



# Simplifying Seamless Redundancy

Norman Finn  
Cisco Systems

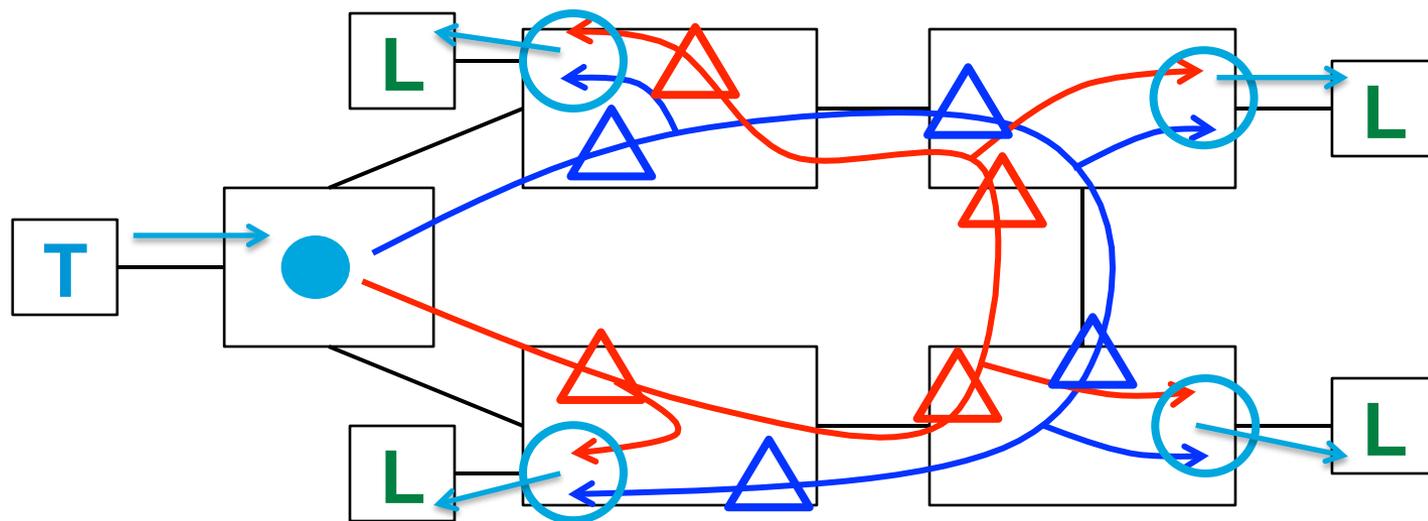
Version 02

January 19, 2016

# Problem

- **This document supports a ballot comment on P802.1CB D2.1.**
- At present, configuring P802.1CB requires a massive amount of configuration.
  - Every single Stream's identification (typically {VLAN, MAC address} pair must be configured on every port on which it might be received on every bridge in the network.
  - Every port on which each Stream might be output twice has to be configured with a sequence recovery function and its parameters (e.g. timeout value).
  - Every bridge through which the Stream passes should be configured with an Individual recovery function, requiring the configuration of Stream Identification and the IRF.
- This is a real problem – it will make bringing up a new Stream take too long to meet many needs. The volume of configuration information, itself, will be an impediment to adoption by users.

# What we want:



- Each Listener gets only one copy. Packets are not discarded until just before output to the Listeners, so that latent errors can be detected.

