U.S. National Cybersecurity

Understanding the Internet

William J. Perry
Martin Casado · Keith Coleman · Dan Wendlandt

MS&E 91SI
Spring 2004
Stanford University

Announcements

- Axess
- Forum
- Bios/Photos
- Law School Event

Goal: Provide Working Knowledge of the Internet

(as it relates to this class!)

The Internet is …

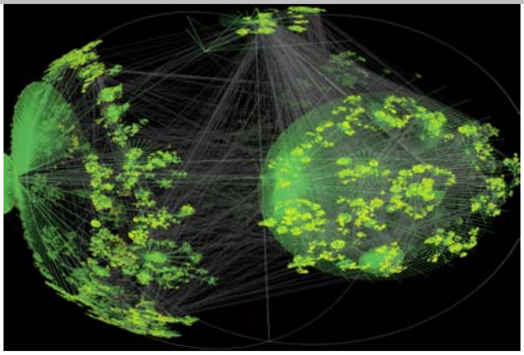

And …

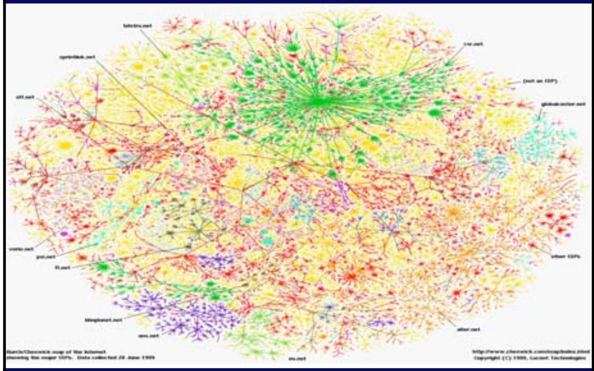

Burch/Cheswick map of the Internet
showing the major ISPs. Data collected 29 June 1999

http://www.cheswick.com/ches/map/index.html
Copyright (C) 1999, Lucent Technologies

At Present:

- More than **500** million computers
- **287.5** Million English Users
- **516.7** Million Non English Users
- Over **38** million active domains
- So … how does it all work?

First: A Story

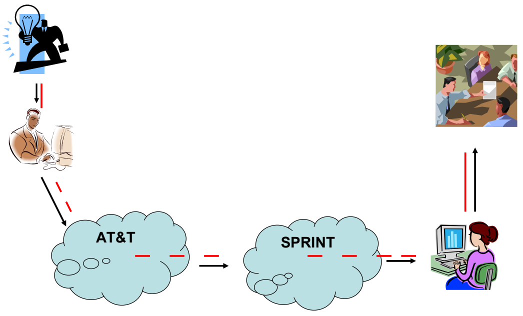

A Network Is ...

- Computers
- Wires Connecting computers

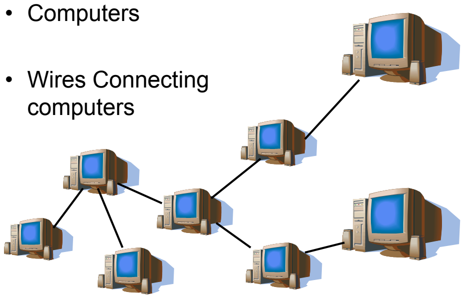

“Internet” Really a Network of Networks (ISPs)

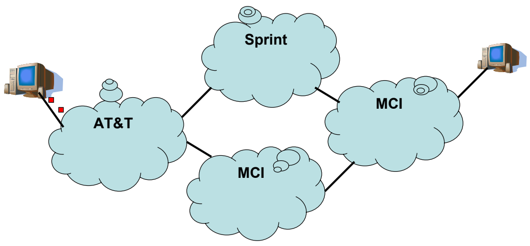

To Complex for my Brain.. (and not really modular)

