

Why Isn't the Internet Secure Yet, Dammit?

Peter Gutmann
University of Auckland

Introduction

Standards for secure email and network communication
have been around for years

Every Windows desktop has S/MIME, SSL, IPsec

US export controls are (effectively) gone

Why isn't the Internet secure yet?

- Existing mechanisms are too hard to use
- Existing mechanisms solve the wrong problem

What this talk won't cover

Viruses / worms / trojan horses

Buffer overflows / rootkits / IIS hacks

Intrusion detection / network monitoring

Spam filtering

DRM / TCPA / DMCA

Microsoft

What this talk will cover

User identification

(Some aspects of) spam, along with phishing,
viruses/trojans

- This is a blended threat that goes beyond pure spam

Credit card fraud

Opportunistic encryption

Key continuity management

User Identification /Authentication

Allow users to sign up for online information (mailing lists, web sites)

- Fraudsters sign up in other people's names
 - Used for DoS, not just pure fraud
- Bots sign up large numbers of addresses for spam purposes

Email-based Identification

Use the ability to receive mail as a form of (weak) authentication

- Sign up using an email address
- Server sends an authenticator to the given address
- Address owner responds with the authenticator to confirm the subscription

Widely used for password resets, mailing list subscriptions

- Good enough unless the opponent is the ISP

Email-based Identification (ctd)

Self-auditing via email confirmation

- Attempting to use the account results in the legitimate owner being notified
- Changing the email address should result in a notification being sent to the original address

Outlook

- More of the same
- Low-value authentication, but relatively difficult to defeat

Identifying Humans

Prevent bots from signing up for online accounts

- Reverse Turing test
 - Turing test: Can't distinguish between human and machine
 - Reverse Turing test: Distinguishes between human and machine
 - Also known as a Human Interactive Proof (HIP)

Reverse Turing Test example


- Display distorted/noisy picture of a word
- User has to enter the actual word
- Text recognition in the presence of noise is extraordinarily difficult for computers



Identifying Humans (ctd)

Harness humans to defeat it

- Pay people in third-world-country Internet cafes to perform the text recognition
- Grant access to porn sites in exchange for performing text recognition



Outlook

- RTT will at least slow the flood
- Not a perfect solution, since you can always use real humans to defeat it

Spam

Spam mechanisms

- Open relays
- Pink contracts
 - Expensive, but avoid ToS
- Gypsy accounts
 - Set up, spam, leave
- Wireless drive-by spamming
- Hacked PCs

Spam (ctd)

```
procedure arms_race
  cobegin
    write spam_filter;
    acquire spam_filter;
    tune spam to avoid filter;
  coend
```

Anti-Spam Legislation

You can call me spam queen, I don't really care. As long as I'm not breaking any laws, you don't have to love me or like what I do for a living

— Wall Street Journal spammer interview

Theory: Some spammers may not want to become criminals just to send spam

- Much spam already violates the law (every kind of fraud, indecency laws, electronic trespass/hacking, etc etc)
- Like passing a law requiring bank robbers to wear ties

Anti-Spam Legislation (ctd)

What is spam?

- “I’ll know it when I see it”

Defining spam also (implicitly) defines non-spam

- Spammers will alter spam to qualify as non-spam
- Legislation will (inadvertently) *legitimise* spam!

Anti-Spam Legislation (ctd)

Opt-in legislation is only marginally better

- At what point does a sender require your permission to send mail?

US corporates are experts at end-runs around opt-in legislation

- “We won’t sell your details to third parties
– Rent or exchange them
- Medical insurers: “Sign this unlimited waiver or go elsewhere for medical care”
- These techniques are easily translated from snail to email

US CAN-SPAM Act

Spammer-endorsed spam legislation

- You know this one's going to be good...

