Internet Protocol Security CS155 Computer and Network Security

Stanford University



What is the Internet?

Global network that lets hosts communicate Internet provides best-effort delivery of *packets* between hosts **Packet**: a structured sequence of bytes Header: metadata used by network **Payload:** user data to be transported Packets are forwarded by routers from sender to destination host

Internet Protocol (IP)

look like to be processed by routers

Every host is assigned a unique identifier ("IP Address")

Every packet has an IP header that indicates its sender and receiver

Routers forward packet along to try to get it to the destination host

Rest of the packet should be ignored by the router

- **Internet Protocol (IP)** defines what packets that cross the Internet need to

IP Addresses

IPv4: 32-bit host addresses Written as 4 bytes in form A.B.C.D where A,...,D are 8 bit integers in decimal (called dotted quad) e.g. 192.168.1.1

IPv6: 128 bit host addresses Written as 16 bytes in form AA:BB::XX:YY:ZZ where AA,...,ZZ are 16 bit integers in hexadecimal and :: implies zero bytes *e.g.* 2620:0:e00:b::53 = 2620:0:e00:b:0:0:53

Instruct routers and hosts what to do with a packet All values are filled in by the sending host



IPv4 Header

Destination Address

Sender sets destination address Routers try to forward packet to that address



Source Address

Source Address (sender) Sender fills in. Routers due not verify.



Checksum

16-bit Simple Header checksum (filled in by sender)



IP Security

Client is trusted to embed correct source IP

- Easy to override using lower level network sockets
- Libnet: a library for formatting raw packets with arbitrary IP headers

- Denial of Service Attacks
- Anonymous infection (if one packet)

Anyone who owns their machine can send packets with arbitrary source IP

Internet Protocol (IP)

Yes:

Routing. If host knows IP of destination host, route packet to it.

No:

- Fragmentation and reassembly: Split data into packets and reassemble
- Error Reporting: (maybe, if you're lucky) tell source it dropped your packet

Everything else. No ordering. No retransmission. No (real) error checking. No acknowledgement of receipt. No "connections". No security. Just packets.







Networks use a stack of layers Lower layers provide services to layers above Don't care what higher layers do Higher layers use services of layers below Don't care how lower layers implement services Layers define abstraction boundaries At a given layer, all layers above and below are opaque







Packet Encapsulation

Protocol N1 can use the services of lower layer protocol N2 A packet P1 of N1 is encapsulated into a packet P2 of N2 The payload of p2 is p1 The control information of p2 is derived from that of p1



Link Layer

Model assumed that hosts can deliver and accept packets from Internet routers

In practice, hosts not connected directly to router

Link layer provides connectivity between hosts and routers

Most common Link Layer Protocol. Let's you send packets to other local hosts.



At layer 2 (link layer) packets are called *frames*

MAC addresses: 6 bytes, universally unique

Ethernet

EtherType gives layer 3 protocol in payload 0x0800: IPv4 0x0806: ARP 0x86DD: IPv6



Originally broadcast. Every local computer got every packet.



Ethernet

Switched Ethernet

address lives based on MAC source addresses

If switch knows MAC address M is at port P, it will only send a packet for M out port P

to all ports

- With switched Ethernet, the switch *learns* at which physical port each MAC
- If switch does not know which port MAC address M lives at, will broadcast

[zakir@scratch-01:~\$ ifconfig ens160: flags=4163<UP, BROADCAST, RUNNING, MULTICAST> mtu 1500 ether 00:50:56:86:b2:03 txqueuelen 1000 (Ethernet) RX packets 1404151714 bytes 1784388363701 (1.7 TB) RX errors 0 dropped 73 overruns 0 frame 0 TX packets 1155689210 bytes 6010503085464 (6.0 TB)

Ethernet

inet 10.216.2.64 netmask 255.255.192.0 broadcast 10.216.63.255 inet6 fe80::250:56ff:fe86:b203 prefixlen 64 scopeid 0x20<link> TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

Two Problems

Local: How does a host know what MAC address their destination has?

Internet: How does each router know where to send each packet next?

ARP: Address Resolution Protocol

ARP is a Network protocol that lets hosts map IP addresses to MAC addresses

Host who needs MAC address M corresponding to IP address N broadcasts an ARP packet to LAN asking, "who has IP address N?"

Host that has IP address N will reply, "IP N is at MAC address M."

Destination address	Source address (Type ARF	ARP Request or ARP Reply Padding				
6	6	2	28 10				
H	lardware t	ype (2 bytes)	Protocol type	(2 bytes)			
Hardware a length (1	address byte)	Protocol address length (1 byte)	Operation code (2 bytes)				
		Source hard	ware address*				
		Source prote	ocol address*				
		Target hardw	vare address*				
		Target proto	col address*				

* Note: The length of the address fields is determined by the corresponding address length fields

ARP Packet

ARP Security

to be another host on the local network! This is called ARP spoofing

to flow through X (*MitM!*)

Claim N_A is at attacker's MAC address M_X

Claim N_B is at attacker's MAC address M_X

Re-send traffic addressed to N_A to M_A , and vice versa

Any host on the LAN can send ARP requests and replies: any host can claim

- This allows any host X to force IP traffic between any two other hosts A and B

Routing (BGP)

to exchange information about their routing tables

Each router announces what it can route to all of its neighbors.

