



Computer Security (COM-301)

Network security

Slides from Carmela Troncoso
Some slides/ideas adapted from: George Danezis

1

Up to here: attacks on hosts
What about the network?




2

Up to here, in all lectures we talked about the network as if it was a tube between Bob and Alice (or Minion Bob and Gru, or Rick and Morty).

Up to here: attacks on hosts
What about the network?

The network is not a tube!!!



3

The network is actually not a tube, but a much more complex infrastructure.

Messages between Bob and Alice are converted into packets that are routed across **routers**, first in their **Local Area Networks (LAN)**, and then between networks (using **Wide Area Networks, WAN**, protocols).

Routers between local area networks are organized in **Autonomous Systems (AS)**, which follow a unique routing policy. Each AS is typically owned by one entity (e.g., and **Internet Service Provider, ISP**).

Desired properties

Confidentiality, Integrity, Availability,
Authentication, Authorization?

Naming security: The association between lower level names (eg. network addresses) and higher level names (e.g. Alice / Bob) must not be influenced by the adversary

Routing security: The route over the network and the eventual delivery of messages must not be influenced by the adversary

Session security: Messages within the same session, cannot be modified (keep ordering and no adding/removing messages)

Content security: The content of the messages must not be readable or influenced by adversaries

In order for Alice to securely send a string of messages to Bob, the following is needed:

1) Know the address of the receiver. Humans do not remember machine-readable addresses; and also machine-readable addresses change over time. Thus we need Naming services to map high-level names to low-level names. It is important that we can grant the security of this naming services. They must:

- Provide Integrity: the (name, address) association must not be modified by unauthorized entities
- Authentication: users need to be able to verify that the (name,address) association come from an authoritative source (i.e., a source entitled to set this association)
- Availability: the naming service must be available to grant service to the users.

2) Choose a route (i.e., a chain of routers) between Alice and Bob. It is important that this route cannot be controlled by the adversary:

- The integrity of the route must not be modified by unauthorized parties
- The route must only go through the chosen routers, thus the identity of the routers must be authentic
- The route (and the routers therein) must be available to route packets

3) All messages / packets within a session must arrive to the receiver without modification, in

order or number (i.e., no messages can be added or removed): the integrity of the messages must be maintained, and they must come from the authentic source.

4) Finally, the content must not be influenced (its integrity must be protected), and many times not readable (i.e., we need confidentiality), by non-authorized parties.

Desired properties

Confidentiality, Integrity, Availability,
Authentication, Authorization?

Naming security: The association between lower level names (eg. network addresses) and higher level names (e.g. Alice / Bob) must not be influenced by the adversary	Integrity Authentication Availability (naming service)
Routing security: The route over the network and the eventual delivery of messages must not be influenced by the adversary	Integrity Authentication Availability Authorization
Session security: Messages within the same session, cannot be modified (keep ordering and no adding/removing messages)	Integrity Authentication
Content security: The content of the messages must not be readable or influenced by adversaries	Confidentiality Integrity

In order for Alice to securely send a string of messages to Bob, the following is needed:

1) Know the address of the receiver. Humans cannot remember machine-readable addresses; and also machine-readable addresses change over time. Thus we need Naming services to map high-level names to low-level names. It is important that we can grant the security of this naming services. They must:

- Provide Integrity: the (name, address) association must not be modified by unauthorized entities
- Authentication: users need to be able to verify that the (name,address) association come from an authoritative source (i.e., a source entitled to set this association)
- Availability: the naming service must be available to grant service to the users.

2) Choose a route (i.e., a chain of routers) between Alice and Bob. It is important that this route cannot be controlled by the adversary:

- The integrity of the route must not be modified by unauthorized parties
- The route must only go through the chosen routers, thus the identity of the routers must be authentic
- The route (and the routers therein) must be available to route packets

3) All messages / packets within a session must arrive to the receiver without modification, in

order or number (i.e., no messages can be added or removed): the integrity of the messages must be maintained, and they must come from the authentic source.

4) Finally, the content must not be influenced (its integrity must be protected), and many times not readable (i.e., we need confidentiality), by non-authorized parties.

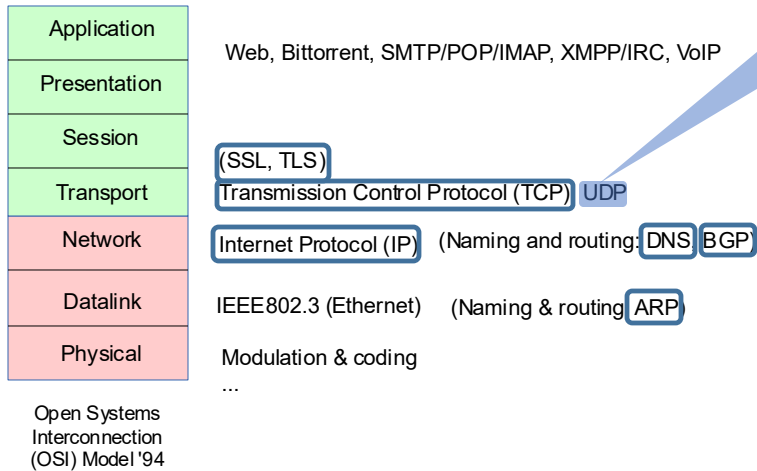
Network security in COM-301

- Do deployed network protocols provide the desired properties?
 - Naming security
 - Routing security
 - Session security
 - Content security
- What are the existing solutions to improve network security?

Where are the problems?

In this lecture security issues on...

We will not see UDP, similar concepts, but differences in the implementation



Computers communicate with each other at different layers. These are typically modeled by the Open System Interconnection (OSI) Model. This model covers from the physical layer, where pulses that codify bits are sent, to the application layer where programs talk to each other.

In this lecture we will cover protocols at different layers, and see the security problems and potential solutions.



