

MEF E-NNI Support for Out-Of-Footprint UNIs

Version 2

Stephen Haddock

May 1, 2009


802.1 Interim, Pittsburgh

Table of Contents

- Metro Ethernet Forum background
- Hairpin Switching
- Virtual UNI
- Potential Provider Bridging modifications

Metro Ethernet Forum background

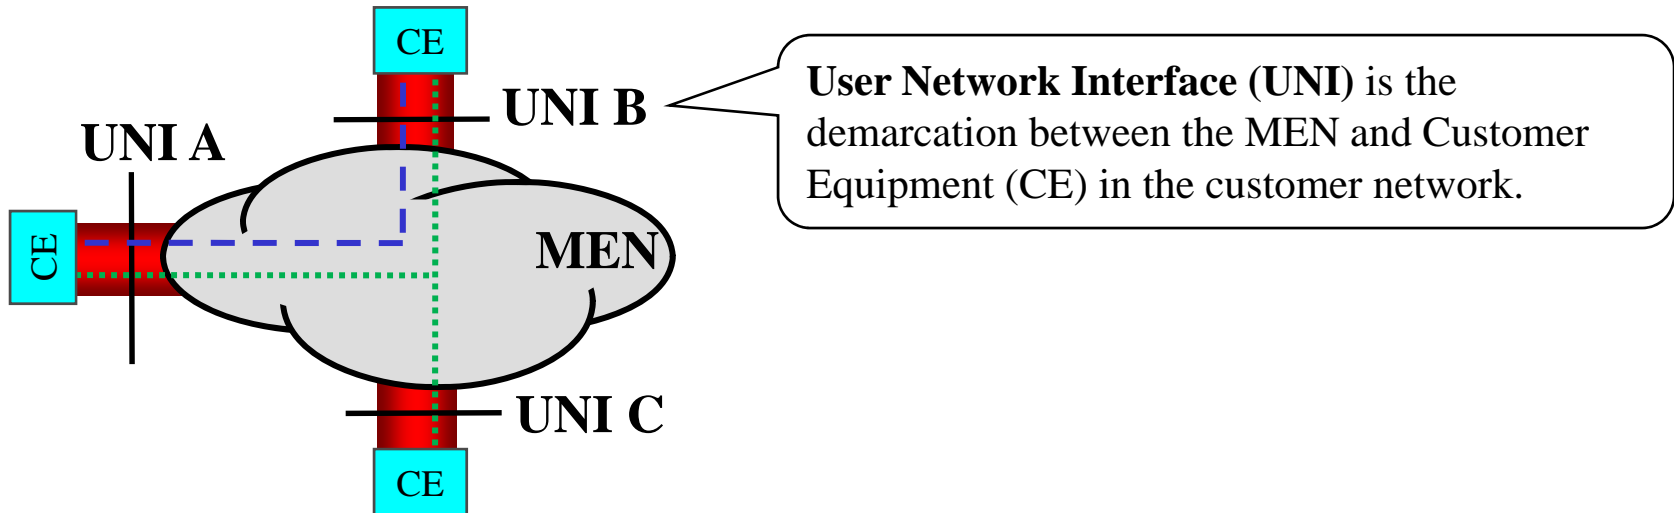
MEN, UNI, and EVC



Metro Ethernet Network (MEN) or Carrier Ethernet Network (CEN) used interchangeably for Provider or Operator owned networks that provide connectivity services to customers.

The MEF specifies the **services** provided by the MEN, the **interfaces** to the MEN, and the **attributes** that characterize the services and interfaces. The MEF does not specify the technology used within the MEN to implement the services. If the MEN is implemented with 802.1 technology then the MEN is equivalent to a Provider Bridged Network (or possibly Provider Backbone Bridged Network) in 802.1 terminology.


MEN, UNI, and EVC



The physical medium at the UNI is a full-duplex 802.3 LAN. The frame format at the UNI is an untagged or C-tagged Ethernet frame. UNIs may be:

- “**All-to-One-Bundling**” where all customer frames are mapped to a single service instance. If the MEN uses 802.1ad technology then this is a **Port-based service interface**.
- “**Service Multiplexing**” where customer frames are mapped to a service instance (or filtered) based on the C-VID. If the MEN uses 802.1ad technology then this is a **C-tagged service interface**.

