

Unsolvable Problems in Computer Security*

Peter Gutmann
University of Auckland

* With a bit of a bias towards cryptography

Unsolvable Problems

You mean *unsolved*, right?

- Nope
- Some problems really 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

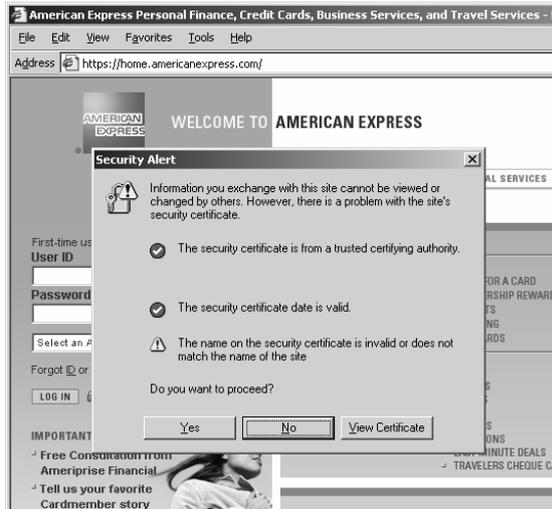
Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



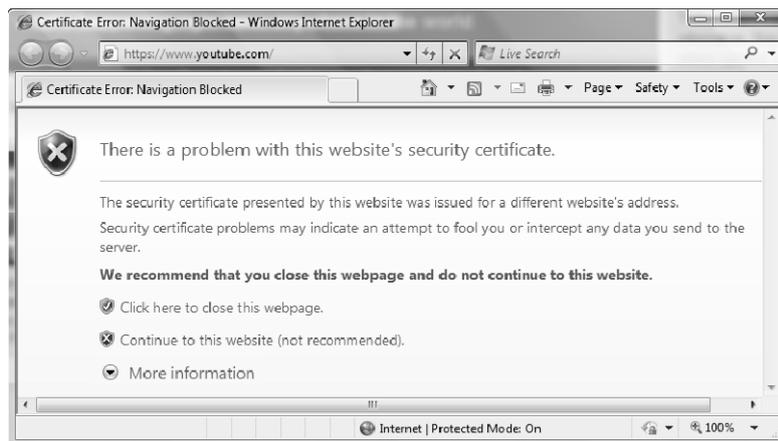
Example: Secure Bootstrapping of Comms

What about PKI?



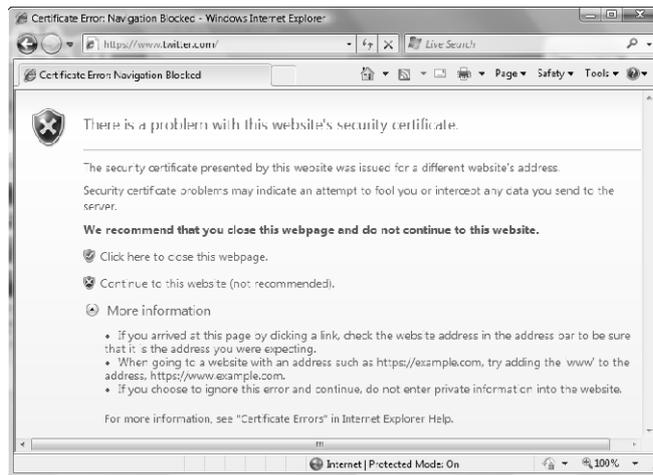
Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



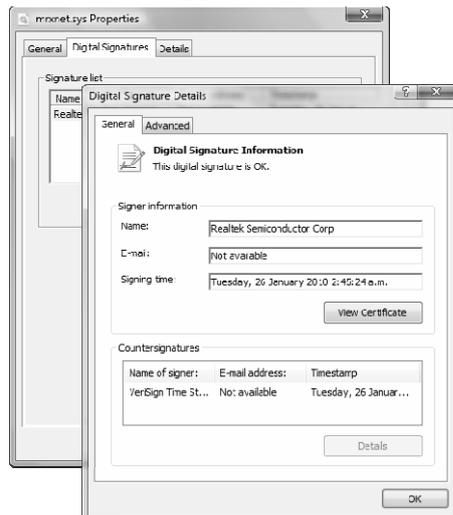
Example: Secure Bootstrapping of Comms

What about PKI?



