

A Survey of Ethernet LAN Security

Timo Kiravuo, Mikko Särelä, and Jukka Manner

Abstract—Ethernet is the survivor of the LAN wars. It is hard to find an IP packet that has not passed over an Ethernet segment. One important reason for this is Ethernet's simplicity and ease of configuration. However, Ethernet has always been known to be an insecure technology. Recent successful malware attacks and the move towards cloud computing in data centers demand that attention be paid to the security aspects of Ethernet.

In this paper, we present known Ethernet related threats and discuss existing solutions from business, hacker, and academic communities. Major issues, like insecurities related to Address Resolution Protocol and to self-configurability, are discussed. The solutions fall roughly into three categories: accepting Ethernet's insecurity and circling it with firewalls; creating a logical separation between the switches and end hosts; and centralized cryptography based schemes. However, none of the above provides the perfect combination of simplicity and security befitting Ethernet.

Index Terms—Ethernet networks, Security, Data security

I. INTRODUCTION

A USB stick with a company logo is found in the company parking lot. A helpful employee picks it up and inserts it into his workstation to find out who it belongs to. Visibly, nothing happens, but underneath, a virus is delivered to the workstation residing in the local corporate network. From there, it can monitor the, typically insecure, Ethernet segment, invade new hosts, or cause harm in many other ways.

In June 2010, a new virus was discovered. Named Stuxnet, it targeted the Siemens' programmable logic controllers used in industrial automation, known to be used in the nuclear program of Iran. Stuxnet travels in removable drives and, upon activation, spreads to hosts found in the target Local Area Network (LAN) using several sophisticated methods for entry. USB key drives enable Stuxnet to enter protected networks, not connected to the Internet, and once there, LAN technologies such as Ethernet offer all the hosts to the virus on a silver plate [1].

Ethernet is the infrastructure for the Internet that everybody uses without further thought. It can be found in offices, homes, and computing centers and it is being extended to carrier, automotive, avionics, and industrial uses [2]. Originally Ethernet was designed to be a flexible, decentralized and low cost LAN with 3 Mbps capacity for up to 256 hosts [3]. The capacity has grown since, but the original qualities have ensured that Ethernet in various forms is virtually the only wired LAN technology in use today. This ubiquitous work-horse became popular in 1980s when its 10 Mbps capacity was shared by all the hosts hanging on the same co-axial cable. Today's 1 Gbps,

switched, full-duplex, collision free Ethernet is considered the standard low cost LAN solution for workstations and laptops, and 10 Gbps is commonly available for servers and high throughput hosts.

Ethernet segments are also being expanded, both in distance and capacity, using various techniques [4]. Network operators are starting to design edge to edge Ethernets, using the layer 2 Ethernet internally to replace higher layer activities like IP routing and addressing, leaving them to be considered only at the edges of the network. The motivation for having larger Ethernet segments is to handle traffic at a layer lower than the IP layer. Replacing IP routers with Ethernet switches makes network configuration easier and faster, thanks to Ethernet being self-configuring. Switches can be simpler than routers, which lowers both cost and energy consumption. An IP router must locate the IP header from the frame and perform a longest-prefix matching based on the destination address, decrease the time to live field and recalculate the checksum. An Ethernet switch just needs to find recipient's Media Access Control (MAC) address in the MAC table.

The architectural security of Ethernet has received little attention from the academic research community. Most of the existing research, especially the early work, has focused on developing cryptographic solutions to the perceived problems [5]–[9]. Equipment vendors and the hacker community have performed most of the practical work related to Ethernet security. These results are published in vendor documentation, at hacker conferences, on security related web pages, and on the individual web pages of interested people. Common security textbooks [10]–[13] pay little or no attention to Ethernet specific aspects and even network security books focus mostly on higher layer protocols [14], [15]. Vendor publications discuss Ethernet's security issues, but focus on solving them using vendor products [16], [17]. The goal of this paper is to collect the information available and analyze it in a way that lets us understand the security related properties of Ethernet technology.

The security of Ethernet should be of interest to many kinds of people. Data center managers benefit from the self-configurability and flexibility of Ethernet. For security professionals, Ethernet is one of the technologies they have to manage and evaluate. Networking technicians should be aware of the limitations and security issues of the most popular LAN technology in use. For the research community, the combination of simplicity and security presents an interesting research problem.

Attackers have various motivations for targeting the vulnerabilities of the Ethernet layer. The attacker may try get access to information, perhaps change it, even encrypt and hold data hostage or prevent its use by legitimate users. The potential

Manuscript received 30 November 2011; revised 8 June 2012.

The authors are with Aalto University, Department of Communications and Networking, Finland

Digital Object Identifier 10.1109/SURV.2012.121112.00190

Preamble	Destination address	Source address	Type or length	Payload	CRC
8	6	6	2	46-1500	4

Fig. 1. Ethernet frame format, units in bytes.

benefits from targeting the Ethernet network grow with the amount and importance of data reachable through such an attack. The resources an attacker has depend on who they are; attackers may vary from unsatisfied employees through to competing companies, from organized crime to governmental organizations.

Since the network itself does not hold any data, Ethernet should be seen mostly as a medium for attacks to hosts. However, the attacker might be satisfied by just having access to the traffic on the network or by being able to disrupt services.

This work focuses on pure Ethernet, which is practically the only wired LAN technology today. The security of wireless technologies, e.g., 802.11 Wireless LAN (WLAN), is out of the scope of this survey. See, e.g., [18], [19] for a survey on WLAN and wireless sensor security. However, many of our findings can be applied to wireless LAN technologies, especially the IEEE 802.11 WLAN. We also think these findings are relevant to the new Ethernet based technologies like Provider Backbone Bridging, Global Open Ethernet or Transparent Interconnection of Lots of Links, which enable ceration of larger Ethernet segments [2], [4].

Layer 1 (physical) issues are generally outside the scope of this article. So is the behavior of the higher layers, except when it is closely tied to the behavior of Ethernet.

This paper is organized as follows. Next, in Section II, we describe the technologies that together form what is known as “corporate Ethernet”. Section III describes the various security threats against Ethernet and Section IV presents the existing solutions and security related technologies. Potential future solutions are discussed in Section V and Section VI concludes this review.

II. ETHERNET TODAY

To set a baseline against which new improvements to Ethernet are evaluated, we define “Plain Ethernet” as a full-duplex, twisted pair based Ethernet consisting of hosts joined by multiport bridges (switches). The star configurations of individual switches may be joined to form a physical mesh, but logically the network is a tree. The Virtual LAN (VLAN) technology can be used to split the network into independent logical networks that may form their own trees.

The standards body responsible for Ethernet is the IEEE’s 802 committee, which is responsible for packet networking. The working group 802.3 is in charge of Ethernet transmission standards. Most of the security related work is done by the working group 802.1. Many Ethernet switch features have been created by vendors and are not standardized. The Internet Engineering Task Force (IETF) is also active in those areas of Ethernet that relate to the use of Internet Protocol (IP) over Ethernet.




Fig. 2. An Ethernet segment with switches and hosts; addresses are abbreviated; the MAC table for bb:bb is shown.

A. The Ethernet Frame

Each node on an Ethernet segment has the hardware or software capability to send and receive Ethernet *frames*. The structure of the frame is presented in Fig. 1. A frame starts with an eight octet preamble of zeros and ones, enabling synchronization for the receiving host. It is followed directly by the six octet destination address, after which comes the sender address, both known as Media Access Control (MAC) addresses. The addresses are locally unique; the same address can not be in use twice in the same LAN, but it may be used in other LANs.

The address of the network interface used to be hard-coded in the Network Interface Card (NIC), but on modern equipment it can be changed using software. Locally generated addresses should have the value 1 in the second least significant bit of the first octet of the address. Most NICs today also support user generated frames that can contain arbitrary addresses. Frames can be sent to unicast, multicast, and broadcast addresses. Multicast addresses have the least significant bit of the first octet set to 1. A broadcast address has all bits of the address set to 1.

After the addresses, the frame has a two byte ethertype field that describes the length or type of the payload. The type field is followed by the payload, usually a higher layer like an IP packet. If the payload is shorter than 46 bytes, padding is added to bring it up to size. The frame ends with a four byte checksum of the contents, called the Frame Check Sequence (FCS) or Cyclic Redundancy Check (CRC).

B. An Ethernet Switch

The switch connects the edge nodes to each other through its multiple input ports (interfaces) and internal switching fabric, which can move frames from one port to another. The switch is not initially aware of each node’s MAC address. As a node sends a frame, the switch learns from which port this frame arrived and adds the sender’s MAC address and port number to a table.

The switch looks for the recipient’s address from its memory and, if it is not found, the frame is resent out to all the

other ports (flooded). Flooding the frames works, because end nodes ignore frames in which the recipient's address does not match their own and only the correct recipient accepts the frame up to higher layers for processing.

When further frames are sent on the network, the switch rapidly builds a table of port and MAC address pairings. As only one frame is required to identify each host, a table that maps the switch's partial view of the network's topology is built quickly and unicast frames are sent out to only the one port corresponding to the recipient's address. Fig. 2 shows an Ethernet segment where switch `bb:bb` has learned the MAC addresses of neighboring switches and two hosts. One of the hosts is behind another switch, thus from `bb:bb`'s viewpoint these are connected to the same port.

Switches can also be connected to each other. To a switch, another switch is just multiple MAC addresses behind one port. It is possible and often desirable to build a mesh topology of switches to increase redundancy. The links connecting the switches to each other are often called *trunk links*, especially when carrying traffic belonging to multiple VLANs.

