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. It works through the issues in some detail. Some important conclusions have been updated from the initial August 2015 revision of this note. Table summarizes the use of various addresses by the PAEs that establish secure connectivity in various interoperability scenarios.

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 SecY transmits have to be able to reach the other(s), without any additional prepended tag (of any type).
2) The EAPOL frames each SecY 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. Such frames might, for example, support routing and configuration protocols. If they do 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 Annex A). 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 Annex B).

To guard against the accidental creation of CAs that exclude intended recipients, group addresses with 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. Care needs to be taken to ensure those constraints are 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, since 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] 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. The equipment should be designed to make key extraction economically infeasible even if it has been stolen.

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.

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 Provider Edge Bridge (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 CA to (b) at the lower sub-layer existing within the CA to (f). Of course, CAs can also 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.

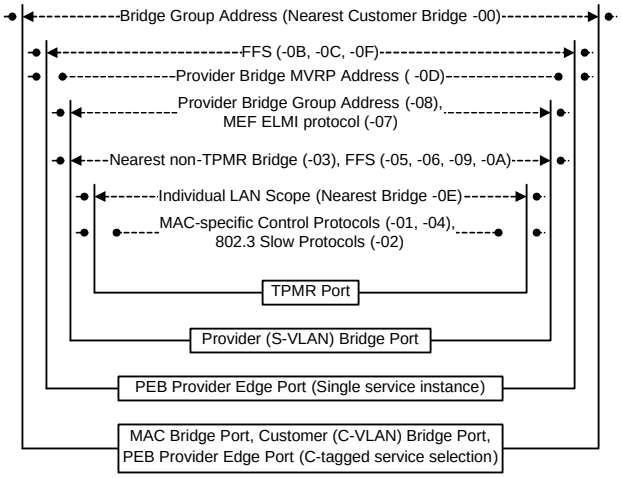

802.1Q Reserved addresses are the block of 16 beginning 01-80-C2-00-00-00

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

Addresses (inc. -07) 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).

The Provider Bridge MVRP Address (-0D) is filtered by S-VLAN components using MVRP as well as by C-VLAN components, so may not reach the latter.

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

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

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. If, as expected, each EDE-CC is connected to a provider network by the Customer Edge Port of a Provider Edge Bridge, then none of the 802.1Q Reserved addresses will be forwarded by the edge port's C-VLAN component unless they are all associated with a single provider service instance[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 could 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. 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.

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 could shift some of the responsibility for detecting incorrectly placed EDE-CCs from the final authorization steps to prior per-CA configuration of addresses. This extra level of configuration might reduce as there is nothing secure about the approach. The use of matching CAK/CKN tuples with PPKs already 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?

6.1 Access protection

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:

—an EDE-T between the EDE and the PBN.

—the PEB providing the EDE with a PBN interface.

The Nearest non-TPMR Bridge group address could be used (see Section 9), though in an EDE-T that could interfere with true non-TPMR device discovery, (e.g. of the PEB). The Individual LAN address might also be used, but that would be filtered by any TPMR specification media converters in the link between the EDE-CC and the EDE-TT. An new Reserved Address allocation is desirable for EDE-T for use in scenarios

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

where both these issues need to be avoided and the Reserved Address ending -0A is suggested.[14]

6.2 PBN attached Customer Bridges

If the EDE-CC is attached to an S-tagged interface, or indeed any type of PBN or PBBN interface other than a C-tagged or port-based interface, the network has been misconfigured. The EDE-CC PEP's PAE cannot analyse that misconfiguration and 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 the same PBN.

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.

