

Network and Internet Security

A Musical Tour

Steven M. Bellovin

AT&T Labs -- Research

smb@research.att.com

<http://www.research.att.com/~smb>

“There but for fortune could go you or I.”

Phil Ochs

The Net is a Bad Neighborhood

- ◆ There are hackers out there; more and more everyday.
- ◆ They're getting better at it:
 - Canned templates to exploit buffer overflows.
 - Implementations of active attacks.
 - Cryptographic protection for their tools.
 - They're not playing crypto games because there's too little to attack -- thus far.
- ◆ We don't know for whom they work.

Who Are the Hackers?

- ◆ Many are joy hackers.
- ◆ Some businesses report targeted attempts.
- ◆ Several political protests (including U.S. government sites).
- ◆ Are foreign governments involved, beyond the “Cuckoo’s Egg” incident?
 - Given “Eligible Receiver”, some governments would be negligent to ignore preparations.

“You never give me your money”

The Beatles

Hacking for Profit

- ◆ A vendor reports prices changed on a Web page.
- ◆ One ISP was hacked by a competitor
- ◆ At least two customers on pay-per-packet nets were targets of packet storms.

Denial of Service Attacks

- ◆ Attacks don't break in, but they deny you access to your own resources.
- ◆ Several recent incidents reported; more are likely.
- ◆ Defending against such attacks is *very* hard. If it's cheaper for the attacker to send a message than for you to process it, you lose.

Denial of Service Attacks

- ◆ SYN flooding
- ◆ “Smurf”
- ◆ “Teardrop”
- ◆ “Land”

SYN Flooding

- ◆ Bombard a host with TCP open request packets, from non-existent sources
- ◆ Half-open connection queue fills up; legitimate open requests are dropped.
- ◆ Mostly solved: use cheaper data structure for queue, plus random drop when queue is full.

The “Smurf” Attack

- ◆ Attacker sends “ping” to intermediate network’s broadcast address.
- ◆ Forged return address is target machine.
- ◆ All machines on intermediate network receive the “ping”, and reply, clogging their outgoing net and the target’s incoming net.
- ◆ Firewalls at target don’t help -- the line is clogged before it reaches there.

“Teardrop” and Related Attacks

- ◆ Teardrop
 - Send overlapping IP fragments.
 - Destination machine doesn’t handle the overlap properly, and crashes.
- ◆ Ping of Death
 - Send very large IP packet, fragmented into many smaller ones.
 - Length wraps around, crashing target.
- ◆ Both can get through some firewalls.

The “Land” Attack

- ◆ Send TCP packet where the source and destination addresses are that of the target machine, and the port numbers match.
- ◆ Target sees this as an attempt to connect a socket to itself, and gets terminally confused.
- ◆ Can be blocked by anti-spoofing filter.

“Back In The U.S.S.R.”

The Beatles

Political Protests

- ◆ Many government sites targeted, here and abroad.
 - CIA, Air Force, Australian political party, the ACLU, Indonesian government, etc.
 - Some may be pranks by joy hackers.
- ◆ What happens when *serious* terrorist groups start going after the Net?

“Everybody's Got Something To Hide
Except Me And My Monkey”
The Beatles

Who Are the Targets?

- ◆ Popular organizations.
 - Someone always wants to take them down.
- ◆ Unpopular organizations.
 - The more enemies you have, the more trouble you're in...
- ◆ More or less anyone.
 - New folks on the net have less experience, and are easier targets.

Robbing the Poor

- ◆ 2600 Magazine has already carried stories on how to eavesdrop on cable TV-based networks.
- ◆ @Home warns against sharing file systems and printers.
- ◆ AOL hackers social-engineer passwords and credit card numbers from naïve users.

Hacker Trends

- ◆ Increased sophistication of attacks.
- ◆ Copious “cookbooks” and packaged kits.
- ◆ Great emphasis on operational security, including use of encryption.
- ◆ Most “hackers” aren’t worthy of the name.
 - A few are *very* good.
- ◆ The hackers share tools and knowledge more than the good guys do.