Typically a switch has three layers of structure:

- The *data plane* forwards frames from one port to another (or others) and is usually implemented in hardware. The interconnection fabric, or back plane bus, has much higher throughput than the I/O ports of the switch; this enables several simultaneous flows to pass at full line capacity.
- The *control plane* handles the frames that need processing, like the frames whose addresses are not listed in the address table or Spanning Tree Protocol (STP) messages, which configure the network of switches. This is usually done by the switch's Central Processing Unit (CPU).
- The *management plane* is used to configure the switch's features like the VLAN networks. The implementation includes usually a TCP/IP stack, a command line interface and a Simple Network Management Protocol (SNMP) agent.

The switch's MAC address table lists known MAC addresses, and for each address the port from which the latest frame was received. Depending on the implementation, there may also be a VLAN identifier for this address, a timestamp for the time out feature, and additional information. This table is named Content Addressable Memory (CAM) when it is implemented using associative memory hardware that enables fast searches. If a host moves to a different port, the address table entry for that MAC address will be updated when the first frame from the new location reaches the switch.

C. Spanning Tree Protocol

The IEEE 802.1D *Spanning Tree Protocol* (STP) is a method for avoiding loops in the LAN [20]. Ethernet protocol has no internal mechanism for avoiding loops. When connected in a mesh topology, switches would receive the same frame over several links and have to decide which port to enter into the MAC address table. Sooner or later (more likely sooner) the individual MAC tables would form a loop together and frames would start to circulate within the network congesting it.




Fig. 3. Creating the spanning tree; switch `bb:bb` becomes the root switch.

STP is a solution to this problem. One of the switches is initially selected to act as a root node and broadcasts Bridge Protocol Data Units (BPDUs), which have a cost (in practice, a hop counter that takes link capacity into account). Each switch increments the cost and floods the frame out of the other ports. If a switch receives a BPDU from two ports, it blocks the port with the higher cost. Thus, as shown in Fig. 3, the mesh converges to a tree configuration and switch `bb:bb` is the root of the tree. If an active link between two hosts is lost, they (or one of them) send a Topology Change Notice (TCN) BPDU to the root switch, which broadcasts a TCN message to all switches, and the tree reconfigures. Several versions of STP exist, Rapid STP improves performance from the original, Multiple STP supports separate spanning trees for each VLAN, and vendors have developed their own versions for similar needs.

Ethernet, as presented up to this point, is a self-configuring network. The switches identify each other from STP messages and form a tree structure. By monitoring the sender's addresses in incoming frames, the switches learn the relative locations of each host. If a network link between the switches is removed or added, STP reconfigures the network. If a host's relative position in the network topology is changed by this, frames sent by the host will overwrite old entries in the switches' memories along the communication path and the network accommodates the changes.

D. Layer 3 Adaptation Protocols

Two protocols are needed for the IP version 4 (IPv4) to operate over Ethernet. These are sometimes called layer 2.5 protocols. *Dynamic Host Configuration Protocol* (DHCP) [21] is used to request an IP address for a host. When an IPv4 host without an IP address becomes active on an Ethernet segment it sends a request for DHCP servers using an Ethernet broadcast. Upon receiving one or more unicast replies, the host selects one server and requests an IP address with a unicast message and, upon success, receives a lease for an IP address and additional information, such as the netmask and the gateway's (router's) IP address. IP addresses may also be configured statically at the host, in which case DHCP is not required.

Address Resolution Protocol (ARP) [22] is needed for IP to operate on shared media like Ethernet, as the MAC addresses need to be mapped to corresponding IP addresses. When a host wishes to communicate with another host in the LAN, like the gateway, it sends a broadcast message requesting a MAC




Fig. 4. Switches are configured to know which VLANs are in which ports.

address that corresponds to the IP address in the message. The host, which has the IP address in use, responds with a unicast message. The recipient stores the IP and the MAC address pair in a table (the ARP cache) for some duration (30 seconds to 5 minutes is common, depending on the operating system). Implementations, especially older ones, are often stateless and thus a host will cache ARP replies even if they have not been requested. This feature is used in gratuitous ARP when a host that has changed its NIC wants to update its address in other hosts' tables.

IPv6 has similar functions. Hosts are found with *Neighbor Discovery Protocol* [23], which uses Ethernet multicast (switches can implement Ethernet multicast as a broadcast, but the use of multicast addresses helps hosts' NICs to filter the frames) or by using DHCPv6 [24]. IPv6 routers are found by listening for multicast Router Advertisements, from which a host can create its own IPv6 address and use Neighbor Discovery to verify its uniqueness.

Switches may also include other layer 3 functionality, which is not further discussed in this paper except for IP multicast handling. IP routers require registration for a multicast group with IPv4 *Internet Group Management Protocol* (IGMP) [25] or IPv6 *Multicast Listener Discovery* (MLD) [26]. However, the multicast packets are typically sent to the Ethernet layer as broadcast frames, thus flooding the network. A switch can snoop into the layer 3 registration messages and build a table of multicast listeners, thus forwarding multicast frames only to those ports where a listener is active. This vendor dependent feature is usually found only in the high-end switches.

Several other infrastructure protocols that depend on Ethernet broadcast or multicast transmission are not discussed here. For example, service discovery protocols such as Universal Plug and Play or Bonjour are not relevant to the basic operations of the Ethernet layer, and thus fall outside the scope of this paper.

E. Virtual LAN

VLANs are used to separate a physical network into several logical networks. Each switch in the network keeps a table associating its ports with the various VLAN identifiers in use. Fig. 4 shows such a network, with two VLANs in use. The motivation for the VLAN mechanism is to increase efficiency by limiting the size of the broadcast domain, but it is used also for security purposes [27]. Hosts in different VLANs can not send frames to each other directly.

To separate the frames when transiting between the switches, the IEEE 802.1Q VLAN mechanism adds a four byte VLAN tag inside the Ethernet header, between the sender-MAC address and the ethertype fields [28]. IEEE 802.1ad adds a second tag to create separate local and provider VLAN domains [29]. The VLAN tag's first two bytes contain the value 0x8100 to notify VLAN capable switches. Older switches can process these frames transparently as the first two bytes of the tag match the position of the ethertype field. The purpose of the other two bytes is to identify which particular VLAN this packet belongs to. Switches enforce the boundaries of these LANs, providing additional security. The tag is added when a frame arrives from the host to the first switch and is removed at the final switch before delivery to destination. The frame can also be delivered to the host with the tag if desired, for example if the host contains virtual hosts.

VLAN management protocols can be used to configure switches and match their parameters. *Multiple VLAN Registration Protocol* (MVRP) [30] is an IEEE replacement for Cisco proprietary protocols *VLAN Trunking Protocol* (VTP) and *Dynamic Trunking Protocol* (DTP).

F. Layer 2 Control Plane Protocols

Besides the previously mentioned protocols, there are several more protocols that are linked to Ethernet. These protocols have their own security issues that will not be discussed further beyond this section. While related to Ethernet, these protocols themselves are implemented in higher layers and thus mostly not relevant to Ethernet architecture.

The redundancy protocols enable critical nodes like gateway routers to communicate their state to their backups. These protocols are often used in a separate (VLAN) segment. *Hot Standby Router Protocol* (HSRP) is a Cisco proprietary protocol designed for multiple redundant routers to communicate on the active and standby roles. The routers share a virtual MAC and IP address. The messages are sent using IP multicast and are authenticated with a clear text password [31]. *Virtual Router Redundancy Protocol* (VRRP) is a standardized protocol that was designed to replace HSRP [32].

Network topology discovery protocols are used mostly by network management systems to find out how the network is organized. Vendors are now switching to the IEEE *Link Layer Discovery Protocol* [33]. These protocols send Ethernet multicasts in which they report on the node's connectivity, addresses and capabilities. Other similar protocols include *Cisco Discovery Protocol* and Microsoft's *Link Layer Topology Discovery*.

Link aggregation technologies combine multiple physical links to appear as one link for traffic, especially for STP to utilize them as one high-capacity link when forming the tree. Several vendor methods exist and also the IEEE 802.1AX [34] standard defines *Link Aggregation Control Protocol* for this. Link aggregation can be configured statically or switches can probe their links to see if they have multiple links to other switches. These links can be aggregated and treated as one link.

III. ETHERNET THREATS

This section describes known Ethernet related security threats. We focus on an Ethernet segment bordered by layer 3 routers terminating the layer 2 Ethernet traffic. The key historical reason for the security vulnerabilities of Ethernet is that security has never been a major consideration in its design. The whole architecture reflects the (highly useful and proven) goal of a cheap and easily deployable LAN.

Ethernet's security, and lack of it, is fundamentally tied to its self-configuring nature. It is a great advantage to be able to install and expand a LAN just by connecting switches and computers together with cabling and have it work automatically. However the features that enable this, like MAC table learning, STP and ARP together with the underlying broadcasting mechanism, are also key vulnerabilities. Furthermore VLANs do not provide sufficient separation between segments, as a switch has no automatic way of knowing if it is connected to another switch or to an end host pretending to be a switch.

The basis for attacks is gaining access to the target Ethernet segment. The attacker may be an insider with full access rights, may have found an Ethernet connection in a public space [35], or may have taken control of a workstation using a malware application, or other methods [36].

In the rest of this Section, we shall describe the most prominent methods for attacking Ethernet segments. However, before that, a few words about attacker motivation.

