

MACsec hops

Mick Seaman

MAC Security (MACsec) cryptographically protects frames on a hop-specific basis. IEEE Std 802.1AE describes a number of use cases, and specifies the position of the MACsec shim in interface stacks and the necessary control plane addressing for each case. When traversing a provider bridged or provider backbone bridged (PBN or PBBN) network, a ‘single hop’ can be supported by a number of intervening bridges. This note explains why these bridges can be treated as a ‘single hop’, why the restriction to hop-specific operation is important, and how to deal with some additional cases. It has been written mainly for an audience that is not familiar with the original development of both MACsec and other 802.1 protocols, but would like to explore additional possibilities.

1. Overview

Ignoring MACsec’s hop-specific design would allow additional uses of the MACsec protocol¹. However a narrow protocol-centric view of the possibilities risks setting aside data rate, control plane performance, and interoperability considerations, and does not take into account the security risks or the existence, needs, and evolution of other network protocols. This note explains why IEEE 802.1AE–2006² does not include multi-hop possibilities that might appear ‘interesting’ at first glance, and why it may prove difficult to procure equipment that uses MACsec in such custom configurations. It also describes additional use cases that support or exploit more recent developments in IEEE 802.1 bridging technology.

The ‘single hop’ restriction is a result of adopting a coherent approach to the considerations alluded to above, more specifically:

- a) MACsec protects all frames transmitted and received by authenticated and authorized protocol participants, ensuring that network operation is not compromised by unauthorized modifications or additions (3).
- b) The communicating MACsec peers that cryptographically protect and validate any given frame have to operate on the same frame fields, and the same frame field data. Modification of any frame field has the potential to disrupt network operation³, so the entire frame (including the addresses, protocol types, and header data) has to be protected and validated (3).
- c) The use of MACsec should be possible (and useful) when authenticated and authorized bridges add or

remove tags, or make permitted changes to frame fields, such as the Priority Code Point and Drop Eligible Indicator in the VLAN tag.

- d) Performing MACsec processing naturally involves accessing and (if confidentiality protection is being provided) encrypting packet data. To avoid the cost of additional memory bandwidth it is desirable to locate the MACsec processing within chips/modules that are already concerned with moving the data, such as network interfaces (4.1).
- e) Achieving the necessary data rates at reasonable cost limits the number of SAKs (the secret keys used to protect data) used at any one time (4.1).
- f) Network control plane protocol performance should not be impacted by MACsec. Specifically, it should not be necessary to agree and install fresh SAKs if the network paths are reconfigured (4.2).

Each of these points is discussed in detail below. [Section 2](#) provides a brief review of fundamental architectural concepts, [Section 3](#) discusses the security threats MACsec has to handle, and [Section 4](#) how connectivity is secured. Additional use cases are presented in [Section 5](#). [Section 6](#) contains some ideas for improvements to the standards, and [Annex A](#) and [Annex B](#) provide space for FAQs and further technical detail respectively.

¹Or derivative and similar protocols. Past proposals have commonly suggested that a single protection operation extend over multiple hops so that MACsec capability/compatibility can be claimed for existing bridges without the need to develop or deploy MACsec upgrades for those bridges.

²See IEEE 802.1AE-2006 7.3.2, and in particular NOTE 1 and NOTE 2 in that clause.

³Beyond the very limited impact of simply causing the frame to be lost.

2. Bridging architecture

A brief review of architectural concepts⁴ follows, as an aid to our subsequent examination of detailed scenarios that illustrate what MACsec (and possible alternatives) can and cannot do well.

2.1 Layering and the ISS

Layered protocol entities communicate with their peer or peers using the service provided by the protocol entities in the layer below. We are really concerned with just one service—the MAC Internal Sublayer Service (ISS)⁵.

Each and every bridge’s job can be summarized as supporting one or more instances of the ISS, with the desired efficiency/extent/manageability/etc. Each instance provides connectivity⁶ between a set of end stations and/or bridges. At the lowest layer, in end stations or intermediate systems (bridges), the ISS is mapped to (provided using) the particular media access method supported by the physical LAN. Each bridge concatenates two or more instances of the ISS. Some bridges, Provider Edge Bridges (PEBs) and Backbone Edge Bridges (BEBs) for example, include protocol entities whose explicit function is to provide one instance of the ISS over another. The protocol entities of others (Provider Bridges, for example) are configured to relay all the frames transmitted by others (Customer Bridges, in this example) so what the latter see as a single instance of the ISS is supported by two or more instances concatenated at a lower (sub-)layer.

Virtual LANs (VLANs) are really just a way (using a VLAN tag field added to each frame) of separating multiple instances of the ISS. The EISS (Enhanced Internal Sublayer Service) used by VLAN-aware bridging components is a compact way of describing that multiplexing at a service interface. An EISS service access point functions just as a number (potentially 4094) of ISS service access points in parallel⁷. Similarly the MAC address encapsulation provided by Backbone Edge Bridges separates and provides address independence between a higher and lower layered instances of the ISS.

The layered architecture of a Bridge is often drawn as in Figure 1. This shows the bridge’s MAC Relay entity below the level of the MAC Service used by higher

layer protocols (supported by LLC) in the end stations to the left and right—emphasizing the fact that the relay is transparent to those service users. For our present purposes it is more convenient to use diagrams that show the interface stacks supporting relay and higher layer entities in more detail (e.g. Figure 2).

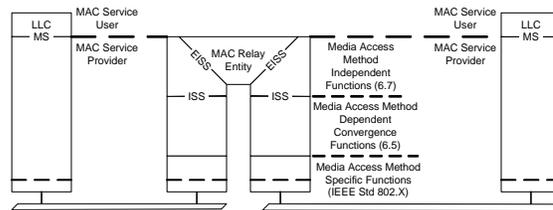


Figure 1—A VLAN bridge and end stations

2.2 Systems, networks, and components

Many protocols have been specified purely from the point of view of the rules governing the interaction of entities whose sole function is to provide the protocol (or worse by describing only the frame formats and—just possibly—individual field processing), thus assuming that the relationship between those protocol entities and the rest of the system is obvious.

In a layered system or network where the intent is often to provide or extend a service transparently (leaving the interactions between individual service users unchanged while increasing network throughput, physical extent, or the number of users) the same protocols can be used at many different layers. So can protocols and entities (such as those adding and removing VLAN tags) that are not truly transparent but serve to select between different instances of transparent service. For a standard to be useful—promoting the availability of equipment and interoperability between items of equipment developed by different organizations—it has to specify (or at least suggest) the layering relationship between (and concomitant configuration aspects of) protocol entities.

When very similar functions have to be performed at different sub-layers it is more useful to re-use an existing protocol entity than to invent a new one, and the same applies to entire combinations of protocol

⁴The architectural concepts and terms are described in some detail in IEEE Std 802.1X Annex D.

⁵The service provided by a LAN, stripped of the peculiarities introduced by one or other media access methods, but with the explicit inclusion of parameters necessary for describing the process of forwarding the frame—which might otherwise be thought particular to the way the service is provided, and not of concern to upper layers in the end stations using the MAC Service.

⁶IEEE Std 802.1AE and 802.1X-2010 Annex D.8 formalizes this notion as a Connectivity Association (CA), following RFC 787.

⁷ IEEE Std 802.1Q specifies trivial protocol entities that can be used to split/recombine an EISS interface into/from component ISS interfaces so that other protocol entities specified just for use with the latter can be used without respecification. Note that PBBN technology allows many more than 4094 VLANs to be supported by a single network, while still allowing each service instance to be explicitly identified (by the ISID).

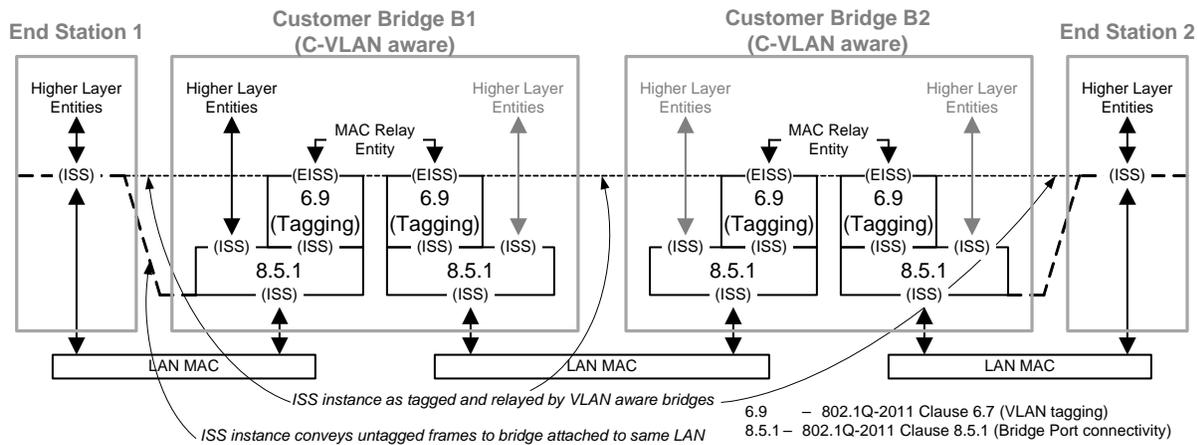


Figure 2—Interface stacks for VLAN-aware bridges and end stations

entities. The functionality of a Provider Edge Bridge is (for example) conveniently expressed (see Figure 3) as the concatenation of (a number of) C-VLAN aware components with an S-VLAN aware component, each of these internal system components having functionality that could be instantiated in separate

systems (compare the left and right sides of Figure 3).⁸ Part of the strength of this component based design is its natural inheritance and preservation of the essential arrangements for protocols that are not the immediate focus of the designer.