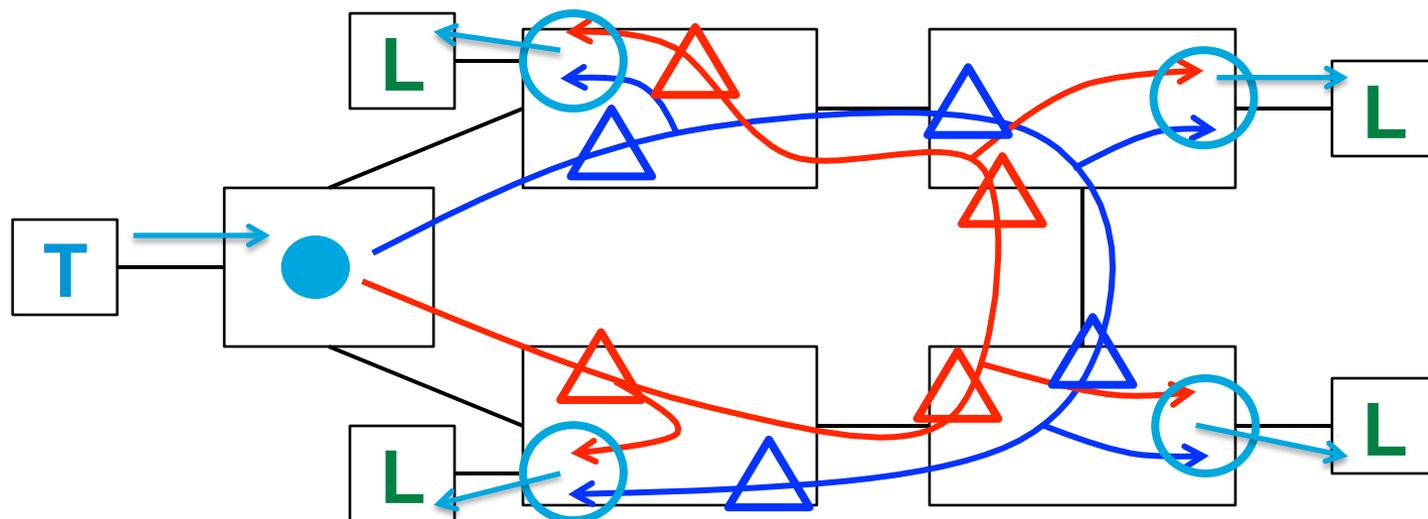
# What we want

- We would like Sequence Recovery and Individual Recovery to auto-configure themselves.
- This means that an instance of one or both of these state machines pops into existence whenever a new Stream (for Individual recovery) or a new Seamless Stream (for Sequence recovery) is seen, using default parameters specified by (a little bit of) configuration.
- When the timeout expires, sending a state machine into “accept any” mode, the state machine can be destroyed.
- We must identify the **Seamless Stream** to auto-configure **Sequence Recovery** functions, and we must identify the individual **Stream** for **Individual Recovery** functions.
- We must also distinguish (as we will see) on which ports recovery functions are to be applied.

# Problem 1: On what do you base auto- configuration of Sequence and Individual Recovery?



# Which streams are paired?



- There are > 2 streams in the network! How do you know that **red** and **blue** are a pair (belong to the same Seamless Stream)?

# On what do we base auto-configuration?

	Sequence Recov. based on	Individual Recovery based on	Comments
1.	{VLAN, Dest}	{VLAN, Dest, Port}	Hard to debug. Can't use current topology protocols to set up paths; they're not trees.
2.	{VLAN, Dest}	{VLAN, Dest, A/B tag}	Not easy to debug. Can't use current topology protocols to set up paths; they're not trees.
3.	{Dest}	{VLAN, Dest}	<i>I think</i> that the number of VLANs are not an issue.
4.	{Source, VLAN}	{Source, VLAN, Dest}	Per-source sequencing, not per-Stream.
5.	{Source, VLAN}	{Source, VLAN, A/B tag}	Per-source sequencing, not per-Stream. Destination not used by auto-config.

- NOTES: Always, only frames with CB-tags are processed.
- “A/B tag” means HSR-like bit(s) in the sequence tag identifying the path taken.

# Choosing an auto-configuration solution

- In this author's opinion:
- 1 & 2 (almost identical packets) are not good, because they have identical (or identical except for the CB-tag) packets criss-crossing through the network. Among other difficulties, this is incompatible with reverse path checking for multicast pruning.
- 3 & 4 (DA identifies Stream within the group) conflates forwarding with stream identification. The DA must be a multicast if one Listener receives  $> 1$  stream. While this is necessary for Bandwidth reservation, 802.1CB can be useful without bandwidth reservation. 3 seems best if per-Stream sequencing is required.
- 5. Author's favorite: Although this uses per-source sequencing, it allows multiple Streams per Listener without requiring a multicast DA, thus separating 802.1CB from MSRP.
- Summary on next slide.