The attacker may utilize the network access for: (1) learning about the private network topology and the network traffic for use in a later attack, (2) gaining control over switches, routers, or servers in the LAN, (3) eavesdropping, (4) manipulating information, or (5) disrupting the availability of the network.

The attacker's motives and methods may vary and while the whole attack may be entirely technical in nature, it should also be noted that all five are useful as a part of a social engineering attack. As an example, knowledge of the MAC address of a particular host (or the ability to cause a problem in the network just before one complains about that over the telephone) may help convince IT support staff of the authenticity of a caller on the telephone [35].

Although much work has been done to secure protocols and applications on upper layers, many systems designed for internal use in organizations still rely on the assumption that the network is secure. The usefulness of motives (2–4) is partly tied to this, while (5) is a more universal network level problem that cannot be resolved at the upper layers [37].

A. Network and System Access

Access to the network is a prerequisite for attacks [38] and a necessity for all types of attackers. Access can be achieved by connecting equipment to the network or by gaining control of existing resources.

Besides targeting Ethernet's features directly, an attacker may use Ethernet to attack other targets. These attacks are made possible by implementation issues and are not caused directly by the architecture of Ethernet. It should be also noted that protocols mentioned in Section II-F facilitate topology discovery and their implementations may be prone to breaks or denial of service (Dos) attacks.




Fig. 5. VLAN double tagging attack; Attacker A's frames reach target T.

1) *Unauthorized Joins*: Ethernet has been designed to be easy to deploy and to require minimal administration overhead. Hence anybody can connect to an Ethernet segment by gaining access to an unconnected port on a switch. This can be done by: gaining physical access to the switch, gaining access to a wall socket, removing the cable from a computer and plugging it into another computer, or plugging in a switch between the existing computer and the socket. It is also fairly common that, when a host is removed from the network, the connection from the socket to the switch is not disconnected, thus leaving wall sockets connected to switches.

2) *Unauthorized Expansion of the Network*: The architecture of the Ethernet allows users to expand the network by installing their own switches or wireless access points, which in turn allows other people join the network. These will automatically be allowed to join the network, unless the switches are configured to prevent it by limiting new MAC addresses as described in next section.

3) *VLAN Join*: If a switch listens for VLAN management protocols on host ports, a host can act as a switch and join all VLANs. On some switches the ports can be configured to not transmit (advertise) VLAN management protocols, but they will still listen for these protocols. An attacker can probe the switch for hidden features like this.

4) *VLAN Tagging and Hopping*: An attacker can create Ethernet frames that have a VLAN tag and thus inject frames to VLANs to which they are not supposed to have access. There are several variations of this attack. In the "double tagging" attack shown in Fig. 5, the attacker's host belongs to VLAN 1 and the trunk link from the local switch is configured as also belonging to VLAN 1. The attacker creates a frame which has the target host's MAC address as the recipient and contains a VLAN 1 tag followed by a VLAN 2 tag. The local switch notices the first tag. Since the frame's MAC address directs it towards the trunk which also belongs to VLAN 1, the switch strips the tag off and pushes the frame to the trunk link of VLAN 1, where the receiving switch notices the second tag and processes the frame as belonging to the target VLAN [39]. The double tagging attack does not provide return traffic capability, but additional spoofing can do this, too [40]. Even without the replies, various attacks can be performed over the unidirectional flow.

VLAN hopping can also be achieved when a layer 3 device, such as an IP router, is serving several VLANs and is reachable through all of them. An attacker can send a frame with the router's LAN port's MAC address and the IP address

of a host in another VLAN, thus using layer 3 to bypass layer 2 restrictions. Depending on configuration, the router will receive the frame and forward it to the IP layer, inspect the IP address and resend it to the correct recipient on a VLAN other than the attacker's [39].

Some Voice over IP (VoIP) telephones use a specific VLAN to indicate the need for Quality of Service (QoS) [41]. The VLAN ID can be detected and after detection it can be possible for a workstation to join the VLAN segment [42].

5) *Remote Access to the LAN:* Access to an Ethernet segment can be achieved by gaining higher layer access to a host on the segment, for example by using social engineering and to get a user at the target network to open a remote system administration service, which then connects to a host on the Internet and enables the attacker to access the Ethernet layer.

6) *Topology and Vulnerability Discovery:* An attacker can probe the network to find hosts and services in them by sending messages and analyzing the replies. The goal can be to map the network's topology and services in hosts or to find vulnerabilities for further attacks. A similar mapping can be performed by network management systems for administration purposes or for security analysis.

The simple network topology can be mapped from the messages that a host sees and more information can be requested from the network nodes. Broadcast ARP requests reveal the IP addresses in use and servers or gateways to which other hosts connect to. The IP address range in use can be detected from this or the information can be requested from the DHCP server. Then connections can be attempted to the transport layer ports of hosts. This scanning process can be very detailed and will reveal plenty of information on hosts and their software, including operating systems, services, and versions, which leads to the identification of potential vulnerabilities.

Typically the goal of a scanning attack is to gain knowledge of the applications and services on a host and is not relevant to the Ethernet layer itself.

7) *Break-Ins:* An attacker can use the Ethernet network as a medium to attack other hosts and switches on the network. These attacks typically target vulnerabilities on higher layer network software, like the TCP/IP stack and especially server applications. They can lead to the capture of a host or a switch, which can be used for further attacks.

The attacker can also target the Ethernet firmware in the NIC and software at the host and attempt to get control of the interface.

8) *Switch Control:* As previously mentioned, switches are shipped with default or no passwords and the password can usually be physically reset. If an attacker gains control of a switch, traffic can be rerouted by switching links down, claiming the STP root by rising the priority of the switch or DoS selected links. However, as a switch is not a general purpose computer, its software limits the attacker's ability to eavesdrop on the traffic or generate spoofed frames; control of a workstation is needed for these attacks. In co-operation with a connected host the switch can be used to turn on mirroring for eavesdropping and, depending on what management protocols are operational on the network, potentially gain access to any VLAN in use.




Fig. 6. The forged MAC address lets A change the port assignment for T's address in the MAC table in the switch.

B. Traffic Confidentiality

Traffic on the network can be useful in itself and also serve the attacker in search of targets. An attacker gains information being transmitted, but also authentication information like passwords and network topology information that can be used for further purposes.

The original co-axial Ethernet was an easily eavesdroppable bus, where every station received every frame. Modern bridged Ethernet filters most of the traffic and a host receives only its own traffic, broadcasts, and random frames flooded at the switch after a MAC table timeout [43].

Passive eavesdropping is possible if an attacker can attach a listening device to a cable connecting a host to a switch or between two switches. Traffic between hosts can be captured this way. Equipment exists for passively tapping into electrical or optical cabling or a switch or multiport computer can be used. Passive eavesdropping is fairly difficult to detect.

If a switch does not know where to forward a frame, it floods it out of all of its ports. With software an attacker can easily generate enough frames with random addresses to overwrite an entire MAC table and make the switch flood all data frames to all ports for eavesdropping [39]. On most switches this *MAC flooding attack* affects all VLANs, even if the attack originates within one VLAN [16].

Fig. 6 shows how sending a frame with a forged sender address overwrites the correct entry in the MAC table and redirects traffic to the attacker. This *MAC spoofing* attack becomes more useful, if the real owner of the MAC address can be disabled or is known to be off-line. Otherwise a race condition exists between the two hosts and traffic will flip-flop between them. If the real host can be made to go off-line on demand, the spoofing host may not only receive traffic intended to the target host, but take over existing sessions of higher layer protocols [44].

Many switches have a port mirroring feature to support diagnostics or intrusion detection systems. If the attacker has control of a switch, this may be activated. Switches are shipped with standard or no passwords, and even if a password has been set, it is usually possible to reset the password if the switch can be accessed physically [41].

C. Traffic Integrity

The next step for an attacker is to modify traffic on the network. For example, an attacker can imitate a bank's web server to a user, and imitate the user to the bank's server, and

ARP spoofing




Fig. 7. IP address capture with a forged ARP message; A modifies T's ARP table.

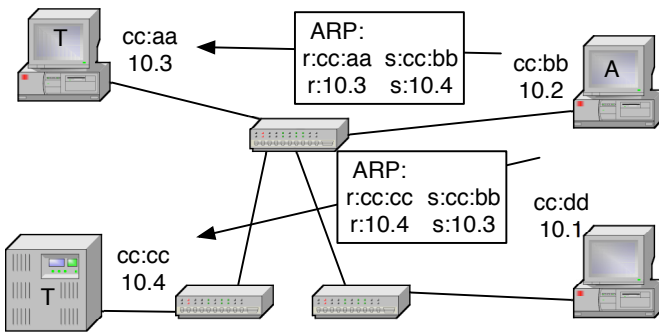


Fig. 8. By poisoning the ARP caches of both targets A can intercept all traffic between the two hosts.

gain temporary control of the user's bank account, as described in a how-to tutorial [45].

1) *ARP and DHCP Poisoning*: ARP is a stateless protocol and most operating systems will accept ARP replies even when not requested, as hosts tend to send these gratuitously whenever a link goes temporarily down. This enables a host to capture traffic intended for another host just by sending an ARP message to the sender with the intended receiver's IP address and the attacker's MAC address [46]. Fig. 7 shows an ARP message with forged sender IP address (MAC and IP addresses abbreviated to two octets).

In a similar way, a host can detect broadcast DHCP server requests and race the server to reply them first; upon success the attacker can assign a gateway (router) and DNS servers to the target host, along with its IP address, and control the host's traffic at will.

