supports only one (rooted multipoint) service, and that root ports supporting multiple services can be trusted not to transmit on the *I* and V_G VIDs. This is fully supported by IEEE Std. 802.1Q (and its amendments).
- Since the approval of IEEE Std. 802.1ad-2005, 802.1Q has also fully supported this more advanced configuration for Provider Bridges (as opposed to Customer Bridges).
- In order to support multiple services, including at least one rooted multipoint service, on a single suspect leaf port or root port, a Customer Bridged network must support two VID filters, one for input VIDs and one for output VIDs.
  - IEEE 802.1Q does not provide for a separate input VID filter; it provides only one in/output VID filter.
  - It is currently anticipated that this capability will be provided in future revisions of 802.1Q, via P802.1aq.

For multiple (tagged) services per leaf port:

- Each rooted multipoint service *n* requires:

  One "root VID" R_n.

  One "individual VID" I_n, if there are any individual leaf ports, else VID I_n is not needed.

  One "group VID" $V_{n,g}$ for each leaf group *g* (if any) in the service.

- Any number of individual leaves can be supported with two VIDs R_n and I_n. Three leaf groups plus any number of individual leaves can be supported with 5 VIDs R_n, I_n, $V_{n,1}$, $V_{n,2}$, and $V_{n,3}$.
- Each set of VIDs *n* use a separate shared learning group.
- This new feature is needed only for Customer networks; Provider networks have supported multiple tagged services per leaf port since 2005.

Configuration: all ports are tagged

Configuration parameter	root port	group n,G leaf port	individual leaf port
Port VID Set: Bridge transmits only these VIDs (before output VID translation).	all R_n, I_n, $V_{n,g}$	all R_n, $V_{n,G}$	all R_n
Input Port VID Filter: Bridge discards these [& unused] VIDs. Group port uses R_n (or $V_{n,G}$).	all I_n, $V_{n,g}$	all I_n, $V_{n,G}$ (all I_n, R_n)	all I_n, $V_{n,g}$
Untagged VID set: VIDs transmitted untagged.	None	None	None
Filter VIDs on input: Bridge filters any input frame not in the Port VID Set.	No	No	No
Port VID: Bridge assigns this VID to any untagged frame received.	one R_n	one $V_{n,G}$	one I_n
MVRP: Port attracts these VIDs via MVRP.	all R_n, I_n, $V_{n,g}$	all R_n, $V_{n,G}$	all R_n
Output VID translation: Root or Individual port sees only VID R_n, Group port sees R_n (or $V_{n,G}$).	I_n, $V_{n,g} \rightarrow R_n$	$V_{n,G} \rightarrow R_n$ ($R_n \rightarrow V_{n,G}$)	none
Input VID translation: Root or Individual port sees only VID R_n, Group port sees R_n (or $V_{n,G}$).	none	$R_n \rightarrow V_{n,G}$ (none)	$R_n \rightarrow I_n$
Admit tagged frames: Admit tagged frames received on port.	Yes	Yes	Yes

VLAN pruning: MVRP or IS-IS

- VLAN pruning prevents unnecessary flooding.
- IEEE 802.1 provides (or soon will provide) three means for VLAN pruning:
  - MVRP: Multiple VLAN Registration Protocol, for use with RSTP or MSTP spanning tree protocols.
  - IS-IS: For IEEE 802.1aq Shortest Path Bridging networks.
  - Administrative configuration: For PBB-TE networks.
- Both MVRP and IS-IS:
  - Propagate pruning information everywhere.
  - Control the *transmission* of VIDs, not the *reception*.

