

draft-ietf-detnet-dp-sol-00

Issues

Detnet working group, IETF100, November, 2017

Norman Finn, Huawei

Issues

- 42 comments recorded in dp-sol-00.
- Other email discussions.
- We cannot go into every comment, here.
- Following is a selection of topics suitable for discussion here, most of which span multiple individual comments.
- Questions for the WG to answer in blue.

Generalized Associate Channel (GACh)

- Will the GACh work with PREF?
- All uses?
- To Be Determined and added to dp-sol.

Load Sharing

- We've had email discussions about "Load Sharing". That is, distributing (*not replicating*) a flow's packets over multiple links, because no one link has sufficient bandwidth to carry the flow.
- We would, presumably, use PREF to restore the packets to the right order.
- This would need to be documented in at least the architecture and dp-sol drafts, and perhaps be added to the use cases and/or problem statement.
- Is this a DetNet requirement, a distraction, or something in-between?

What goes in what draft (to be RFC)?

- Some parts of dp-sol, especially parts of sections 1 and 4, need to be eliminated, reduced, or incorporated into architecture draft.
- dp-sol includes two solutions: IPv6 encapsulation, and something similar to pseudowires.
- Should dp-sol be split into two drafts, describing the two different solutions, or remain a single draft?

Which particular “pseudowire-ish solution”?

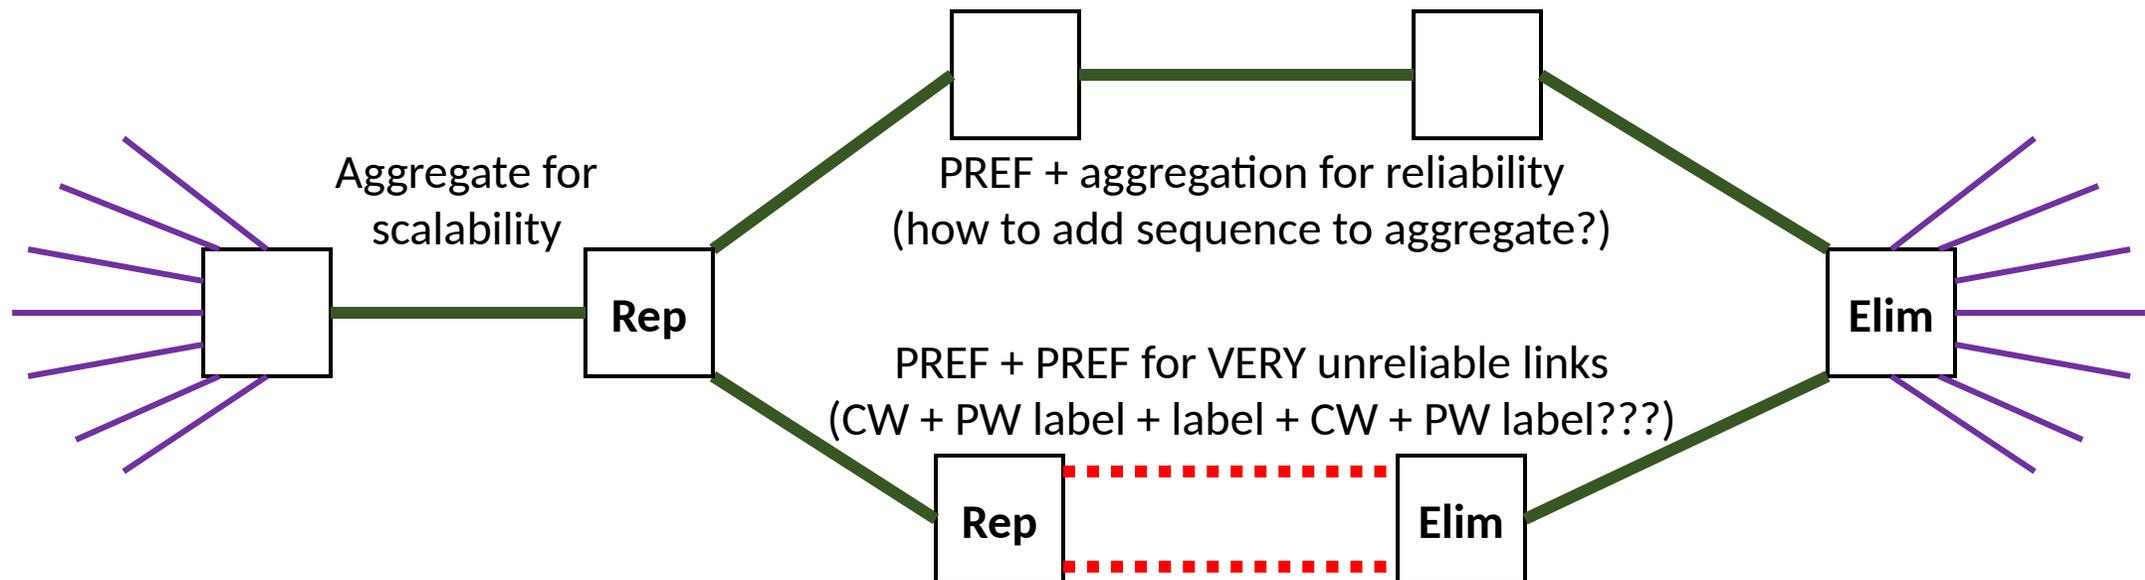
- Three ways to describe the pseudowire-ish solution are proposed:
 - dp-sol is based on RFC6073 Multi-Segment Pseudowires (MS-PW).
 - It could be based on RFC7432 BGP MPLS-Based Ethernet VPN (EVPN).
 - It could describe a brand-new construct, the “DetNet Wire” (DW).
- How much do these choices differ in the implementation of the data plane?

