Network Protocols and Vulnerabilities

John Mitchell

Outline

- ◆Basic Networking (FMU)
- ◆Network attacks
  - Attack host networking protocols
    - SYN flooding, TCP Spoofing, …
  - Attack network infrastructure
    - Routing
    - Domain Name System

This lecture is about the way things work now and how they are not perfect. Next lecture – some security improvements (still not perfect).

Internet Infrastructure

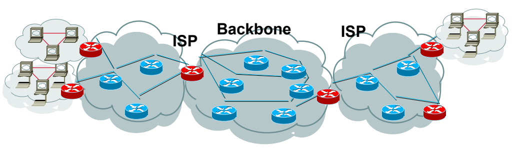

- ◆Local and interdomain routing
  - TCP/IP for routing, connections
  - BGP for routing announcements
- ◆Domain Name System
  - Find IP address

TCP Protocol Stack

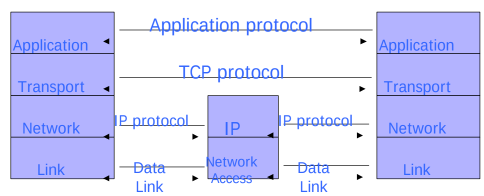

Data Formats

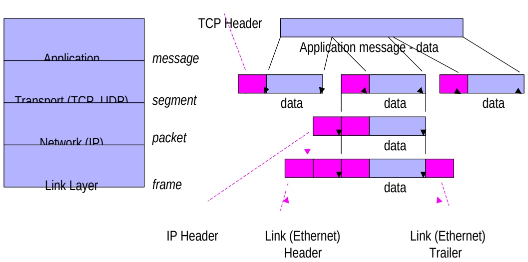

IP

Internet Protocol

- ◆Connectionless
  - Unreliable
  - Best effort
- ◆Transfer datagram
  - Header
  - Data

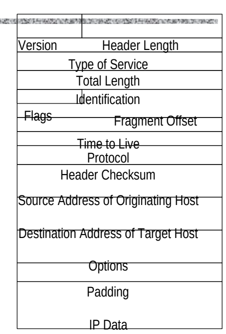

IP Routing

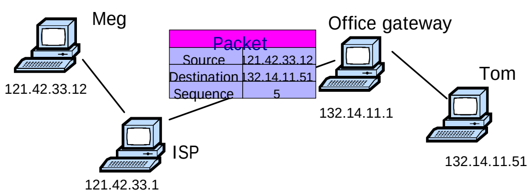

- Internet routing uses numeric IP address
- Typical route uses several hops

IP Protocol Functions (Summary)

- Routing
  - IP host knows location of router (gateway)
  - IP gateway must know route to other networks
- Error reporting
  - IP reports discards to source
- Fragmentation and reassembly
  - If packets smaller than the user data

UDP

User Datagram Protocol

- IP provides routing
  - IP address gets datagram to a specific machine
- UDP separates traffic by port
  - Destination port number gets UDP datagram to particular application process, e.g., 128.3.23.3, 53
  - Source port number provides return address
- Minimal guarantees (... mice and elephants)
  - No acknowledgment
  - No flow control
  - No message continuation

TCP

Transmission Control Protocol

- Connection-oriented, preserves order
  - Sender
    - Break data into packets
    - Attach packet numbers
  - Receiver
    - Acknowledge receipt; lost packets are resent
    - Reassemble packets in correct order

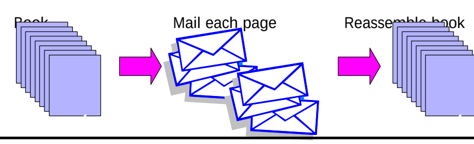

ICMP

Internet Control Message Protocol

- Provides feedback about network operation
  - Error reporting
  - Reachability testing
  - Congestion Control
- Example message types
  - Destination unreachable
  - Time exceeded
  - Parameter problem
  - Redirect to better gateway
  - Echo/echo reply - reachability test
  - Timestamp request/reply - measure transit delay

Basic Security Problems

- Network packets pass by untrusted hosts
  - Eavesdropping, packet sniffing
- IP addresses are public
  - Smurf
- TCP connection requires state
  - SYN flooding attack
- TCP state easy to guess
  - TCP spoofing attack

Packet Sniffing

- Promiscuous NIC reads all packets
  - Read all unencrypted data
  - ftp, telnet send passwords in clear!

