# Provider Bridges and MMRP

Version 03

Stephen Haddock

May 16, 2013

# Outline

1. Review of the problem presented in March 2013 in Orlando.

2. Review of feedback from March 2013 presentation.

3. Divide the problem/solution into two elements.
   - Focus on just one of them in this presentation
4. Brainstorm for potential solutions and categorize potential solutions to each element.
5. Narrow the potential solution space (ideally to a single proposed solution).

# “Bundled” Services Are Problematic:

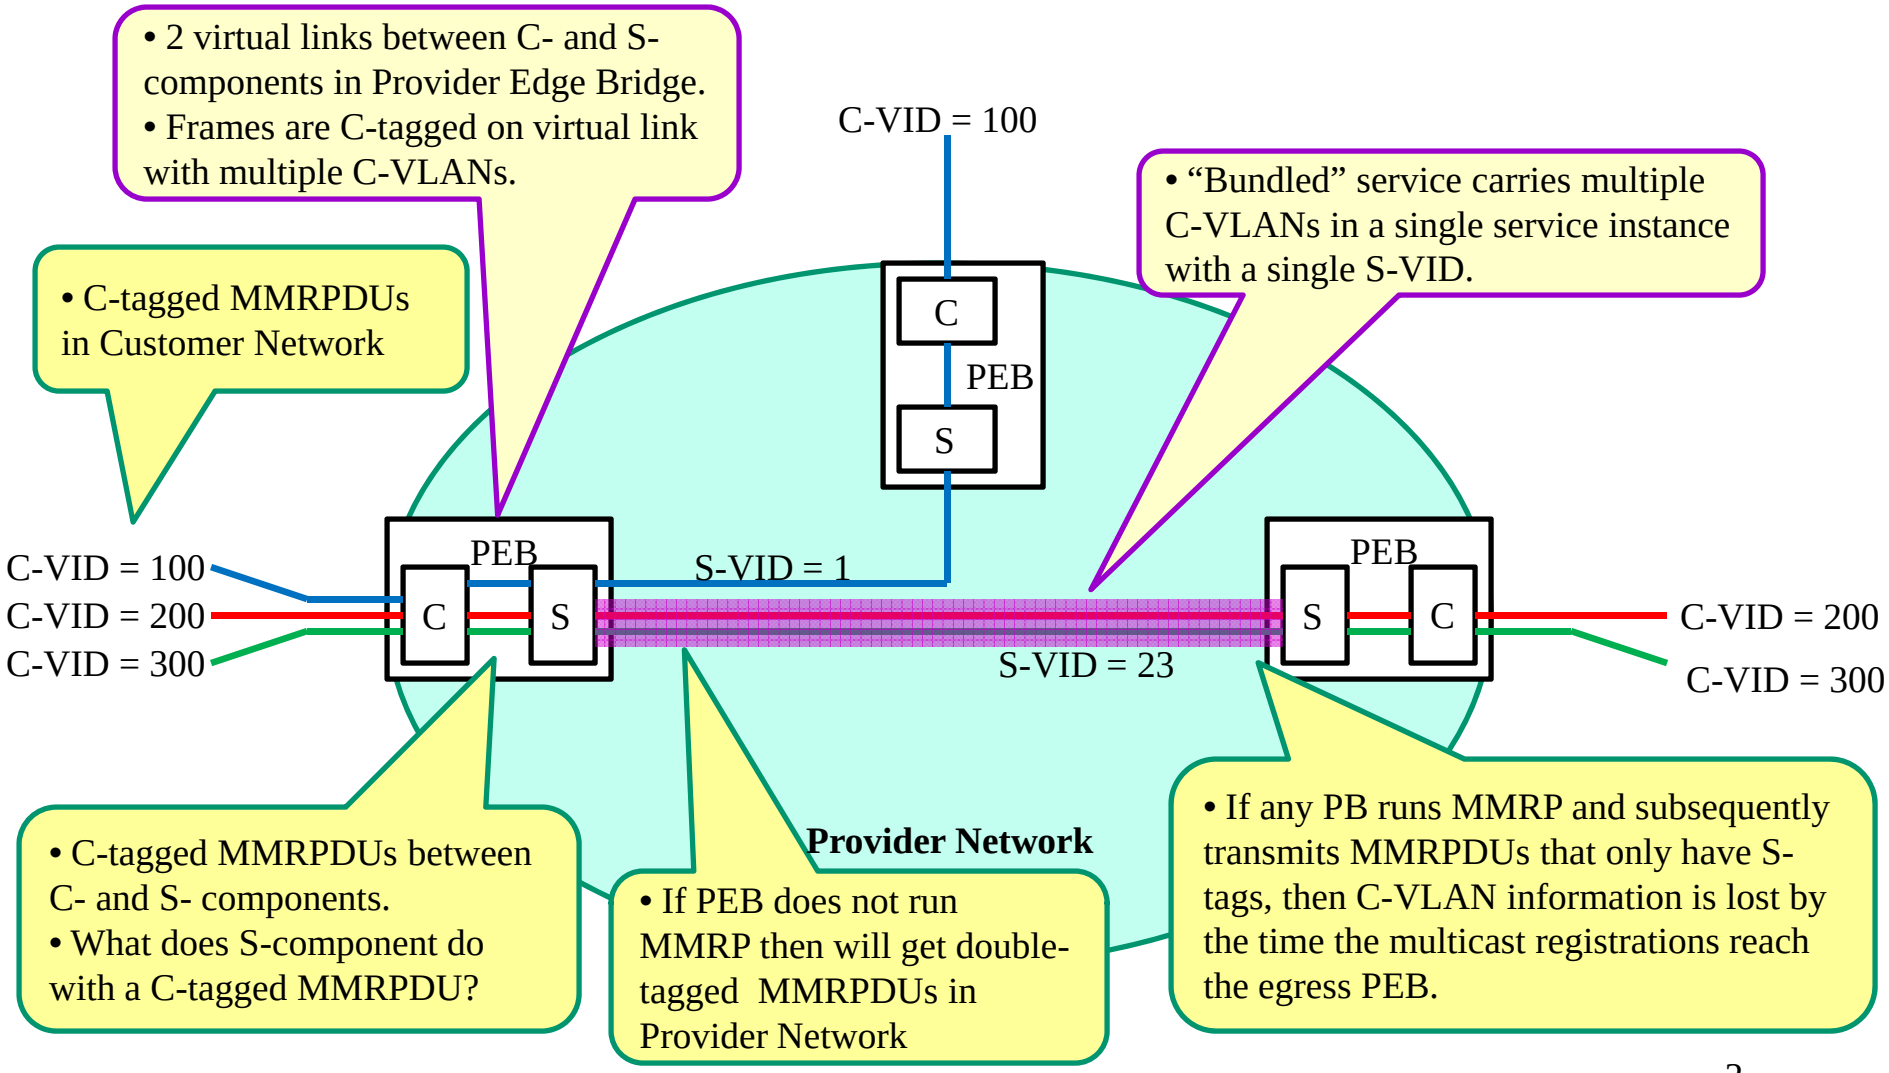

# Three promising approaches

1. Provider MMRP processes double-tagged MMRPDUs, keeps attribute state per S-VID and C-VID, transmits double-tagged MMRPDUs.
   - I think this works, but keeping state per C-VID in Provider Core is ugly.
2. Provider MMPR forwards double-tagged MMRPDUs and also processes them ignoring the C-tag, keeps state per S-VID only, transmits S-tagged MMRPDUs.
   - This might work. Seems ok for join operations, but there are (solvable?) issues with leave operations.
3. MMRP at CNPs forwards and processes C-tagged MMRPDUs, keeps state per S-VID and C-VID, transmits C-tagged MMRPDU to Customer Network and S-tagged MMRPDU toward Provider Network. MMRP in Provider Core simply forward double-tagged MMRPDUs.
   - Seems promising. Keeps Customer state isolated to edge of Provider Network at the expense of special processing at CNPs.

- None of these require changes to the PDUs, protocol operations or state machines, although number 3 may work better if Provider and Customer MMRP used different reserved DAs.

