CS 55: Security and Privacy

Packet sniffing and spoofing

Adapted from Computer and Internet Security by Du unless otherwise noted





1. Network basics

2. Sending and receiving packets

- 3. Sniffing packets
- 4. Spoofing packets

The TCP/IP protocol stack uses layers to transmit data over unreliable networks

Conceptual network layers

	Interacts with application programs	the "protocol stack
Application	Interpretation of data is outside the scope of this laver	OSI model has 7
	Examples: HTTP, SSH, FTP, SMTP	layers (application
	Provides host-to-host message delivery in message pa	called layer 7) ckets
Transport	May provide error control, flow control, application ad Examples: TCP (connection-oriented), UDP (connection	ldressing (ports) nless)
	TCP provides sequencing, dropped packet resend, traf	fic congestion routing
Network (Internet)	Moves packets (datagrams) across local area network Each computer identified by an IP address (IP v4 or v6) Also called Laver 3 or IP laver (ICMP is here)	boundaries (routing))
,	Also called Layer 5 of it layer. (ICIVIT is field)	
Link (Data Link)	Moves frames within a local area network (switching)	
	Also called Layer 2, MAC layer, or Ethernet layer	
	How data is physically transmitted	
Dhysical	 Transmitter converts logical 1's and 0's to electrical 	/light pulses or

phase/amplitude of radio frequency (RF) and sends down wire or over air

I'll refer to this as

•htm Receiver converts electrical/light or RF back to logical 1's and 0's

Adapted from https://www.guru99.com/tcp-ip-model.htm

Messages travel down the stack from the transmitter and up the stack at the receiver



Adapted from http://www.tcpipguide.com/free/t_IPDatagramEncapsulation.htm

Small network example shows how components work together



https://www.igcseict.info/theory/4/netsetup/index.html

At the Transport Layer, TCP is connectionoriented, UDP is not

Transmission Control Protocol (TCP)



TCP establishes a connection between machines with a 3-way handshake

User Datagram Protocol (UDP)

UDP does not establish a connection



- Block 3 retransmitted
- Messages are broken into packets
- Each packet given sequence number
- Packets reassembled in seq order
- Missing packets are resent

Adapted from http://www.steves-internet-guide.com/tcp-vs-udp/



UDP Transmission Illustration

- Messages are broken into packets
- Packet received or not, no resend
- Faster than TCP, smaller headers
- Use for streaming

IP addresses contain the network and device ID

Classless Interdomain Routing (CIDR, 1993+)

IP address is 32 bits (4 bytes)



Split wherever we'd like

- Specify how many bits make up the network ID, the rest of the address specifies a device on that network
- 128.230.115.0/24 means left 24 bits are network ID (8 bits are used to identify a device on that network, ranging from 0 to 255)
- Split does not have to occur on 8-bit boundaries
- 128.230.0.0/9 means left most 9 bits are the network ID

Some IP addresses are reserved for private use

Private IP address

10.0.0.0/8 172.16.0.0/12 192.168.0.0/16 ISP issues a single IP address Use private network for devices in home

Non-routable on Internet (multiple homes may use the same address inside)

Lookback address

127.0.0/8 Commonly used 127.0.0.1 (localhost)

Network Address Translation (NAT) handles translation between inside and outside



The *ping* command is useful to see if a computer is reachable (and on homework)