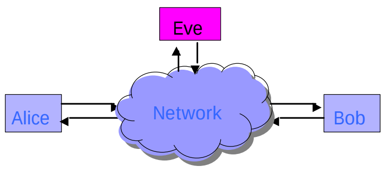

Sweet Hall attack installed sniffer on local machine

Prevention: Encryption, improved routing (Next lecture: IPSEC)

Smurf Attack

- Choose victim
  - Idea: Flood victim with packets from many sources
- Generate ping stream (ICMP Echo Req)
  - Network broadcast address with spoofed source IP set to victim
- Wait for responses
  - Every host on target network will generate a ping reply (ICMP Echo Reply) to victim
  - Ping reply stream can overload victim

Prevention: Turn off ping? Authenticated IP addresses?

TCP Handshake

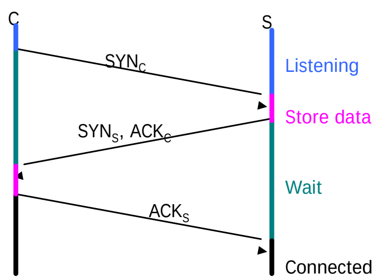

SYN Flooding

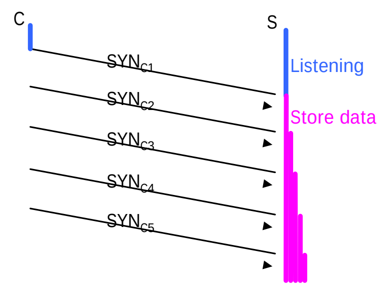

SYN Flooding

- Attacker sends many connection requests
  - Spoofed source addresses
- Victim allocates resources for each request
  - Connection requests exist until timeout
  - Fixed bound on half-open connections
- Resources exhausted ⇒ requests rejected

Protection against SYN Attacks

[Bernstein, Schenk]

- Client sends SYN
- Server responds to Client with SYN-ACK cookie
  - sqn = f(src addr, src port, dest addr, dest port, rand)
  - Server does not save state
- Honest client responds with ACK(sqn)
- Server checks response
  - If matches SYN-ACK, establishes connection

See http://cr.yp.to/syncookies.html

Random Deletion

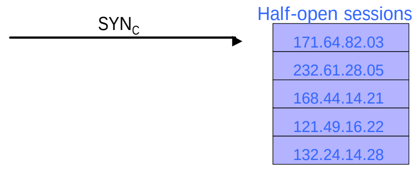

- If queue is full, delete random entry
  - Legitimate connections have chance to complete
  - Fake addresses eventually deleted

Easy to implement, some improvement

TCP Connection Spoofing

- Each TCP connection has an associated state
  - Sequence number, port number
- Problem
  - Easy to guess state
    - Port numbers are standard
    - Sequence numbers often chosen in predictable way

IP Spoofing Attack

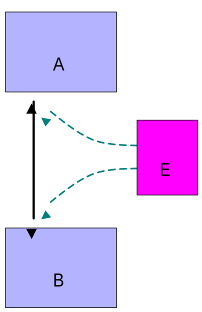

- A, B trusted connection
  - Send packets with predictable seq numbers
- E impersonates B to A
  - Opens connection to A to get initial seq number
  - SYN-floods B's queue
  - Sends packets to A that resemble B's transmission
  - E cannot receive, but may execute commands on A

Attack can be blocked if E is outside firewall.

TCP Sequence Numbers

- Need high degree of unpredictability
  - If attacker knows initial seq # and amount of traffic sent, can estimate likely current values
  - Send a flood of packets with likely seq numbers
    - larger bandwidth => larger flood possible
- Reported to be safe from practical attacks
  - Cisco IOS, OpenBSD 2.8-current, FreeBSD 4.3-RELEASE, AIX, HP/UX 11i, Linux Kernels after 1996
  - Solaris 2.6 if strong seq numbers turned on:
    - Set TCP_STRONG_ISS to 2 in /etc/default/inetinit.
  - HP/UX , IRIX 6.5.3, ... if so configured

Cryptographic protection

- Solutions above the transport layer
  - Examples: SSL and SSH
  - Protect against session hijacking and injected data
  - Do not protect against denial-of-service attacks caused by spoofed packets
- Solutions at network layer
  - IPSec
  - Can protect against
    - session hijacking and injection of data
    - denial-of-service attacks using session resets

TCP Congestion Control

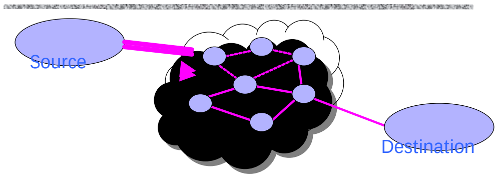

- If packets are lost, assume congestion
  - Reduce transmission rate by half, repeat
  - If loss stops, increase rate very slowly