# Feedback from March 18 presentation to the Interworking Task Group

- Nobody wants Customer (i.e. per C-VLAN) state in the core of the Provider Network!!!
- Approach number 2 has issues with leave operations because it only keeps per-S-VLAN state. Can resolve by keeping per C-VLAN state but then want it only at Provider Edge so it becomes very close to approach 3.
- Approach number 3 hopefully does not need to transmit C-tagged MMRPDU toward the Customer Network (only snoops Customer MMRPDU but does not participate).
- Focus on approach #3.
- Also need to resolve the MAP Context change that occurs when have a D-bridge in a VLAN network.

# Proposal (based on approach #3)

- The basic idea:
  1. C-tagged MMRPDUs forwarded through Provider Network.
  2. CNPs (at edge of Provider Network) can snoop C-tagged MMRPDUs, keep per-C-VLAN state, summarize the customer address registrations into provider address registrations that get propagated with S-tagged (not double-tagged) MMRPDUs.
- Effectively separates the MMRP control planes in the Customer and Provider Networks
  - Much like the separation of Provider and Customer xSTP and MVRP.

# Two elements of the problem/solution

1. Propagate the C-VLAN information between customer service interfaces of the provider network.
   - Without requiring C-VLAN state in the provider core.
   - Maintaining interoperability with existing provider bridge implementations (including those that do and those that do not run MMRP) to the extent possible.
   - Effectively this means separating Customer and Provider MMRP operation into separate regions, where Customer MMRP operates only in C-VLAN context and Provider MMRP operates only in S-VLAN context.
2. Specify a means for customer service interfaces at the edge of the network to create address registrations in a S-VLAN context in the provider network from address registrations in a C-VLAN context in the customer network.
   - Effectively this provides a connection between Provider MMRP and Customer MMRP at the edges of the Provider Network.
   - Ideally CNP only snoops ingress C-MMRPDUs, but may need to snoop egress C-MMRPDUs as well, and may even need to generate C-MMRPDUs.

# Propagating C-VLAN information through Provider Network

- Lots of possible solutions!
  - Many are variations on few concepts that have similar pros and cons.
- Group potential solutions into three broad categories:

1. Use separate destination addresses for C-MMRP and S-MMRP (as we do with xSTP and MVRP).
2. Tunnel C-MMRP frames/information through Provider Network.
   - Modify C-MMRP frames (or create new frames from information in C-MMRP frames) so they are forwarded through Provider Network to all customer service interfaces.
3. Specify S-VLAN component behavior such that C-MMRP frames (with C-tags) are forwarded unchanged.
   - "Unchanged" at least meaning no frame encapsulation or destination address translation. Variations include changing or not changing the source address.

# Separate Customer/Provider MMRP DA

- Variations:

  **1a)** Allocate new address to be used for C-MMRP in Customer Networks.

  **1b)** Allocate new address to be used for S-MMRP in Provider Networks.

  (Third possibility is categorized as a variation of tunneling: Use same DA for C-MMRP in Customer Network and for S-MMRP in Provider Network, but allocate new address to be used for C-MMRP frames as they traverse a Provider Network.)

- Pros and Cons:
  - **Pro:** Follows precedent of Customer/Provider xSTP and MVRP.
  - **Con:** Maximizes interoperability issues with existing implementations. Only "works" if upgrade entire Customer or Provider Network. Severity of "not working" depends on which Network uses new address and how existing S-component implementations treat double-tagged MMRP frames.
  - **Con:** Ideally want S-MMRP to be blocked by C-component so it doesn't leak into another Provider Network (is this as important with MMRP as it is with xSTP and MVRP?) Means allocating one of three remaining addresses that are filtered by C-components but not S-components, so not a solution that could be applied to other MRP applications.

# Tunnel Customer MMRP through PBN

- Variations:

  **2a)** Create a shim to Encap/Decap C-MMRP frames at Provider Edge.

  **2b)** Create a shim to translate DA of C-MMRP frames at Provider Edge.

  **2c)** Provider Edges run Customer MMRP (maintaining state per S-VLAN per C-VLAN), but propagate C-VLAN information in MMRP frames that address only other Provider Edges.

- Pros and Cons:
  - **Pro:** Only Provider Edges require new behavior.
  - **Con:** Requires that all Provider Edges participate or else all Customer MMRP information gets lost in Provider Network. Unrealistic in multi-operator networks, so need to include ENNIs as Provider Edges that have this behavior.
  - **Pro:** Assuming all Provider Edges participate, then works even when other Provider Bridges in network would not forward C-tagged MMRPDUs.
  - **Con:** All variations require a DA that is filtered by C-components, forwarded by S-components in Provider Core, and filtered by S-components at Provider Edge. Means allocating one of three remaining addresses that are filtered by C-components but not S-components, so not a solution that could be applied to other MRP applications.

