Provider Backbone Bridges and Congestion Notification

A way to use both PBB and CN in a Data Center Bridged network

Version 1

Norman Finn

Cisco Systems

“C-tagged I-component” is legal in 802.1ah

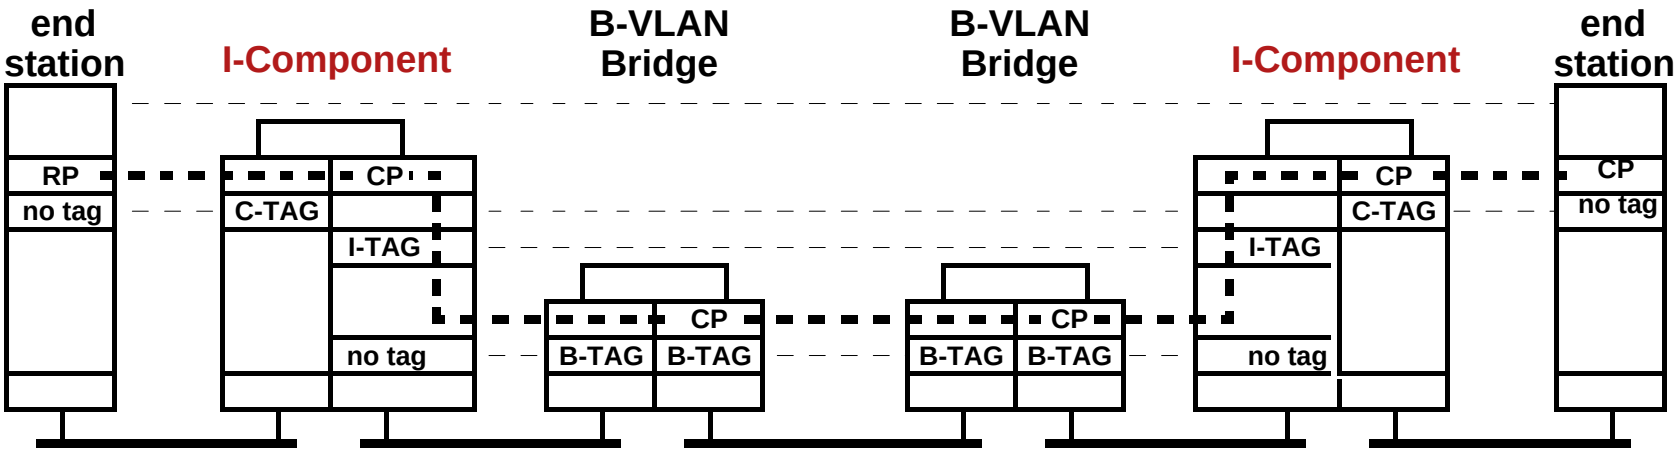

CP and RP as peers

Lower half of Figure 30–5 of P802.1Qau D2.1.

- “I-component” should be labeled “I-component + C-component”.
- C-component has one C-untagged virtual port to the I-component per C-VLAN.
- I-component has one S-untagged VIP to its PIP per S-VLAN. Overall effect is as illustrated.

Ways to combine PBB and CN

1. **Don't** specify how to do it.
2. **P802.1Qau/D2.1** says that a backbone CP encapsulates the frame header and returns the CNM to the PIP, which has a translation function that uses the encapsulated header to find the Customer MAC address of the RP, and edits the CNM.
3. **Smart backbone CP** could peek over the I-tag and generate a CNM aimed at the customer RP without help from the PIP, with very little more effort than in P802.1Qau/D2.1, now, and no added effort at the PIP.
4. **I-component RPs** can be added.

Don’t

- PBBs are likely to be of considerable use in a Data Center Bridged Network, because of the large number of virtual end stations expected, each with its own MAC address.
- While possible not to mention PBBs in the DCB context, this would be a disservice to the users of 802.1Qau.

