

Threat Pragmatics

Steven M. Bellovin

<https://www.cs.columbia.edu/~smb>

Targets

- Many sorts of targets:
 - Network infrastructure
 - Network services
 - Application services
 - User machines

What's at risk?

Network Infrastructure

- Routers (and routing protocols)
- Switches and other network elements
- Links

Links

- Primary risk is wiretapping
- Easily defeated by encryption—but are people using it?
- Most encryption doesn't protect against traffic analysis—but that isn't in everyone's threat model
- Link-layer encryption protects against most traffic analysis, but it has to be done on every vulnerable link

Traffic Analysis

- Looks at *external* characteristics of traffic: who talks to whom, size of messages, etc.
- *Very* valuable to intelligence agencies, police, etc.
 - Who works with whom? Who gives orders to whom?
- Not generally useful for ordinary thieves, though a few sophisticated attackers could use it to find targets

Solutions

- Use VPNs or application-level encryption
- Use link encryption for high-risk links (e.g., WiFi)
- Also use link encryption for access control (especially WiFi)
- Don't worry about traffic analysis—unless your enemy is an intelligence agency

(Is WiFi Safe?)

- Inside an organization, WiFi+WPA2 Enterprise is generally safe enough without further crypto
 - However, it's harder to trace an infected host that's doing address-spoofing
- For external WiFi, *always* use crypto, preferably VPNs
 - Make sure you do mutual authentication
- There is some residual risk if your VPN doesn't drop unencrypted inbound traffic

Switches and the Like

- Compromised switches can be used for eavesdropping
- Special risk in some situations: reconfigured VLANs
 - VLANs provide good traffic separation between user groups
 - Especially useful against ARP- and MAC-spoofing attackers
- Other danger point: the monitoring port

ARP and MAC Spoofing

- ARP maps the IP address desired to a MAC address
- Switches learn what MAC addresses are on what ports, and route traffic accordingly
- If a malicious host sends out traffic with the wrong MAC address, the switch will send traffic to it
- If a malicious host replies to an ARP query for some other machine, the malicious host will receive the traffic, but this might be noticed

Address-Spoofing Happens

- A few years ago, someone spoofed the IP and MAC addresses of my department's FTP server
- The attacking machine was in another building but on the same VLAN
- No one had noticed the intermittent failures of the FTP service
- Our switches log MAC and IP addresses, and we keep those logs for a long time
- The machine had been penetrated 6 months earlier....

Defenses

- Harden switch access
 - ACLs
 - ssh-only access, and only using public/private key pairs; no passwords
- Hosts should use crypto and cryptographic authentication

Routers

- Routers can be used for the same sorts of attacks as switches
- Because routers inherently separate different networks, they always defend against certain kinds of address spoofing
 - This makes them targets
- Worse yet, routers can launch *routing protocol attacks*


Routing Protocol Attacks: Effects

- Traffic is diverted
 - Attacker can see the traffic and do traffic analysis
 - Attacker can modify packets
 - Attacker can drop packets
 - Attacker can hijack prefixes
- End-to-end crypto can protect the packets' contents, but can't stop traffic analysis or denial of service

Why is Routing Security Different?

- Most security failures are due to buggy code, buggy protocols, or buggy sysadmins
- Routing security problems happen when everything is working right, but some party decides to lie. The problem is a dishonest participant
- Most routers can lie via any routing protocols they're using

A Routing Attack



Z claims that it has a shorter path to B than Y does (1 hop versus 2). X believes Z.

Defending Against Routing Attacks

- Must *know* authoritative owner of prefixes
 - Generally done with a certificate signed by the address space owner
 - Being rolled out today as RPKI
- All routing announcements must be digitally signed
 - Each router needs a route-signing certificate
 - All signatures must be over the full path; signatures are thus *nested*
 - In the IETF process as BGPSEC

An (Oversimplified) Example

- $W \rightarrow Y: \{B, \{W, Y\}\}_W$
- $Y \rightarrow X: \{\{B, \{W, Y\}\}_W, X\}_Y$
- W announces “I can reach B; I’m sending this to Y”; the message is signed by W
- Y announces “I can reach B by going through W’s signed path”; this message is signed by Y

Network Services

Certain core services are ubiquitous—and frequently attacked

- DNS
- SMTP
- Assorted local services: file servers, printers, LDAP, and more

These are the means, not the goals of the attackers

DNS

DNS responses are easily spoofed by attackers

- Cache contamination
- Query ID guessing
- Deliberate tinkering by ISPs, nation-states, hotels, etc.

Because responses are cached, client/server authentication can't solve it.

Must have *digitally signed* records (DNSSEC)

SMTP

- Historically, a major attack target; principle implementations were very buggy
- Today, the big problem is spam; must keep attackers from spamming your users, and from using you to spread spam
- Secondary issue: separate inside and outside email systems—inside email often has sensitive information

Encrypted Email

- Email messages themselves can be encrypted: useful for end-to-end security
 - But S/MIME and PGP are hard to use, and their *absence* will not be noticed
- SMTP can be encrypted, too
 - Not that crucial for site-to-site relaying (but eavesdroppers do exist); *very* important for authenticated email submission
 - Your users *must* authenticate somehow—via IP address if inside; via credentials if roaming—before sending mail through your outbound SMTP server

Local Services

- Rarely directly accessible from the Internet; (ab)used after initial penetration
 - Virus spreading
 - File contents, in targeted attacks
 - Privilege escalation
- Quite often buggy, but there's little choice about running them; they're necessary for scalability and productivity

Application Services

- Data center-resident: deliver services to the outside world
- Obvious example: HTTP
- But—HTTP is generally a front end for a vital database
- A prime target

Targeting Application Services

- Generally exposed to the outside—and you can't firewall them, because they *must* be exposed to the outside
- The server can be used for the bad guys' content: phishing servers, “warez” sites, more
- The database often holds very valuable information, like credit cards
- There are usually connections from these servers back into the corporation

User Machines

Ordinary desktops are targets, too

- Plant keystroke loggers to steal passwords, especially for financial sites
- Turn into bots—bandwidth is what matters
- Turn into spam engines; use machine's privileges (generally based on network location) to send out spam through the authorized SMTP server

Users

- Users make mistakes
 - They click on things they shouldn't
 - They visit dangerous sites
 - They mistake phishing emails for the real thing
 - They don't keep their systems up to date
 - “PEBCAK”: Problem Exists Between Chair and Keyboard
- (It's not even their fault; our systems are horribly designed)

Social Engineering

- Try to trick people into doing things they shouldn't
- People *want* to help
 - Walk in the door dressed as a delivery or repair person
 - Call and sound like an insider: “Chris, could you reset my password on server #3 in rack 7? Its connection to the RADIUS server is hung.”
- A very different skill than purely technical stuff—but *very* useful