Ethernet Data Encryption device interoperability

Mick Seaman

This note discusses various aspects of EDE[1] interoperability, both between EDEs across various types of bridged network and between EDEs and other types of bridging systems including Provider Edge Bridges.

1. Connectivity

For two (or more) MAC Security Entities (SecYs) to interoperate and participate in the same secure connectivity association (CA):

1) The MACsec protected data frames that each transmits have to be able to reach the other(s), without any additional prepended tag (of any type).
2) The EAPOL frames that each transmits to support MACsec Key Agreement (MKA) and EAP[2] also have to be able to reach the other(s), also without any additional prepended tag and with mutually acceptable destination addresses.

It is also vital, when a CA is created, that the network a whole functions as intended and that the secure connectivity created does not prevent non-participants in the CA from receiving and processing frames that they legitimately expect to receive. Such frames might, for example, include those for routing and configuration protocols. If they did not arrive at their intended destinations, encoded as expected by those destinations, then the network might behave in undesirable and unpredictable ways, and might not be usable (see A.1). A sound approach is to aim at securing only existing connectivity (rather than using MACsec to create additional connectivity) while explicitly planning for the continued unsecured use of other desirable connectivity (see A.2).

To guard against the accidental creation of CAs that exclude the intended recipients of frames, group addresses that have an intentionally limited scope (or reach) are used as the destination address of EAPOL frames. If, for example, the 'PAE group address'[3] is used the CA will not extend through Customer or Provider Bridges[4]. These bridges filter that address, and others that a PAE might use that are identified as Reserved Addresses in IEEE Std 802.1Q. This filtering also reduces each PAE's exposure to attacks launched from other parts of the bridged network, requiring an attacker to gain access[5] to a LAN whose path to the target PAE is not filtered.

2. Authorization

Of course EAPOL connectivity between PAEs is a necessary but not necessarily sufficient condition for CA creation. The PAEs have still to mutually authenticate and the connectivity needs to be authorized. However, considerable care needs to be taken if those constraints are to be effective against accidents and attacks[6]. Consider, for example, an organization that simplifies its secure network provisioning by using pre-placed keys (PPKs), installing the same CAK,CKN tuple on a large number of EDEs. Any pair of those EDEs that have mutual MACsec and EAPOL connectivity [(1) and (2) above] will then create a CA — mutual CAK, CKN possession implies prior authentication and authorization[7]. Using pair-wise PPKs is a more robust approach, but naturally requires more administration. Using periodic EAP authentication and subsequent authorization provides the network administrator with additional ways to ensure that the CAs created are limited to those intended. In an ideal world, the network administrator will have a reliable, automated, network map and the authorization decisions that precede an Authentication Servers transmission of an EAP MSK[8] will be informed by the location of the two PAEs.

[1] Ethernet Data Encryption device as specified in P802.1AEcg.

[2] Extensible Authentication Protocol.

[3] 01-80-C2-00-00-03 (identified by IEEE Std 802.1X as the 'PAE group address' and by IEEE Std 802.1Q as the 'Nearest non-TPMR Bridge group address').

[4] It might exclude such a bridge entirely from a CA, but will not carry data through the bridge while making it blind to control frames.

[5] The access might be physical, though access through some decapsulating or other packet transforming trojan device or software entity should also be considered.

[6] Of course attackers might attempt to force accidents, especially configuration errors, so they may be much more likely than first thought.

[7] If one of the EDEs is stolen (and some EDEs, such as those providing Network Interface Device or NID functionality, will be designed to be small and thus eminently portable) the organization may have a serious problem. The possibility of equipment theft (from installed locations or by diversion) should be explicitly considered in any network design. Naturally the equipment should be designed to make key extraction economically infeasible even if it has been stolen.

[8] MSK - Master Session Key - from which the CAK and CKN used by MKA is derived.