Perfectly accurate VLAN pruning via MVRP, IS-IS, or manual configuration

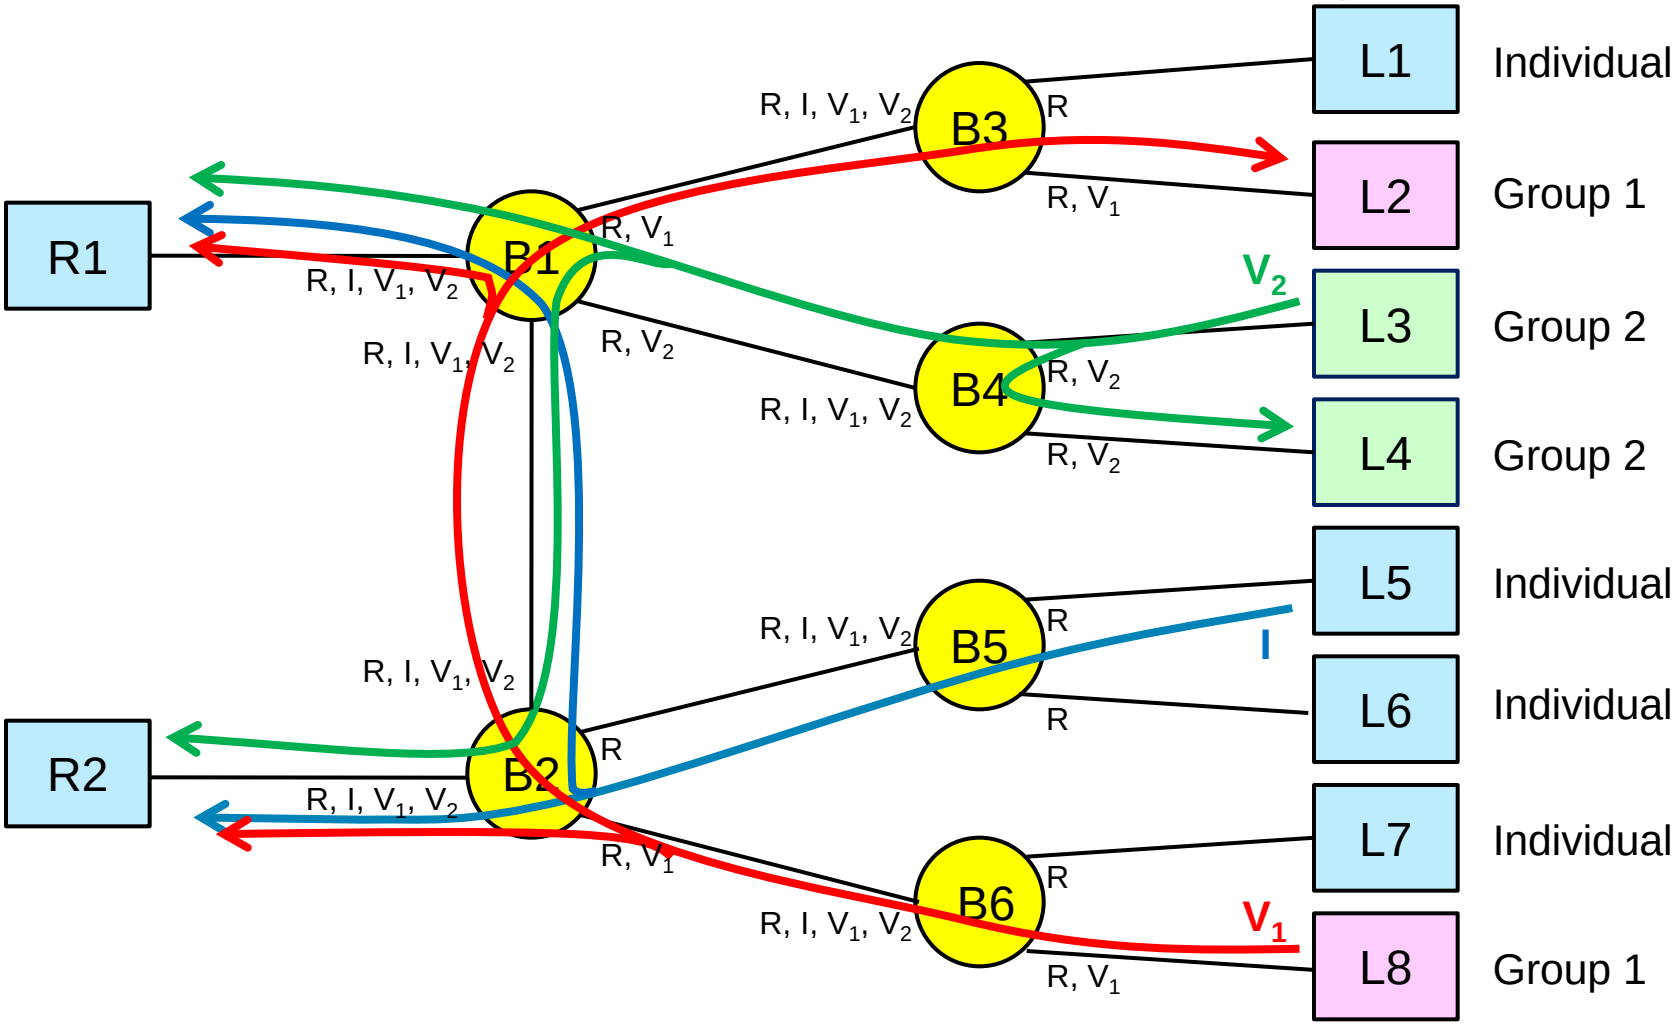

- Each port marked with VLANs that can exit that port.

Rooted multipath services: Multiple VIDs vs. split horizon

Criterion	Split horizon + “leaf bit” in 802.1Q tag	Multiple VIDs
Code point space	Consumes half of the available Priority or VID code points, or eliminates the Drop Eligible field.	Adds one code point per Rooted Multipath service (more if Group Leaf Nodes used).
Capabilities	Supports only Individual Leaf ports.	Supports both Individual and Group Leaf ports.
Standardization	New standards work required.	Standardized since 2005.
Implementation	Existing implementations must be modified.	Already supported by all implementations compliant with IEEE 802.1Q-2005 or later.
Multiple services per leaf port, including at least 1 rooted multipath service.	New standards work, new implementations.	Already supported by all implementations compliant with IEEE 802.1ad-2005 or later.