P802.1Qau/D2.1

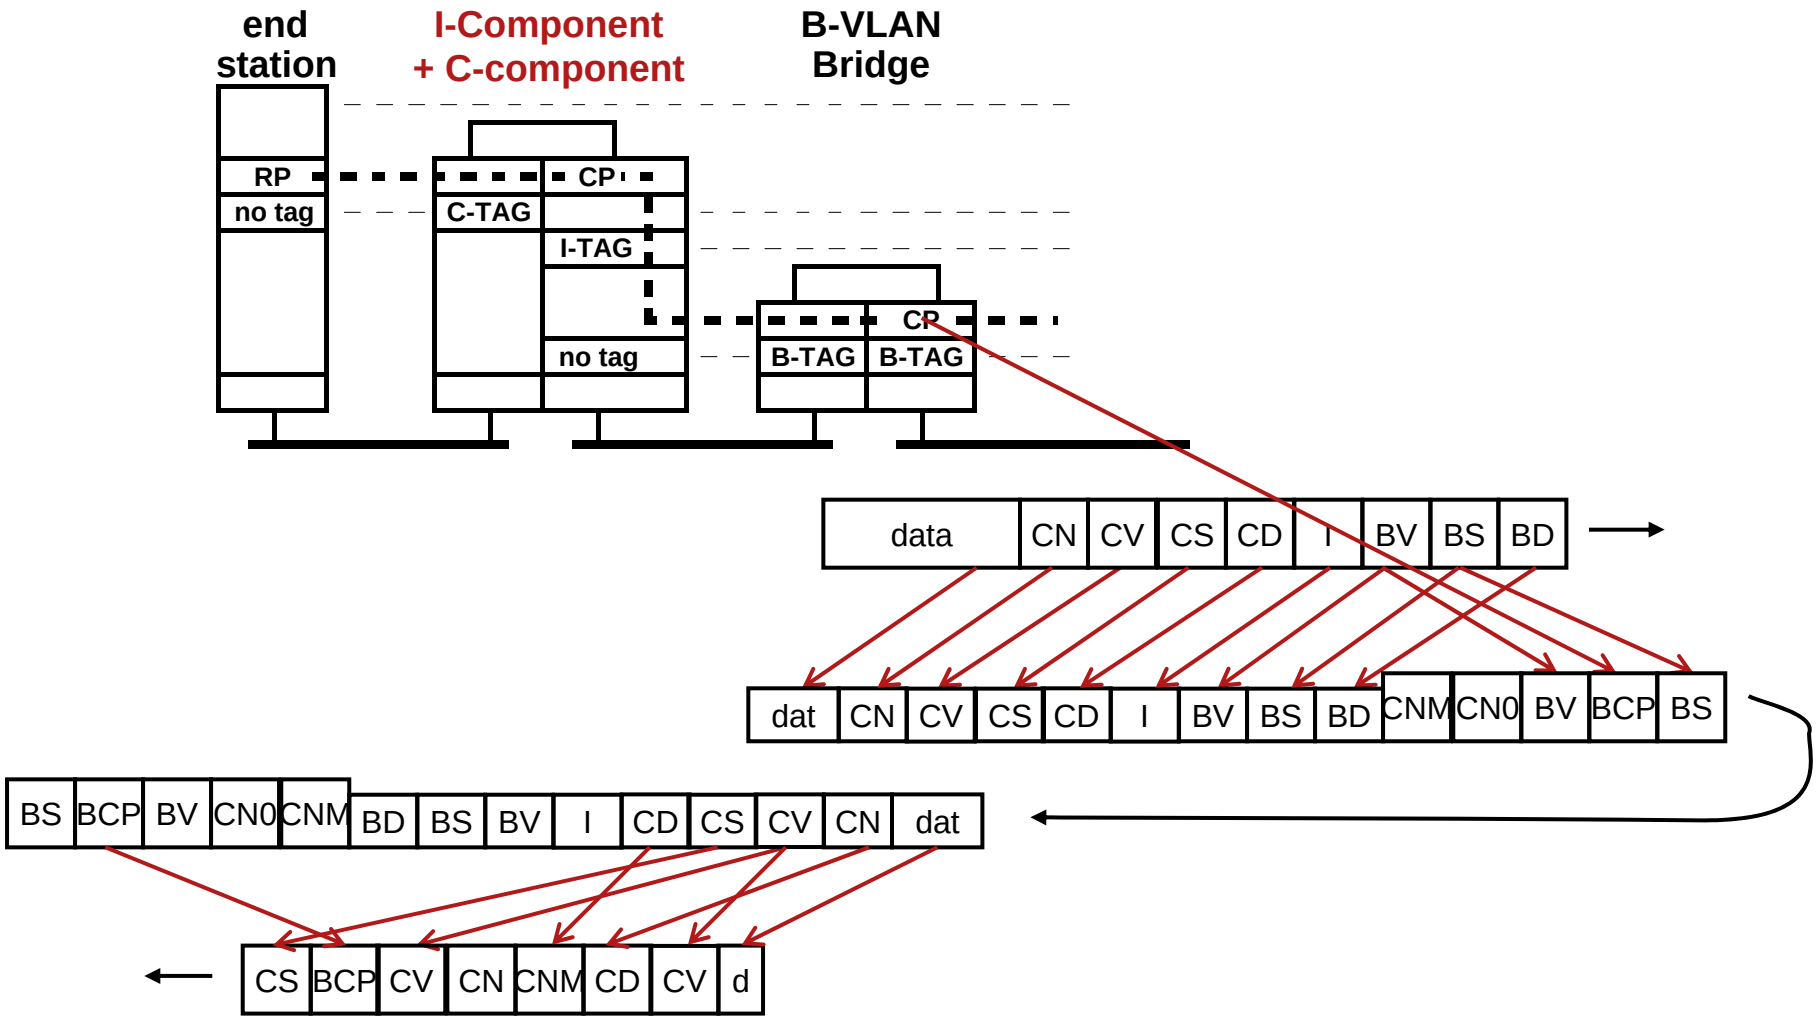

- Two transformations.

