



NETWORK SECURITY
KNOWLEDGE AREA
(DRAFT FOR COMMENT)

AUTHOR: Sanjay Jha – University of New South Wales

EDITOR: Andrew Martin – Oxford University

REVIEWERS:

Gene Tsudik – University of California

Other reviews in progress

KA Infrastructure Security: Network Security

Sanjay Jha

12th October 2018

INTRODUCTION

The ubiquity of the Internet allows us to connect all sorts of devices to the network and gain unprecedented access to a whole range of applications and services anytime, anywhere. However, our heavy reliance on networking technology also makes it an attractive target for malicious users willing to compromise security of our communications and/or cause disruption to services that are critical for our day-to-day survival in a connected world. In this chapter, we will endeavour to explain the challenges associated with securing a network under a variety of attacks for a number of networking technologies and widely used security protocols, along with emerging security challenges and solutions. An understanding of basic networking protocol stack and TCP/IP suite is assumed. Basic networking text books explain the fundamentals of the 7-layer ISO OSI model and Internet Protocol [1]. Applications require some or all of the five key standard security features, namely, authentication, confidentiality, integrity, availability, and non-repudiation. As we examine the security of the Internet, we will discover how the original protocols were not designed with security in mind and that additional protocols and mechanisms have been developed to meet the security requirements.

CONTENT

1 Network Security

1.1 Internet Architecture

A complex system such as distributed applications running over a range of networking technologies is best understood when viewed as layered architecture. Figure 1 shows the 7-layer protocol ISO OSI stack and the interaction between the various layers. The model also allows us to understand the security issues at each layer and the interplay between them. As the Internet is the predominant architecture today, we will use the five-layer protocol stack in figure 1 in this chapter. The Presentation and Session layers shown in the dotted box are remain optional in the IP protocol stack and some or all of the functions can be custom built on application requirements. Network security requires cryptographic techniques such as public and symmetric keys for encryption and signing, block and stream ciphers, hashing and digital signature, as described in the Cryptography Knowledge Area ?? . We will take an applied approach to understand how these techniques help build a secure network.

1.2 Network Protocols and Vulnerability

Typically, Dolev-Yao, [2] adversarial formal model is used for a formal analysis of security protocols in the research literature. The Dolev-Yao model assumes that an adversary has complete control over the entire network, and concurrent executions of the protocol between the same set of 2-or-more parties can take place. The Dolev-Yao model describes the worst possible adversary: depending on the context, real adversaries may have limited capabilities.

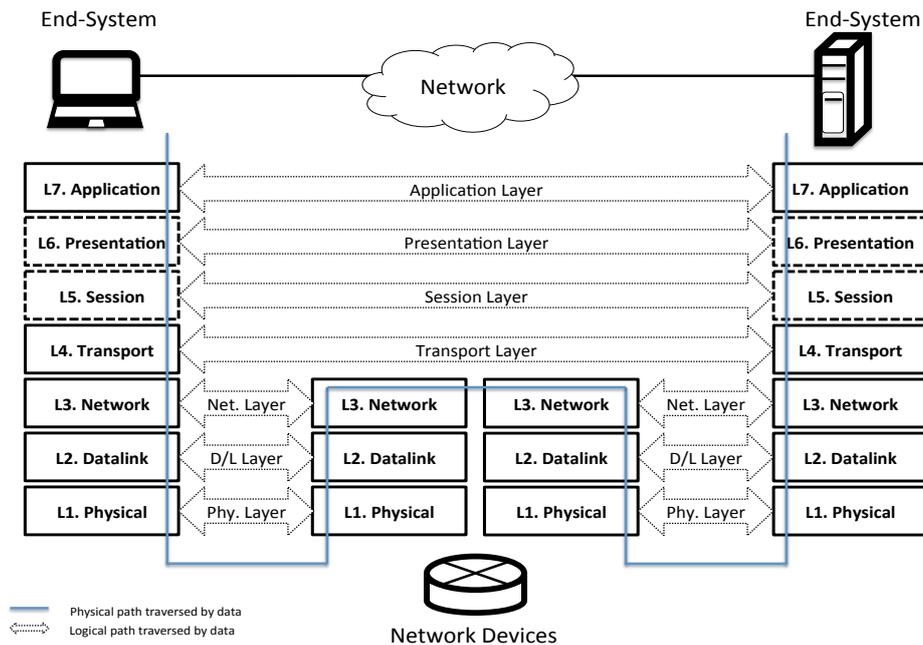


Figure 1: 7 Layer Protocol Stack

The popular characters Alice and Bob from the security literature are willing to exchange messages securely. Web servers and clients, two email clients, two people using video-conferencing are examples of real world Alices and Bobs. The hackers, an eavesdropper called Eve, and a malicious attacker called Mallory are waiting to compromise their communications. We examine a few common network security attacks. Messages sent by Alice and Bob over a network can be captured by Eve using packet sniffing tools. This allows Eve to inspect each packet and possibly extract confidential information such as passwords, credit card details or many other types of sensitive information. Broadcast networking technologies such as wireless LAN or cable modem make it relatively easy to sniff packets. Man in the middle attack (MiTM) is another common security threat where Mallory, an attacker, places itself between Alice and Bob. For example, a compromised gateway/router/access-point, malware present in the user's device, or a server can potentially capture all of the packets being exchanged between two parties, add/modify/delete information and carry out other malicious activities. Denial of Service (DoS) attack is a technique where an attacker sends an avalanche of bogus packets to a server. This would either keep the server constantly busy or clog the access link resulting in disruption of service to legitimate users. Typically, a large number of compromised hosts (bots) are used to launch a distributed DoS attack aka DDoS. DoS and MiTM are not disjoint, many DoS attacks are also MiTM, and vice-versa. Mirai [3] is an example of a malware, first found in 2016 which launched a DDoS attack by compromising Linux-based consumer devices, aka Internet of Things (IoT) devices, such as IP cameras, utility meters, home routers and others. The IoT devices were turned into bots by exploiting the weak authentication configurations including use of default passwords. The bots were then used from a command and control centre to attack several high-profile websites. The use of IoT devices allowed the attackers to circumnavigate traditional security measures.

In an IP Spoofing attack, an attacker tries to impersonate as an authorised user by crafting a packet with a forged IP address and adjusting certain other fields to make it look legitimate. Having looked at examples of network attacks, we will now examine the security at each layer of the protocol stack.

1.3 Application Layer Security

Alice and Bob wish to use email. In a simplistic scenario, Alice and Bob would decide to use an encryption technique such as AES ?? with a 128 or 256 bit key to encrypt their messages. This meets their requirement of confidentiality as the message cannot be decrypted by anyone other than Alice and Bob. However, this would require Alice and Bob to agree on a shared key. Distributing this key over the network makes the secret key an easy target for Eve or Mallory. Also, the above scenario fails to provide integrity and origin authentication. The message can be altered as it traverses the network. Alice and Bob (in this instance, their email clients) must use additional measures to provide message integrity and origin authentication. In a variant of this setting, it is also likely that Alice and Bob do not care about the confidentiality of their messages, but they want assurance that their messages will not be tampered with in transit. Alice could calculate the hash of her message using the SHA-3 algorithm ?? and send it to Bob. On receiving this message, Bob would recalculate the hash and verify whether there is a match. However, a potential attacker could easily replace the genuine message with a forged one and a matching hash. Bob cannot tell whether the message sent by Alice has been altered since the hash matches. One possible solution for Alice is to use a pre-negotiated symmetric key to encrypt the hash. Bob now decrypts this hash using the pre-negotiated symmetric key and verifies the integrity of message received. This also authenticates that the message was sent by someone who shares a key with Bob, in this instance Alice.

We highlighted the challenges of key distribution over the network. See the Cryptography KA for details of public key cryptography ?. We will ignore the confidentiality requirement for the moment. Alice signs the hash of her message using her private key. Bob then decrypts the message using Alice's public key. This allows for an integrity check and authentication at the same time, as no one but Alice knows her private key. We avoided pre-negotiation or sharing of keys. So, how does Bob get Alice's public key and trust that Eve or Mallory are not using a forged public/private key to perform MITM. We provide a brief introduction to key management in the context of public key cryptography in the next section as it is used by a number of network security protocols.