Issues driven by MS-PW / EVPN / DW answer

- The rotating sequence number in the control word skips over 0, because 0 means, “sequencing not done; don’t eliminate me”.
 - Skipping over 0 is not done in three existing L2 technologies with which we wish to interwork, and dropping duplicates is not optional.
 - [Is this behavior an exception to existing RFCs, or do we describe a new DW?](#)
- There are terminology issues among these choices, with the IPv6 encapsulation, and with the architecture document.
 - The architecture draft necessarily uses neutral terms different from all of the above.
 - dp-sol is inconsistent in its choices to use terms from the architecture draft, or terms from its basis RFC6073 MS-PW.
 - [Which RFCs’ terminology do we use in dp-sol?](#)

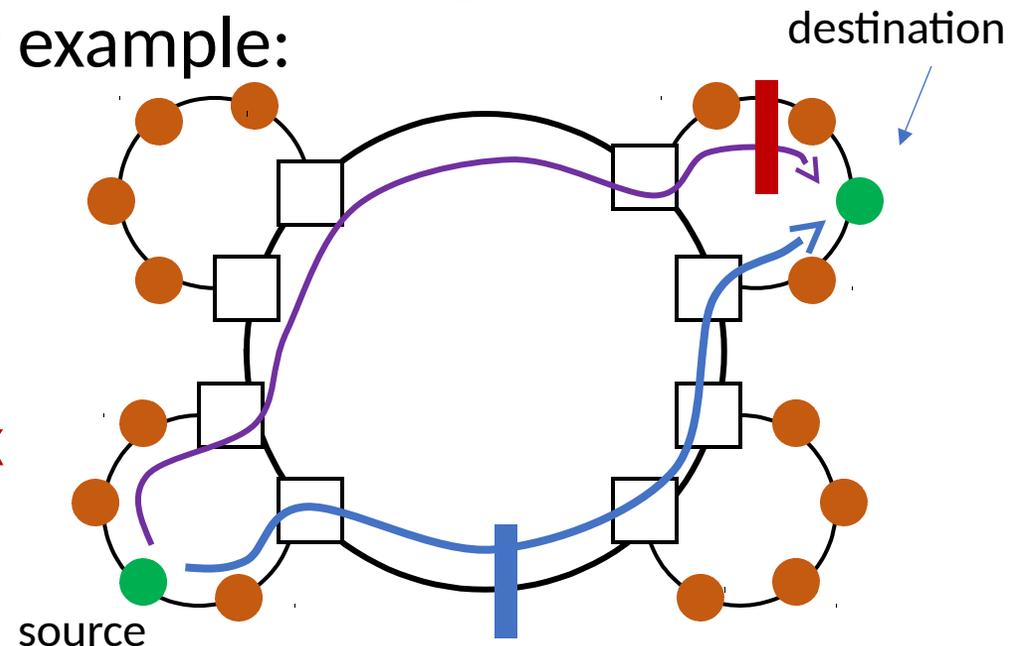
Aggregation encapsulation techniques

- Not well explored in current dp-sol draft. Additional work needed.
- Aggregation is required to scale up the number of flows receiving bounded latency service. Adding an outer MPLS label is all we need.
- **Do we need multi-layer PREF?** Examples:



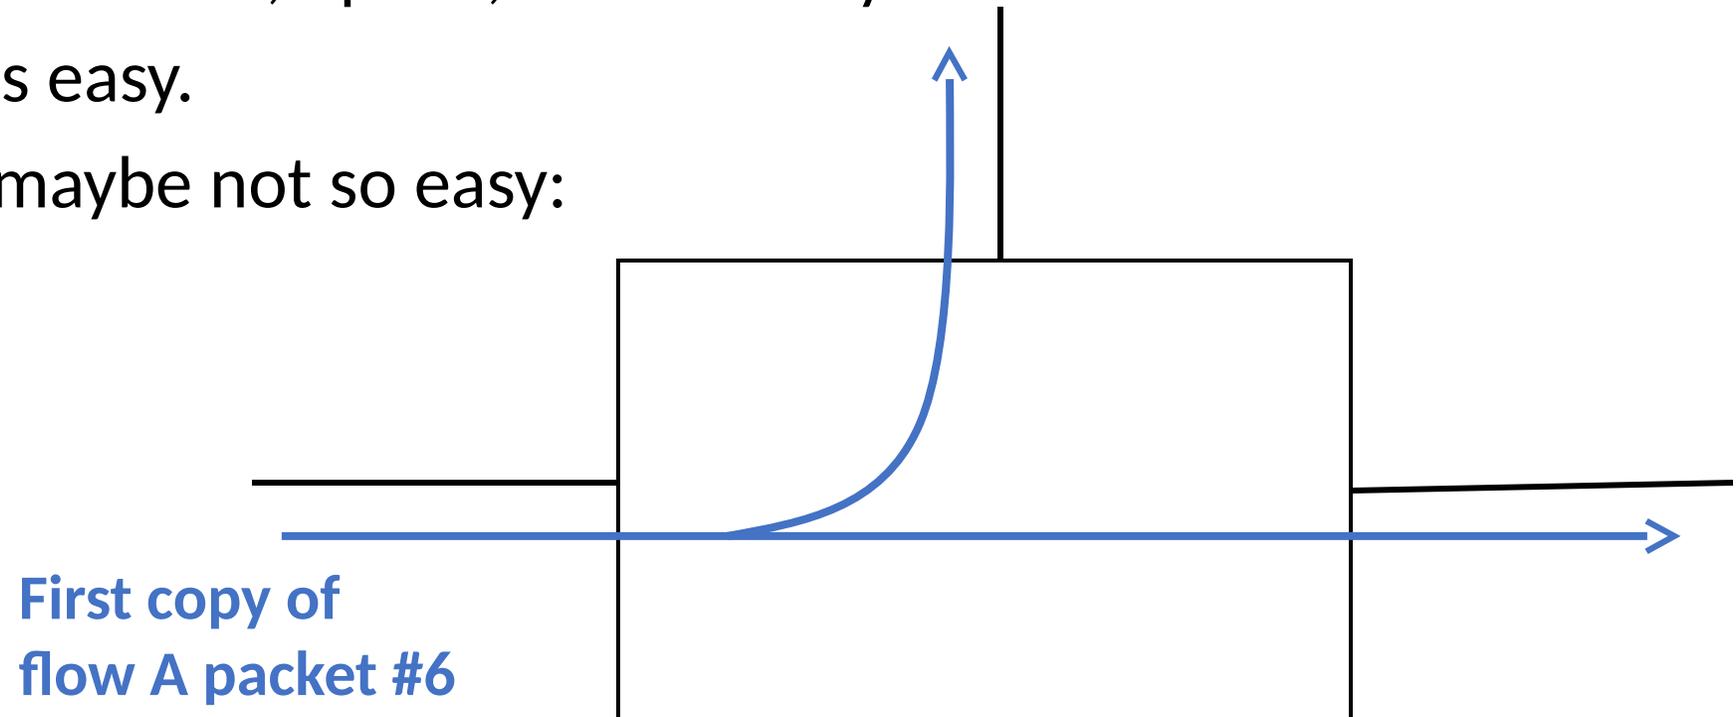
End-to-End PREF vs. “ladders”