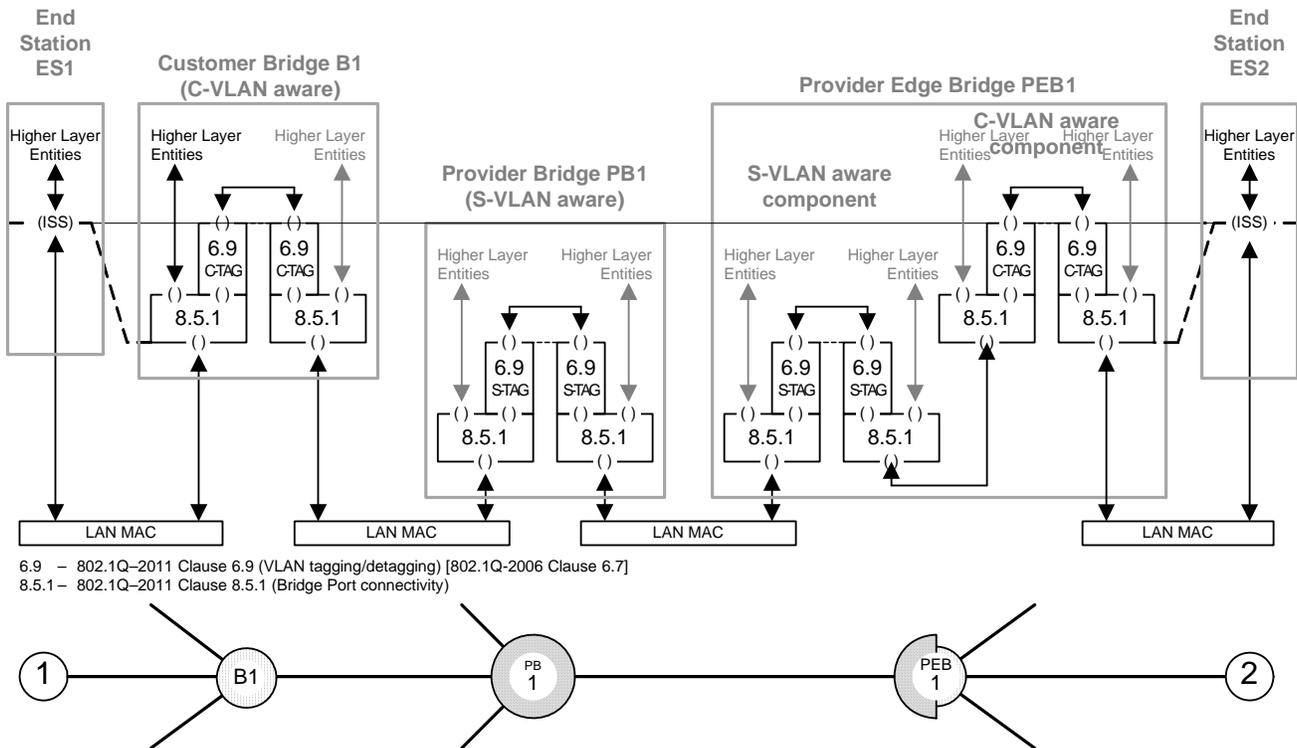


Figure 3—Interface stacks for a path through a Provider Bridged Network

The lower part of Figure 3 is a plan diagram of the connectivity between the systems and their components. This view ‘from above’ serves to

emphasize that the interface stack picture is best at showing a single path, and that other end to end paths can join and share part of this path. In some cases it is

⁸There was some initial reluctance to follow this approach—a fear that it would mandate more functionality than necessary—but experience has shown its value. Some functionality that was initially explicitly removed has had to be reinstated to address real needs, and it is now a useful way of specifying new functionality, and of actually providing that functionality by interconnecting appropriately configured systems that are already available—rather than requiring custom engineering that might prove uneconomic in cases where relatively few systems are required.

MACsec hops

convenient to restrict the properties of one of the system components (requiring a component to serve a single customer, or to only have two ports, for example). Such a restriction permits some adjustment of the way that configuration protocols are supported by the system.

2.3 Connectivity and address scopes

In general, a configuration protocol needs to know its immediate peer neighbours in a network, and to be able to transmit and receive particular frames to and from those neighbours, while the same frames should not reach more distant participants in the protocol⁹. When different instances of the same protocol is to be deployed at different layers within the architecture, the set of protocol entity peers for each protocol instance naturally differs, and has to be kept separate. This separation is enforced by using different destination MAC addresses for each protocol instance. Specific group MAC addresses are used for this purpose—it is

either impossible or impractical to manually configure the individual addresses of each entity's peers before protocol operation begins. This use of group addresses is a general feature of LAN-based protocols, but the use of reserved group addresses that are always filtered by particular types of bridges is specific to the layered architecture of bridged networks.

Each MAC Relay entity includes a Filtering Database (FDB). FDB entries are used to ensure that frames with given destination MAC addresses (or given combinations of MAC address and VLAN ID) are not forwarded. FDB entries can be created by management, by the operation of network configuration protocols (such as ISIS-SPB), or by learning the (relative) location of bridges by observing the source MAC addresses of frames. Permanent FDB entries are made for the Reserved Addresses used by protocol entities to discover their peers (at the appropriate layer), see [Table 1](#) and Figure 4.

Table 1—Reserved addresses for bridge components

Value	Assignment	Filtered by		
		C- ¹	S-B- ²	T-
01-80-C2-00-00-00	Bridge Group Address, Nearest Customer Bridge group address ³	Y		
01-80-C2-00-00-01	IEEE MAC-specific Control Protocols group address	Y	Y	Y
01-80-C2-00-00-02	IEEE Std. 802.3 Slow_Protocols_Multicast address	Y	Y	Y
01-80-C2-00-00-03	Nearest non-TPMR Bridge group address ⁴	Y	Y	
01-80-C2-00-00-04	IEEE MAC-specific Control Protocols group address	Y	Y	Y
01-80-C2-00-00-05 01-80-C2-00-00-06	Reserved for future standardization - media access method specific	Y	Y	
01-80-C2-00-00-07	Metro Ethernet Forum ELMI protocol group address ⁵	Y	Y	
01-80-C2-00-00-08	Provider Bridge Group Address	Y	Y	
01-80-C2-00-00-09 01-80-C2-00-00-0A	Reserved for future standardization	Y	Y	
01-80-C2-00-00-0B 01-80-C2-00-00-0C	Reserved for future standardization	Y		
01-80-C2-00-00-0D	Provider Bridge MVRP Address	Y		
01-80-C2-00-00-0E	Individual LAN Scope group address, Nearest Bridge group address ⁶	Y	Y	Y
01-80-C2-00-00-0F	Reserved for future standardization	Y		

¹Filtered by C-VLAN aware components in Customer Bridges and Provider Edge Bridges, and by VLAN-unaware MAC Bridges (IEEE 802.1D).

²B-components (in Backbone and Backbone Edge Bridges) behave exactly as S-components (in Provider and Provider Edge Bridges). The MAC address encapsulation provided by PEBs separates the address spaces for these components.

³As stated in 802.1Q-2011 (clause 13.39, and Table 8-1) a C-VLAN component (within a Provider Edge Bridge) that relays frames from a single Customer Edge Port to a single Provider Edge Port (see 802.1Q-2011 clause 15.4) may forward (not filter) frames with this destination address.

⁴Also known as the 'PAE group address' in 802.1X-2010, which recommends its use as the default for the PAE's EAPOL clients and the KaY.

⁵This address is not exclusively reserved for this purpose; other uses are reserved for future standardization.

⁶It is intended that no IEEE 802.1 relay device will be defined that will forward frames that carry this destination address. Protocol uses include controlling Power over Ethernet.

⁹The protocol might need to transmit frames that reach all participants in a particular instance of the protocol as well. In this case they are all treated as neighbours.