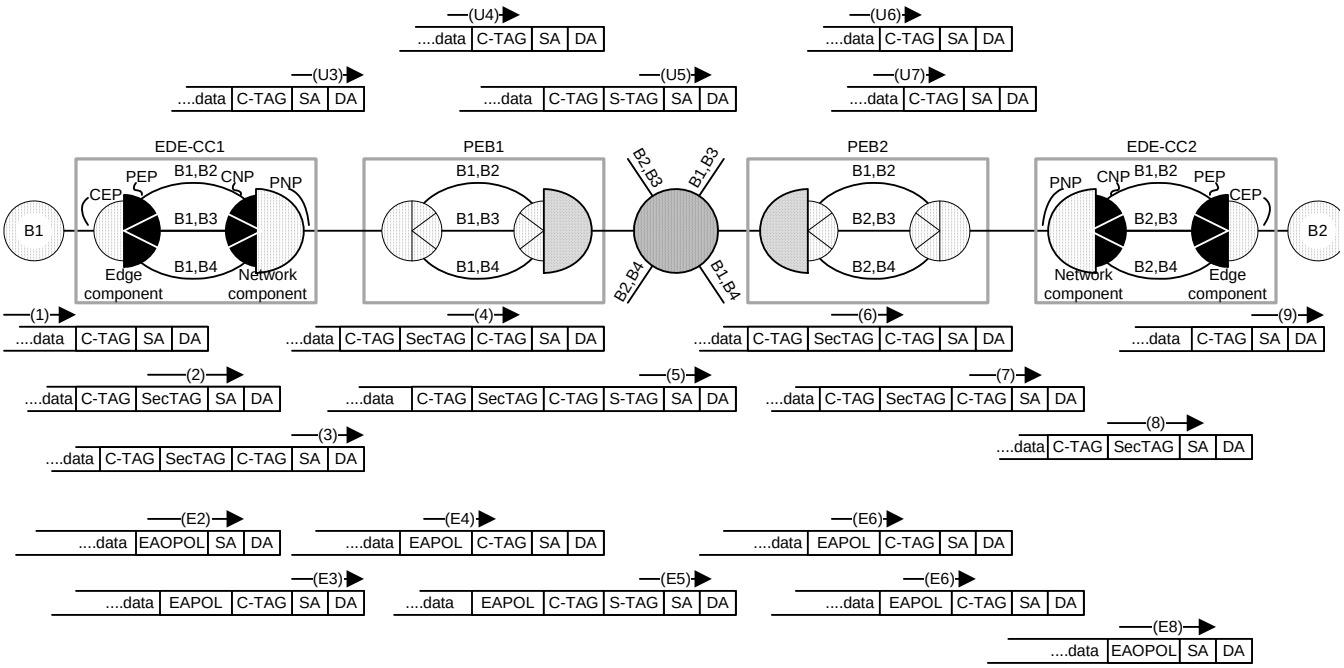

Figure 2—Unprotected data, data, and PEP EAPOL frames in a network with EDE-CCs

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[15] (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

[14]i.e. 01-80-C2-00-00-0A. Normally we would not allocate or suggest an address assignment until the close of Sponsor Ballot, but here there are few candidates, and we would not wish some pre-standard devices to use one and others to use the other, potentially causing future problems.

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

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 an 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 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[16], 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.

6.3 EDE-CSs

An EDE-CC attached to a PEB that provides a C-tagged interface to a provider network could peer with one or EDE-CSs attached to an S-tagged interfaces to that network, though that it is unclear why such a scenario would be configured. For this to work the EDE-CS would have to assign each relayed C-VLAN to a distinct S-VLAN, forwarding frames through the provider network without a C-VLAN tag immediately following the S-VLAN tag. The Nearest Customer Bridge address could not be used for EAPOL frames, while the address allocated for EDE-CC use would be suitable, so the EDE-CS PEP PAEs would have to be configured to use that address.

7. EDE-SS PEP PAE addresses

The addresses used by EDE-SS PEP PAEs need to be able to pass through Provider Bridges[17], 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 the -0B 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.

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. The need for these EDE types may be limited to organizations that are administering their own provider networks, or are using other's provider networks to implement their own PBBNs. Is it worth allocating one of the limited number of Reserved Addresses for EDE-SS use? An alternative would be to allocate a further group address, mandating its filtering in EDE-SS's, EDE-CS's, and EDE-CC's, and recommending filtering in other C-VLAN components, hoping that would pick up over time. However that is not the best solution, and since PBN specification is now generally regarded as a 'solved problem' it would be a better use of resources to allocate a reserved address, leaving the remaining two C-VLAN component-specific reserved addresses to handle future requirements.

8. EDE-CS PEP PAE addresses[18]

The addresses used by EDE-CS PEP PAEs needs to be able to pass through Provider Bridges, and ideally not through C-VLAN components. The Nearest Customer Bridge address can be used. If an EAPOL frame sent by the EDE-CS is relayed by the C-VLAN component for a PEP (a Provider Edge Port on a Provider Edge Bridge) supporting a single service instance to an attached Customer Bridge, that bridge will not relay the frame further. If the attached port was MACsec (or at least 802.1X) capable and the EAPOL frame untagged, a CA could be created; otherwise the frame would be simply discarded.

An EDE-CS PEP could also interoperate with an EDE-CC PEP, using the (to be allocated) standard group MAC address for the latter.

