Privacy considerations in bridged networks

Mick Seaman

This note is a result of a dry run of the P802E/D1.1 (Recommended Practice for Privacy Considerations for IEEE 802 Technologies) suggestion of a self-assessment by the developers of each 802 standard, with the results recorded in a Privacy Considerations annex. I chose IEEE Std 802.1AS as the target because the service it provides is rather different from that offered by 802.11 Access Points to the users of mobile personal devices.

I found the scope of what I was trying to do naturally expanding—it is not possible to describe IEEE Std 802.1AS privacy exposures without some description of how it operates within a bridged network, which brings in a description of that network's operation. In turn that leads to consideration of what 802.1AS is being used for, and to the privacy exposures inherent in recognizing flows through a bridge network, so timed gates can be used to shaped traffic in a network. A preliminary conclusion is that 'the standard' so far as P802E is probably the set of 802.1 standards. Of course any adversary trying to violate personal privacy is not bound to confine attacks to the scope of any particular standard or set of standards, but we have to draw boundaries somewhere to make work practicable.

I have cast the result in the form of an informative annex, that might conceivably be attached to a single 802.1 standard, most likely 802.1Q. *However I should stress that this note is not a proposal for a PAR, and the result of any PAR might be a very different approach.* For the present the goal is to inform the development of P802E/D1.1 and show one way in which its recommendations might be result in a feasible/tractable amount of work. An 802.1Q informative annex would not prevent further detailed work within specific standards, and any change to a standard's mandatory or optional requirements and recommendations should be in that standard.

Notes:

1) The proposed annex draws on P802E discussions in the 802.1 Security Task Group and with the contributors to that work. I have attempted to be consistent with that work, but have not referenced it in annex text since a sponsor ballot comment on 802.1AR-2018 resulted in removing an acknowledgment to the P802E work in progress. The annex should also be intelligible without reference P802E (see next point).
2) Each of the clauses in the annex that discuss particular standards include a brief summary of what that standard is about, with a focus on what the privacy oriented reader most likely needs to know or be interested in. I believe it is essential to include such material. The problems we have to cope with are generally broader than the expertise of any of the participants and a goal of having a critical mass of privacy and security experts read and understand all the 802.1 standards is no more realistic than expecting all the 802.1 experts to review all the relevant security standards. Even if that were possible we would still be faced with the difficulty of keeping the participants on the same page, or more accurately on the same line item, if that were not explicitly written in such an annex. Without a summary of what we are working on that passes through the approvals process it would also not be possible to assess the coverage of the result.
3) I have avoided making any suggestion that any standard should be changed. Such changes might result from a privacy considerations study but would have to be carried out by an amendment of the base standard (if desired). In reviewing the referenced standards I have found a number of things that I believe ought to be changed, but they have nothing to do with privacy.
4) I have avoided some privacy terms that I don't think are particularly relevant in our context.
5) No one else has had the opportunity to review this first draft, so it should not be taken as representing any group opinion.

Annex Y

(informative)

Privacy considerations in bridged networks

This informative annex describes privacy considerations related to the use, design, and deployment of bridged networks based on IEEE Std 802.1Q and related standards (IEEE Std 802.1X, IEEE Std 802.1AB, IEEE Std 802.1AE, IEEE Std 802.1AR, IEEE Std 802.1AS, IEEE Std 802.1AX, IEEE Std 802.1BA, IEEE Std 802.1BR, IEEE Std 802.1CB, and IEEE Std 802.1CM).

The unintentional or unauthorized disclosure of personal information arises from a combination of the following factors:

a) The use of personal devices that are attached to, or form part of, the network (Y.1)
b) The type of information that adversaries might wish to acquire (Y.2)
c) The efficient operation and management of the network (Y.3)
d) The use of security protocols for authentication, authorization, integrity, and confidentiality (Y.5)
e) The frame fields that contain information useful to an adversary, the sophistication of, and the network access afforded to, that adversary (Y.5)

Privacy considerations particular to a given referenced standard are discussed in Y.6.