MACsec hops

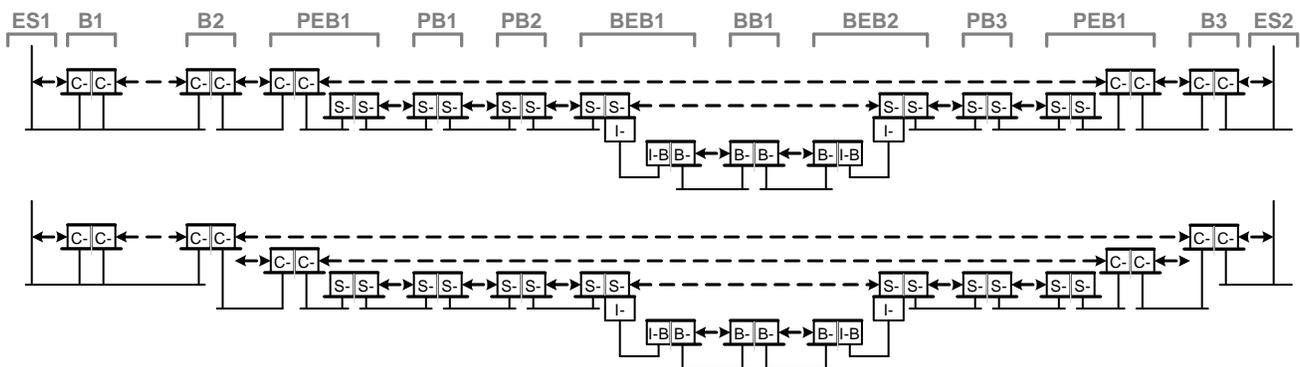


Figure 4—Address scopes in a Bridged Local Area Network

Figure 4 depicts a cross-section through a Bridged Local Area Network, using a simplified interface stack diagrams to illustrate the various address scopes provided by [Table 1](#). TPMRs (Two Port MAC Relays) and the corresponding scope (bounded by the limits of real physical media and not passing through any bridge component) are not shown. Two sets of interface stacks are shown, with the lower depicting the additional scope that can be provided if the C-VLAN components of the Provider Edge Bridges

forward the Nearest Customer Bridge group address (see footnote 3 to [Table 1](#)). Backbone Edge Bridges (BEBs) do not currently support a similar scope for the directly attached Provider Bridges (PBs), and this could be consider a deficiency of the current standards since BEBs are necessarily administered by the backbone network provider while the Provider Bridges might well be administered by an entirely separate organization.

3. Security threats and requirements

3.1 Protecting the infrastructure

As 802.1AE-2006 says in its opening paragraph:

“IEEE 802@ Local Area Networks (LANs) are often deployed in networks that support mission-critical applications. These include corporate networks of considerable extent, and public networks that support many customers with different economic interests. The protocols that configure, manage, and regulate access to these networks typically run over the networks themselves. Preventing disruption and data loss arising from transmission and reception by unauthorized parties is highly desirable, since it is not practical to secure the entire network against physical access by determined attackers.”

and (later in 1.1 Introduction):

“MACsec protects communication between trusted components of the network infrastructure, thus protecting... network operation. MACsec cannot protect against attacks facilitated by the trusted components themselves, and is complementary to, rather than a replacement for, end-to-end application-to-application security protocols. The latter can secure application data independent of

network operation, but cannot necessarily defend the operation of network components, or prevent attacks using unauthorized communication from reaching the systems that operate the applications.”

Thus, although MACsec can provide confidentiality, it has more to do than just hiding the data transmitted and received on behalf of network users from prying eyes. Indeed, because end-to-end transmission is usually supported by IP, there is no way that security at or just above the MAC layer could ensure that user data is accessed by only the original transmitter and the final receiver.

3.2 Protecting all frames

A significant motivation for the standardization of MACsec was the desire to avoid the need to design a protocol-specific security mechanism for each and every protocol used as part of network control and configuration. It can be readily appreciated that such a protocol-specific security approach would most likely lead to a delay in the development of the necessary protocols, or in forcing a choice on every network designer between being able to use the latest (but insecure) or secure (but older) technology.¹⁰

MACsec hops

Satisfying the desire for MACsec to be capable of protecting all our MAC layer control protocols, including those yet to be designed, without serial development or deployment delays imposes additional requirements:

- the MACsec specification itself needs to remain unchanged when it is to be incorporated into a new type of system (or network);
- the specification of the new system needs to naturally provide the right interfaces and opportunities for the addition of MACsec, even if no thought has been given to the subject by the system's designers;
- existing MACsec capable systems, within the network or available off the shelf, should not require modification to operate in a network with the new control protocols or systems unless they explicitly need to use the new protocols.

One consequence of the above is that MACsec needs to protect all the traffic transmitted and received on a given LAN, if the traffic is to be protected at all. A more or less general way might be devised to subset this protection (identifying those protocols that do have their own security mechanisms and skipping protection for those frames, for example) but would not reduce the frame protection and validation performance requirements (high throughput, very low delay) since such frames constitute a very low percentage of the potential load.

3.3 Data frames can have control effects

The point has to be made that if (as is the case for the vast majority of bridges¹¹) a bridge learns from the source MAC address of forwarded frames, then each frame forwarded is de facto a control frame, potentially altering the configuration of the network, as well as being a data frame. If such frames are not validated by the forwarding bridge, an attacker with LAN access can selectively deny service by transmitting frames (with a source address that is being used by legitimate traffic) on the 'wrong' LAN. A crude DoS attack, aimed at simply denying all service and possibly carried out by sending large

number of frames to overwhelm a switch's control processor, might be easily detected, but a learning attack offers more possibilities. The attacker might, for example, attack a source only when it transmitted frames secured by IPsec (or some other protocol) and thus persuade the frustrated user to turn security off.

NOTE—This attack could be carried out with frames with an Ethertype reserved for an 'end-to-end' authentication protocol if individual LANs on the path are not secure—existing bridges will (and should) learn from the source addresses of such frames.

3.4 Protecting part of the path

While MACsec cannot protect end-to-end if IP routers lie along the path, the requirement is often to protect only part of the path—even if each end uses the MAC address of the other directly as the destination in the frames it transmits. If part of the path is known to be physically secure (within a cage in a co-location facility, for example) there is no particular need to require MACsec capability on the end equipment (which might be a router without MACsec capability, for example). In such cases there is a positive requirement to protect only those LANs in the path that an attacker might be able to access¹². In general a network might comprise a number of trusted regions, each under the secure control of a single administration, connected by LANs or LAN services (such as provider bridged network) that may be controlled by a different administration and that are not (or are not trusted to be) physically secure. Requirements naturally arise to secure particularly exposed LANs in any network, to authenticate and secure connectivity between different administrations, or to secure connectivity 'end-to-end' where the ends are those of a path provided by a subcontracted administration and each 'end' of that path lies within equipment administered by the same organization—be that the organization providing the connectivity or the organization using it. IEEE 802.1AE–2006 Figure 11–12 provides some examples.

While end-to-end protection can be provided by IPsec, the use of IPsec tunnels can be inconvenient and over elaborate, particularly when the requirement is to provide protection for a network where much of the traffic can safely travel in the clear. A large campus

¹⁰This would mirror the early experience of those wishing to use MIBs as part of operational practice (i.e. to really manage their networks) in an era of MIB development as a separate arcane skill, quite separate from the rest of system and network protocol design. Under these conditions MIB development only starts when everything else is almost complete—ensuring that the first release product is only fully manageable through the console interface.

¹¹Backbone Bridges that only support PBB-TE or that only support shortest path operation using ISIS-SPB would be an exception, but even in a backbone one or two B-VLANs are likely to be dedicated to providing local management connectivity or other services that use station location learning. Even in 'exclusively routed' networks learning bridges (switches) can be found, playing a valuable (if largely transparent) role expanding interface port counts etc.

¹²This definite requirement has been a problem for proposed layer 2 and-to-layer 2 end schemes, such as the (never deployed and now withdrawn) 802.10 'interoperable LAN security'. In that standard bridges that sought to terminate the scope of protection on a path had to acquire the (secret) cryptographic keys from the end stations that authenticated their mutual communication, with the burden of a large number (potentially thousands) of keys in bridges that connected trusted and untrusted regions of the network. These keys might have to be acquired in a hurry if a network reconfiguration resulted in end to end paths traversing different bridges.

MACsec hops

with just a few exposed links between building or between clusters of buildings provides one example where it may prove much simpler to protect traffic on those links without constraining the network design.

3.5 Confidentiality

Today much traffic that should be secured is protected by https. This leaves IP addresses and port numbers, and usually the domain name (though not the detailed URL for the accessed web page) exposed. Since quite a lot can be inferred by observing the domain name

and packet exchanges and sizes, this is a real issue¹³. Keeping this detail confidential on highly exposed links is useful.

In some cases confidentiality is far more important than delivery (if the network itself is under attack the user will choose other means of delivering data) or the operation of the network is the responsibility of a separate organization, and is secured independently of the user data conveyed. Such cases are discussed later in this note.

4. Securing connectivity

This section discusses how connectivity is to be secured, given the bridging architecture, systems, and networks (2. above), and the threats and requirements (3. above). It reviews relevant aspects of the available cryptographic technology (4.1), particularly the use of secret keys, before considering how those keys are to be used to protect frames (4.2), and how such protection is to be included within the network and system architecture (4.3).

4.1 Cryptographic technology

At the data rates of interest, from a few megabits per second to hundreds of gigabits or even terabits per second, the only feasible security is provided by per frame symmetric (secret) key cryptography. This uses a block cipher together with a ‘mode’ of operation that allows the successive use of that block cipher to protect variable length data (such as a frame or packet). IEEE 802.1AE and 802.1X have been deliberately designed to allow the addition of Cipher Suites (each specifying MACsec’s use of a block cipher and mode) to take advantage of future developments in cryptography. At present all conformant Cipher Suites use the AES block cipher and the GCM (Galois Counter Mode)¹⁴ as documented by NIST. The crucial characteristic of GCM is that it is relatively efficient and parallelization, enabling high throughput implementation (to and above 100 Gb/s).