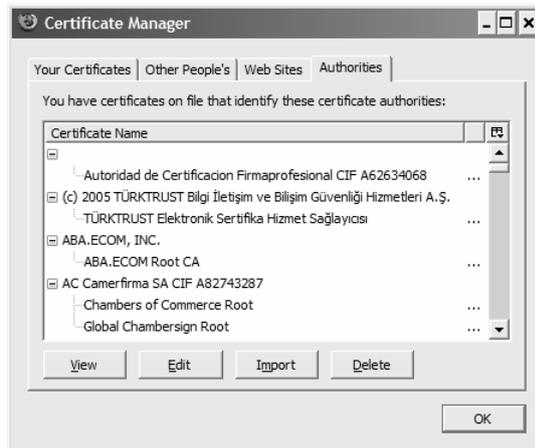
Example: Secure Bootstrapping of Comms

What about PKI?



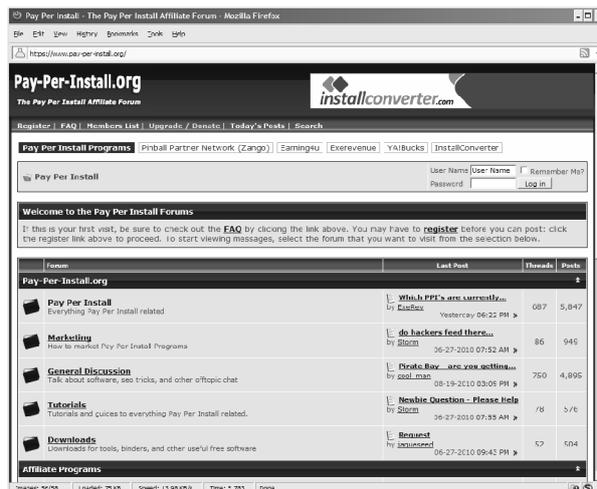
Example: Secure Bootstrapping of Comms

What about PKI?



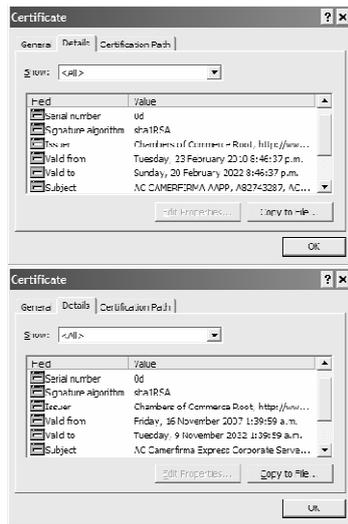
Example: Secure Bootstrapping of Comms

What about PKI?



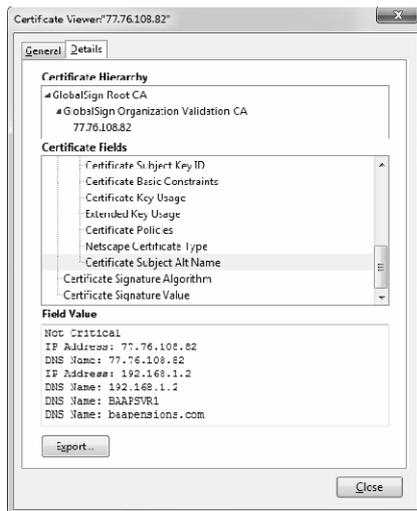
Example: Secure Bootstrapping of Comms

What about PKI?



