

Network Defenses

KAMI VANIEA
21 JANUARY

Similar statements are found in most content hosting website privacy policies.

What is it about how the internet works that makes this statement necessary to have?

"By submitting, posting or displaying User Content on or through the Service, you grant Piazza a worldwide, non-exclusive, royalty-free license (with the right to sublicense) to use, copy, reproduce, process, adapt, modify, publish, transmit, display and distribute such User Content. This license is limited to providing the Service. Your User Content will not be used for publicity, advertising or any public statements without your prior consent."

— Piazza Terms of Service

First, the news...

- The silencing of KrebsOnSecurity opens a troubling chapter for the Internet
- <http://arstechnica.co.uk/security/2016/09/why-the-silencing-of-krebsonsecurity-opens-a-troubling-chapter-for-the-net/>

Tutorials

- Tutorials start in week 3
- Path is wrong
- The ITO currently has 13 tutorials for this course when they should have 5 tutorials and 5 labs.
- We are trying **really** hard to get this sorted by next week

Courseworks

- There are three courseworks
- Deadlines are on the website
- The deadlines the ITO has are wrong, again we are working to get this resolved

Today

- Open System Interconnect (OSI) model
- Firewalls
- Intrusion detection systems (IDS)
- Time allowing:
 - Network Address Translation (NAT)

Open Systems Interconnect model

A good way to think about networking steps logically
Not how software is actually built

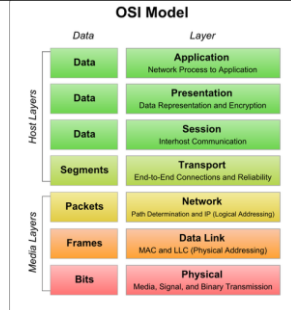
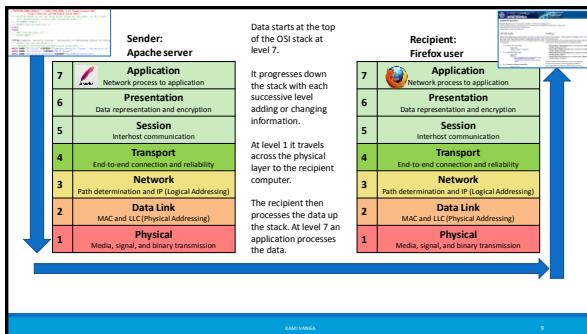




Image from: <http://www.tech-faq.com/osi-model.html>

OSI in terms of debugging errors




- Levels 7 and 6 involve the internal representation of the message
- Levels 5 and 4 involve setting up the connection
- Levels 3, 2, and 1 add header (H) and tail (T) information to each packet


Information is added to the message as it travels down the OSI levels




Header data on a packet



Frame header data on a packet




IP header data on a packet



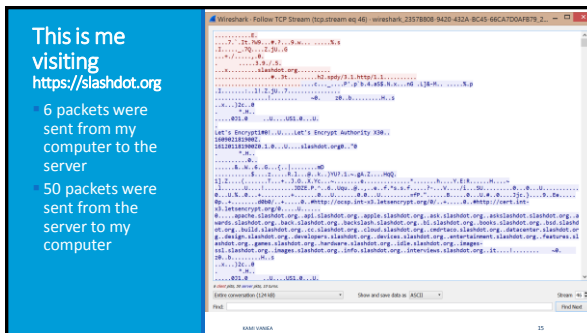
- Levels 7 and 6 involve the internal representation of the message
- Levels 5 and 4 involve setting up the connection
- Levels 3, 2, and 1 add header (H) and tail (T) information to each packet

Information is added to the message
as it travels down the OSI levels



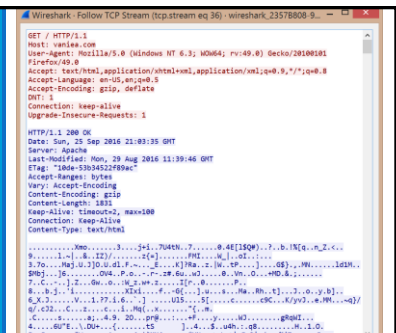
This is me
visiting
<https://slashdot.org>

- 6 packets were sent from my computer to the server
- 50 packets were sent from the server to my computer



This is me
visiting
<http://vaniea.com>

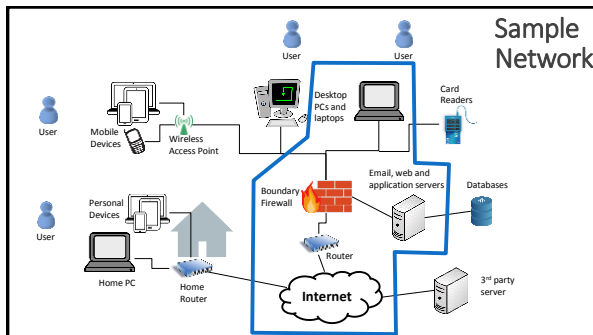
- Note the lack of https
- Why does the text look garbled anyway?