Smart backbone CP

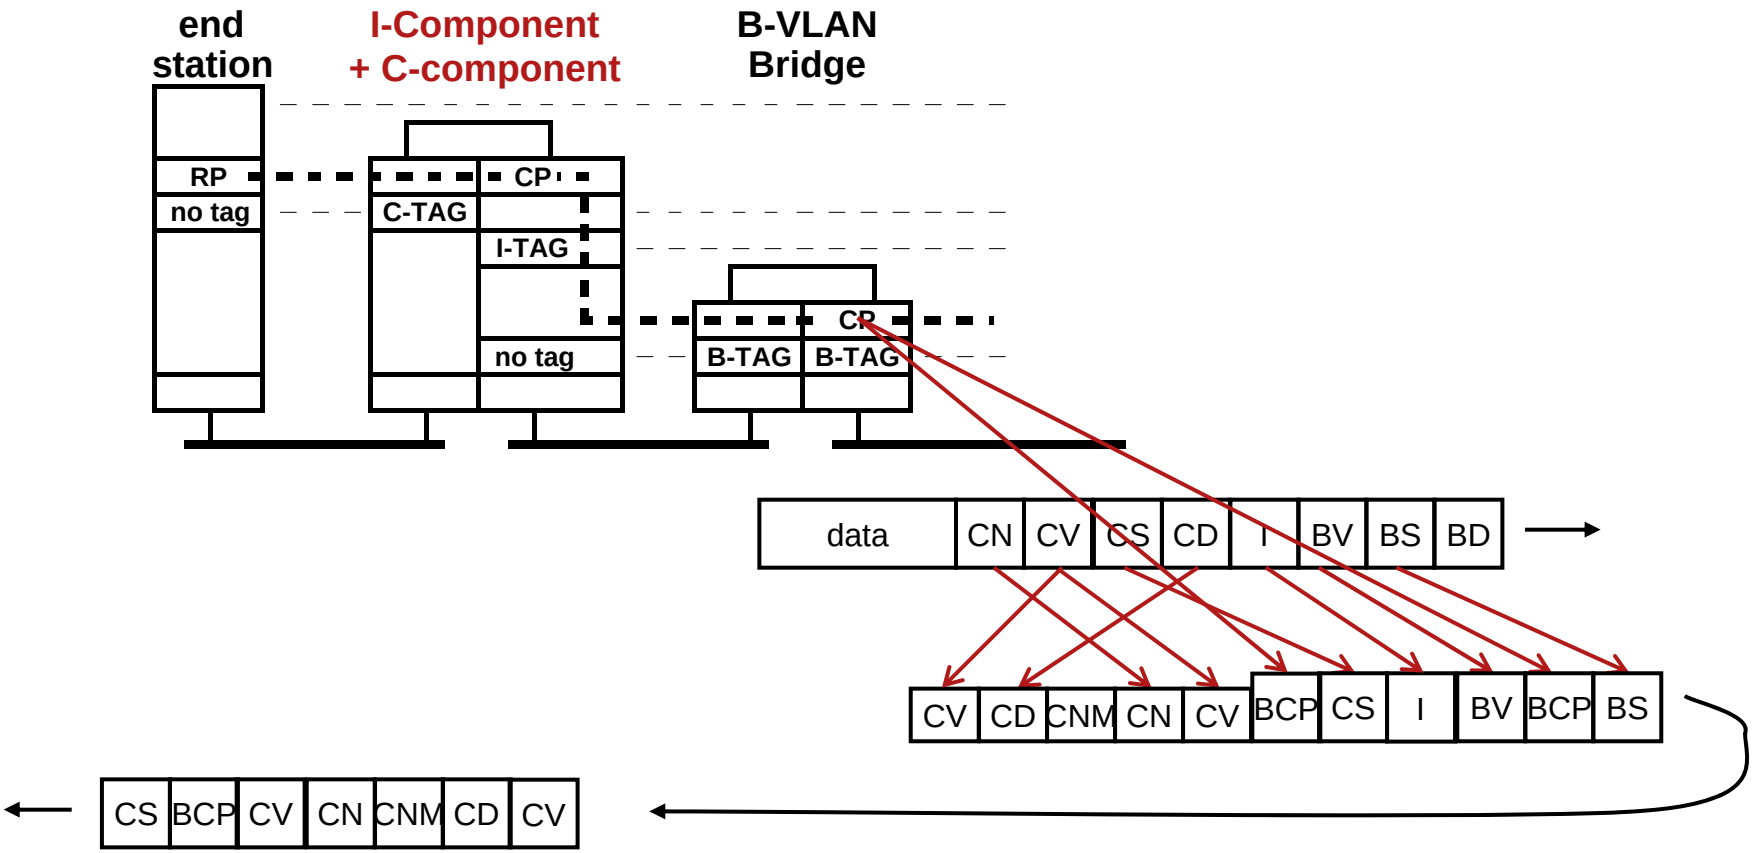

- CNM is treated as ordinary data.

There is a problem with addressing ...

- What is the appropriate source MAC address for the final CNM that goes to the original end station?
- Using the MAC address of the backbone CP (as shown in the preceding two diagrams) is incorrect, because the MAC address of the CP is not valid outside the backbone. (It could be a conflicting locally-administered address.)
- Using the MAC address of a port in the I-component is incorrect, because it is not, in fact, the source of the CNM, and there is no way for the end station to figure out what CP sent the CNM.
- **The fundamental problem is that the RP and the CP are in two different address planes.**