Design assumes routers blindly obey this policy

Competition

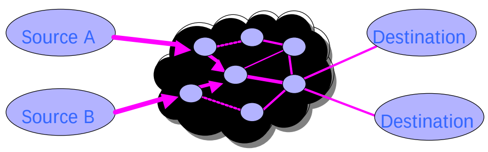

- Amiable Alice yields to boisterous Bob
  - Alice and Bob both experience packet loss
  - Alice backs off
  - Bob disobeys protocol, gets better results

TCP Attack on Congestion Control

- Misbehaving receiver can trick sender into ignoring congestion control
  - Receiver: duplicate ACK indicates gap
    - Packets within seq number range assumed lost
    - Sender executes fast retransmit algorithm
  - Malicious receiver can
    - Send duplicate ACK
    - ACK before data is received
      - needs some application level retransmission – e.g. HTTP 1.1 range requests … See RFC 2581
  - Solutions
    - Add nonces – ACKs return nonce to prove reception

See: Savage et al., TCP Congestion Control with a Misbehaving Receiver

Routing Vulnerabilities

- Source routing attack
  - Can direct response through compromised host
- Routing Information Protocol (RIP)
  - Direct client traffic through compromised host
- Exterior gateway protocols
  - Advertise false routes
  - Send traffic through compromised hosts

Source Routing Attacks

- Attack
  - Destination host may use reverse of source route provided in TCP open request to return traffic
    - Modify the source address of a packet
    - Route traffic through machine controlled by attacker
- Defenses
  - Gateway rejects external packets claiming to be local
  - Reject pre-authorized connections if source routing info present
  - Only accept source route if trusted gateways listed in source routing info

Routing Table Update Protocols

- Interior Gateway Protocols: IGPs
  - distance vector type - each gateway keeps track of its distance to all destinations
    - Gateway-to-Gateway: GGP
    - Routing Information Protocol: RIP
- Exterior Gateway Protocol: EGP
  - used for communication between different autonomous systems

Interdomain Routing

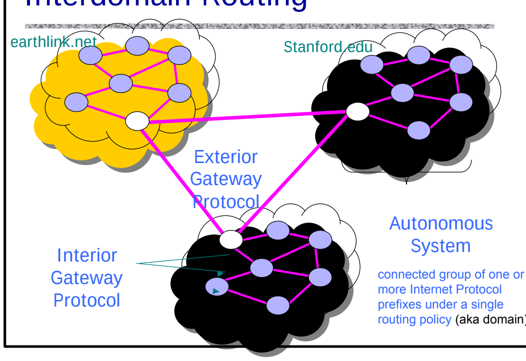

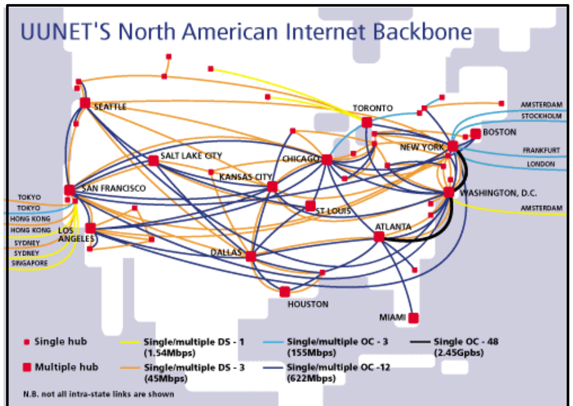

BGP overview

- Iterative path announcement
  - Path announcements grow from destination to source
  - Subject to policy (transit, peering)
  - Packets flow in reverse direction
- Protocol specification
  - Announcements *can* be shortest path
  - Nodes allowed to use other policies
    - E.g., "cold-potato routing" by smaller peer
  - Not obligated to use path you announce

BGP example [D. Wetherall]

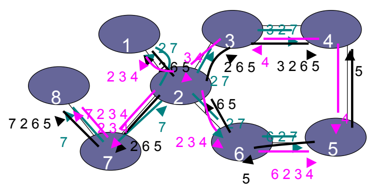

- Transit: 2 provides transit for 7
  - 7 reaches and is reached via 2
- Peering: 4 and 5 peer
  - exchange customer traffic

Issues

- BGP convergence problems
  - Protocol allows policy flexibility
  - Some legal policies prevent convergence
  - Even shortest-path policy converges slowly
- Incentive for dishonesty
  - ISP pays for some routes, others free
- Security problems
  - Potential for disruptive attacks