3. Potential peers

In the context of the layered 802.1 architecture, the SecYs and PAEs that meet the connectivity requirements discussed above (Section 1) are peers, with all the protocol entities on the path between them being part of lower sub-layers. Each SecY and its associated PAE are instantiated within an interface stack, and its sub-layered position within that stack is constrained only by constraints on the protocol entities above and beneath it. For example, a PAE and SecY in a VLAN-unaware MAC Bridge Port might reasonably peer with a PAE and SecY in any of the following:

a) The TPMR (or an EDE-T) acting as a Network Interface Device (NID) for a provider network.

b) A Provider (S-VLAN) Bridge Port.

c) A Customer Edge Port on a Provider Edge Bridge (PEB).

d) A Provider Edge Port on a Provider Edge Bridge.

e) Another MAC Bridge Port.

f) A Customer (C-VLAN aware) Bridge Port.

In the first three of these scenarios, MACsec is being used to protect access to a provider network. In the fourth (d) it is protecting connectivity across the provider network to a distant PEB. In the last two, (e) and (f), connectivity between the two ports might be provided by a LAN without any intervening bridges or, alternatively, by a provider network. While not all these scenarios are of equal interest to most users of MACsec, all are possible. In 802.1Q, for example, the distinction between ‘provider’ and ‘customer’ is a technical one, not a definition of a business relationship. An organization with a large number of customer bridges might operate its own provider network, thus bringing the peers in the examples (d) and (e) above under a single administration.

A given port might participate in more than one of these secured connectivity scenarios, with more than one PAE/SecY pair in its interface stack (whether a particular item of equipment is capable of doing so is another issue). 802.1AE-2006 Figure 11-13 describes the simultaneous support of both (b) and (f) above. The two CAs created can be described as nested, with the scope of the CA at the lower sub-layer, to (b), lying within the CA to (f). CAs can also, of course, be nested along a network path without sharing a bridge port. For example: two Provider Bridges might participate in a CA forming part of a path independently secured with a CA created by Customer Bridges.

4. Address scopes

The VLAN-unaware MAC Bridge Port’s PAE can use an appropriately scoped group address as the MAC destination address of EAPOL frames, to be certain of communicating with the intended peer PAE. Figure 1 shows 802.1Q Reserved addresses that can be used by protocol entities associated with various port types. Each of the addresses shown between a pair of like ports is filtered by the MAC Relay Entity associated with that port. So (for example) frames transmitted by one Customer Bridge Port, with the Nearest Customer Bridge group destination MAC address, can reach another Customer or MAC Bridge Port (or a Provider Edge Port on a PEB that is supporting C-tagged service selection) only if it does not have to pass through another of those ports en route.

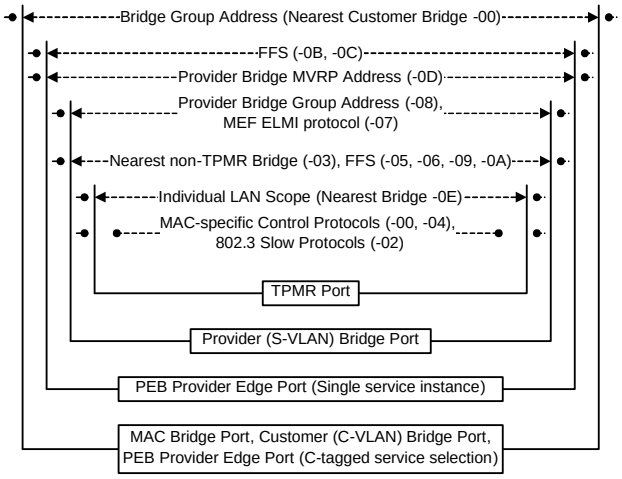

802.1Q Reserved addresses are the block of 16 beginning 01-80-C2-00-00-00

FFS - for future standardization: -05,-06,-09,-0A,-0B,-0C,-0F

Addresses (inc. -08) are not assigned for the exclusive use of particular protocols, EtherTypes are still required as other protocols can use each address, with possible constraints (e.g. for -02).