MEN, UNI, and EVC




EVCs may be point-to-point, multipoint-to-multipoint, or rooted-multipoint.

If the MEN uses 802.1ad technology then the EVC is a **service instance** implemented by an S-VLAN and identified by an S-VID.

Service Provider, Operators, and E-NNI


External Network Network Interface (E-NNI) is a reference point representing the boundary between two Operator MENs that are operated as separate administrative domains.



In the MEF model there is a Service Provider responsible for the end-to-end service offered to a customer. The Service Provider may contract with one or more Operators, each responsible for a MEN, to realize the service. The Service Provider may (or may not) be the same business entity as one of the Operators.

E-NNI and OVC

Operator Virtual Connection (OVC) is an association of external interfaces (UNIs or E-NNIs) of a single Operator MEN such that any ingress customer frame mapped to an OVC at one interface may be delivered to any or all other interfaces that have mappings to the same OVC.




The physical medium at the E-NNI is a full-duplex 802.3 LAN. The frame format at the E-NNI is an **S-tagged** 802.3 frame. The S-VID is (roughly speaking) a service identifier that allows the operator on either side of the E-NNI to map frames to the appropriate OVC End Point.

An EVC is an end-to-end (UNI-to-UNI) service instance. An OVC is a local (to one Operator MEN) service instance. In many cases there is a one-to-one relationship within a given Operator MEN between an OVC and an EVC, however this is not true in all cases.


Hairpin Switching at E-NNI

OOF-UNIs



In this case the Service Provider is the same as Operator A, and is providing a multipoint EVC to a customer with several sites. Two of the customer sites, UNI D and UNI E, are “Out-of-Footprint” (OOF) meaning they are not reachable by the Operator A MEN. Therefore the SP obtains a point-to-point OVC in Operator B’s MEN and in Operator C’s MEN to connect each OOF-UNI to an E-NNI.

OOF-UNIs and Hairpin Switching



Same scenario except each OOF-UNI is reachable through MEN B. Therefore the SP obtains two point-to-point OVCs in Operator B's MEN to connect each OOF-UNI to the E-NNI. In theory the SP could obtain a single multipoint OVC in Operator B's MEN, however for business purposes the SP does not want to disclose or delegate to Operator B any of the details of the service being provided to the customer.

In MEN B these are two completely unrelated OVCs. At the E-NNI frames to/from each OOF-UNI are identified by different S-VID values. But in MEN A these frames map to the same OVC. A particularly problematic case is where a frame from UNI E destined to UNI D needs to be received by MEN A at the physical port that is the E-NNI, and transmitted on the same physical port but with a different S-VID. This is “**hairpin switching**.” To make this operate exactly as the previous case where frames to/from UNI E and UNI D came into MEN a on different ports, MEN A needs to use the S-VID value to create different **virtual ports** at the E-NNI.

Comparison to EVB scenario

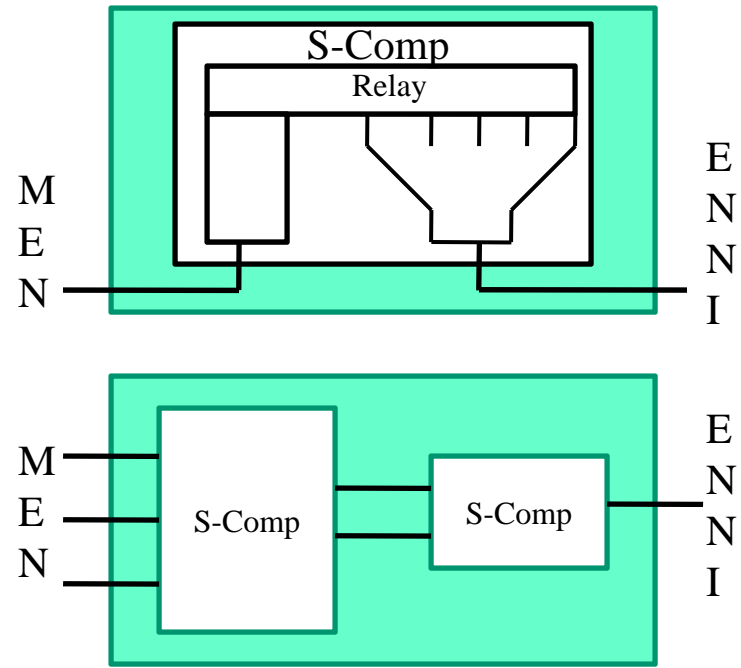
- Similarities to the Edge Virtual Bridging scenarios being discussed in the Data Center Bridging Task Group:
 1. Both call for a solution where a bridge forwards packets between two entities that are connected to the bridge through the same physical port.
- Differences from EVB:
 1. Use of the S-tag as the virtual port selector.
 - a) In the MEF E-NNI environment, the presence of the S-tag is a given making it the logical choice.
 - b) Using an additional tag would not provide any benefit, and would unnecessarily require the OOF network to treat OVCs that go to a virtual port at the E-NNI differently from other OVCs.
 2. No local switching, including multicast replication.