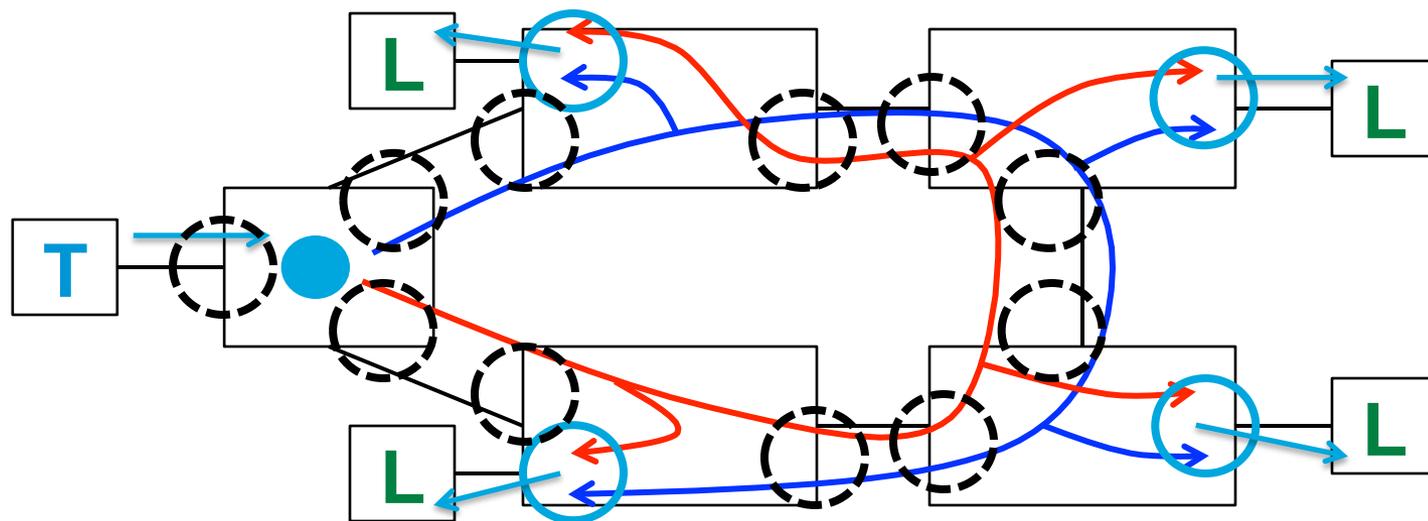
# Author's opinion

	Sequence Recov. based on	Individual Recovery based on	Comments
1.	{VLAN, Dest}	{VLAN, Dest, Port}	<b>Non-starter.</b> Identical packets are criss-crossing through the network.
2.	{VLAN, Dest}	{VLAN, Dest, A/B tag}	<b>Non-starter.</b> Almost-identical packets are criss-crossing through the network. (The A/B bits in the tag are not used by forwarding.)
3.	{Dest}	{VLAN, Dest}	→ <b>contender.</b> SR cannot be used except on multicast streams, and multicast DAs are required (now) only for bandwidth reservation.
4.	{Source, VLAN}	{Source, VLAN, Dest}	No. Requires multicast DAs without gaining advantages of per-Stream sequencing.
5.	{Source, VLAN}	{Source, VLAN, A/B tag}	→ <b>Compatible with HSR/PRP. Seamless Redundancy not locked to MSRP.</b>

