## Hard and Not-necessarily-hard Problems in Cryptography

Peter Gutmann University of Auckland

#### Hard Crypto Problems

Some crypto problems have no known general solution

• This may be the first talk ever to admit that there exist security problems for which adding more cryptography isn't the answer

Why are you telling us about them if there's no solution?

- To warn you about them so you can try alternatives
- To let you know that if you're finding it hard to deal with them then it's not your fault

# Example: Secure Bootstrapping of Comms How do you securely initiate communications with an entity that you've never communicated with before? The killer problem The elephant in the room The mixed metaphor of Internet security protocols We simply have no way of doing this











#### Wicked Problems

So why is this so hard?

This (and many other issues) are examples of wicked problems

• Concept from the field of social planning



• Proposed in the 1970s as a means of modelling the process for dealing with social, environmental, and political issues

#### Wicked Problems (ctd)

Amongst a wicked problem's weaponry are such diverse elements as...

Lack of any definitive formulation of the problem

Lack of a stopping rule

• One of the core requirements for dealing with a wicked problem becomes not deciding too early which solution you're going to apply

Solutions that are rateable only as "better" or "worse" and not true or false

- Particularly bad for security geeks
- There are only two options, absolutely secure or absolutely insecure



#### Wicked Problems (ctd)

A wicked problem presents...

- No clear idea of what the problem is
- No clear idea of how to get to a solution
- No easy way to tell whether you've reached your goal or not
- All of the participants are pulling in different directions

#### Example: High-performance Sports Cars

Fit a more powerful engine



- Adds extra weight
  - Slows it down again
- Adds size
  - If taken to extremes leaves little room for anything else, including a driver



## Example: High-performance Sports Cars (ctd)

Use exotic materials like carbon fibre to decrease weight



• Raises the price and again discourages buyers

















#### Example: Crypto Woo-Woo (ctd)

"Smart" meters

- Digital signatures!
- X.509 certificates!
- CRLs!
- The whole PKI shebang!

MSP430F148 CPU

- 8 MHz 16-bit CPU
- 16-bit multiplier as external functional unit (no divide)
- 2kB RAM, 48kB flash



• Additional analog/digital circuitry for a power meter You can guess how much PKI this actually implements...



#### Wicked Problems

This perfectly illustrates the characteristics of a wicked problem...

No definitive formulation of what's required for a sports car

No stopping rule to tell you that you've definitely reached your goal

• Running out of money is one oft-encountered stopping rule

The various options can only be rated in terms of tradeoffs against each other

continues...

#### Wicked Problems (ctd)

#### ...continued

It's not obvious which steps are the best ones to take in getting to your goal

## All manner of externalities

- Participants' opinions of which option is best
- Bikeshedding comes as an automatic built-in



• Externally-applied materials and regulatory constraints on what you can and can't do

#### Problem: Secure Ops on Insecure Systems

Trying to perform safe operations on an untrusted system has historically been mostly of academic interest

- "Programming Satan's Computer", Anderson and Needham, 1995
  - Satan's Computer in 1995 was positively benign compared to what's hiding inside many current PCs
- On the off chance that your machine gets compromised, reformat and reinstall from clean media

Today the sheer scale and scope of the problem has made this approach all but impossible



Problem: Secure Ops on Insec.Systems (ctd)
What about trusted computing?

Yeah, any year now...

Even if it could be deployed, it can protect only a small part of the system, typically the OS core
A fully-protected computer on which you can't make any changes isn't terribly useful





#### Problem: Secure Ops on Insec.Systems (ctd)

That's between one and two million bugs in the OS

- Ignores the additional code that'll be added in the form of userinstalled device drivers and other kernel components
  - A majority of Windows OS crashes are due to these additional drivers

Ignores the perpetual churn of updates

• TPMs weren't really designed for a constantly-changing code base

So your TPM-verified boot guarantees that you're loading an OS core with only a million bugs

• As opposed to a tampered one with a million and one bugs



that's done on a PC and treat it purely as a router?