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

[17]The Provider Bridge MVRP Address does not meet this criterion, even though 802.1Q Clause 8 does not specify its filtering in all S-VLAN components.

[18]The text and conclusions in this section have been revised from the initial August revision of this note. The positive effects of the changes are: (a) the use of only one address drawn from the 802.1Q Reserved Addresses set (to support EDE-SS's); (b) interoperability (with constraints) between EDE-CS's and MACsec-capable MAC Bridges attached to port-based/single service instance provider network instances, between EDE-CS's and EDE-CC's. The downside is that an EDE-CS configured to interoperate with an EDE-CC introduces the same risk of an incorrectly scoped CA traversing a customer network as the EDE-CC.

9. 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. 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. However that might interfere with other uses of the -03 address, and allocation of a further -0A is recommended (see 6.1).

MAC Bridge or Customer Bridge Port (-03 default, -07, -08, -0A, -00, A1)	**LAN**	-03											
	PBN port i/f	X[1,2]	-00[3]										
	PBN C-tagged i/f	X[1,4,5]	X[1,4,5]	X[1,4,5]									
EDE-CC (Provider Edge Port) (A1 default)	**LAN**	X[2,6]	X[7]	X[7]	A1[1,8]								
	PBN port i/f	X[2]	A1[6]	X[6]	A1[1,8]	A1[1,8]							
	PBN C-tagged i/f	X[4]	A1[1,8,9]	X[6]	A1[1,8]	A1[1,8,9]	A1						
PEB (Customer Edge Port) (-03 default, -07)	**LAN**	-03	X	X	X[10]			X					
EDE-CS (Provider Edge Port) (-00 default, A1)	**PBN i/f S-tagged**	X[1,2]	-00[6,9]	X[6]	A1[1,8,9]	A1[1,8,9]	A1[1,8,9]	-00					
EDE-SS (Provider Edge Port) (-0B default)	**PBN or PBBN port-based i/f**	X	X	X	X	X	X	X	-0B[1,11]				
	PBN or PBBN S-tagged i/f	X	X	X	X	X	X	X	-0B[1]	-0B			
	PBBN transparent i/f	X	X	X	X	X	X	X	X	X	-0B[12]		
Provider Bridge (PB) Port	**(S-)LAN**	-03 -08	X	X	X	X	X	X	X	X	X	-03 -08	
EDE-T	**LAN**	-0A -03[13] -0E[14]	X	X	X	X	X	X	X	X	X	-0A -03[13] -0E[14]	-0A -03[13] -0E[14]
	Connectivity	**LAN**	**PBN port i/f**	**PBN C-tagged i/f**	**LAN**	**PBN port i/f**	**PBN i/f C-tagged**	**PBN i/f S-tagged**	**PBN or PBBN port i/f**	**PBN or PBBN S-tagged i/f**	**PBBN[15] transparent i/f**	**(S)-LAN**	**LAN**
		MAC or Customer Bridge Port			**EDE-CC (Provider Edge Port)**			**EDE-CS (PEP)**	**EDE-SS (Provider Edge Port)**			**PB Port**	**EDE-T**

Additional scenarios are include to provide context for EDE operation. X indicates misconfiguration scenarios where MACsec connectivity is undesirable or impossible. In some cases there is a fine line between 'X' with advice on address selection to avoid connectivity and describing constraints on partial connectivity.

-00, -03, -07, -08, -0A, -0B: use Nearest Customer Bridge, Nearest non-TPMR bridge, E-LMI, Provider Bridge Group Address, EDE-T address, or EDE-CS PEP address respectively for EAPOL. A1 indicates use of the (tba) EDE-CC PEP PAE address.