2) *Man in the Middle*: If an attacker can direct traffic to pass through his node and that traffic is not protected by an integrity verification mechanism, the attacker can easily modify the traffic. These Man in the Middle (MITM) attacks against higher layer protocols are performed relatively easily on an Ethernet segment. IP being the most common higher layer protocol on Ethernets, the previously mentioned ARP and DHCP poisoning attacks can be deployed to redirect traffic to go through the attacker's host for modification or just eavesdropping. Fig. 8 illustrates how a double ARP spoofing attack can redirect traffic between two hosts to transit via the attacker's machine.

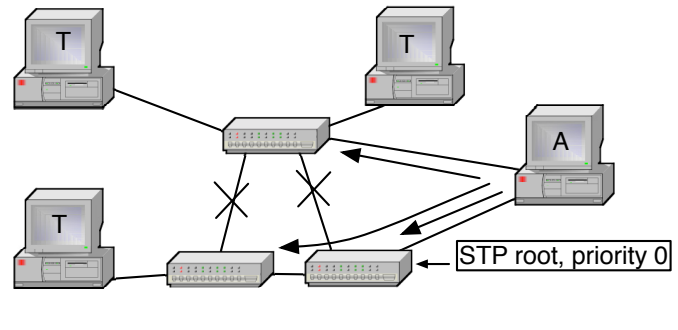


Fig. 9. STP root capture splits the network and leaves A in the middle for MITM.

On the Ethernet layer itself a MITM attack is harder, but can be done using STP. If a host is connected to two switches, it can act as the Root Bridge in the STP environment and create a tree topology, where part of the traffic goes through this host [47]. Shown in Fig. 9, a host sets its STP priority to highest and becomes the root, splitting the network in two and gaining access to traffic between the two halves. This attack requires gaining a connection to two switches. The attacker can also utilize two hosts connected to different switches and connect them to each other using out-of-band methods, like WLAN.

The router redundancy protocols (HSRP or VRRP) may also be used to masquerade as a router and gain access to traffic; however, these protocols are usually used in a dedicated segment.

3) *Session Hijacking*: Ethernet is a stateless protocol, but many higher level protocols create a session. Once a session is set up, it is often assumed to be trusted and no further traffic verification is made. If an attacker can eavesdrop on, or otherwise gain enough information about a session (IP addresses, TCP ports and sequence numbers, and application data, like an HTTP authentication cookie), the attacker can re-create the session and act like one endpoint [48], [49].

If one endpoint of the session can not be diverted, it might partake in the communications and disrupt the sessions. If the attack to be performed is speedy enough, this might not matter. Gratuitous ARP can be used to direct the local endpoint's traffic to a bogus MAC address and incoming traffic from the gateway router to the attacker's host. One endpoint can be silenced with a DoS attack. With the correct timing, a session may be brought up to date with the correct application messages or by trusting TCP to discard packets that appear to be duplicates based on the sequence number.

4) *Replay*: A message eavesdropped on earlier can be sent again. As the message is not modified, it can be authenticated or encrypted by the original sender without affecting the attack – the attacker just needs to guess at the content of the message to consider whether it is worth resending.

Within the Ethernet domain useful messages to resend would be small, stateless control messages that fit within one frame. Typical messages for targeting a resend attack could be routing notifications or SNMP "set" or "trap" messages.

D. Denial of Service

The attacker's motivation for DoS attacks is not to gain access to data but to prevent its use. The motivation may be revenge by a dissatisfied employee or blackmail by criminals. In a military environment, a well-timed DoS attack may have great significance, e.g., blocking the enemy's command and control.

A DoS attack can be performed in several ways. The attacks can cause total loss of service or degradation of service. An attack can be implemented on layer 1 by cutting links physically or by damaging the circuitry with electricity. Depending on the switch design this may affect more than one port, as one chip serves several ports. Making the entire switch inoperable is less likely but possible. However, layer 2 attacks can cause much more damage.

1) *Resource Exhaustion Attacks*: Resource exhaustion attacks can target the control and management planes of a switch by sending frames that require additional processing and handling. For example, if a switch is set to log certain types of frames, flooding the switch with such packets will likely overload the logging functionality. Status requests or VLAN configuration changes will stress the CPU of a switch. Overloading can slow traffic, block it completely, or stop certain functions, like MAC table updates. Multiple switches can be affected by the STP by sending and withholding STP root announcements and thus causing the spanning tree to fluctuate.

Unknown unicast flooding is a method for sending frames with a receiver address that does not exist in the network. As the CAM tables do not have this address, the frame is broadcast over all links (within the VLAN). This is in effect the same attack as MAC flooding but the intention is to congest the network and success depends on being able to cause sufficient traffic – while the goal in MAC flooding is to allow normal traffic but make the switch broadcast it. With resource exhaustion attacks an attacker's success depends on the ability to create a sufficient traffic volume and may depend on co-operation between several hosts or being connected to a high capacity link.

2) *Protocol Based DoS*: The STP that makes a tree out of a mesh network is designed to be self-configuring. An attacker that controls a node on the network can send STP messages and pretend to be a switch. The whole switching network can be brought to halt by flooding it with STP TCNs or other STP control messages [47].

E. Systems Security

Several threats are not tied to the architecture of Ethernet itself but to its implementation and use.

1) *Configuration and Installation Issues*: Besides the features of Ethernet technology, there is the practical issue of using the technology correctly and the level of skill and attention required to implement a secure solution.

Faulty, lacking, or incorrect configuration of the network switches can enable an attacker to get access to more of the network's resources than intended. This can especially contribute to a VLAN hopping attack (see Section III-A4). Vendors often ship their products with default settings, which

assume that every port can potentially be connected to another switch, thus listening to STP and other interswitch messages and processing VLAN tags.

On complex multipath networks the sheer complexity may overwhelm administration and create unforeseen consequences. Even when using network management tools, vulnerabilities are usually invisible by their nature and thus hard to notice.

2) *Implementation Issues and Vendor Extensions*: The standards leave room for implementation. Specifically, the control and the management plane processes are not defined in detail. An attacker can study a particular implementation and likely find unforeseen features, especially in the higher level areas of the switch. In a similar way, extensions to the existing architecture, such as access control lists at switch ports or service discovery protocols, may enhance or weaken security.

3) *Issues with Legacy Technology*: As Ethernet has been designed by providing incremental additions to existing technology, many Ethernet installations have equipment and software from various eras. This may make deploying modern solutions difficult.

Some security solutions, like IEEE 802.1X and 802.1AE described in the next section, require support from both the hosts and switches in the network that may not be available in the legacy environment. Operating with legacy equipment can thus leave holes in the security perimeter.

4) *Architectural Issues*: Ethernet is generally considered layer 2, but it is intimately involved with mechanisms such as ARP and DHCP that are not quite part of Ethernet, but relevant to security. These layer 3 base services assume that the layer 2 they are operating on is not hostile, i.e., the whole Ethernet segment is inside the protection domain.

The architectural paradigm is "fail open", i.e., it errs on the side of message delivery and ignores security issues. This is showcased by the basic switch design where a frame to an unknown address is flooded out to all ports. An alternative design would be to fail to closed position, i.e., not deliver anything unless the authenticity of the receiver is established.

Mixing user and control planes creates a fundamental problem: any frame may have control plane data and the switch must pick it out from the traffic. This offers attackers access to the control plane and, as on the control plane frames usually require more processing, possibility for exhaustion based DoS attacks.

5) *Freely Available Software for Attacks and Exploits*: Many of the attacks presented here are freely available in easy to use software. The incomplete list below illustrates the volume of the available tools:

- Network sniffers: Wireshark, Ngrep, Tcpdump, Snoopy
- Port scanners: Nessus, Nmap, Saint, Satan
- Packet crafting: packETH, Bit-Twist, Mausezahn, Hping, Nemesis, Scapy, Yersinia, THC Parasite, Ettercap, Macof
- Ettercap can also be used for MITM attacks
- Capture, edit, and replay: Packetsquare

IV. EXISTING SECURITY SOLUTIONS

The security of Ethernet has been improved by standardization organizations, equipment vendors, and the research

community. This section reviews the existing and proposed solutions and discusses the remaining security gap between the Ethernet and IP layers.

Traditionally, Ethernet's lack of security has been solved by defining any Ethernet segment as unsecure and requiring it to be placed inside a protected domain: behind a firewall in a secure building with trusted staff. Higher layer cryptographic solutions like IPsec and TLS are used to solve the remaining issues. Cryptography carries its own costs (mostly key management) and thus can not be considered a universal solution.

When looking for security in the Ethernet layer itself, it is clear that the switches form the core of the solution. A major problem is that a switch has no way of knowing if each of its ports is connected to: one computer (a host); a host with several virtual hosts (and virtual MAC addresses); a hub; a silent switch (that does not talk STP and other topology revealing protocols); a regular switch; or a switch that has other switches behind it. This dynamic ambiguity makes the issue challenging.

A. Router Based Security

We start by presenting how replacing one central Ethernet switch with an IP router would affect security. This provides, in our opinion, an useful baseline for evaluating Ethernet. The IP router partitions the rest of the Ethernet network into several segments. Each new segment is a separate broadcast domain. ARP, STP, VLAN, and MAC address table based attacks are no longer possible between the segments. Inside the segments the same attacks remain feasible, unless each switch is replaced with a multiport router.

