

# Securing Layer 2 in Local Area Networks

Hayriye Altunbasak<sup>1</sup>, Sven Krasser<sup>1</sup>, Henry L. Owen<sup>1</sup>,  
Jochen Grimming<sup>2</sup>, Hans-Peter Huth<sup>2</sup>, and Joachim Sokol<sup>2</sup>

<sup>1</sup> Georgia Institute of Technology, Atlanta, GA 30332-0250, USA  
{hayriye, sven, owen}@ece.gatech.edu

<sup>2</sup> Siemens AG, CT IC2 Corporate Technology, 81730 Munich, Germany  
{jochen.grimming, hans-peter.huth, joachim.sokol}@mchp.siemens.de

**Abstract.** Network security problems have been well known and addressed in the application, transport, or network layers. However, the Data Link Layer (Layer 2) security has not been adequately addressed yet. To secure Local or Metropolitan Area Networks, the IEEE 802.1AE Media Access Control (MAC) Security Task Group has proposed the IEEE P802.1AE Standard for Local and Metropolitan Area Networks: MAC Security (MACsec). MACsec introduces a new tag field, Security TAG (SecTAG), in Layer 2 frames. In this paper, we discuss the security concerns in Layer 2 and summarize some of the possible attacks in Layer 2 in Internet Protocol (IP) over Ethernet networks. We also provide an overview of the MACsec. Lastly, we propose to incorporate additional fields into the SecTAG to improve security in local area networks.

## 1 Introduction

Network security has become a concern with the rapid growth of the Internet. There are several ways to provide security in the application, transport, or network layer of a network. However, the network security is only as strong as the weakest section. Since the Data Link Layer security has not been adequately addressed yet, the weakest section may be the Data Link Layer (Layer 2) [1]. Layer 2 enables interoperability and interconnectivity in networks. However, a compromise in Layer 2, which enables internal attacks, may not be detected by the upper layers. In this paper, we focus on the Layer 2 security issues in IP over Ethernet networks.

### 1.1 Security Concerns in Layer 2

Layer 2 in IP over Ethernet networks is prone to several attacks. Three most commonly known Layer 2 sniffing attacks are Address Resolution Protocol (ARP) poisoning, Media Access Control (MAC) flooding, and port stealing.

ARP is a network layer protocol used to map an IP address to a physical machine address recognizable in the local network, such as an Ethernet address. When a host machine wishes to find a physical address for an IP address, it broadcasts an ARP request, which includes the IP address, on to the network.

The host that owns the IP address sends an ARP reply message with its physical address. Each host machine maintains a table, called ARP cache, used to convert MAC addresses to IP addresses. Since ARP is a stateless protocol, every time a host gets an ARP reply from another host, even though it has not sent an ARP request for that reply, it accepts that ARP entry and updates its ARP cache [2]. The process of updating a target host's ARP cache with a forged entry is referred to as poisoning.

Content-Addressable Memory (CAM) tables store MAC addresses, switch port numbers, and Virtual Local Area Network (VLAN) information at switches. They are fixed sizes. In the MAC flooding attack, the attacker floods the switch with MAC addresses using forged gratuitous ARP packets until the CAM table is full. Then, the switch goes into hub-like mode and starts broadcasting the traffic without a CAM entry.

The port stealing attack uses the ability of the switches to learn to bind MAC addresses to ports. When a switch receives traffic from a port with a MAC source address, it binds the port number and the MAC address. In this attack, the attacker floods the switch with forged gratuitous ARP frames with the target host's MAC address as the source address and the attacker's MAC address as the destination address. Since the target host sends frames as well, there is a race condition. If the attacker is fast enough, frames intended for the target host are sent to the attacker's switch port and not to the target host.

In addition to these attacks, there are Layer 2-based broadcasting, Denial of Service (DoS), MAC cloning, and hijacking attacks. In the broadcasting attack, the attacker sends spoofed ARP replies to the network. These ARP replies set the MAC address of the network router to the broadcast address. This causes all the outbound traffic to be broadcast enabling sniffing. This type of attack also affects the network capacity. In the Layer 2-based DoS attacks, the attacker updates the ARP caches in the network with non-existent MAC addresses. The MAC address of each network interface card in a network is supposed to be globally unique. However, it can easily be changed enabling MAC cloning. The attacker uses a DoS attack to disable the network connection of the victim and then uses the IP and MAC addresses of the victim. In the Layer 2-based hijacking attack, an attacker takes control of a connection between two computers in the network. For instance, the attacker takes control of a telnet session after the victim logs in to a remote computer.