# Problem 2: On which ports do you perform Sequence Recovery?

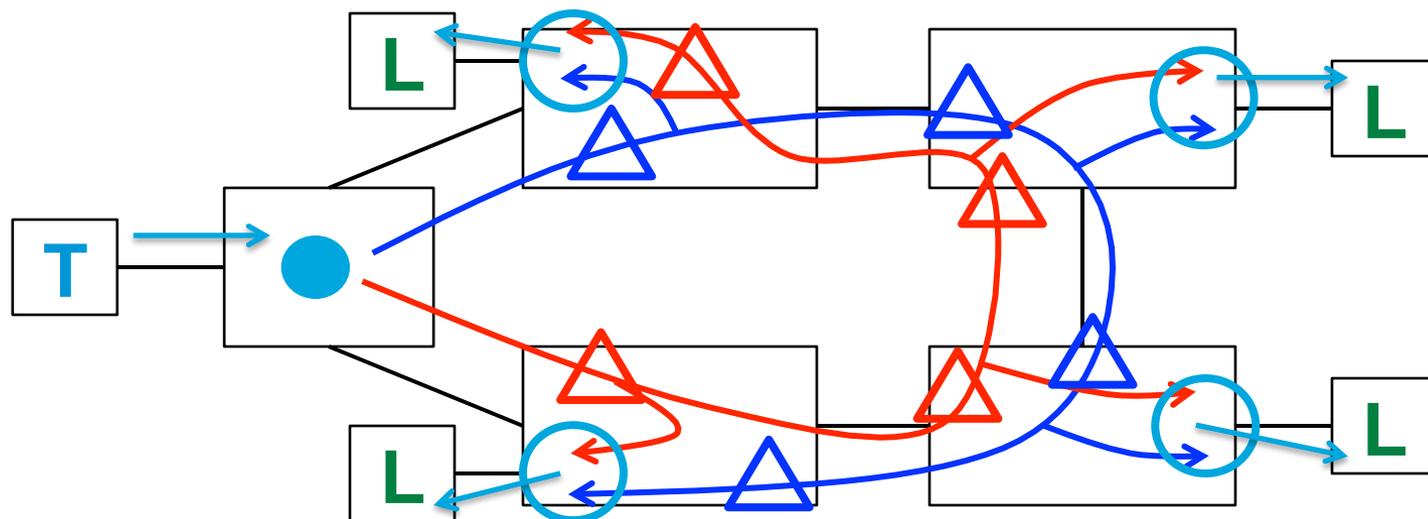


# Which ports are sequence filtered?



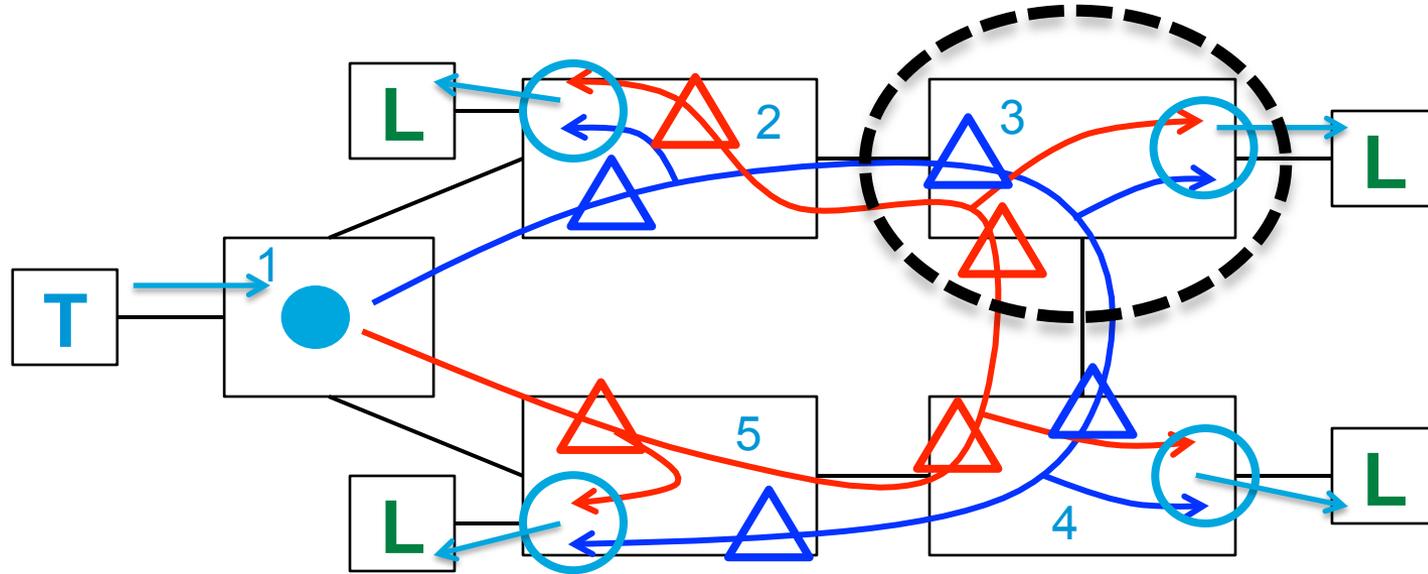
- **You can auto-configure Sequence Recovery on every port!**
- The dotted circles are functional, but never discard a packet, because the two streams **do not both egress** on those ports.