Every router maintains a global table of routes

- BGP (Border Gateway Protocol): protocol that allows routers

Pakistan hijacks YouTube

network

advertisement for 208.65.153.0/24

Youtube offline.

On 24 February 2008, Pakistan Telecom (AS 17557) began advertising a small part of YouTube's (AS 36561) assigned

PCCW (3491) did not validate Pakistan Telecom's (17557)

From Packets to Streams

Most Internet applications want a data stream abstraction (not best-effort packets)

Application on host X wants to send a sequence of bytes to application on host Y. Wants reliable, in-order delivery of data

Transmission Control Protocol (TCP) provides a data stream abstraction using a best-effort packet transport (IP)

Have: network that will deliver packets Packets may be dropped, re-ordered, duplicated

Want: Bytes delivered reliably and in-order

Abstraction of a stream of bytes between applications on different hosts

Each application is identified by a *port number*

TCP connection established between port A on host X to port B on host Y Ports are 1–65535 (16 bits)

Some destination port numbers used for specific applications by convention

Ports

Port		Applic
8	0	HTTP
44;	3	HTTP
2	5	SMTP
6	7	DHCP
22	2	SSH (s
514	4	RSH (I
23	3	Telnet

Ports

- cation
- (Web)
- S (Web)
- (mail)
- (host config)
- secure shell)
- remote shell)

Bytes in application data stream numbered with a 32-bit sequence number

in a single IP datagram

There are two logical data streams in a TCP session, one in each direction

Data sent in segments: sequences of contiguous bytes sent

Sequence number in packet header is seq. number of first byte of payload Acknowledgement number is seq number of next expected byte of stream in opposite direction

- Sender sends 3 byte segment
- Sequence number indicates where data belongs in byte sequence (at byte 401)
 - Note: Wireshark shows relative sequence numbers

TCP ACKs

Receiver acknowledges received data

 Sets ACK flag in TCP header
 Sets acknowledgement number to indicate next expected byte in sequence

receiving acknowledgement

Sender may send several segments before

- Sender may send several segments before receiving acknowledgement
- Receiver always acknowledges with seq. no. of next expected byte

What if the first packet is dropped in network? Receiver always acknowledges with seq. no. of next expected byte

- next expected byte
- Sender retransmits lost segment

 What if the first packet is dropped in network? Receiver always acknowledges with seq. no. of

- What if the first packet is dropped in network?
- Sender retransmits lost segment
- Receiver always acknowledges with seq. no. of next expected byte

- ACKs may be piggybacked on data flowing in opposite direction or sent without data
- All packets after initial connection setup will carry an acknowledgement
- Sequence numbers wrap around:
 ..., 2³²-2, 2³²-1, 0, 1, 2, ...

Client

State changes to SYN-SENT

State changes to ESTABLISHED

SYN-ACK seq: 200 ack:101

Starting a Connection

Sends packet with FIN flag set Must have ACK flag with valid seqnum

Peer receiving FIN packet acknowledges receipt of FIN packet with ACK FIN "consumes" one byte of seq. number

Eventually other side sends packet with FIN flag set: This terminates the TCP session

Ending a Connection

TCP Connection Reset

previous connections)

If a connection exists, it is torn down Packet with RST flag sent in response

TCP designed to handle possibility of spurious TCP packets (e.g. from

- Packets that are invalid given current state of session generate a reset
- If a host receives a TCP packet with RST flag, it tears down the connection

TCP Connection Spoofing

Can we impersonate another host when *initiating* a connection?

Off-path attacker can send initial SYN to server but cannot complete three-way handshake without seeing the server's sequence number

1 in 2³² chance to guess right if initial sequence number chosen uniformly at random

TCP Reset Attack

Can we reset an *existing* TCP connection?

Need to know port numbers (16 bits) Initiator's port number usually chosen random by OS Responder's port number may be well-known port of service

There is leeway in sequence numbers B will accept Must be within window size (32-64K on most modern OSes)

1 in 2¹⁶⁺³²/W (where W is window size) chance to guess right

UDP (User Datagram Protocol)

Sometimes we do only want best-effort delivery

wrapper around IP

Adds ports to demultiplex traffic by applications

header (parts of IP header)

- User Datagram Protocol (UDP) is a transport layer protocol that is essentially a

Application-layer protocols (and people) usually refer to Internet host by host name (e.g., google.com)

DNS is a delegatable, hierarchical name space

DNS

A DNS server has a set of records it authoritatively knows about

\$ dig bob.ucsd.edu

;; Got answer: ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 30439</pre> ;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 6 ;; QUESTION SECTION: ;bob.ucsd.edu. INA ;; ANSWER SECTION: bob.ucsd.edu. 3600 INA 132.239.80.176 ;; AUTHORITY SECTION: ucsd.edu. 3600 INNS ns0.ucsd.edu. ucsd.edu. 3600 INNSns1.ucsd.edu. ucsd.edu. 3600 IN NS ns2.ucsd.edu.