“It’s All Happening at the Zoo”
Simon and Garfunkel

Types of Attacks

- ◆ Sniffers
- ◆ Active attacks
- ◆ IP spoofing
- ◆ Buffer overflows
- ◆ Race conditions

Sniffers

- ◆ Password collection has been going on since at least late 1993.
- ◆ Other uses are possible:
 - NFS file handle collection
 - Credit card numbers
 - DNS spoofing

Active Attacks

- ◆ IP spoofing.
- ◆ Session hijacking possible with canned programs.
 - Requires eavesdropping ability.
 - Canned programs seem to be available.
- ◆ Cryptographic stunts.
 - None yet, but...

More Active Attacks

- ◆ DNS cache contamination
 - Exploit script widely available
 - Was once done for commercial purposes; resulted in a Federal indictment.
- ◆ False route advertisements
 - Given well-publicized accidental incidents, a deliberate version seems likely.
 - We don't have good defenses.

Routing Attacks

- ◆ Routers advertise their own local nets, plus what they've learned from their neighbors.
- ◆ Routers believe even dishonest neighbors.
- ◆ Routers further away must believe everything they hear.
- ◆ Authentication must be end-to-end, not just hop-by-hop.
- ◆ Theoretical solutions just starting to appear in the literature.

IP Spoofing

- ◆ Attack described in a 1985 paper by Morris.
- ◆ First known use against Tsutomu Shimomura -- but it's hard to detect.
- ◆ Cryptographic authentication is a strong defense, but is rarely used.
- ◆ A simpler defense has been developed, but it is not yet widely deployed.

Implications of Active Attacks

- ◆ Remote login is no longer secure, even when protected by hand-held authenticators.
- ◆ Login through a firewall is not safe, either.
- ◆ Other protocols are subject to similar attacks.

“Too Much Lovin’ ”

Ray Drew

Buffer Overflows

- ◆ C uses character arrays for strings.
- ◆ It doesn't check bounds (and the language design makes such checking hard).
- ◆ Too many programmers say “this array is big enough” -- and it is, for normal purposes...
- ◆ N.B. Technique first introduced in the Internet Worm of 1988 -- but we still see new examples.

Race Conditions

- ◆ Mostly local attacks to gain root privileges.
- ◆ Low probability of success each try -- but attempts are cheap, and the attacker only has to win once.
- ◆ Most common variety: temporary files being created in /tmp or other world-writable directory.

“Fixing A Hole”
The Beatles

What are the Causes?

- ◆ Not enough cryptography.
- ◆ Buggy code.
- ◆ Complex code -- see above.

Not Enough Cryptography

- ◆ The amount of encrypted traffic on the net is almost unmeasurably small.
- ◆ Most sites use COTS equipment; few vendors support cryptography.
- ◆ If you don't use cryptography, I can't use it to talk to you.

Misplaced Cryptography

- ◆ SSL is a security blanket: “Our Web site is secure because we use cryptography”.
 - “For your convenience, we’ll store your credit card number, too.”
- ◆ SSL is precisely the wrong layer; it doesn’t sign orders, and it requires changes to all applications.
- ◆ *But* -- it was deployable.

Bad Cryptography

- ◆ Misuse of encryption modes.
- ◆ Home-brewed ciphers.
- ◆ Pseudo-random one-time pads.
- ◆ Domestic sites (of big companies) that only accept 40-bit keys.

“Please allow me to introduce myself.”

The Rolling Stones

Worthless Certificates

- ◆ Most users don't know what certificates are.
- ◆ Most certificates' real-world identities aren't checked by users.
- ◆ Why should Dow, Jones own the `www.wsj.com` certificate? Is that certificate good for `interactive.wsj.com`?
- ◆ Is it `NASA.COM` or `NASA.GOV`?
`MICROSOFT.COM` or `MICR0S0FT.COM`?
- ◆ Effectively, we have no PKI for the Web.

Encryption

- ◆ Starting to be deployed.
- ◆ Standards still in a state of flux, though that is improving rapidly.
- ◆ Has been held up by patent issues and export restrictions.
- ◆ *Not* a panacea; an encrypted channel to a buggy program will still let hackers in.

Attacking Crypto

Ignoring the algorithm....

- ◆ How good is the PKI?
- ◆ Where do the random numbers come from?
- ◆ Will the information be resent over other links?
- ◆ Is the OS secure?
- ◆ Is the cryptographic program correct?