There are several ways to mitigate these types of attacks. One of these actions is to enable port security on switches. Port security ties a physical port on a switch to a MAC address/es. A change in the specified MAC address/es for a port or flooding of a port can be controlled in many different ways through switch administration. The port can be configured to shut down or block the MAC addresses that exceed a specified limit. The recommended best practice is to shut down the port that exceeds the limit [1]. Port security prevents MAC flooding and cloning attacks. However, port security does not prevent ARP spoofing. Port security validates the MAC source address in the frame header, but ARP frames contain an additional MAC source field in the data payload, and clients use

this field to populate their caches [3]. Another recommended action is to employ static ARP entries. Static ARP entries are permanent entries in an ARP cache. It prevents most of the attacks. However, this method is impractical. Furthermore, it does not allow the use of some Dynamic Host Configuration Protocol (DHCP) configurations. The third method of defense is to utilize Intrusion Detection Systems (IDSs). These can be configured to listen for high amounts of ARP traffic. However, IDSs are prone to reporting false positives. There are also tools specifically designed to listen for ARP traffic on the networks. It is possible to utilize Reverse ARP to detect MAC cloning as well. In addition, there are methods to detect machines in promiscuous mode on the network.

In local networks, VLANs are employed as a security measure to limit the number of clients susceptible to attacks. VLANs create network boundaries, which ARP traffic cannot cross. Then again, VLANs are not always an option and have their own set of vulnerabilities. VLAN Hopping, Spanning Tree, and Private VLAN attacks are some of the possible attacks in VLANs.

VLAN hopping attacks allow attackers to bypass a Layer 3 device when communicating from one VLAN to another. The attack works by taking advantage of an incorrectly configured trunk port [1]. Trunk ports are generally used between switches to route traffic for multiple VLANs across the same physical link. Since the basic VLAN hopping attack is prevented in the new versions of switches, attackers have developed Double Encapsulated VLAN Hopping attacks [4]. This attack uses the fact that switches perform only one level of decapsulation. To mitigate this type of attack, administrators should disable Auto-trunking, use a dedicated VLAN ID for all trunk ports, disable unused ports and put them in an unused VLAN, and avoid using VLAN 1 (only the defaults are allowed in VLAN 1).

Spanning Tree Protocol (STP) is a link management protocol that provides loop-free topologies in a redundant Layer 2 infrastructure. The STP elects a root bridge to prevent loops in a network. In the STP, messages are sent using Bridge Protocol Data Units (BPDUs). The Standard 802.1D STP takes about 30-50 seconds to deal with a failure or root bridge change. The attacker sends BPDUs to force these changes creating a DoS condition in the network [5]. There are two features on switches used to mitigate this type of attack: BPDU Guard and Root Guard. BPDU Guard disables ports upon detection of a BPDU message on the interface. Root Guard disables interfaces that become the root bridge due to their BPDU advertisement.

Private VLANs (PVLANS) are used to create distinct networks within a VLAN. PVLANS work by limiting which ports within a VLAN can communicate with the other ports in the same VLAN. The attacker sends a frame with a rogue MAC address (the one of the router) but with the IP address of the victim. Switches do not forward the frame to the victim, but the router forwards the packet to the victim. To mitigate this attack, an ingress Access Control List (ACL) can be setup on the router interface or VLAN ACL (VACL) can be used.

Lastly, Dynamic Host Configuration Protocol (DHCP) is used to dynamically allocate IP addresses to computers for a time period. It is possible to attack DHCP servers causing DoS in the network or impersonate a DHCP server. For instance, in the DHCP starvation attack, the attacker requests all of the available DHCP addresses. This results in a DoS attack on the network. The attacker can also use a rogue DHCP server to provide addresses to the clients. The attacker can point the users to a different default gateway with the DHCP responses. Authentication of the DHCP messages is required to prevent this type of attack.