This annex is informative. It does not modify the mandatory or optional provisions or the recommendations contained in any referenced standard.

Y.1 Personal devices

Privacy, in the context of bridged networks, relates to the use of personal devices i.e. devices used by one person or a small group of people. Information that identifies a personal device or is associated with that device identification can thus yield information about the location and activities of a person.

Shared service devices, in contrast, support applications for a large enough group of people such that correlation between any given person and the observable behavior of the device is weak. Other devices, e.g. sensors in industrial networks, have no direct correlation with a person.

In general IEEE 802.1 standards are applicable to both personal devices and shared service and other devices. However some protocol roles, e.g. Grandmaster in IEEE Std 802.1AS Timing and Synchronization for Time-Sensitive Applications, are unlikely to be associated with personal devices in other than the smallest bridged networks, and are even more unlikely to be associated with mobile personal devices.

Y.2 Goals of adversaries

An adversary can be interested in the following personal information:

a) Who is using a personal device (identification)
b) Where are they (location, and location tracking)
c) What are they doing (activity, application use)
d) With whom are they associated (communicating, shared interest).

The information on all, or indeed on any of these, need not be complete to be useful to an adversary. The adversary can, for example, be interested in facts such as:

- an identified person appears to be engaging in the same, unknown, activity as a group of unknown persons at another identified location
- there appears to be no one at an identified location.

The information obtained need not be particularly accurate to be useful to an adversary. It is sufficient that the cost of acquiring the information is less than the benefit expected from its use, allowing for the probability that it is incorrect and any costs associated with the use of incorrect information. Use of incorrect information can negatively affect a targeted person.

Y.3 Network operation

Bridged networks support frame based transmission, with variable length frames and without requiring attached stations (except for certain time-sensitive network applications) to adhere rigidly to a clocked transmission schedule. Stations are not obliged to transmit when there is nothing to transmit and frames are not all padded to the same length, so the use of network resources benefits from statistical multiplexing. At the same time some network applications have requirements for timely delivery that cannot be met simply by relying on that multiplexing and increasing transmission speeds but require signaling, to bridges in the network, of the differential service requirements of individual frames. Time-sensitive network applications with more stringent delivery requirements require bandwidth allocation, supported by end station protocols or management configuration, and sufficient information in individual frames for bridges to associate each frame with an allocation (and thus with an individual end station and a particular type of end station application). Bridged networks provide more bandwidth than is available from each of their constituent individual LANs by restricting data frames to paths to their intended destinations. One of a number of alternate paths to a given destination end station or set of end stations can be used to further increase the available bandwidth, but common network application frame ordering requirements constrain the distribution of frames amongst such paths to those that bridges can distinguish as belonging to separate application flows.

Bridges in the network can distinguish between application flows using each frame’s destination MAC address (DA), source MAC address (SA), the VLAN identifier (VID) and priority code point (PCP encoded in the VLAN tag (if present), the EtherType (or LSAP) identifying the higher layer protocol conveyed by the frame, and the initial fields of that protocol. Protocols that operate over the bridged network and are used by personal devices to support network applications and to communicate with application servers and other devices (as opposed to reserving network resources for that communication) typically use the Internet Protocol (IP). It is rare for two personal devices to communicate without transmitting frames via one or more intervening routers. Any given IP subnet is often supported by a single VLAN, so bridges that support parallel paths for routed application flows from individual end stations typically use the source and destination IP addresses, the conveyed protocol type (IP, UDP, or SCTP), and source and destination ports for that protocol (see 9.1.5 of IEEE Std 802.1CB-2017).

Some IEEE Std 802.1 protocols, e.g. IEEE Std 802.1Q Stream Reservation Protocol (SRP), transmit frames with group destination MAC addresses. These addresses identify the type of the intended recipient protocol entity and allow bridges to use address filtering to restrict those frames to an appropriate scope, reaching only the nearest bridge, for example. Some group addresses support a particular type of application, and thus associates the source MAC address (and the station that is using it) with that application.