- End-to-end PREF, where only one point (at or near the source) replicates packets, and only one point (at or near the destination) eliminates packets, makes the solution much simpler.
- However, the authors feel that there are use cases where **protection against just one failure is inadequate**. For example:
 - Factory floor: Main ring + attached rings.
 - **Planned break** in main ring is required to add/delete attached rings.
 - **Planned break in main** + **unplanned break in attached ring** = **end-to-end failure**.
 - 3 or 4 paths is a big burden on the host.



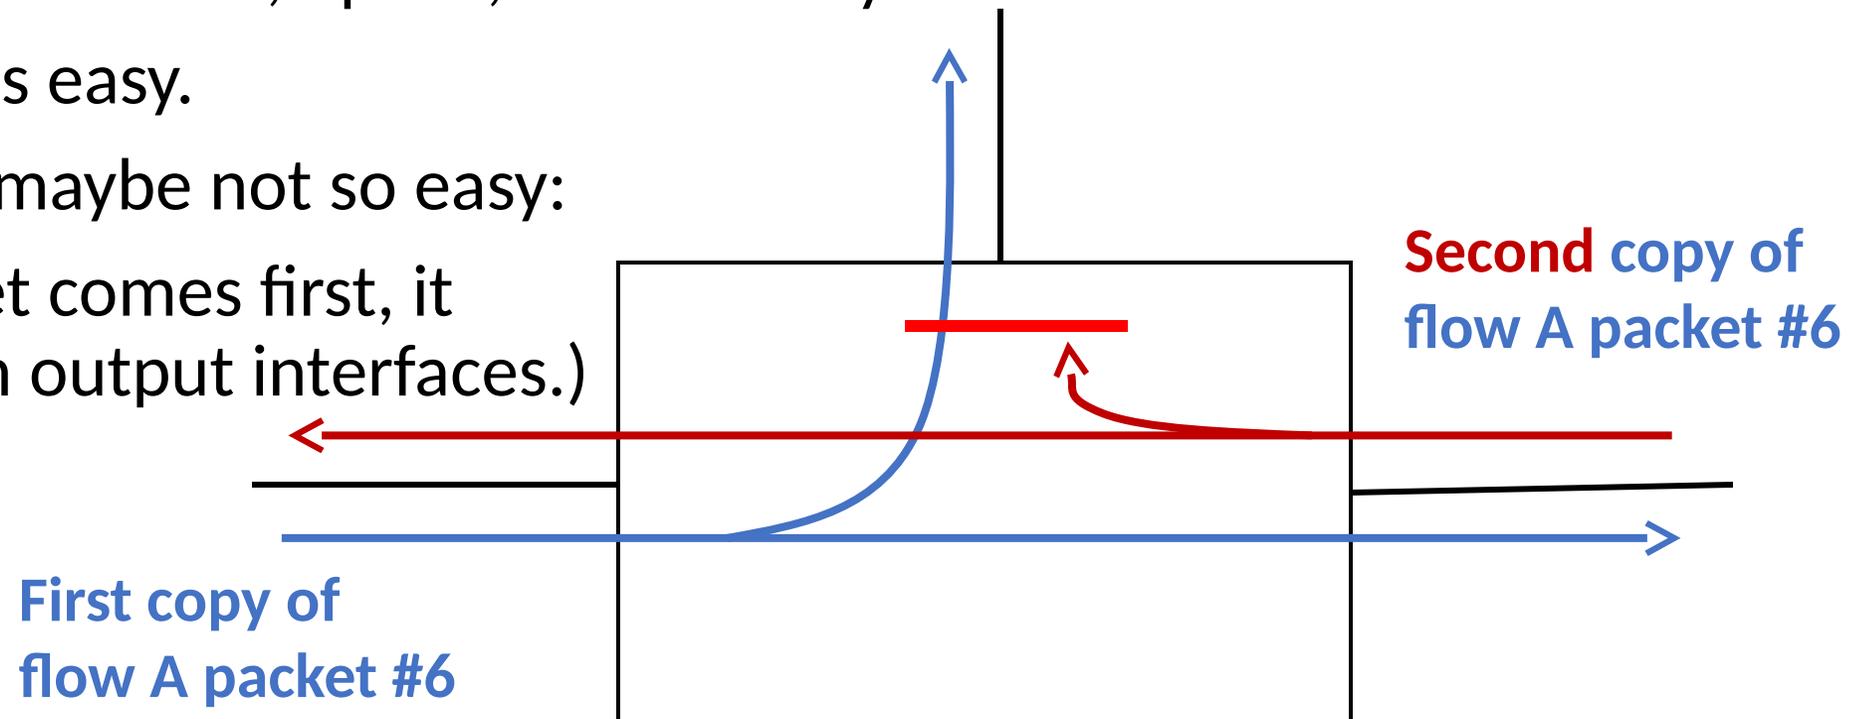
PREF in complex, federated routers

- Various aspects of packet replication and elimination are discussed in all five of the adopted DetNet drafts: architecture, problem-statement, use-cases, dp-sol, and security.
- Replication is easy.
- Elimination maybe not so easy:



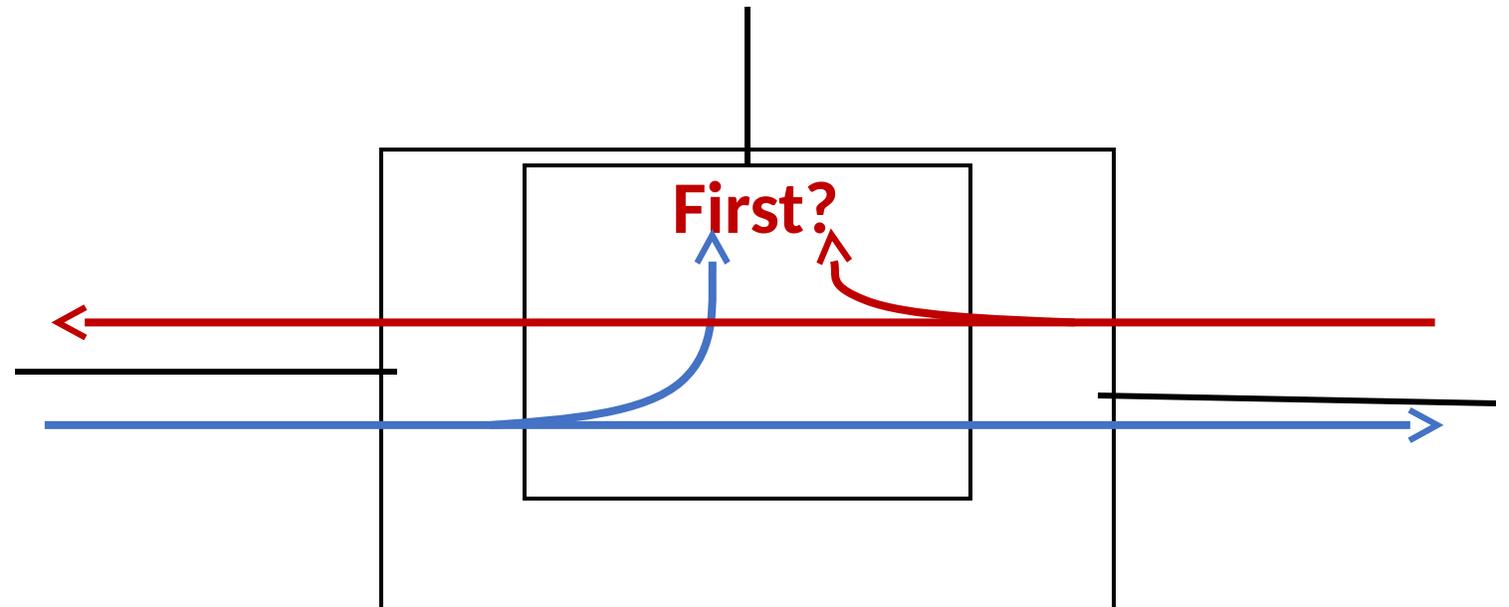
PREF in complex, federated routers

- Various aspects of packet replication and elimination are discussed in all five of the adopted DetNet drafts: architecture, problem-statement, use-cases, dp-sol, and security.