1.3.1 Public key infrastructure

Public key infrastructure (PKI) provides a solution for registering and managing a trustworthy public key. Government agencies or standard organisations appoint registrars who issue keys, and keep track of the public certificates of entities (individuals, servers, routers etc). The registrars, a large number of which are private companies, themselves have a registered public/private key pair with the relevant authorities. The idea is similar to registering your motor number plate with an authority. Alice generates a pair of public/private keys for herself using her computer. She then presents her identity proof to one of the registrars. The registrar then issues a certificate to Alice. This certificate is signed by the registrar's private key and can be verified by anyone using the registrar's public key. Typically, a user's identity, public-key and CA information is used as an input to hash function. The hash is then signed with the CA's private key to produce a Public Key Certificate (PKC). The fields of a certificate include a unique identifier/serial number, a signature algorithm used by the CA and the period of validity. The IETF RFC1422 and ITU-X.509 standards have prescribed the format and standard for managing PKI [4]. Organisations can also manage their own private PKI. CAs also publish a list of revoked certificates which have either expired or been revoked. The web of trust is an alternative scheme where users can create a community of trusted parties by mutually signing certificates without the needing a registrar. Continuing with our email example, Alice could send her certificate to Bob along with her email message. Bob is now able to check the validity of the certificate presented by Alice. In our trivial example, Alice and Bob could use these techniques to build a secure email system. Pretty Good Privacy (PGP) was one of the early email systems to propose the security approach described above, albeit using the web of trust for certificates. Generally, in order for systems to be compatible across platforms and between vendors, application developers make use of the standard application layer protocol, the Simple Mail Transfer Protocol (SMTP) for

exchanging messages between mail servers. The content itself is formatted based on a set of standards called Multipurpose Internet Mail Extensions (MIME). As the original Internet protocols lacked security features, a secure version SMIME was developed for the addition of an integrity check and certificates in the email header. The functions of certificate verification and the checking revocation list are automatically performed by Alice's and Bob's mail agents.

1.3.2 DNS Security Extensions

The Internet design philosophy mandates keeping the Internet core functions implemented in the backbone routers to be simple along with other supporting functions to be deployed at the edge. For most people, human cognition means that it is easier to remember host/server names, e.g. cnn.com, over an IP address 151.101.1.XX. Internet routing and other protocols, however, function using IP addresses. The IETF has designed an application layer protocol, the Domain Name Service (DNS), which performs the translation between a host name and the corresponding IP address. This mapping is performed and maintained by a hierarchy of name servers. There have been a number of DDoS attacks in recent years [5]. In a MiTM, Mallory can impersonate a DNS server, return a bogus address and divert traffic to a malicious server, thus allowing it to collect user passwords or other credentials. A DNS poisoning attack allows attackers to plant bogus addresses, thus diverting a user request to malicious servers. However, the robust distributed design of the DNS has fortunately saved us from total collapse of the Internet. Learning from these attacks, the IETF introduced a secure version called DNS Security Extensions (DNSSEC). DNSSEC uses techniques similar to our secure email example above by sending a response signed by the private key of a DNS server. The authenticity of the DNS records is proved by the fact that a responding server signs the record using its private key, which a requester can verify using the corresponding public key. In addition, a digital signature also provides the integrity of the response data. An astute reader may note that confidentiality is not a significant issue for this transaction. More than half a dozen IETF RFCs cover DNSSEC. DDoS defence is not part of DNSSEC. We will look at defence mechanisms in IDS/IPS in section 1.8.

1.3.3 Hyper Text Transfer Protocol Secure (HTTPS)

The most prominent application layer protocol, the Hyper Text Transfer Protocol (HTTP), was designed without any security considerations. The popularity of HTTP and its wide adoption for e-commerce imposed strict security requirements on the protocol. A Secure version called HTTPS was introduced by using security services from the transport layer which allows one to encrypt URL, content, forms and cookies during communication. We discuss the secure transport layer protocols in the next section. A new version HTTP 2.0 has further enriched the security features of HTTP1.0. Although not mandated, most browsers support confidentiality by encrypting data. New features like header compression and flow control require servers to maintain additional state information. An attacker could send a large number of empty or tiny frames and keep the server busy processing frame headers. Servers must employ a threshold on the number of connections being processed to limit such attacks.

1.4 Transport Layer Security

In the previous section, we saw how applications could build security features by using cryptographic primitives. Data sent over the TCP/IP protocol were not safe and hence each application had to take care of security itself. Ideally, if the transport layer could provide confidentiality, integrity and authentication mechanisms, the application layer could be relieved from the burden of security and use the transport layer services instead. This would also provide compatibility across platforms/vendors. Researchers designed a shim layer between the application and transport layers called the Secure Socket Layer (SSL). A standard application programming interface (API), similar to the socket API, allows applications bootstrap secure connections and to send/receive data securely. IETF started to develop the Transport Layer Security (TLS) borrowing most of the ideas from the SSL 3.0 protocol.

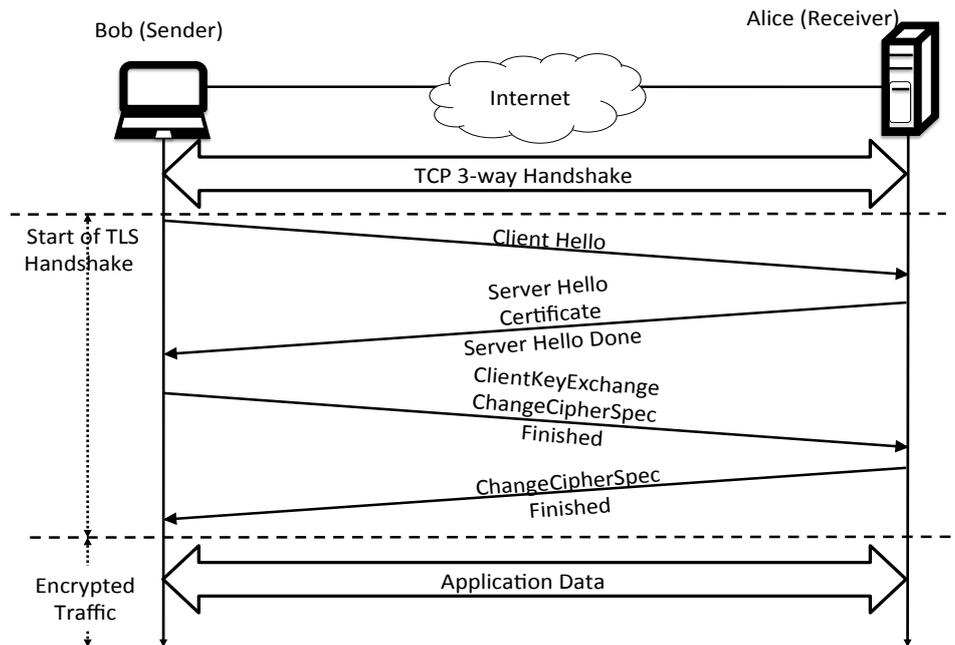


Figure 2: TLS Handshake

Most prominent web browsers have started to support the latest TLS 1.3 standardised in 2018. In this section our discussions will relate to a simplified version of the security features in order to understand the basics of the TLS protocol. The exact syntax and semantics of the protocols and a rich set of configurations are described in hundreds of pages of RFCs.

We will now bring Alice (server) and Bob (client) back into the action. Alice is configured with her public/private key pairs as described in 1.3.1. It is worth emphasising that some of these basic techniques are also used in security protocols at other layers that we will discuss in this chapter. The TLS protocol has 3 phases: Handshake, key-derivation and data transfer, as shown in figure 2.

1.4.1 Handshake

1. First Bob and Alice exchange the three-way TCP SYN, SYNACK and ACK messages. It should be noted that this step is not part of TLS/SSL.
2. Bob then sends a ClientHello message to Alice along with the cipher suites (ciphers and the hash functions it supports) and a NONCE.
3. Alice responds with a ServerHello message along with her choice from the cipher suites (e.g. AES for confidentiality, RSA for public key, SHA2 for MAC), a Certificate containing her public key and a NONCE. Additionally, she could also request for the client's certificate and parameters for other TLS extensions
4. Bob checks validity of the certificate and is assured that it belongs to Alice. He initiates the ClientKeyExchange message. It can use a range of key exchange methods e.g. RSA or the Diffie-Hellman (and variants) to establish a symmetric key for the ensuing session. For example, using RSA, Bob could generate a 48-bit Pre-master secre (PMS) and encrypt it with Alice's public key obtained using the steps above and send it to Alice.