Computer Security (COM-301)

Network security

ARP Spoofing

Slides: Carmela Troncoso
Some slides/ideas adapted from: George Danezis

8

Routing: routing IP on an Ethernet LAN




- **Ethernet:**

- Local area network (LAN) technology
- Machines have a “unique” 48 bit MAC address (Medium Access Code)

- **Internet Protocol (IP) on the LAN**

- Hosts communicate using the IP protocol
- Each machine has an IP address (4 bytes in IPv4).
 - Part of the address denotes the network and part the host



RFC791 "INTERNET PROTOCOL" <http://www.ietf.org/rfc/rfc791.txt>

A LAN is a network that spans a (typically) small geographic area, and connects devices that are all local in this area.

Inside the LAN, these devices have two addresses:

- Address at the link level: a unique address called the **MAC address**.
- Address at the network level: the **IP address**

Routing: routing IP on an Ethernet LAN

Refresh


How does IP routing work?

- Alice needs:
 - Her own IP address (eg. 192.128.5.130)
 - Bob's IP address (eg. 192.128.5.125)
 - Her "subnet mask" (eg. 255.255.255.0)
 - Her "gateway" (eg. 192.128.5.1)
- Option 1: Alice and Bob are on the same subnet
 - Address Alice AND mask = Address Bob AND mask
 - Route through the LAN
- Option 2: they are on different subnets
 - Send to gateway
 - Route through the WAN (Wide Area Network)



Alice

Send this packet from address 192.128.5.130 to 192.128.5.125



10

To route an IP packet (format on the bottom right) Alice needs to know her own address, and Bob's address.

Furthermore, Alice needs to know her own *subnet*, which tells her on which LAN she and Bob are; and her *gateway* which she needs to use to communicate with Bob if he is not on the same subnet.

First, Alice checks if she and Bob are on the same subnet by using the mask. If they are, Alice needs Bob's MAC address to route his message. If they are not, Alice needs the MAC address of the gateway to route her message out of the LAN.

Routing: routing IP on an Ethernet LAN

Refresh

How does IP routing work?

and inside the LAN?

- Alice does not (want to) know network details
- Alice does not know Bob's MAC address



Alice

Send this packet
from address
192.128.5.130
to 192.128.5.125

How can she learn about Bob's MAC?

ARP: "translation" between IP address and MAC address

- Each host maintains a cached table of IP ↔ MAC mappings
- If not available: broadcast an ARP request to query for target IP
- An ARP reply responds with the MAC address for that IP

```
*-----*
| HTYPE (2 bytes) |
| PTYPE (2 bytes) |
| HLEN (1) | PLEN (1) |
| OPERATION (2) |
| Sender HA (HLEN) |
| Sender PA (PLEN) |
| Target HA (HLEN) |
| Target PA (PLEN) |
*-----*
```

11

ARP (Address Resolution Protocol) enables machines to find mappings between IPs and MACs inside.

HTYPE (Hardware Type): Defines the network link type. For Ethernet/Wi-Fi, this is set to 1.

PTYPE (Protocol Type): Defines the protocol address type. For IPv4, this is set to 0x0800.

HLEN (Hardware Length): The length of the hardware address in bytes. For a MAC address, this is 6.

PLEN (Protocol Length): The length of the protocol address in bytes. For an IPv4 address, this is 4.

OPERATION: This specifies what the packet is doing:

1 = **ARP Request** ("Who has this IP?")

2 = **ARP Reply** ("I have that IP, here is my MAC address.")

Sender HA (Hardware Address): The MAC address of the device sending the packet.

Sender PA (Protocol Address): The IP address of the device sending the packet.

Target HA (Hardware Address):

In a *Request*, this is zero (or ignored) because the sender doesn't know it yet, that is

what they are asking for!

In a *Reply*, this is filled in with the destination's MAC address.

Target PA (Protocol Address): The IP address the sender is looking for.

It works as follows: every machine maintains a list of the known mappings (IP,MAC). When the machine needs to send a message to a new IP in the subnet, it broadcasts a message (format on the right, H stands for hardware, P for Protocol). The sender broadcasts its Hardware Address (HA) and the target Protocol Address (PA) in this case IP Address. The target machine responds with a packet directed to the sender with the requested MAC address.

Routing: routing IP on an Ethernet LAN

Naming security: The association between lower level names (eg. network addresses) and higher level names (e.g. Alice / Bob) must not be influenced by the adversary

**Integrity
Authentication**

Does ARP provide naming security?

ARP: "translation" between IP address and MAC address

- Each host maintains a cached table of IP ↔ MAC mappings
- If not available: broadcast an ARP request to query for target IP
- An ARP reply responds with the MAC address for that IP

```
*-----*
| HTYPE (2 bytes) |
| PTYPE (2 bytes) |
| HLEN (1)       | PLEN (1) |
| OPERATION (2)  |
| Sender HA (HLEN)|
| Sender PA (PLEN)|
| Target HA (HLEN)|
| Target PA (PLEN)|
*-----*
```

No Integrity check, nor Authentication



12

Note that the packet format **does not include signatures or message authentication code**. Therefore the receiver cannot check:

- the *integrity* of the packet, i.e., that it has not been modified in transit.
- The *authenticity* of the sender, i.e., that there is no adversary impersonating the sender.

ARP spoofing

If nobody checks...

You can impersonate! (provide the identity of others)

What can you achieve?

- **Just impersonation is bad**
- **Man in the middle:** provide two hosts (sender/receiver) with your MAC address
 - Monitor communication or tamper with it
- **Abuse resource allocation**
- **Denial of Service:** avoid that packets arrive to one host

Also bad for

Routing security: The route over the network and the eventual delivery of messages must not be influenced by the adversary

Integrity
Authentication
Availability
Authorization

13

The modification of ARP packets by changing the content or impersonating other computer is called **ARP spoofing**.