- Replication is easy.
- Elimination maybe not so easy:
- (If red packet comes first, it goes to both output interfaces.)



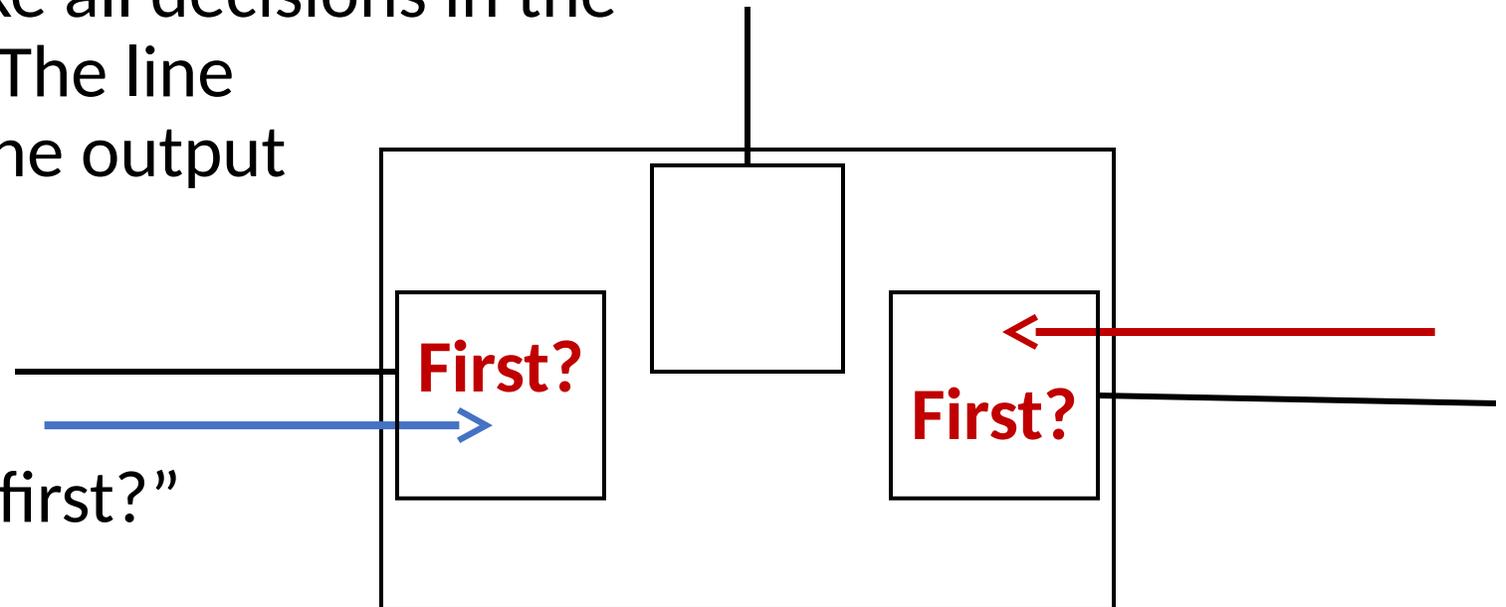
PREF in complex, federated routers

- If you have a “one-chip” router, then PRE may be easy to implement.