5. Bob sends a ClientCipherSpec and a Finished Message suggesting that the key generation and authentication are complete.
6. Alice also has the shared key at this point. She responds with ChangeCipherSpec and Finished message back to Bob.
7. Bob decrypts the message with the negotiated symmetric key, verifies the MAC.

After successfully completing the above steps, a secure tunnel is established and the encrypted application data can now be sent as shown at the bottom of figure 2. Details of the protocol exchange and message processing can be found in [4].

1.4.2 Key-Derivation

The client NONCE, server NONCE, and PMS are input into a pseudorandom function to produce a master secret. All the other key data for this connection are derived from this master secret in conjunction with the additional parameters. Following four common keys are derived at both ends:

1. Session encryption key for data sent from Bob to Alice (client encryption key).
2. Session encryption key for data sent from Alice to Bob (server encryption key).
3. Session MAC key for data sent from Bob to Alice (client MAC key).
4. Session MAC key for data sent from Alice to Bob (server MAC key).

Bob and Alice derive separate keys for encryption and integrity in each direction for enhanced security. Generating these ephemeral keys allows for perfect forward secrecy, as these keys cannot be reused in future sessions. For example, Eve could capture every communication between Alice and Bob. She could pretend to be Bob and repeat the sequence of commands sent by Bob later in the day. This attack is called a *connection replay* attack. The TLS and most other protocols use a session specific NONCE, a random number, to avoid this attack. The PMS generation algorithm uses NONCE in the mix. The connection replay attack will fail, as Alice would have a different set of keys from Eve for the new session due to this new NONCE.

1.4.3 Data-Transfer

TCP is a byte oriented transport protocol where Application layer data are sent as a stream of bytes. Integrity check algorithms require fixed length data for MAC calculation. If applications have to collect and pass fixed length data to these algorithms, further delay will be incurred. Hence, TLS defines a record format, as shown in figure 3, where the length of the data sent in each record can be indicated along with the type of record (data or control). A MAC is also appended at the end of each record. For example, if data are sent from Bob to Alice, the session MAC key for the data sent from Bob to Alice are used to generate this MAC. Further, the data plus the MAC are encrypted using the session encryption key for data sent from Bob to Alice.

As the TCP sequence number is not encrypted, a possible MiTM attack could simply capture the TCP segments and swap the TLS records between these segments. A receiver would not be able to detect this attack as the integrity of the TLS records remains unchanged. TLS provides a separate mechanism where the sender and receiver keep track of the record sequence number without explicitly exchanging it. However, the MAC calculations at both ends use this sequence number in the mix. Any MiTM rearrangement of records will fail the integrity check.

Having discussed the technical details of the TLS, we will now consider a few attacks and see how TLS provides security for users and applications. In a Password Sniffing attack, Eve captures a few

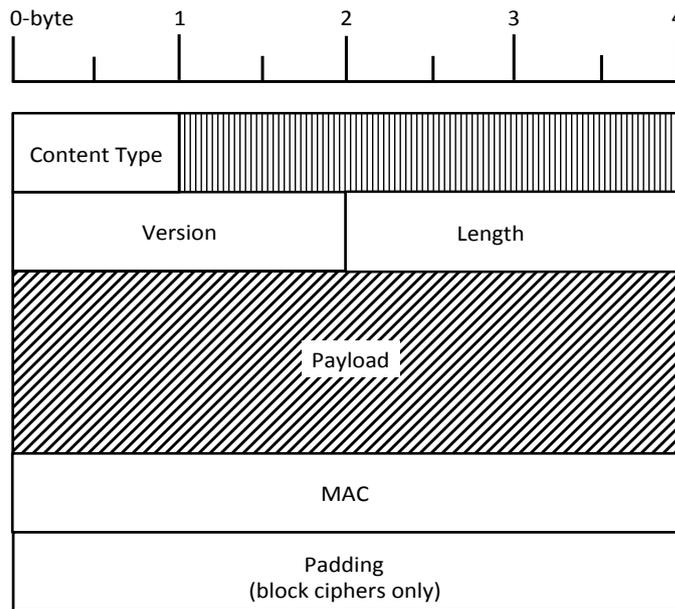


Figure 3: TLS Record Structure

packets and wishes to get passwords in HTTPS or other application traffic. As the user data are encrypted, the password can not be sniffed. In the IP Spoofing attack, Mallory uses a forged IP addresses to fool Bob into accepting bogus data. Mallory must be in possession of the secret key as well as the forged IP address to succeed. A MiTM attack is prevented by using public key certificates to authenticate the correspondents.

We note that in a related transport layer attack called a *SYN Flooding DDoS attack*, a group of attacking machines keeps sending TCP SYN messages to request a connection and let the server allocate resources. However, this type of attack can be handled by TCP and hence is not duplicated in TLS. A defence known as SYN Cookies has been implemented in many operating systems [RFC4987]. The server does not half open a connection right away upon receiving a TCP connection request. It selects an initial sequence number (ISN) using a hash function over source and destination IP addresses, port numbers of the SYN segment, as well as a secret number only known to the server. The server then sends the client this ISN aka Cookie in the SYNACK message. If the request is from a legitimate sender, the server receives an ACK message with a new sequence number which is ISN plus 1. Once this validation is done, the server opens a TCP connection. A DDoS sender would either not respond with ACK or would not have correct ISN in its response. Hence no TCP resource has been wasted.

The current version of SSL (and TLS) has evolved through experiencing several attacks and vulnerabilities found in earlier versions. *SSL Stripping attacks* remove the use of SSL/TLS altogether by modifying unencrypted protocols which request the use of TLS. The *BEAST attack* exploits the predictable initialisation vector of TLS 1.0 implementation due to use of Cipher Block Chaining (CBC). This allows an attacker to decrypt parts of a packet, e.g. HTTP cookies. A long list of known attacks and mitigation were discussed in RFC 7457. Many of these vulnerabilities are also attributed to either an improper implementation or poor understanding of the protocol suite rather than a lack of proper specifications.

1.5 Network Layer Security

The reader would be curious as to why we need more security if application and transport layers protect us. First, higher-layer security mechanisms do not necessarily protect an organisation's internal network links from malicious traffic. If and when malicious traffic is detected at the end-hosts, it is too late as the bandwidth has already been consumed. The second major issue is that the higher-layer security mechanisms described earlier (e.g. TLS) do not conceal IP headers. This makes the IP addresses of the communicating end-hosts visible to eavesdroppers.

Additionally, many organisations prefer their traffic to be fully encrypted as it leaves their network. In the early days of networking, several private networks were in use. However, maintaining a private network may not be cost effective. An alternative solution is to make use of the Internet to connect several islands of private networks owned by an organisation. Additionally, employers and employees want a flexible work environment where people can work from home, or connect from a hotel room or an airport lounge without compromising their security. We have already determined that the Internet is unsafe. The concept of a virtual private network (VPN) over the public Internet requires a set of network layer security mechanisms that we will explore in this section. We start our discussion with security additions to the network layer IP protocol called IPSec. Figure 4 shows that an employee working from home accesses a server at work, the VPN client in their host encapsulates IPv4 datagrams into IPSec and encrypts IPv4 payload containing TCP or UDP segments, or other control messages. The corporate gateway detects the IPSec datagram, decrypts it and decapsulates it back to the IPv4 datagram before forwarding it to the server. Every response from the server is also encrypted by the gateway as well. We note that encryption is not mandatory in IPSec. Figure 4 is one of several modes of operation for IPSec. For example, there could be two corporate networks, each with their own IPSec gateway communicating over the open Internet.