Spam volume *increased* after the law was passed

- Use non-promotional content to classify it as non-spam
Fact of the day
BUY V1AGRA NOW!
- Provide (invalid) unsubscribe link
 - Use it to validate live email addresses
 - Use it to plant viruses/trojans
Click here to become infected

US CAN-SPAM Act (ctd)

- Many variations of unsubscribe trickery exist
 - Opt-out address actually works, but simply returns a request to submit the request in writing (“the unsubscribe facility is temporarily unavailable”)
- Disguise the snail-mail address
 - Use HTML rendering tricks to create legitimate-appearing addresses that can't be seen by filters

At best these features are completely ineffective

- At worst they aid the spammers (unsubscribe trickery)

US CAN-SPAM Act (ctd)

Compliance with CAN-SPAM

- January 2004 Introduced
- April 2004 3%
- June 2004 1%
- August 2004 0.5%

Also known as the I-CAN-SPAM Act

US CAN-SPAM Act (ctd)

National legislation pre-empts stricter state legislation

- Its main effect is to legitimise spam designed to exploit it
- High-profile initial prosecutions will encourage pseudo-compliance to avoid further prosecutions

– Like antibiotics creating super-bugs

– Actually it isn't even encouraging pseudo-compliance...

CAN-SPAM is a toothless tiger that nullifies most aspects of every state's anti-spam legislation and leaves spam victims without meaningful legal recourse

— Dan Appelman, Heller Ehrman, White & McAuliffe, LLC

UK Spam Legislation

UK law is no better

- Riddled with loopholes
Spammers can fill up people's work email with adverts for Viagra, child porn and money laundering scams without their permission
— Guardian newspaper
- European spam gangs are moving to the UK because of its spam-friendly environment
 - A number of major spam brokers are based in the UK because of thisA fine has never been handed down and, according to insiders, is unlikely to be
— Guardian newspaper

The Spam Business

Buy CDs with harvested addresses

- Prices vary depending on the quality
- Vacuum-cleaner for ~\$50, verified for \$x00

Send mail via spam brokers

- \$1 buys 1000–5000 credits
- Broker handles spam distribution via open proxies, relays, compromised PCs, ...
 - Sending is usually done using broker-provided software from the client PC, using information from the broker
- Credit is deducted when spam is accepted by the target MTA

This is a completely standard commercial business

- They even have their own trade associations

“I have a solution...”

Build a wall around the Internet and only let the good guys in

Can never work: In order to have perimeter security, you first need a perimeter

- The Internet is 800 million Manchurian candidates waiting to activate
- Largest observed bot-net had 11,000 members
- In late '04 these were growing at 30,000 machines per day
- Peak rate was 75,000 per day during the MyDoom/Bagle virus group wars

We have met the enemy and he is us...

Spam Technical Mechanisms

The net needs to be secure at every point

- Nothing in here but us chickens...
- Spammers already routinely break into legitimate user's PCs to send spam

Email security firm MessageLabs reports that

- *Two thirds* of the spam it blocks is from infected PCs
- Much of the spam comes from ADSL/cable modem IP pools
- Distributed Server Boycott list reports 350,000 compromised hosts on the US RoadRunner network alone

Spam Technical Mechanisms (ctd)

Worms act as special-purpose spam relays (e.g. Backdoor.Hogle, MyDoom.*)

- Infected PCs (“fresh proxies”) are traded in spammer forums

Worms act as reverse HTTP proxies

- Provide a distributed fault-tolerant “web site” for spammers
- Backdoor.Migmaf changed the “site” every 10 minutes

Other worm functions

- Disable anti-virus/firewall software (ProcKill, Klez)
- Block access to anti-virus vendor sites (MTX)
- Re-enable unsafe defaults in software, e.g. MS Office (Listi/Kallisti)

Spam Technical Mechanisms (ctd)

Other compromised host functions

- Email address harvesting (several)
- DDoS on spam-blockers (numerous)
- Run a SOCKS proxy for spammers (BID 9182 MSIE hole)

Monoculture paper: Computing power is moving to the (insecure) web periphery

- Centralised vulnerable servers → distributed, hard-to-hit servers

Convergence of Spam and Virus Threats

Publicity virus: Written by bored script kiddies

- Poorly tested, often barely work

Spam/phishing virus: Written by paid profession programmers

- Well-tested, can be quite sophisticated
- Spam vendors are employing professional linguists to bypass filters
 - They have better experts than we do!

Virus writers provide a continuously-replenished supply of compromised servers

Phishing Attacks

Have gone from amateurish to very sophisticated, professionally-run operations

- Perfect copies of the original site
- Links lead back to the real site

Convince users to go to the fake site

- Promotions: \$1,000 top-up for your account, reduce your home mortgage rate, win a car, ...
 - Mirror existing promotional efforts by banks
- Update/verify your account information
- Visit our new, more secure web site (!!)