ARP spoofing enables an adversary to:

- Deploy a *man in the middle (MITM)* attack. For instance, Mallory is able to convince Alice that Mallory's MAC is Bob's MAC, and to convince Bob that Mallory's MAC is Alice's MAC. From then on, Mallory will receive all messages from Alice to Bob and vice-versa. This gives Mallory access to the content of the message and even modify their content.
- Deny service to a machine in the network. This could be done by launching a MITM attack and not relaying messages, or by just ensuring that any ARP request to the victim of the attack is always responded with a wrong MAC address. This way, no sender will find the correct mapping (IP,MAC) for the victim and the victim cannot receive messages.
- The possibility to impersonate other computers in the network also enables adversaries to steal resources allocated to others. For instance imagine a LAN in which bandwidth is limited per MAC address. By spoofing others' address you can enjoy more bandwidth than initially allowed.

ARP spoofing

If nobody checks...

You can

What can you do?

- Just in

- Man in

- Abuse resource allocation

- Denial of Service: avoid that packets arrive to one host

Also bad for

Routing security: The route over the network and the eventual delivery of messages must not be influenced by the adversary

Integrity
Authentication
Availability
Authorization

The same happens in DNS, IP, Ethernet,...
No network protocol was (initially) designed with security in mind!

BECAUSE OF A VERY NAÏVE THREAT MODEL:
outsiders are bad, insiders behave, trust them!

address

14

This lack of security in ARP stems from a bad threat model. This protocol assumes that only externals (i.e., not the members of the network) may act maliciously.

As we will see in the rest of the lecture, this underestimation of the threat model is common to all network protocols and enables most of the attacks that nowadays threaten the security of the internet.

ARP spoofing - Defenses

- Use of static, read-only entries for critical services in the ARP cache of a host
- Use ARP spoofing detection and prevention software
 - check if one IP has more than one MAC or one MAC reported by multiple IPs
 - certify requests by cross-checking
 - sends email if IP-MAC association change



Separation of privilege: force the adversary to gain control of more entities

15

The problem of ARP spoofing is rooted on the fact that there is no authentication of who is providing a mapping IP-MAC.

One solution would be to not ask in the network (where spoofing is possible) but to ask in an out of band channel, e.g. asking in person where authentication is easy. Then one can store the learned mappings on a static read-only mapping. While this is a very secure means: authentication is granted, it does not scale and it is not flexible: everytime there is a change in an association one needs to manually learn the new mapping and update the table. Thus, this is typically only used for critical services whose failure may have big repercussions. For instance, to secure the mapping IP-MAC for the gateway.

For the others, one can prevent spoofing by:

- Instead of taking the first IP-MAC association returned by ARP as truthful, check whether there is an inconsistency: there is more than one IP associated to this MAC in your cache table, there is an IP for which you have seen more than one MAC, you observe packets with more than one MAC associated to the IP.

- Instead of taking the first IP-MAC association returned by ARP as truthful, ask other members of the network if they have the same observation, i.e., cross-check the discovery.
- Send an email to the user, or the system administrator, if one observes a change in an IP-MAC association. Only validate the change if a person confirms it is correct.

Note that the two last methods effectively implement the concept of separation of privilege by requiring that the adversary compromises more than one entity (more than one machine in the network in the first case, and one human in the second case).



Computer Security (COM-301)

Network security

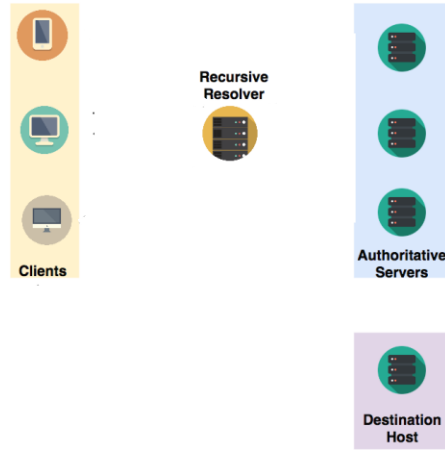
DNS Spoofing

Slides from Carmela Troncoso
Some slides/ideas adapted from: George Danezis

16

Domain Name Service (DNS)

Refresher



17

In order to resolve a domain, a client sends a DNS query to a *recursive resolver*, a server typically provided by the ISP with resolving and caching capabilities. If the domain resolution by a client is not cached by the recursive name server, it contacts Root Server, Top Level Domain server, and then *authoritative name server* which hold a distributed database of domain names to IP mappings. The recursive resolver traverses the hierarchy of authoritative name servers until it obtains an answer for the query, and sends it back to the client. The client can use the resolved IP address to connect to the destination host.

DNS spoofing – Attacks

DNS Spoofing

Cache poisoning: corrupt the DNS resolver with fake pairs (IP, domain)

DNS Hijacking: corrupt the DNS responses with fake pairs

What can you achieve?

- **Denial of Service:** avoid that packets arrive to one host → censorship
- **Redirection:** reroute clients to malicious host
 - Malicious host attacks client (e.g., serving malware...)
 - Malicious host act as man in the middle (e.g., monitoring)

19

By being able to influence the associations (domain, IP) observed by a victim the adversary can:

- Deny service to a machine by avoiding that anybody learns its IP. This can be seen as a form of *censorship* in which no user can access a given service whose IP is continuously spoofed.
- Put herself in a Man in the middle position to mediate connections between client and domain (e.g., to *monitor* the actions of the victim)
- Redirect clients to her machine to attack the client machine (e.g., serve malware)

DNS spoofing – Defenses

Domain Name System Security Extensions (DNSSEC)

- Extensions to DNS that provide **origin authentication**
 - DNS responses are **digitally signed by authoritative name server** – prevents poisoning!
 - DNSSEC responses are not encrypted – **does not provide confidentiality!**
- 1st attempt ([RFC 2535](#)) 99-01: impractical, non-scalable, complex key management
- Nowadays (*DNSSEC-bis* [RFC 4033](#)): simplified messages and key management

DNS-over-HTTPS (DoH) ([RFC8484](#))

- Since 2019 – DNS queries over HTTPS connection (confidentiality & integrity)
- Deployed by Cloudflare (integrated in Firefox), Google, others

Others: DNS-over-TLS, DNSCrypt, DNSCurve

20

The attacks exploit the fact that there is no protection of the associations or the responses.