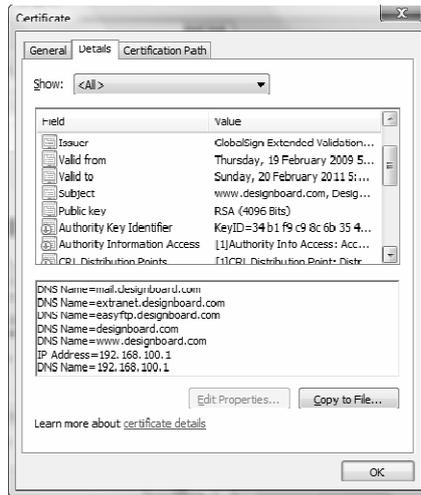
Example: Secure Bootstrapping of Comms

What about PKI?



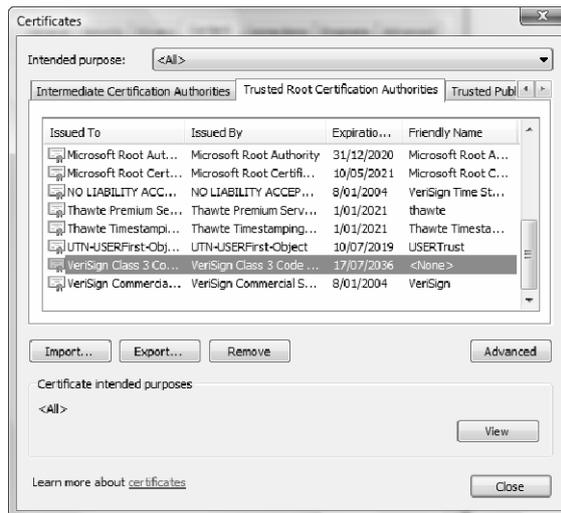
Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



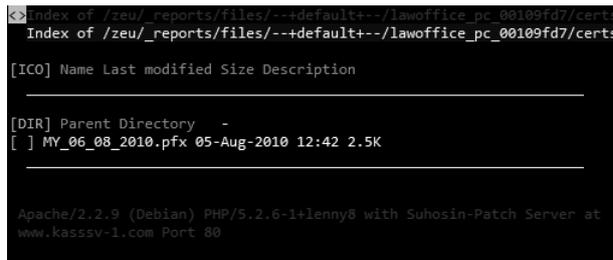
Example: Secure Bootstrapping of Comms

What about PKI?



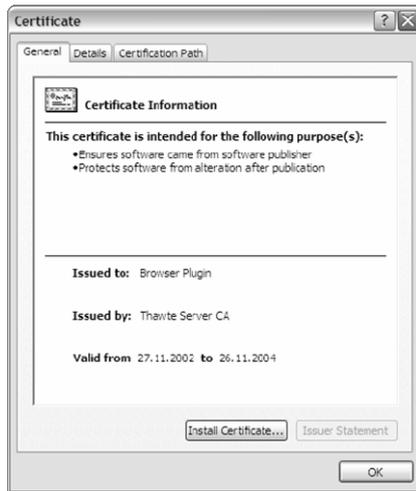
Example: Secure Bootstrapping of Comms

What about PKI?



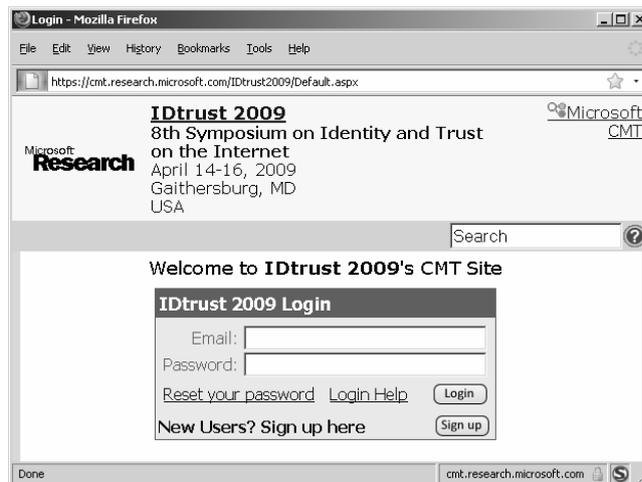
Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?



Example: Secure Bootstrapping of Comms

What about PKI?

What about it?

Example: Secure Bootstrapping of Comms

What about SSH?

Do SSH Fingerprints Increase Security?