We have presented eleven possible Layer 2 attacks in this section. However, this list is not comprehensive. Other attacks worth mentioning are Multicast Brute-Force Failover Analysis, Random Frame Stress attack, and attacks based on proprietary protocols. Furthermore, most of the network management protocols are insecure causing additional vulnerabilities.

## 2 MACsec

Recently, the 802.1AE Media Access Control (MAC) Security Task Group has been formed to secure Local or Metropolitan Area Networks. The IEEE P802.1AE Standard (Draft) for Local and Metropolitan Area Networks (LAN/MANs): MAC Security specifies how all or a part of a network can be secured transparently to peer protocol entities that use the MAC Service provided by IEEE 802 LANs to communicate. The draft defines MAC security (MACsec) entities in end stations that provide connectionless user data confidentiality, frame data integrity, and data origin authenticity. However, the standard's scope does not include the key management and the establishment of secure associations [6].

MACsec provides security services on a frame by frame basis without introducing any additional frames. MACsec introduces an additional transit delay due to the increase in the MAC Service Data Unit (MSDU) size. MACsec defines how a MAC Security Entity (SecY) operates with a MAC Security Key Agreement Entity (KaY). Each KaY discovers the KaYs present in other stations attached to the same LAN, mutually authenticates, and authorizes those stations, and creates and maintains the secure relationships between the stations that are used by the SecYs to transmit and receive frames [6]. However, MACsec does not specify how the KaY works.

There is only one Connectivity Association (CA) per LAN service. In [6], the abbreviation LAN is used exclusively to refer to an individual LAN specified by a MAC technology without the inclusion of bridges. Each SecY only participates in a single CA at any one time. Each CA is supported by Secure Channels (SCs). There is one SC for secure transmission of frames from one of the systems to all the others in the CA. All the SCs use the same cipher suite at any one time. Each SC comprises a succession of SAs, each SA representing a single value of the transient session key(s) used for a period by the cipher suite to support the SC. Each SA is identified by the SC Identifier (SCI) concatenated with an Association Number (AN). The Secure Association Identifier (SAI) thus

created allows the receiving SecY to identify the SA, and thus the session key(s) to be used to decrypt and authenticate the received frame. When the service guarantees provided include replay protection, the MACsec protocol requires a separate replay protection sequence number counter for each SA, as well.

The SecY provides both secure and insecure services to the users of its Controlled Port and Uncontrolled Port respectively, which are part of the IEEE 802.1X. The SecY operates without integrity, origin, or confidentiality protection if the Null Cipher Suite is selected. The services provided by the SecY when a cipher suite is selected include the MAC Service Data Unit (MSDU) encryption, Integrity Check Value (ICV) calculation to protect the MAC Protocol Data Unit (MPDU), and inclusion of a SC field. The SC presents the address where the encryption is applied. In a multipoint or Provider Bridge network, the MAC source address (SA) and destination address (DA) are not the addresses of the intermediate devices that are encrypting and decrypting, they are the original addresses, the end-to-end addresses. If the SecY is a part of the Bridge stack, its address will not be seen at the end points. In that case, in order to make the knowledge of where the encryption is applied available, the SC is used to provide the address of the Bridge port.




Fig. 1. Ethernet frame format with MACsec

The SecY can include a Security TAG (SecTAG) in the initial octets of the provider MSDU, prior to the user data and ICV. The MACsec protocol specifies the mandatory cipher and integrity suites as Null, Galois/Counter Mode-Advanced Encryption Standard (GCM-AES), and GCM as Message Authentication Code (GMAC). The cipher suites, except the Null Cipher Suite, provide confidentiality or integrity or both confidentiality and integrity. The addition of a SecTAG is required for the cipher suites, except the Null Cipher Suite. KaY associated with the SecY provides and periodically updates the keys for the cipher suite. If confidentiality is required, the user data parameter (MSDU) is encrypted. The destination address, source address, and the SecTAG fields are not encrypted. If data integrity is required, an Integrity Check Value (ICV) is calculated over the destination address, source address, SecTAG, and user data (after encryption, if applicable). A simple frame format for Ethernet using MACsec is presented in Figure 1.