AES does not use the secret key directly, but expands it into a number of keys, one for each round¹⁵. Eleven round keys are needed for 128-bit encryption, and fifteen for 256-bit. High performance GCM–AES implementations store¹⁶ the round keys for each secret key in active use, rather than recomputing them for each protected message. Thus, when a GCM–AES implementation has to authenticate (and possibly decrypt) a received frame, it needs rapid access to more data than is carried in a minimum sized Ethernet frame. The fact that it is neither desirable nor practicable to control (either on transmission or reception) which of the currently in-use secret keys will be required by the next frame places a high premium on keeping the maximum number of such keys low. If the number is sufficiently low, all secret-key and other secure association specific data can be kept with the logic¹⁷ that is to perform the frame protection and verification. The alternative is to increase the logic’s external memory bandwidth requirements, perhaps reducing the number of network interfaces supported.

While it is notoriously difficult to advance implementation dependent arguments in standardization efforts, IEEE 802.1AE aims at widespread deployment in LAN switching equipment. This goal can (on past experience) be met only if the incremental cost of including the security capability, and the performance impact of using it, is negligible. So it had to be possible to incorporate MACsec within

¹³A recent newspaper article reported an academic study suggesting that a fair amount of information on financial accounts, including the number of them held at a particular institution and the approximate balance of each, could be deduced from this exposed information.

¹⁴NIST SP 800-38D. GCM–AES is also used in IPsec, SRTP (Secure Real-time Transport Protocol), and Fibre Channel Security Protocols (FC-SP) amongst other uses. http://en.wikipedia.org/wiki/Galois/Counter_Mode provides a useful introduction. Many detailed implementation studies are publicly available.

¹⁵See http://en.wikipedia.org/wiki/Advanced_Encryption_Standard and http://en.wikipedia.org/wiki/Rijndael_key_schedule for a brief description.

¹⁶This description may over simplify, but contrasting the 176/240 octets (for 128 and 256-bit secret keys respectively) of the round keys against the 96 octets of a minimum sized-packet (with MACsec SecTAG and ICV) makes the point. A small amount of additional key-specific (or at least SA specific which amounts to the same thing) management data is also needed.

¹⁷Perhaps immediately accessible through multiplexing logic, or close by on chip with wide bus access.

MACsec hops

a design that meets the cost/performance goals of a purchaser who, initially at least, has no interest in MAC security. Avoiding any need for the equipment vendor to offer, much less design, distinct security capable variants is vital. MACsec satisfies this essential requirement. The capability has been included in a number of Ethernet interface devices and switches, irrespective of whether it is to be deployed by the end user or not, simply to avoid proliferation of variants. The performance of these implementations, both in terms of wire speed throughput and low latency (of the order of a minimum packet size¹⁸) has kept pace with increases in transmission speed. They necessarily protect or verify frames as they are being transferred in and out of memory (or between memories). Otherwise memory bandwidth would have to be greater than that required for non-MACsec capable implementations, as an inevitable consequence of encrypting or decrypting (and thus changing) frame data as well as adding or removing the SecTAG (security header) and ICV (trailer). This increase would be in addition to any due to using large numbers of secret keys¹⁹.

The fact that MACsec processing can be incorporated in high volume and cost effective network interface components, not intruding upon the rest of an end station's or switch's design and making implementation relatively easy, also lessens the likely availability of alternate processing components²⁰. This may make it more difficult to insert a MACsec shim at all permissible points in an interface stack (see 4.3) in complex systems: while it may be possible to carry instructions as to where and how to insert a header in the frame, this demands significant flexibility both within the interface implementation and the rest of the system, transferring system intelligence to the interface. Attempts to procure custom MACsec capable systems may reasonably be met by refusals or special engineering costs, even if the required interface stack arrangements meet 'the letter of the law' for use of the MACsec shim. On the other hand the functionality of complex systems, e.g. Backbone Edge Bridges (BEBs), can often be realized economically by separate connected systems, mirroring their standard specification. In all events we attempt to work toward a short catalog of useful systems and interworking arrangements.

¹⁸The (small) added delay can be made quite predictable. This low jitter can be important in time sensitive networking applications. The parallelization capability of GCM means that it is entirely feasible to use a pipelined implementation to maintain full wire speed for indefinite back-to-back minimum frame size transmission/reception, though the usual 65 byte frame considerations may make this a commercial consideration.

¹⁹At the time of writing a few ten's of keys is a typical upper limit. This is sufficient to support virtual ports on a shared LAN serving end stations.

²⁰Though not of gate designs for inclusion in custom chips. High performance end station implementation using Intel processors has been facilitated by the new instructions specifically targeted at GCM-AES.

²¹This is a vital part of a bridge's handling of network reconfiguration or end station movement.

Implementing LAN security in a way that imposes the lowest possible cost on those who are either uninterested or marginally interested also means paying attention to detailed security requirements, or sometimes the lack of them. See 3 above. In particular links directly connected to end stations may not be exposed, so the task of deploying on the most numerous devices can be avoided or postponed.

4.2 Protecting frames

Consider the network fragment shown in Figure 5 (a). End stations ES1 and ES2 are connected by VLAN bridges B1 thru B3. Each of the connecting LANs/links is shown in red, to indicate that they are exposed to attack, carrying data that ought to be protected. How should this be done?

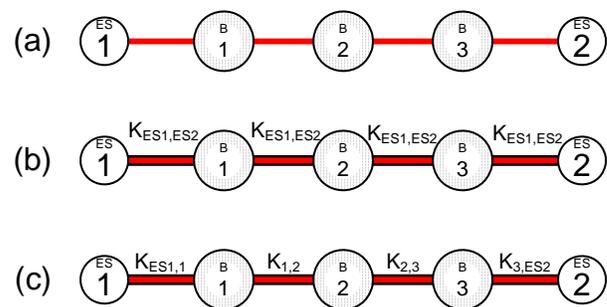


Figure 5—Protection along a path

An 'end-to-end' approach (Figure 5 (b)) protects the frame (shown in the figure by covering the links along the path with black) with a key ($K_{ES1,ES2}$) agreed between (and known only to) the two end stations. This has a few downsides:

- The bridges cannot check the frame integrity, and thus safely learn²¹ (or refresh) station locations from the MAC source address (3. above), unless each of them also possesses $K_{ES1,ES2}$.
- A bridge cannot make (and protect) a change to the frame that would be permitted in (and may be essential to the operation of) an unsecured network, unless it also possesses $K_{ES1,ES2}$. For example, ES1 might transmit a frame without a VLAN tag, with B1 adding it (with the appropriate VLAN ID). Other stations, in different parts of the network, might be assigned to different VLANs, with all these frames (including their VLAN IDs) being destined for the ES2. In one typical network arrangement ES2 is a

MACsec hops

router, and the VLANs correspond to IP subnets. Two stations, on different VLANs/subnets, might have the same MAC address. The assigned VLAN IDs have to be protected if one station is to be prevented from masquerading as another.

MACsec uses the ‘hop-specific’ approach shown in Figure 5 (c). A different key is used for each hop (LAN), with each participant validating the frame using the key for the reception LAN, and reprotecting the frame for onward transmission. While receiving station has to trust not only the transmitting end station, but also the intervening bridges this is also the case for the end-to-end approach (b)—once the latter is extended to permit and protect normal bridge functions (learning, VLAN assignment,...). In both cases a secure infrastructure has to be established, and the users of the network provided by that infrastructure need to trust it.

The difference between these approaches is readily apparent when communication from ES1 to an additional station, ES3, is considered (Figure 6).

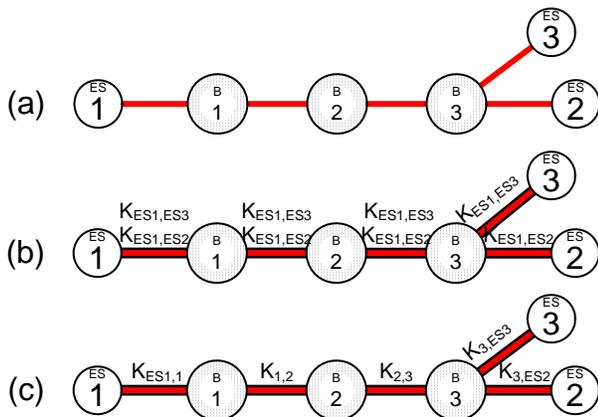


Figure 6—Protection to multiple stations

The end-to-end approach (b) requires a key for each communicating pair of stations, and each such key needs to be known by intervening learning bridges or by any bridge that needs to modify the frame²². The hop-specific approach (c) does not.

The hop-specific approach also facilitates incremental deployment. Initially it may be important to secure one link in the network (B3–B4 in Figure 7 (a)), while

others are considered immune from attack even if not explicitly secured (shown as green in the figure).

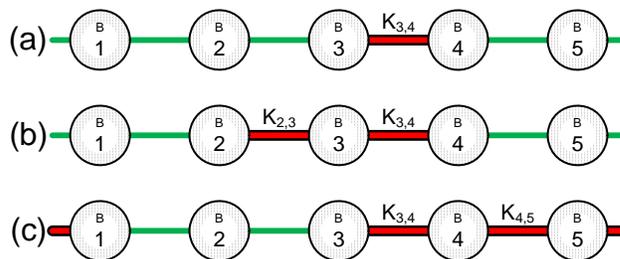


Figure 7—Protected and trusted LANs

The protected region or portion of the network path can then be extended, as in Figure 7 (b), though protected traffic within the region cannot be trusted unless appropriate controls/policies are applied to traffic entering it. If part of the network is truly immune from direct interference (the LANs connecting B1, B2, and B3 in Figure 7 (c) might be completely contained in a locked closet, for example) then they can form part of the trusted region.