More optimal data distribution

Optimal distribution of data: split horizon

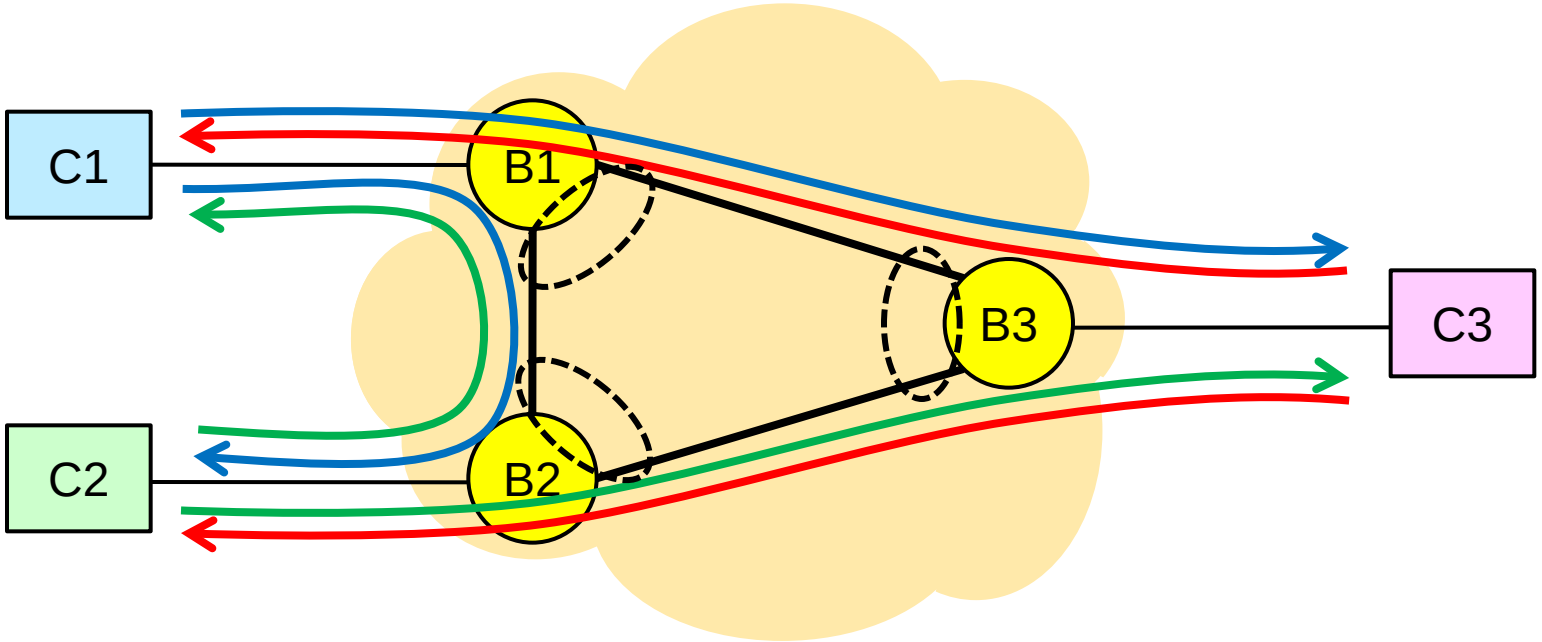

- The interior (heavy) links can be made part of a split-horizon port group.
- Then, **if** a full mesh of links among B1, B2, and B3 is maintained, data can be sent on the optimal path, as shown.