Figure 1—802.1Q Reserved addresses

An important property of frames with these destination addresses is that all bridges, not just bridges that implement MACsec, cooperate to limit their propagation. Thus two MACsec-capable Customer Bridges cannot create a CA that spans an intervening Customer Bridge, preventing the latter from understanding the RSTP BPDUs that they transmit. The effects and conditions associated with the possible creation, or of an attempt at creation, of a CA with an inappropriately extended reach can be obscure. If a network administrator is being careful problems will be detected before there are any adverse network consequences. However it is bad policy to assume that a configuration mistake will be cleared up by scrupulous attention to other details.

A bridge (at any sub-layer) cannot simply fall-back on EtherType filtering if it is unable to use the destination group MAC address to prevent inadvertent forwarding of EAPOL frames whose scope should have been constrained:

—The EAPOL EtherType might be hidden behind a VLAN tag, with that intervening tag being subsequently removed by another bridge.

—In the absence of destination information, any given EAPOL frame might relate to a perfectly reasonably scoped CA that is meant to extend through the bridge.

—The EAPOL frame might be protected by MACsec already. In that case it might be part of protocol exchanges that are setting up or maintaining a reasonable nested CA or an inappropriately scoped CA. In either case a bridge forwarding a confidentiality protected EAPOL frame can't tell that it is an EAPOL frame.

All of the above potential scoping deficiencies (and possibly others) might be remedied or the likelihood of their occurrence reduced by controls in other bridges in the network. A particularly important case is that of the choice and use of a group address for communication between Provider Edge Ports in EDE-CCs. None of the 802.1Q Reserved addresses are forwarded by a PEB's C-VLAN component if the PEB is providing a C-tagged service interface, unless all the frames that the PEB forwards are associated with a single provider service instance (equivalently, with a single PEP[9]) just as if the PEB were providing a port-based interface but with the possibility of filtering some C-VLANs. The likely candidate address for forwarding would be the Nearest Customer Bridge address, but this address is used by spanning tree protocols and its selective forwarding could cause protocol failure[10].

5. EDE-CC PEP addressing

Bearing in mind the discussion immediately above, the possible addressing choices for frames exchanged between the PAEs of EDE-CC PEPs include: (1) use of individual addresses; (2) allocation of a single additional group address; (3) allocation of a block of group addresses for selective use; (4) use of one or more administrator selected group addresses. The scoping provided by these would be enhanced by selective filtering in EDE-CCs and (if possible) in other bridges whose configuration can be controlled or influenced by the EDE-CC administrator. Each of these potential choices is discussed below. One further possibility can be dismissed at the outset: allocating one of the existing 802.1Q Reserved addresses and convincing sufficient PBN equipment suppliers and service providers not to filter it.

1. The use of individual addresses would require per CA EDE configuration, prior to EAP authentication. It would also not protect against the accidental misplacement of the EDEs[11] resulting in incorrectly scoped CAs. EAPOL and other frames transmitted by the frames might also traverse unnecessary links, possibly advertising more information about the network or security configuration than strictly necessary, particularly in the period before addressing learning in the bridged network has located the EDEs. EAPOL frames could be sent to a PAE from almost anywhere in the bridge network, increasing the potential for a DOS attack or similar mischief. Filtering of all individually addressed EAPOL frames by each EDE-CC's edge C-VLAN component would reduce but is not guaranteed to eliminate the potential for incorrect CA scope extension, while undesirably restricting the future use of EAPOL in other scenarios:

—it would be possible for a CA to extend from an EDE-CC PEP through a PBN interface, out through another interface to the PBN, across the customer network, back through a further interface into the same or another PBN, and then out through a PBN interface to the paired EDE-CC.