Table 1—EDE Interoperability and Addressing

[1]Probable or definite misconfiguration.
[2]The bridge on the external LAN can use -03 or -07 to avoid undesirable CA creation, these addresses will be filtered by intervening components.
[3]A better choice than A1 if connectivity is desired, otherwise use -03 or -07.
[4]Customer bridge uses -03, -07, or -00 which are filtered avoiding undesirable CA.
[5]Outer SecTAG/EAPOL type defeats service selection, so one service instance/if, probable misconfiguration so do not change to use A1.
[6]A CA might be created for just one EDE PEP per PBN port i/f using Customer Bridge, or the latter might act as if supporting a multi-access LAN scenario. Accidental creation of this scenario is avoided by not configuring the MAC/Customer Bridge Port PAE use of the A1 address.
[7]Undesirable connectivity might be provided with restrictions[6] if both devices use A1. Clearly there is some confusion about how one of them has been connected.
[8]Misconfiguration likely to be harmless if untagged VLAN not protected and only standard protocols used.
[9]A hub-and-spoke scenario with the EDE at the hub. One service instance per C-VLAN, without a C-VLAN tag following the S-VLAN tag required for successful operation.
[10]One PEP to might secure connectivity for management protocol which might include LLDP (.1AB) and E-LMI on an untagged VLAN passing through the EDE-CC, this would require new explicit support on the PEB probably using -00, and the EDE would have to forward the -07 address to support use of the E-LMI protocol.
[11]A port-based interface will discard all S-tagged frame, so there is little point to its use with an EDE-SS, except as a spoke with an S-tagged interface to the hub.
[12]A transparent interface passes all frames through to a single further port, so there is little point to its use with an EDE-SS.
[13]Might interfere with other uses of -03.
[14]Will not work if there is any intervening TPMR.
[15]No EDE type has been defined to make direct use of a PBBN I-tagged interface: a backbone provider can provide an S-tagged interface (which that provider supports by encapsulation and I-SID mapping under his control) allowing the customer to use an EDE-SS at the interface. Equally the backbone provider could use an EDE-SS to secure the backbone (B-VLANs) after the I-tag has been added to the data of the backbone frame.

A. Incorrectly-scoped CAs

This section (A) illustrates points made at the beginning of this note—the creation of a CA that excludes legitimate recipients of frames can cause undesirable and/or unpredictable network behavior, and such CAs can result from inappropriate forwarding of EAPOL frames (or alternatively can be prevented by appropriate use of EAPOL group addresses and filters). Sometimes other protocols will prevent, or at least identify, the danger. However it is always foolhardy to assume that one administrative mistake will be remedied by careful attention to other administrative details. While not minimal, examples may be simplified to the point of appearing unrealistic. However little (other than the availability of large sheets of paper and patience) stands in the way of using basic ideas to construct large but equivalent network scenarios.

A.1 A potential oscillator

The upper half of Figure 3 shows a small network. Bridges BL and BR (each presumably having further ports, not shown) are connected to bridges PL and QL, and to PR and QR respectively. The P's and Q's are two port bridges, each with MACsec capability on one (with that side of the bridge shown in black to protection of frames and their encapsulation by MACsec).

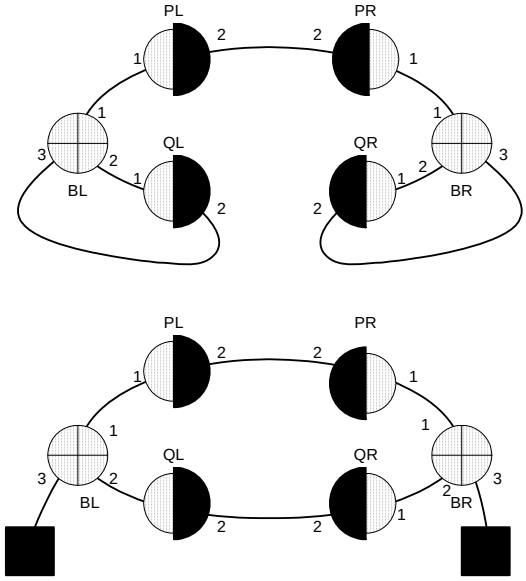

Figure 3—A potential oscillator

The network administrator's intention is that ports PL2 and PR2 participate in one CA , while QL2 and QR2 participate in another. In pursuit of that goal he has assigned one group address G_p (say) for use by the PAE's for PL2 and PR2, and another G_q (say) for QL2 and QR2. Naturally PL and PR filter G_p, so a system on the red (port 1) side of PR cannot send EAPOL frames to PL2, in an attempt to disrupt EAPOL/MKA operation. However PL and PR don't filter G_q, and QL and QR don't filter G_p, while BL and BR filter neither of the assigned group addresses.