Optimal distribution of data: P802.1aq

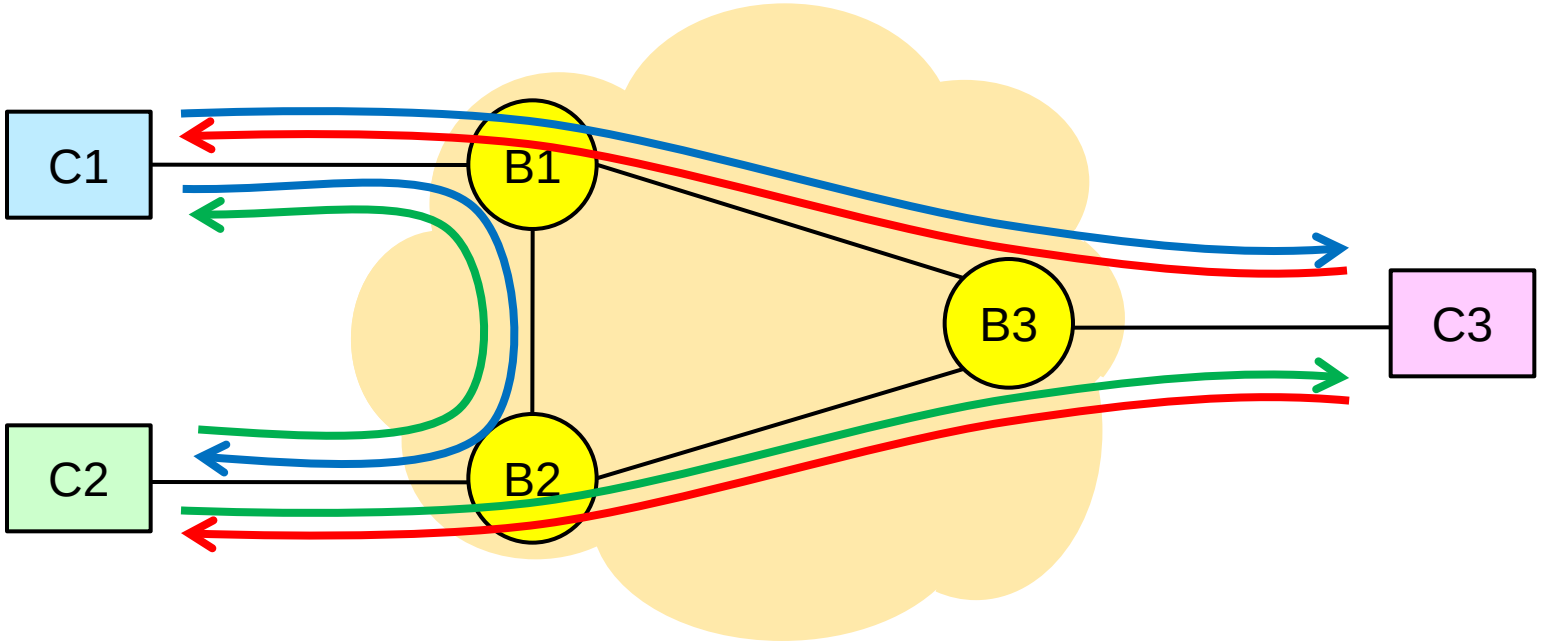

- Interior links are no different than any other link.
- P802.1aq's version of IS-IS can send data along the same (most optimal) paths as split horizon.