A first attempt to protect DNS is DNSSEC, which protects the integrity and authenticity of origin of the responses by making the authoritative name server *sign* associations. This way, when the client receives an association she can be sure that there has not been cache poisoning.

A major problem with DNS is that, even if the integrity and the authenticity of the responses, as those are not encrypted an adversary observing the communication can still *monitor* the users' actions. For this reason nowadays new protocols encrypt DNS, e.g. by tunneling requests and responses into an HTTPS or TLS connection, which provide all authenticity, integrity and confidentiality as we see later in the lecture.



Computer Security (COM-301)

Network security

BGP Spoofing

Slides by Carmela Troncoso
Some slides/ideas adapted from: George Danezis

21

If we fix DNS, do we solve the routing problem?


Routing security: The route over the network and the eventual delivery of messages must not be influenced by the adversary

Integrity
Authentication
Availability
Authorization

If we fix DNS, do we solve the routing problem? BGP (Border Gateway Protocol) ([RFC 4271](https://www.rfc-editor.org/rfc/rfc4271))

Refresher

- BGP **constructs the routing tables** between AS - Autonomous Systems with independent routing domains
 - Routers maintain tables of (IP subnet → Router IP, cost)
 - Routes change (faults, new contracts, new cables) - BGP updates constantly
 - Cost is **crucial**: BGP chooses the routes with lowest cost (real money!)



23

Border Gateway Protocol (BGP) is the protocol that enables ASes to know what is the best route to get to another AS that hosts the target IP. Basically, BGP enables ASes to advertise routes to IPs. These routes are modified fairly often and a main criteria for them is cost, which does not necessarily depend on infrastructure, but on the agreements between ASes.

Here is some more details on the protocol: <https://blog.cdemi.io/beginners-guide-to-understanding-bgp/>

<https://iamroot.tech/asndatabase/by-country.aspx?country=CH>

What are AS? Operators of a network, but there are lot of laws and paperwork to become an AS with a associated number on the internet, at some level you have a responsibility, it is a cooperative system.

Principle of recording.

BGP Security

Weak authentication mechanism between routers ([RFC 2385](#)):

- Aimed at preventing DoS
- Short shared secret (up to 80 bytes of ASCII)
- Ad-hoc message authentication code based on the weak algorithm MD5

Does this guarantee the integrity of the advertised routes?

NO!! BGP Hijacking!

- An adversary controls or compromises a BGP router *somewhere* on the Internet
- Injects false low-cost routes to redirect portions of traffic to themselves
- The routing information propagates to routing tables until it expires

What can you achieve?

- **Redirection:** surveillance, injection, modification, or censorship.

24

BGP provides some authentication mechanism, that enables routers to know that they are talking with other routers. The mechanism is very basic: uses a short secret and a broken hash algorithm.

Even if this mechanism would be secure, it only makes sure that routers speak with other routers, but does not provide any assurance on the information the routers provide. Thus, an adversary controlling a router can advertise any route they want, e.g., a very cheap route that will attract all traffic to a particular destination. This advertisement will propagate to other routers.

As in the previous cases, by rerouting traffic, the adversary can monitor or deny service.


Note that this attack **in itself does not** enable the adversary to redirect traffic to their own server (as ARP spoofing or DNS spoofing do). The attack allows the adversary to only modify the route, not the destination. If the adversary then can run an on-path attack, and modify the packets, then they can run an IP-spoofing attack or other means to redirect.

Example 1: Belarus hijacks internet (2013)

- Global traffic redirected to Belarusian ISP GlobalOneBel.
 - Daily basis throughout February 2013
 - Changing set of victims: major financial institutions, governments, and network service providers.
 - Affected countries included the US, South Korea, Germany, the Czech Republic, Lithuania, Libya, and Iran



Example 2: BGP hijacking as censorship

- 2000: Pakistan tries to censor YouTube (and accidentally shuts it down...)
 - <https://www.wired.com/2008/02/pakistans-accid/>
- 2014: Turkey bans Twitter by hijacking DNS provider routes (after direct DNS hijacking stopped working)
 - 
 - <https://www.bgpmon.net/turkey-hijacking-ip-addresses-for-popular-global-dns-providers/>
- 2017: Iran censors a number of webpages (mostly porn)
 - <https://bishopfox.com/blog/bgp-hijacking-technical-post-mortem>
- 2021: Myanmar tries to censor Twitter
 - <https://www.manrs.org/2021/02/did-someone-try-to-hijack-twitter-yes/>

BGP Spoofing - Defenses

Filtering help alleviating
(some routes should really not come from some routers).

But... there is no authority to guarantee the correctness of routes (all contractual).

Fundamental flaw (again): Design did not consider insiders as adversaries!

BGPsec

Each AS is given a certificate that links its verification key to its IP blocks.
Updates are only accepted if they are signed by the authority for the AS/IP Block.
Delegation is possible
Effort started in 2003 ([RFC8205](#)) -- weakly deployed

27

A possible defense would be, as in ARP, to try to identify inconsistencies in the routes. A packet traveling from Switzerland to Italy should not need to go through the US. Yet, there is no real definition for “correct route” as routes do also depend on economical factors.

There is a standard BGPsec, that similar to DNSsec aims at mitigating hijacking but ensuring that routes are signed by authoritative ASes, i.e., those hosting the IPs.

Spoofing: lesson to be learned



1. The network is hostile!

Routing security attacks, facilitated through **poor association of high level and low level names & addresses** (IP to Ethernet MAC / Route to router).

- **Threat model:** assumes network “insiders” are trusted to provide authoritative information.
- Also **no** integrity or confidentiality.

2. The solution is intimately linked to cryptography

Why? There is **no centralized authority** to act as either (a) originator of policy or (b) provide a trusted computing base

- Cryptography allows mutually distrustful actors to achieve some collective security properties
- Asymmetric cryptography (certificates and signatures) particularly useful for all to verify name and route associations!

But also... Who has authority?