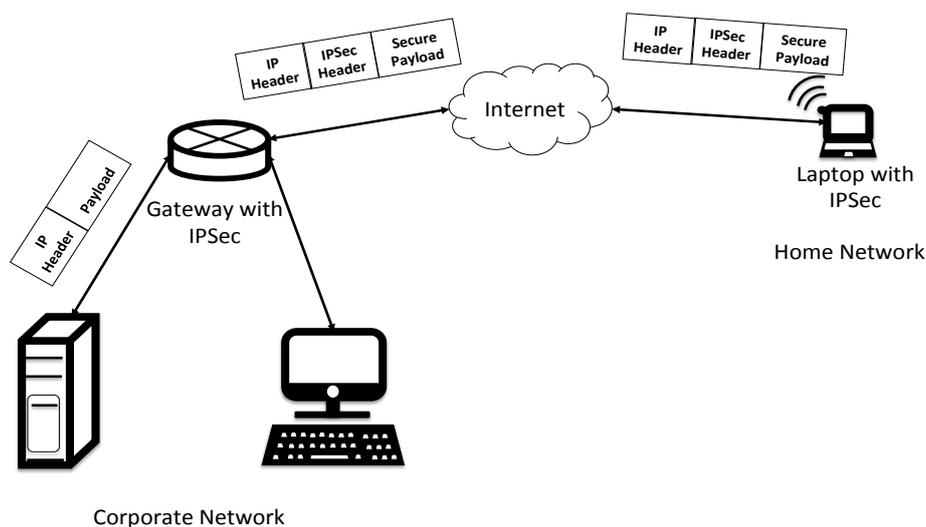


Figure 4: Example IPSec Client Server Interaction

We started off with a simple example showing data confidentiality using encryption. However, IPSec also provides data integrity, origin authentication and replay attack prevention. Again, the set of modes/configurations/standards provided by IPSec is extensive, interested readers should access

the relevant IETF RFCs.

IPSec supports Tunneling and Transport modes of operation. In Transport mode as shown in figure 5, the original IP header is used but the rest of the payload gets encrypted. In the rest of this section, we will discuss the widely used alternate Tunneling mode in detail. If the edge devices (routers/gateways) of two networks are IPSec aware, the rest of the servers/hosts need not worry about IPSec. The edge devices perform the encapsulation of every IP including the header. This virtually creates a secure tunnel between the two edge devices. The receiving edge device then decapsulates the IPv4 datagram and forwards within its network using standard IP forwarding. Other possible configurations for a tunnel could involve one IPSec aware host and an IPSec aware gateway (as in figure 4). A tunnel between two IPSec aware hosts is also possible without involving edge routers. The tunneling mode remains in widespread use due to its simplicity as it does not require IPSec protocol support in the end hosts. Also, key negotiation is simplified as two edge devices can handle connections on behalf of multiple hosts in their respective networks. An additional advantage is that everything, including, IP source/destination address, gets encrypted, thus making traffic analysis harder.

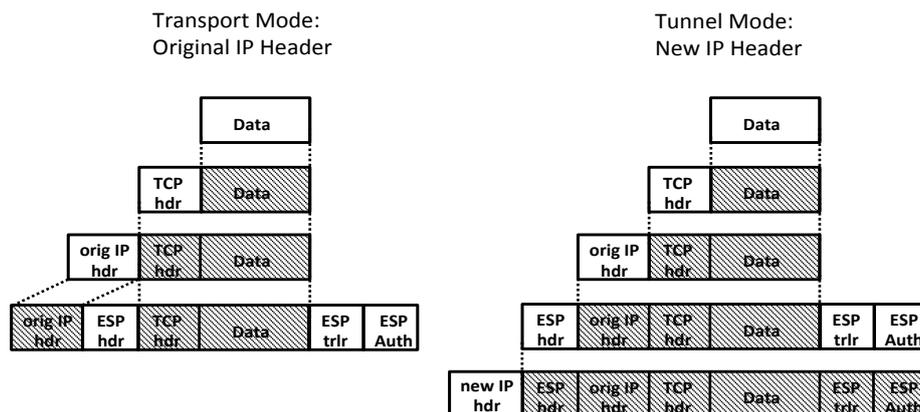


Figure 5: Transport and Tunnel Mode Encapsulation

IPSec supports a set of formats. The Encapsulation Security Protocol (ESP) supports confidentiality, data integrity and source authentication. If an application does not require confidentiality, it may simply use the Authentication Header (AH) protocol, which supports data integrity and source authentication. In total, we get four different options for communication: Transport mode with ESP, Transport mode with AH, Tunnel mode with ESP and Tunnel Mode with AH. Since VPN tunnels are fully encrypted, the Tunnel mode with ESP remains the protocol of choice.

Two entities participating in IPSec communication establish *security association (SA)* for each direction of the link. Essentially a number of variables are recorded in a database called the Security Association Database (SAD) for lookup during IPSec protocol processing, somewhat similar to the TCP connection state. Some of the state information includes the type of encryption used (e.g. AES or 3DES), the encryption key, the type of integrity check used (e.g., SHA-2 or MD5), the authentication key and so on. An Anti-Replay Window is also used to determine whether an inbound AH or ESP

packet is a replay.

When a large number of end-points use IPsec, distributing the keys becomes challenging. The RFC5996 defines the Internet Key Exchange protocol (IKE). Readers will observe a similarity between TLS 1.4 and IKE, in that IKE requires each entity to be in possession of a public key so that they can exchange certificates during the handshake process and negotiate the encryption as well as the authentication algorithms. We will skip the details of a complex two-phase protocol exchange which results in the establishment of a master key. These master keys are used to generate the keys used during a session, as we recall IPsec SAs. The Internet Security Association and Key Management Protocol (ISAKMP) is a framework that defines the procedures for authenticating the communicating peer, creation and management of Security Associations, and the key generation techniques. Additionally, it can also provide threat mitigation against DoS and replay attack. Defined in RFC 2408, ISAKMP is widely used by IKE for key exchange.

1.5.1 BGP Security

So far, we have primarily focussed on the security of data being sent using the TCP/IP protocol suite. However, a network can easily be disrupted if either the routers themselves are compromised or they accept spurious routing exchange messages from malicious actors. In recent years, we have seen YouTube service disruption globally. The entire Internet had an outage in another country due to either mis-configuration or the malicious advertising of bogus BGP updates. Either way, it highlights the security weakness in the BGP protocol. The vulnerability arises because of a lack of integrity and authentication for BGP messages.

The Internet uses a hierarchical system where each autonomous system (AS) managed by an ISP, exchanges routing information with other ASes using the Border Gateway Protocol (BGP). Readers are encouraged to read the details of the BGP protocol in RFC1163 and RFC1267. After receiving an IP prefix, the reachability information, from its neighbour, the router checks the newly received information against its stored knowledge to see if there is a better path to reach a destination network. This information is updated locally and propagated to its immediate neighbours. The distributed system allows networks to reach each other globally. In what is known as a *BGP route hijacking attack*, a malicious router could advertise an IP Prefix saying that the best route to a service is through its network. Once the traffic starts to flow through its network, it would then choose to drop all the packets to this service for a variety of reasons, including censorship. It could also read all of the un-encrypted packets. Additionally, the attacker could divert traffic through an unsuspecting AS, suddenly increasing their load. In a *BGP denial-of-service (DoS) attack*, a malicious router would send an avalanche of BGP traffic to a victim AS, keeping its border router busy so that it could not process any valid update. The attacker could also propagate spurious BGP updates and corrupt routing tables so as to prevent traffic from reaching its intended destination.

IETF is currently working on a standard called BGPsec to address these security concerns. The work is based on a prior proposal called S-BGP [6]. The core of the scheme lies in the use of PKI to verify the signatures of the neighbours sending the updates. Two neighbouring routers could use IPsec mechanisms for point-point security to exchange updates. We will now look at a simple example where a BGP router receives a path ZZZ YYY XXX. The BGP router verifies the signature of AS XXX using PKI mechanisms that we have learned earlier. It then verifies the signature generated by YYY and subsequently for ZZZ. This allows us to verify the origin and authenticity of the whole chain of updates. However, this approach entails large overheads. Signature verification comes at a cost. With some of the core routers keeping a global routing table of more than half a million IPv4 prefixes and an average of four AS hops, a router would need to verify two million signatures on booting. Additional crypto hardware and memory would certainly help keep performance on track.

1.6 Link Layer Security

In this section, we will discuss security measures that can be implemented at the link layer of the protocol stack. We will start our discussion with the prominent 802.1X Port-based Authentication followed by link layer security issues in Ethernet Switched LAN and Wireless LAN environments.