• Forward encrypted/authenticated content from a remote server to a self-contained device via USB or Bluetooth or NFC

#### Problem: Secure Ops on Insec.Systems (ctd)

This "solution" gets re-invented every six to twelve months by academics and vendors





#### Problem: Secure Ops on Insec.Systems (ctd)

PC-as-a-router for secure tokens doesn't work too well as a general solution...

- Expensive
- Requires deployment of specialised hardware
- Requires custom protocols and mechanisms both on the client and the server in order to handle the constraints imposed by the attached device
- A royal pain to use

Usefulness/inconvenience ratio is just too big



Availability concerns dictate that in the case of a problem the system allows things to continue

• Security concerns dictate that in the case of a problem the system doesn't allow things to continue

This is an umbrella problem that encompasses several other unsolvable sub-problems (covered later) as subclasses

- Unattended key storage
- Upgrade a product or device after a security breach







Management's response to this was to have the safety interlocks on the generators disabled

- Might get an extra five or ten minutes out of them
- Could potentially ride out a power outage that they wouldn't otherwise have survived



#### Problem: Security vs. Availability (ctd)

Preferable to run the generators to destruction than to risk having the data centre go down

• You won't find this in the MCSE or CCNA training material



High-availability systems (e.g. SCADA) cannot go down

• Ever

MTBF requirements of ten years are not uncommon

• May be run for decades Often can't be patched





#### Problem: Security vs. Availability (ctd)

Problem: CA-issued certificates are valid for one year

- MTBF 12 months << MTBF 10 years
- Ignore certificate expiry
  - In any case it's just a CA billing mechanism
  - Certificate that's perfectly fine on day n doesn't become completely insecure on day n + 1
- Issue your own certificates with infinite lifetimes

Problem: Certificates may suddenly stop working due to revocation

• Ignore CRLs and OCSP

Problem: Devices don't have DNS names, or even fixed IP addresses

- Identity = address often doesn't hold in any case
- Device can be identified by IEEE EUI64 ID (OUI + unique ID)
- ID the device by EUI64 or similar after you connect

The more checking you do, the greater the chance of something breaking

• Disable the safety interlocks nuisance checks to make sure things keep running



Random number generation

- /dev/random blocks until enough entropy is available
- /dev/urandom doesn't

#### Tinfoil-hat response

• Only ever use /dev/random

Linux kernel provides a system call getrandom()

- In the kernel for years but wasn't supported in glibc
- Google "ulrich drepper"
- Some support finally added in late 2016

## 

Sometimes you can reach a compromise...

Microsoft did this with the Windows XP SP2 firewall settings

• Finally, *finally* turned on by default in XP SP2

Found that home networks in which a computer acts as a file/print server were broken by having ports closed by default

• Open the ports required for print and file sharing... ... but only for the local subnet

#### Problem: Security vs. Availability (ctd)

Home users are unlikely to be running computers on multiple subnets

• Anyone sophisticated enough do this will presumably know what a firewall is and what to do with it

Protects home users from Internet-based attacks without breaking their existing network setup





#### Security vs. Availability at Design Level (ctd)

You can write an implementation that ignores MAC failures, decryption errors, and invalid signatures, and be fully standards compliant

• As long as you deal with a small number of obscure corner cases specifically called out as MUST NOTs

Look at the spec for your favourite protocol after the talk

- Someone else's problem?
- It was never even considered?
- The experienced driver will usually know what's wrong?

