MACsec in the car

Mick Seaman

This note documents, and expands upon, discussion in the recent 802.1 Security Task Group interim meeting in Budapest. Errors and omissions are my own. Much of what is said here is obvious, but it does need to be written down at least once.

---

1. Preamble

The potential use of MACsec in the car[1] is of considerable technical interest. At the same time few of us can claim any standing in the automotive industry, and there is a general feeling (from other unrelated activities) that 802 may not be the best place for a strong industry vertical focus[2]. At the same time at least some of us feel uneasy at the continuing durability of 'first we'll build it, then we'll secure[3] it' networking approaches, and the idea that 'security' in some sort of absolute, simple label that can be affixed to a design[4]. So, even though our discussions are somewhat speculative, they may help others with knowledge of the particular problems to be solved[5]. First, they illustrate how 802.1 technology can, might, or cannot be used[6]. Second, they illustrate the issues that have attracted our particular attention, which might help in refining a threat analysis. Having a clear and comprehensive threat analysis is of paramount importance if we are not to end up with 'solutions' that are easily circumvented by an intelligent attacker.

2. Network model

The car (for our purposes) is assumed to comprise a number of Ethernet attached devices connected in a simple topology, probably one of the following:

a) A simple ring (with each device connected to two neighbours).

b) A small number of rings, interconnected by one (or just possibly two, for redundancy) device(s), with the rings laid out geographically (covering quadrants of the vehicle) and with (probably) an increased level of functionality associated with the interconnecting devices.

c) A small number of rings, one dedicated to especially critical functions (braking, steering, motor control, suspension control, for example) and others laid out geographically.

The driving force[7] behind the chosen network topology is assumed to be the cost[8] of the wiring harness, which is assumed to be of low cost fiber optic. It is seriously hoped that the link bandwidth's will be 1 Gb/s or more[9], given the anticipated devices and their bandwidth, delay, and jitter requirements.

The attached devices are likely to include some 'low level' common logic, at least for ring transmission and reception, but otherwise encompass a wide variety of functions—from sensors and actuators, switches, motors (for windows, mirrors, trunk openers, shades, car seat adjustments, as well as more basic automotive functions), to hifi (sources and speakers), and cameras (one or more per corner, possibly more elsewhere). In all it would seem wise to provide for scaling to at least a hundred devices.

---

[1] I resist the temptation to be cute and talk of 'automotive applications' as the latter could include car to car, car to signals, car manufacturing and possibly a number of other things which we haven't been discussing explicitly—at least so far.

[2] This note is not the place to discuss why, what has worked in terms of collaborative efforts in other areas, what might work in the future, and the differences in approaches between various industry forums and standards bodies like IEEE 802.

[3] The layout of some of these may challenge our topology assumptions (there can be a large number of switches and motors in a car door, for example, particularly if the seat controls are sensibly located there, as they are in one of my cars, rather than on the side of the seat where a seated occupant can't see them, as in the other) though several may be aggregated into a single ring attached device.

[4] This can so easily degenerate into 'security theater', a demonstrative performance designed mainly to avoid responsibility.

[5] Putting it differently I'm sticking my neck out here in the hope I'll learn more than how it feels to have your head shot off.

[6] However hard we try it's always going to be possible for someone to jump to a conclusion that the standards ought to be used in a way that we might consider suboptimal.

[7] Excuse the pun.

[8] Installed.

[9] The time taken to produce and agree standards has to be taken into account when considering target topologies, capabilities, and devices. What we have today' is hardly relevant in standards terms unless that is what we are going to have in 3 years time and longer.

3. Network use

The attached devices can use the network in a number of different ways:

a) Replicated transmission with duplicate elimination, for reliability to support safety critical functions.

b) Reserved TSN[10] time critical streams for video and audio.

c) Simple best effort transmission (there may be nothing in this category).

Elements of (a) and (b) can be mixed, particularly as the functions provided by cameras (using video streams) fall increasingly into the safety category (lane departure warnings, pedestrian detection, collision avoidance, etc.).

In a single ring topology, redundant transmission can be a simple matter (for unicast) of each device transmitting in each ring direction, with all the devices on the ring forwarding received packets other than those destined to their own address(es). Duplicates can be eliminated using the P802.1CB header, or by some other frame field not explicitly known by the network. In more complex topologies, where safety critical traffic traverses multiple linked rings, duplication and duplicate elimination may be needed en route to the destination and require an explicit network visible sequence field in each frame.