# Modify S-component to forward C-MMRP

- Variations:

**3a)** Modify "Address Filtering": Instead of filtering only on DA/S-VID, specify filtering based on DA/S-VID/C-tag/Ethertype. Only MMRP frames containing C-tags would be forwarded, all other MMRP frames filtered.

**3b)** Tunnel within S-component: Create a shim on each port that hides C-MMRP frames from address filtering by translating DA (or encapsulating) on ingress and translating back (or decapsulating) on egress.

**3c)** LTM Forwarder model: Create a shim that redirects C-MMRP frames to a new forwarding entity that figures out which ports should be egress ports and re-inserts the C-MMRP frames in the egress streams of those ports.

**3d)** Higher Layer Entity forwards C-MMRP: Stay with current model that all MMRP frames are filtered in MAC Relay. The MMRP application recognizes C-MMRP (containing C-tags) and regenerates these frames at each potential egress port active on the S-VLAN (call it "forwarding" or not as you like). Could allow the SA to be transmitted unchanged, or insist that it be changed to the egress port address (doesn't affect MMRP operation but does affect the transparency of the Provider network as observed by the Customer).

# Modified S-component continued:

- Note: All of the above variations (3a – 3d) have the same externally observable behavior.
  - If we decide that one of these is what will be specified, an implementation could do any of them.
  - There is potentially an observable difference if think one variation should change the SA while another shouldn't, but any of the variations could be specified either way.
- Pros and Cons:
  - **Pro:** Does not change how C-MMRP frames appear "on the wire". Minimizes interoperability issues with existing implementations.
  - **Pro:** Each Provider Bridge can do this unilaterally. Its does not require all bridges in network, or even all provider edges, to do it together.
  - **Pro:** Does not require assignment of a new address.
  - **Con:** Have to specify some means for a S-component to make filter/forward decision based on something more than DA/S-VID.

# Back-Up Slides

# Overview

- MMRP creates a common, flat control plane for pruning multicast trees across Customer Networks and Provider Networks.
  - This is unlike xSTP and MVRP that use different Reserved DAs to create separate, hierarchical control planes for Customer Networks and Provider Networks.
- This works when there is a one-to-one map between a C-VID in the Customer Network and an S-VID in the Provider Network.
- It only sort-of works when there is a many-to-one map of C-VIDs to S-VID.
  - It can be made to work with actions that are local to each individual bridge (don't require changes to the PDUs exchanged), but
    - The standard is silent on how to do this so it must be discovered by each implementer independently.
    - It requires saving Customer state in the core of the Provider Network.
  - There may be better solutions if we are willing to make changes to the standard.

