E-NNI Redundancy
(Considerations and Musings)

Stephen Haddock

November 19, 2009

802.1 Plenary, Atlanta

Introduction

- This presentation is kind of a random walk considering aspects of providing redundancy at external interfaces (UNI, ENNI) and implications on the overall network behavior and service delivery.
- The goal is to try to zero in on cases where it is OK for the redundancy design choices to determine the network/service behavior, and when the desired network/service behavior needs to drive project objectives which then drive the redundancy design choices.

Consideration 1:
An E-Line service traversing an ENNI with redundancy

E-Line service over an ENNI

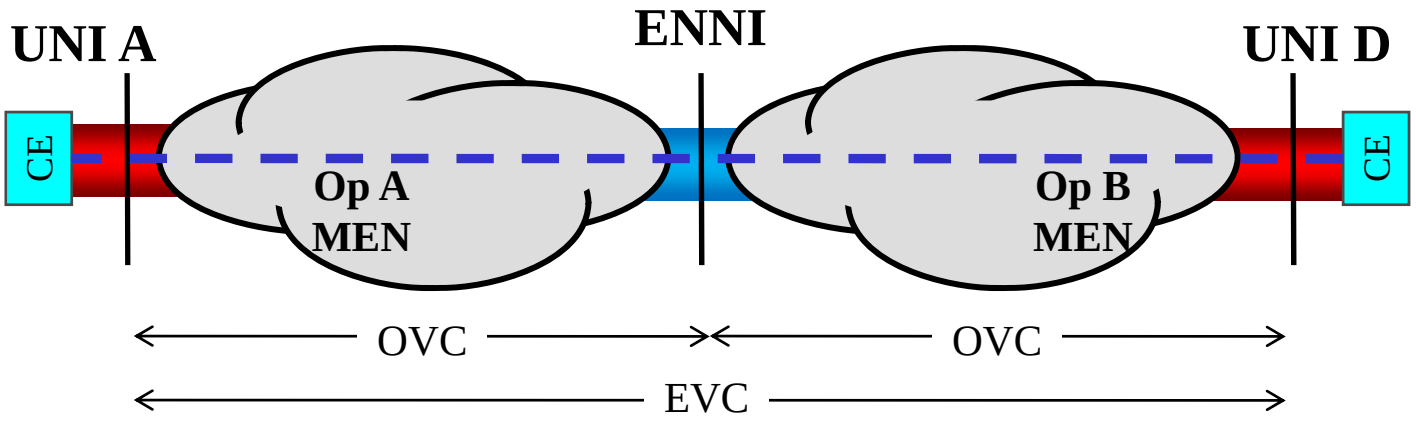

An E-Line service is defined as being supported by a point-to-point EVC (the EVC being the entire service from UNI to UNI). The MEF has not yet defined how specific service types are supported at the OVC level although that work started about two weeks ago (the OVC being the portion of the EVC that is in a single operator's network). The MEF has also not defined or specified any redundancy mechanisms at the ENNI except Link Aggregation, but there is interest in having other redundancy mechanisms.

The assumption so far has been that an E-Line service would be supported by a point-to-point OVC, but that has not been examined in conjunction with potential redundancy mechanisms.