The case where the physical connectivity to B2 is actually exposed to attack, and only unprotected because B2 itself lacks the capability, is more difficult. If B2 does not modify any of the frames it forwards (and those it originates are readily identifiable and subject to sufficient ingress policy controls by the adjacent bridges) it is tempting to protect the path from B1–B2–B3 with a key agreed by B1 and B3 (Figure 8 (a)). However this can pose problems. Almost all network designs provide alternate paths to protect against device or link failure (as in (Figure 8 (b)). A failure of the link B2–B3 should divert traffic from B2–B3–B4 to B2–B30–B4, and may well be supported by rapid reconfiguration protocol mechanisms (aimed at meeting or bettering a 50 millisecond service restoration time). Notifying B1 of the failure and having B1 and B3 agree and install the new key²³ is not currently part of such mechanisms and is unlikely to fit within the time budget.

However if B2 really does not modify forwarded frames it may be possible to treat it as operating at a lower (sub-)layer, as a Provider Bridge (Figure 8 (c), for example). In that case B1, B3, and PB30 are really one hop apart from the point of view of their network configuration protocols—they are all attached to the (virtual) shared medium supported by PB22—and can

²²Addition of a tag is not the only potential modification. The priority bits in the tag can also be changed as part of normal class of service handling. After such a modification the frame would need to be reprotected with a key acceptable to the recipient.

²³In some cases the path B2–B30–B4 would not exist prior to the failover while the required key cannot be agreed until it exists, in other cases the path only exists for the purposes of supervisory traffic, which would not naturally report to B1, and in any case does not support key agreement protocol.

MACsec hops

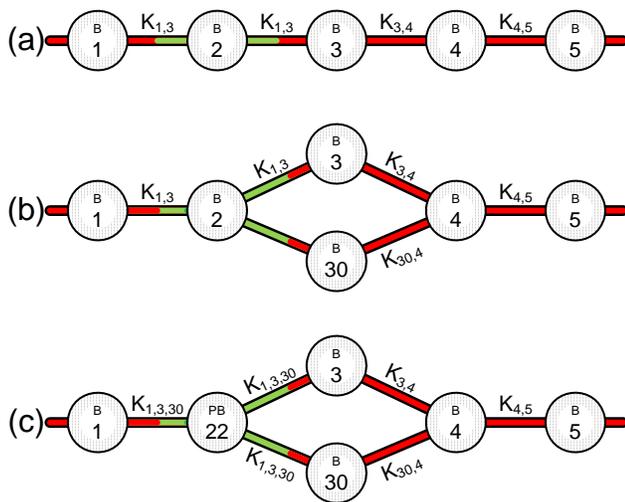


Figure 8—Network reconfiguration

agree a group key. Similarly (though not shown in the figure) B4 could also be treated as a Provider Bridge.

Two systems may not be immediate neighbours for an instance of a configuration protocol in which they both participate, but the path protected by that configuration protocol may be constrained to pass between the two systems—if it passes through either of them. In that case it may be possible to omit the intervening systems from the configuration protocol, effectively placing them at a lower (sub-)layer, as in the above, and making the systems immediate neighbours. Simply forwarding Nearest Customer Bridge group addressed frames (see [Table 1](#)) through the C-VLAN components of Provider Edge Bridges has that effect.

4.3 The MACsec shim

The paradigm of connectionless networking, in which communicating peers can exchange data without previously participating in an explicit exchange to setup a connection (as required by X.25 or TCP), is now so prevalent as to pass without comment²⁴. The notion of a ‘connectivity association’ as an a priori association between communicating peers—that is to say an association that is assumed to exist and that has been created without their explicit knowledge²⁵—remains useful. What concerns us is the connectivity association between neighbouring peers at a given (sub-)layer—the simple ability of a set of protocol entities to exchange frames without the

need for the frame to be relayed at that (sub-)layer. This is the ‘single hop’ that we wish to secure, and in general we wish to secure it without the explicit involvement of the communicating peer protocols—otherwise we would fail in our goal (see [3](#) above) to keep pace with the new protocol development that continues largely independently of security concerns.

This secure connectivity association (CA, as defined by 802.1AE) formalizes our notion of a secured instance of the ISS, defining its extent and participants. The connectivity association implied by the existence of the ISS at any point in a protocol interface stack can be secured by the insertion of protocol entities whose operation is transparent to that of the existing protocol entities above and below. Such transparent protocol entities are known as ‘shims’. Figure 9 (compare to Figure 2) provides an example.

Three connectivity associations (ES1–B1, B1–B2, B1–ES2) have been secured by the addition of the MAC Security Entities (SecYs). In addition to cryptographically protecting frames that pass between its upper (Controlled, or secured) port and its lower (Common Port), each SecY also supports transmission and reception of unprotected frames (through an Uncontrolled Port) so that companion protocol entities (PAEs and KaYs) specified in 802.1X–2010 can authenticate, reauthenticate, and agree keys with the other participants or potential participants in the CA²⁶.

4.4 CA Scope

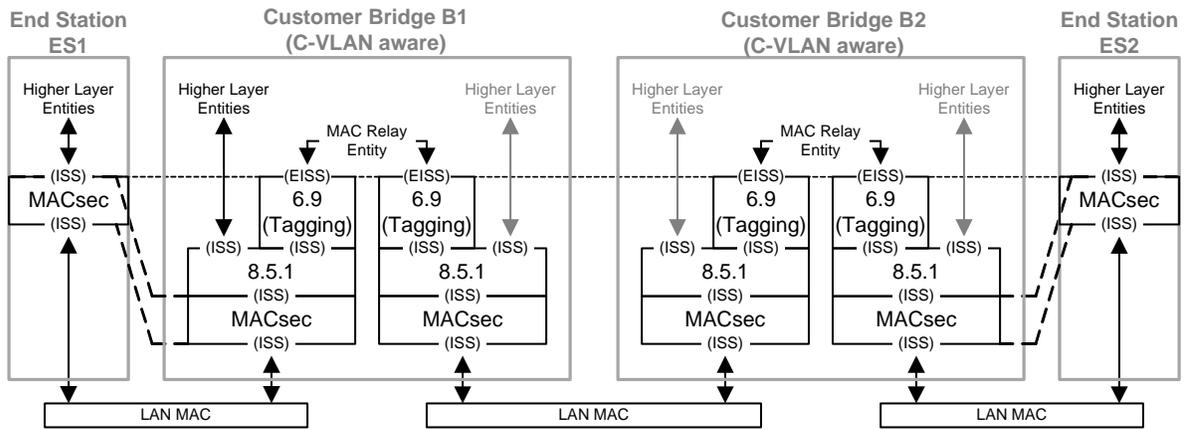
The scope of a CA (if connectivity is permitted at all) is thus determined by the scope of the addresses used for authentication and key agreement. Management controls for the SecY determine whether insecure connectivity is permitted, or indeed whether the connectivity is to be secured at all. Moreover received frames are (at least notionally) passed both to the Controlled Port (for possible relay, if the system is a bridge) and the Uncontrolled Port (for use by authentication and key agreement). The destination address of the frame determines whether the relay’s FDB will allow forwarding, and whether the PAE or KaY (Key Agreement Entity, see 802.1X–2010) will wish to process the frame.

²⁴See RFC 787 for a useful tutorial.

²⁵The term ‘a priori’ does not simply mean ‘prior’, nor is it restricted to discussions of probability. See the Wikipedia discussion of ‘a priori’ (knowledge independent of experience) and ‘a posteriori’, and A.C.Grayling’s ‘An Introduction to Philosophical Logic’ (Chapter 3). In the present case examples of events and actions outside the media and subnetwork independent experience of protocol entities include plugging an Ethernet cable into a network, and setting up an ATM connection that will subsequently carry UDP packets. All that the protocol entities know is that they can, once active, transmit packets.

²⁶Clearly no participant is to be given a key to participate in secure communication until mutual authentication has taken place. Authentication thus either implies, or is followed by, authorization. Authorization may result in changes to the management variables of other protocol entities—permitting or denying access to certain VLANs, for example.

MACsec hops



NOTE—This figure has been simplified to show the addition of MACsec functionality as a single shim, comprising both the MAC Security Entity (SecY) that carries out the frame protection and validation operations, and the associated PAE and KaY (see 802.1X-2010) that facilitate authentication and key agreement. The latter make use of an Uncontrolled Port provided by the SecY. This figure and similar figures in this note show only each SecY's Common Port (below MACsec) and its Controlled Port (above MACsec) and allows communication between peer PAEs and KaYs to be shown as communication between peer MACsec entities.

Figure 9—Securing the ISS in VLAN-aware bridges and end stations

Consider the connectivity between Customer Bridges B1 and B2 in Figure 9. The MACsec PAE and KaY for B1's right-hand port, and those for B2's left-hand port, can be configured to transmit and receive using either the Bridge Group Address (also known as the Nearest Customer Bridge group address) or the PAE group address (also known as the Nearest non-TPMR Bridge group address). Either will work in this scenario: both are filtered by the MAC Relay Entities of B1 and B2, restricting EAPOL (EAP over LANs) and MKA (MACsec Key Agreement) exchanges to the two ports.