“Masters of War”

Bob Dylan

Why Don't Vendors Ship Encryption?

- ◆ Little demand -- users think they're safe.
- ◆ Patent issues have limited university development of free code.
- ◆ Export controls -- vendors don't want two different product sets and two different architectures.
- ◆ *Export restrictions have hurt the security of U.S. computer networks.*

“Don't Bug Me”
Jimmy Buffet

Buggy Code

- ◆ 85% of CERT Advisories describe problems that cannot be fixed with cryptography.
- ◆ Most of these are bugs in code
- ◆ But writing correct code is the oldest -- and probably the most difficult -- problem in computer science. We're not going to solve it any time soon -- and possibly not ever.

Preventing Bugs

- ◆ 30-40% of newly-reported holes are due to buffer overflows -- better languages or libraries (or programmers) can solve this.
- ◆ Structuring code properly can help -- isolate the security-critical sections.
- ◆ But Orange Book-style security kernels are obsolete -- we have too many operating systems (browsers, word processors, etc.)

WebOS

- ◆ Untrusted and mutually suspicious places are sending me programs.
- ◆ These programs are allowed access to only a few files.
- ◆ I must allocate CPU time, memory, etc., fairly.
- ◆ Conclusion: my browser is -- or should be -- an operating system.
- ◆ But my word processor has the same problem, as does my spreadsheet, my slide maker, etc. All of these are COTS.

Can We Build WebOS, WordOS, etc.?

- ◆ Hypothesis: Structuring these tools the way we do operating systems will make them more secure.
- ◆ Hypothesis: Operating systems need the capability to create “client operating systems”.
- ◆ Can we build an OS that handles multiple views of access?
- ◆ Can users manage the permissions?

File Access in WebOS

- ◆ Many programs have tried -- and failed -- to implement access controls based on file name patterns.
- ◆ Real operating systems don't rely on patterns.
- ◆ Can we map a Web server's name space into the underlying OS's permission space?
- ◆ Can we do that for a Web client?

“Your debutante has what you need
but I have what you want.”

Bob Dylan

Complex Code

- ◆ When was the last time a vendor *deleted* features when shipping a new release?
- ◆ When did you see an ad bragging that some Web browser *doesn't* have Javascript?
- ◆ People don't understand how to use what's already there -- so vendors add even more complexity to help people find the knobs and buttons.

Why is Complexity Bad?

- ◆ Complexity implies more code, and hence more bugs.
- ◆ Different pieces interact with each other; interactions grow as the square of code size.
 - Example: setuid programs + shared libraries + environment variables = hole.
- ◆ “Ship first; test later” -- first-to-market often wins the war.

The Web: Threat or Menace

- ◆ Everyone uses the Web.
- ◆ The Web is now the universal graphical interface. Applications that don't have Web-based GUIs today will by tomorrow.
- ◆ But the Web is not well-designed from a security perspective.

Web Complexity: Client Problems

- ◆ The server is telling the client what to do.
- ◆ Bogus URLs can exploit buggy code.
- ◆ Plug-ins, active content, etc.

Active Content

- ◆ Outsiders supplying code to be executed on user's machine.
- ◆ Can this code be trusted?
- ◆ Can it be contained?
- ◆ How can we give active content enough power to be useful, while still keeping it safe?
 - Can users administer fine-grained controls?

Java

- ◆ Nominally runs in a “sandbox”
- ◆ Relies on very complex model to ensure security.
 - But at least Sun did try to address the problem.
- ◆ Many bugs have been found.
- ◆ Code signatures being added.

ActiveX

- ◆ No execution-time protection.
- ◆ Sole security is digital signature.
 - Is the provider really trustworthy?
 - Was the provider hacked?
 - Was the certificate checked?
- ◆ Signatures provide accountability, not protection.

Javascript

- ◆ Javascript can do almost anything the end-user can do -- the human is out of the loop.
- ◆ No simple protection model.
- ◆ Both design and implementation bugs have occurred, by both Netscape and Microsoft.
- ◆ Java + Javascript is a particularly dangerous combination.