Phishing Attacks (ctd)

Anti-virus firm MailFrontier reports that

- 28% of test subjects fell for phishing email
- 20% regarded genuine email confirming a purchase as fake

Phishing Attacks (ctd)

Attacker controls the DNS

- Server compromise
- Bribing/blackmailing ISPs
- Virus changes the victim's DNS server entries
 - Can be used to disable security updates
 - (Fake) windowsupdate.com: Your system is up to date and doesn't need any security fixes

Phishing Attacks (ctd)

Trojans control the victim's PC

- Sniff keystrokes, mouse clicks, images of graphical “virtual keyboards”
- Render copies of genuine bank pages from the browser cache

Trojan installs itself as a browser help object (BHO)

- Watches for access to a who's who of banking sites around the world
- Captures banking details before they go into the SSL layer

A Lesson from History: The Numbers Racket

The numbers racket = Lotto before the government took it over

- Bets were for cents
- Choose a 3-digit number
- “Drawn” using last 3 digits of the total amount bet on pari-mutuel racetrack betting machines
- Run through barber shops, groceries by local operators

Seen as a harmless vice, no-one paid much attention to it

A Lesson from History (ctd)

Then organised crime moved in...

- Dutch Schultz took over from existing operators
- They weren't career criminals and were intimidated by explicit death threats

Dutch hired mathematician Otto "Aba Daba" Berman to fix the numbers racket

- Ensure that heavily-played numbers never won
- No-one had ever considered this level of attack
 - c.f. spammers hiring professional linguists

Once organised crime became involved, everything changed

- A trivial/nuisance became a major criminal enterprise

Political Problems with Build-a-Wall

Everyone would have to upgrade to modern email clients

[Security plugins] are better than a root canal, but not better than a regular filling

— Matt Hamrick, Cryptonomicon

Who manages access control?

- Blacklist could block thousands of domains (e.g. a Class C block) because of one open relay
- Existing blacklists work because they're voluntary
- System requires policies, appeal procedures, etc etc etc

Requires a global secure access control mechanism

- Something like a PKI, but not as simple

Political Problems with Build-a-Wall (ctd)

Unheard-of restriction on freedom to communicate

- Even the most oppressive regimes still allowed you to send letters

Unacceptable overhead for authentication

- Existing mail systems barely cope

Other Problems

Closed communities

- Refuse to accept mail from someone you don't know
You don't exist, go away
- Need to predict in advance everyone who'll ever send mail to you
 - Change of address
 - Using someone else's PC to send mail
 - Sales inquiry
 - etc etc etc
- Single large list: Spammers will be on it
- Many small lists: Too hard to manage, no-one can talk unless you're on the "good list"

Build-a-Wall Example: Reverse Lookup

Prevents use of (outright) forged addresses

- Various proposals exist
 - Designated Mailers Protocol (DMP)
 - Reverse Mail Exchanger (RMX)
 - Sender Permitted From/Sender Policy Framework (SPF)
- Assumes that domains are associated with static IP addresses
 - Many systems can't provide the requested reverse MX
 - Vanity domains
 - Mobile users
- Various kludges possible
 - Controlled relaying for known hosts/users

Build-a-Wall Example : Reverse Lookup (ctd)

Predictions of SPF et al (in)effectiveness

- Myself: Will be rendered ineffective in 6-12 months
 - Spammers will simply move to compromised known hosts
- Anti-virus researcher: Will be ineffective within weeks
 - Spammers can adapt far more quickly than that

We were both wrong...

- Spammers are adopting SPF faster than legitimate users
 - More spam passes SPF checks than legitimate mail
- It failed in negative time!