Not a cryptographic question! related to name resolution & security policy

28

The main lessons from all these attacks are:

- 1) A bad threat model leads to many problems. In this case a threat model that does not consider insiders as potential adversaries with interest in modifying the information they provide, or rely; or simply monitor or stop the information leads to a whole suite of protocols being susceptible to attacks that endanger Internet transactions.
- 2) As there is no TCB that can enforce a policy, most of the solutions rely on cryptography to ensure integrity, authentication and confidentiality.

Yet, cryptographic solutions fall short in some cases as they still rely on a given definition of authority (e.g., authority to provide a BGP route, authority to provide a domain-IP mapping). How to define these authorities is not a security problem, but a policy one.

Computer Security (COM-301)

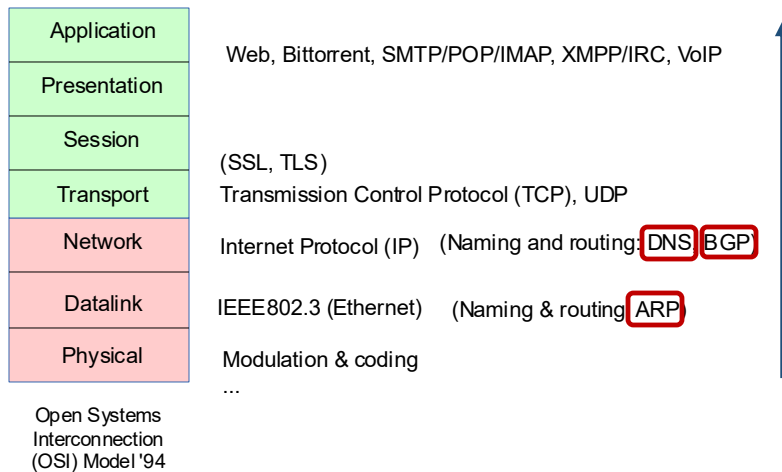
Network security

IP

Slides by Carmela Troncoso
Some slides/ideas adapted from: George Danezis

29

Where are the problems?



30

Computers communicate with each other at different layers. These are typically modeled by the Open System Interconnection (OSI) Model. This model covers from the physical layer, where pulses that codify bits are sent, to the application layer where programs talk to each other.

In this lecture we will cover protocols at different layers, and see the security problems and potential solutions.

So what about IP?

Refresher

```

0           1           2           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
-----|-----|-----|-----|
|Version| IHL |Type of Service|   Total Length   |
-----|-----|-----|-----|
|           Identification           |Flags|   Fragment Offset   |
-----|-----|-----|-----|
| Time to Live | Protocol |   Header Checksum   |
-----|-----|-----|-----|
|           Source Address           |
-----|-----|-----|-----|
|           Destination Address      |
-----|-----|-----|-----|
|           Options                   | Padding |
-----|-----|-----|-----|
|           Packet Data               |
-----|-----|-----|-----|
Example Internet Datagram Header
```

31

The Internet Protocol (IP) enables addressing and routing of packets between hosts in a network (<https://www.cloudflare.com/learning/ddos/glossary/internet-protocol/>). For the purpose of this lecture, the most important factor is that IP packets, as all the other protocols do not include any protection (just a non-cryptographic checksum to detect transmission errors) to preserve the integrity of the source and destination address, nor the integrity and confidentiality of the data.

This enables adversaries to:

- 1) Tamper with the IPs (IP spoofing)
- 2) Learn the destination of packets (privacy invasion)

IP spoofing

```
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
-----
|Version|  IHL |Type of Service|  Total Length  |
-----
| Identification | Flags | Fragment Offset |
-----
| Time to Live | Protocol | Header Checksum |
-----
| Source Address |
-----
| Destination Address |
-----
| Options | Padding |
-----
| Packet Data |
-----
Example Internet Datagram Header
```

No integrity or authentication mechanism for Source Address

What can we do?

- **Impersonation:** for instance to steal resources
- **Man in the middle:** monitor, intervention, deny service
- **Denial of Service:** fake source IP so that others send packets to a target victim with that IP

32

By being able to spoof source IPs we can achieve:

- **Impersonation:** we can fake being others. This could be useful to take advantage of others' resources, for instance in networks that allocate resources based on IP.
- (Nowadays honest ISP will have their routers filter ingress traffic, maybe mention reverse lookup?)
- **Man in the Middle:** if we can impersonate others, we can launch a man-in-the-middle attack and monitor, change, or deny communication. Note that for a man in the middle we may need to also hijack the TCP connection to achieve a link between MAC and IP (see TCP hijacking below)
- **Denial of service:** we can use spoofing to trigger the sending of packets from e.g., a server, to a victim computer. This is done by telling that server that packets originate from the victim's computer IP.

IPSec - Internet Protocol Security

- Cryptographic security properties at the IP level
 - Key exchange based on public key cryptography or shared symmetric keys
 - **Authentication Header (AH)**: authentication & integrity (HMAC), protection from replay attacks (sequence number)
 - **Encapsulating Security Payload (ESP)**: can add confidentiality on top of auth
- Two modes:
 - **Transport**:
 - protects IP packet payload using AH/ESP
 - sent with the **original IP headers**
 - **Tunnel**:
 - protects the whole packet (Headers + Payload) is protected and placed inside another packet

33

Internet Protocol Security (IPSec) complements IP with security features using cryptography. It enables key exchange (using public keys) or the use of a symmetric key if it exists a pre-shared one.