When a Cipher Suite is selected (except the Null Cipher Suite), the SAI decoded from the SecTAG of a valid MPDU, the MAC destination and source addresses, the octets of the SecTAG, the octets of the Secure Data, and the ICV are presented to the Cipher Suite implementation. The Cipher Suite implementation identifies the validation parameters associated with the SA for the received frame using the SAI. Then, (using the validation parameters) it validates the

addresses, the SecTAG, and the User Data and decrypts (if encrypted) the User Data. If any of the parameters are invalid or MSDU extraction fails, the received frame is discarded, and no indication is made to the user of the SecY.

The ICV is encoded in the last eight octets of the MACsec PDU. It authenticates the MAC destination and source addresses as well as all the fields of the MPDU (including the SecTAG if present). The ICV is computed over the encrypted or clear text data. There are significant advantages of computing the ICV over the encrypted text.

MACsec provides point to point integrity, but not global integrity [6]. There is no protection against legitimate users in ARP spoofing. It is a case of a legitimate user with bad intentions. MACsec does offer the ability to identify the bad user. The ICV is recomputed every time the frame presented to the MAC layer, so a man-in-the-middle (legitimate user) can make unauthorized changes and it will pass the integrity check. If the security transform is applied only to data, not to control frames, it will not offer protection against disclosure due to ARP spoofing, in which an attacker sends gratuitous ARP messages claiming to have IP addresses of other stations. The attacker can then intercept, read, and alter messages between any two points in a point to point topology. The ICV will be recalculated and the altered data will pass the integrity check. If the message is cryptographically protected above Layer 2, for example with IPsec, the data cannot be read or changed by the attacker. This threat occurs at the intersection between the Layer 2 and Layer 3 layers [6].

### 3 Layer 2/Layer 3 Impact

Layer 2 switches/bridges are used to provide connectivity in LANs, whereas Layer 3 routers are typically used to provide connectivity between LANs. For instance, a switch forwards Ethernet frames based on the MAC addresses or frame headers. It does not alter the MAC addresses when it forwards a frame. On the other hand, a router extracts the network layer header to match the destination address in the routing table, identifies the new MAC destination address for the packet, and transmits the packet to the port associated with the destination address in the routing table. As a result, when a frame/packet goes through a router, the Layer 2 header of the frame/packet is changed. Hence, the information regarding the original MAC source and destination addresses is lost.

The goal of MACsec is to secure LAN/MANs. However, most of the LAN/MANs are composed of bridges/switches and Layer 3 devices (routers). Routers support the Internet Protocol (IP). When a router receives a data frame destined to itself, it removes all the link layer (Layer 2) headers and adds a new link layer header before transmitting it. MACsec mentions an optional Secure Origin Authenticity (SOA) field. Nevertheless, it is not clear how this field is used at switches/bridges. In addition, MACsec does not consider the case when a Layer 3 device forwards a frame. Since MACsec is end-to-end in Layer 2, in the case that the router is the end point for a transmission in Layer 2, all the information regarding the origin of security (original MAC source address) is removed. This

limits the capability of tracking spoofed IP/MAC packets/frames because the IP and MAC address pair in the outgoing packet/frame at the router does not provide any information regarding the original MAC source address used. In addition, it removes the IP source address and the MAC source address binding created in the original frame. When MACsec is used with a SecTAG, it provides security for the data frame and binds the MAC and IP addresses via the ICV. The ICV is calculated over the MAC destination and source addresses, SecTAG, and user data. Thus, it prevents unauthorized modifications of the MAC and IP addresses. The binding between the IP and MAC source addresses is critical to provide security and aid in billing procedures. Note that a DHCP server may assign a different IP address to the same subscriber each time it joins the network even though the MAC address of the subscriber stays the same.

We propose to transmit the original MAC source address in the Ethernet frame when a router transmits the frame in a LAN/MAN. If MACsec is not used to secure the ARP messages, the inclusion of the original MAC source address may be used to prevent the ARP attacks in a LAN/MAN. At a router, before forwarding a frame, the original MAC source address should be included in the frame in addition to the SecTAG. The maximum size of the data field in a frame should be reduced to convey the original 48 bit MAC source address field. In addition, this field may be made optional to conserve bandwidth in a network.