The traffic between segments thus becomes impossible to eavesdrop on from other segments or to be redirected for a MITM attack. Ethernet's MAC headers are dropped at the router and traffic is guided by the IP addresses and router's IP table. The router blocks Ethernet's control plane protocols (ARP and STP) from transit between the segments. We assume a well configured router to set the bar a little bit higher; this means that DHCP attacks are valid only within one segment and that the router does not listen to routing protocol messages from the Ethernet segments.

Replacing a switch with a router will incur some costs. An IP router requires configuration, such as address allocation and default route configuration. Occasionally this can be automated, as is the case for residential broadband where the topology is clear to the access router.

The router also prohibits easy mobility. A host may move in the Ethernet network and keep its IP and MAC addresses, the MAC address tables in switches are updated automatically. IP routers can support mobility, but this usually requires additional protocols to manage the location of a host. A router also splits the broadcast domain. Autodiscovery protocols are blocked (unless the router includes application layer support for these) and thus services such as file sharing or printing in the other segments are no longer reachable.

Compared to an Ethernet switch an IP router provides a considerable amount of protection against other users connected to the same router. An attacker may target the protocols and

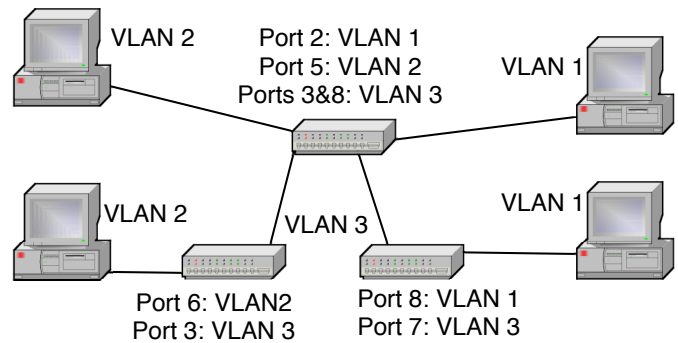


Fig. 10. VLAN segmentation.

implementations on higher layers (IP, transport, and application). DoS is also possible. However, in general, most of the attacks described in Section III become infeasible.

B. Access control

An attacker needs access before being able to perform any attacks. Untrusted entities can be kept out by limiting access to the network or requiring authentication. Limiting the access capabilities of trusted entities reduces the threat potential even further.

1) *Physical Protection of the Network*: Network equipment can be located in locked cabinets and racks and wiring installed inside walls to prevent unauthorized access. However, access is needed for the network to be useful and physical protection is of limited value.

2) *Segmentation and VLANs*: Limiting the size of an Ethernet segment limits the area vulnerable to attacks. A segmentation method external to Ethernet would be a higher layer device, such as a router or firewall. Inside Ethernet the IEEE 802.1Q virtual LAN mechanism provides a way to limit broadcasts and other traffic to specific segments. VLANs are logically separate within the same physical installation and define security domains inside one network. Fig. 10 shows how the switches are configured to assign VLANs 1 and 2 to specific ports and use VLAN 3 as a trunk. Hosts on VLAN 1 are not able to reach hosts on VLAN 2 on layer 2.

Vendors recommend using VLANs for security. However, VLAN based security depends on proper switch configuration and vendor documentations also note that the default settings of switches are not secure, thus enabling, e.g., VLAN hopping [27].

3) *Individual VLANs*: Each host on the Ethernet can also be placed into its own VLAN, using IEEE 802.1ad Q-in-Q double tagging or vendor provided Private VLAN (PVLAN) switch configuration. This is useful in networks where hosts communicate to only one or a few other nodes, which is typical for Ethernet-based access networks.

Q-in-Q is mainly a specification for extending the VLAN 14 bit identifier space by adding another VLAN tag. The PVLAN technique uses switch configuration to isolate hosts and only let their traffic pass to one "promiscuous" port, typically connected to a router and to the Internet [50]. Each host sees only itself and the host(s) connected to the promiscuous port (or a chain of such ports) and only a few VLAN IDs are needed at the trunks to indicate PVLAN traffic.




Fig. 11. A 802.1X authentication session; the host presents its credentials to the switch, which verifies them from a central server.

4) *Authentication Based Access Control*: Being able to authenticate the user or host connecting to a port at the switch is a step forward from plain physical access control.

IEEE 802.1X port authentication supports several types of authentication credentials, such as a user-name and password pair, or a certificate and corresponding private key [51]. 802.1X requires supporting client software in the end host, software in the switch, and a centralized authentication database server. The host communicates with the switch and the switch verifies the authentication from the database, as shown in Fig. 11. 802.1X uses Extensible Authentication Protocol (EAP) that has broad support for different types of authentication methods and structures (cryptographic exchanges related to certificates are more complex than just supplying a user-name and password) [52], [53].

802.1X authenticates a host at the beginning of a session, attaching the MAC address to a specific port in the switch. If the port senses a link disconnection (electrically), the association should be dissolved and a new authentication required. When hosts are connected directly to an 802.1X capable switch, it provides protection from MAC spoofing and flooding attacks; however, ARP poisoning and other attacks remain possible.

An attacker may place a hub or switch between the authenticated host and the authenticating switch. After authentication the authenticated host can be disconnected without losing the electrical connection to the authenticating switch and another host, configured with same MAC address, be used in the network [54].

Authentication can also be used between switches to form a trusted inner network. This can be used to prevent attacks where a host acts as a switch.

Besides 802.1X, the individual LAN techniques mentioned previously (Q-in-Q and PVLAN) can be used to prevent any traffic from a host besides the initial (DHCP) configuration and traffic to an authentication server, until the host has authenticated itself using some higher layer service, such as Web using HyperText Transfer Protocol (HTTP). Switches may also contain a higher layer authentication functionality (e.g., HTTP) that performs functions similar to 802.1X but without requiring the host to have 802.1X client software [55]. These types of services may have their own authentication databases or use the existing 802.1X structure.

5) *Access Control Lists*: Access Control Lists (ACLs) are not part of the Ethernet specification; switch vendors have added various types of capabilities by themselves. The Ethernet frame does not have many features: for a simple Ethernet frame ACL the usable attributes are the sender's or receiver's MAC address or the Ethertype field. Access can be limited based on MAC addresses, but several service specific ACLs are commonly implemented.

Port security lets the administrator limit access to a port in a switch, based on the number of MAC addresses [55], [56]. This blocks MAC flooding and can make it more difficult to expand the network by adding switches without authorization. Typically the functionality has detailed control features. Besides just blocking new MAC addresses when their number exceeds a limit, port security may also be set to block a port from existing MAC addresses. Old MAC addresses may expire after a period, or be static until the port is reset manually. MAC addresses may also be attached to the port where they are first connected and the same address will be blocked at other ports, preventing mobility and MAC address spoofing.

Packet storm protection limits the amount of frames per time unit a switch will allow from a port. This will prevent MAC flooding attacks if the limit is low enough but is more commonly used to guard against packet storms, where a host sends large numbers of frames, intentionally or because of a malfunction.

A *BPDU guard* blocks all STP messages from a port and can be used to designate a port that will not form a part of the mesh network. An *STP root guard* indicates a port which is part of the STP network, but can not become the STP root. These features are also used for performance reasons, for example to make sure that only the fastest links are used to form the tree topology of the network [56].

6) *Control and Management Plane Overload Protection*: Control and management planes can be protected from overload by limiting the amount of traffic on these planes. Control Plane Policing (CoPP or CPP) achieves this by providing a set of filters, based on rate limiting methods and addresses to prevent the overloading of control plane functionalities. The filters are configured to allow only a certain amount of control and management plane data packets to reach the CPU – everything else is blocked before reaching the CPU level. This protects against intentional CPU exhaustion attacks, with the drawback that during an attack legitimate messages are also lost. The filters can usually be set separately for different types of control frames.

7) *Centrally Managed LAN Security*: Several approaches to collecting information from a LAN to a central point and using this to manage security have been presented by the research community in recent years.

SANE is a clean slate design where hosts publish services and require access to each other from a centralized controller [57]. SANE focuses on implementing “natural” (organization centric) security policies on the LAN level and attempts to provide comprehensive protection for Ethernet.

Ethane is a continuation of SANE that drops the requirement for new software at the hosts [58]. Policies are still held in the central controller and hosts must be authenticated.

Ethane switches direct all new flows initially to the controller that decides whether to allow or deny them and configures the switch to act accordingly.

OpenFlow follows the ideas of Ethane and presents a design for a switch that replaces the MAC table with *flow tables* and an external controller [59]. The flow tables are based on several attributes of the packet, including Ethernet, IP, and transport layers. Frames can be directed to ports or the central controller, thus ARP requests can be separated from other traffic. OpenFlow represents an architecture called *Software Defined Networking* (SDN) that separates the control function from the switching function and usually implements control only at the beginning of flow. The difference from SANE and Ethane is that OpenFlow does not define any particular purpose for itself, but can be used for performance, security, routing algorithm testing or other purposes. However, security is one of the research directions for OpenFlow technologies [60]–[62].

C. Secure Protocols

Access controls limit the availability of targets to attackers. The targets can also be made harder to reach by adding security features to protocols and protocol implementations

1) *Encryption and Integrity Verification*: Cryptography can solve integrity and confidentiality requirements. IEEE 802.1AE MACsec forms encrypted connections between hosts and switches, protecting confidentiality and integrity of the content in the frames [63]. Deployment requires software installation and configuring authentication for each participating network entity. MACsec uses 802.1X authentication information as its basis, but leaves several issues such as key management outside its specification, leaving it to vendors to implement them [64].