The deployment and operational costs of bridged networks have been considerably reduced by the use of protocols that volunteer device information (e.g. IEEE Std 802.1AB) even when those protocols are not used to support full ‘plug-and-play’ operation. Management and trouble shooting of faulty devices or apparently incorrect network behavior depends on the recording of device location and gathering statistics on network

use. Stations implementing protocols whose operation depends on the presence of a reachable collaborating peer (e.g. IEEE Std 802.1AS gPTP time synchronization) typically advertise their capabilities, either by using IEEE Std 802.1AB or by sending their own messages.

Y.4 Network security and privacy

As described above (Y.3) efficient use of network resources, particularly for data frames that require other than best effort delivery, depends on the bridges in the network being able to identify end stations and (for some applications) service characteristics (priority, bandwidth and delay) required by their network applications. Where physical access and attachment to the whole or part of a bridged network is restricted to authorized personnel, confidentiality protection can be limited to that provided by higher layer protocols, notably TLS or IPsec. This leaves all the identifying information specified in IEEE 802.1 standards exposed to an adversary that does gain access to the network media.

The end stations and bridges in bridged networks are typically connected by IEEE Std 802.3 Ethernet links. MACsec (IEEE Std 802.1AE) can be used to provide both confidentiality and integrity protection hop by hop, leaving (in the most common configuration) just the MAC source and destination addresses, frame length, and frame transmission timing visible to an adversary with access to the network media. MACsec adds fields to each frame, but an adversary can recover the original frame length. MACsec operation can affect frame timing, but implementations suitable for use in time-sensitive networks impose a small fixed delay so as not to degrade the operation of IEEE Std 802.1AS time synchronization or IEEE Std 802.1Q timing gates supporting traffic shaping and bandwidth allocation. The frame to frame timing relationships that an adversary might observe remain unaltered. Where MACsec is used with Ethernet frame preemption and in-order delivery of preemptable and (separately) of preempting frames is enforced, an observer can distinguish these two classes of frames. Privacy considerations particular to IEEE Std 802.1X (Port-Based Network Access Control) support of MACsec are described below (Y.6.2).

Unlike IEEE Std 802.11 operation in which a mobile end station participates in observable protocol to discover and select a suitable service it is rare for an end station to be connected to an Ethernet link that does not provide the expected service. Authentication, authorization, and confidentiality protection of subsequent data frames, if required, typically occurs before additional end station information is disclosed. For exceptions see Y.6.2, Y.6.3.

Y.5 Privacy exposures

A personal device can be identified explicitly by a single frame field, notably by using a universal MAC address as the source address of transmitted frames.

A station can use a locally assigned MAC address, chosen randomly from the entire local address space or from a subset large enough to yield a sufficiently low probability of address assignment collision, or explicitly assigned by a higher layer protocol. However once a local MAC address has been assigned to a station and is being used to support higher layer protocols (such as IP), to restrict data frames to the path to that station, and to reserve resources in bridges along the path, any further MAC address change can be expected to interrupt or degrade the MAC Service. Moreover the disappearance of one address coupled with rapid appearance of another facilitates correlation of the two addresses and cannot be expected to reduce an adversary’s ability to infer information from the frame fields and other characteristics of persistent flows.

Where an individual frame field does not directly identify a personal device, either persistently as in the case of a universal MAC address or temporarily while the device is continuously active, an adversary can correlate those frame fields and other frame characteristics to identify (to an acceptable probability) the frames and frame flows associated with a single device and even to ascribe a permanent identity to that

device or the particular network applications and activities supported by the device. Such a correlation is called a ‘fingerprint’, and the process of obtaining it ‘fingerprinting’. Fingerprinting does not necessarily require a detailed understanding of the protocols used by a device, but can use general correlation and machine learning techniques to find any persistent pattern in the behavior of a device. Indeed a fingerprint can use device characteristics, such as the persistent scheduling of a transmission by one activity immediately after transmission for another activity, that do not appear in protocol specifications. In the absence of information that all the personal devices of a given type in widespread use consistently use the same network applications in the same way (and consequently exhibit indistinguishable network behavior) it has to be assumed that devices and activities can be distinguished by a sufficiently interested adversary.