—if a pair of EDE-CCs was configured 'the wrong way round' they might create a CA across the customer network (see A.1 for undesirable consequences). If just one of the EDE-CCs was reversed then a variant on the prior scenario is possible.

—it would not be possible to use EAPOL directly to create a MACsec protected connection between, for example, a pair of router ports across the bridged network. While using MACsec in this way explicitly violates the security-independent connectivity criteria stated in 802.1AE, it is of current pragmatic interest[12]. Although such use is

[9]See 802.1Q clause 13.41.

[10]Each of the attached Customer Bridges would be unaware of the fact that its spanning tree protocol entity would be exchanging BPDUs with at most one of its peers while potentially transmitting and receiving data frames from all of them.

[11]Even with supposedly automated systems physical systems can be swapped and sent to the wrong location, or their configurations transposed—which amounts to the same thing. I once received a pair of iPhones with their allocated numbers and other identifiers swapped in a way that took a (dedicated and intelligent) service rep an hour to disentangle. If they had been shipped to different locations resolution would have been much more difficult.

[12]Because the performance of current highest performance commercially available MACsec implementations surpasses those for of IPsec.

clearly outside the scope of the current EDE project, it would be foolish to write-off future interest—possibly coupled to a detailed analysis of the restrictions needed to avoid ill-effects. Forcing, or attempting, the allocation of a further EtherType for this purpose would be bound to encounter stiff and well-reasoned opposition, not least because other protocols might also emerge as candidates for moving scoping attributes from addresses to protocol identifiers with implications for the EtherType allocation rate. To the degree such end-to-end EAPOL use could be facilitate by simply tagging EAPOL frames we come back to the difficulty of actually determining whether a tagged or protected frame is carrying an EAPOL EtherType (as already discussed above).

2. Allocation of a single group address for EDE-CC PAE use would most closely resemble the current use of Reserved addresses. At a minimum the scope of this address would be restricted by filtering by the edge C-VLAN component of EDE-CCs. In that case some of the misconfiguration scenarios described above when considering the use of individual addresses would still exist. If the network was correctly configured, that is to say individual EDE-CCs were correctly placed in the network and all paths from each PBN to customer bridged regions of the network passed through EDE-CCs or other bridges that filtered the allocated group address, then EAPOL frames would not traverse unnecessary links and the potential for mischief making through attacks from a distance would also be reduced. No prior per-CA configuration of the group address would be required, as it would be well-known (standardized). Overloading the address to (weakly) express authorization would no longer be possible. When EAP is being used that would place all authorization under the control of the centrally administered authentication and subsequent authorization process, with some residual probability of incorrect CA creation if all of the following are true: (a) EDE-CCs have been exchanged so they are not in fact in the expected network location; (b) the network is not correctly configured so an extended EAPOL path does exist; and, (c) the authorizing server is unaware of the relative location of the EDE-CCs. These conditions would also need to occur (with the exception of the last) if pre-placed keys (PPKs[13]) were used.

3. Allocation of a block of group addresses would allow some of the responsibility for detecting incorrectly placed EDE-CCs to be shifted from the final authorization steps to prior per-CA configuration of addresses. However the gain from this extra level of configuration is small, and only helps reduce the probability of accidents since there is nothing secure about the approach. It also does not help if PPKs are used. The use of matching CAK/CKN tuples with PPKs provides a secure method of pairing EDEs, without the extra complication of address configuration.

4. The use of administrator configured group addresses that are not drawn from a standardized block would seem to offer all the disadvantages of the standardized block with the additional downside of lowering the probability that these addresses would be filtered in all the EDE-CCs in a network.

Of the four approaches discussed above none is perfect, but it would appear that the second is the clear winner. The question is then whether the EDE-CC specification should forbid, allow, or require implementations that (to be allocated) standard address to be over-ridden by administrative configuration. In the next section we consider an EDE-CCs potential peers to help answer that question.