MACsec provides protection against intruders to the network, preventing reading and modification of data frames. However, authorized hosts may misbehave. MACsec defines one perimeter of protection and inner entities are not protected against each other. For example an authorized host may use ARP to gain control of a host's traffic [65]. DoS and traffic analysis attacks remain a possibility.

2) *Securing Address Resolution Protocol*: ARP creates a major vulnerability in the Ethernet architecture. Information gained by *DHCP snooping* can be used to prevent ARP spoofing attacks, by tying MAC addresses to their corresponding IP addresses and ports, based on information gained from DHCP messages. DHCP snooping suffers from a potential lack of scope, as a single switch can not see the allocations made to hosts whose path to the DHCP server does not pass through this switch [66].

The research community has mostly focused on cryptography based solutions, such as S-ARP [67], which adds an authentication field to ARP messages and provides a corresponding key management structure, [64] that uses cryptographic name space binding, or [68], which extends MACsec's reach to endpoint to endpoint and multicast protection.

3) *Control and Management Plane Logical Protection*: Protecting the higher functions in a switch from misuse relies on controlling the access to the switch. Control plane functions

(e.g., MAC learning or STP) have to be connected to the user plane and can be protected as described previously.

The simplest protection is to curtail management plane functions to a separate physical or virtual management network. Besides this, encrypted connections are used to protect management data (usually SSH for command line and TLS or SSL for WWW interfaces). Authentication can be based on passwords or cryptographic credentials.

Simple Network Management Protocol (SNMP) is usually used to monitor switches and may or may not use password protection. Version 1 of SNMP does not support encryption, and when available for version 3 and certain varieties of version 2, it may not be used as long as SNMP is used just for monitoring.

4) *Replay Protection*: The basic Ethernet frame has no protection against replay attacks, leaving it to higher layers. MACsec and many higher layer protocols include timestamps or nonces (non-repeating values) to thwart replay.

D. Security Monitoring

Previous sections describe security techniques that are mostly proactive and, once set, do not require active participation from external systems or human interaction. Active technologies enhance the protection of the network.

1) *Ethernet Firewall and Deep Packet Inspection*: Firewalls are used to limit traffic between network segments and can be considered more complex cases of access control lists, including stateful features. Firewalls can also employ deep packet inspection (DPI) and application layer session recreation for inspection purposes. Current firewall products can operate on all network layers and thus the concept of an "Ethernet firewall" lacks separate meaning. The switches' ACLs can be used to limit traffic on the Ethernet layer and standard firewall products can control the higher layers.

DPI means analyzing the contents of the packet at the application level, beyond the headers (just looking at the headers is occasionally referred to as "shallow packet inspection"). DPI is out of the scope of this paper, except for those protocols that are relevant to Ethernet, like ARP and DHCP, where DPI can be used to protect the control and management planes (this is in effect a type of ACL).

2) *Intrusion Detection and Prevention Systems*: Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) systems use DPI to identify network attacks, usually from a signature library of known attacks. They require access to the network traffic that can be gained by placing an IDS/IPS device directly between two endpoints (typical when used as a firewall or to enhance a firewall) or they can monitor traffic from a switch via the *port mirroring* feature. Port mirroring copies traffic to and from selected ports to a listening port, where the monitoring device is located. If needed, a separate network can be used to keep the monitoring devices separate from the protected network.

Some other features in the switches can help in detecting malicious behavior, like *MAC address notification* which can send an SNMP trap message when a host moves in the network. Several SNMP Management Information Base (MIB) definitions can be useful for IDS purposes, like the Remote

Network Monitoring MIB (RMON) [69], [70] and its switch extensions (SMON) [71].

Besides passive monitoring, active measures can be used to detect malicious behavior. Frames with known expected behavior can be injected into the network and the results monitored to detect ARP spoofing attacks [72], [73].

3) *Planning, Configuration, and Administration*: Good network administration practices can considerably influence the aspects of an Ethernet network. Several of the technical solutions mentioned previously require configuration and constant adjustment as the network topology changes. As there is no reliable mechanism to automatically separate a trunk network from connections to the leaf nodes, administrators must configure this information into the switches. Separating management information to a dedicated VLAN and limiting control plane functionality and data flows enhance the security of a network to a level that can be close to an IP router based network.

Vendors have seen the security issues related to Ethernet and have provided their own guidelines for configuring an Ethernet based network [66], [74]. This paper has already covered the recommendations in these guidelines.

Various network management systems exist to ease the workload of configuring the switches on a network. These managers maintain the topology picture of the network and reduce mistakes by automating tasks. However, they require that switches are compatible with the management software and are configured to work together.

Network administration duties can also include active network scanning, probing, and testing to detect vulnerabilities.

E. Discussion

It appears that the level of security increases with the efforts of administration and that there is no simple technological way to add self-configuring transparent security to the Ethernet layer.

Plain, self-configured, out-of-the-box Ethernet is clearly not secure against any threats. MAC flooding, ARP spoofing, and STP attacks are easy to perform.

An Ethernet network configured according to vendor's instructions fares better. ARP spoofing is still possible (unless the switches feature DHCP snooping and ARP inspection). Switches still leak frames at MAC address table time outs. However, the key problem is that any mistake by network administrators will compromise the security. Each software update at switches or configuration change must be analyzed against the security policy. The attacker can poll and probe the network's defenses freely until a vulnerable point is located.

Adding 802.1X or other authentication does not secure the network by itself. It just limits the potential attackers to known users, while adding to the workload of the administration. Additional software may be needed at the hosts and the management of authentication information adds a considerable workload.

MACsec or other frame encryption mechanisms solve eavesdropping issues, including frames broadcast at MAC table time out, and negates MITM attacks. DoS and traffic analysis are still possible, as are attacks through the network

to higher layers. The workload is roughly the same as using authentication without MACsec.

Intrusion detection and prevention systems can detect several attacks. Some attacks, such as VLAN double tagging, are easily identifiable. DHCP snooping pairs MAC addresses to IP addresses and thwarts ARP spoofing.

V. RESEARCH DIRECTIONS

We mentioned in the introduction that Ethernet is a growing technology family [4] of many uses [2]. Security is needed in many areas. Cloud computing means that data centers will be larger and have an increasing number of independent and potentially hostile tenants. Operators are running Ethernet edge-to-edge and inter-operator Ethernet segments is not an impossible thought. Stuxnet brought out the need for protecting legacy networks, especially as many industrial Supervisory Control And Data Acquisition (SCADA) and automation networks contain old hardware and software that can not be easily upgraded.

A. New and Existing Areas of Use

Multiple security domains within one Ethernet segment have been traditionally solved with VLAN technology. As shown, VLAN is not a very secure separator, unless configured carefully and without mistakes. VLAN also suffers from the 14 bit identifier space, partially relieved by Q-in-Q double tagging. There is a need for a method that provides the robust separation of security domains in large Ethernet installations. This method should not affect performance of the switching fabric and should allow mobility for the virtual hosts.

Legacy automation systems often use Ethernet for supervisory control and monitoring, occasionally even for real-time needs. The lifetime of such systems is often measured in decades and the legacy equipment might have been designed before security was a consideration in networking. Legacy equipment can also be sensitive to all kinds of modern features, such as large data frames or traffic volumes. These systems are often connected directly or indirectly to the Internet and require protection from attacks that can penetrate the first layer of defense.

B. Architectural issues

Ethernet is self-configuring and the mechanisms that support this cause it to not be very secure. Increasing security while maintaining the self-configuring nature is a worthy goal. The key issue is that data and control planes are mixed and that the authenticity of the participants is not verified. Several concrete paths for research are presented by recent activities.

1) *Software Defined Networking*: SDN and OpenFlow have potential to unify the switches in what can be considered one large, flow-based switch. This allows implementing a centrally controlled LAN, as SANE and Ethane have demonstrated. This, in turn, could be used to analyze the control messages coming from each host and to compare them to central knowledge. Having the information from ARP and DHCP snooping, or replacing these mechanisms with the controller, should solve most of the security issues presented in this paper.

2) *Removing Broadcasts*: If the broadcasts could be removed from Ethernet, its security would improve considerably. Distributed Hash Tables (DHT) have been proposed as a replacement solution for locating a host [75], [76]. Another solution is re-engineering the control plane [77], [78]. The motivation for removing broadcasts is usually to extend the size of the Ethernet segment, while avoiding moving to the IP layer and thus maintaining zero-configurability. However, removing ARP or all broadcasts from the Ethernet layer would have implications for security too.

3) *Cryptographically Generated Addresses*: RFC 3972 Cryptographically Generated Addresses (CGA) uses the rightmost 64 bits of the IPv6 address to store a hash of the host's public key [79]. Ethernet addresses with 46 or 47 bits of significance could be created by hashing from the public keys of a host. These addresses would be compatible with legacy equipment, but other hosts could verify the identity of an endpoint when needed. Participating equipment could use its public/private key pair to sign control layer frames, thus enabling switches to monitor the identities of hosts.

47 bits might not be a long enough hash space to protect against a brute force attack but it might be long enough for most practical purposes.

Potentially this system could be made auto-configuring, by using the "leap of faith" used in e.g., SSH protocol, where the server is not authenticated when connected the first time [80] or the "resurrecting duckling" concept, where a device stores the credentials of the first device it encounters, imprinting itself strongly like a duckling to its mother [81].

C. New Vulnerabilities and Threats

Have the vulnerabilities of current LAN Ethernet been analyzed thoroughly? The literature surveyed so far indicates that findings have been mostly found by a random process and reported as individual cases. Unrecorded weaknesses might exist in the current architecture and in the implementation of technologies.