The pattern of frame sizes transmitted and received by a personal device can fingerprint application activity and reveal details of that activity. The Ethernet MAC does, however, impose a minimum frame size, and higher layer protocols include fields that allow them to determine the applicable data length. To support Ethernet bridging of frames to and from media without the minimum size requirement, MACsec can encode the short length of those frames, but short frames that have been padded prior to being protected with MACsec will appear to be of uniform length, thus depriving an adversary of the opportunity of fingerprinting application types using the small frame sizes that can be used in initial capability advertisement.

NOTE 1—Frame size patterns have been used to identify banking applications for specific financial institutions, approximate account balances, and whether money is being added to or removed from the account.

Static personal devices, e.g. desktop computers and home routers, typically connect to bridged network using an individual wired IEEE Std 802.3 Ethernet connection. An adversary that can gain access to that wired connection has usually already identified (knowledge of home occupancy, etc.) the person or people associated with such a device and there is no question of tracking device movement. However the pattern of device activity (e.g. turning on security cameras when there is nobody at home) can reveal important personal location information.

NOTE 2—This annex does not detail privacy exposures resulting from media access control method operation, but notes that they can exists. For example, PoE (Power over Ethernet) use can reveal the identity and software version of some consumer electronics devices even when the adversary is restricted to observing the neighboring electromagnetic field.

Bridged networks are typically intraconnected with Ethernet links. Where these are wholly on private premises, access by an adversary can be prevented or at least made so difficult and expensive as to limit the targets to previously identified persons. Where personal device traffic to and from those private premises passes through an IP router, the privacy considerations are those applicable to the use of IP.

NOTE 3—At the time of preparation of this annex, discussion of the extension of TSN capabilities beyond the scope of bridged networks to the use of IP under the heading of ‘DetNet’ (deterministic networks) was still at an early stage. The privacy impacts of explicit flow identification and resource allocation described in this annex can be expected to apply.

IEEE Std 802.11, non-standard wireless connectivity, and in-home electrical power wiring can also be used to connect devices to personal bridged networks and to connect bridges within those networks. Where IEEE Std 802.11 is used to connect to an access point (AP) operating as an IP router, the security considerations applicable to 802.11 and IP apply. Where non-standard wireless connectivity and electrical power wiring are used, an adversary located sufficiently close as to be able to intercept the wireless signal or access power wiring outside possibly secured premises can be assumed to have access to MAC address, frame size, and frame timing information at a minimum with the further possibility of access to all the resource allocation and flow identification information conveyed. Frames with specific group and individual MAC addresses can be filtered by bridges in the network and do not necessarily traverse those links.

An adversary with management access to bridges in the network will have access to resource allocation and flow identification information, but not (at least with standardized objects) the sizes of specific frames and their transmission timing. Such adversaries can include organizations that have a business relationship with the targeted person and are considered trustworthy by that person.

Y.6 Standard specific considerations

This clause summarizes particular ways in which each of the bridged network related standards can, when supporting personal devices, expose information that can be used to fingerprint the device's identity or use of network applications. Unless otherwise stated the general considerations described above (Y.3, Y.4, Y.5) also apply to the use of each standard. The brief summary of each standard's capabilities is intended to provide the context for privacy considerations, and is not a substitute for the text of each referenced standard.

Y.6.1 IEEE Std 802.1Q Bridges and bridged networks

The general considerations described above (Y.3, Y.4, Y.5) all apply to the use of IEEE Std 802.1Q.