Checking identity papers is a complete waste of time. If anyone can be counted on to have valid papers, it will be the terrorists

— Colonel Mathieu, “The Battle of Algiers”

Build-a-Wall Example : Reverse Lookup (ctd)

The work [is] an excellent example of how to not design security protocols. This was all about marketing, commercial interests, patent claims, giving interviews, spreading wrong information, undermining development, propaganda. It completely lacked proper protocol design, a precise specification of the attack to defend against, engineering of security mechanisms. It was a kind of religious war

— Hadmut Danisch on the cypherpunks list

Sender-pays Mail

Spam is effective because it's free

- To make it less effective, make it non-free

Hashcash

- Sender: I have some mail for you
- Receiver: Please submit the solution to the following problem
 - Receiver computes in $O(1)$ time
 - Sender computes in $O(1000)$ time
- Receiver-controlled rate limiting
- Sender pays in CPU time to send mail
 - Requires large bot-nets to defeat
 - Needs to be used in conjunction with whitelists for mailing lists

Sender-pays Mail (ctd)

Only works if everyone does it

- The fax machine effect
- Need to convince sendmail, Microsoft, qmail, Postfix to implement it
 - Others would be forced to follow

Who manages the billing?

- 15 years of work on micropayments haven't produced any (practically) useful results

Breaks mailing lists

- Use white-lists for trusted partners
- Drop unpaid mail into quarantine

Sender-pays Mail (ctd)

Proof of resource consumption just wastes resources

- Cycles should be applied usefully
- “I'll only talk to you on the phone if you prove you've burned a \$20 note”

Attack: Use someone else's CPU time to send your spam

- Hordes of bots/zombies sending spam for you

Users really, *really* hate paying for email

- Email is effective because it's free
- There's a reason why everyone uses email and not Telex, EDI, Compuserve, ...

Virus Throttles

Limits the damage caused by compromised hosts

- Limit outgoing connections to 0.5–1 connection per second (cps)
 - Code Red ran at 200 cps
 - SQL Slammer peaked at 30,000(!) pps (using UDP)
- Suspend programs that make too many connections at too high a rate

HP Labs studied this in great detail

- No noticeable effect on users
 - Only sites like ad servers were affected
 - Non-MSIE browsers block these sites anyway
- Virus damage was massively reduced

Virus Throttles (ctd)

THIS SHOULD BE MICROSOFT'S #1 PRIORITY SECURITY FIX

- Electronic equivalent of a firebreak: You'll never be able to prevent the problem, but you can at least limit the damage when it occurs
- MSIE was 5 years behind everyone else in supporting ad blocking
- Adding virus throttling would be an admission that Windows is a petri dish
- MS attempted to add it in Windows XP SP2, but got it wrong (the "4226 bug")
 - Broken implementation will train users to disable it
- Maybe we'll finally see it in Longhorn...

VoIP Spam

Use compromised PCs to phone out sales pitches

- Outlawed in the US over POTS lines
- Marketers would *love* to have this re-enabled
- VoIP spam isn't covered by current regulations such as do-not-call lists

Qovia's proof-of-concept VoIP mass-caller makes 1,000 synthetic calls every five seconds

VoIP mail boxes dutifully record every message that they receive

- Humans will hang up within seconds
- VoIP providers will need to massively expand storage to store VoIP voicemail spam

VoIP Spam (ctd)

Blind net users are already affected by this

- Deleting spam via text-to-speech is unworkable

VoIP could use a Reverse Turing Test to weed out automated callers


“Please enter the following four-digit number...”

- No worse than the usual voice-mail maze

Instant Messaging spam (spim)

- Much like email spam, sent from compromised PCs, etc

Spam-fighting Timeline



Current situation can only be addressed via legislation

- Neither users nor vendors have any natural incentive to fix things

Spam Threat Model

Spam comes from legitimate (if unwitting) users

→ Any us-vs-them approach is doomed to failure

Spammers operate from jurisdictions where prosecution is unlikely/impossible

→ Going after the spammers directly is unlikely to be very effective

Solution to the threat is to address inadvertent spamming via legitimate users (open proxies, compromised hosts, etc)

Effective Anti-spam Legislation, Part 1

Most spam is sent via unauthorised channels

- The dress-code-for-bank-robbers legislative approach will never work

Pass legislation to close the unauthorised channels

- Most people only wore seatbelts/helmets when it was required by law
- Penalise vendors for selling spam-enabling software (MS Outlook, via viruses/worms)
- Penalise users for running software in a spam-enabling manner (open relays)
- Equivalent to existing corporate governance legislation (auditing/environmental/due diligence/etc requirements)