I-component RPs

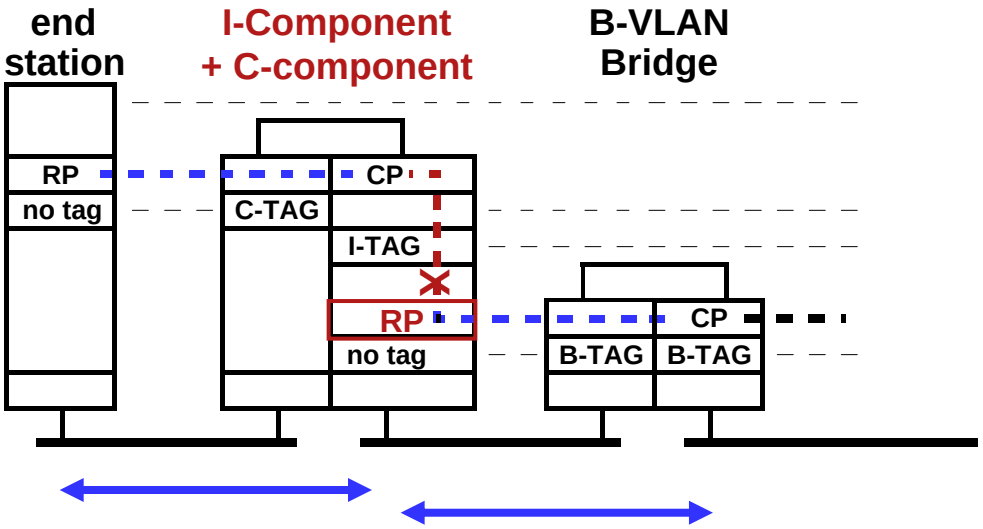

- If you install an RP in the VIP/PIP, then you eliminate the peering problem, and therefore the addressing problem.
- Depending on whether there is a network between the end station or not, the end station – I-component relationship could be priority flow control.
- The CN-tag can help the I-component RP figure out which I-component CP or customer-facing port is supplying the flow that triggered the CNM from the backbone CP.

Conclusion

- This author does not believe that the solution in P802.1Qau/D2.1 is satisfactory, because it is not actually possible to peer a backbone CP with an outside-the-backbone RP because of addressing problems.
- Intermediate RPs in the I-component can solve this problem.