Replacing the LAN connecting the two ports with a provider bridged network service²⁷ introduces a further possibility. The intent may be to secure connectivity across the PBN, between B1 and B2, as in Figure 10. In this case each PAE and KaY should use the Bridge Group Address for EAPOL exchanges (if required) and to agree keys (using MKA). Frames with this address will be forwarded by the MAC Relay Entities of the intervening Provider Bridges.

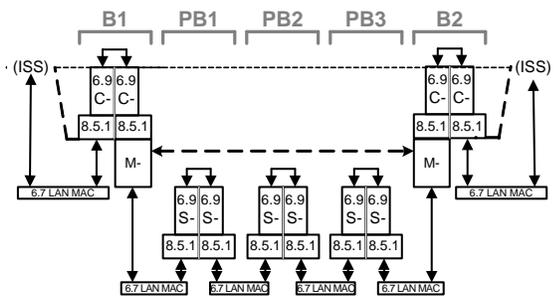


Figure 10—MACsec across a PBN

Alternatively the intent may be to secure connectivity to the provider bridged network, from both or either of B1 and B2. See Figure 11. In this case the Nearest non-TPMR Bridge group address should be used, and EAPOL and EAP exchanges occur between each Customer Bridge and the Provider Bridge that supports its interface to the PBN.

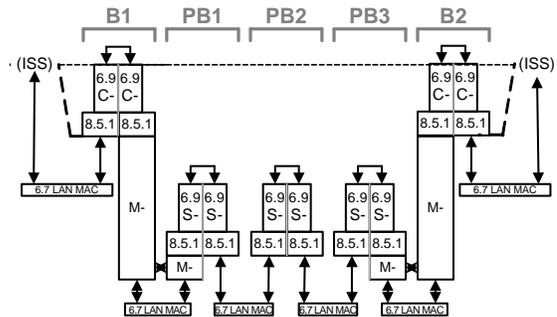


Figure 11—MACsec to a PBN

As a further alternative, the MACsec shim (and its associated control protocol components) may appear twice²⁸ in the Customer Bridge Port's interface stack: securing communication to the provider's network and (separately) securing communication across it as described in 802.1AE 11.7 and Figure 11-13.

4.5 MACsec end-to-end

It is possible (in theory at least) to operate MACsec end-to-end, and this possibility is not infrequently suggested by those new to MACsec. This note is largely about the reasons why this is rarely a good idea. It does make sense when the 'end-to-end'²⁹

²⁷Port-based service interfaces are shown to simplify the diagram and discussion.

²⁸I do not know if this configuration is commonly supported by single systems (the considerations outlined towards the end of 4.1 would suggest otherwise) but access to the PBN could easily be provided by a separate MACsec capable NID (Network Interface Device). If the NID design follows 802.1 standards it is likely to be a TPMR, in which case EAPOL and MKA might well use the Individual LAN Scope group address (see 802.1Q-2012 Table 8-1).

MACsec hops

connectivity can be considered equally to be a single hop, as described above. Very briefly: the traffic to be protected flows between the two ends or not at all (specifically there is no reason for sending it from either end if it is not to arrive at the other), and is not destined for (nor has any effect upon) any other system on any potential path from one to another. These criteria are met by customer traffic when PBN point-to-point services³⁰ or PBBN services are used—the traffic is segregated by S- or B-VLAN. Thus the reconfiguration protocol concerns raised in [4.2](#) above are handled by ensuring that the ends are not affected, and MACsec (or at least that instance of it operated by the ends) does not have to play its primary role of protecting the network infrastructure between the ends. At best that protection is now achieved by the usual careful configuration of provider services and the possible use of further layered instances of MACsec within the provider network, see [4.4](#), at worst it has simply been declared an SEP (someone else's problem).

Of course this end-to-end approach may still result in one or both of the end systems having to handle a very large number of secured end-to-end connections, each with its own secret keys. Whether that is feasible depends on the relative effort of supporting those connections and the rest of the system's workload.

4.6 MACsec end station-to-end station

The end systems for end-to-end connectivity could, so far as the MACsec protocol itself is concerned, be the MAC destination and source of the frame. This satisfies the reconfiguration protocol concern (the frame is constrained to arrive at the destination, or not to arrive at all). However it leaves open the question of what (unprotected) protocols are to be used so that the ends can discover each other and decide to set up the secured connection. Broadcast and multi-cast frames would have to be protected separately, if protected at all, so some further higher layer by higher layer protocol threat analysis would be required. Additional

destination address and protocol type specific mechanisms and controls would need to be specified to assign frames to the Uncontrolled or Controlled Ports and to Secure Channels (SC)³¹.

MKA could easily use individual rather than group destination addresses (once the addresses have been discovered) but the use of EAP and EAPOL seems unlikely, if for no other reason than that the scaling mechanisms developed to support EAP Authenticators are no longer a natural fit. Use of a Kerberos/IKE based alternative would maximize the benefit of past experience, and indeed if only IP traffic is to be protected it should be possible to borrow the design wholesale from the control protocols used to support IPsec. The question is, of course, what would be the point? Such a web of point-to-point MACsec connections would not protect the MAC layer switching infrastructure itself (without an additional hop-specific MACsec sublayer, as mentioned as a possibility for PBNs above, in the last paragraph of [4.4](#)) and is unlikely to perform significantly better than IPsec. Each of the end-points of each MACsec protected part of a complete IP end-to-end path would necessarily be an IP router or IP end station, removing the flexibility to use intermediate devices for protection and to protect just exposed links—the two capabilities that really make MACsec attractive.

There is a further potential issue with attempting to use MACsec end station-to-end station at present. Unless there is some intervening tagging (C-VLAN, S-VLAN, or PBBN address encapsulation) or a further sublayer of MACsec is being used to protect each hop, then any SecYs in the intervening bridge port interface stacks will discard the protected frames (or at best remove their SecTAGs and ICVs) as their SAs will not be recognized. This is not an issue for C-Tagged or S-Tagged service interfaces (see [5.1](#), [5.3](#), and [5.4](#) below) but may affect the use of port-based service interfaces. For further detail see [B.1](#) below.

²⁹The ends of the 'end-to-end' connectivity may refer to ultimate source and destination of the communication, probably coincident with the ends as viewed from IP. Contrariwise in telco-speak the 'ends' simply refer to the points beyond which an organization responsible for carrying traffic no longer has financial or legal responsibility.

³⁰PBN point-to-point services do not have to learn from frame's source addresses. The criteria are also met for all PBN services when all the frames accepted by the provider have been protected over the access link, denying an attacker the opportunity of spoofing MAC source addresses. PBBNs provide address independence via encapsulation. Traffic engineered services may provide other opportunities for connectivity to be declared 'effectively single hop'.

³¹So the (significant) parts of 802.1AE that deal with MAC Security Entity (SecY) operation and its management are not directly applicable.

5. Additional use cases

Clause 11 of 802.1AE–2006 specifies how MAC Security is incorporated within the architecture of end stations, within systems that incorporate link aggregation or the Link Layer Discovery Protocol (LLDP), within VLAN-unaware and VLAN-aware Bridges, and within Provider Bridges. It shows how to secure connectivity between Customer Bridges attached to PBNs (independently of PBN operation) and how to secure connectivity from a Customer Bridge to the first Provider Bridge in such a network, but only for port-based interfaces to the PBN. It also shows how to provide independently secured access for multiple end stations connected to the same LAN. Clause 7 of 802.1X–2010 provides additional detail.

It might not be apparent, to those not familiar with 802.1Q's specification of provider bridged networks, how 802.1AE 11.7 applies to tagged as well as port-based service interfaces, so this section provides additional detail (5.1), before considering the following additional cases:

- PBBN port-based service interfaces(5.2)
- PBBN S-tagged service interfaces(5.3)
- PBBN I-tagged service interfaces(5.4)

5.1 Tagged PBN service interfaces

The following have to be taken into account when designing the secure connectivity solution:

- a) The C-VID³² is naturally behind the SecTAG³³, and thus may be encrypted and in any case will not be accessed by an ordinary Provider Bridge³⁴.
- b) Any given provider service instance/connection across a PBN is not constrained to use the same type of service interface at both/all interfaces. An organization's central office may, for example, use a tagged interface to distinguish between connections to a number of branch offices. The branch offices might use a port-based interface, and in any case the interface equipment has no need to understand the central office's tagged interface numbering scheme.

A natural approach to these problems is for security conscious customers to use an S-tagged provider service interface and supply their own provider edge bridge functionality, either in a separate or a combined

system, as in Figure 12 (compare to Figure 3). The S-TAG components (S-VID, PCP, DEI) can then be policed, modified, translated, or removed as required by the service provider, while the integrity of the customer's C-VID is guaranteed on delivery.

Some explanation of Figure 12, covering the basic operation of a PBN and its interfaces as well as the use of MACsec, may help: A plan view of part of the network, showing three service instances connecting customer owned and operated bridges and Provider Bridges, is shown below an interface stack diagram. The latter necessarily follows the path taken by just one service instance—between B1 and PEB2—though (in this case) only the identities, rather than the interface stacks, of some components would differ if another service instance were to be depicted. On the left the MACsec processing and the addition of the S-VID that identifies is shown split over two systems (B1 and PB1), while on the right a single PEB3 provides both C and S component functionality³⁵. The header format for each frame transmitted, from B1 on any of the service instances, is also shown.