Effective Anti-spam Legislation, Part 2

Today's dumb-terminal equivalent is more capable than the departmental server of 10 years ago

- 90+% of them are used as little more than dumb terminals (Web3270's)

Example: "Blaster Revisited", ACM Queue magazine

- The task: Electricity-bill payment terminal
 - Enter name, address, amount, hit Enter
- The tool: Windows XP Pro with all network services running on a live Internet connection
 - The required tool: a VT52 (1970's dumb terminal)
 - (A model 33 teletype or an 029 keypunch would do too)
- The solution: Change the oil in the ambulance at the bottom of the cliff and wait to be hit again

Effective Anti-spam Legislation Outlook

Unlikely to pass in the US due to software industry lobbying

- Would require an Enron-style debacle to pass

This is a social problem that can't be fixed using technology

- Technical solutions are a band-aid on a sucking chest wound



No solution in sight

Online Credit Card Purchases

Convince consumers that it's safe to buy online

- SSL protects credit cards in transit
- If you see the padlock, you're safe

Security theatre

-  Vendors put padlock GIFs on their web sites to provide extra reassurance 
- Protects against a man-in-the-middle (MITM) attack
 - Zero recorded instances in 10 years of online credit card use
- No protection for cards once they're at the merchant's server
- Black-market price for 1,000,000 stolen credit cards
 - A few years ago: \$100
 - Today: < \$1

Obtaining CC/banking Information

Crooks obtain user credit card/banking information by

- Breaking into poorly-secured servers
 - Large-scale, tens to hundreds of thousands of cards collected
- Phishing/social engineering
 - Set up a fake web site, convince users to hand over details
 - Medium-scale, more effort required but can collect more information
- Dishonest restaurant/bank/hotel employees
 - Very small-scale, but provides the most information

Small-scale Credit Card Fraud

Information distributed via web sites

Card checks performed via IRC bots

- `!chk cardno expiry`
- `!cclimit cardno`
- `!cvv2 cardno expiry`
 - CVV is the 3-4 digit crypto checksum on the back of the card
 - Required as an extra check by some merchants
- This is more sophisticated than many merchants!

Small-scale Credit Card Fraud (ctd)

Complete identity (card number, expiry date, CVV, name, address, Social Security Number (SSN), mother's maiden name) sells for ~\$10

- Sellers claim to work for banks, hotels, restaurants

User identities are hidden via IRC proxies (bouncers) on hacked PCs

The trade of BotNets on compromised machines is becoming an industry in itself. Organised crime is making use of this industry

— Detective Chief Superintendent Les Hynds,
head of the UK National Hi-Tech Crime Unit

Large-scale Credit Card Fraud

See “Small-scale fraud”

Obvious: Crime rings get 25 PCs shipped to eastern Europe

- Countermeasure: Merchants refuse to ship internationally
- Merchandise is shipped via US middlemen
 - “Earn big bucks working from home!”

Large-scale Credit Card Fraud (ctd)

Less obvious: Use online auctions for money laundering (triangulation)

- Advertise new \$1000 digital camera on ebay for \$800
- Buy with stolen card, get sent to ebay buyer
- Collect \$800 cash
- Buyer countermeasure
 - Watch out for auctions asking for cash-equivalents (money orders, wire transfers)
 - Ask for product serial number before buying (requires middleman to hold payment)
- Merchant countermeasure
 - Require shipping address on file with issuing bank

The problem with Credit Card numbers

Credit cards confuse identification and authorisation

- Credit card must be both public (identification) and private (authorisation)


Properly-designed mechanisms separate the two

- Username and password
- London Underground ID card and pass ticket

Credit Card security mechanisms should enforce this separation

Preventing Fraud (in theory): SET

SET (Secure Electronic Transactions)



Acquirer gateway is an Internet interface to the established credit card authorisation system and cardholder/merchant banks

SET (ctd)

Card details are never disclosed to merchant

- Encrypted purchase instruction (PI) can only be decrypted by the acquirer
 - In practice the acquirer usually reveals the card details to the merchant after approval, for purchase tracking purposes
- PI is cryptographically tied to the order instruction (OI) processed by the merchant
- Client's digital signature protects the merchant from client repudiation