IPSec provides two protocols:

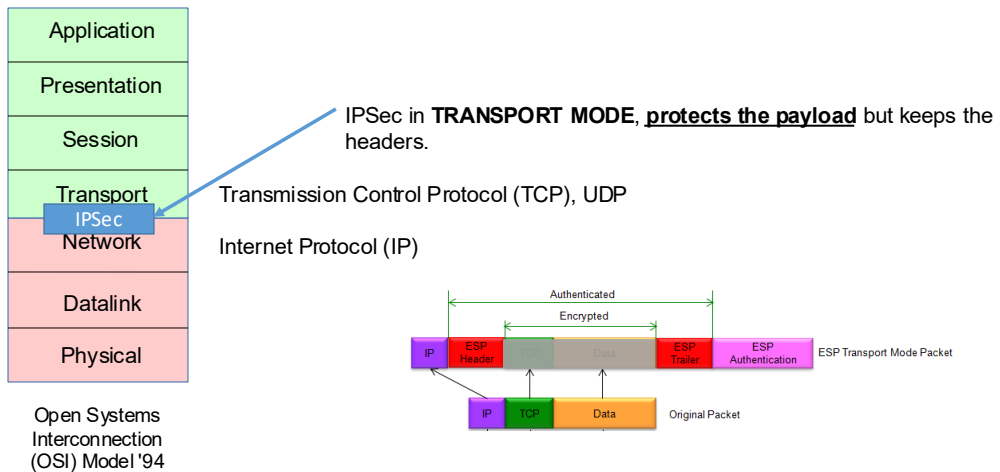
Authentication Header protocol (AH): provides a mechanism for header authentication only. AH provides data integrity, data origin authentication, and, optionally, replay protection. Data integrity is ensured by including a digest generated by a secure message authentication algorithm such as HMAC-SHA. Data origin authentication stems from the use of a shared secret key (remember from the cryptography lecture that this provides authentication for the communication parties, but not repudiation against third parties). Replay protection relies on a sequence number field with the AH header. AH authenticates IP headers and their payloads, with the exception of certain header fields that can be legitimately changed in transit, such as the Time To Live (TTL) field. AH does **not** provide confidentiality of content. Also AH breaks NAT, why?

Encapsulating Security Payload protocol (ESP): provides data confidentiality (encryption) and authentication (data integrity, data origin authentication, and replay protection). It can be configured to provide only confidentiality, only authentication, or both. When ESP provides authentication functions, it uses the same algorithms as AH, but protects different parts. AH authenticates the entire IP packet, including the outer IP header, while ESP only authenticates the IP datagram portion of the IP packet.

ESP+Null can be sort of used to avoid spoofing. You lose the strict "Anti-Spoofing" protection on the IP header itself. You are trusting that if the *data* inside is signed correctly by the sender, the packet probably came from them, even if the IP header was modified by a router along the way.

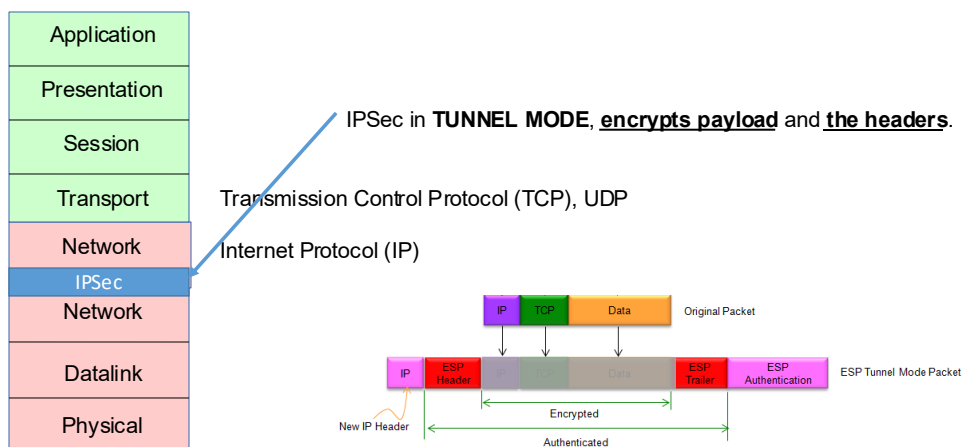
These two protocols can be used in two modes (see next slides)

Where does IPSec in Transport Mode happen?



Transport mode: it protects the packet payload (authentication and/or confidentiality depending on the choice). The original IP header is left intact. This means that the original information, such as origin and destination, are still available and observable.

Where does IPSec in Tunnel Mode happen?



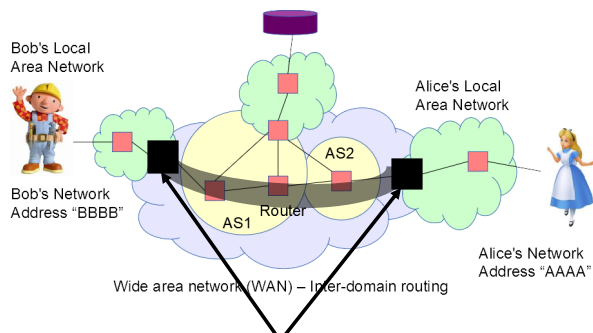
Open Systems Initiative
(OSI) Model '94

http://www.sharetednote.com/html/IP_Network_IPSec_ESP.html

35

Tunnel mode: tunnel mode creates a *new header*. The full packet is encrypted and authenticated. This way, the original source and destination of a packet are protected from potential observers.

Use of IPSec: Virtual Private Network



- IPSec in tunnel mode. The VPN
 - Looks like one single network
 - Routing internally
 - Inside VPN "tunnel" fully protected packets: confidentiality, authentication, integrity, reply

36

Using IPSec in tunnel mode is one of the means to create a **Virtual Private Network (VPN)**. A VPN is a private network that uses a public network (typically, the internet) to connect remote hosts. The VPN uses "virtual" connections to route from one host to the other. Inside the VPN, routing is local (in the slide, Bob can locally route to Alice using local routing. Inside the VPN tunnel, packets are fully protected by IPSec.

When Bob wants to send a packet to Alice, instead of doing (IP header simplified):
 $\{src_IP, dst_IP, Data\} = \{IP_Bob, IP_Alice, Data\}$, Bob prepares an encapsulated packet such that $\{IP_Bob, IP_VPN, Enc\{IP_Bob, IP_Alice, Data\}\}$.

This packet is sent to the entry of the VPN, that transmits it to the exit. The exit decrypts the packet and then can send it to Alice.

In transit, an adversary between Bob and the entry to the VPN cannot see the destination or the payload. An adversary between the entry and the exit to the VPN cannot see the origin, the destination or the payload.


Use of IPsec: Virtual Priv

Does it protect against Denial of Service?