Frames for several different Customer VLANs (C-VLANs) can be carried on any one of the provider connections (since each of these is providing a point-to-point service there is no need to learn from the source addresses of customer frames within the provider network, and hence no need to speculate as to whether the customer's MAC addresses are unique amongst C-VLANs carried by a single service connection). Following a frame from left-to-right, it is:

- Forwarded by B1 to just one of the right-facing ports of B1, as the FDB is configured to permit egress for the frame's C-VLAN only on that port and one or more of the left-facing ports of B1 (not shown).

If that C-VLAN is the only one to be carried on the connection and the C-TAG is not required to carry the frame's priority it could be removed and the frame passed down the stack and across the connection untagged.

- Transmitted through one of B1's ports³⁶ protected by, and with the SecTAG provided by, the

³²The Customer's (C-VLAN's) VLAN Identifier.

³³The MACsec header.

³⁴Either MACsec unaware or with a MACsec interface stack as specified by 802.1AE–2006.

³⁵This difference has been introduced merely to indicate this possibility, the relative positions of B1+PB1 and PEB2 could equally well be reversed, or both ends of any provider service instance/connection over the same network could be supported by separate (or combined) systems.

³⁶In general 'port' refers to the interface stack associated with an ISS interface, though it is often convenient to emphasize the associated processing by using the term 'interface stack' and to emphasize the physical connectivity by using the term 'port'. See 802.1X-2010 D.4 or 802.1Q–2011 6.1.4

MACsec hops

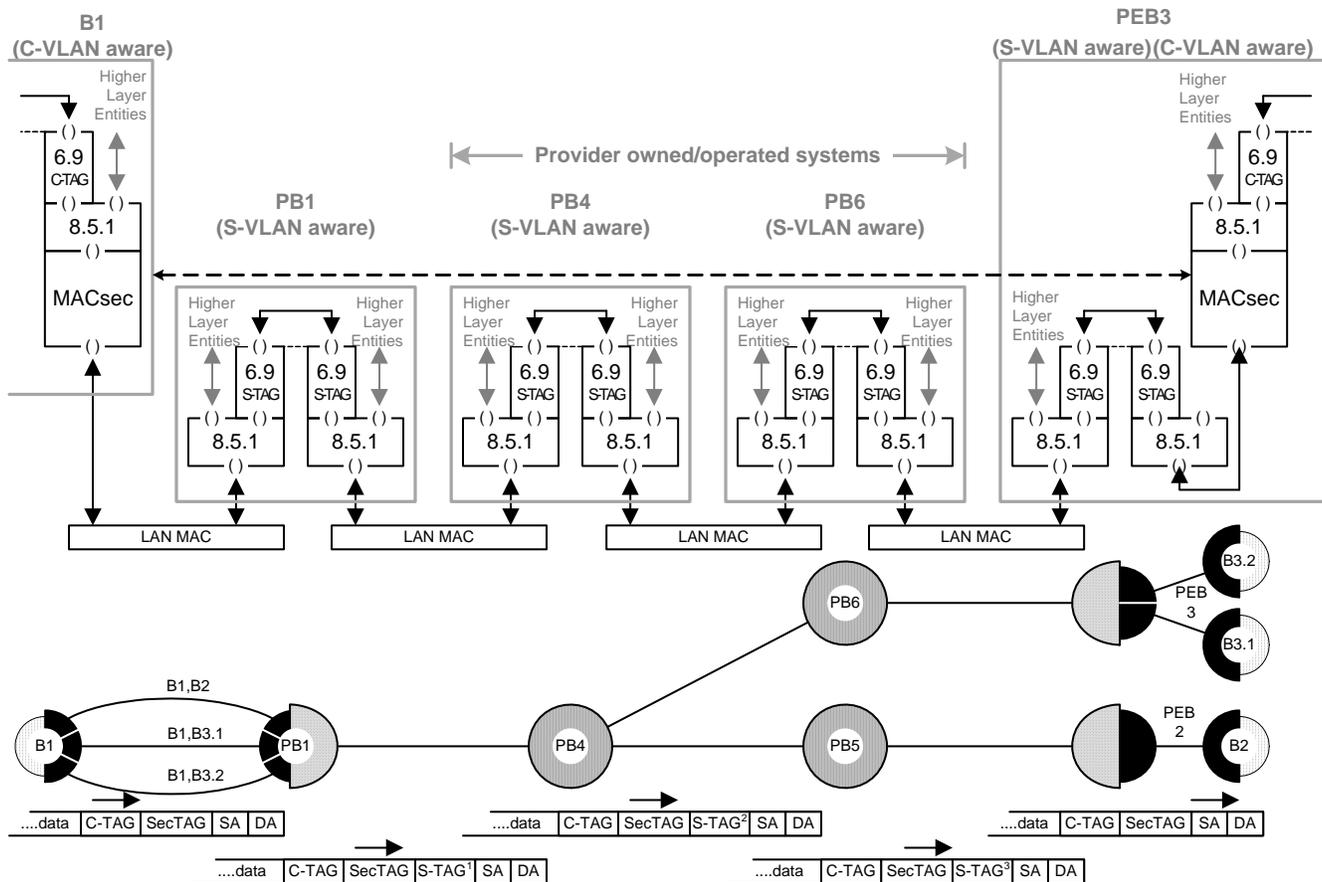


Figure 12—Interface stacks and paths for PBN tagged service interfaces

MACsec shim for that port onto one of the three LANs connecting B1 to PB1.

- Received by PB1 and assigned to an S-VLAN using the receiving port's PVID³⁷. Thus B1's egress port decision has identified the connection for the frame on the basis of the latter's C-VLAN.
- Forwarded by PB1 to the PBN, and through the PBN on the basis of the S-VLAN to PEB 3.

The S-VID value is policed and typically translated by PB4 on ingress to the PBN, allowing the provider to organize or reorganize the S-VIDs used within the network without impacting any customer.

- Forwarded by the S-VLAN component of PEB3 (on the basis of the S-VLAN) to one of its ports, which removes the S-TAG, and transmits the frame
The port concerned lies within the PEB and can be realized by any technology that meets the requirements of 802.1Q 6.14 (Support of the ISS within a system).

- to a port on C-VLAN component B3.2, whose MACsec shim validates the frame and removes the SecTAG.
- Forwarded by B3.2 to its other port(s) on the basis of its C-VLAN and destination MAC Address.

The PAE and KaY³⁸ associated with the MACsec shims of B1's and B3.2's ports communicate over the same S-VLAN tagged path, using the Nearest Customer Bridge group address which is not filtered by any of the intervening S-VLAN components (see [Table 1](#)), as the destination address of each frame.

The interface stack and path just described do not provide a way for B1 to communicate the ISS's priority parameter information to PB1 on a frame by frame basis unless that capability is an inherent part of LAN MAC—Ethernet lacks this. IEEE 802.1Q 6.13 (Support of the ISS for attachment to a Provider Bridged Network) remedies this deficiency by allowing the C-VLAN bridge component to priority tag each frame with an S-TAG³⁹. Figure 13 shows the

³⁷Port-based VLAN Identifier (the receiving port treats the frame as untagged).

³⁸The PAE (Port Access Entity) and KaY the Key Agreement Entity, part of the PAE, see 802.1X-2010)

³⁹Priority tagged frames convey have 0 in the VID field. Note that the interface provided by the 6.13 shim is the ISS, not the EISS.

MACsec hops

interface stack for each of the C-VLAN bridge components of Figure 12 with the addition of this capability. There is no need for the S-VLAN components of PB1, PEB2, or PEB3 to communicate this information to the C-VLAN components, as the latter will recover the original C-TAG's priority field from the frame after MACsec validation.

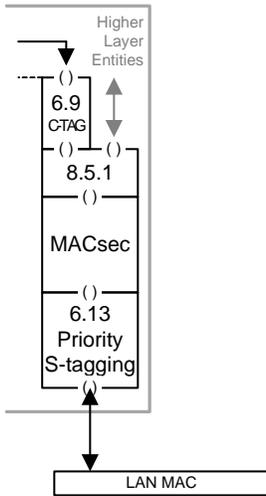


Figure 13—Communicating priority from a secured C-VLAN component

5.2 PBBN port-based service interfaces

These interfaces are described in 802.1Q 25.3 “The PBBN Port-based interface provides the same type of service as the PBN Port-based interface...” and illustrated in 802.1Q Figure 25-4. All S-tagged frames presented to the interface are required to have a null VID, i.e. from an S-VLAN components point of view they are untagged or priority tagged. A suitable interface stack for the attached customer equipment is that shown for B1 in Figure 12, or that shown in Figure 13 if priority selection is occur on a frame-by-frame basis. The frame is forwarded over the backbone without an S-TAG, though this is of little concern to the customer who delivers and receives his MACsec protected frame to and from the service just as specified for a PBN (see [5.1](#) above).

The backbone adds I-TAGs (conveying an ISID, or service identifier) and B-TAGs (adding encapsulating addresses) but these are outside the scope of the customer's MACsec protection and are removed

before the frame is delivered to the receiving customer equipment.

The Nearest Customer Bridge group address used by the PAE and KaY is encapsulated over the backbone, so authentication and key agreement proceed just as before.