New technologies such as TRILL and SDN very likely include new vulnerabilities. Even if the TRILL problem statement [82] states "TRILL solutions should not introduce new vulnerabilities compared to traditional bridged subnets." the focus of TRILL work is on path efficiency and concentration, not security. SDNs have the potential to be very complex systems and thus have more room for vulnerabilities.

VI. SUMMARY

The major strengths of Ethernet are its simplicity and zero-configurability. These are also the cause for its weaknesses, the dynamic features that allow it to self-configure can be misused.

Ethernet can be secured to a reasonably high level by administering all switches, hosts, and users centrally and applying cryptographic methods. However, this means the loss of simplicity and zero-configurability. Also existing solutions are fairly granular and do not protect very well against misuse from authorized users.

A reasonable level of protection can be reached by administering the switches and maintaining separation of the switching

core and leaf nodes. This could almost be considered the equivalent to the protection provided by replacing switches with IP routers, except for ARP broadcasts which leave the system vulnerable to misuse.

Several research approaches provide the potential for enhancing the security of Ethernet. Centralized management can be streamlined and made more granular. However, even if adding end nodes is allowed, the core network will be in hands of the administration and adding any switches requires interaction with central management.

Another potential solution is to select a source of trust and leverage this to authenticate and authorize known entities. This could require changes to end nodes and protocols but could be done without human intervention and would save the zero-configurability aspect, while losing some of the simplicity.

A third option would be to remove the ARP broadcasts, which would solve a major security issue while maintaining the desirable aspects of Ethernet.

ACKNOWLEDGMENTS

This work was performed within the ECEWA project, funded by the Finnish National Technology Agency Tekes and industry partners. The authors would like to thank Markus Peuhkuri and Nuutti Varis for their comments.

REFERENCES

- [1] N. Falliere, L. Murchu, and E. Chien, "W32.Stuxnet dossier." Symantec Security Response, 2010, retrieved Oct 24, 2011 Online: http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.
- [2] J. Sommer, S. Gunreben, F. Feller, M. Kohn, A. Mifdaoui, D. Sass, and J. Scharf, "Ethernet - a survey on its fields of application," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 2, pp. 263–284, 2010.
- [3] R. Metcalfe and D. Boggs, "Ethernet: Distributed packet switching for local computer networks," *Communications of the ACM*, vol. 19, no. 7, pp. 395–404, July 1976.
- [4] R. Sofia, "A survey of advanced ethernet forwarding approaches," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 1, pp. 92–115, 2009.
- [5] G. King, "A survey of commercially available secure LAN products," in *Computer Security Applications Conference, 1989., Fifth Annual, 1989*, pp. 239–247.
- [6] R. Housley, "Encapsulation security protocol design for local area networks," in *Local Area Network Security*, ser. Lecture Notes in Computer Science, T. Berson and T. Beth, Eds. Springer Berlin Heidelberg, 1989, vol. 396, ch. chapter 10, pp. 103–109.
- [7] F. Poon and M. Iqbal, "Design of a physical layer security mechanism for CSMA/CD networks," *Communications, Speech and Vision, IEE Proceedings I*, vol. 139, no. 1, pp. 103–112, Feb 1992.
- [8] M. Soriano, J. Forné, F. Recacha, and J. L. Melús, "A particular solution to provide secure communications in an Ethernet environment," in *CCS '93: Proc. 1st ACM conference on Computer and communications security*. New York, NY, USA: ACM Press, 1993, pp. 17–25.
- [9] M. El-Hadidi, N. Hegazi, and H. Aslan, "Implementation of a hybrid encryption scheme for Ethernet," in *Computers and Communications, 1995. Proceedings., IEEE Symposium on*. IEEE Comput. Soc. Press, Jul 1995, pp. 150–156.
- [10] M. Bishop, *Introduction to Computer Security*. Pearson Education, 2005.
- [11] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Wiley Publishing, 2008.
- [12] W. Stallings and L. Brown, *Computer Security: Principles and Practice*, 2nd ed. Pearson Education, 2012.
- [13] J. Pieprzyk, T. Hardjono, and J. Seberry, *Fundamentals of Computer Security*. Springer-Verlag, 2003.
- [14] A. Singh, B. Singh, and H. Joseph, *Vulnerability Analysis and Defence for the Internet*, ser. Advances in Information Security. Springer US, 2008, vol. 37.
- [15] S. McClure, J. Scambray, and G. Kurtz, *Hacking Exposed*, 10th ed. McGraw-Hill, 2009.

- [16] E. Vyncke and C. Paggen, *Lan switch security: what hackers know about your switches*. Cisco Press, 2007.
- [17] G. Schudel and D. J. Smith, *Router Security Strategies: Securing IP Network Traffic Planes*. Cisco Press, 2008.
- [18] J. P. Walters, Z. Liang, W. Shi, and V. Chaudhary, "Wireless sensor network security: A survey," in *Security in Distributed, Grid, and Pervasive Computing*. CRC Press: Boca Raton, FL, USA, 2007.
- [19] D. J. Welch and S. D. Lathrop, "A survey of 802.11a wireless security threats and security mechanisms," *United States Military Academy, West Point*, 2003.
- [20] *Media access control (MAC) Bridges*, IEEE Std. 802.1D, 2004.
- [21] R. Droms, "Dynamic Host Configuration Protocol," RFC 2131 (Draft Standard), Internet Engineering Task Force, Mar. 1997, updated by RFCs 3396, 4361, 5494.
- [22] D. Plummer, "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware," RFC 826 (Standard), Internet Engineering Task Force, Nov. 1982, updated by RFCs 5227, 5494.
- [23] T. Narten, E. Nordmark, W. Simpson, and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)," RFC 4861 (Draft Standard), Internet Engineering Task Force, Sep. 2007, updated by RFC 5942.
- [24] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)," RFC 3315 (Proposed Standard), Internet Engineering Task Force, Jul. 2003, updated by RFCs 4361, 5494, 6221, 6422.
- [25] E. Mannie and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control," RFC 4606 (Proposed Standard), Internet Engineering Task Force, Aug. 2006, updated by RFC 6344.
- [26] H. Holbrook, B. Cain, and B. Haberman, "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast," RFC 4604 (Proposed Standard), Internet Engineering Task Force, Aug. 2006.
- [27] "VLAN security white paper," 2002. [Online]. Available: http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml
- [28] *Virtual Bridged Local Area Networks*, IEEE Std. 802.1Q, 2005.
- [29] *Virtual Bridged Local Area Networks - Amendment 4: Provider Bridges*, IEEE Std. 802.1ad, 2005.
- [30] *Multiple Registration Protocol*, IEEE Std. 802.1AK, 2007.
- [31] T. Li, B. Cole, P. Morton, and D. Li, "Cisco Hot Standby Router Protocol (HSRP)," RFC 2281 (Informational), Internet Engineering Task Force, Mar. 1998.
- [32] S. Nadas, "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6," RFC 5798 (Proposed Standard), Internet Engineering Task Force, Mar. 2010.
- [33] *Station and Media Access Control Connectivity Discover*, IEEE Std. 802.1AB, 2009.
- [34] *Link Aggregation*, IEEE Std. 802.1AX, 2008.
- [35] K. D. Mitnick, *The Art of Intrusion*. Wiley, 2005, page 127.
- [36] S. Stasiukonis, "Social engineering, the USB way," June 2006, retrieved Jan 17, 2011. [Online]. Available: <http://www.darkreading.com/security/article/208803634/index.html>
- [37] T. Koponen, S. Shenker, H. Balakrishnan, N. Feamster, I. Ganichev, A. Ghodsi, P. B. Godfrey, N. McKeown, G. Parulkar, B. Raghavan, J. Rexford, S. Arianfar, and D. Kuptsov, "Architecting for innovation," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 3, pp. 24–36, Jul 2011.
- [38] "Global security report 2011," white paper, Trustwave, 2011. [Online]. Available: https://www.trustwave.com/downloads/Trustwave_WP_Global_Security_Report_2011.pdf
- [39] S. Convery, "Hacking layer 2: Fun with Ethernet switches," Black Hat USA, August 2002, retrieved Feb 15, 2011. [Online]. Available: <http://www.blackhat.com/presentations/bh-usa-02/bh-us-02-convery-switches.pdf>
- [40] A. A. Vladimirov, "Making unidirectional VLAN and PVLAN jumping bidirectional," Full-disclosure mailing list, Dec 2005, retrieved May 15, 2011. [Online]. Available: <http://lists.grok.org.uk/pipermail/full-disclosure/2005-December/040333.html>
- [41] "Catalyst 3550 multilayer switch software configuration guide, 12.1(19)ea1," Manual, Cisco Inc, Oct 2003.
- [42] "VoIP hopper software," Online, 2009, retrieved May 15, 2011. [Online]. Available: <http://voiphopper.sourceforge.net/>
- [43] O. Arkin and J. Anderson, "Etherleak: Ethernet frame padding information leakage," Online: http://www.rootsecure.net/content/downloads/pdf/atstake_etherleak_report.pdf, 2003, retrieved Feb 3, 2011.
- [44] R. Spangler, "Packet sniffing on layer 2 switched local area networks," December 2003, retrieved Feb 10, 2011. [Online]. Available: <http://packetwatch.net/documents/papers/layer2sniffing.pdf>
- [45] A. Coffman, "Simple DNS man in the middle attack for password phishing," October 2011, retrieved Oct 24, 2011. [Online]. Available: <http://thecoffman.com/2011/10/20/simple-dns-man-in-the-middle-attack-for-password-phishing/>
- [46] C. Abad and R. Bonilla, "An analysis on the schemes for detecting and preventing ARP cache poisoning attacks," in *Distributed Computing Systems Workshops, 2007. ICDCSW '07. 27th International Conference on*. IEEE, Jun 2007, p. 60.
- [47] M. Marro, "Attacks at the data link layer," Master's thesis, University of California Davis, 2003. [Online]. Available: http://seclab.cs.ucdavis.edu/papers/Marro_masters_thesis.pdf
- [48] E. Norris, "Analysis of a telnet session hijack via spoofed MAC addresses and session resynchronization," Mar 2001, retrieved Feb 10, 2011. [Online]. Available: <http://www.scribd.com/doc/18397013/Analysis-of-a-Telnet-Session-Hijack-via-Spoofed-MAC-Addresses>
- [49] O. Zheng, J. Poon, and K. Beznosov, "Application-based TCP hijacking," in *Proc. Second European Workshop on System Security*, ser. EUROSEC '09. New York, NY, USA: ACM Press, 2009, pp. 9–15.
- [50] S. HomChaudhuri and M. Foschiano, "Cisco Systems' Private VLANs: Scalable Security in a Multi-Client Environment," RFC 5517 (Informational), Internet Engineering Task Force, Feb. 2010.
- [51] "Port-Based Network Access Control," IEEE Std. 802.1X, 2010.
- [52] B. Aboba, L. Blunk, J. Vollbrecht, J. Carlson, and H. Levkowitz, "Extensible Authentication Protocol (EAP)," RFC 3748 (Proposed Standard), Internet Engineering Task Force, Jun. 2004, updated by RFC 5247.
- [53] B. Aboba, D. Simon, and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework," RFC 5247 (Proposed Standard), Internet Engineering Task Force, Aug. 2008.
- [54] S. Riley, "Mitigating the threats of rogue machines 802.1X or IPsec?" Microsoft TechNet, August 2005, retrieved Oct 11, 2011. [Online]. Available: <http://technet.microsoft.com/en-us/library/cc512611.aspx>
- [55] "ProCurve series 6120 switches access security guide," Manual, Hewlett-Packard Development Company, 2009.
- [56] I. Dubrawsky, "Safe layer 2 security in-depth," white paper, Cisco Inc, 2004. [Online]. Available: http://www.cisco.com/warp/public/cc/sc/cuso/epso/sqfr/sfblu_wp.pdf
- [57] M. Casado, T. Garfinkel, A. Akella, M. Freedman, D. Boneh, N. McKeown, and S. Shenker, "SANE: A protection architecture for enterprise networks," in *USENIX Security Symposium*, 2006.
- [58] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. Gude, N. McKeown, and S. Shenker, "Rethinking enterprise network control," *IEEE/ACM Trans. Netw.*, vol. 17, no. 4, pp. 1270–1283, Aug 2009.
- [59] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow, enabling innovation in campus networks," *SIGCOMM Comp Commun Rev.*, vol. 38, no. 2, pp. 69–74, Mar 2008.
- [60] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, and S. Shenker, "NOX, towards an operating system for networks," *SIGCOMM Comp Commun Rev.*, vol. 38, no. 3, pp. 105–110, Jul 2008.
- [61] R. Clark, N. Feamster, A. Nayak, and A. Reimers, "Pushing enterprise security down the network stack," Georgia Tech, Tech. Rep., 2009. [Online]. Available: <http://hdl.handle.net/1853/30782>
- [62] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker, "Onix: a distributed control platform for large-scale production networks," in *Proc. 9th USENIX conference on Operating systems design and implementation*, ser. OSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–6.
- [63] "Media Access Control (MAC) Security," IEEE Std. 802.1AE, 2006.
- [64] H. Altunbasak and H. Owen, "An architectural framework for data link layer security with security inter-layering," in *SoutheastCon, 2007. Proceedings. IEEE*. IEEE, Mar 2007, pp. 607–614.
- [65] H. Altunbasak, S. Krasser, H. Owen, J. Grimminger, H. Huth, and J. Sokol, "Securing layer 2 in local area networks," *Networking-ICN 2005*, vol. 3421, pp. 699–706, 2005.
- [66] "Hardening ProCurve switches," white paper, Hewlett-Packard Development Company, 2007.
- [67] D. Bruschi, A. Ornaghi, and E. Rosti, "S-ARP: a secure address resolution protocol," in *Computer Security Applications Conference, 2003. Proceedings. 19th Annual*. IEEE, Dec 2003, pp. 66–74.
- [68] K. Wahid, "Rethinking the link security approach to manage large scale Ethernet network," in *Local and Metropolitan Area Networks (LANMAN), 2010 17th IEEE Workshop on*, 5-7 2010.