6. EDE-CC PEP PAE peers

The prior section (5. above) considered the choice of destination MAC address for frames exchanged between the PAEs of EDE-CC PEPs. What other potential peers might an EDE-CC PEP PAE have?

A PAE and SecY associated with the EDE-CCs Provider Network Port (PNP), and not with the PEP, would be used if it was also thought advisable to protect connectivity in the limited nested scope between an EDE-CC and either:

—a TPMR or EDE-T between the EDE and the PBN.

—the PEB providing the EDE with a PBN interface.

If the EDE-CC is attached to an S-tagged interface, or indeed any type of PBN or PBBN interface other than a S-tagged or port-based interface, the network has been misconfigured. There is little or nothing to be gained from trying to use the EDE-CC PEP's PAE to analyse that misconfiguration, and hence no requirement to change its PAE address to match that of the connected equipment. It is reasonable to ask whether an EDE-CC PEP might peer directly with a MACsec capable Customer Bridge, or indeed an end-station attached to such an interface.

[13]CAK/CKN tuples configured in the EDE-CCs. The term 'pre-placed key' has been used to express a subset of 'pre-shared keys' (PSKs), the latter may be considered as including CAK/CKN tuples acquired by authentication/authorization protocols other than EAP.

As a prelude to that inquiry consider the use of EDE-CCs to secure a PBN providing with C-tagged interfaces. Figure 2 shows the path between two bridges attached to such a network, with the tagging of data and EAPOL frames along the path. The figure supposes that C-tagged service interfaces are being used to facilitate point-to-point communication between four bridges, but only two of these need to be shown to facilitate the present discussion.