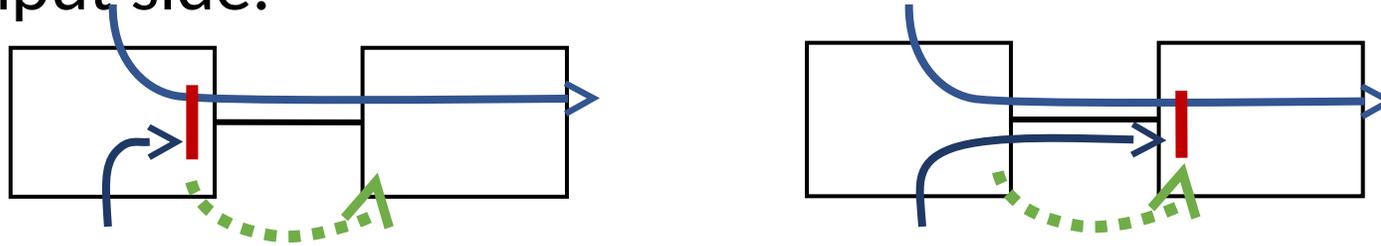
PREF in complex, federated routers

- If you have a “one-chip” router, then PRE may be easy to implement.
- If you have a “**federated**” router, with multiple line cards and an interconnect, PRE may be relatively more difficult to implement.
- It is common to make all decisions in the “ingress” line card. The line cards are stupid in the output direction.
- It is not easy to distribute the question, “Which is first?”

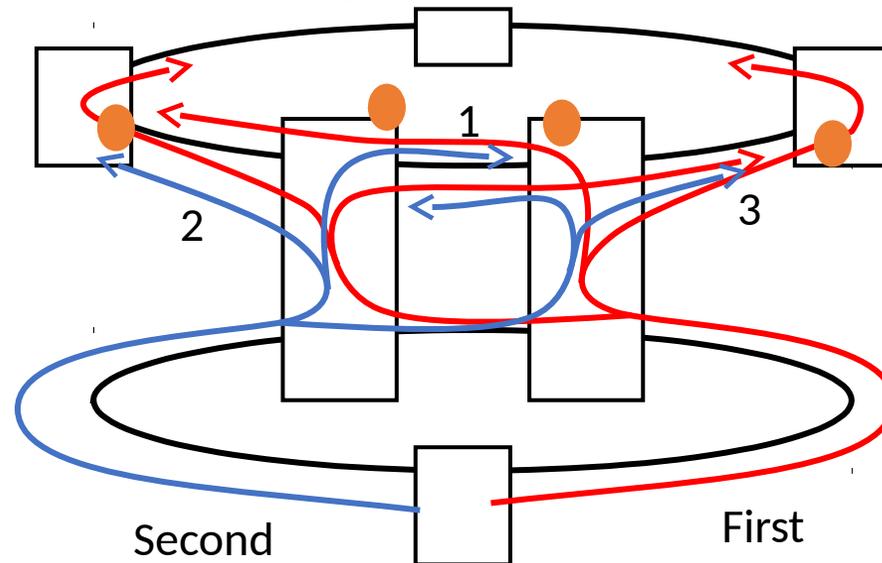


PREF in complex, federated routers

- We could allow the **packet elimination** function to be **moved** to the other end of the wire by defining an elimination function that works on the input side.



- If you have the **bandwidth** to carry all copies of all the flows that use a given link, **this can work**.
- (E.g. IEC 62439-3 “quad box” with input-only elimination uses **2x** bandwidth on link 1 and **2x** or **3x** bandwidth on links 2 and 3.)



Potential document actions

A number of potential remedies have been discussed:

1. Drop PREF as a DetNet concept.
2. Divorce PREF from bounded latency, placing PREF into separate set of architecture, data plane, and security documents.
3. Restrict PREF to one-chip routers.
4. Expand architecture and dp-sol to include input-only elimination.
5. Mention this issue in the architecture document.
6. Mention this issue in the dp-sol document.
7. Leave it to the implementer to figure out.

Further discussion

- More on the above items?
- Any specific comments from draft-ietf-detnet-dp-sol-00?
- New comments?