- [69] S. Waldbusser, "Remote Network Monitoring Management Information Base," RFC 2819 (Standard), Internet Engineering Task Force, May 2000.
- [70] —, "Remote Network Monitoring Management Information Base Version 2," RFC 4502 (Draft Standard), Internet Engineering Task Force, May 2006.
- [71] R. Waterman, B. Lahaye, D. Romascanu, and S. Waldbusser, "Remote Network Monitoring MIB Extensions for Switched Networks Version 1.0," RFC 2613 (Draft Standard), Internet Engineering Task Force, Jun. 1999.
- [72] V. Ramachandran and S. Nandi, "Detecting ARP spoofing: An active technique," in *Information Systems Security*, ser. Lecture Notes in Computer Science, S. Jajodia and C. Mazumdar, Eds. Springer Berlin Heidelberg, 2005, vol. 3803, ch. chapter 18, pp. 239–250, 10.1007/11593980_18.
- [73] N. Hubballi, S. Roopa, R. Ratti, F. A. Barbhuiya, S. Biswas, A. Sur, S. Nandi, and V. Ramachandran, "An active intrusion detection system for LAN specific attacks," in *Proc. 2010 international conference on Advances in computer science and information technology*, ser. AST/UCMA/ISA/ACN'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 129–142.
- [74] S. Convery and B. Trudel, "SAFE: A security blueprint for enterprise networks, white paper," white paper, Cisco Inc, 2000. [Online]. Available: http://www.cisco.com/en/US/prod/collateral/wireless/wirelessw/ps1953/product_implementation_design_guide09186a00800a3016.pdf
- [75] C. Kim, M. Caesar, and J. Rexford, "Floodless in SEATTLE, a scalable Ethernet architecture for large enterprises," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 3–14, Oct 2008.
- [76] N. Varis and J. Manner, "Minimizing ARP broadcasting in TRILL," in *GLOBECOM Workshops, 2009 IEEE*. IEEE, Nov 2009, pp. 1–6.
- [77] J. Rexford, A. Greenberg, G. Hjalmytsson, D. Maltz, A. Myers, G. Xie, J. Zhan, and H. Zhang, "Network-wide decision making: Toward a wafer-thin control plane," in *Proc. HotNets*. ACM Sigcomm, 2004, pp. 59–64.
- [78] A. Myers, E. Ng, and H. Zhang, "Rethinking the service model: Scaling Ethernet to a million nodes," in *Proc. HotNets*. ACM Sigcomm, 2004.
- [79] T. Aura, "Cryptographically Generated Addresses (CGA)," RFC 3972 (Proposed Standard), Internet Engineering Task Force, Mar. 2005, updated by RFCs 4581, 4982.
- [80] T. Ylonen and C. Lonvick, "The Secure Shell (SSH) Protocol Architecture," RFC 4251 (Proposed Standard), Internet Engineering Task Force, Jan. 2006.
- [81] F. Stajano and R. Anderson, "The resurrecting duckling: Security issues for ad-hoc wireless networks," in *Security Protocols*, ser. Lecture Notes in Computer Science, J. Malcolm, B. Christianson, B. Crispo, and M. Roe, Eds. Springer Berlin / Heidelberg, 1999, vol. 1796, pp. 172–182, 10.1007/10720107_24.
- [82] J. Touch and R. Perlman, "Transparent Interconnection of Lots of Links (TRILL): Problem and Applicability Statement," RFC 5556 (Informational), Internet Engineering Task Force, May 2009.



Timo Kiravuo Timo Kiravuo (born 1965) is a postgraduate student in Aalto University. He received his MSc. (1999) from Helsinki University of Technology. After a career in private sector his work focuses on Internet security and related matters. Currently he is researching the security of Ethernet technologies.



Mikko Särelä Mikko Särelä works as a post-doctoral researcher at Aalto University. Prior to joining Aalto University, he worked at Nomadyclab, Ericsson on information centric network architecture and future Internet. His current research interests include ethernet scalability and security, energy efficiency, and mobility.



Jukka Manner Jukka Manner (born 1972) received his MSc. (1999) and PhD. (2004) degrees in computer science from the University of Helsinki. He is a full professor (tenured) of networking technology at Aalto University, Department of Communications and Networking (Comnet). His research and teaching focuses on distributed systems and various networking aspects, most recently on the development of the Internet and its services, particularly in topics related to energy efficient ICT, networking beyond IP, multipath connectivity and transport protocols.

He is the Academic Coordinator for the Finnish Future Internet research programme. He is an active peer reviewer and member of various TPCs. He has contributed to standardization of Internet technologies in the IETF for over 10 years, and was the co-chair of the NSIS working group. He has been principal investigator and project manager for over 15 national and international research projects. He has authored over 80 publications, including several IETF RFCs. He is a member of the ACM and the IEEE.