Particular care has then to be taken when mixing (a) and (b)[11] (above) on the same media. There is effectively no pure (a) approach: explicit reservations are needed for each of the replicated streams, so that the high reliability transmissions do not impact the time critical streams. Buffering is also required at elimination nodes to accommodate the spread of arrival times, so that frames can be forwarded on a single onward reservation. The complexity involved may be better handled by simply duplicating the high reliability streams, allocating each of the duplicates to one of two maximally redundant paths[12], and eliminating duplicates at the destination(s).

4. Device model

A simple, but explicit, general device model should help guard against the sort of mistakes that might arise from focusing on only part of the problem. Apart from the obvious connectors, memory, etc., a device is assumed to comprise[13]:

a) MACsec protection/validation functionality for frames transmitted/received on its physically connected ring ports, and possibly on frames that might be transmitted on either or both ports.

b) Secure storage for cryptographic keys for MACsec and other security functions (see later), credentials, and secure means of operating on those keys.

c) Filters that can be applied to frames, both before and after transmission[14].

d) Some securely stored, possibly unmodifiable code, that is used (unconditionally) to verify any other code running on the device.

e) General purpose code, possibly segmented by function and ability to use other device features.

f) External (non-ring) connectivity.

However obvious it seems, it has to be said that if a device has external connectivity that an attacker can use to compromise the device, injecting new code that can be run without invalidating existing credentials and cryptographic keys, and that code can transmit and receive arbitrary protected frames, then there is nothing MACsec use can do to protect the system. Further, in the constrained environment of the car, the times at which we can rely on external connectivity to validate credentials are limited, so any change to the device (such as modifying running software) are probably best viewed as invalidating any keys and credentials that can be used directly.

If an attacker can only compromise part of the device functionality described above, then we can use MACsec to our advantage. We consider both partly compromised and entirely foreign, rogue, devices that have been surreptitiously added to the car network[15].

[10] TSN—Time Sensitive Networking.

[11] See P802.1CB/D2.4 E.9 "FRER and reserved bandwidth".

[12] The maximally redundant trees algorithm used in IEEE Std 802.1Qca provides a general purpose answer, though for two rings redundantly connected (at adjacent nodes) an almost trivial solution is as follows. Frames for clockwise streams on one ring are forwarded, by the first interconnecting node they encounter, clockwise on the other ring B; while anticlockwise frames are similarly forwarded anticlockwise. The interconnecting nodes have to avoid forwarding frames back onto the original ring, this can be done in a number of ways.

[13] Someone somewhere has a much better structured description of this sort.

[14] This might be general ACLs but I've avoided that acronym here as they may be much more specific.

[15] The most prominent reports of car hacking concern attacks made over the Internet, and are unlikely to be addressed by MACsec use without significant other measures being taken on the device facilitating the attack, but some attacks that involve physical access to the car and addition of components/devices have also been reported.

5. Protecting communication

In broad terms our goals are to protect communications between trusted (authenticated and authorized) devices, preventing untrusted communication from:

- being mistaken for the genuine article, and
- preventing trusted communication.

There are various ways that we might think of using MACsec and our devices to meet these goals. Some discussion of each of these follows.

MACsec was designed to protect frames on a hop-by-hop[16] basis, even though (in some cases) what a 'hop' is might be a little complicated. In the car network that has the advantage (as was intended in networks in general) of localizing attacks to the hop where they are injected into the network. A rogue device inserted into one of the car's rings[17] would only be able to disrupt traffic on the directly attached sections of the ring—either by simply failing to forward traffic or by injecting bogus traffic that would consume bandwidth to the same effect before being discarded by its nearest neighbours. If only one such device was added to the network, a replicated stream traversing the ring in the other direction would maintain connectivity for all the other devices. In any event the network disruption caused could equally well be achieved using wire cutters or timed explosive charges, there is nothing particular to the use of Ethernet as a replacement for a traditional wiring harness. The MACsec frame protection capabilities for such a device could be designed so that frames added to the ring (though not those forwarded from one ring port to another) always used the fixed MAC source address of the device, and might constrain the parameters of protected frames in other ways—such as constraining the VLAN Identifier used. The important thing, of course, is not that these restrictions be burnt into a component at manufacturing time, but that any credential or cryptographic key representing the final device as authenticated and authorized be invalidated if these parameters are changed[18]. These restrictions can then be made available to the intended peers of the device, so they can validate (using MACsec) that the frame has not been modified since originally transmitted, and then (by filtering) that it was transmitted by an appropriate peer.