Firewalls

Firewalls

- Firewalls divide the untrusted outside of a network from the more trusted interior of a network
- They run on dedicated devices
 - Less possibilities for compromise – no compilers, linkers, loaders, debuggers, programming libraries, or other tools an attacker might use to escalate their attack
 - Easier to maintain few accounts
 - Physically divide the inside from outside of a network



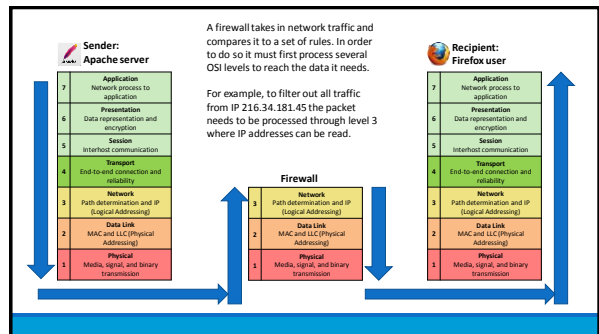
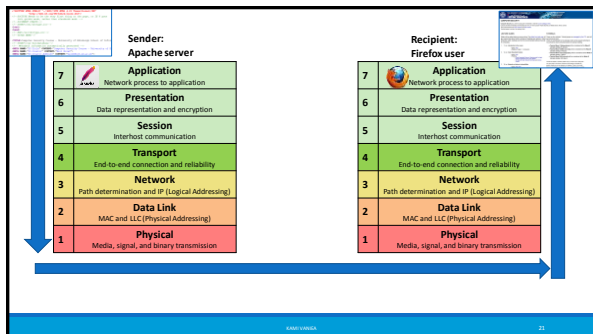
Questionable things come from the Internet AND from the local network

Firewall applies a set of rules

Based on rules, it allows or denies the traffic

Firewalls can also act as routers deciding where to send traffic

Rule	Type	Source Address	Destination Address	Destination Port	Action
1	TCP	*	192.168.1.*	22	Permit
2	UDP	*	192.168.1.*	69	Permit
3	TCP	192.168.1.*	*	80	Permit
4	TCP	*	192.168.1.18	80	Permit
5	UDP	*	192.168.1.*	*	Deny



Firewall ruleset from a custom home router

Taken from an ARSTechnica article

```

root@ars-router: ~
##### Service rules
# OpenVPN
-A INPUT -p udp -n udp --dport 1194 -j ACCEPT

# ssh - drop any IP that tries more than 10 connections per minute
-A INPUT -p tcp -n tcp --dport 22 -n state --state NEW -n recent --set --name DE
FAULT --mask 255.255.255.255 --rsource
-A INPUT -p tcp -n tcp --dport 22 -n state --state NEW -n recent --update --sec
nd -p --hitcount 11 --name DEFAULT --mask 255.255.255.255 --rsource -j LOGGING
-A INPUT -p tcp -n tcp --dport 22 -j ACCEPT

# www - accept from LAN
-A INPUT -i p1p1 -p tcp --dport 80 -j ACCEPT
-A INPUT -i p1p1 -p tcp --dport 443 -j ACCEPT

# DNS - accept from LAN
-A INPUT -i p1p1 -p tcp --dport 53 -j ACCEPT
-A INPUT -i p1p1 -p udp --dport 53 -j ACCEPT

# default drop because I'm awesome
-A INPUT -j DROP

##### Forwarding ruleset
  
```

Image: <http://arstechnica.co.uk/gadgets/2016/01/numbers-dont-lie-its-time-to-build-your-own-router/>

There are many types of Firewalls

Key differences include:

- How implemented
 - Software – slower, easier to deploy on personal computers
 - Hardware – faster, somewhat safer, harder to add in
- Number of OSI levels of processing required
 - Packet size (level 1)
 - MAC (level 2) and IP (level 3) filtering
 - Port filtering (level 3)
 - Deep packet (level 4+)

Today we will talk about:

- Packet filtering gateway
- Stateful inspection firewall
- Application proxy
- Personal firewalls

Packet filtering gateway or screening router

- Simplest – compares information found in the headers to the policy rules
- Operate at OSI level 3
- Source addresses and ports can be forged, which a packet filter cannot detect
- Design is simple, but tons of rules are needed, so it is challenging to maintain

