A Proposal for Co-existence of MACsec and Priority/Pre-emption Features

Brian Weis
Cisco Systems
10/30/14

1 Introduction

IEEE 802.1AE [AE06], known as MACsec, provides for the protection of Ethernet frames. In the process it also creates an ordered set of frames, where a sequence number documents the order. Receivers of the frames use this sequence number as a way to enforce replay protection. When the ordering of those frames is adjusted by any network service (including priority queuing and preemption features), the MACsec receiver will receive the frames out of order. In some MACsec use cases this can cause the MACsec replay protection method to discard out of order frames.

This paper proposes a general method of avoiding the discarded frames but fully retaining the value and semantics of the replay protection method.

2 MACsec replay protection

A good description of the MACsec replay protection method is available in Clause 4 of [S14]. Briefly, a MACsec sender sends a stream of packets including a monotonically increasing sequence number called a Packet Number (PN) in the MACsec SecTAG.

Each receiver maintains a "window" of acceptable PN values, which is simply a range of PN values where the highest value is the most recently received and validated PN. Any frame that arrives with a PN with a value greater than the low end of the window is accepted after its integrity check verification successfully completes. Note that the MACsec receiver does not maintain any record of previously received PN values, so it is possible that duplicate frames that successfully complete integrity check verification will be accepted.

Receivers can control the likelihood of accepting duplicate frames by defining the smallest necessary window. Indeed, Clause 10.7.8 of [AE06] indicates a default value of zero, which would absolutely mitigate replays. This is often a preferred setting for MACsec users, who tend to be conscious of threats to their network. Indeed, their security policy may prefer the occasional discarded frame to the occasional replay that might occur if the window was open (even with a value of one, which would allow for frames with the most recently received PN to be replayed.

Replay protection can be disabled (through the use of the replayProtect control), but this document assumes that such a setting is not acceptable in any meaningful network use case.

3 Network services that can cause MACsec-protected frames to be re-ordered

Although a window size of zero is preferred, in practice there are ramifications to its use. As stated in Clause 6.10 of [AE06]: "The reception of misordered frames can cause MACsec implementations to discard additional frames depending upon the configuration of replay protection parameters." Field use of MACsec has shown that there are several cases where this misordering is likely to occur, and one additional case is expected to be a problem in the near future. These use cases are described in the following sections.

3.1 Provider Bridge Re-Ordering

MACsec is used to protect frames through Provider Bridge Networks (PBN) (e.g., an Ethernet Virtual Circuit (EVC)), which represents an untrusted network between customer bridges. Field use has shown that some provider bridge implementations unexpectedly prioritize select MAC destination addresses, typically addresses associated with routing protocols. This behavior is entirely out of control of the MACsec receiver, but results in it processing the frames out of order.

Unfortunately, the frequency of discarded frames goes beyond the tolerance of “an occasional discarded frame” which might be tolerable to sites with a security policy preferring a window size of zero. It has been observed that the prioritized frames are dispersed between a set of non-prioritized frames, so opening up the replay window by a small number often seems to mitigate the problem. Still, this mitigation is not the most optimal solution for sites with a security policy of discarding all replays.

3.2 Class-Based Quality of Service

Any time a stream of MACsec frames are encapsulated with a C-VLAN or S-VLAN component and the Priority Code Point (PCP) encoding differs, there is a risk of MACsec frames being delivered to the MACsec receiver out of order. The use of PCP values is common across PBNs, where some frames are marked “high priority” and queued in bridges independently from frames marked “low priority”.

For example, the current draft of P802.1AEcg [AEcg14] describes a use case where prioritization is used with MACsec. Figure 15-4 in that document graphically shows an Ethernet Data Encryption device (“EDE-M1”) includes an S-PRI tag preceding the SecTAG, which allows the PBN to deliver frames based on priority.

Unlike the case of provider bridge re-ordering, the ordering and frequency of different classes of service assigned to MACsec frames is less likely to be predictable. That is, a “low priority” MACsec frame may be followed by a stream of “high priority” MACsec frames, but the number of “high priority” frames is unpredictable. If the PBN prioritizes the packets as indicated in the PCP, the “low priority” frame may be delivered to the MACsec receiver after the “high priority” frames. The MACsec receiver has no option in this case but to configure a replay window with a relatively high value, guessing the number of high priority frames that may be delivered out of order. This is not a satisfying security posture for a site with a strict security policy.

3.3 802.3 Preemption

P802.3br [3br14] describes a MAC Merge sublayer, which includes “a pair of full-duplex MACs and a single PHY”. One of the MACs is described as a “preemptable MAC”, and the other an “express MAC”. The MAC Merge allows frames transmitted on the “express MAC” to interrupt frames transmitted by the “preemptable MAC”. P802.1Qbu [Qbu14] supports the preemption method.

As shown in Clause 5 of [S14], this can result in MACsec protected frames to be processed on the MACsec receiver out of order, insomuch that it will process the expressed MACsec protected frames before it processes the merged preempted frame.

Similar to the Class-Based Quality of Service use case, the number of expressed (“high priority”) frames is indeterminate, and whether or not the PN of the preempted (“low priority”) frame is within the replay window depends on the size of the replay window. As in that use case, the only effective mitigation to discarded frames is to increase the size of the replay window to a large number, which (as pointed out earlier) is not a satisfying security posture for a site with a strict security policy.

3.3.1 Previous Proposal for Mitigation

[S14] proposes an alternative method for mitigating the discarded frame problem, which subtly changes the replay protection method when preemption is in use.