In the Data Center Bridging EVB environment there is a desire not to have the bridge with the virtual ports do multicast replication for the virtual ports (and thus send multiple copies of the multicast packet on the same physical link). Rather there is a desire to devise a system that allows multicast replication to be done closer to the other end of the virtual link. In the MEF environment this is explicitly not desirable, because to do so is antithetical to the premise that the OOF network operator knows nothing about the details of the service, including the full connectivity.
 3. No “Port Expander”.

In the DCB EVB environment there is some kind of device that multiplexes traffic from virtual interfaces on to a single physical link connecting to the bridge implementing the virtual ports. In the MEF scenario there is no such device. The OOF network does not distinguish UNIs that connect to virtual ports at the E-NNI from UNIs that do not. The OOF network is completely unaware of the port virtualization.


Possible solutions: Hairpin Switching

1. Expand VLAN-tagging shim (6.9) so that it may present multiple Virtual Bridge Ports to the MAC Relay, similar to the Provider Instance Port (6.10) for Backbone Edge Bridges.
2. Define a new type of Provider Edge Bridge that is similar to the current PEB except uses S-Components to demultiplex based on S-VID where the current PEB uses C-components to demultiplex based on C-VID.
3. Expand VLAN-tagging shim (6.9) so the S-VID translation table supports many-to-one S-VID translation, and add functionality for local switching and multicast replication.