Issue command \$ ping 129.170.171.44



Domain Name System (DNS) converts between IP addresses and names

Name	IP
www.google.com	172.217.3.110
www.dartmouth.edu	129.170.171.44

Humans are not good at remembering numeric IP addresses

Give a computers a name humans can remember

DNS matches name with IP address

Demo: Wireshark shows the protocol stack in action

	800	*enp0s3															
					3	Q 🤇	>	٦ (1							
	📕 Apply	/ a display f	ilter <	Ctrl-/>												 Expression. 	+
	No.	Time				Sou	Irce			De	stinati	on	F	rotocol	Length Ir	nfo	<u> </u>
		2 2021	02-17	04:2	29:39.	2 10	.0.2.	15	_	75	5.75	75.7	5 [NS	84 S	tandard	
		3 2021	02-17	04:2	29:39.3	2 75	.75.7	5.75	-	10	9.0.2	2.15	L I	NS	195 S	tandard	
		4 2021	02-17	04:2	29:39.2	2 75	. /5. /	15.75)	10		7 221	00 T	INS ICP	2075	tandard	
		6 2021	-02-17	04:2	19.39. 19.39.	2 10 2 34	107	15 221	82	16	1.107 0 0 2	, 221) 15	.02		74 0 60 8	$0 \rightarrow 5168$	
		7 2021	-02-17	04:2	9:39	2 34	.0.2.	15	02	34	1.107	7.221	.82]	CP	54 5	1680 → 8	
		8 2021 -	02-17	04:2	29:39.	2 10	.0.2.	15		34	1.107	7.221	.82	TTP	348 G	ET /succ	
		9 2021·	02-17	04:2	29:39.	2 34	.107.	221.	82	10	0.0.2	2.15	٦	СР	60 8	0 → 5168	Ų
	▶ Fran	me 8: 3	848 bv	tes o	n wire	e (27	84 bi	ts),	34	8 bv	tes	captu	red (27	84 bi	ts) on	interface	0
	▶ Ethe	ernet I	I, Sr	c: Pc	sComp	u f2:	d2:08	(08	:00	:27:	f2:d	2:08)	, Dst:	Realt	ekU_12	:35:02 (52	:54:0
	▶ Inte	ernet F	rotoc	ol Ve	rsion	4, S	rc: 1	.0.0.	2.1	5, D	st:	34.1	7.221.8	2	_		
	▶ Trai	nsmissi	on Co	ntrol	Prote	ocol,	Src	Port	: 5	1680	, Ds	t Por	: 80,	Seq:	150479	078, Ack:	64002
Network	▶ Нуре	ertext	Trans	fer P	rotoco	51											
(ID) Javor	Ι.												Lin				
(IP) layer		ransp	ort						\pp	lica	tior			`			
(3)		ayer (4)					li	ave	r (7)		lay	er (2			
	•																
	0000	52 54	00 1	2 35	02 08	00	27 f2	d2	08 (08 00	9 45	00	RT5.	'.	E.		ĥ
	0010	01 4e	9e 8	e 40	00 40 50 08	06 8	3T 4T	0a	00 (92 01 Fa 01	F 22	6D	.N@.	ø0	о"К		
	0020	72 10	0 d 0	000 00 d 00	00 47	45	20 eo 54 20	00 2f	73	75 63	2 50	10 65	.к r	 ЗЕ Т	/succe		
	0040	73 73	2e 7	4 78	74 20	48 !	54 54	50	2f :	31 20	e 31	0d	ss.txt	н тт	P/1.1.		\cup
	0050	0a 48	6f 7	3 74	3a 20	64	65 74	65	63	74 70	9 6f	72	.Host:	d et	ectpor		
	0060	74 61	6c 2	e 66	69 72	65 (66 6f	78	2e (63 61	F 6d	0d	tal.fi	re fo	x.com.		
	0070	0a 55	73 6	5 72	2d 41	67 (65 6e	74	3a 2	20 40	d 6f	7a	.User-	Ag en	t: Moz		
	0080	69 6C	6C 6	1 21	35 2e	30	20 28	58	31 3	31 31	5 20	55	111a/5	.0 (X11; U		

13



- 1. Network basics
- 2. Sending and receiving packets
 - 3. Sniffing packets
 - 4. Spoofing packets

Typically, the OS handles packet construction



Application has data that it wants to send

Passes data to OS which generates a packet

Packet sent out over network via NIC

How does the application hand data to OS? OS provides socket API system calls Recall sockets from CS10 and 50

Each NIC has unique MAC address Sees all traffic on network

Drops packets not addresses to itself (unless promiscuous or monitor mode)

Applications must provide protocol and destination information



You must decide which transport protocol to use:

- UDP faster, but delivery not guaranteed
- TCP slower but delivery guaranteed

You must also provide destination information

- IP address
- Port

OS adds router's MAC and checksum

Sending a packet is not difficult with Python

send_udp.py



Sending a packet is not difficult with Python

send_udp.py

import socket

dest_addr = "127.0.0.1" port = 9090 data = b'Hello world!'