# 1:1 mapping of C-VID to S-VID works

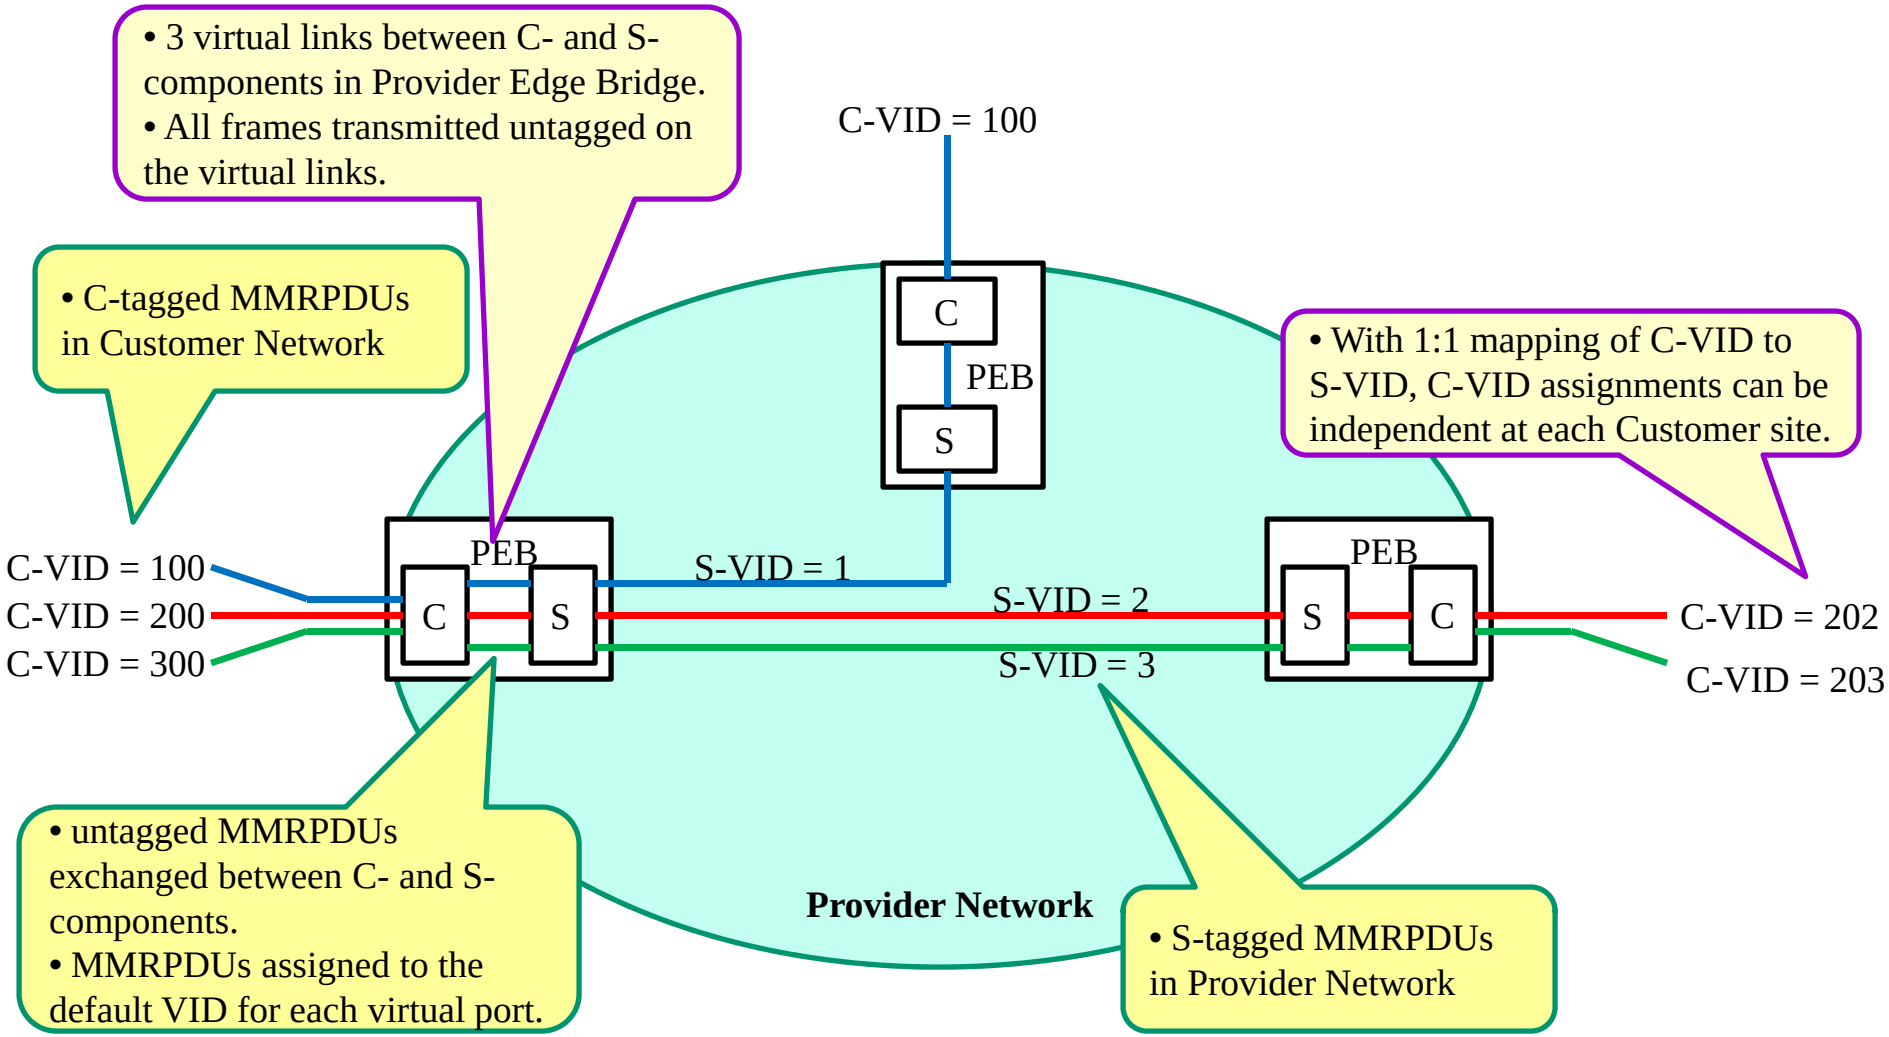

# 1. C-VLAN state in Provider Core

- Virtually no change to standard; just an “interpretation” of the MAP context.
  - A Provider Bridge supports a MAP context of “per S-VLAN” using S-tagged MMRPDUs **and** a MAP context of “per S-VLAN per C-VLAN” using double-tagged MMRPDUs.
- Advantages and Disadvantages:
  - Maintaining Customer specific state in the core of a Provider Network is a significant disadvantage.
  - When the Customer state is present, it allows a Provider Bridge to filter data frames based on both S-VID and C-VID.
    - Filtering multicast using both S-VID and C-VID was requested in ballot comments during 802.1ad. The comments were rejected because there was a strong desire not to require Provider Bridges to process both tags. Having the Customer MMRP (and MVRP?) state would allow this level of filtering without requiring it.

# 3. C-VLAN state at Provider Edge

- Customer specific state contained to the edge of PBN by having CNPs “snoop” ingress Customer MMRPDUs.
  - “Snoop” means the original frame is forwarded, but the MMRP application also processes it and keeps C-VLAN specific state. CNPs do not generate MMRPDUs(?)
  - Applies to all Customer Network Ports at Port-Based Service Interfaces and C-tagged Service Interfaces. Other ports in Provider Network (PNPs) keep only S-VLAN specific state, and generate only S-tagged MMRPDUs
- Advantages and Disadvantages
  - No Customer specific state in the core of the Provider Network.
  - If MMRP is supported anywhere in Provider Network then must be supported at all CNPs(?)