No! Your IP still exists

Does it solve the authentication problem we have studied?

No! Only authentication at network level. It cannot authenticate programs or applications




- IPsec in tunnel mode. The VPN
 - Looks like one single network
 - Routing internally
 - Inside VPN "tunnel" fully protected packets: confidentiality, authentication, integrity, reply

37

VPNs provide confidentiality, authentication, integrity, and reply protection *for the packets*. However it still does not protect any higher layer: authentication problems at the application layer, or denial of service problems (see end of the lecture) cannot be solved with IPsec.

Virtual Private Network - other common configuration



VPN out of Bob's LAN (VPN as a Service)

AS1 can see the connection from the VPN to the server (if that connection is in the clear, it can spy)

38

VPNs do not necessarily need to join two local networks. They can be used to simply leave the local network from another exit point than the gateway. This is the typical use of a VPN to change one's IP (e.g., to user pay TV services on another country ;).

The Return Address Problem (Why they must hide your IP). Rewriting/Hiding your IP isn't just a privacy feature of a VPN, it is necessary for the VPN to work.


Think of an IP address like a Return Address on an envelope.

- **Without VPN:** You send a letter to Google. You write *your* return address on it. Google reads the letter and mails the reply directly to you.
- **With VPN:** You put your letter inside a secure, steel box (the encrypted tunnel) and ship it to the VPN company.
 - The VPN company opens the box.
 - They take out your letter.
 - **Crucial Step:** If they mailed your letter to Google with *your* return address still on it, Google would mail the reply directly to your house (bypassing the VPN). This would break the security/encryption.
 - **The Fix:** The VPN company crosses out your return address and writes **their own** return address.
 - Now, Google mails the reply to the VPN company. The VPN company puts it


- back in the steel box and ships it to you.

Is a VPN the same as a proxy?

No! They both hide the IP from the receiver but they offer very different properties!



Encrypted traffic end-to-proxy
Proxy separates two networks



Encrypted traffic end to end
Acts as one network

Even the last use of a VPN may remind us of a proxy, a proxy and VPN provide very different security properties. They do not protect users in the same threat model.

Proxies are at the boundary between two networks, and rely traffic from one network to the other. In a way, they act as man in the middle and can see what happens in both sides. The traffic is not encrypted end to end, but only end-ro-proxy.

A VPN creates one network! The traffic is encrypted end-to-end. There is no middle-man that can see what happens in both side.

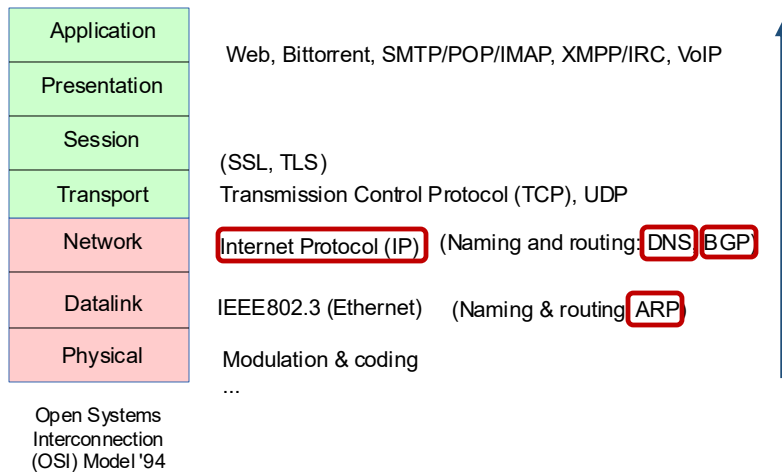


Computer Security (COM-301)
Network security
TCP

Slides by Carmela Troncoso
Some slides/ideas adapted from: George Danezis

40

Where are the problems?



41

Computers communicate with each other at different layers. These are typically modeled by the Open System Interconnection (OSI) Model. This model covers from the physical layer, where pulses that codify bits are sent, to the application layer where programs talk to each other.

In this lecture we will cover protocols at different layers, and see the security problems and potential solutions.

IP limitations

Refresh

- **No reliability:** messages can get dropped, there is no mechanism to ensure a message was received
- **No congestion/flow control:** no mechanism to avoid congestion either in the network or the end hosts
- **No sessions:** no way to associate messages together (and in both directions) into one logical “session”
- **No multiplexing:** no way to associate messages to a network address to specific applications / users on host.

The Transmission Control Protocol (TCP)

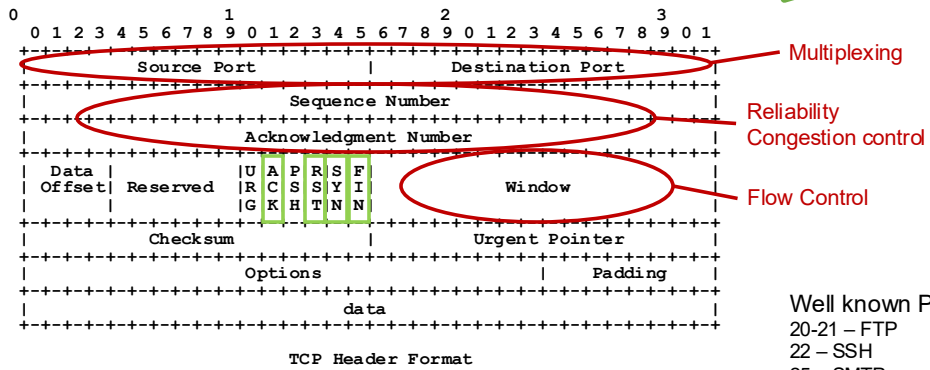
- Protocol run “inside/above” the IP protocol
- Addresses the issues above

42

The goal of IP is to facilitate routing, but it does not feature many other desirable properties to have reliable, robust, efficient communications. Those are provided at the transport layer by the **Transmission Control Protocol (TCP)** that runs just above IP in the OSI model.