Y.6.2 IEEE Std 802.1X Port-Based Network Access Control

IEEE Std 802.1X specifies a general method regulating access to a network, both by systems that are the source and destination of frame carried by the network and by relay systems that are to be connected to multiple other systems in the network and that forward frames between those connections. In both cases each of the system's ports either participates in a mutual authentication exchange with the neighboring system or proves the success of past authentication and authorization to access the network. This clause discusses potential privacy exposures arising from the use of the media-independent capabilities of IEEE Std 802.1X with Ethernet, for privacy considerations related to the use of IEEE 802.11 connections to or within bridged networks see IEEE Std 802.11.

Extensible Authentication Protocol (EAP, IETF RFC 3748) messages are encapsulated in EAP over LANs (EAPOL) PDUs so they can be sent between a Supplicant port (also referred to as a Peer in IETF EAP RFCs), that wishes to gain access to the network, and a Authenticator port, on a system that provides network access. EAP is an authentication framework, not a specific authentication mechanism, and more than 40 specific authentication methods have been defined. An Authenticator is typically supported by an Authentication Server (AS) that executes the particular method or sequence of methods selected. The authentication credentials supported by different methods can differ, as can the degree to which they expose the identity claimed by a Supplicant. EAP messages between the Authenticator and the Authentication Server are typically encapsulated in the RADIUS (IETF RFC 3579) or Diameter (IETF RFC 4072) protocols. IEEE Std 802.1X-2010 mandates the use of mutual authentication methods, and requires support for EAP-TLS (IETF RFC 5216) if integration with the use of IEEE Std 802.1AR is claimed.

Following EAP authentication, RADIUS or Diameter server can provide the Authenticator with attributes that include access controls appropriate to the authorization accorded to the Supplicant and information that supports subsequent reattachment of a device to the network without repetition of the full authentication exchange and authorization process. These attributes can include persistent identifiers, e.g. the EAP-Key-Name (the IEEE 802.1X secure Connectivity Association Key Name, CKN) and Network-Id-Name (2.2 and 2.7 of IETF RFC 7268). Privacy considerations relating to communication between the Authenticator, the Authentication Server, a RADIUS or Diameter Server, and any Online Certificate Status Protocol (OCSP) Server are described in the relevant IETF RFCs.

If data transmission, following successful authentication and authorization, between the Supplicant and Authenticator ports is protected by MACsec, the MACsec Key Agreement protocol (MKA) is used to distribute the succession of Secure Association Keys (SAKs) used to provide confidentiality and integrity protection. MKA uses keys derived from a secure Connectivity Association Key (CAK) and the CKN to integrity protect MKPDUs and to confidentiality and integrity protect (using AES Key Wrap) distributed SAKs. The contents of MKPDUs (other than distributed keys) are not confidentiality protected to support network monitoring and debugging without needing to share the CAK or derived keys. The CAK and CKN

can be derived from an EAP authentication or can be pre-shared by other means, including local device management. The initial octets of each MKPDU contain the CKN, so a peer MACsec capable system knows which (if any) of its key to use to verify that the MKPDU has been transmitted by a previously authenticated system. A device can be configured to attempt, or require, EAP authentication each time it is connected to the network, thus obtaining a fresh CAK and CKN. Shared service infrastructure devices typically need to be capable of restoring connectivity to their neighbours without re-authentication, since neither they or their neighbors are guaranteed to have connectivity to an Authentication Server or other supporting services.

EAPOL frames, and integrity protected MKPDUs which are carried in EAPOL frames, can convey network announcements (Clause 10 of IEEE Std 802.1X-2010). These can be used by personal devices, but are expected to be transmitted by shared service devices.