Peter Gutmann

Department of Computer Science

University of Auckland

Abstract

No.

Example: Secure Bootstrapping of Comms

What about *insert-pet-mechanism-here?*

- No, that won't solve it either

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 ratable only as "better" or "worse" and not true or false
- No clear idea of a which steps or operations are necessary to get to the desired goal
- A variety of ideological and political differences among stakeholders

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)

Reduce weight by fitting a lighter engine



- Have to make the car lighter to compensate for the less powerful engine

If taken to extremes leads to a car that's little more than an exoskeleton with a motorcycle engine

- Has limited appeal to the general market

Example: High-performance Sports Cars (ctd)

Use exotic materials like carbon fibre to decrease weight



- Raises the price and again discourages buyers

Example: High-performance Sports Cars (ctd)

Strip out as many weight-adding features as possible



- Trade-off between performance and comfort
- Some jurisdictions have safety regulations that affect what you can and can't do
- Tradeoff between being able to sell the car in a particular market and making performance-reducing changes

Example: High-performance Sports Cars (ctd)

High-end audio is like
high-performance
sports car design,
only much sillier



Example: High-performance Sports Cars (ctd)

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...

Example: High-performance Sports Cars (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
- 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", 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 Insecure Systems

What about trusted computing?

- Yeah, any year now...

Even if it's deployed, protects only a small part of the system, e.g. the OS core

- A fully-protected computer on which you can't make any changes isn't terribly useful

Even for the portions that it does protect, all it guarantees is that they're unchanged from the state they were in when the TPM initially examined them

- Just because it's TPM-verified doesn't mean that it's safe

Problem: Secure Ops on Insecure Systems

A typical industry figure for code defects is about twenty bugs in every thousand lines of code (KLOC)

- Feel free to substitute your own pet value at this point
- The important thing isn't the absolute value but the rough order of magnitude for estimation purposes

Widely-used operating systems like Linux and Windows weigh in at 50-100 million lines of code

- Again, depending on which version and what you count as being part of "Linux" and "Windows"

Problem: Secure Ops on Insecure Systems

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

- Ignores the additional code that'll be added in the form of user-installed 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

Problem: Secure Ops on Insecure Systems

Various other cat-and-mouse games are possible

- Graphical data-entry mechanisms
 - The malware takes a screenshot and/or records mouse actions
- Trying to hook the system at a lower level than the rootkit
 - Malware engages the application in a race to the bottom of the driver chain
 - The malware invariably wins
- And so on

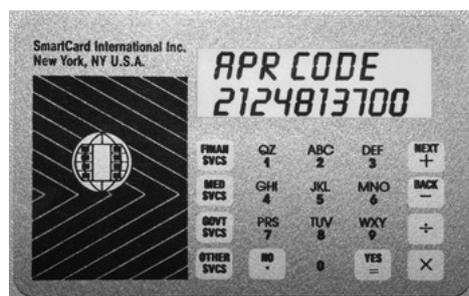
Problem: Secure Ops on Insecure Systems

Accept the fact that you can never really trust anything 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

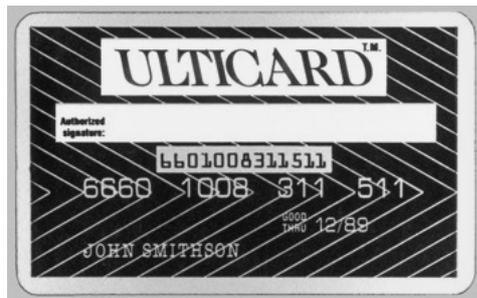
Problem: Secure Ops on Insecure Systems

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



Problem: Secure Ops on Insecure Systems

The process has been ongoing for at least twenty years now



- It's a remarkably persistent meme

Problem: Secure Ops on Insecure Systems

Doesn't work too well as a general solution

- Expensive
- Requires deployment of specialised hardware
 - Unless you use a cellphone as your external device
 - Need to create an application that installs and runs on a range of completely incompatible platforms that their owners tend to swap out every year or two
- Requires custom protocols and mechanisms both on the client and the server in order to handle the constraints imposed by the attached device