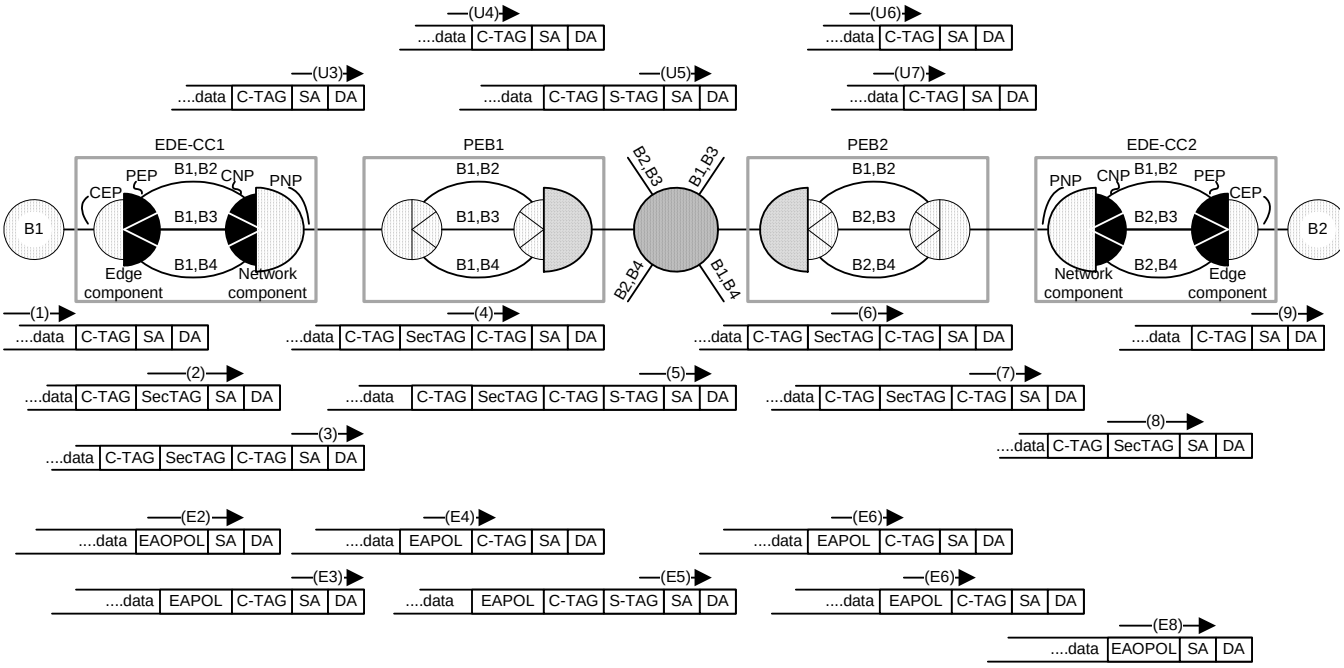

Figure 2—Unprotected data, data, and PEP EAPOL frames in a network with EDE-CCs

If B1 and B2 were directly connected to the PEBs providing the C-tagged interfaces, the tagged frames would look like that shown in U3 through U7[14] (towards the top of the figure), with the following caveats. U3 shows a C-tagged frame transmitted by B1 to PEB1. That frame might be untagged, in which case PEB1's edge component would classify it into a VLAN before forwarding it on the internal LAN selected by that classification, thus assigning it to one of the service instances shown—B1,B2; B1,B3; or B1,B4. The frame could be untagged at B4 (U4), though each service instance can only carry one VLAN without a C-TAG. A frame that lacks a C-TAG at U4 would also lack that C-TAG when transmitted on the PEB2 internal VLAN corresponding to the B1,B2 service instance (U6), but the edge component of PEB2 would have to assign each received frame to a C-VLAN and tag frames for all but one VLAN if B2 is to be able to distinguish between VLANs supporting connectivity to B1, B3, and B4.

EAPOL frames exchanged by the B1,B2 service instance PEPs of the EDE-CCs are shown in E2 through E8 (towards the bottom of the figure). They are transmitted untagged (E2) on the B1,B2 internal LAN of EDE-CC1 and need to be received untagged (E8) on the B1,B2 internal LAN of EDE-CC2 (and vice versa for EAPOL frames transmitted in the opposite direction). Removal of the outer C-TAG by the Network component of EDE-CC2, prior to transmission on the internal LAN to its Edge component, is also required for the MACsec protected data frames—so we have consistency between the treatment of data and EAPOL frames. For B2, in the figure, to distinguish between EAPOL frames for the secure associations between B1 and B2, B2 and B3, and B2and B4, those EAPOL frames would need to be tagged; and in order for B2 to associate each of these secure associations with a port-like construct it would have to add a structure very much like of the EDE-CC to its secured port. There is little point therefore in elaborating on how B2 can peer with a distant EDE-CC PEP in the case where B2 requires multiple secure associations, since the answer to that question is that B2 needs to incorporate and EDE-CC of its own, and the PEPs on that EDE-CC are doing the peering. However in the special case where B2 only requires a single secure association, as would occur if the distant EDE-CC is providing the hub for

[14]Numbers chosen to align with points on the path numbered for the inclusion of EDE-CCs.

point-to-multipoint secure connectivity, then B2 could attach to PEB2 without an EDE-CC. PEB2 would have to be configured to deliver and accept untagged (that is to say not C-tagged) traffic to and from B2. Since the PEB at the hub would filter the Nearest Customer Bridge group address[15], B2 would have to use a different group address for its EAPOL frames, and might as well use the one to allocated for EDE-CC PAE use, as discussed above.

7. EDE-SS PEP PAE addresses

The addresses used by EDE-SS PEP PAEs need to be able to pass through Provider Bridges, but (preferably) not through C-VLAN components. The address to be used should also not be passed through an EDE-SS, to avoid the accidental creation of CAs with overlapping scopes. This rules out the use of the Nearest Customer Bridge group address (-00). Use of either the -0B or -0C 802.1Q Reserved Address would be convenient. which ever address is to be allocated would no longer be available for use by C-VLAN components in any network where EDE-SS's might be present.