DNS

Domain Name System

- Hierarchical Name Space

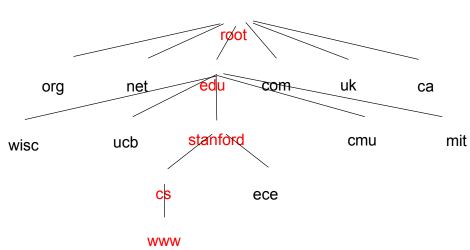

DNS Root Name Servers

- Root name servers
- Local name servers contact root servers when they cannot resolve a name

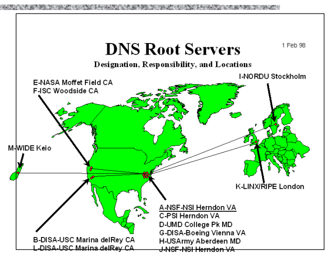

DNS Lookup Example

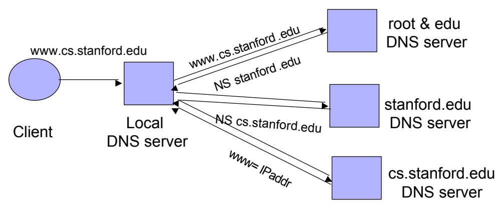

Caching

- DNS responses are cached
  - Quick response for repeated translations
  - Other queries may reuse some parts of lookup
    - NS records for domains
- DNS negative queries are cached
  - Don't have to repeat past mistakes
  - E.g. misspellings, search strings in resolv.conf
- Cached data periodically times out
  - Lifetime (TTL) of data controlled by owner of data
  - TTL passed with every record

Subsequent Lookup Example

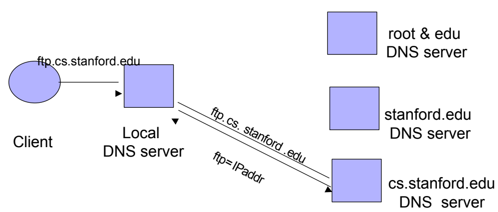

DNS Implementation Vulnerabilities

- Reverse query buffer overrun in BIND Releases 4.9 (4.9.7 prior) and Releases 8 (8.1.2 prior)
  - gain root access
  - abort DNS service
- MS DNS for NT 4.0 (service pack 3 and prior)
  - crashes on chargen stream
  - telnet ntbox 19 | telnet ntbox 53

Inherent DNS Vulnerabilities

- Users/hosts typically trust the host-address mapping provided by DNS
- Problems
  - Zone transfers can provide useful list of target hosts
  - Interception of requests or compromise of DNS servers can result in bogus responses
  - Solution – authenticated requests/responses

Bellovin/Mockapetris Attack

- Trust relationships use symbolic addresses
  - /etc/hosts.equiv contains friend.stanford.edu
- Requests come with numeric source address
  - Use reverse DNS to find symbolic name
  - Decide access based on /etc/hosts.equiv, …
- Attack
  - Spoof reverse DNS to make host trust attacker

Reverse DNS

- Given numeric IP address, find symbolic addr
- To find 222.33.44.3,
  - Query 44.33.222.in-addr.arpa
  - Get list of symbolic addresses, e.g.,

1	IN	PTR	server.small.com
2	IN	PTR	boss.small.com
3	IN	PTR	ws1.small.com
4	IN	PTR	ws2.small.com

Attack

- Gain control of DNS service for domain
- Select target machine in domain
- Find trust relationships
  - SNMP, finger can help find active sessions, etc.
  - Example: target trusts host1
- Connect
  - Attempt rlogin from compromised machine
  - Target contacts reverse DNS server with IP addr
  - Use modified reverse DNS to say addr is host1
  - Target allows rlogin

Defense against this attack

- Double-check reverse DNS
  - Modify rlogind, rshd to query DNS server
  - See if symbolic addr maps to numeric addr
- Use another service besides DNS
  - Network Information Service (NIS, or YP)
  - Only works if attacker cannot control NIS …
- Authenticate entries in DNS tables
  - Relies on some form of PKI?
  - Next lecture …

Summary (I)

- Eavesdropping
  - Encryption, improved routing (Next lecture: IPSEC)
- Smurf
  - Turn off ping? Authenticated IP addresses?
- SYN Flooding
  - Cookies
  - Random deletion
- IP spoofing
  - Use less predictable sequence numbers

Summary (II)

- Source routing attacks
  - Additional info in packets, tighter control over routing
- Interdomain routing
  - Authenticated routing announcements
  - Other issues
- DNS attack
  - Double-check reverse DNS
  - Use another service besides DNS
  - Authenticate entries in DNS tables