Python hides a lot of complexity!

Note: we do not set a source port Will see soon the OS picks one for us

if __name__ == '__main__':

sock = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
sock.sendto(data,(dest_addr,port))

Listen in another terminal on port 9090 using netcat

\$ nc -luv 9090 Listening on [0.0.0.0] (family 0, port 9090) Hello world!^C

Sending a packet in C is a little more involved but gives you more control

send_udp.c



,

Receiving a packet is the reverse of sending a packet



Packets can be easily received in Python

receive_udp.py

import socket **Computer might have multiple NICs** Zeros means accept packets from any of them ip addr = "0.0.0.0"port = 9090Set up socket for UDP as before if name == ' main ': sock = socket.socket(socket.AF_INET, socket.SOCK_DGRAM) sock.bind((ip addr,port)) Bind socket to listen on ip_addr and port while (True): **Block until packet arrives** data, (ip, port) = sock.recvfrom(1024) recvfrom gets data then print("Sender: {} and Port: {}".format(ip,port)) print packet print("Received message: {}".format(data)) Use netcat in another terminal to send UDP packets to IP address on port 9090 \$ nc –u 127.0.0.1 9090 hello world Type in messages to send, sends after 21 return key pressed

OS picks a source port for the sender

receive_udp.py	We did not pick a source port on the sender (but did
import socket	pick destination port!) OS chooses one randomly for you
ip_addr = "0.0.0.0" port = 9090	When this is run, see source port other than 9090 This way OS can sort out replies if multiple instances of application are running at the same time

```
if __name__ == '__main__':
    sock = socket.socket(socket.AF_INET, socket.SOCK_DGRAM)
    sock.bind((ip_addr,port))
```

```
while (True):
    data, (ip, port) = sock.recvfrom(1024)
    print("Sender: {} and Port: {}".format(ip,port))
    print("Received message: {}".format(data))
```

A receiver written in C is not too different from one written in Python

receive_udp.c

const int MAX_SIZE = 1500; const int port = 9090;

void main() {
 struct sockaddr_in server;
 struct sockaddr_in client;
 int clientlen;
 char buf[MAX_SIZE];

// Create the socket int sock = socket(AF_INET, SOCK_DGRAM, IPPROTO_UDP); memset((char *) &server, 0, sizeof(server)); server.sin_family = AF_INET; server.sin_addr.s_addr = htonl(INADDR_ANY); server.sin_port = htons(port);
Bind to every interface (like Python with "0.0.0.0")
if (bind(sock, (struct sockaddr *) &server, sizeof(server)) < 0) { Server.sin_printf ("Binding error!"); return; }
Bind to every interface (like Python with "0.0.0.0")

```
// Getting captured packets
while (1) {
    bzero(buf,MAX_SIZE);
    recvfrom(sock, buf, MAX_SIZE-1, 0, (struct sockaddr *) &client, &clientlen);
    printf("%s\n",buf);
```

close(sock);

}

Client info filled when packets arrive



- 1. Network basics
- 2. Sending and receiving packets
- 3. Sniffing packets
 - 4. Spoofing packets

Promiscuous (monitor) mode sniffs all frames



Sniffing all frames can take a lot of CPU time!

sniff_raw.c



Most of the time we will want to filter out unwanted frames to reduce processing

Layer 2 overview



Berkeley Packet Filter (BPF)

- Creates filter at link-level
- Only passes frames matching criteria to protocol stack
- Does not DMA to memory
- Compile filter and set with setsockopt

BPF is somewhat complicated Add to raw socket with SO_ATTACH_FILTER Not portable between OSes PCAP API is much easier!