# Security vs. Availability at Design Level (ctd) Any security RFC ever • "Here is a security protocol. Whatever it happens to defend against is the threat model" The Inside-Out Threat Model Here is some crypto. If it happens to do what you want, go ahead and use it — Matt Blaze?, Protocols Workshop



• In other words, within the bounds that are defined by the defender







#### Security vs. Availability at Design Level (ctd)

TLS designers: No, of course not. Everyone knows that

- Except 99.99% of all web users everywhere
- "If you can see the padlock/green bar, you're safe"
- TLS has no threat model
  - Nor does SSH, S/MIME, or PGP
  - DNSSEC model was reverse-engineering from the spec a decade after it was published
  - IPsec was similar

# Security vs. Availability at Design Level (ctd)

IEEE standards require a rationale

• Explanation for why the standard does something

No major security RFC (TLS, SSH, PGP, S/MIME, IPsec, etc) contains one

- There are things in RFCs that cannot be rationally explained
- Seriously! When the authors were asked, they had no idea why their protocol design required XYZ

Rationale serves two purposes

- Guidance to implementers on how to apply a feature
- Sanity check on why the spec requires an action









Capability will only be exercised in extremely rare circumstances

- Complex, error-prone code that's never really exercised
- Has to sit there unused (but resisting all attacks) for years until it's needed
- Has to work perfectly the first time

How do you safely load a replacement algorithm into a remote device when the existing algorithm that's required to secure the load has been broken?



Example: TLS 1.2

• Deployment lagged for years because the change from TLS 1.1 to 1.2 was far bigger than from SSL to TLS

Scan carried out by a browser vendor in mid-2010 found exactly two public web servers supporting TLS 1.2

• Both were specially set-up test servers

Even in 2016, the most widely-encountered TLS version was 1.0 from 1999

• In SCADA/embedded, TLS 1.0 will probably be around forever



- If the TLS 1.2 experience (15+ year lag to general deployment) is anything to go by, we could see general adoption of 1.3 (outside of Google, Facebook, Akamai, etc) by 2030 or 2035
  - For SCADA/embedded, that date could be "never"
    - HTTP/2 was explicitly forked, HTTP/2 for large Silicon Valley Internet companies, HTTP 1.1 for everyone else
  - They're still planning the move to 1.2 within the next 5-10 years



Remote boxes administered from a central server

- Boxes communicate their state to the central server
- Central site loads it into a new device that gets sent out
- Used by some VoIP boxes rented from a provider

Watch out for supply-chain attacks!



#### Problem: Upgrading Insecure Crypto (ctd)

Opportunistic upgrade of algorithms

- If the other side presents a certificate with algorithm n + 1 then switch all communication with the certificate owner to n + 1 as well
- Lots of fun for security geeks to play with

May be subject to rollback attacks



#### Problem: Key Storage for Unattended Use

Another variant of security vs. availability

Storing keys in plaintext form is a cardinal sin in cryptography

• A user is expected to enter a password or PIN to unlock or decrypt keys so that they can be used

How do you do this for devices that have to be able to operate unattended?

How do you recover from a crash/power outage/OS upgrade/VM migration without explicit human intervention?



#### Problem: Key Storage for Unattended Use

None of the "obvious" general-purpose solutions to this problem actually solve it

TPMs can only store the fixed storage-protection key that's required to decrypt the real key

- TPMs are just repurposed smart cards and don't have the horsepower to perform anything more than lightweight crypto themselves
- Can't offload the overall encryption processing to them

For unattended operation they have to release their secrets without a PIN

• Merely provide plaintext key storage with one level of indirection



#### Problem: Key Storage for Unattended Use

Not very secure against an attack that compromises the host system

• All the HSM does is move the key from the compromised machine into an external box that does anything that the compromised host tells it to

Useful however for meeting auditing or regulatory requirements

- Adds an auditable physical artefact to the process
- "The box is still there, so we'll assume that the key is also still there"



If you're really concerned about security, move more of the security functionality into the HSM

- Instead of acting as a yes-box for crypto ops, implement whole portions of the underlying security protocol in the HSM
- Takes a large amount of programming effort





#### Hard and Not-so-Hard Problems (ctd)

Move the problem to an easier space and then solve that

Accept the fact that there are no perfect solutions

- Well, OK, there are perfect solutions
- Smart cards, PKI, biometrics, quantum anything, ...
- Monorails, cold fusion, power too cheap to meter, ...