Virtual UNI at E-NNI

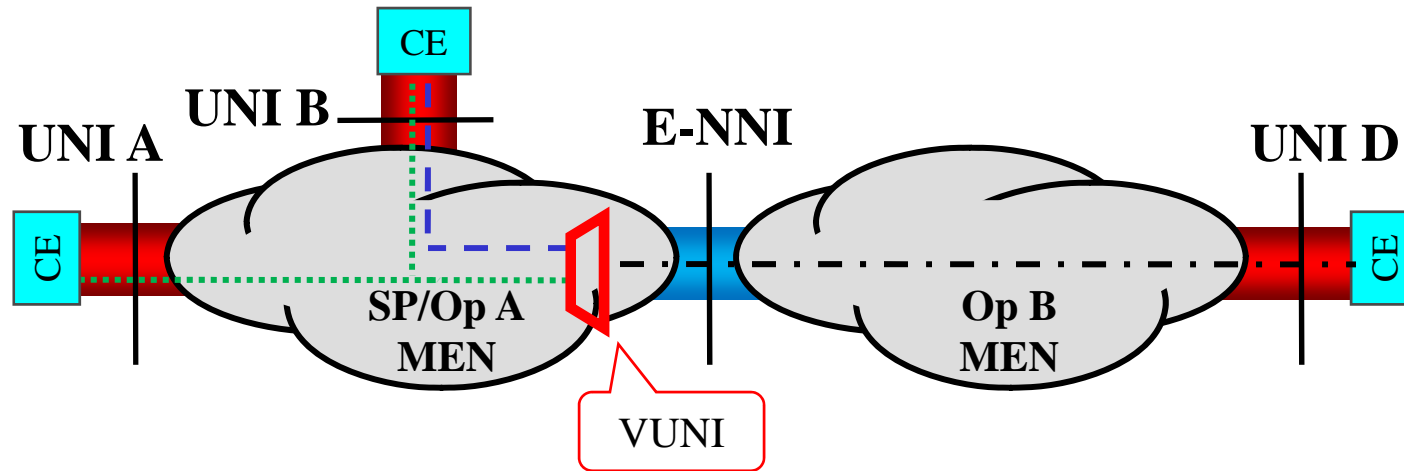
Service Multiplexing OOF-UNI



In this case the Service Provider is providing two (or more) EVCs to UNI D which is Out-of-Footprint. The straight-forward solution is for the SP to obtain an OVC per EVC from Operator B and have Operator B perform the service multiplexing functionality as shown. The SP finds this solution undesirable for several reasons:

1. Requires obtaining multiple OVCs from Operator B.
2. Requires disclosing and delegating details of the service being provided to Operator B.
3. Requires coordinating with Operator B whenever there is a change in the number of EVCs or attributes of the EVCs being provided at UNI D.

Virtual UNI (VUNI)

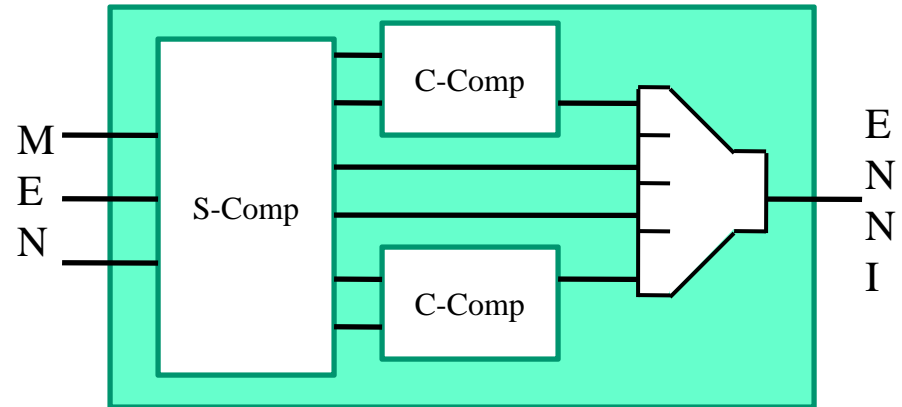


Alternatively the Service Provider obtains a single OVC from Operator B to transport all frames between UNI D and the E-NNI. The SP creates a Virtual UNI (VUNI) on the MEN A side of the E-NNI that performs the service multiplexing and other UNI functions.

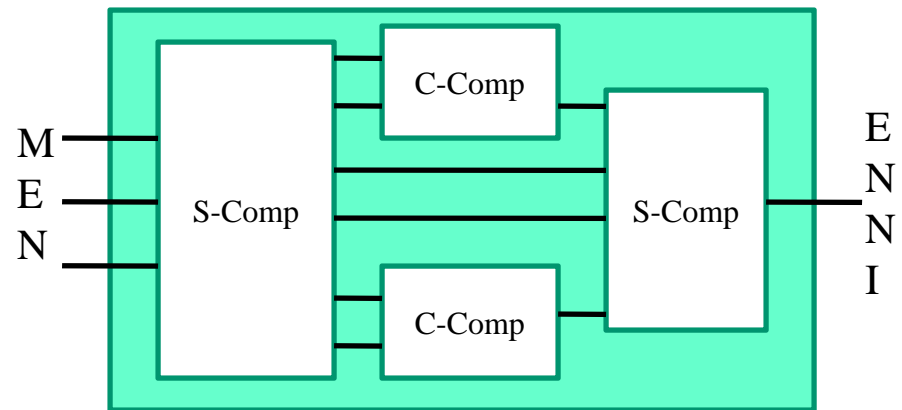
If MEN uses 802.1 technology, implementing the VUNI at the E-NNI requires something that will first demultiplex frames received at the E-NNI based on the S-VID, and then perform the normal Provider Edge Bridge function of mapping frames to EVCs based on the C-VID. This could be done with an S-VLAN Bridge at the E-NNI connected to a separate PEB, but it is desirable to provide the functionality in a single device.