Organized Into Hierarchy by Function

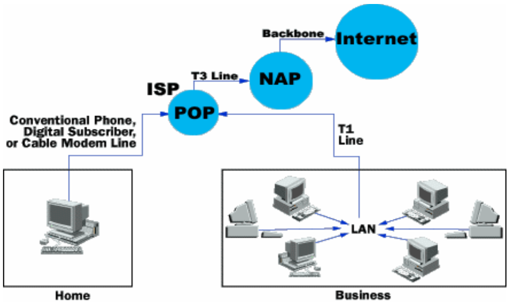

Separated into “Layers”

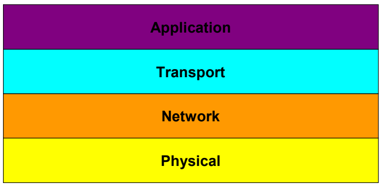

Application

Transport

Network

Physical

The Physical Layer

The Physical Layer

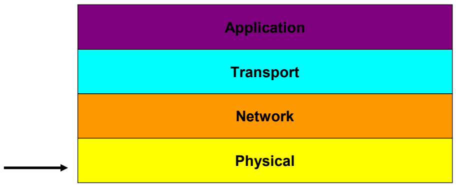

Physical Layer

The hardware that makes up the Internet

- Anything you can “touch”
- The physical computers
- The wires connecting computers

Physical Layer

- Computers are physically located in some-ones jurisdiction (whose? implications?)
- Must be physically protected (destroyed computers don’t work, nor to clipped wires)
- One wire, can carry a lot of data … better to use less? (2 fiber lines across Rockies)
- Hard limitation (about 5 ways in and out of US)
- Often overlooked, though a serious component in security!!!

The Network Layer

Network Layer

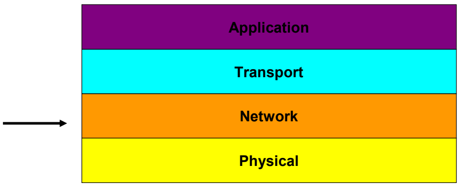

IP Addresses

Every reachable computer (not really) on the Internet is given a unique identifier called an IP address

171.67.71.18

IP Packets

Information leaving computers is broken into discrete segments or packets, marked with the IP address of the destination and the IP address of the source.

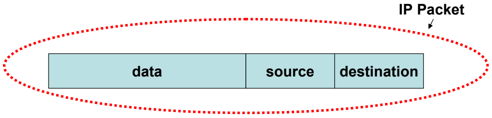

Routing

Computers on the Internet that lie between sending and receiving computers (called “routers”) forward received packets to connecting computers.

The choice of wire to send a packet out of is based solely on the destination of the packet.

Packet Flow on Internet

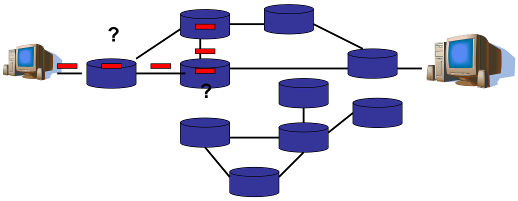

Network Layer : (Overview)

- Computers are identified by globally unique address (32 bits) called and IP address (note: this is somewhat of a white lie)
- Data is broken into small “chunks” called packets
- Packets flow between computers over specialized computers. “routers”
- Each router makes its own decision where to send a packet
- Routers ONLY make the decision via the packets destination

A Look at ISPs

- Carry their customers traffic to anywhere in the globe
- What kind of power does an ISP have?
- What factors determine routing decisions?
- How can ISPs trust each other?
- Why would an ISP want to limit attacks on the network?
- What is the potential damage of a rogue ISP?

A Quick Digression: Domain Name System

Domain Name System

- When trying to contact a computer (www.google.com) do not use IP addresses …
- Instead use DNS … converts “names” (can remember) to IP addresses (cannot remember)

Domain Name System

- Convert Name (www.foo.com) to 32 bit value (134.114.223.91)
- Must ask special machine (name server) for answer (problems here?)
- Has good and bad properties! (as we will see!)