E-Line service over a redundant ENNI

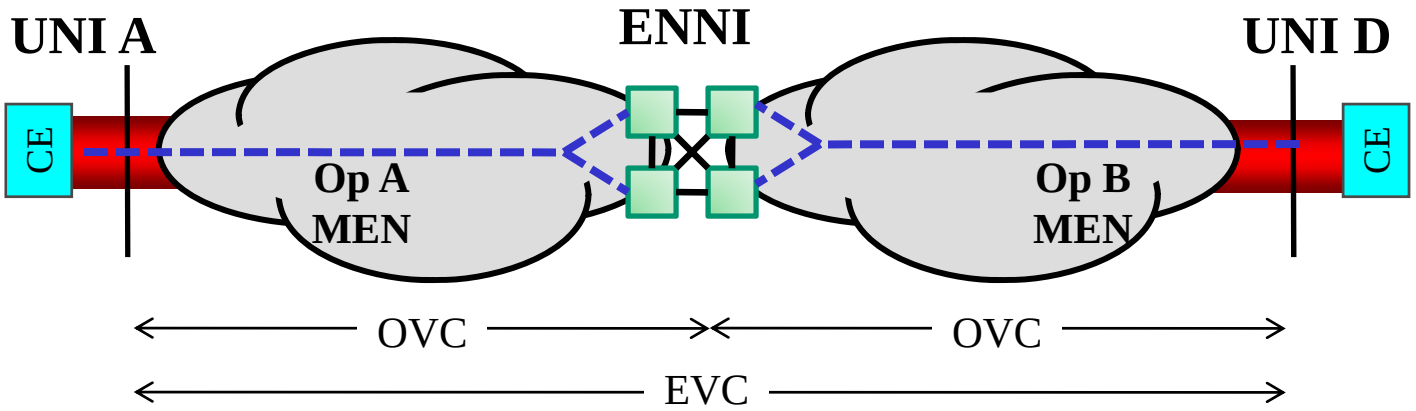

Does a redundant ENNI necessarily imply the E-line service be supported by a multipoint OVC?

There are advantages to keeping the OVC point-to-point. Moderate advantages with Provider Bridge technology, but PBB incorporates significant optimizations for point-to-point Backbone Service Instances. (PBB-TE wouldn't support a multipoint EVC, but PBB-TE has it's own protection mechanisms.)

Is it possible to make it appear that there is a single OVC end-point at the ENNI, so it appears to be a point-to-point OVC?

One possibility is to connect the OVC to only one of the ENNI nodes and "hide" the local connection between the nodes. (This assumes a local connection between nodes.)

Another is to connect to the OVC to both ENNI nodes, but make it appear that the two nodes are a single switch.

Consideration 2:
What assumptions/requirements are made of the “local” (within one operator’s network) connection between the redundant ENNI nodes?

How are ENNI nodes connected?

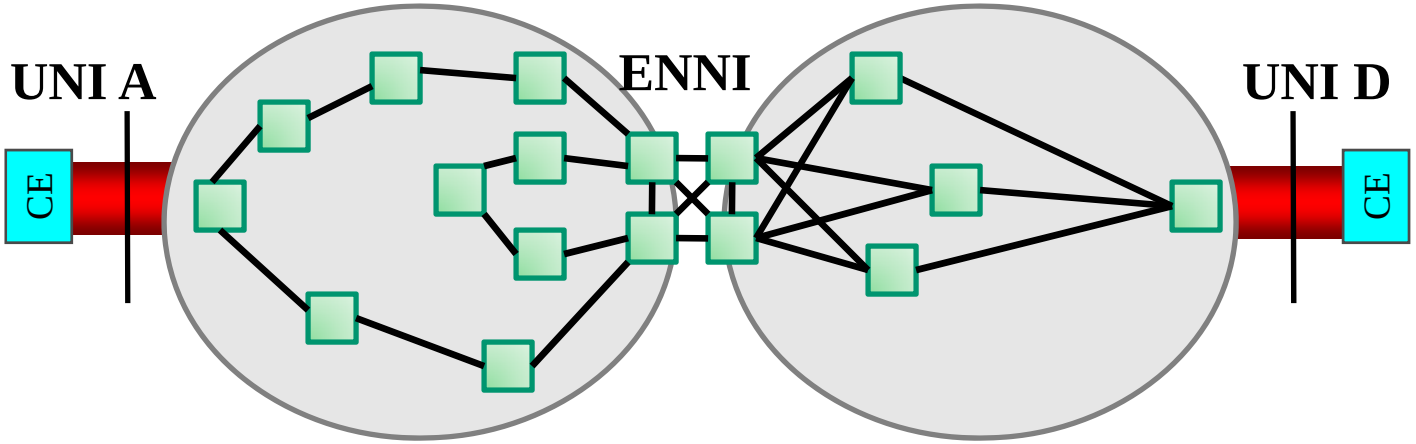

Does a local connection imply a direct connection?

For control packets only? For data as well? For transit data or only packets that will traverse ENNI?

How are switches in network connected to ENNI nodes?

Does there need to be a path to each of the ENNI nodes? Does there need to be a redundant path to each node?

For ENNI redundancy protocol control packets:

Does there need to be a direct connection between the nodes? If not, does the control path connectivity need to be monitored? How? If with CFM, does it need to be monitored per service? If not, are we back to a direct connection restriction to avoid overlapping MDs?