DNS Record

DNS Root Name Servers

In total, there are 13 main **DNS root servers**, each of which is named with the letters 'A' to 'M'.

HOSTNAME	IP ADDRESSES	MANAGER
a.root-servers.net	198.41.0.4, 2001:503:ba3e::2:30	VeriSign, Inc.
b.root-servers.net	199.9.14.201, 2001:500:200::b	University of Southern California (ISI)
c.root-servers.net	192.33.4.12, 2001:500:2::c	Cogent Communications
d.root-servers.net	199.7.91.13, 2001:500:2d::d	University of Maryland
e.root-servers.net	192.203.230.10, 2001:500:a8::e	NASA (Ames Research Center)
f.root-servers.net	192.5.5.241, 2001:500:2f::f	Internet Systems Consortium, Inc.
g.root-servers.net	192.112.36.4, 2001:500:12::d0d	US Department of Defense (NIC)
h.root-servers.net	198.97.190.53, 2001:500:1::53	US Army (Research Lab)
i.root-servers.net	192.36.148.17, 2001:7fe::53	Netnod
j.root-servers.net	192.58.128.30, 2001:503:c27::2:30	VeriSign, Inc.
k.root-servers.net	193.0.14.129, 2001:7fd::1	RIPE NCC
l.root-servers.net	199.7.83.42, 2001:500:9f::42	ICANN
m.root-servers.net	202.12.27.33, 2001:dc3::35	WIDE Project

Caching

- DNS responses are cached Quick response for repeated translations NS records for domains also cached
- DNS negative queries are cached Save time for nonexistent sites, e.g. misspelling
- Cached data periodically times out Lifetime (TTL) of data controlled by owner of data TTL passed with every record

DNS requests sent over UDP

Four sections: questions, answers, authority, additional records

Query ID: 16 bit random value Links response to query

DNS Packet

Response

Authoritative Response

	←	32 bits					
	Ve	T	hlen	TOS		pk	
		identification			flg frag		
IΡ		TTL		protocol	head		
	src IP = 64.170.162						
	dst IP = 68.94.156						
Р		<pre>src port = 53</pre>			dst po		
IJ		UDP length			UDI		
\langle	QID = 43562			1 Op	•0 1 C		
	Question count = 1			Answer			
	1 A	Authority count = 2				ll. Re	
	Qu	Wł	nat is	A record for	r www	.unixw	
$\langle \cdot \rangle$	An	www.unixwiz.net A = 8.7.25.9				25.94	
	Au	unixwiz.net NS = linux.u				nixwiz	
	Au	unixwiz.net NS = cs.unixwiz.r				wiz.ne	
	Ad	linux.unixwiz.net A = 64.170.				.170.1	
	Ad	cs.unixwiz.net A = 8.7.25.					

DNS Security

- Users/hosts trust the host-address mapping provided by DNS Used as basis for many security policies: Browser same origin policy, URL address bar
- Interception of requests or compromise of DNS servers can result in incorrect or malicious responses
- DNSSEC Fixes, but nobody uses. Use TLS!!

DNS Spoofing

Scenario: DNS client issues query to server

Attacker would like to inject a fake reply Attacker does not see query or real response

How does client authenticate response?

DNS Spoofing

How does client authenticate response?

UDP port numbers must match Destination port usually port 53 by convention

16-bit query ID must match

DNS Cache Poisoning

resolution process for each query

Lifetime of record determined by record TTL Could also be evicted from cache due to limited memory

cache poisoning No protocol-defined way for to refresh cached record

- Recursive resolvers cache records to avoid repeating recursive
- Injecting spoofed records into a resolver's cache is called DNS

Early Attack Strategy

Root

4. The BadGuysAreUs name server responds with an IP address but adds a false IP address for a completely different Web site, www.paypal.com.

Cache

5. The targeted name server stores the false IP address for paypal.com.

6. When people using this name server attempt to go to www.paypal. com, they are directed to a Web site that looks like PayPal's but works only to harvest their user names and passwords.

Kaminsky Attack

Try Again! Victim machine visits attacker's web site, downloads Javascript a.bank.com $QID=x_1$.com response 256 responses: Random QID $y_1, y_2, ...$ **NS** bank.com=ns.bank.com A ns.bank.com=attackerIP attacker

Defenses

Increase QueryID. But how? Don't want to change packet. Randomize src port, additional 11 bits - Now attack takes several hours

The network is out to get you.

Conclusion