5.3 PBBN S-tagged service interfaces

Again this interface (described in 802.1Q 25.4) behaves, from the customers point of view, in the same way as an S-tagged PBN service interface, and the interface scenario follows that for [5.1](#) above.

PBBNs do not offer a C-tagged service interface, so present none of the minor complications that arose in considering this interface for PBNs ([5.1](#)).

5.4 PBBN I-tagged service interfaces

A customer of a PBBN I-tagged service interface (802.1Q 25.5) is most likely to be the operator of a peer attached PBBN uses a Backbone Edge Bridge to connect. The ISID supplied by the customer is mapped 1-1 to an ISID within the PBBN, and is not otherwise carried through the PBBN. Since the intent is that the ISID supplied across the interface be changed, or at least open to change by the PBBN, there is no sense in which it is desirable to ensure its integrity when delivered at the distant end of the connection through the PBBN. A common reason for interconnection of PBBNs in this way is to allow two operators to extend their joint area of service availability and it is most likely that the PBBN will deliver the frame not to the I-tagged service interface user but to some common customer who has no interest in the particular ISID value, but may simply want to confirm (by MACsec operating above the level of the PBBN, and not protecting the I-TAG or B-TAG) that he has received a (possibly encrypted) frame that his authenticated peer originally transmitted.

Protection of communication directly between the BEB attached to the service interface and the PBBN itself is a different issue, and can simply be achieved over this very simple hop by using MACsec at the lowest layer in the interface stack as can be done for PBN interfacing (see 802.1AE-2006 Figure 11-13).

6. Follow-up

This note has pointed out areas in current standards that could be expanded or improved. Ideas on what might be done follow.

6.1 BEB S-VLAN component addresses

The S-VLAN components of Backbone Edge Bridges (BEBs) have the same reserved addresses as those of Provider Bridges (see 2.3 final paragraph), just as the C-VLAN components of Provider Edge Bridges (PEBs) have the same reserved addresses as Customer Bridges. However in the latter, C-VLAN case, the PEB component can simply forward frames for one of these addresses (the Nearest Customer Bridge group address) when provider network is supporting their transmission with a single provider service instance⁴⁰. This address can be safely forwarded, under this constraint, because the C-VLAN component is forwarding frames between just two ports and thus does not need to appear as a node in the topology as viewed by protocols using the address. Thus, in the C-VLAN case, the PAEs and KaYs of two customer-owned Customer Bridges can set up a single MACsec hop across the entire provider network, including the provider-owned PEBs. No corresponding reserved-address-forwarding exception exists for S-VLAN components.

As a consequence, PAEs associated with the provider network facing ports (see 802.1Q Clause 15) of customer-owned Provider Bridges or Provider Edge Bridges (PEBs) cannot use the Provider Bridge Group Address, or any other currently assigned Reserved Address (see Table 1) to communicate with their peers participating in the same provider service instance. Such frames would need to pass through the S-VLAN components of the intervening provider owned Provider Bridges or Backbone Edge Bridges (Figure 14 (a)).

Some real world justification is required for supporting this security use case — it is not an absolute given that administrative boundaries for security will align with the interworking interfaces specified for PBBNs. There is a strong prima facie case, and this note takes it as sufficient. Exactly the same requirement (for ‘single hop’ communication between some S-VLAN components through a network of other S-VLAN components) would arise if the provider backbone network service was actually provided by access provider networks, as illustrated by the (potentially) secured communication between PB1 and PB2 in Figure 14 (b).

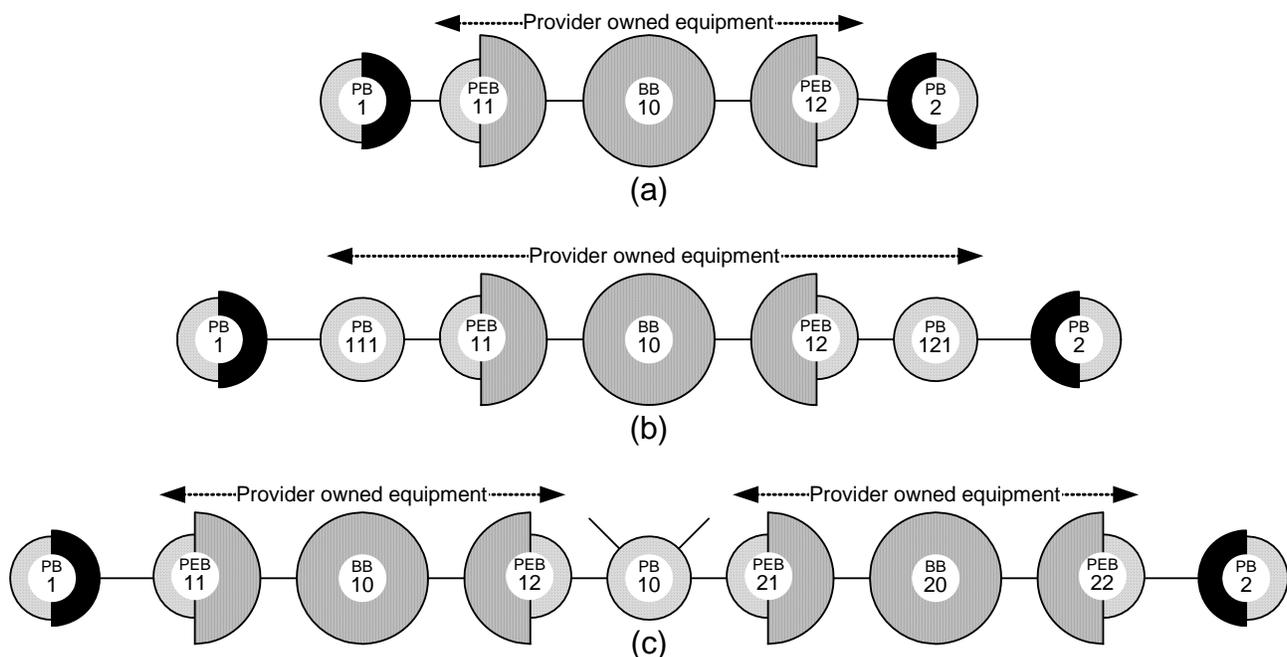


Figure 14—Secure S-VLAN connectivity across a PBBN

⁴⁰See footnote 3 to Table 1 for precise conditions.

MACsec hops

In all cases, the possibility of treating the provider (backbone) network as a single hop depends on that ‘hop’ functioning as a single provider service instance — a single CA⁴¹. The point is that, under this condition, the internal structure of the provider network can be modeled either as a point-to-point link or as ‘virtual’ shared media — all the protocols (control and data) used by bridges attached to the service are capable of handling this connectivity without requiring any active protocol ‘node’ (real or ‘virtual’) to be embedded in the service.

Five currently unallocated reserved addresses might be suitable for communication between the PAEs of PB1 and PB2 in Figure 14 (a) and (b). Three of these addresses are filtered by C-VLAN components but not by S-VLAN components. Two are filtered by both C-VLAN and S-VLAN components. However these addresses are a precious resource, and what is really wanted is an address that would be filtered by the S-VLAN components in the customer network, but not by those in the provider network. Such an address would avoid any possibility of accidentally creating the secure connectivity illustrated in Figure 14 (c)—from a customer attachment to one provider service instance, through an insecured customer network attachment to that provider service instance, through the customer network, through a further insecured attachment to a second provider service instance, and finally to a second secured provider instance. The requirement to use a reserved address to prevent this accidental, undesirable, connectivity is weak. Plug and play considerations are not paramount

in this application, and a perfectly ordinary group address could be allocated. That allocation could be standard or private, using existing group address allocation methods.

If EAPOL frames (sent to this ordinary group address) do accidentally discover the connectivity between PB1 and PB2 illustrated in Figure 14 (c), there are existing mechanisms that can prevent this being turned into secured connectivity. First, parameters used in EAP exchanges and methods can identify the undesirability of the connection. Second, if a CAK is derived from an EAP exchange or provided as a pre-shared key (PSK), the key management domain (KMD) parameters included in EAPOL announcements and MKA frames can identify PB1 and PB2 as belonging to two different domains.

What should be considered for inclusion in some future revision or amendment to 802.1X?:

- a) A description of this use case.
- b) The allocation of a standard group address for EAPOL to support this use case. This address would not be one of the existing 16 ‘reserved addresses’ but would be filtered by all bridges that actually used it (e.g. by PB1 and PB2 in Figure 14). Filtering by other S-VLAN components in customer networks would be optional.
- c) An augmented description of the use of KMD parameters to prevent undesirable connectivity, and possibly a reference to the use of parameters in EAP exchanges and methods for the same purpose.

A. FAQs

This section is currently a place-holder for answers to questions arising from review of this document.

B. Additional technical detail

B.1 Transparent SecY operation

802.1AE’s specification of received frame processing and management controls currently do not allow a SecY to be configured as totally transparent: i.e. to operate as if it were not present in the interface stack.

<<Why this was a deliberate choice, wisdom or otherwise of making this choice flexible, limits of plug-and-play vs limits of safety in configuration.>>

⁴¹The sufficient and necessary condition is as follows. If A, B, and C are interconnected by the service, then if A can communicate with B and C using the service, then B and C can also communicate using the service (with only brief inconsistencies as the service becomes available/unavailable). The condition is trivially met by point-point connectivity. The condition is violated if B and C are communicating through some other path.