

Network and Internet Vulnerabilities

David Aspinall
School of Informatics,
University of Edinburgh.
3rd February 2005

<http://www.informatics.ed.ac.uk/teaching/courses/cs>

2005/02/03 11:43:40

Contents

- Some attacks:
SYN flooding Smurfing DDos UDP RPC, NFS, NIS
Spamming Spoofing Routing X Window SNMP
- Some vulnerabilities:
 - Next lecture: some defences

Internet protocols were designed for robust transport, not robust security. Most current problems are due to unsecured protocols and/or exploiting software bugs.

Internet attacks and defences

- General pattern in serious internet security incidents:
 1. Somebody discovers an exploit (majority are due to software bugs)
 2. Exploit is seen in the wild, possibly to large effect
 3. Short-term workarounds; attack-specific detection/recovery
 4. Proper repairs to software or protocols are issued
 5. Over time, majority of sites implement repair mechanisms
 6. Remaining sites may be black-listed and outlawed by abiding sites.
- The last step happens in the most serious cases, especially where security flaws may be exploited to attack other, well-managed sites.
- Internet security is a **distributed community-wide responsibility**. Black-listing is a socioeconomic countermeasure. Black lists may be useful for crackers as well as good guys (they list hosts which may have security holes), so systems which are not repaired find themselves being attacked and isolated from the rest of the network.

SYN flooding

- Recall the basic three-part handshake used by Alice to initiate a TCP connection with Bob, and send initial sequence numbers:

$A \rightarrow B : \text{SYN}, X$

$B \rightarrow A : \text{ACK}, X + 1; \text{SYN}, Y$

$A \rightarrow B : \text{ACK}, Y + 1$

- A DoS attack is **SYN Flooding**. Alice sends many SYN packets, without acknowledging any replies. Bob accumulates more SYN packets than he can handle. Large-scale attacks were seen in 1996.
- A protocol implementation fix called **SYNcookie**, is for Bob to send out Y as encrypted version of X , so he doesn't need to keep state. This is implemented in Linux and some other systems.

Spoofing: forged TCP packets

- **Responses to attacks**. Victim and Internet community at large want to find attack source, so that the corporate network administrators or ISP can be notified and given the chance to prevent it (or risk being isolated). Tracing an attack also allows the ultimate possibility of raising legal actions against attack perpetrators.
- Tracing network packets to their source is therefore important. But unfortunately it is easy to **forge the source address** on an IP packet.
- SYN flooding attacks usually have forged source addresses. The ACK is either impossible (address not reachable) or targets another machine, at the least wasting bandwidth. sending meaningless ACK packets.
- The SYNcookie fix doesn't itself prevent flooding. As a countermeasure to assist tracing, network providers are urged to implement **ingress filtering** on edge routers (RFC 2267). Packets entering the Internet are checked to have source addresses within the network fragment they came from, restricting forged packets.

Smurfing (directed broadcast)

- The **smurfing** attack exploits the ICMP (*Internet Control Message Protocol*) which allows remote hosts to respond to echo packets to say they're alive (e.g., sent with ping). Some implementations respond to pings to broadcast address (idea was: ping a LAN to see which hosts alive). A collection of hosts that do this gives a *smurf amplifier*.
- Attack: make packet with forged source address containing the victim's IP number. Send it to smurf amplifiers, who then swamp target with replies. Often used to take control of an IRC chatroom.
- Fix: standards change August 1999, ping packets sent to broadcast addresses aren't answered. As this is implemented, no. of smurf amplifiers is reducing. Black-listing: "concerned sysadmins" at <http://www.netscan.org> publish name-and-shame list of misconfigured nets, used by diligent sysadmins and hackers alike.
- A **fraggle**: similar attack with UDP packets (port 7, or other ports). Also attacks using TCP.

netscan.org on 5th Feb 2004

netscan.org

Current count: **10,901** broken networks.
Average amplification: **3x**

Welcome to **netscan.org**. This site contains a searchable and browsable list of broadcast ICMP ("smurf") amplifiers.

- 16th Feb 2003: 46k broken networks
- In UK a few nets, including on nets hosted by BT regional nets. (But really fault of individual sysadmins, not ISPs).
- Today (3rd Feb 2005): only 2k broken networks reported.

- 7 -

DDoS attacks

- In a **distributed denial of service attack**, a large number of machines are subverted with malicious code (e.g., via worm or virus), and then synchronized to attack a target together.
- Specific defences:
 - Distribute servers over server farm (expensive)
 - Dynamically relocate network under heavy attack (tricky/ineffective)
 - *Pushback*: try to dynamically restrict likely DoS packets, by rate-limiting with a *congestion signature* (experimental)
- General defences, by improving packet integrity/traceability:
 - Community responsibility: filtering out forged source addresses
 - Have routers add extra *ICMP traceback messages* with a low probability, e.g., 1 in 20,000. Then sysadmins can trace large-scale attacks back to responsible machines (even if IP spoofing is used).

- 8 -

DNS attacks

- Many protocols, including most email and web protocols (e.g., smtp and http) assume that lower levels are secure. The most they will do to authenticate is check source or destination addresses using DNS look-ups of hostname or reverse look-ups of IP addresses.
- If the DNS can be corrupted somehow, DNS checks may be unreliable, leading to **address forgery**, **spam**, in general, powerful **spoofing attacks**.
- The attack called **DNS cache poisoning** is based on feeding false information into locally cached DNS tables. It means that, within some network portion, a web site can be redirected elsewhere, for example, completely outwith the web-site server's control.