Unfortunately QL2 and QR2 have not been connected to each other, but to BL3 and BR3 respectively. Once PL2 and PR2 have exchanged EAPOL PDUs and secured their connection, QL2 and QR2 can also exchange EAPOL PDUs over the path QL2-BL-PL-PR-BR-QR2, securing their own connectivity. From the point of view of BL and BR, the network topology might then resemble the lower half of Figure 3. The frames that BR3 receives from QR2 are unintelligible (since they have been confidentiality protected), and any (apart from those originating from QL2) that it sends to QR2 will be dropped. If BR decides to forward frames (whose eventual destination lies through BL) via BR2 rather than BR2, then the following scenario can occur. Once the BR1-PR-PL-BL1 forwarding ceases, the EAPOL path from QR2 to QL2 is disrupted and, in consequence, the BR-BL path via QR-QL will disappear, causing BR to revert to using its BR1-BL1 path. This causes the QR-QL path to reappear, and to be chosen again, repeating the cycle of oscillation.

Fortunately, this oscillation will not occur if the standard spanning protocol is being used to select the paths, because spanning tree BPDUs are transmitted to the Bridge Group Address. They will be filtered by the B's, P's, and Q's, since the destination MAC address—which determines the addressing scope of the frame—is not changed by MACsec. If the P's and Q's do not use this address, the path BL-QL2-BL-PL-PR-BR-QR2-BR will not be discovered since both BL and BR will filter the group address that they are using for the spanning tree protocol.

Unfortunately, if QL and QR are using the original Spanning Tree Protocol rather than RSTP, or RSTP has not been configured with non-default settings of AutoIsolate (to TRUE) and AutoEdge (to FALSE), then QL and QR can conclude that their CA is providing connectivity to an end system rather than a bridge, causing the formation of a loop (as a result of the non-reception of BPDUs on the QL-BL and QR-BR links) and the network will melt down.

B. Secured and unsecured connectivity

The Uncontrolled Port provided by the MACsec SecY is used to transmit unprotected frames to the port's peers. It is used by specific protocol entities that reside in the same system as the port, e.g. by the port's PAE to transmit EAPOL frames. Other examples include its use by IEEE 802.1AB LLDP (Link Layer Discovery Protocol), which can make use of both the Controlled and Uncontrolled Ports. Each of the local protocol entities making use of the Uncontrolled Port are usually the ultimate source and/or sink of the frames passing through that port, with one notable exception: 802.1X describes the Selective Relay of Wake-on-LAN (WoL) frames (see 802.1X-2010 clause 7.1.3, Figure 7-4, Figure 7-13, Annex F). These frames can be received through a SecY Controlled Port, supporting one port of a bridge, and subsequently relayed through an Uncontrolled Port associated with a SecY for another port.

An EDE that provides an encrypting interface to a PBN poses a similar relaying requirement. Assume that an E-LMI (Ethernet Local Management Interface) supported by the PBN is being used by an attached Customer Bridge, and the intention is to insert an EDE-CC (say) between the PEB and the Customer Bridge without changing the use of the E-LMI. One way to do this would be to introduce an further special purpose relaying entity to convey E-LMI frames from the red-side of the EDE to the black-side (and vice versa) without protecting them on the black-side. This places an additional burden on the Customer Bridge to ensure that none of these black-side-unprotected frames leak further into the red-side network. Its idea of what constitutes an E-LMI frame better match (or superset) the EDE's. A preferable, and more general purpose, approach is to use a VLAN to segregate E-LMI traffic. The edge component of the EDE can direct frames for this VLAN that are received on its Customer Edge Port to a Provider Edge Port that transmits traffic unprotected, and the frame can be sent to the PEB (via EDE-CC's network component) simply tagged with the appropriate VLAN (or indeed untagged).

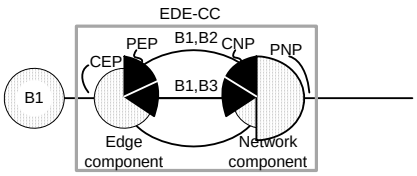

Figure 4—EDE-CC with in-clear E-LMI frames

Figure 4shows an EDE-CC that is similar to EDE-CC1 in Figure 2 above, but with its lowest PEP dedicated to supporting a VLAN (or VLANs, as determined by the 802.1Q member set for the PEP) for unprotected traffic. Equally this PEP could support protected communication with a CA that included just the PEP and the Provider Edge Bridge, thus securing E-LMI exchanges independently of the B1, B2 and B1, B3 traffic. This is likely to be a more attractive arrangement than that shown in 802.1X-2010 figure 7-17 where two layers of MACsec are used - one to secure access *to* the PBN and the other to secure access *across* the PBN. The advantage of the latter is that the PBN never accepts traffic that is not sent by the EDE, the advantage of the distinct singly MACsec protected CAs is avoiding the implementation complexity of double protection while still ensuring that only traffic sent by peer EDEs is received from the PBN.