PCAP makes it easy to filter frames at a low level

PCAP (packet capture)

- Originally written for tcpdump (powerful sniffer)
- Supported by multiple platforms
 - Linux: libpcap
 - Windows: WinPcap and Npcap
- Written in C
- Other languages generally provide a wrapper around C version
- Basis used by other sniffers
 - Tcpdump (of course)
 - Wireshark
 - Scapy
 - Nmap
 - Snort

Use PCAP and compile BPF filter to filter out unwanted frames

void got_packet(u_char *args, const struct pcap_pkthdr *header, const u_char *packet) { printf("Got a packet\n"); Prints got packet when UDP or ICMP packet arrives Last parameter has packet details Parsing packets int main() { in C is tedious **Open PCAP session** pcap_t *handle; Sniff on interface *enpOs3* (use ifconfig to find) char errbuf[PCAP ERRBUF SIZE]; 3rd parameter sets promiscuous mode to true struct bpf program fp; char filter exp[] = "udp or icmp"; bpf u int32 net; const int MAX SIZE = 8192; // Step 1: Open live pcap session on NIC with name enp0s3 handle = pcap open live("enp0s3", MAX SIZE, 1, 1000, errbuf);

if (handle == NULL) { printf("Error on open\n"); printf("errbuf %s\n",errbuf); return (1); }

// Step 2: Compile filter_exp into BPF psuedo-code
pcap_compile(handle, &fp, filter_exp, 0, net);
if (pcap_setfilter(handle, &fp) != 0) { printf("set filter error"); return(1); }

// Step 3: Capture packets
pcap_loop(handle, -1, got_packet, NULL);

pcap_close(handle); //Close the handle
return 0;

Start sniffing Set callback function to *got_packet*

Scapy makes it easy to use PCAP from Python to sniff packets with a filter





- 1. Network basics
- 2. Sending and receiving packets
- 3. Sniffing packets



Demo: packet spoofing is easy with Scapy

udp_spoof.py

False flag: it looks like this packet came from another computer!

#!/usr/bin/python3 from scapy.all import *

Demo

- Start secondary VM with address 10.0.2.5 (confirm with ifconfig)
- Start Wireshark on secondary VM
- Run this code from primary VM (sudo python3 udp_spoof.py)
- Wireshark shows packet came from 1.2.3.4!



Sometimes we want to sniff, then spoof a reply

sniff_spoof_icmp.py

from scapy.all import *

def spoof_pkt(pkt):

if ICMP in pkt and pkt[ICMP].type == 8:
 #listen for ICMP request packets (type 8)
 print("Original Packet.....")
 print("Source IP : ", pkt[IP].src)
 print("Destination IP :", pkt[IP].dst)

```
#spoof a reply, even if the request wasn't for us
#must reverse source and destination on reply!
ip = IP(src=pkt[IP].dst, dst=pkt[IP].src, ihl=pkt[IP].ihl)
icmp = ICMP(type=0, id=pkt[ICMP].id, seq=pkt[ICMP].seq)
data = pkt[Raw].load
newpkt = ip/icmp/data
```

print("Spoofed Packet......")
print("Source IP : ", newpkt[IP].src)
print("Destination IP :", newpkt[IP].dst)

send(newpkt,verbose=0)

Demo

- Start Wireshark on secondary VM
- Run this code from primary VM (sudo python3 sniff_spoof_icmp.py)
- From secondary VM ping 10.0.2.6 (doesn't exist)
- Wireshark shows packet came from 10.0.2.6

Steps

- Receive packet
- Extract details
- Spoof a new packet using extracted details

We have created a phantom computer on the network that responds to pings from 10.0.2.15!