1.6.1 IEEE 802.1X Port-based Authentication

The IEEE 802.1X is a port-based authentication for securing both wired and wireless networks. Before a user can access a network at the link layer, it must authenticate the switch or access point (AP) they are attempting to connect to, either physically or via a wireless channel. As with most standards bodies, this group has its own jargons. Figure 6 shows a typical 802.1X setup. A user is called a Supplicant, a switch or AP is called an Authenticator. The architecture requires an Authentication Server (AS) that can be implemented using one of the existing protocols: Remote access dial-in user service (RADIUS), DIAMETER, Kerberos, Lightweight Directory Access Protocol (LDAP) or Active Directory (AD) and others. The AS function can also be co-located with the Authenticator. We will now consider RADIUS as our example AS. Typically, the AS and authenticator are pre-configured with a shared secret. Using the RADIUS protocol as an example, once a supplicant request is received, the authenticator sends a RADIUS Access Request message to the RADIUS server, requesting authorisation to be granted to access the resources.

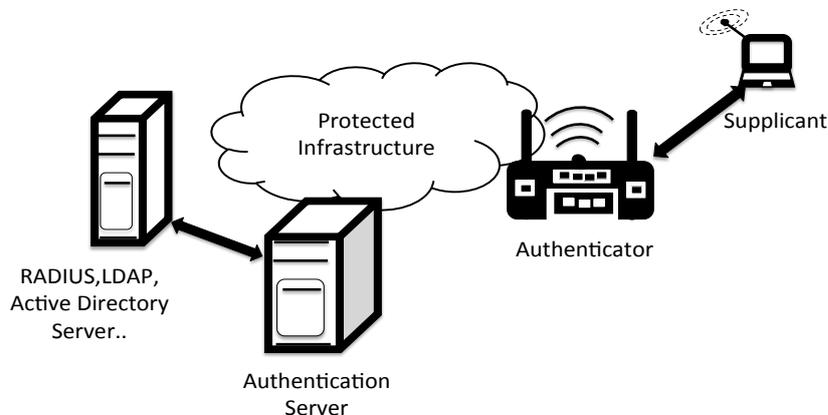


Figure 6: 802.1X Port-Based Authentication Architecture

Supplicant software is typically available on various OS platforms or can also be provided by chip-set vendors. A supplicant (client) wishing to access a network must use the Extensible Authentication Protocol (EAP) to connect to the AS via an authenticator.

1.6.2 Extensible Authentication Protocol (EAP)

The EAP is an end-end client to authentication server protocol. From supplicant to authenticator, it is sent over Layer2 protocols, i.e., EAP over LAN (EAPOL). There is no need for higher layer protocols.

Since the authenticator is connected to the AS using a trusted link with a shared secret, a higher layer protocol such as RADIUS/DIAMETER over UDP can be used on this side of the link.

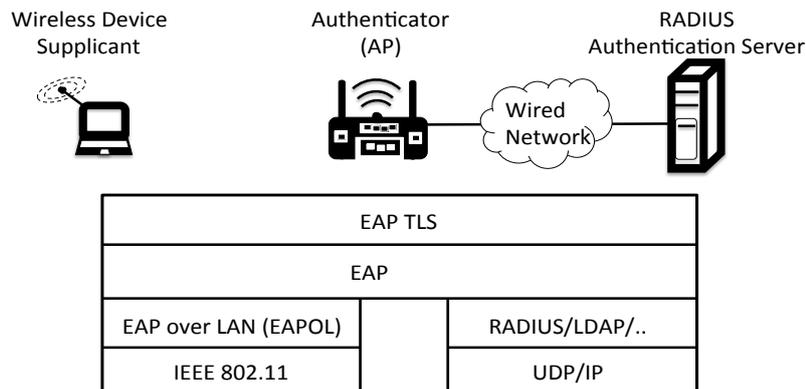


Figure 7: Extensible Authentication Protocol (EAP)

When a new client (supplicant) is connected to an authenticator, the port on the authenticator is set to the 'unauthorised' state, allowing only 802.1X traffic. Other higher layer traffic, such as TCP/UDP is blocked. The authenticator sends out the EAP-Request identity to the supplicant. The supplicant responds with the EAP-response packet, which is forwarded to the AS. This typically includes the supplicant's credentials (username and hash of password). Upon verification, the AS returns one of the following responses: Access Accept, Access Reject, Access Challenge for extra credentials. If the result is Access Accept, the authenticator unblocks the port to let higher layer traffic through. When the supplicant logs off, the EAP-logoff to the authenticator sets the port to block all non-EAP traffic.

Sending the supplicant's credentials in plaintext opens a can of worms. To safeguard against any eavesdropping, the EAP uses a Tunnel for authentication and authorisation. A whole range of EAP Tunneling protocols are available. EAP-Transport Layer Security (EAP-TLS), EAP for GSM Subscriber Identity (EAP-SIM) and EAP Protected Authentication Protocol (EAP-PEAP) are some of the examples for establishing a secure tunnel. EAP-PEAP, also known as 'EAP inside EAP' is one of the most popular protocols. If we dig deeper, before the port is unblocked, a complex process is used to generate a dynamic encryption key using a 4-way handshake. Essentially, all of these protocols establish a TLS tunnel but differ in choice of hash algorithms, the type of credentials used, whether a client-side certificate is used etc. Most protocols would use a server side certificate.

Once the supplicant and AS mutually authenticate, they together generate a Master Key (MK). As we have already discussed, the authenticator plays the role of a relay up to this point. During this process, the supplicant derives a Pairwise Master Key (PMK). The AS also derives the same PMK and sends this to the authenticator. From this point, the supplicant and authenticator use the PMK to derive the Temporal Key (TK) used for the message encryption and integrity. The key derivation process is similar to what we learnt in the TLS earlier. We will revisit key generation and the relationship

between various keys in detail later in the Robust Secure Networking (RSN) section.

1.6.3 Attack On Ethernet Switch

Although the research literature has primarily focused on higher layer security, the Stuxnet [7] attack has demonstrated that an innocuous looking USB drive could easily wreck havoc in a local area network (LAN) environment without any need for an Internet connection. Ethernet technology is built around self-learning and configuring protocols. This allows for ease of management, but at the same time introduces several security vulnerabilities [8]. We provide a brief review of some of the possible attacks here.

Media Access Control Attack: Switch Poisoning Attack

Ethernet switches keep forwarding table entries in a Content Addressable Memory (CAM). As a switch learns about a new destination host, it updates the table and for all future communications, this table entry is looked up to forward a frame. Unlike broadcast Ethernet or Wireless LAN, these frames are not accessible to hosts attached to other ports. However, if the switch does not have a mapping to a new MAC address, i.e., which port to forward a new frame to, it will flood the frame on all of its outgoing ports. An attacker could craft several frames with random addresses to populate an entire CAM. This would result in the switch flooding all incoming data frames to all outgoing ports, as there is no space available to enter a new mapping. This makes the frame available to the attacker attached on one of these ports. As a consequence a MAC flooding attack would also affect all the VLANs filling their CAM. However, this kind of attack requires an insider to directly connect to an Ethernet switch or possibly used unattended Ethernet wall sockets which are still connected to a port. Mitigating this kind of attack would require authentication and verification of MAC addresses from some local database of legitimate addresses before populating the forwarding table entry.

MAC Spoofing: attacks occur when an attacker eavesdrops on a link and detects the MAC address of a target host. It then masquerades as a legitimate host by altering the MAC address of its host to match the newly detected MAC address. The attacker floods the network with the newly configured MAC address directing traffic to itself by altering the switch forwarding table entry. The switch is now tricked into forwarding frames destined for the target host to the attacking host.

The MAC address is not expected to be used for security. The 802.1X discussed earlier is a good starting point to prevent unauthorised users from accessing any service on a network. As a side issue, a user may choose to spoof his or her MAC address in order to protect their privacy. Most popular operating systems support MAC address randomisation to avoid devices being tracked based on MAC address.

Address Resolution Protocol (ARP) Spoofing: attacks occur when an attacker sends a fake ARP message over a LAN binding a target's IP address to its own MAC address. Once it manages to compromise the ARP table, it will start receiving any data intended for the target's IP address. ARP spoofing can also be used for DoS attacks by populating the ARP table with multiple IP addresses corresponding to a single MAC address of a target server, for example. This would then redirect unnecessary traffic to the target, keeping it busy processing these messages. ARP Spoofing is also helpful in session hijacking and MiTM attacks.

Some vendors use a verification process where they inspect the MAC address and IP address information in ARP packets against the MAC-IP bindings contained in a trusted binding table. This allows for any ARP packets that do not have a entry in the binding table to be discarded. The binding table must be updated frequently to avoid blocking legitimate updates.