# Interoperability with current implementations

- Depends upon what current implementations do with double-tagged MMRPDUs

A. If Provider MMRP simply discards C-tagged MMRPDUs:
   - Customer and Provider MMRP is defeated: Registration information is not propagated within Provider Network or between the Customer Network sites.

B. If Provider MMRP simply forwards C-tagged MMRPDUs:
   - Customer MMRP works but Provider MMRP is defeated: Registration information is propagated between Customer Network sites, but not within the Provider Network.
     - Will data frames with the group address get through the Provider Network?

C. If Provider MMRP processes MMRPDU ignoring C-tag:
   - Customer MMRP is defeated, but Provider MMRP might work.
     - Provider MMRP “registration” operations will work, but will “leave” operations?

# What should an S-VLAN component do with MMRPDUs containing C-tags?

I. S-VLAN Component does not run MMRP

MMRPDU is forwarded on the S-VLAN

II. S-VLAN Component runs MMRP

MMRPDU is intercepted (relay filters; MMRP application processes)

A. If MMRP does not handle state per C-VLAN, options are:

1. MMRP discards the PDU
2. MMRP forwards the PDU
3. MMRP processes the PDU ignoring the C-tag, and generates MMRPDUs without C-tags.
4. MMRP forwards and processes the PDU ignoring the C-tag (and generates MMRPDUs without C-tags?)

B. If MMRP handles state per C-VLAN

1. MMRP processes the PDU including the C-tag, and generates MMRPDUs with C-tags.
2. MMEP forwards and processes the PDU including the C-tag (and generates MMRPDUs without C-tags?)