Web Complexity: Servers

- ◆ Complex administration: easy to get wrong.
 - It took one site I know of three tries to get even simple access controls correct.
- ◆ Complex structure
 - the servers try to validate source addresses; check passwords; parse file names; implement access restrictions; switch uids (which means they must run as root); etc.
 - Scripts...

WWW Scripts

- ◆ Scripts are, in essence, programs that provide network services. Are they secure?
- ◆ Most such scripts are written by information providers, not security specialists...
- ◆ The languages used to write these scripts are often inappropriate. Perl5, for example, has security problems.
- ◆ The existence of these scripts implies the need for these interpreters (and for programs they invoke, especially for shell scripts) to be accessible to the Web servers.

The Web and Credit Cards

- ◆ Sniffing is easy; not everyone uses encryption.
- ◆ Even if the number is protected in transit, it's sitting on a Web server, in a file accessible to a Web script...

“Tunnel of Love”
Bruce Springsteen

IPsec and Virtual Private Networks

- ◆ Firewall to Firewall
- ◆ Host to Firewall
- ◆ Host to Host

IPsec: Firewall to Firewall

- ◆ Implement VPNs over the Internet.
- ◆ Deployment already in progress; may some day largely replace private lines.
- ◆ Caution: still vulnerable to denial of service attacks.

“2000 Light Years From Home”
The Rolling Stones

IPsec: Host to Firewall

- ◆ Primary use: telecommuters dialing in.
- ◆ Also usable for joint venture partners, clients, customers, etc.
- ◆ But today's firewalls grant permissions based on IP addresses; they should use certificate names.

IPsec: Host to Host

Attractive, but...

- ◆ It's not widely available. (But NT 5.0 should have IPsec.)
- ◆ Can we manage that many certificates?
- ◆ Can servers afford it?
- ◆ Can today's hosts protect their keys?

“Your Mother Should Know”

The Beatles

Limits to IPsec

- ◆ Encryption is not authentication; we must still control access.
 - Firewalls can't peek inside encrypted packets
- ◆ Traffic engineers want to look inside packets, too.
- ◆ New techniques for handling unusual links -- satellite hops, wireless LANs, constant bit rate ATM, etc. -- require examining, replaying, and tinkering with packets.
- ◆ NAT boxes incompatible with end-to-end IPsec.
- ◆ Use key recovery technology?

“It Ain’t Me, Babe”

Bob Dylan

Naming: IPsec and Certificates

- ◆ Users specify hosts by name: `www.nsa.gov`.
- ◆ IPsec operates on IP addresses (`135.207.32.62`).
- ◆ We must use DNSsec to protect the mapping between the two. (It isn't deployed yet.)
- ◆ But IP addresses are increasingly transient, given DHCP, dial-up users, and IPv6 renumbering. How do we name endpoints?

“Light My Fire”