Le mieux est l'ennemi du bien

- Voltaire (and others)

The good you can have right now, the perfect you'll have to wait forever for



Hard and Not-so-Hard Problems (ctd)		
And that's not just effective on lesser-known blogs		
CodingHorror is already a pretty popular blog, but it works just as well for major targets like this one		
Leave a comment Name (required):		
E-mail Address:		
URL:		
Fill in the blank: the name of this blog is Schneier on (required):		
Comments:		
• "What was the colour of the Lone Ranger's white horse?"		



#### 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 to obtain accounts 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
- Sometimes known as double opt-in

Widely used for password resets, mailing list subscriptions, blog registration

• 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

Enhanced version: Get users to set up a separate emailauth-only account

- Not used for anything else
- Not publicly visible
- Little chance of being phished

#### Email-based Identification (ctd)

Low-value authentication, but relatively difficult to defeat compared to what it's protecting

• An attacker who goes to the trouble of compromising your email account probably isn't interested in using it for mailing list access or blog spam





Comment/Link Spam (ctd) OK how else to deal with this			
	Authentication Code:		
	Or, use your backup method: Email → $\leftarrow$ Back to Bang! Boom! Pow!	Source: Wordpress	
• Go full crypto on them			











#### **Opportunistic Encryption**

After twenty years of 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"

#### STARTTLS/STLS/AUTH TLS

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

Most commonly encountered in SMTP/POP/IMAP

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

#### STARTTLS/STLS/AUTH TLS (ctd)

A year after first 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 has mostly displaced straight SMTP
- The fact that it helps authenticate/authorise users no doubt helped

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
- See a near-infinite number of papers, blogs, and articles on the failure of web PKI to prevent any real attacks



#### Assurance through Continuity

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

- McDonalds primary product line 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

Image removed following copyright infringement claim from the Coca Cola Company

#### Assurance through Continuity (ctd)

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, the 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, the client software verifies that the current server key matches the initial one

• Warn the user if it changes

#### Key Continuity in SSH (ctd)

Do SSH Fingerprints Increase Security?

Peter Gutmann Department of Computer Science University of Auckland

Abstract

No.

OK, so the fingerprint part doesn't work so well, but the continuity does











#### Key Continuity in SSL (ctd)

Finally got it right at the second attempt

RFC 7469: Public Key Pinning Extension for HTTP

• Only accept one of the following set of certificates for the next time period *x* 

Well, they tried...

- Tied to HTTP, so doesn't work for any other SSL use
- Google Chrome is the only major browser to support it

   Guess who wrote the spec?

#### 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







Used by PyPI (Python Package Index) to authenticate packages

Link to package is posted as http://pypi.python.org/packages/foo.tar.gz#sha1=23cb[...]e5fc

- Link contains hash needed to check the package
- Packed into the HTTP fragment identifier

Python install tools automatically verify its integrity on download

#### Self-Authenticating URLs (ctd)

Deals with the global PPI (Per-Per-Install) industry

• Repackage existing distros to include malware

Transforms the problem of

- Signing every package
- Managing a PKI
- Providing client-side software capable of interpreting the codesigning data

#### То

• Providing a secure location to post URLs

That's moving a very large goalpost!

#### Self-Authenticating URLs (ctd)

BitTorrent uses something a bit like this

- Actually just a fragment identifier to identify a piece of a large file
- Has the convenient side-effect that the torrent metadata also provides something like a self-authenticating URL

Other P2P protocols similarly use hashes to uniquely identify content online

More generally, DHTs use them to create self-certifying named objects

- Get me the object with this hash
- Does this object correspond to the hash I've got for it



A general form was proposed as link fingerprints

Attempts to standardise it were torpedoed due to concerns that it sapped and impurified the precious bodily fluids of URLs

- Added to Firefox, but removed again due to concerns that people might actually use it
- Seriously!

Supported in various plugins and download managers



• Redefine the problem to make it solvable