Problem: Security vs. Availability

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

Problem: Security vs. Availability (ctd)

Availability concerns can be a powerful motivator

Data centre was built with marine diesel generators for backup power

- Marine diesels come with a built-in cooling system
- That would be "the ocean"

Less concern about overheating than conventional generators

- Makes them more compact than standard generators
- Space was a concern for the data centre in question

Problem: Security vs. Availability (ctd)

The data centre wasn't anywhere near the ocean

- Used a stand-in consisting of a large water cistern whose contents were flushed through the generators' cooling systems
- When the cistern had emptied, the generators' thermal cut-outs shut them down

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

Problem: Security vs. Availability (ctd)

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 to 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

Problem: DRM and Copy Protection

Ugh, where to begin...

DRM is an attempt to fix a social and business-model problem using technology

- Functions even more poorly than most other cases where this is being attempted

No-one involved with the technology or its application has any interest in seeing it work

- Content vendors get the benefits
- Manufacturers and consumers carry the cost
- Variation of “moral hazard” from economics

Problem: DRM and Copy Protection (ctd)

Consider a manufacturer building some sort of DRM system into their product

- Best-case possible outcome, if the technology works one hundred percent perfectly, is that the user doesn't notice anything
- Net gain to the manufacturer from the expense and overhead of adding DRM is zero

Since no technology is ever perfect, a number of users will run into problems caused solely by the DRM

- Customer dissatisfaction leading to product returns and lost sales is a net loss for the manufacturer

Problem: DRM and Copy Protection (ctd)

Any manufacturer who doesn't bother with the DRM won't encounter these problems

- Creates a strong incentive for manufacturers to pay the barest lip service to DRM
- If they don't cheat then someone else will

The more functional (less crippled) products from the other manufacturer will prevail in the marketplace

Problem: DRM and Copy Protection (ctd)

Example: HDCP strippers

High-bandwidth digital content protection is a protection scheme for HDMI/DVI/etc

- Glitches or implementation bugs can prevent legitimate content from being sent to legitimate devices
- High overhead of HDCP handshake can introduce long delays when a signal is switched from one port to another

To most users this indicates a faulty device

Problem: DRM and Copy Protection (ctd)

A few boutique vendors used to sell expensive HDCP strippers



- Quickly vanished under an avalanche of legal threats

Problem: DRM and Copy Protection (ctd)

You can still get HDCP strippers, they're just not advertised or sold as such

- Lesser-known manufacturers of HDMI switches and repeaters made the economically rational decision to omit HDCP in their outputs
- Accept incoming HDCP signals and pass on HDMI-only signals
- Cheapest devices of this kind retail for around US \$15
- Solve problems like sending Blu-ray output to a non-HDCP device

Problem: DRM and Copy Protection (ctd)

Revocation mechanisms don't work in practice

- Try tracking down a small and ever-changing army of fly-by-night manufacturers of no-name HDMI switches

HDCP keying material was quite probably purloined from legitimate products by third-shift manufacturing practices

- Revoking the keys could potentially brick large numbers of legitimate products

Problem: DRM and Copy Protection (ctd)

DRM's killer problem: No-one involved in manufacturing the devices that use it or using the devices that contain it has any interest in it actually working except for the minimum that's enforced by legal threats

- (Also, "trying to make digital files uncopyable is like trying to make water not wet" — Bruce Schneier)

Problem: DRM and Copy Protection (ctd)

No end of companies will be all too happy to sell you as much DRM as you can afford

- There's an awful lot of money to be made in this area
 - None of these vendors will provide liability cover for you if their DRM fails

Do the minimum necessary to comply with what the lawyers want and make sure that you can't be sued when it breaks

- DRM compliance engineering
- Eventually it will break but if you've demonstrated sufficient diligence in your pursuit of compliance then you won't be held liable

Problem: DRM and Copy Protection (ctd)

If your business model is crucially dependent on functioning DRM then consider getting into another business

If your business model merely requires the appearance of functioning DRM then that's fine

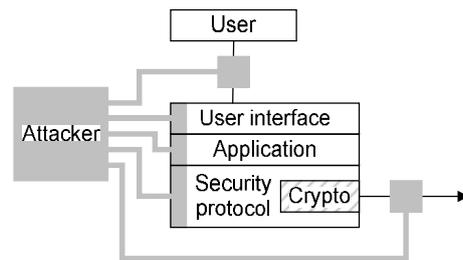
- You're getting that anyway no matter what DRM measures you employ

Problem: Upgrading Insecure Crypto

How do you recover from the catastrophic compromise of a security system?