The Transport Layer

Transport Layer

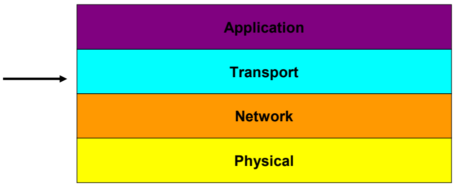

Transport Layer

Uses “IP packets” to send information from computer A to computer B

e.g. TCP

A Reliable method of sending information to someone

“hi there mom → “hi there mom”

TCP

- 99% of Internet Traffic
- Must set up a connection before hand

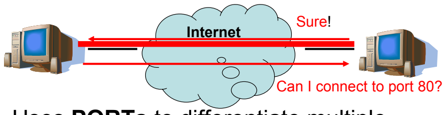

Internet

Sure!

Can I connect to port 80?

- Uses **PORTs** to differentiate multiple connections per machine
- Once established can be reasonably assured you are talking to a real machine!

(http://www.ja.net/CERT/Morris/r.t.morris-TCP.html)

Transport Layer

- ICMP, UDP : Other methods of sending information
- Not “seen” by normal users
- Nuts and bolts are good for understanding attacks and vulnerabilities

The Application Layer
(finally some familiarity)

Application Layer

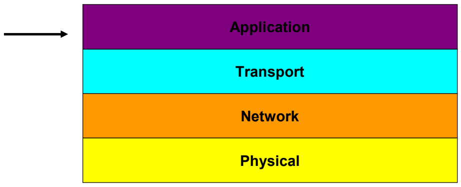

Application

Transport

Network

Physical

The Application Layer

- Email
- World Wide Web
- SSH
- Telnet
- FTP
- Applications we love and use every day!

The Application Layer

Applications use transport layers (such as TCP) to communicate across the Internet

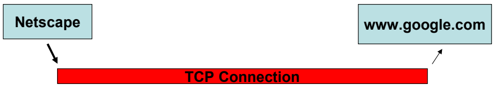

Netscape

www.google.com

TCP Connection

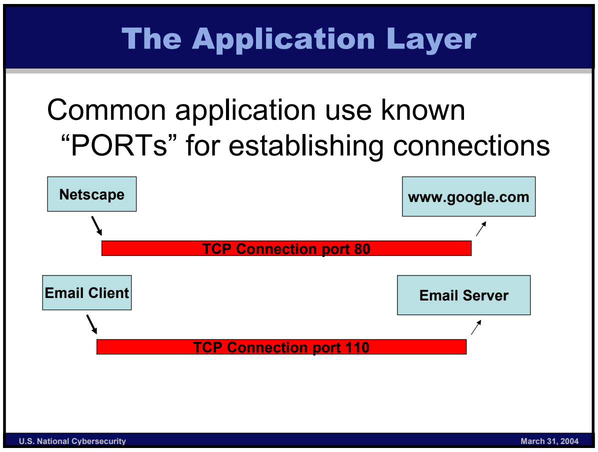

The Application Layer

Common application use known "PORTs" for establishing connections

Netscape

www.google.com

TCP Connection port 80

Email Client

Email Server

TCP Connection port 110

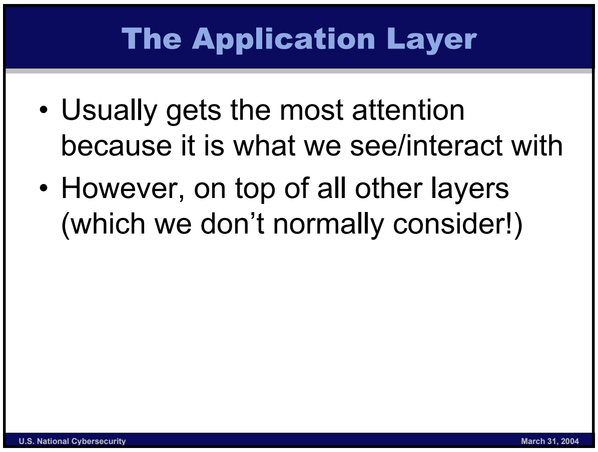

The Application Layer

- Usually gets the most attention because it is what we see/interact with
- However, on top of all other layers (which we don't normally consider!)