A significant advantage of this hop-by-hop MACsec with frame field (source address, possibly VLAN) approach is that all the protocols that are being used to configure the network are protected, including those responsible for establishing timing and bandwidth reservation at each node.

One alternative approach, though not one I would wish to advocate, would treat the entire network as virtual shared media. In its simplest form all the devices would be provisioned with the same CAK, with any device being able to communicate with any other. Any device would also be able to spoof any other, though this deficiency could be remedied by only provisioning (with the CAK) those devices with a trusted source MAC address enforcement component (with automatic CAK invalidation on any change). In this case, simple source address checking prior to transmission would be preferable to more elaborate frame field checks, as that can be done after MACsec protection even if the frame is encrypted[19]. The apparent advantage of the virtual shared media approach is that each network device can forward traffic without having to validate it, then protect it again. However MACsec implementations are quite capable of supporting full line rate fixed delay processing.

Such a virtual shared media network could also support multiple independent CAs, each with a different CAK. A device that needed to participate in more than one of these would use a different virtual port (with its own SCI). If there is no intent to support insecure connectivity in a given CA and the topology comprised a single ring, there would no need for any additional tagging before the SecTAG. A visible VLAN TAG would not be needed for filtering, as the frame should reach all devices on the ring[20]. Of course each member of a CA would still have to check that frames received for that CA contained appropriate field values to guard against compromise of other CA members, and source address filtering on transmission might still be desirable. Apart from the failure to localize attacks and protect other configuration protocols, a problem with the virtual shared media

[16]See "MACsec hops" http://www.ieee802.org/1/files/public/docs2013/ae-seaman-macsec-hops-0626-v03.pdf for rationale and an extended discussion.

[17]Or a compromised device designed so that the compromise would necessarily invalidate existing credentials and cryptographic keys.

[18]Some permanently trusted or unmodifiable component needs a way of reading the components configuration or posture.

[19]All standardized MACsec Cipher Suites allow for integrity protection only, and that would seem to be desirable in this application, but security has so long been associated with obscurity that assuming encryption (integrity protection plus confidentiality protection) seems the safest choice.

[20]Forwarding of unicast frames can be suppressed easily at the destination, so a given frame does not have to traverse all the links in the ring in both directions. The MACsec PN could be used to suppress multicast duplicate traversals, though if done without validation would allow damage frames to suppress counter rotating frames, and would require part of the PN recovery logic for XPN Cipher Suites and knowledge of SAIs for all frames if not just performed at intended receivers (this can be extracted from MKPDUs without participating in the CA).

model would be the so far unprecedented numbers of CA members[21], and the total number of keys and credentials required.

6. Key management[22]

The actual authentication and authorization process needs to follow the work flow and requirements normally associated with adding or replacing car components/devices.

The goals of this activity include:

a) Validating the component, is it what it appears to be (as opposed to a Trojan device or inferior substitute).

b) Tracking the history of the component—was this previously installed in another car and if so were there any problems associated with its installation, use, or de-installation.

c) Provisioning the component with the necessary CAKs, and with any other details (such as MAC Addresses) that it needs to filter/verify frames from other components.

d) Provisioning/re-provisioning other car components with any other details they need to know about the added component.

The initial stages of this process might use a component Secure Device Identifier (as specified by 802.1AR) and a process very similar to the enrollment process, checking that the component did come from the presumed manufacturer and is as expected. This could then result in a Local DevID that could be used by the device in subsequent authentication exchanges.

A new device could be configured, prior to installation, with a single CAK common to all the devices in the car, though when used with the devices neighbours it would result in independent SAKs for each of the links. It’s not clear whether the possibility of compromising a device and exfiltrating a CAK for use in a subsequent attack is an issue, or whether any compromise would be overwhelming likely to be used immediately.

The usual EAP model, of a Supplicant communicating with an immediately neighboring Authenticator that makes use of a centralized Authentication Server to derive an EAP MSK, and from that a MACsec CAK, could be applied. If each device is capable of acting as Supplicant or Authenticator that process could work its way around a new ring successfully. Offloading as much of the Authenticator functionality to a provisioning device temporarily attached to the car would seem desirable. This process would result in CAKs specific to each link.

[21]Especially if multiple transmit SCIs are used to support differential stream/priority forwarding at ring nodes while retaining strict replay protection.
[22]The easy stuff last :-). The following represents just a fraction of the possible approaches to authentication, authorization, and key management.