KIM VANDER

25

Stateful inspection firewall

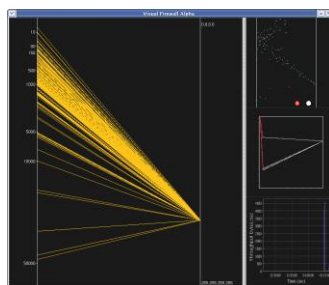
- Maintains state from one packet to another
- Similar to a packet filtering gateway, but can remember recent events
- For example, if a outside host starts sending packets to many internal destination ports (aka a port scan) a stateful firewall would record the number of ports probed and once it is over the threshold specified in the policy it would block all further traffic

KIM VANDER

26

Port scan

- An attacker is looking for applications listening on ports
- A single IP address (right) is contacting many ports (left) to see if any respond

Image: <http://chrislee.dhs.org/projects/visualfirewall.html>

27

Firewall ruleset from a custom home router

- Taken from an ARSTechnica article

```
root@ars-roulen:~#
##### Service rules
# OpenVPN
-A INPUT -p udp -m udp --dport 1194 -j ACCEPT

# ssh - drop any IP that tries more than 10 connections per minute
-A INPUT -p tcp -m tcp --dport 22 -m state --state NEW -n recent --set --name D1
-FULL --mask 255.255.255.255 --rsource
-A INPUT -p tcp -m tcp --dport 22 -m state --state NEW -n recent --update --seconds 60 --hitcount 10 --name DDEFAULT --mask 255.255.255.255 --rsource -j LOGDROP
-A INPUT -p tcp -m tcp --dport 22 -j ACCEPT

# www - accept from LAN
-A INPUT -i p1p1 -p tcp -m tcp --dport 80 -j ACCEPT
-A INPUT -i p1p1 -p tcp -m tcp --dport 443 -j ACCEPT

# DNS - accept from LAN
-A INPUT -i p1p1 -p tcp --dport 53 -j ACCEPT
-A INPUT -i p1p1 -p udp --dport 53 -j ACCEPT

# default drop because I'm awesome
-A INPUT -j DROP

##### forwarding ruleset
```

Image: <http://arstechnica.co.uk/gadgets/2016/01/numbers-dont-lie-its-time-to-build-your-own-router/>

28

Application proxy

- Simulates the (proper) effects of an application at OSI level 7
- Effectively a protective Man In The Middle that screens information at an application layer (OSI 7)
- Allows an administrator to block certain application requests.
- For example:
 - Block all web traffic containing certain words
 - Remove all macros from Microsoft Word files in email
 - Prevent anything that looks like a credit card number from leaving a database

KIM VANDER

29

Personal firewalls

- Runs on the workstation that it protects (software)
- Provides basic protection, especially for home or mobile devices
- Malicious software can disable part or all of the firewall
- Any rootkit type software can disable the firewall

KIM VANDER

30

Think-pair-share

Imagine you want to put a firewall in front of the email server

- Why is deep packet inspection easier to do on email than on normal network traffic?
- As a malicious actor, how might I go around your email firewall?

Intrusion Detection Systems (IDS)



Firewalls are preventative, IDS detects a potential incident in progress

- At some point you have to let some traffic into and out of your network (otherwise users get upset)
- Most security incidents are caused by a user letting something into the network that is malicious, or by being an insider threat themselves
- These cannot be prevented or anticipated in advance
- The next step is to identify that something bad is happening quickly so you can address it

Signature based

- Perform simple pattern matching and report situations that match the pattern
- Requires that admin anticipate attack patterns in advance
- Attacker may test attack on common signatures
- Impossible to detect a new type of attack
- High accuracy, low false positives

Heuristic based

- Dynamically build a model of acceptable or "normal" behavior and flag anything that does not match
- Admin does not need to anticipate potential attacks
- System needs time to warm up to new behavior
- Can detect new types of attacks
- Higher false positives, lower accuracy

Number of alarms is a big problem

- In the Target breach the IDS did correctly identify that there was an attack on the Target network
- There were too many alarms going off to investigate all of them in great depth
- Some cyberattack insurance policies state that if you know about an attack and do nothing they will not cover the attack.
- Having a noisy IDS can potentially be a liability

Network Address Translation (NAT)

IPv4

- Version 4 of the Internet Protocol
 - 192.168.2.6
- There are less than 4.3 billion IPv4 addresses available
- We do not have enough addresses for every device on the planet
- Answer: Network Address Translation
 - Internal IP different than external IP
 - Border router maps between its own IP and the internal ones

Questions