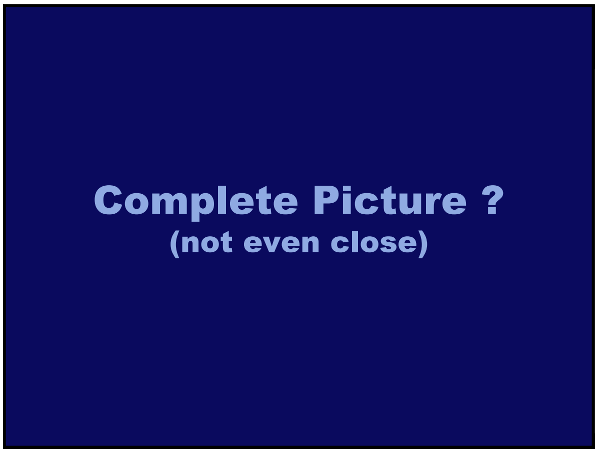

Complete Picture ?

(not even close)

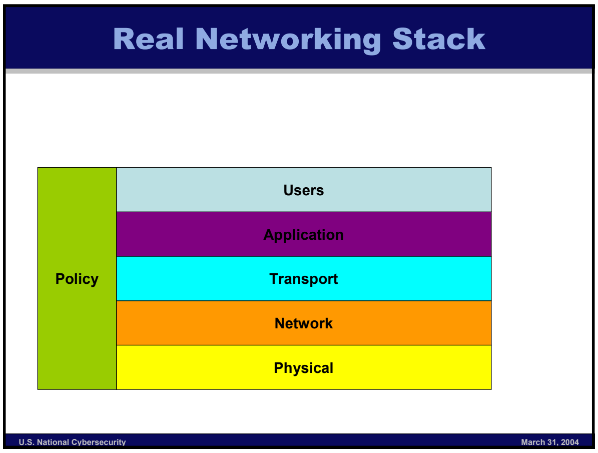

Real Networking Stack

Policy

Users

Application

Transport

Network

Physical

Why do we need to know this technology ?

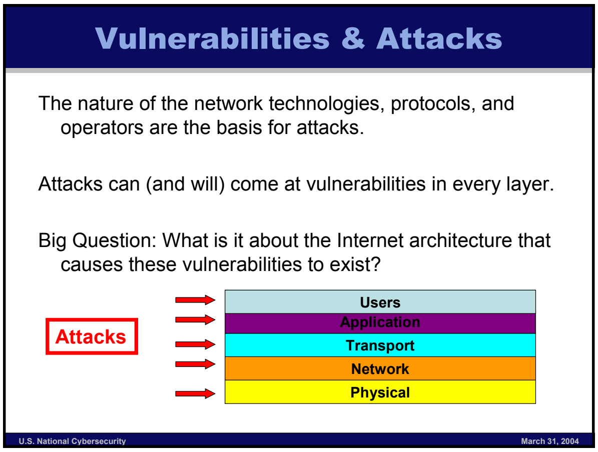

Vulnerabilities & Attacks

The nature of the network technologies, protocols, and operators are the basis for attacks.

Attacks can (and will) come at vulnerabilities in every layer.

Big Question: What is it about the Internet architecture that causes these vulnerabilities to exist?

Attacks

Users

Application

Transport

Network

Physical

Definitions

In cybersecurity:

def. vulnerability (n):
any avenue for attack.

def. attack (n)
Any action that without authorization exposes, modifies, utilizes or denies the availability of an Internet related resource.

Attacks on the Internet

Why do attacks matter?

Attacks affect the Internet's ability to function as a reliable and secure critical infrastructure.

Scanning & Fingerprinting

What is it?