Authorisation request includes the consumer PI and merchant equivalent of the PI

- Acquirer can confirm that the cardholder and merchant agree on the purchase details

SET (ctd)

Capture can take place later (e.g. when the goods are shipped)

- User can perform an inquiry transaction to check the status

The whole SET protocol is vastly more complex than this

Why SET failed : Complexity

SET is the most complex (published) crypto protocol ever designed

- > 3000 lines of ASN.1 specification
- 28-stage (!!) transaction process
 - “The SET reference implementation will be available by, uhh...”
 - Although SET was specifically designed for exportability, you couldn’t export the reference implementation
- Interoperability across different implementations is a problem
- SETco charged a huge amount of money for compliance testing of implementations
 - Hard on small companies, who were doing the implementation work

Why SET failed : Practicality

Huge numbers of merchants use the credit card number as a primary key for customer databases

- “Solved” by making the card number visible to merchants
- Defeats the major purpose of SET (protecting the CC number)

SET requires

- Custom wallet software on the cardholders PC
- Custom merchant software
- Special transaction processing software (and hardware) at the acquirer gateway

Some (small-scale) fraud is still possible

- Break into merchant, steal CC → Break into client PC, steal SET keys

Why SET failed : Politics

All the liability was carried by the issuing bank

- All the benefit was obtained by the acquiring bank
- Some attempt to mitigate this by splitting the costs

VISA / MasterCard didn't care if SET succeeded or not

- SET was a counter to Cybercash, Mondex, etc
- When those didn't go anywhere, SET was superfluous
- Credits cards over SSL were seen as far more profitable, since they're charged as card-not-present transactions

Avoiding Fraud (in practice)

Merchants will only ship to the CC billing address

- Sledgehammer approach inconveniences many customers
- Fraudsters haven't had to seriously attack this measure yet

Attempts to resurrect SET

- Verified by VISA — roll-your-own SET
- Everyone gets to independently reinvent the wheel...
... badly
- Design target seems more to impress VISA's auditors than to provide real security

Avoiding Fraud: Outlook

Banks aren't too worried, merchants carry the cost

- Consumers pay via increased prices

Each step back buys 1-2 years

- Ship-to-billing-address is the last line of defence

TAN-based approach

- One-time password per transaction

~~UZDCOQwG~~

~~XeSnvszE~~

0uZVSJ2U

– Send new TANs when the old list is about to expire

Avoiding Fraud (ctd)

Used by European banks for online banking

- Marginal cost is close to zero
 - TANs sent out with bank statements
- Remarkably effective against online credit card theft/fraud
 - The one thing you can't do online is intercept paper mail
- Requires participation by banks
 - Non-European banks haven't got past username + password (or SSN, mother's maiden name, ...)
 - Currently the pain isn't sufficient to motivate changing the CC authorisation system

Avoiding Fraud (ctd)

Still not 100% effective

- Virus reads bank's page from browser cache
- Pops up window asking user to re-authenticate due to session timeout
 - Users are conditioned to accept this
 - Too many banks use Javascript pop-ups, aggressive session timeouts during online banking
- Username + TAN go to eastern Europe, user's session continues as normal

European banks are switching to challenge-response calculators in response to this type of attack

Opportunistic Encryption

After 10-15 years effort, S/MIME and PGP use is lost in the noise floor

- Most mail clients include S/MIME support
- Many (OSS) clients include PGP support
- Usage is virtually nonexistent
 - It's too much bother for most people

The vast majority of users detest anything they must configure and tweak. Any really mass-appeal tool must allow an essentially transparent functionality as default behaviour; anything else will necessarily have limited adoption

— Bo Leuf, “Peer to Peer: Collaboration and Sharing over the Internet”

Opportunistic Encryption (ctd)

Encrypt data using keys managed via key continuity (see next section)

- Completely transparent to end users
- Requires no extra effort to use
- Effectively free (except for the slight CPU overhead)

Most commonly encountered in SMTP/POP/IMAP

- Protects mail in transit
- Authenticates sender/prevents unauthorised relaying/spamming

STARTTLS/STLS/AUTH TLS

What is it?