Optimal distribution of data

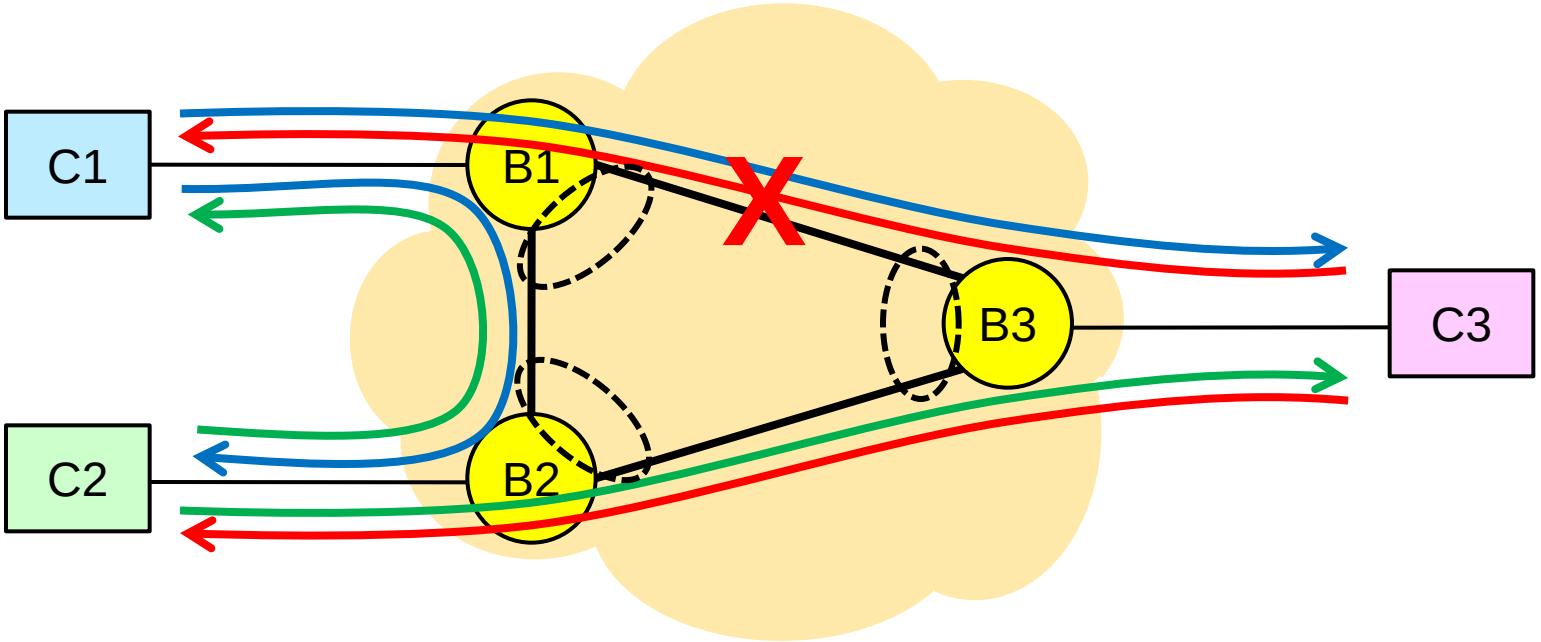

- What happens if a link fails (in this case, between C1 and C3) ?
- In split horizon, connectivity is lost.
- In P802.1aq, the traffic is rerouted through B2.