The Doors

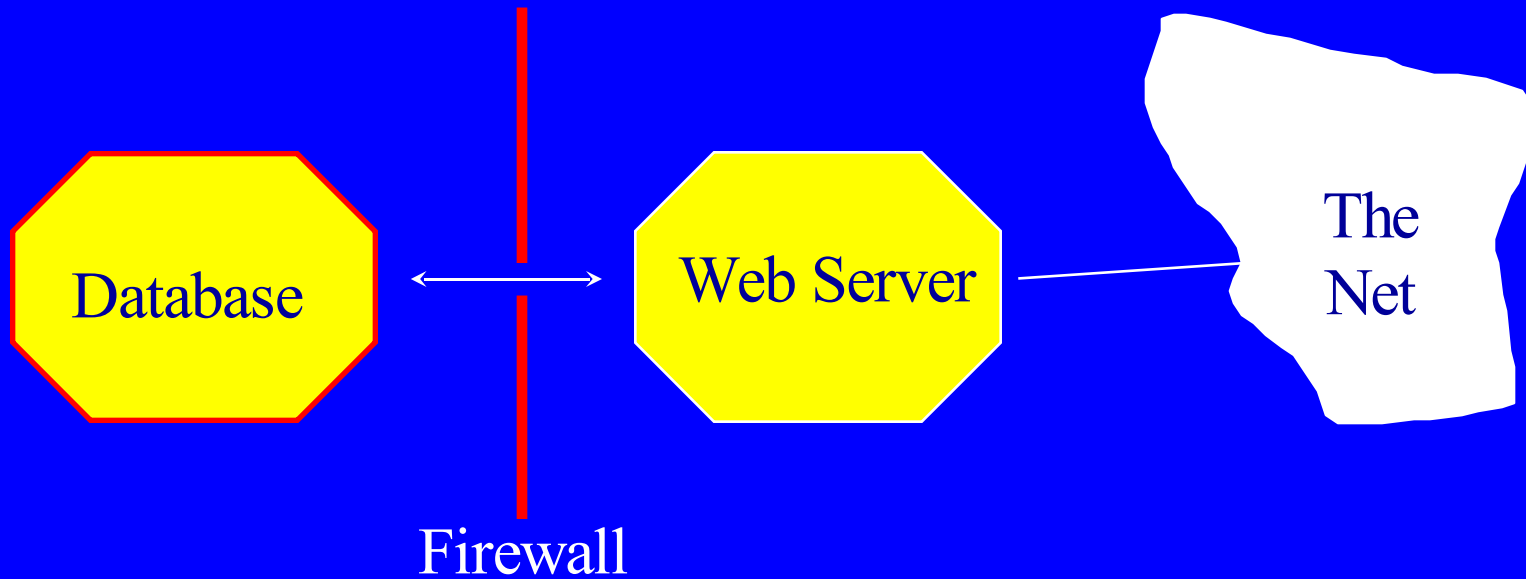
Firewalls

- ◆ A barrier between “us” and “them”.
 - “They” may be another part of the same company.
- ◆ Limit communication to the outside world.
- ◆ Firewalls work because only a few machines running a few services are exposed to attack.
- ◆ *Firewalls are the network’s response to the host security problem.*


How to Use Firewalls

- ◆ Large corporate-scale firewalls are dinosaurs.
- ◆ They are best used as one element of a total security structure.
 - Shield legacy systems and system components that cannot economically protect themselves.
- ◆ Placement is critical.