# Which ports are sequence filtered?



- You can auto-configure an Individual Recovery function in every forwarding function, separate from any particular port.
- IF identical packets entering different ports are the same Stream.

# Sequence Recovery on every port $\neq$ Sequence Recovery in Relay Function



- We do want to filter packets going to the Listener.
- We do not want to filter packets on the ring – that can lead to latent errors.

# Sequence Recovery on every port $\neq$ Sequence Recovery in Relay Function

- In the previous slide, in bridge 3, you
  - **Do** want to discard duplicates sent to the Listener.
  - **Do not** want to discard duplicates on the ports to Bridges 2 and 4.
- A single Sequence Recovery function in the bridge 3 relay function that does not pay attention to individual ports **could**:
  1. Receive red packet #10 from bridge 4, and relay it to L and bridge 2.
  2. Receive blue packet #10 from bridge 2, discard it as a duplicate, and not send it either to L or to bridge 4.
- In one sense, this is no problem, because we know that bridge 4 has already received packet #10. It also makes more bandwidth available for best-effort traffic.
- But, if link 1—5 fails, we don't know for sure that packet #11 will reach bridge 5; there may be a “**latent error**”, e.g. a misconfiguration that blocks blue packets on link 4—5.

# Problem 3: Do we need Stream Confluence?



# The (non-)need for Stream Confluence

- At present, there is no requirement for a Bridge to configure a Stream Merge function. (Draft 2.1 Table 8-1.)
- The only reason for a Bridge to configure a Stream Split function is to satisfy Comment #7 on Draft 2.0 (see [cb-nfinn-seamless-issues-1015-v02.pdf](#)). If that problem is ignored and we reverse that decision, then we can eliminate the Stream Confluence sublayer entirely from P802.1CB.
- Note that the issue of Draft 2.0 Comment #7 is the inability (without Stream Splitting) to take in one packet from an end station and generate two copies of that packet on one output port with two different encapsulations.

# The (non-)need for Stream Confluence

- Some text will still be required in Annex E, to show that it will be commonly the case that the Streams comprising a Seamless Stream will be all fed through a single instance of the Sequence Recovery function on an output port, after which some or all of the Streams' encapsulations will be changed in order that a single Stream is output on that port for the Seamless Stream.

# Suggested Remedy



# Suggested remedy

- Put managed objects into Clause 9 to support auto-configuration:
  - Administrator selects a Default Stream/Seamless Stream identification and Default sequence number encapsulation plan for the relay system.
  - Sequence Recovery is then auto-configured on all ports as shown in Figure 8-1. (This can still be represented as a function inside the relay as in Figure 8-2. It just means that the in-the-relay function has to pay attention to the entry and exit ports.)
  - Individual recovery is also auto-configured for all streams, without regard to port. Note that this function can ignore entry vs. exit port.
  - Support for plan 5 (above) using the HSR tag is required. Plan 3 using the CB tag is optional.
- An implementation may (optionally) support the detailed per-Stream configuration of the current Clause 9, with auto-configuration, if turned on, handling the Streams that miss the detailed configuration.
- Eliminate the Stream Confluence sublayer, making it an option to have the ability to output multiple copies of a packet with different encapsulations.

Thank you.