Extremely rare in properly-designed systems

- Attackers target the implementation, the way it's used, or some other aspect unrelated to the crypto
- “Crypto is bypassed, not attacked”



Easiest approach is to ignore it and hope that it never occurs

Problem: Upgrading Insecure Crypto (ctd)

Consider a system that uses two authentication algorithms in case one fails

- Device receives a message authenticated with algorithm A saying “Algorithm B has been broken, don't use it any more”
- Device also receives a message authenticated with algorithm B saying “Algorithm A has been broken, don't use it any more”
- Device may also receive a third message saying “All Cretans are liars”

Could address this with fault-tolerant design concepts

- Voting protocols for algorithm replacement

Adds a huge amount of design complexity and new attack surface

Problem: Upgrading Insecure Crypto (ctd)

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

Problem: Upgrading Insecure Crypto (ctd)

Security geeks to want to replace half the security infrastructure that you're relying on as a side-effect of any algorithm upgrade

- c.f. TLS 1.2
- Deployment has 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

Problem: Upgrading Insecure Crypto (ctd)

Situation-specific solutions are possible...

Small number of high-cost units

- Courier out replacement devices that clone their state from the existing one
- Used by some hardware security modules (HSMs)

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

Need to secure records over long periods of time

- Use standard document-retention mechanisms
- Crypto is little more than a complex gimmick for this
 - “We’ll solve this problem with cryptography”
 - Now you have two problems

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

Various cat-and-mouse games possible

- Poke hierarchies of keys into various locations
- Use them to decrypt other keys
- Hope that an attacker can't work their way back far enough to grab the real keys

For unattended operation at some point you need to fall back to a fixed key stored in plaintext-equivalent form that can survive a crash or reboot

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

Add custom encryption hardware and perform all of the crypto in that

- Most manufacturers will be reluctant to add \$500 of specialised encryption hardware to a \$50 embedded device
- Scaled up to PC terms, a \$20,000 hardware security module (HSM) to a \$2,000 server
 - If the HSM vendor has particularly good salespeople they’ll sell the client at least two \$20,000 HSMs (each storing a single key) for disaster recovery purposes.

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

May be adequate to meet auditing or regulatory requirements though

- Adds an auditable physical artefact to the process

Problem: Key Storage for Unattended Use

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

IBM used to sell a fully programmable high-security crypto coprocessor

- Almost no-one took advantage of its programming capabilities

A good idea in theory but practical experience has shown that few users will make the effort

Problem Structuring Methods

There is one problem-solving approach that may be usable here

- Soft operations research, or problem-structuring methods

Operations research dates back to just before WWII

- UK scientists got together to look at ways of optimising the management of military materiel and manpower
- Known in the UK as operations research and in the US as management science

In the 1970s and 1980s practitioners of traditional OR realised that it didn't work for wicked problems

- Came up with soft OR, or problem-structuring methods (PSMs)

Problem Structuring Methods (ctd)

Methods include...

- Interactive Planning
 - Notable principally because its first step is “Mess Formulation”
- Soft Systems Methodology (SSM)
- Strategic Assumption Surfacing and Testing (SAST)
- Strategic Choice Approach (SCA)
- Strategic Options Development and Analysis (SODA)
- Many, many more

To be tested at a future Auckland ISIG meeting :-)

Unsolvable Problems Redux

So how do we fix this?

Uhh, go back to the slide that introduces wicked problems

If do you run into any of the problems discussed here then you're not going to find any (general) solution to them

- Design your systems to deal with this
- Try something else that avoids the need to solve a wicked problem