Optimal distribution of data

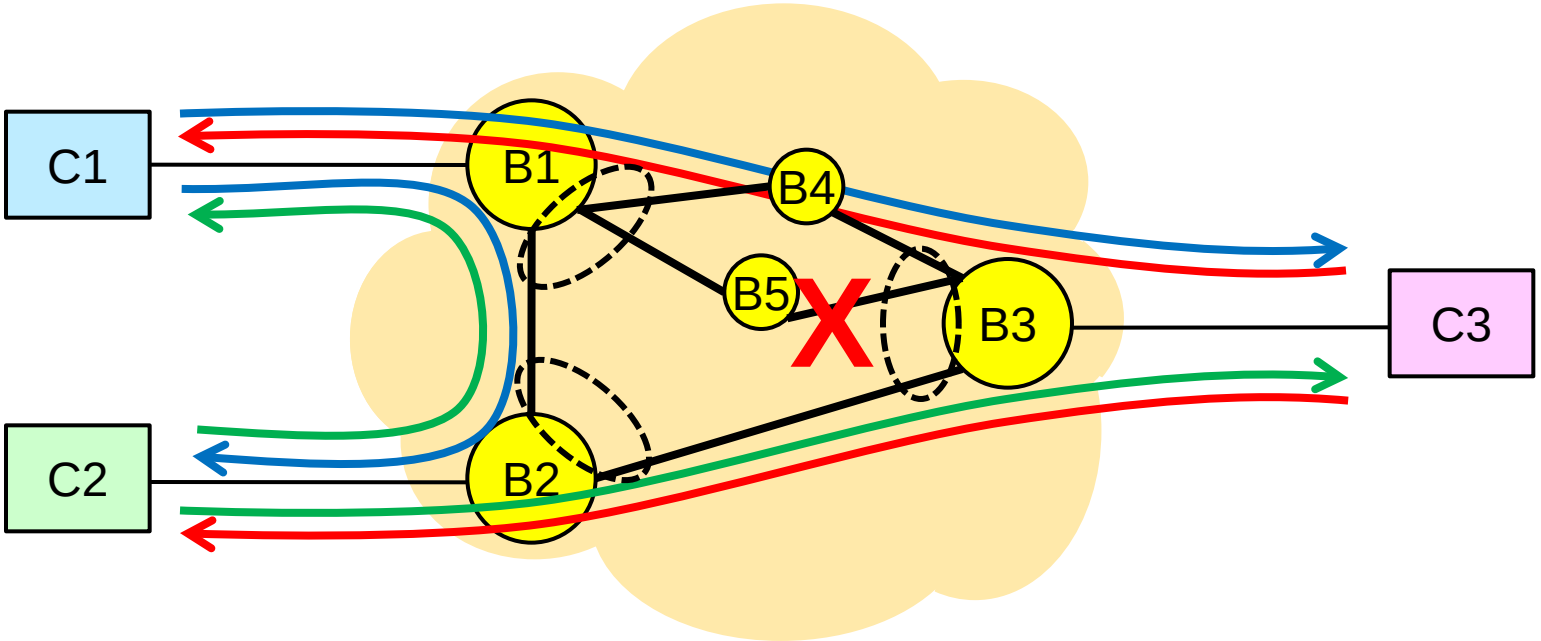

- “But,” says the proponent of split horizon, “the B1-B3 link is a protected link, so the data still flows!”
- To which the P802.1aq proponent answers, “There is no architectural reason why P802.1aq cannot run over protected links, as well, thus providing the same failover time for expected errors, and some connectivity in more cases.

Optimum data distribution: Multiple VIDs vs. split horizon

- Anything engineered traffic pattern that can be done with a full mesh of point-to-point links among split-horizon port groups can be done with exactly the same physical topology using links among bridge ports running P802.1aq.
- To get the same quick response to failures of the full mesh, the same links can be protected using protection switching, so that IS-IS recovery times are not invoked.
- Please note that the above is not a comparison of the relative merits of IS-IS vs. protection switching – that argument does not involve port groups!
- This comparison shows that one can think of IS-IS as establishing, automatically, exactly the same data distribution that can be established manually using split-horizon port groups, and that IS-IS further insures against catastrophes.
- **Note for both 802.1 and SG15:** there is no architectural reason why both protection switching and bridging cannot make use of each other as network components, simply a lack of supporting standards.