Possible solutions: VUNI

1. Create a demultiplexing interface stack that is not part of a bridge component (i.e. not attached to a single MAC Relay). Each “Virtual Port” may attach to Bridge Ports on separate bridge components.



2. Define a new type of Provider Edge Bridge that is similar to the current PEB except adds an S-component to demultiplex based on the E-NNI S-VID with internal connections to Customer Edge Ports (for VUNI) or Customer Network Ports (for “normal” traffic and hairpin switching).



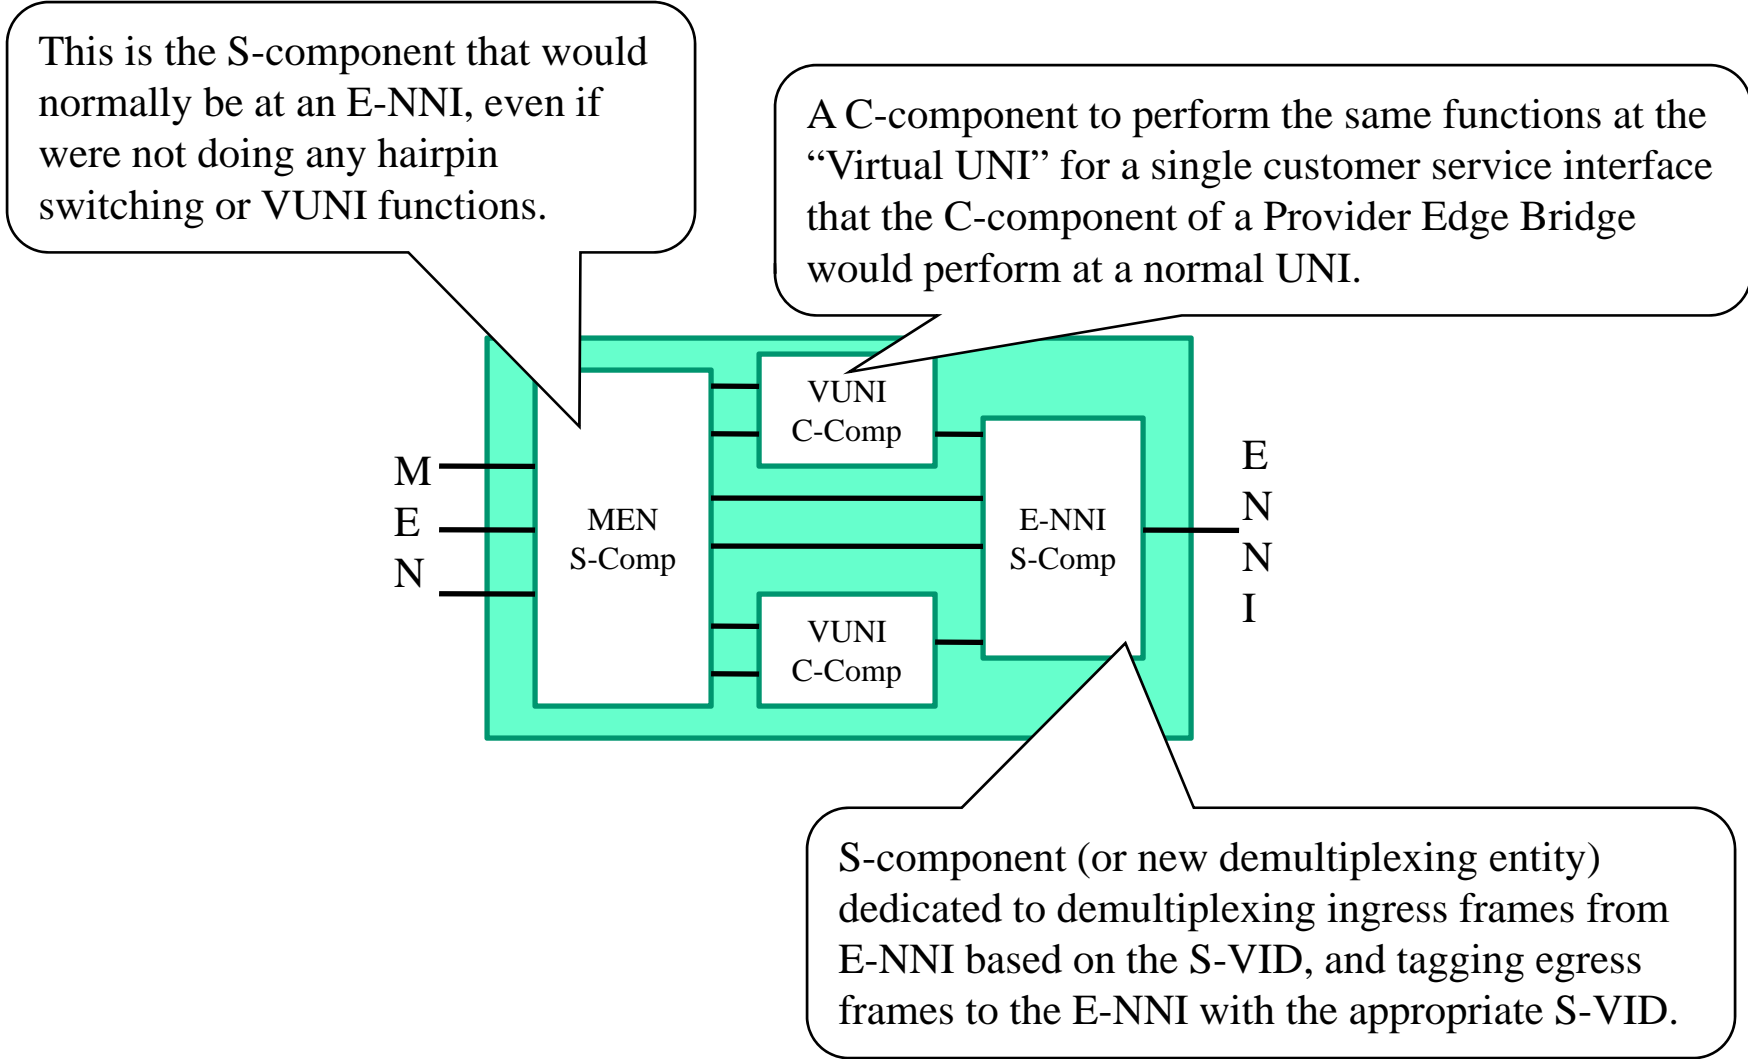
Potential Provider Bridging modifications