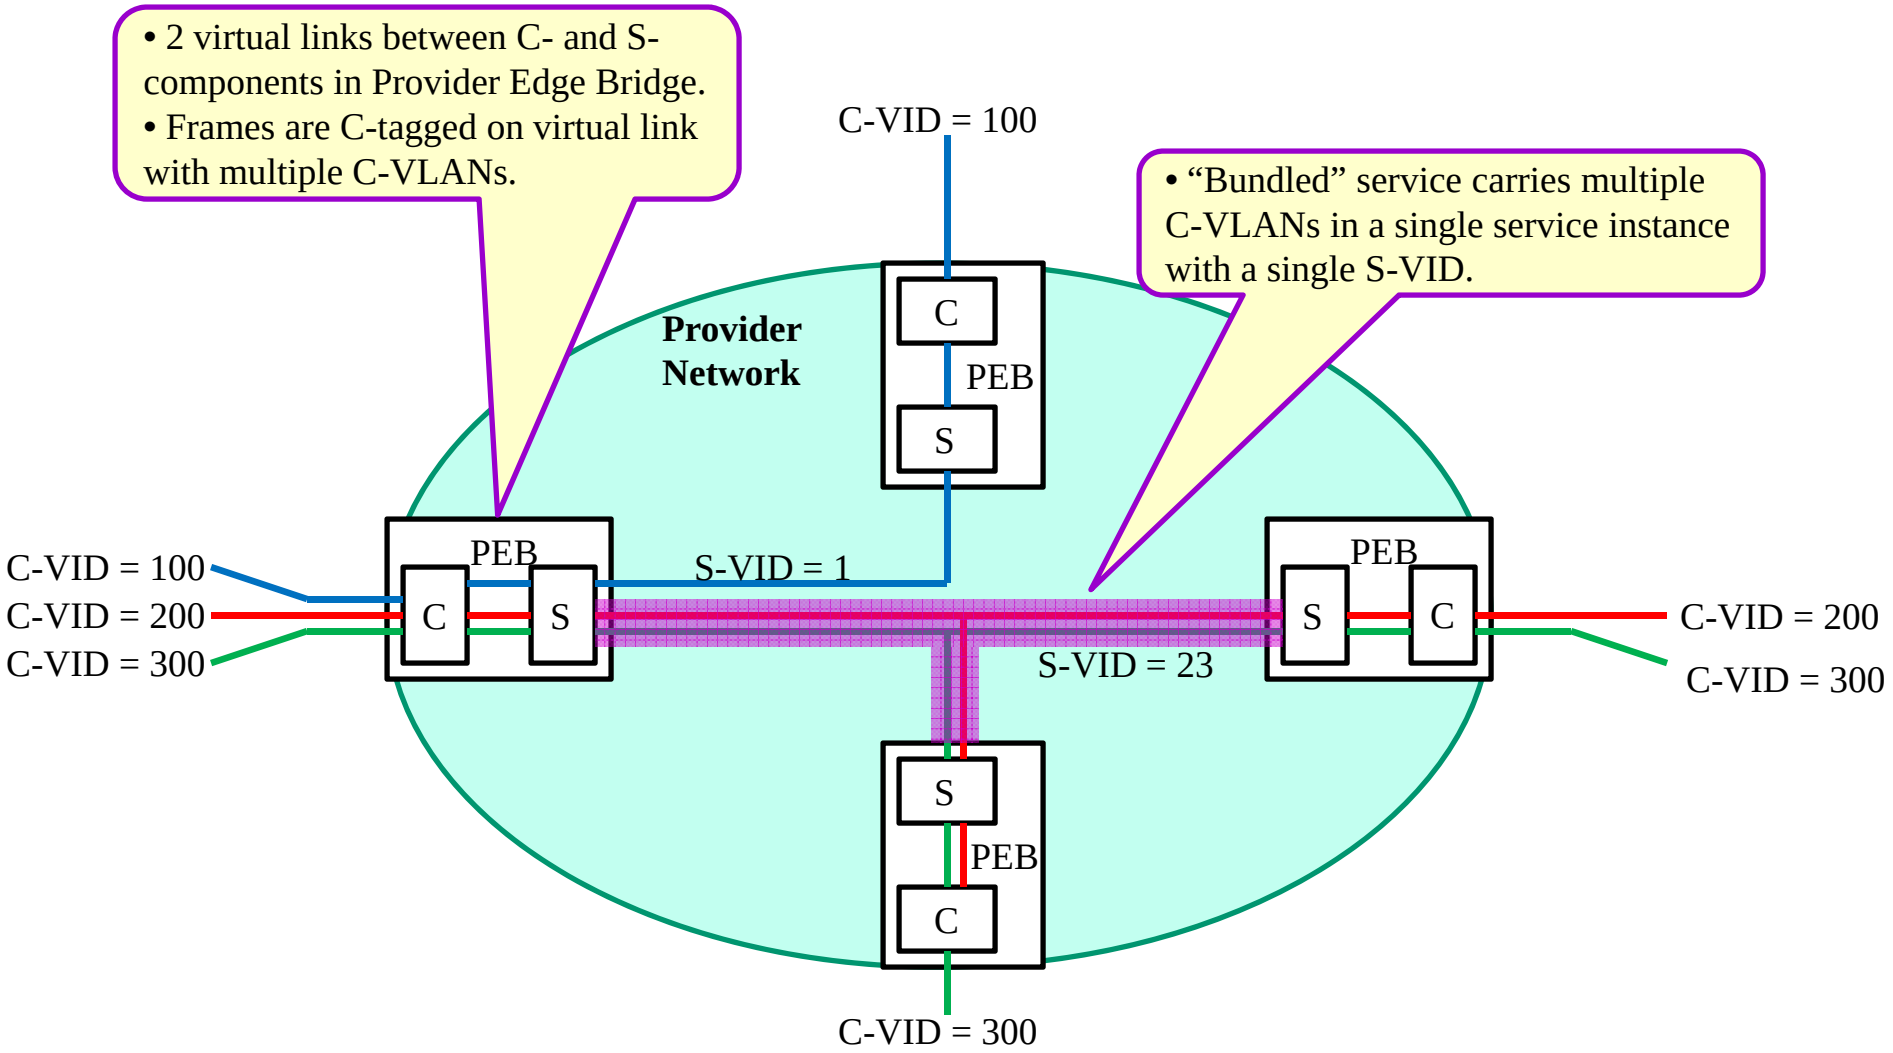
• 2 virtual links between C- and S-components in Provider Edge Bridge.
• Frames are C-tagged on virtual link with multiple C-VLANs.
• "Bundled" service carries multiple C-VLANs in a single service instance with a single S-VID.
C-VID = 100
Provider Network
C
PEB
S
PEB
C-VID = 100
C-VID = 200
C-VID = 300
C
S
S-VID = 1
S-VID = 23
S
C
PEB
C-VID = 200
C-VID = 300
S
PEB
C
C-VID = 300