- Opportunistic encryption for SMTP/POP/IMAP/FTP
220 mail.foo.com ESMTP server ready
EHLO server.bar.com
250-STARTTLS
STARTTLS
220 Ready to start TLS
<encrypted transfer>
- Totally transparent, (almost) idiot-proof, etc
- Protects more mail than all other email encryption protocols combined

STARTTLS/STLS/AUTH TLS (ctd)

Outlook

- A year after appearing, STARTTLS was protecting more email than all other email encryption protocols combined, despite their 10-15 year lead
- Just as SSH has displaced telnet, so STARTTLS may displace (or augment) straight SMTP
 - Auckland Uni turned off unencrypted mail to local servers after STARTTLS appeared, just as they turned off telnet after SSH appeared
- Not perfect, but boxes attackers into narrower and narrower channels

Key Continuity Management

Where's the PKI?

It's too...

- Expensive
- Complex
- Difficult to deploy
- Doesn't meet any real business need
- etc etc etc

Key Continuity Management (ctd)

The only visible use of PKI is SSL

- This is certificate manufacturing, not PKI
- Once a year, exchange a credit card number for a pile of bits
- Three quarters of *all* SSL server certs are invalid (SecuritySpace survey, December 2003)
- No-one notices...

Assurance through Continuity

Continuity = knowing that what you're getting now is what you've had before/what you were expecting

- McDonalds food is the same no matter which country you're in
- Coke is Coke no matter what shape bottle (or can) it's in, or what language the label is in

Continuity is more important than third-party attestation

- Equivalent to brand loyalty in the real world
- Businesses place more trust in established, repeat customers

Use continuity for key management

- Verify that the current key is the same as the one you got previously

Key Continuity in SSH

First app to standardise its key management this way

- On first connect, client software asks the user to verify the key
 - Done via the key fingerprint, a hash of the key components
 - Standard feature for PGP, X.509, ...
- On subsequent connects, client software verifies that the current server key matches the initial one
 - Warn user if it changes

Concept was formalised in the resurrecting duckling security model

- Device imprints on the first item it sees
- Device trusts that item for future exchanges

Key Continuity in SIP

Same general model as SSH

- First connect exchanges self-signed certificates
- Connection is authenticated via voice recognition

Same principle has been used in several secure IP-phone protocols

- Users read a hash of the session key over the link

Key Continuity in STARTTLS et al

SMTP/POP/IMAP servers are usually configured by sysadmins unconcerned about browser warning dialogs

- Remember the initial certificate, warn if it changes
- Using self-signed certificate avoids having to pay a CA

The ideal key continuity solution

- Automatically generate self-signed certificate on install
- Use key continuity to warn if the certificate changes

Key Continuity in STARTTLS et al (ctd)

Currently still somewhat haphazard

- Many open-source implementations support it fully
- Some still require tedious manual operations for certificate management
- Commercial implementations often require CA-issued certificates, an even more tedious (and expensive) manual operation

Key Continuity in S/MIME


S/MIME has a built-in mechanism to address the lack of a PKI

- Include all signing certificates in every message you send
- Lazy-update PKI distributes certificates on an on-demand basis

S/MIME gateways add two further stages

- Auto-generate certificates for new users
- Perform challenge-response for new certificates they encounter

Key Continuity in S/MIME (ctd)



Msg0 gets remote certificates to local server (as standard S/MIME message)

Msg1 gets local certificates to remote user

Msg2 proves possession of remote server keys/certificates
(Variants, e.g remote server sends challenge in Msg3)

Key Continuity in S/MIME (ctd)

- Provides mutual proof of possession of keys and certificates to both sides
 - In practice has a few extra tricks to avoid various attacks
- Both parties now have verified keys for the other side
- Fully automatic, no human intervention required

Outlook

- Invented/reinvented as needed by implementors
 - Not specified in any formal standard
 - Standards groups are still waiting for PKI to start working
 - Present in many apps, but needs standardisation to unify approaches

Conclusion

Simple human-in-the-loop solutions can be remarkably effective against large-scale automated harvesting attacks

Spam is a social problem that can't be fixed using technology

- Current attempts to fix it via legislation are ineffective or nonexistent

Card fraud: The banks had better come up with something quickly

Opportunistic encryption has achieved more penetration in one year than traditional methods did in 10