Y.6.3 IEEE Std 802.1AB Station and Media Access Control Connectivity Discovery

The Link Layer Discovery Protocol (LLDP) allows a station to advertise, to others attached to the same LAN, the station’s management address and major capabilities. The receiving stations allow management access to received LLDP information to support network topology discovery and configuration checking. The point of LLDP would be lost if the advertised attributes were to be temporary or unavailable to intended recipients. Standard attributes include a system name and description. The range of attributes has been extended by other standards and organizations such as equipment suppliers.

LLDP is a one way protocol: it does not contain mechanisms for soliciting or confirming receipt of information. The destination address of each LLDPDU is usually one of the reserved group addresses specified in IEEE Std 802.1Q and filtered by bridges to limit the scope of its propagation through the network. This filtering allows a management application to use the information received by end stations and bridges in the network to build a map of the network topology. The filtering also restricts exposure of any station’s advertised attributes to adversaries that have access to the individual LANs traversed by the LLDPDUs that station transmits, or that have management access to their recipients or to the management application. IEEE Std 802.1AB-2016 mandates support for the Nearest Bridge group address (01-80-C2-00-00-0E, also referred to as the Individual LAN Scope group address). This address is filtered by all bridges.

Where port access is controlled by IEEE Std 802.1X, IEEE Std 802.1AB mandates Controlled Port support for LLDP exchanges, thus providing confidentiality (on the LAN) if MACsec is used. Unprotected transmission using the Uncontrolled Port is permitted.

Y.6.4 IEEE Std 802.1AE MAC Security

The exposure of personal information, including information that can contribute to fingerprinting a device or activity, conveyed in frames that are confidentiality protected by MAC Security (MACsec) can be reduced as described above (Y.4). The potential exposure of personal device information by the supporting IEEE Std 802.1X MACsec Key Agreement protocol (MKA) is discussed in Y.6.2.

MACsec protects communication between neighboring systems, but the scope of that protection depends on what each system considers to be a potential neighbor. By default frames conveyed by the IEEE Std 802.1X Port Access Control Protocol (PACP) that encapsulates the Extensible Authentication Protocol (EAP, IETF RFC 3748) are transmitted to the Nearest non-TPMR Bridge group address (also referred to as the IEEE Std 802.1X PAE address), so any intervening TPMR cannot access confidentiality protected frame fields (see Y.4). However MACsec can also be used to secure a point-to-point connection across a Provider Bridge Network exposing any priority information required by PBN systems to provide the desired class of service, and to secure connectivity where the PBN uses VLAN tag information to select a provider service

instance (15.4 and 15.5 of IEEE Std 802.1AE-2018). Where a Provider Backbone Bridge (PBB) is used, the source MAC address of the originator of the frame is encapsulated and confidentiality protected. A PBB is not, itself, likely to be a personal device.

Y.6.5 IEEE Std 802.1AR Secure Device Identity

IEEE Std 802.1AR specifies Secure Device Identifiers (DevIDs) for use with IEEE Std 802.1X and other industry authentication, provisioning, and authorization protocols. Privacy consideration for use of DevIDs are discussed in 6.5 of IEEE Std 802.1AR-2018.

Y.6.6 IEEE Std 802.1AS Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks

The generalized precision time protocol (gPTP), state machines, and algorithms specified in IEEE Std 802.1AS support time-sensitive applications such as audio, video, and time-sensitive control, by maintaining synchronized time across packet networks, including bridged networks, comprising interconnected time-aware systems. Each time-aware system exchanges messages with its immediate neighbor to measure the link propagation delay experienced by packets forwarded by that neighbor. Time-aware end stations receive time information, either directly or indirectly via one or more time-aware relay systems, from a grandmaster that is the source of time information in a network domain. Each system adjusts the time information received to account for the link propagation delay, and in the case of time-aware relays for the residence time of the information in the relay prior to forwarding. The current grandmaster, and the port used to receive information from that grandmaster, is selected by a best master clock algorithm (BMCA) that constructs a time-synchronization spanning tree throughout the network domain with a spanning tree priority vector that allows each time-aware system to select its best port for receiving (and in the case of a time-aware relay, as the basis for forwarding) timing information.