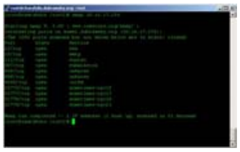

Reconnaissance technique to explore networks, classify + analyze connected hosts, and identify potential vulnerabilities.

Example: nmap security scanner

Exploits

What is it?

The use of vulnerabilities in or misconfiguration of software or hardware to gain access to information or resources on a system. Exploits may be manual or automated.

example: Blaster worm exploits RPC bug

Denial of Service

What is it?

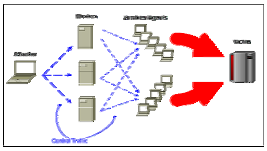

The malicious consumption of resources in order to make a system incapable of fulfilling its designed role. Attacks are often "distributed" to increase resource consumption.

example: SYN flood against Yahoo

Social Engineering Attack

What is it?

Any attempt that employs non-technical means to attack a system. Often the attacker uses information gleaned from outside sources to produce false credentials. Attacks are often hybrid, relying on human and technical factors.

example: Beagle virus used email domain name to pose as a message from the user's ISP.

Infrastructure Attack

What is it?

An attack against the core systems that operate as the Internet infrastructure. Attacks can be either physical or virtual, often focusing on central points of failure.

example: Attack on root DNS servers.

Sniffing Traffic

What is it?

Using access to a link or infrastructure system to examine the contents of Internet traffic. Similar to a phone tap.

example: ISP's potential for information gathering

Why is this Interesting?

The existence of each of these attacks point out a number of fundamental issues with the Internet that are potential problems for its use as a critical infrastructure.

Considering these issues will be the beginning for us to develop a framework and understanding for key points that will apply across many of the guest lectures.

Discussion Questions

Attributability

For traffic on the Internet, can we determine who a packet come from?

Two levels:

- Can we tell what computer sent a given packet? (what are the implications of source spoofing?)
- Can we attribute a packet to a human?

- What does this say about our ability to catch and prosecute perpetrators of online attacks? What about active response?

Determining Intent

Can you infer intent from analyzing network traffic? What about at the application level?

- What is the different between a denial of service attack and normal overwhelming usage?
- What is more important, the intent or the result of Internet traffic?
- What about 'enablement' versus 'use'?

Information Containment

Is it possible to contain information on the Internet?

- What can someone ‘on route’ potentially do? How can you trust the integrity of what you see?
- What are the dangers of arbitrary routes between points A and B?
- Lifetime of data; can data be provably destroyed?
- What about online caching engines (Google?) Security sensitive documents have been posted on the web and removed.
- Can we tell if sensitive information is being leaked? (what is the “entropy bandwidth” of the Internet?

Infrastructure Attacks

How vulnerable is the actual Internet infrastructure to attacks?

- Could a single group bring down the Internet? What kind of resources would it take?
- How reliant is the Internet on a relatively few critical systems?
- What happens when you rely on the security of infrastructure that you have absolutely no control over? As a company? As a country?

Determining Identity

How can we trust an Internet entity is who they say they are?

- Why is this process more difficult than it is in the “brick & mortar” world?
- How important is this for a critical infrastructure?
- Do our solutions for providing identity scale to the millions of actions on the Internet?

Overwhelming Complexity

What does the extreme complexity of the Internet mean for our ability to secure it?

- Are there just too many things that could go wrong to ever possibly be able to completely rely on it?
- In what way does the complexity impact our ability to educate average users? Is user education necessary? Is effective user education even possible?
- Will the Internet become more or less complex to manage in the future?

Why is this so hard?

What are the major barriers to providing perfect security for a system on the Internet?

- What are the weak links for security systems?
- Can we ever really secure a usable Internet computer system? (e.g. directed attack)
- If we know what the major vulnerabilities are, why is Internet security in the state it is today?

Monoculture

Much of the Internet operates on the same software and hardware, what does this mean for security?

- What are the advantages of monoculture?
- What are the drawbacks?
- Do we even have a choice about this, or is this just an uncontrollable parameter of the system?