TCP header

Refresher



- Well known Ports:
- 20-21 – FTP
 - 22 – SSH
 - 25 – SMTP
 - 53 – DNS
 - 80 – HTTP
 - 110 – POP3
 - 143 – IMAP
 - 443 – HTTPS

RFC793 (1981) <http://www.ietf.org/rfc/rfc793.txt>

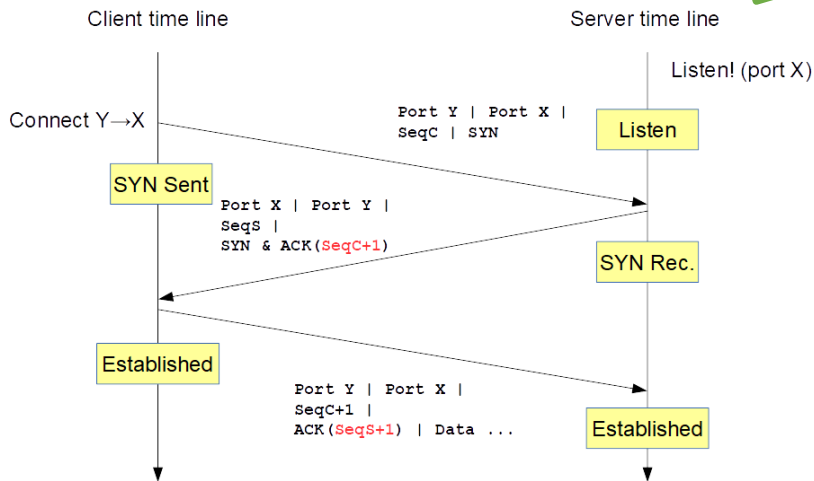
TCP helps hosts establishing and maintaining a connection to exchange a stream of messages. It determines how the data set by the application is divided into packets that can be routed through the network. It manages flows, retransmission, and order.

For the purpose of this lecture we care about:

- Ports: these are used to multiplex applications over a connection.
- Sequence, window, and acknowledgement numbers: these are used to manage the flow number and guarantee error-free transmissions
- Flags: ACK (the packet is an acknowledgements), RST (terminate a connection), SYN (acknowledge the begin of a session), FIN (another way of terminating a connection). [<https://www.geeksforgeeks.org/tcp-flags/>]

TCP 3-way handshake

Refresher



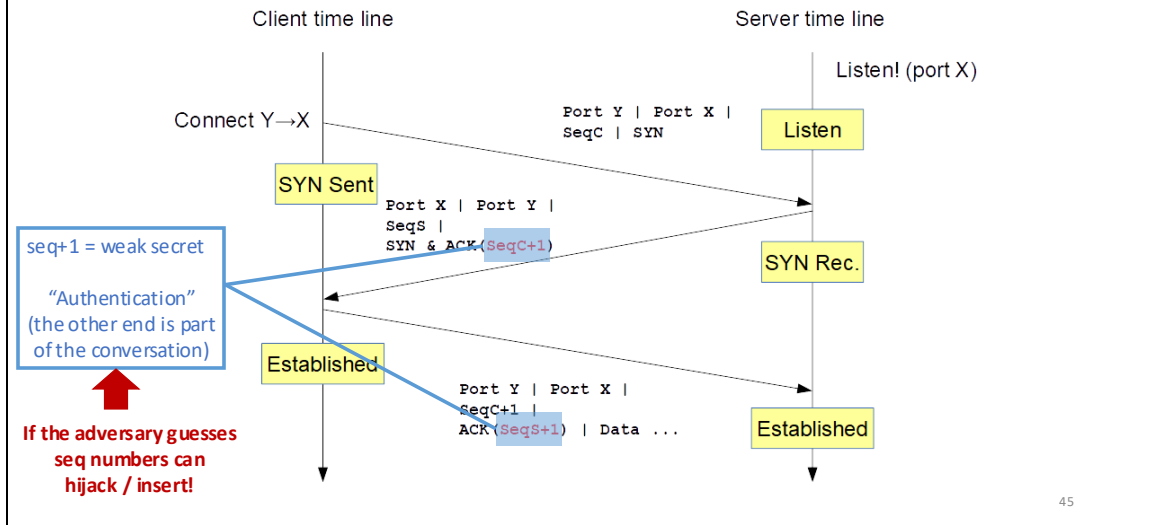
44

Step 1 (SYN) : In the first step, client declares the desire to establish a connection with the server using the SYN flag and sending a sequence number SeqC with which it will start sending segments with

Step 2 (SYN + ACK): the server responds with both SYN-ACK signal flags set. ACK responds the request of the client (using the received sequence number increased by one unit – SeqC+1), and SYN indicating which sequence number will the server user.

Step 3 (ACK) : Finally, the client acknowledges the response of server and they both establish a reliable connection with which they will start the actual data transfer

TCP 3-way handshake – Security considerations



TCP, like the previous protocols does not include any integrity or authentication mechanism. The sequence numbers act as very weak secret that is used as authentication. Both client and server will accept packets as long as they have sequence numbers that correspond to the current session.

TCP 3-way handshake – Security considerations

Can the adversary guess???

- Weak random numbers generation
- Observation (if connection in the clear)

Example attack

- The (historical) “rsh” UNIX utility that provides a remote shell implemented **authentication and authorization on the basis of remote IP address** only! (**Bad idea**)
- The Robert Morris Attack:
 - 1) Send a SYN packet **spoofed** as if it was from authorized host.
 - 2) Guess server SeqS and send an ACK with SeqS+1 and some data.
 - 3) The data is interpreted as a shell command and executed!

<http://www.techrepublic.com/article/tcp-hijacking/>

46

Therefore, if an adversary can guess the sequence numbers, she can hijack the connection and deny service, man in the middle, or modify packets.

How can the adversary guess?

- If the adversary is on path, the adversary can directly observe the sequence number and use this observation
- In some cases, hosts use a weak random number generator (e.g., based on current time), that enables the adversary to make a good guess with high likelihood.

The Robert Morris attack exploits this, together with the fact that IP is easy to spoof to take over remote sessions opened using the “rsh”.