Each time-aware system port that supports gPTP is identified by a sourcePortIdentity, comprising a clockIdentity and a portNumber. The clockIdentity identifies the clock being used by a specific time-aware bridge or end station for a particular instance of distributed time and is constructed using an NUI-48 or NUI-64 (see IEEE Std 802c): i.e. while it is an identifier and not a protocol address it is constructed in the same way as MAC address, is intended to be unique within a network, and it is possible to tell by examining one of the bits derived from the NUI in the construction (the bit corresponding to the U/L bit when the an NUI is used as a MAC Address) whether the clockIdentity is intended to be locally or globally unique.

The media-independent specification of gPTP is supported by media-dependent procedures. Neighboring time-aware systems connected by full-duplex point-to-point links, such as those specified by IEEE Std 802.3, use gPTP messages to measure the propagation delay and convey timing information. Each message includes the transmitter’s sourcePortIdentity. If the connection is confidentiality protected by MACsec, this message field will only be visible to the communicating systems.

Neighboring IEEE Std 802.11 stations, whether AP capable or not, do not use gPTP messages to measure propagation delay and convey timing. They use the IEEE 802.11 MAC Layer Management Entity (MLME) which generates, timestamps, and consumes measurement frames to provide timing information.

NOTE 1—For privacy considerations related to the IEEE 802.11 MLME see IEEE Std 802.11.

The BMCA spanning tree conveys a trace of each port’s sourcePortIdentity on the best path (for timing distribution) from each potential grandmaster. A personal device attached to the network is thus aware of, and receives a permanent identifier for each system that is part of, that path. The BMCA protocol does not propagate path information in the reverse direction (i.e. towards the grandmaster root of a timing tree): the sourcePortIdentity of a personal device that has a single port attached to the network is only conveyed to its immediate neighbor.

NOTE 2—While use of the redundant grandmasters and the BMCA allows the precision timing service provided by IEEE Std 802.1AS to be resilient in the face of system and link failures, it is highly desirable that the network remain stable and the standard provide priority values for grandmaster selection and timing path selection that discriminate against devices that are not permanently part of the network and powered on. A personal device, and particularly a mobile personal device, is therefore unlikely to find itself in the position of propagating BMCA path trace information including the sourcePortIdentity of one of its ports, even if it has more than one port.

Time-aware stations connected by media for which gPTP is supported by media-independent procedures send Signaling messages (10.4 of IEEE Std 802.1AS-2011) that signal the stations ability to participate in the protocol together with station dependent parameters that control aspects of protocol operation (e.g. message interval request). For stations connected by IEEE Std 802.11 media this capability is provided by the IEEE 802.11 MLME.

Y.6.7 IEEE Std 802.1AX Link Aggregation

Link Aggregation allows parallel point-to-point links to be aggregated to form a Link Aggregation Group (LAG) that is treated as a single link. A bridge or end station port generally distributes frames amongst the links so as to preserve frame ordering within flows (see Y.3 above). The distribution algorithm and parameters for its use can be specified by using Conversation-sensitive Collection and Distribution (CSCD). A further capability Distributed Resilient Network Interface (DRNI), that provides system level redundancy by allowing two cooperating systems to mimic the behavior of a single system terminating a LAG, is unlikely to be used by personal devices.

The addition and removal of links to and from a LAG is facilitated by the operation of the Link Aggregation Control Protocol (LACP) in each of systems they connect. To ensure that the candidate links for a given LAG do connect the same pair of systems, LACP exchanges a System Identifier that is a combination of System Priority and System MAC Address. This System MAC Address needs to be unique amongst any set system capable of aggregating links with each other, but does not have to be globally unique and is not necessarily (except for any conditions imposed to avoid the profligate assignment of unique identifiers) the address used as a source MAC address by transmitted frames originating from the system. Other LACP parameters, because of potential system to system differences, can contribute to system fingerprinting though not in such a clear way. LACPDUs are transmitted to a group address selected to limit their propagation within the network, typically the Nearest non-TPMR Bridge group address (01-80-C2-00-00-03).