Scapy vs. C

Scapy

Pros

- Easier to write
- Sets reasonable default fields
- Calculates values (checksums)
- Focus on fields interested
- Every layer is an object, can easily stack together
- Increased productivity

Pros

C

- Runs much faster! (50-100X!)
- More control over crafting packets

Scapy vs. C

Scapy

Pros

- Easier to write
- Sets reasonable default fields
- Calculates values (checksums)
- Focus on fields interested
- Every layer is an object, can easily stack together
- Increased productivity

Cons

Runs slowly

Pros

C

- Runs much faster! (50-100X!)
- More control over crafting packets

Cons

- Tricky to write, must get the byte offsets right
- Tedious

Scapy and C can work together in a hybrid approach

Hybrid approach

- For some uses speed is critical Scapy is slow
- Example:
 - Send 1000 UDP packets
 - Scapy: 9.4 seconds
 - C : 0.25 seconds
 - C is 37X faster than Scapy in this case!
- Sometimes racing to reply before the "real" computer Scapy too slow

Idea:

- Use Scapy to create packets, and save to file
- Read packets in C (might to adjust some fields)
- Send using C



Endianness

Endianness: a term that refers to the order in which a given multibyte data item is stored in memory

- Little Endian: store the most significant byte of data at the highest address
- Big Endian: store the most significant byte of data at the lowest address
 - Little endian always seems backwards to me

Store 0xDEADBEEF

Big endian

DE	AD	BE	EF	
----	----	----	----	--

Increasing addresses ———

Little endian

EF BE	AD	DE	
-------	----	----	--

Endianness in network communication

Computers with different byte orders will misunderstand each other

- Solution: agree upon a common order for communication
- This is called "network order", which is the same as Big Endian order
- But Intel computers ("hosts") use Little Endian
- Must convert data between "host order" and "network order"

Macro	Description
htons()	Convert unsigned short integer from host order to network order.
htonl()	Convert unsigned integer from host order to network order.
ntohs()	Convert unsigned short integer from network order to host order.
ntohl()	Convert unsigned integer from network order to host order.

Classful addressing scheme (1981 - 1993)

Class A: 0.0.0.0 – 127.255.255.255 👞

Class B: 128.0.0.0 – 191.255.255.2555

Class C: 192.0.0.0 – 223.255.255.255

Class D: 224.0.0.0 – 239.255.255.255

Class E: 240.0.0.0 – 255.255.255.255

Used to help routing, 32 bits total

Left bit always 0, if see you a zero in the left bit, you know it belongs to a Class A network

Network ID is 2⁷ = 127 possible address for first octet

If you had Class A network (first 8 bits set) you have 2²⁴ = 16.7M possible addresses that can be given to devices

Issued to very large organizations

Classful addressing scheme (1981 - 1993)

Class A: 0.0.0.0 – 127.255.255.255

Class B: 128.0.0.0 – 191.255.255.2555 🍗

Class C: 192.0.0.0 – 223.255.255.255

Class D: 224.0.0.0 – 239.255.255.255

Class E: 240.0.0.0 – 255.255.255.255

Used to help routing **32** bits total

Network ID is left 16 bits

- Left bit = 1
- Left second bit = 0

If you see left bits are 10, you know this is a Class B network

If you had Class B network (first 16 bits set) you have 2¹⁶ = 65.5K possible addresses that can be given to devices

Issued to large organizations

Classful addressing scheme (1981 - 1993)

Class A: 0.0.0.0 – 127.255.255.255

Class B: 128.0.0.0 – 191.255.255.2555

Class C: 192.0.0.0 – 223.255.255.255

Class D: 224.0.0.0 – 239.255.255.255

Class E: 240.0.0.0 – 255.255.255.255

Used to help routing **32** bits total

Network ID is left 24 bits

- Left bit = 1
- Left second bit = 1
- Third left bit = 0

If you see left bits are 110, you know this is a Class C network

If you had Class C network (first 24 bits set) you have 2⁸ = 256 possible addresses that can be given to devices

Issued to medium sized organizations

Classful addressing scheme (1981 - 1993)

Class A: 0.0.0.0 – 127.255.255.255

Class B: 128.0.0.0 – 191.255.255.2555

Class C: 192.0.0.0 – 223.255.255.255

Class D: 224.0.0.0 - 239.255.255.255

Class E: 240.0.0.0 – 255.255.255.255

Used to help routing **32** bits total

Class D for multicast (special purpose)

Class E reserved

Problem:

We were using up IP addresses too quickly (lots of unused addresses if you own a Class A network)

Happened to us, we had a class C and had to give up some addresses