Evaluating the possible Hairpin Switching
and VUNI solutions

Narrowing the solution space

- Of the potential Hairpin solutions, #3 is least promising.
 - Even without the VUNI functionality, achieving hairpin switching, multicast replication, and many-to-one S-VID translation makes the shim very complex.
 - With VUNI functionality the shim would need to replicate all functions of a C-component in a Provider Edge Bridge, including the control protocols.
- Hairpin solution #1 can be generalized to VUNI #1.
 - Hairpin solution #2 cannot accommodate the VUNI functionality, however separating the multiplexer from a specific bridge component (as in VUNI solution #2) can accommodate both Hairpin switching and the VUNI.
- VUNI solutions #1 and #2 are very similar
 - The distinction comes down to determining how a newly defined S-VID multiplexer would differ from a full S-component.
 - For packet forwarding the S-component would be configured to behave very much like (exactly like?) the multiplexer.
 - The primary difference between a full S-component and a multiplexer is likely to be in things like control protocols and CFM.

A more detailed look



Common Elements

- Both solutions demultiplex frames arriving at the E-NNI to different internal links, and thus to different Bridge Ports on the MEN S-component, based on the E-NNI S-VID.
 1. This allows the normal relay function of the MEN S-component to perform both hairpin switching and multicast replication without echoing any frames back to their source.
 2. This allows frames destined for a VUNI to be directed to a C-component.
 3. This allows normal operation of control protocols (e.g. RSTP, MVRP) on the MEN S-component.

Goal

- **Enhance the specification of the Provider Bridging S-tagged service interface such that ...**
 - **when this interface exists between two independently administered Provider Bridged Networks, and ...**
 - **all traffic to/from a customer location attached to one (access) PBN is identified on the interface by a single S-VLAN identifier, then ...**
 - **the Provider Bridge in the other (primary service provider) PBN can treat that customer traffic as if the customer were directly attached using a Port-based or C-tagged service interface.**