When MACsec is used to protect communication between neighboring system, the MAC Security Entity is instantiated in the interface stack associated with each of the individual aggregatable links (see 11.5 of IEEE Std 802.1AE-2018) and thus can confidentiality protect both the conversations carried over those links and operation of LACP.

Y.6.8 IEEE Std 802.1BA Audio Video Bridging (AVB) Systems

IEEE Std 802.1BA specifies the selection of specific features and options from IEEE Std 802.1Q, IEEE Std 802.1AS, and LAN MAC/PHY standards that facilitate manufacture of AVB-capable components. A person not skilled in networking can use those components to build networks that provide working audio and video services. This standard does not introduce additional privacy considerations beyond those inherent in the referenced standards.

Y.6.9 IEEE Std 802.1BR Virtual Bridged Local Area Networks—Bridge Port Extension

IEEE Std 802.1BR specifies a method for increasing the effective geographical extent of the control parameters of a single bridge by supporting multiple instances of the Enhanced Internal Sublayer Service, each associated with a single bridge port, over a single LAN connected to an External Bridge Port Extender that can support one or more ports attached to LANs, each serving a single end station, and zero or more ports connected to further External Bridge Port Extenders. Bridge Port Extenders also support frame replication for multicast. While Bridge Port Extenders extend the effective extent of a single bridge they do require port extender specific configuration to support time-sensitive network flows.

Bridge Port Extenders were standardized to meet data center bridging requirements and are not expected to be personal devices or to provide services directly to personal devices.

Y.6.10 IEEE Std 802.1CB Frame Replication and Elimination for Reliability

Frame Replication and Elimination for Reliability (FRER) increases the probability that any given packet will be delivered by replicating each of an identifiable sequence of packets, transmitting the replicates on disjoint network paths, and eliminating duplicates where those paths meet. The sequence of duplication, duplicate transmission, and elimination, can be repeated between transmission by the original source of the packets and reception by the eventual destination(s). Resources can be reserved on each of the paths that support duplicate transmission so that TSN delivery objectives (timeliness and extremely low loss) can be met even if LANs or relay systems fail (on all but one of the potential paths between the original source and a destination).

FRER requires, at a minimum, the addition of a sequence number to each packet. IEEE Std 802.1CB specifies a redundancy tag (R-TAG) that adds just that sequence number, and also allows use of the High-availability Seamless Redundancy (HSR) sequence tag or the Parallel Redundancy Protocol (PRP) sequence trailer both specified by IEC 62439-3:2016 (7.8, 7.9, and 7.10 of IEEE Std 802.1CB-2016). None of these contribute significantly to an adversary’s ability to assign each packet to a stream or flow as is necessary, using the contents of other frame fields, by bridges and end stations supporting resource allocation and FRER. IEEE Std 802.1CB does not specify positioning of its processing relative to that carried out by the IEEE Std 802.1AE MAC Security Entity in interface stacks, but the usual considerations place the latter closer to the PHY. MACsec confidentiality protection, where used, will apply to the FRER tags and trailer just as it would to the stream and flow identifying frame fields.

While IEEE Std 802.1CB addresses the requirements addressing from industrial networks it can be used to support personal devices.

Y.6.11 IEEE Std 802.1CM Time-Sensitive Networking for Fronthaul

IEEE Std 802.1CM specifies the selection of specific features and options from IEEE Std 802.1Q, IEEE Std 802.1AC, IEEE Std 802.3, IEEE Std 1588, ITU-T G.8275.1, ITU-T G.8261, ITU-T G.8262, and ITU-T G.8264, to enable the transport of time-sensitive fronthaul streams in Ethernet bridged networks. This standard does not introduce additional privacy considerations beyond those inherent in the referenced standards.