VLAN hopping: VLAN hopping attacks allow an attacking host on a VLAN to gain access to resources on other VLANs that would normally be restricted. There are two primary methods of VLAN hopping: switch spoofing and double tagging.

In a switch spoofing attack, an attacking host impersonates a trunking switch responding to the tag-

ging and trunking protocols (e.g. IEEE 802.1Q or Dynamic Trunking Protocol) typically used in a VLAN environment. The attacker now succeeds in accessing traffic for multiple VLANs. Vendors mitigate these attacks by proper switch configuration. For example, the ports are assigned a trunking role explicitly and the others are configured as access ports only. Also, any automatic trunk negotiation protocol can be disabled. In a double tagging attack, an attacker succeeds in sending its frame to more than one VLAN by inserting two VLAN tags to a frame it transmits. However, this attack does not allow them to receive a response. Again, vendors provide recommended configuration methods to deal with these possible attacks. A comprehensive survey of Ethernet attacks and defence can be found in [8] and vendor specific courses.

1.7 Wireless LAN Security

Wireless LANs are more vulnerable to security risks due to the broadcast nature of media, which simplifies eavesdropping. The Wireless Equivalent Privacy (WEP) protocol, despite being obsolete due to its design flaws, provides several important lessons for learning how not to design a security protocol. The WEP protocol was designed to provide integrity, confidentiality and authentication. It uses a symmetric key encryption method where the host shares a key with an access point using out of band methods, mostly pre-installation by an administrator or a home network user. The sender calculates a 32-bit Integrity Check Value (ICV) using a cyclic redundancy check (CRC) algorithm over the payload. A 104-bit shared key combined with a 24-bit initialisation vector (IV) is fed into a pseudo random number generator (PRNG) such as a RC4 stream cipher. The plaintext payload and the CRC of the frame are then combined with the key sequence generated by the RC4 using bit-wise exclusive-or operation to encrypt the frame. A new IV is used for each frame.

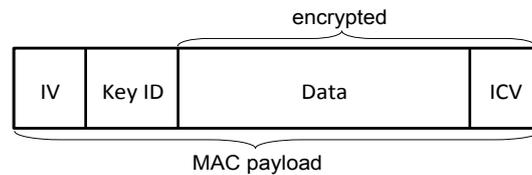
For authentication, the access points (APs) advertise via beacon frames whether authentication is necessary or not. However, not all APs support this feature. If authentication is required, before association a host connecting to an AP would receive a 128-bit NONCE from the AP. It would encrypt the NONCE with the shared key and send it back to the AP. The AP would decrypt this response with the shared key and verify whether it matches the NONCE it sent originally. The receiver would extract the IV received in plaintext, input IV and shared secret key into PRNG, get a keystream, XOR the keystream with the encrypted data to decrypt data + ICV and finally verify the integrity of the data with the ICV.

The WEP protocol has a number of design flaws. First, the use of a 24-bit IV introduces a weakness into the scheme in that 2^{24} or 16 million unique IVs can be exhausted in high speed links in less than 2 hours. Given that IVs are sent in plaintext, an eavesdropper can easily detect this reuse and mount a known plaintext attack. Using the RC4 in WEP allows for the Fluhrer, Mantin and Shamir (FMS) attacks. In the FMS, an attacker can recover the key in an RC4 encrypted stream by capturing a large number of messages in that stream [9].

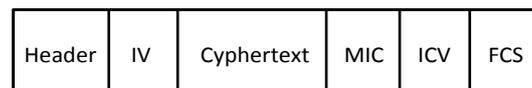
The linear CRC algorithm is good for detecting random link errors but is a poor choice for the malicious modification of messages. Even though the ICV is encrypted, an attacker can modify a message and then adjust the CRC accordingly. Such a frame will pass an integrity check at the receiver. This problem could have been avoided by choosing strong cryptographic techniques such as message authentication code and signatures as discussed in higher layer protocols.

Given the poor security design of WEP, the Wi-Fi Alliance took the job of securing wireless networks. An interim standard called the Wi-Fi Protected Access (WPA) was quickly developed for backward hardware compatibility while WPA2 was being worked out. Figure 8.

The WPA uses the Temporal Key Integrity Protocol (TKIP) but maintains RC4 for compatibility. The pre-shared key (PSK) also known as WPA-Personal is similar to the WEP-Key. However, the PSK is used differently, a NONCE and PSK are hashed to generate a temporal key. Following this, a cryptographic mixing function is used to combine this temporal key, the TA (transmitter MAC address), and the sequence counter resulting in one key for encryption (128 bits) and another key for integrity



a) 802.11 frame with WEP



b) 802.11 frame with WPA

Figure 8: 802.11 WEP and WPA Frames

(64 bits). As a consequence, every packet is encrypted with a unique encryption key to avoid FMS style attacks. Also, the WPA extends the WEP IV to 48 bits, which is used as a packet sequence counter. It would take 100 years to replay the same IV. A packet received out of order, would be dropped by the receiving station. Figure 8 part-b shows, the new Frame check sequence (FCS) field, a CRC-32 checksum for error correction. A hash function based the new field Michael (MIC) is used for integrity check. The WPA had its own share of attacks as reported in literature [10].

The Wifi alliance WPA2 standards derived from the IEEE 802.11i standards were finalised in 2004. WPA2 relies on more powerful hardware supporting a 128-bit AES Counter Mode with the Cipher Block Chaining Message Authentication Code Protocol (CCMP). These methods are discussed in the Cryptography chapter ???. It also provides an improved 4-way handshake and temporary key generation method.

1.7.1 Robust Security Network (RSN)

An earlier section described the evolution of the WEP into the WPA and WPA2. However, the IEEE 802.11i working group came up with the RSN framework to provide the strongest form of security. It adopts the 802.1X based mechanisms for access control as discussed above. Authentication and key-generation are done via the EAP. It continues to use the TKIP and CCMP for various cryptographic functions such as encryption/decryption, integrity check in addition to origin authentication and replay attack detection. Stallings [4] provides a good overview of the RSN protocols and standards.

The RSN Key derivation mechanisms are quite involved, as can be seen in figure 9. We will provide a summary of this as many other protocols (including Cellular GSM) following a similar scheme as the pairwise key scheme provides a mechanism to generate dynamic session keys each time a user starts a new session.

As a starting point, the user device and AP would have a pre-shared key (PSK) using out of band

methods. However, this is not a scalable solution and typically in an enterprise setup using the IEEE 802.1X, a Master session key (MSK) is generated during the authentication phase. With these two options available, a Pairwise master key (PMK) can be generated in the following two ways: use the PSK as the PMK or derive the PMK from the MSK using the Pseudo Random Function (PRF). The PSK also uses the host and AP addresses when generating the PTK providing additional defence against session hijacking and impersonation. Further, a NONCE is used in the mix to achieve good random keying material. The PTK is now split into three ways, thus generating separate keys for each function.

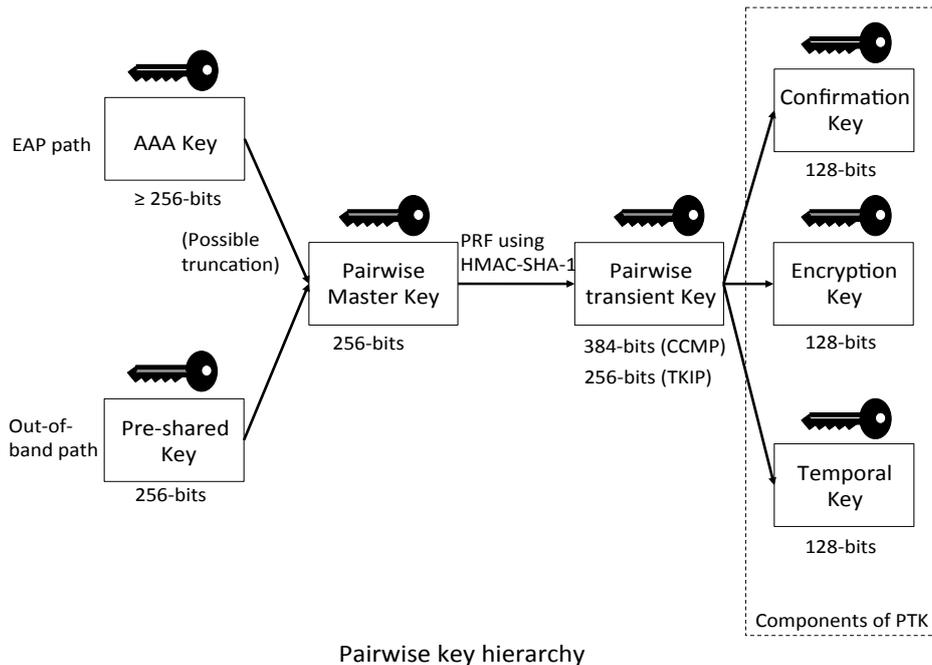


Figure 9: WLAN RSN Pairwise Key Hierarchy

The RSN also caters for a group key generation where a host can communicate with a multicast group, as shown in figure 10. This key is generated by the AP and distributed securely to the hosts associated using the secure pairwise keys derived above. This group key can be changed periodically based on a variety of network policies. The Group Temporal Key generation method is not defined in the standards.