- 9 -

Sequence number attacks

- Suppose Mallory knows Alice and Bob are hosts on a target LAN, and wants to masquerade as Alice to Bob (in one direction). Here's a strategy:
 1. Take Alice down with a DoS attack (optionally)
 2. Initiate a new connection with Bob, by sending a SYN packet.
 3. Mallory doesn't get Bob's ACK, but may be able to **guess the sequence number** Y used by Bob.
- Initial sequence numbers may be somehow predictable, so Mallory could make his own connection with Bob and a while later use a related Y value when masquerading as Alice. Good IP stack implementations (e.g., recent Linux, Solaris, IRIX, FreeBSD) use random increments or random values. Other implementations are either not random at all, or not random enough (e.g., all versions of Windows).
- This is a complex attack, but handily, scripts are available.
- ISN predictability graphically is at <http://lcamtuf.coredump.cx/newtcp...>

- 10 -

ISN Predictability



- Plots in 2002 for WinXP (tl), Linux (tr), OS/400 (bl), UNICOS (br).

- 11 -

Connection hijacking

- An attacker who observes the current sequence number of a connection can inject phony packets.
 1. Alice logs in to a server. Mallory watches the connection.
 2. At the right moment, he takes down or disconnects Alice
 3. Then he takes over the session; provided he gets the sequence number correct, it will be accepted by the server.
- Session hijacking may be detectable, if the acknowledgement packet sent by Mallory cites data that the server never sent. This ought to cause a reset of the connection; instead the server may assume that sequence numbers have wrapped around, and resend its current sequence number and acknowledgement number.
- Can also be done undetectably, by sending synchronized empty packets to Alice instead of disconnecting her, and then letting them reconnect afterward.

- 12 -

Routing attacks

- Protocols like OSPF (*Open Shortest Path First*) let routers tell their neighbours about hosts they can reach, and cost metrics (hops). The *transitivity of trust* in routing protocols makes security difficult.
- Attacker who controls routing protocols can monitor, intercept, and modify much traffic. For example, malicious node *M* announces low cost route to hosts *A* and *B*, and thereby diverts traffic through itself.
- Packet switched networks route return data flow independently. Using network address for authentication implicitly trusts integrity of return path; if subverted, allows *masquerading*. Less of a threat on circuit-switched networks. But switches become new points of trust.
- TCP includes **source routing**, used to bypass network outages. Source-routed packets bypass the (weak) authentication of the return address. Forged ICMP **redirect messages** can have similar effect.

- 13 -

X Window protocol

- In X, a server runs the physical screen, keyboard, and mouse; applications connect and are allocated use of those resources. A malicious application can monitor all keystrokes, dump the screen, scribble on it, etc.
- X has several authentication mechanisms of varying quality:
 1. none
 2. IP address
 3. "magic cookie" — clear-text password
 4. cryptographic mechanismsIf any of these authentication mechanisms are broken by an attacker, he can attach a malicious application to the server.

- 14 -

UDP...RPC, NFS, NIS

- UDP, the *User Datagram Protocol*, is connectionless. There isn't even the weak authentication from a return path, so source addresses **cannot be trusted at all**.
- Protocols built on UDP are therefore immediately at risk, unless they implement their own security mechanisms. Unfortunately, the most important, **RPC** (*Remote Procedure Call*) does not. The ordinary RPC authentication field is unsecure; the RPC crypto option is rarely used.
- RPC is used to implement **NFS** (*Network File System*), and **NIS** (*Network Information Service*).
- NFS and NIS have numerous additional security problems. NFS file-handles can be guessed. NIS may serve up password files, and NIS server responses can be faked.

- 15 -

SNMP

- SNMP, the *simple network management protocol*, is used to configure network devices including routers and switches, and allows servers and devices to report status information.
- Useful for hackers to obtain sensitive info about systems, for example, routing tables.
- Later versions of SNMP have security features (MD5 authentication, DES encryption), but many devices only implement SNMPv1 which sends reports and passwords in clear text.
- Recently new vulnerabilities have been discovered.

- 16 -

Other attacks, mechanisms and tools

- **Packet sniffers** are eavesdropping tools which collect packets passing over the network, typically to skim plaintext login ids and passwords.
- **Port scanning** tools or more generally **vulnerability scanners** can be used to find and investigate network hosts open to particular attacks. Useful to good guys as well as bad guys. SATAN, nmap, etc..
- **Authentication attacks** based on breaking authentication protocols using attacks we've seen before, or brute-force guessing passwords or keys. Can be easy: many network devices have default passwords or hidden "service" accounts.
- **Software bug attacks** exploiting bugs in particular network server (or client) program versions. Very common mode of attack, most incidents on CERT are because of program bugs.

- 17 -

References

- Surveys of network attacks and defences are in [Bel98, CBR03] and Chapter 18 of Anderson [And01]. Of course, there's a wealth of information on the net: see links on the course web page for starting points.
- [And01] Ross Anderson. *Security Engineering: A Comprehensive Guide to Building Dependable Distributed Systems*. Wiley & Sons, 2001.
- [Bel98] Steve Bellovin. Network and internet security. In Dorothy E Denning and Peter J Denning, editors, *Internet Besieged — Countering Cyberspace Scofflaws*. ACM Press, 1998.
- [CBR03] William R Cheswick, Steven M Bellovin, and Aviel D Rubin. *Firewalls and Internet Security Second Edition: Repelling the Wily Hacker*. Addison-Wesley, 2003.

- 18 -