Particularly if ENNI redundancy protocol control packets not confined to a direct connection between the ENNI nodes, what is the interaction between the redundancy control protocol and whatever loop prevention protocol is running in the operator network (RSTP/MSTP, 8032, etc)?

Consideration 3: VLAN Load Balancing and MAC Address Learning

Load balancing at a redundant UNI

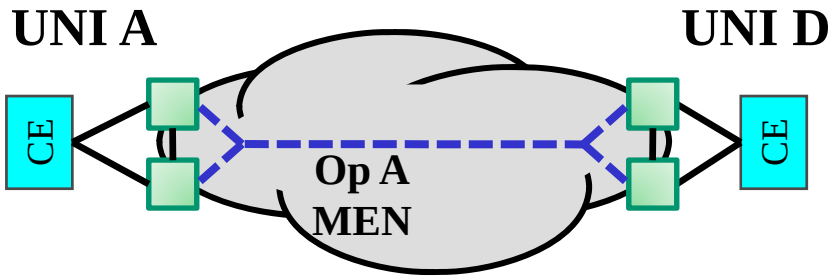

Generally assumed in carrier environment that load balancing will be done by VLAN. For the operator/provider, this is the S-VLAN. In other words both links between the CE and the provider may be active, but the provider will only transmit packets for any given S-VLAN on one of the links.

For the customer, the logical criteria for determining on which link to transmit a given packet is the C-VLAN. A potential problem arises when more than one C-VLAN maps to a given S-VLAN. If the customer uses different links for different C-VLANs that map to the same S-VLAN, then:

- The provider sees both links active for a single S-VLAN, and
- The provider cannot reliably learn through which UNI node a customer MAC address is reachable.

This is readily avoided if the customer uses a single link to transmit packets for the set of C-VLANs that map any given S-VLAN, but the onus is on the customer and the consequences of doing it wrong are manifested in the provider network.

Load Balancing over a PBB Peer ENNI

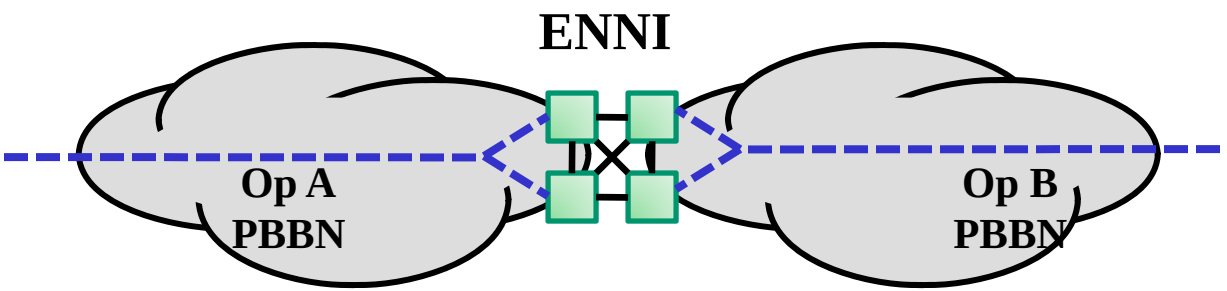

A similar, but more subtle, MAC address learning issue arises at a redundant PBB Peer ENNI. In a PBB network learning is based on a combination of B-SA and B-VID. The B-SA is the address of the I-component that encapsulated the packet. The assignment of I-SIDs to B-VLANs is local to each operator network.

If Operator A chooses the ENNI link on which to transmit a given frame based on either the I-SID or the B-VID, it is possible that Operator B will see frames arriving on both links with I-SIDs operator B wishes to map to the same B-VLAN. If any of those frames were encapsulated at the same I-component, Operator B would then see frames with the same B-SA and mapping to the same B-VID arriving on each link. Operator B then cannot reliably learn which ENNI node a given I-component is reachable.

The worst case constraint to avoid this is that both operators must map the same sets of I-SIDs to a given B-VLAN. It may be possible to relax this constraint depending on the ENNI redundancy algorithm. An algorithm that made both ENNI nodes appear to be a single switch may avoid the problem altogether.

Thank You