Optimal distribution of data: Non-802.1aq

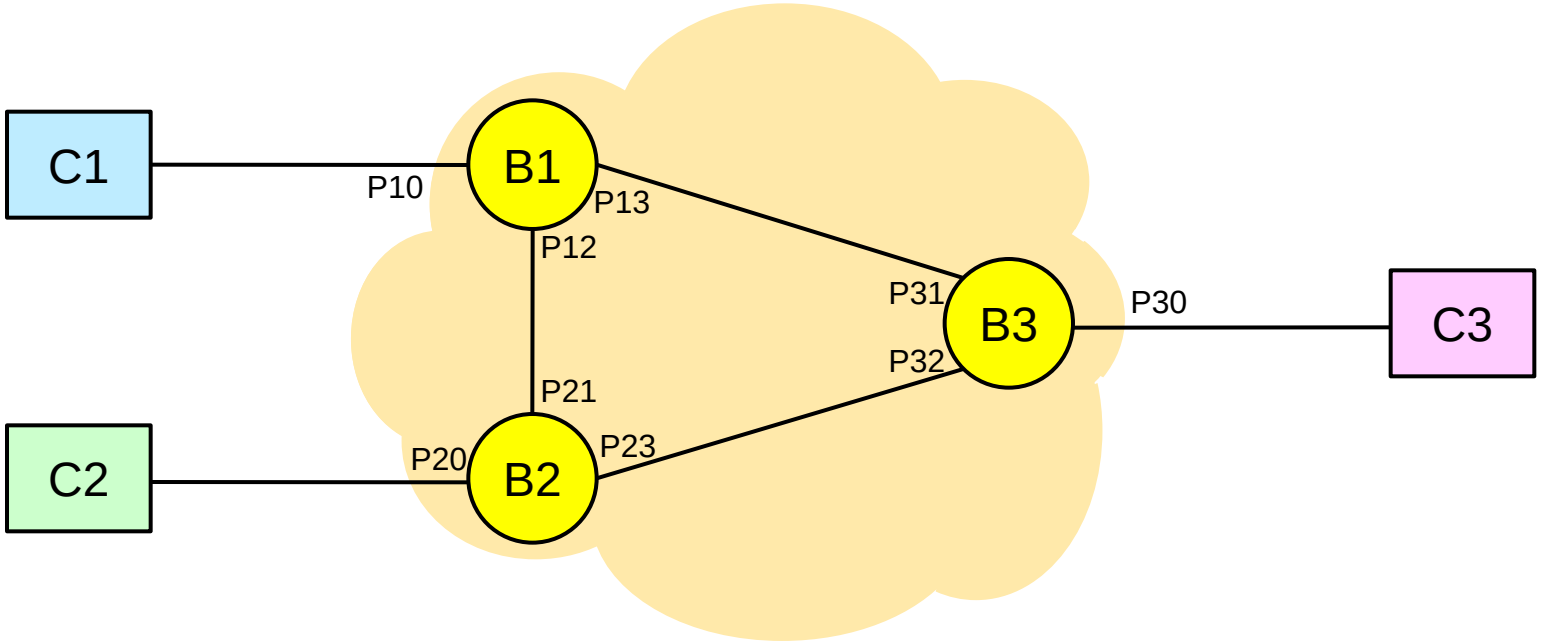

- If 802.1aq is not needed or wanted, the effect of a “port group” can be achieved within current standard bridge forwarding and port filtering capabilities, using VIDs, without introducing the concept of a “port group”.
- In addition to the external service VID *V*, seen by the Customer Equipment, there is a second VID *I* used internally only.

Using VIDs for manually configured optimum data distribution

Bridge	Port	Can transmit (before xlate)	Output translation
B1	P10	*V*, *I*	*I* → *V*
	P12	*V*	*V* → *I*
	P13	*V*	*V* → *I*
B2	P20	*V*, *I*	*I* → *V*
	P21	*V*	*V* → *I*
	P23	*V*	*V* → *I*
B3	P30	*V*, *I*	*I* → *V*
	P31	*V*	*V* → *I*
	P32	*V*	*V* → *I*

- The separate VID *I* prevents forwarding from one “internal” port to another.
- VIDs *V* and *I* use shared learning.

Reinventing bridging

- If one starts with the following premises:
  - I want to do approximately what bridging does;
  - I don’t want to use 802.1 standards;
  - I have been working for a long time with a set of standards (e.g. IETF, ITU-T) with a more-or-less consistent outlook;
- Then, I will come up with something that works, and will be sufficiently, though not completely, compatible with the expectations of current users of bridged networks.
- Does the industry really need three ways (IEEE, IETF, and ITU-T) of doing bridging?

Action items

Action items

- IEEE 802.1Q has a number of capabilities that have been widely deployed, but may not be immediately apparent to the reader. 802.1 will be happy to provide further assistance on like matters.
- IEEE 802.1 requests that ITU-T Q.9/S.G.15:
  - Remove port group functions in 9.1, 9.1.1 and App VII;
  - Replace with appropriate model of existing 802.1 functionality, specifically, bridge forwarding and port filtering capabilities, using VIDs. (The use case is described in 802.1Q Annex B.1.3.)