# Backwards Compatibility Issues

What might existing Provider Bridge implementations do with MMRPDUs containing C-tags?

1. Nothing:
   a) The Provider Bridge does not run MMRP and therefore just forwards all MMRPDUs, including those with C-tags.
   b) The Provider Bridge runs MMRP in an S-VLAN context, but recognizes C-tagged MMRPDUs as a different context for which it does not run MMRP, and therefore just forwards all C-tagged MMRPDUs.
2. Run MMRP in per C-VLAN per S-VLAN context:

   Means the Provider Bridge has to maintain per C-VLAN per S-VLAN state at all ports. Generates double-tagged MMRPDUs. This works!
3. Provider MMPR forwards double-tagged MMRPDUs and also processes them by ignoring the C-tag and keeping state per S-VLAN only. Generates S-tagged MMRPDUs.

   This may work for Join operations but has significant problems with Leave operations. For compatibility considerations the important point is that the C-tagged MMRPDUs get forwarded.

# Backwards Compatibility Issues

What might existing Provider Bridge implementations do with MMRPDUs containing C-tags? (continued)

4. Provider MMRP processes double-tagged MMRPDUs by ignoring the C-tag and keeping state per S-VID only, but does not forward them. Generates S-tagged MMRPDUs.

   Same issues concerning Leave operations as number 3. For compatibility considerations the important point is that the C-VLAN information is not forwarded.

5. Decide double-tagged MMRPDUs are invalid frames and discard them.

Note that these alternatives, and any variations, can be put into two categories: those that propagate the C-VLAN information and those that don't.

# Thank You

.