Header	SecTAG w/LC	Original MAC Source Address	Data	ICV	CRC
--------	----------------	--------------------------------	------	-----	-----

**Fig. 2.** Proposed Ethernet frame format with MACsec

Furthermore, we propose to add a hop/link count field into the SecTAG to track the number of hops/links that a frame travels. This is necessary because switches/bridges do not have MAC addresses. They are invisible in Layer 2/3. Even though the port number is contained in the SecTAG, it is not possible to identify the number of switches/bridges that a frame passes through in a network. The port number in the SecTAG identifies the port number of the last end point bridge/switch. A hop/link count field makes end devices aware of intermediate switching/bridging devices. Moreover, it helps to track the traffic in a network providing the information regarding the number of Layer 2 devices on the path. In conjunction with topology plans, network administrators can also examine whether frames take the expected path and IDSs can utilize this field to recognize spoofed or misguided frames. Note that a hop/link count field denotes the number of hops/links in Layer 2 instead of Layer 3. An example for a similar concept with a different intention would be the IP time to live field at Layer 3. The proposed field should have a fixed size to prevent frame fragmentations later in the network. Each SecY should increment the hop/link count before transmitting a frame. Routers require additional features to transfer this link

layer information between the ports, as well. However, it may always not be desirable to reveal the number of Layer 2 devices in a network. In such a case, the TAG Control Information (TCI) field, which comprises bits 3 through 8 of the octet 3 of the SecTAG, may be used to indicate whether the hop/link count field is being utilized. A simple presentation of the proposed Ethernet frame is shown in Figure 2. The proposed additional fields should be protected by the ICV in each frame. The TCI field in the SecTAG can be used to facilitate the optional use of these additional fields. In Figure 2, LC stands for Link Count and is included in the SecTAG.

## 4 Conclusions

In this paper, we focus on the Data Link Layer (Layer 2) security issues in IP over Ethernet networks. We introduce the security concerns in Layer 2 and summarize some of the possible attacks in this layer. In addition, we provide an overview of the IEEE P802.1AE Standard for Local and Metropolitan Area Networks (LAN/MANs): Media Access Control (MAC) Security. We discuss the Layer 2/Layer 3 impact of MACsec, as well. Finally, as an initial approach to improve security in LANs, we propose to use a hop/link count field in the SecTAG and an original MAC source address field in addition to the SecTAG at a Layer 3 device before transmitting a frame. These proposed additional fields provide some level of visibility for Layer 2 devices and protect the original MAC and IP address binding with additional bandwidth requirements.

Future work should focus on performance of MACsec and Layer 2 security threats with MACsec. Moreover, methods to secure ARP should be investigated. Finally, the IP and MAC address binding problem should be studied and addressed in general.

## References

1. Howard, C.: Layer 2 – The Weakest Link: Security Considerations at the Data Link Layer. Available at [http://www.cisco.com/en/US/about/ac123/ac114/ac173/ac222/about\\_cisco\\_packet\\_feature09186a0080142deb.html](http://www.cisco.com/en/US/about/ac123/ac114/ac173/ac222/about_cisco_packet_feature09186a0080142deb.html)
2. Bashir, M. S.: ARP Cache Poisoning with Ettercap. (August 2003) Available at <http://www.giac.org/practical/GSEC/Mohammad.Bashir.GSEC.pdf>
3. Plummer, D. C.: Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware. RFC 826 (November 1982)
4. Rouiller, S. A.: Virtual LAN Security: weaknesses and counter measures. Available at <http://www.sans.org/rr/papers/38/1090.pdf>
5. Convery, S.: Hacking Layer 2: Fun with Ethernet Switches. (Blackhat, 2002) Available at <http://www.blackhat.com/presentations/bh-usa-02/bh-us-02-converyswitches.pdf>
6. IEEE P802.1AE/D2.0 Draft Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Security. Available at <http://www.ieee802.org/1/files/private/ae-drafts/d2/802-1ae-d2-01.pdf>