1.8 Network Defence Tools

Ideally, attacks should be detected as early as possible, or even predicted before they have started so that they can be prevented altogether. We will discuss a number of approaches.

1.8.1 Packet Filters/Firewalls

The term filter is used for a set of rules configured by an administrator to inspect a packet and perform a matching action, e.g. let the packet through, drop the packet, drop and generate a notification to the sender via ICMP messages. Packets may be filtered by source and destination network addresses, protocol type (TCP, UDP, ICMP), TCP or UDP source/destination port numbers, TCP Flag bits (SYN/ACK), rules for traffic from a host or leaving the network via a particular interface and so on. This was typical of the first generation packet filters, which worked on the inspection of header fields. These filters did not retain any state information about the packets/flows/sessions they belonged to.

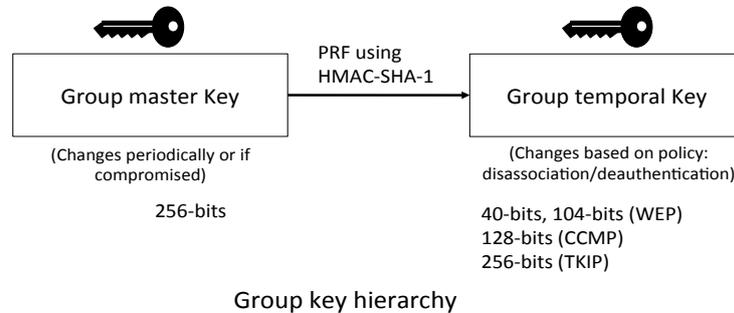


Figure 10: WLAN RSN Group Key Hierarchy

As more computing power and cheaper memory became available, the next generation of packet filters started to track transport layer flow, a chain of packets belonging to a session, known as *stateful* filters.

The packet filters, aka, Firewall system can be co-located with routers or can be implemented as specialised servers. In either case, they are gatekeepers, inspecting all incoming/outgoing traffic. The filters are set based on a network's security policy and the packets are treated accordingly. Although firewalls play a key role in securing a network, taking down a firewall can potentially wreak havoc for organisations dependent on networking technology.

1.8.2 Application Gateway (AG)

As we saw earlier, a firewall can check rules based on the filtering criterion using TCP/UDP protocol headers, port number etc. However, many organisations use application level gateways, aka application proxy, to perform access control, as they facilitate additional requirements of user authentication before a session is admitted. These AGs can inspect information from the full five-layer (Internet) or 7-layer OSI stack, except for encrypted bits. In a typical setting, the application gateway will use a firewall's services after performing authentication and policy enforcement. Both AG and a firewall are also co-located in many deployments. A client willing to access an external service would connect to the AG first. The AG would prompt for authentication before initiating a session to the external server. The AG would now establish the connection with the destination acting as a relay on behalf of the client, essentially creating two sessions: one between the client and the AG and one between the AG and the destination.

Another interesting application of an AG is SSL termination. An incoming webserver SSL connection could be terminated at the AG so that it could do the resource intensive encryption/decryption and pass the un-encrypted traffic to the back-end servers. This allows the workload on these busy servers to be reduced in addition to implementing security measures.

Higher level security provided by an AG comes at the expense of additional hardware/software resources. Further, an AG can slow down the connection as authentication, policy checks and state maintenance are performed to keep track of every session going through the AG. Another complexity involved with an AG is the need to configure it for each application, or possibly implemented as multiple application specific servers.

1.8.3 Circuit-level Gateway(CG)

A circuit-level gateway is a proxy that functions as a relay for TCP connections allowing hosts from a corporate Intranet to make TCP connections over the Internet. The CGs are typically co-located with the firewall. The most widely used CG in use today is SOCKS. For end user applications, it runs transparently as long as the hosts are configured to use SOCKS in place of a standard socket interface. A CG is simple to implement compared to an AG as it does not need to understand application layer protocols.

1.8.4 Intrusion Detection Systems (IDS)

Intrusion detection systems are a must for securing networks. Similar to AGs, they inspect higher layer information and many more attributes of sessions beyond what a packet-filter or a firewall can do. An IDS would monitor network traffic with the help of agents/sensors/monitors on the network and generate alarms when it detects (or thinks it has) suspicious activities. Essentially, the IDS would compare the traffic against what it considers normal traffic and, using a range of techniques, would generate an alert. False alarms are a huge problem for network/security administrators despite decades of research. For example, false positives may be generated by the IDS for legitimate hosts carrying out identical legitimate behaviour that may appear malicious. We will now consider a situation where a legitimate domain accessed frequently by hosts in a network becomes temporarily unreachable. The failed DNS queries to the same domain in this instance would be generated for many hosts and may appear suspicious, but should not be considered malicious activity. Likewise, a false negative would cause classifying malicious activity as benign.

The following are the two main categories of IDS:

- Signature-based intrusion detection systems compare monitored traffic against a database containing known threat signatures similar to virus scan software. The database has to be continually updated, however, or it will not detect new types of attacks. Signatures can be as simple as a source/destination IP address or contain many other protocol headers including certain patterns in the payload. We provide a simple example from an open source IDS Snort below.

```
alert tcp any any -> 192.168.5.7/24 80
(content:"GET"; msg:"WWW GET has been detected";
sid:1000007; rev:1;)
```

In this trivial example, the action is 'alert'. The source is defined for any TCP flow with any address. The destination is defined as 192.168.5.7/24 at port 80. The rule is defined to check if the packet contains a 'GET' string and then generate an alert. The 'sid' or Snort Identifier refers to the Snort rule used. Snort provides a rich set of rules but allows users to define their own.

IDS systems impose heavy workload on an IDS which has to compare against huge numbers of signatures. Speed of detection plays a key role in preventing such attacks. Several systems deploy parallel and distributed detection systems that can cope with high traffic rates on large networks and allow online detection; others exploit parallelism at the hardware level in order

to overcome processing delays so that packets and flows can be processed at high speeds, providing faster results. A lot of research has also focused on faster patterns or rule matching with the goal of reducing packet processing delays.

- Anomaly-based intrusion detection systems use statistical features of normal traffic to compare against the monitored traffic. The criterion to capture normal traffic could be bandwidth usage, protocols, ports, arrival rate and burstiness[11]. For example, a large percentage of port scans would generate an alert. Attacks can be detected by monitoring hosts or networks for behaviour typical of different attacks. A target link flooding attack aims to overwhelm a particular link in the network, thus disconnecting a selected network region or a server from the network. Observing an increase in `traceroute` packets in the network could indicate an upcoming target link flooding DDoS attack.

Despite using machine learning techniques such as Linear Regression, Neural Networks, Deep learning etc, which train a classifier from normal or malicious data and use it to identify the same behaviour within future data, false positive rates of these systems remain high [12].

Another way of classifying IDSes is the point of monitoring for malicious behaviour. A host intrusion detection system (HIDS) runs on individual hosts in the network. Most virus scan software would have this feature where they also monitor inbound and outbound traffic in addition to the usual virus scanning. This can be particularly helpful if the hosts have been compromised and form part of a bot to attack other servers/networks. In contrast, a network intrusion detection system is deployed at a strategic locations within the network to monitor inbound and outbound traffic to and from the devices in various segments of the network.

1.8.5 An Intrusion Prevention System (IPS)

An intrusion prevention system distinguishes itself from an IDS in that it can be configured to block potential threats by setting filtering criteria on routers/switches at various locations in the network.