8. EDE-CS PEP PAE addresses

Similarly, the addresses used by EDE-CS PEP PAEs needs to be able to pass through Provider Bridges, and should not pass through C-VLAN components. The latter restriction should apply even if the C-VLAN component was on a PEB that met the Customer Edge Port to single service instance restriction, so the Nearest Customer Bridge address cannot be used. To be safe it would be desirable to use a group address that is always forwarded by S-VLAN components, including those in EDE-SS's. Note that the Provider Bridge MVRP Address does not meet this criterion, even though 802.1Q Clause 8 does not specify its filtering as a Reserved Address in all S-VLAN components.

9. EDE PAE non-Reserved Addresses

At the present there is a very clear case for EDE-CC deployment, while that for EDE-SS's and EDE-CS's is less clear, as service providers have (reportedly) not made S-tagged interfaces widely available. In the immediate future the need for these EDEs may be limited to organizations that are administering their own provider networks, or are using other's provider networks to implement their own PBBNs. As a result their might be some resistance to standardizing the use of the -0B (for EDE-CS's, say) and -0C (for EDE-SS's, say). In that case we could allocate further Group Addresses, mandating their filtering in EDEs (as appropriate for each type of EDE) and recommending their filtering (as appropriate, and in the same way that -0B and -0C would be filtered) in other bridges. However this is not the best solution, and given the fact that PBN specification is now generally regarded as a 'solved problem' it would be a better use of resources to use -0B and -0C as suggested and resolve any future address demand that might lead to conflict in the future.

10. EDE-T PEP PAE addressing

TPMRs, as currently specified by 802.1Q, forward frames addressed to 'Nearest non-TPMR Bridge group address' (01-80-C2-00-00-03). This is also the group address initially assigned for 802.1X use as the default 'PAE group address', raising the question as to how exactly EDE-Ts are to be used—is the scope of the CAs that they create to be bounded by other TPMRs as well as other devices, or should they function at a level 'slightly above' that of other TPMRs, but below all other devices. The latter seems much more desirable, since a large part of the reason for creating TPMRs was to create a specification that could describe media converters and other devices in a way that they could be uniformly managed and be non-disruptive in a network architecture. By using the -03 address EDE-T's can be incorporated within a chain of TPMRs without their functionality being disrupted by the addition of further media converters and the like.

[15]Since it doesn't make a lot of sense to attempt to run unmodified spanning tree over a p2mp configuration.

A. Example scenarios

A.1 Overlapping CAs

<<This is currently a placeholder for the simple oscillator or an equivalent example:

Two bridges (BL, BR), three ports each, two EDEs (EL, ER). Connect the two bridges, BL1 (BL port 1) to BR1. Connect BL2 to ELred, ELblack to BL3 and BR2 to ERred, ERblack to BR3.

Two network paths across the BL1-BR1 link, although the bridges don't know the paths share the link. The bridges think they are connected by BL1-BR1 and by an independent path that goes BL1-EL-ER-BR. They both see the links to the black side as dark (no intelligible traffic, presumably goes to end station).

So the bridges see the two paths, and a potential loop. If their spanning tree (or similar loop free protocol) cuts the BL1-BR1 link then the other one will get cut as well, so the BL1-BR1 link gets put back - and the other reappears - repeat.>>

A.2 Secured and unsecured connectivity

<<This is currently a placeholder for showing the use of both secured and unsecured connectivity (the latter using the Uncontrolled Port in some scenarios) and the resulting discussion, which may have to be enlarged. One way this could also go is for the untagged traffic (or other specifically selected VLAN) through an EDE-CC (for example) to be directed to a port corresponding to an untagged service instance which is deliberately not secured so that it can support LMI exchanges with the PEB. Such traffic would not be accepted by a distant EDE/MACsec enabled station. In the inbound (to the attached customer network) direction, the VLAN would serve to segregate the LMI traffic (if forwarded at all by equipment attached to the EDE-CC).>>