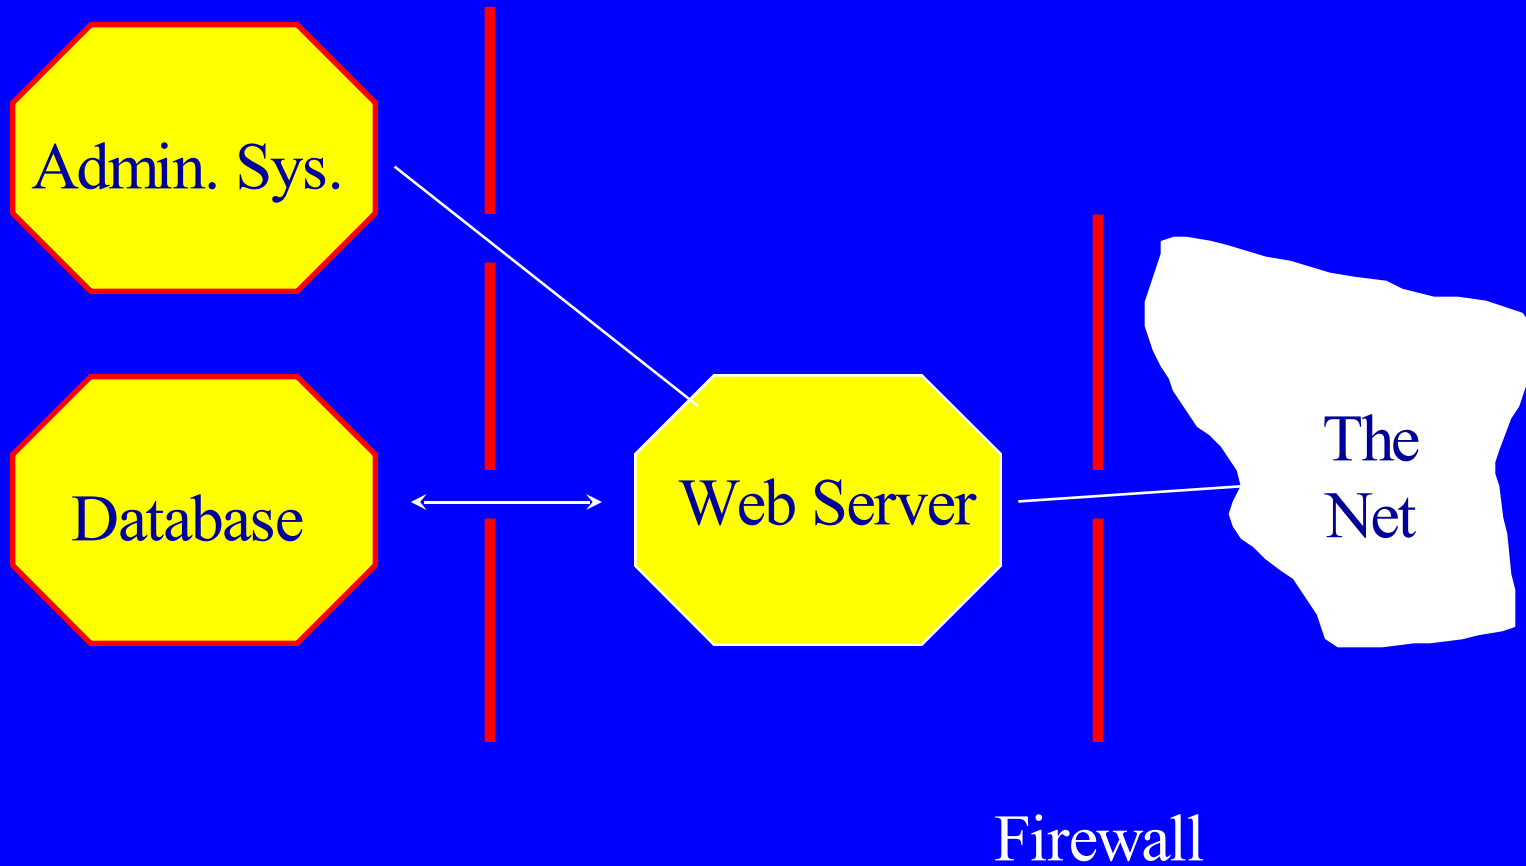
Firewalls and Databases



The Wrong Choice



Other Channels



Limitations of Firewalls

- ◆ Cannot protect against inside attacks.
- ◆ Increased interconnectivity makes attacks from inside -- though not necessarily by *insiders* -- more likely.
- ◆ Cannot block attacks at higher level of the protocol stack.

“She came in through the bathroom window”

The Beatles

Why Are Firewalls Dying?

- ◆ There is too much connectivity that bypasses the firewall.
- ◆ Too many protocols are being allowed through the firewall.
- ◆ There is too much “transitive trust” -- trust of machines that have their own connections to untrustworthy parties.

Typical Versus Secure Firewalls



“You can’t always get what you want,
but if you try sometimes
you might just find
you get what you need.”

The Rolling Stones

What Should Developers Do?

- ◆ Take security seriously
- ◆ Follow good programming practice; avoid fixed-length strings.
- ◆ Design in security from the start.
- ◆ Make security part of the schedule.
- ◆ Structure the program properly.

Structure Example: FTPD

Standard Version

- ◆ Input language is giant YACC grammar.
- ◆ Login and password checking intermixed with other code.
- ◆ Result: most of the program is security-sensitive.

How to Do It Right

- ◆ Do authentication first (100-150 lines of code)
- ◆ YACC grammar handles non-privileged commands only.
- ◆ Result: very little is security-sensitive.

What Should End-Users Do?

- ◆ “Just say no” to dangerous technology.
- ◆ Vote with your feet -- and dollars -- when purchasing software.
- ◆ Use encryption.

Should Organizations Disconnect?

- ◆ There are risks in doing anything. Even doing nothing carries risks; staying off the net is a denial of service attack on yourself.
- ◆ There are no guarantees of absolute safety.
- ◆ The trick is to *manage* the risk.

Where to From Here?

- ◆ We *must* deploy strong cryptography, as soon as possible.
- ◆ We need more secure hosts.
- ◆ Smaller, “point” firewalls will continue to be useful.