IPS systems monitor traffic in real-time dropping any suspected malicious packets, blocking traffic from malicious source address or resetting a suspect connection. In most cases, an IPS would also have IDS capabilities. Several IDS/IPS tools raise an alert for a spam preparation stage, indicated by a rise in DNS MX queries that spam bots generate to discover a mail server before sending spam emails.

The IPS system is proactive and, in theory, it should work autonomously without intervention from a security/network administrator. For example, on inspecting the headers, if an IPS system suspects an email to be unsafe, it could prevent it from being forwarded to an end user. However, the risk of blocking legitimate traffic is a huge burden due to false positives or the mis-configuration of these systems.

1.8.6 Network Perimeter Protection

The idea of network perimeter protection comes from the ancient technique of using walls such as Hadrian's Wall or the Great Wall for protecting a city. In networking parlance, a zone called a demilitarised zone (DMZ), aka a perimeter network, is created. All external untrusted users are restricted from using the services available in this zone. Typically, an organisation's public web server and authoritative DNS would reside in the DMZ. The rest of the network is partitioned into several security zones by a security architect. Each zone is managed by one or more of the IDS, IPS, AG systems based on the significance of the information/infrastructure to be protected. Additionally, a centralised security management system can also collect information from the distributed security entities and make policy/filtering decisions to be implemented at various security check-points in the network.

1.9 Advanced Network Security Topics

1.9.1 Software Defined Network, Virtualisation

Software Defined Networking (SDN) has emerged as a new paradigm for managing and controlling the operation of networks. SDN separates the packet forwarding functionality of the forwarding devices, i.e. the data plane from the control plane. The centralisation of control functions and network intelligence presents opportunities to enhance security networks, including new ways to prevent, detect and react to threats, as well as providing innovative security services [13, 14]. For example, a DDoS attack can be inferred by the central controller more accurately, and a threat mitigation application may dynamically reprogram switches at the network perimeter to drop malicious traffic flows. A user on an infected machine can automatically be routed to a web-server issuing a quarantine notification. Another group of researchers has focussed on securing the SDN platform itself. Zerkane et al [15] have methodically analysed and reported 114 SDN vulnerabilities.

A recent trend in networking has been the use of Network Functions Virtualisation (NFV). Specialised network middleboxes such as firewalls, encoders/decoders, DMZs, deep packet inspection units have typically been closed black box devices running proprietary software [16]. NFV researchers have proposed the deployment of these middleboxes entirely as virtualised software modules and managed via standardised and open APIs. The goal is to reduce capex and allow for the rapid introduction of new services to the market. SDN plays an important role in interconnecting these virtualised network functions in a dynamic and transparent manner. Yang et al [17] provided a comprehensive survey on security issues in NFV.

1.9.2 Internet of Things (IoT) Security

Transport Layer Security (TLS) and Datagram TLS (DTLS) are cornerstones of the IoT security. Prominent IoT application layer protocols adopt either TLS or DTLS as their security protocol in combination with Public Key Cryptography (PKC) or a Pre-Shared-Key (PSK) suite. Such IoT application frameworks fulfill standard security requirements similar to traditional Internet applications. Since TLS requires a TCP connection, DTLS is widely used for limited bandwidth and lower reliability, as it is connectionless and UDP-based. While DTLS is designed to be used in constrained devices with limited communication capability, the End-to-End (E2E) communication manner of the DTLS causes scalability issues in large-scale IoT applications, especially under low-bandwidth standards such as IEEE 802.15.4. Given the emerging characteristics of heterogeneity, energy and performance, scalability, mobility and management, it is obvious that the current PKC with the E2E infrastructure will almost certainly not scale to accommodate future IoT applications [18]. E2E communication causes unscalable communication overheads and delays in large-scale applications. Furthermore, constrained devices are not equipped with adequate performance/memory to process resource-intensive PKC cryptography suites, resulting in performance/security degradation.

CROSS-REFERENCE OF TOPICS VS REFERENCE MATERIAL

This will be completed in the next version

	kopka2003 [19]	gratzer2016 [20]
?? ??				
?? ??	c4			
?? ??	c11,c12			
?? ??	c7	ll.5		
...				

FURTHER READING

REFERENCES

- [1] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach (7th Edition)*, 7th ed. Pearson, 2017.
- [2] D. Dolev and A. C. Yao, "On the security of public key protocols," vol. 29, no. 2, pp. 198–208, March 1983.
- [3] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, "Understanding the mirai botnet," in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, 2017, pp. 1093–1110. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/antonakakis>
- [4] W. Stallings, *Network Security Essentials, Applications and Standards (6th Edition)*, 6th ed. Pearson, 2016.
- [5] S. T. Zargar, J. Joshi, and D. Tipper, "A survey of defense mechanisms against distributed denial of service (ddos) flooding attacks," *IEEE Communications Surveys Tutorials*, vol. 15, no. 4, pp. 2046–2069, Fourth 2013.
- [6] S. Kent, C. Lynn, and K. Seo, "Design and analysis of the secure border gateway protocol (s-bgp)," in *Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00*, vol. 1, Jan 2000, pp. 18–33 vol.1.
- [7] R. Langner, "Stuxnet: Dissecting a cyberwarfare weapon," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 49–51, 2011.
- [8] T. Kiravuo, M. Sarela, and J. Manner, "A survey of ethernet lan security," *IEEE Communications Surveys Tutorials*, vol. 15, no. 3, pp. 1477–1491, Third 2013.
- [9] A. Stubblefield, J. Ioannidis, and A. D. Rubin, "Using the fluhrer, mantin, and shamir attack to break wep," in *NDSS*, 2002.
- [10] E. Tews and M. Beck, "Practical attacks against wep and wpa," in *Proceedings of the Second ACM Conference on Wireless Network Security*, ser. WiSec '09. New York, NY, USA: ACM, 2009, pp. 79–86. [Online]. Available: <http://doi.acm.org/10.1145/1514274.1514286>
- [11] T. T. T. Nguyen and G. Armitage, "A survey of techniques for internet traffic classification using machine learning," *IEEE Communications Surveys Tutorials*, vol. 10, no. 4, pp. 56–76, Fourth 2008.
- [12] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *2010 IEEE Symposium on Security and Privacy*, May 2010, pp. 305–316.
- [13] S. Taha Ali, V. Sivaraman, A. Radford, and S. Jha, "A survey of securing networks using software defined networking," *IEEE Transactions on Reliability*, vol. 64, pp. 1–12, 09 2015.
- [14] A. Shaghghi, M. A. Kaafar, and S. Jha, "Wedgetail: An intrusion prevention system for the data plane of software defined networks," in *Proceedings of the 2017 ACM on Asia Conference*

REFERENCES

- on *Computer and Communications Security*, ser. ASIA CCS '17. New York, NY, USA: ACM, 2017, pp. 849–861. [Online]. Available: <http://doi.acm.org/10.1145/3052973.3053039>
- [15] S. Zerkane, D. Espes, P. Le Parc, and F. Cuppens, “Vulnerability analysis of software defined networking,” in *Foundations and Practice of Security*, F. Cuppens, L. Wang, N. Cuppens-Boulahia, N. Tawbi, and J. Garcia-Alfaro, Eds. Cham: Springer International Publishing, 2017, pp. 97–116.
- [16] Z. G. A. Gember, P. Prabhu and A. Akella, “Toward software defined middlebox networking,” in *In Proceedings of the 11th ACM Workshop on Hot Topics in Networks (HotNets)*, Redmond, WA, 10 2012, pp. 7–12.
- [17] W. Yang and C. Fung, “A survey on security in network functions virtualization,” in *2016 IEEE NetSoft Conference and Workshops (NetSoft)*, June 2016, pp. 15–19.
- [18] J. Y. Kim, W. Hu, D. Sarkar, and S. Jha, “Esiot: Enabling secure management of the internet of things,” in *Proceedings of the 10th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '17. New York, NY, USA: ACM, 2017, pp. 219–229. [Online]. Available: <http://doi.acm.org/10.1145/3098243.3098252>
- [19] H. Kopka and P. W. Daly, *Guide to LaTeX (Tools and Techniques for Computer Typesetting)*. Addison-Wesley Professional, 2003.
- [20] G. Grätzer, *More Math Into LaTeX*. Springer, 2003.