Before describing this solution, the replay protection method needs to be described in more detail. Section 2 of this document describes the MACsec replay protection method, which is enforced after integrity check validation (see the Secure Frame verification process (Clause 10.6 of [AE06]). However, that Secure Frame verification process also includes a preliminary replay check, which is performed prior to cryptographic validation. The intent of this preliminary replay check is to efficiently discard packets that would appear to obviously fail cryptographic validation. As a conventional replay protection method, this check is an optimization – replay protection isn't enforced, and replay protection state is not altered until cryptographic validation has established that the frame is authentic.

[S14] observes that, unlike other priority schemes, the head of the preempted packet is likely to be received by the MACsec receiver ahead of the expressed frames. This implies that the PN in the preempted packet can be evaluated by the preliminary replay check in advance of the express frames. The suggested solution is to remove the corresponding post-cryptographic validation check, depending on only the preliminary replay check to detect replays. Now the preliminary replay protection check (typically considered an optimization) is actually acting as the replay protection enforcement mechanism (although the replay state used for enforcement it not updated until after cryptographic validation).

It is understood that this approach does weaken the replay protection method, and the risk of the method deserves some consideration to judge whether the simplicity of the method warrants a slight less risk. Under normal conditions the method is likely to achieve the same result as the current replay protection method. But replay protection is most valuable when there exists an active attacker in between the MACsec sender and receiver attempting to get replays through the MACsec receiver. It must be assumed that such an attacker has control of all frames delivered by the MACsec sender and can deliver them in any order, including replays of those frames. When the attacker is dealing with express frames it is likely to be able to achieve this even before the MKA delay protection can react.

However, even if the risk is seen as marginal a reasonable question might be to ask whether a special solution is valuable for this situation, when there are other similar use cases for which a solution is also needed.

4 Proposed Solution

It can be observed that in each of the use cases in Section 3 that there are unique classes of MACsec frames. Most frequently this results in exactly two classes: "high priority" and "low priority". (The use of two classes is not universal: three or more PCP markings can be used in the stream of frames marked with a VLAN component. However, to simplify the proposal only two classes are proposed.)

A solution then would be to give each level of priority its own set of PN values so that the ordering of the frames between the classes is irrelevant. The method requiring the least change to MACsec would be to assign each priority its own MACsec SA (including a unique SCI). There is some symmetry here with the P802.3br model, which describes the Merge Service as "pair of MACs".

This solution would have an obvious cost in the number of MACsec SAs that would be consumed on a network port, where MACsec SAs are a limited resource. However, it can be observed in the marketplace that each generation of MACsec supported network ports tends to support higher numbers of MACsec SAs and over time the extra consumption should not be a constraint to

MACsec usage. Additionally, network ports supporting both MACsec and P802.3br would be motivated to support a sufficient number of MACsec SAs to support preemption.

This document uses the term “*priority SA group*” to describe the set of SAs that are linked to a particular MACsec session.

4.1 Impact on IEEE 802.1AE

If each priority level is given an independent MACsec SA, no change is anticipated to IEEE 802.1AE. However, this assumes that the MAC merge and MACsec services are handled in the correct order, where MACsec is the last service before transmission, and the re-assembly of the preempted frame is first service on receipt. Each MACsec SA in the *priority SA group* would have a unique SCI, accomplished by giving each SA a unique Port Number. The System Identifier would need to be the same value.

Implementations would need a more extensive policy for transmit SA assignment, but this policy is not specified in MACsec (see Clause 10.5.1 of [AE06]).

4.2 Impact on IEEE 802.1X MACsec Key Agreement

MACsec Key Agreement (MKA) would need some additional functionality to support *priority SA groups*. So as to not seriously encumber the MKA state machine, all SAs within the *priority SA group* should share the same fate, in the sense that they are created together and deleted together. Furthermore they should use the same AN, are scheduling to begin transmission at the same time, use the same cipher suite, etc. Updates to the MKPDU should be minimal as well.

The following is a preliminary analysis of the changes needed.

A. The *priority SA group* capability and declaration of use by the key server needs to be included in an MKPDU. Possibly this requires the definition of a new *priority SA group* MKA Parameter Set and/or an increment in MKA version number.
B. Creation of two SAs requires two SAKs. When a KS declares that a *priority SA group* is to be used, the distributed SAK should be used to derive the *priority SA group* SAKs. This would update Clause 9.8.1 of [X10].
C. The use of different Port Numbers for priority-based SAs make them indistinguishable from Virtual Ports. However, the existing Virtual Ports paradigm may not fit this use of Port Numbers. This requires further analysis.
D. If Delay Protection is supported, additional Lowest Acceptable PN values (including the most significant 32 bits in the case of XPN cipher suites) would need to be carried in a *priority SA group* MKA parameter set.
E. Declaring a policy that includes the *priority SA group* in Network Announcements may be useful.

5 References

[3br14] P802.3brTM/D0.1, Draft Standard for Ethernet Amendment: Specification and Management Parameters for Interspersing Express Traffic

[AE06] Std. 802.1AE-2006, IEEE Standard for Local and metropolitan area network – Media Access Control (MAC) Security.

[AEcg14] P802.1AEcg/D0.3. Draft Standard for Local and metropolitan area networks Media Access Control (MAC) Security Amendment 3: Ethernet Data Encryption devices.

[S14] Preemption and MACsec replay protection, Mick Seaman. August 2014. Available at http://www.ieee802.org/1/files/public/docs2014/ae-seaman-preemption-0814-v02.pdf.

[Qbu14] P802.1Qbu/D1-1, Draft Standard for Local and Metropolitan Area Networks—
Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks—Amendment: Frame Preemption.

[X10] Std. 802.1X-2010, IEEE Standard for Local and metropolitan area networks – Port-Based Network Access Control.