



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Security II

2013-2014

Dr Frank Stajano and Dr Markus Kuhn
+ guest lecturers

Computer Science Tripos Part II
University of Cambridge

Based on slides by Professor Ross Anderson, with changes

Introduction

Aims

- Give you a thorough understanding of security engineering as a systems discipline
 - Policy (what should be protected)
 - Mechanisms (cryptography, hardware security...)
 - Attacks (malicious code, exploiting users...)
 - Assurance (assessing how secure it is)

Objectives

By the end of the course, you should be able to tackle an information protection problem by

- drawing up a threat model
- formulating a security policy and
- designing specific protection mechanisms to implement the policy.

Broad range of topics


Not an “axiom, theorem, exercise” subject

You must learn to think outside the box

- Human factors, Security Policy, Crypto, Protocols, Incentives etc
- Guest lectures, to broaden horizons further:
 - Dr Richard Clayton, security economics
 - Dr Steven Murdoch, anonymity
 - Dr Sergei Skorobogatov, hardware security
 - Dr Robert Watson, concurrency and security

Resources

- Anderson, *Security Engineering*
 - developed from lecture notes
 - free online from author's page
- Katz, Lindell, *Introduction to modern cryptography*
- Web page for course
- Further reading:
 - Gollmann, *Computer Security*
 - Cialdini, *Influence: science and practice*
 - Stajano, *Security for ubiquitous computing*
 - Kahneman, *Thinking fast and slow*



Further resources

- History:
 - David Kahn: *The Codebreakers*
 - Gordon Welchmann: *The Hut Six Story*
- Use the Source, Luke:
 - Read the original papers (books come after papers)
- Lab:
 - Security seminars, usually Tuesdays, 1500-1600 (new time)
 - Security group meetings, Fridays, 1600-1700

www.cl.cam.ac.uk/teaching/current/SecurityII/materials.html

What is Security Engineering?

Security engineering is about building systems to remain dependable in the face of malice, error and mischance.

As a discipline, it focuses on the tools, processes and methods needed to design, implement and test complete systems, and to adapt existing systems as their environment evolves.

Systems!

- *A system* can be:
 - a product or component (PC, smartcard,...)
 - some products plus O/S, comms and infrastructure
 - the above plus applications
 - the above plus internal staff
 - the above plus customers / external users
- Common failing: policy drawn too narrowly
 - Want a secure system? You need to consider *users*

Security and psychology

Why Johnny can't encrypt

Whitten, Tygar: 'Why Johnny Can't Encrypt', 1999.

- Study of encryption program PGP – showed that 90% of users couldn't get it right give 90 minutes
- Private / public, encryption / signing keys, plus trust labels was too much – people would delete private keys, or publish them, or whatever
- Security is hard – unmotivated users, abstract security policies, lack of feedback...
- Geeky “security experts” would rather deal with machines than with unpredictable people. They miss the point.

Users are not the enemy

Adams, Sasse: “Users are not the enemy”, 1999.

- Insufficient communication with users produces unusable systems
- Users forced to comply with password mechanisms incompatible with work practices will look for workarounds
- Vicious circle:
 - Security departments think users are inherently insecure
 - Users think security departments get in the way of real work
- But “users never motivated to behave securely” is wrong!
- Treat users as stakeholders and they'll cooperate
- Provide feedback, guidance, awareness; and usable security

Think like an attacker

- Mitnick: *The art of deception*, 2003.
(Also *The art of intrusion*, 2005.)
 - You don't have to pick the lock or break into the server:
get someone on the inside to open the door for you
 - Pretext calls surprisingly effective
- Traditional responses:
 - mandatory access control
 - operational security

But why do the attacks work?

Think like a victim

Stajano, Wilson: “Understanding scam victims”, 2009-2011.

- Learn from fraudsters: they know how to push the victims' buttons
- The Real Hustle (BBC3): hundreds of scams recreated for hidden cameras

What makes these scams work?

Understanding scam victims

- **Distraction** – When focused, you forget the rest
- **Herd** - If others do it, it's probably ok
- **Need and greed** - There are things you just can't say no to
- **Dishonesty** - It's illegal, that's why it's such a good deal
- **Kindness** - People are actually nice
- **Authority** - If the man in uniform says so, you do it
- **Time and scarcity** – If there isn't enough of it, it must be good

Exact principles not actually that important...

But their existence is!

Seven principles for systems security

- Principles for sales, scams and other “persuasion” contexts
- Those principles predate computers! Rooted in human nature
- Thus, each also applies to computer *systems* security
 - Phishing: user is fixated on task completion (e.g. finding why new payee on PayPal account)
 - Advance fee frauds (419) take this to extreme lengths!
- You have to accept them, almost like the laws of physics
 - That's the way people work, no matter what you tell them to do
 - “If only it weren't for those gullible users” is arrogant and idiotic

Wise system engineer understands and acknowledges the principles and makes the system robust by preventing their exploitation

Social Psychology

- Solomon Asch, 1951: two-thirds of subjects would deny obvious facts to conform to group
- Stanley Milgram, 1964: a similar number will administer torture if instructed by an authority figure
- Philip Zimbardo, 1971: you don't need authority: the subjects' situation / context is enough

Cfr Herd principle and Authority principle

- The Officer Scott case: a “police officer” phones a fast food restaurant and persuades the manager to strip-search and sexually humiliate an employee
- What about users you can't train (your customers)?

Weapons of influence

Cialdini: *Influence: science and practice*, 5th ed 2009.

Based on undercover study of salesmen “and other compliance professionals”

Salesmen are like scam artists, except legal (perhaps)

- Reciprocation: they'll feel compelled to respond
- Commitment and consistency: “but you previously said X”
- Social proof: like to do what others do
- Liking: want to deal with people they can relate to
- Authority: will defer to authority figure
- Scarcity: less is best and loss is worst

Framing effects include ‘Was £8.99 now £6.99’ and the estate agent who shows you a crummy house first. Take along an ugly friend on a double date.

Phishing

- Started in 2003 with six reported (there had been isolated earlier attacks on AOL passwords)
- By 2006, UK banks lost £35m (£33m by one bank) and US banks maybe \$200m
- Early phish crude and greedy; but phishermen learned fast
- E.g. ‘Thank you for adding a new email address to your PayPal account’
- The banks make it easy for them

Types of phishing website

- Misleading domain name

`http://www.banckname.com/`

`http://www.bankname.xtrasecuresite.com/`

- Insecure end user

`http://www.example.com/~user/www.bankname.com/`

- Insecure machine

`http://www.example.com/bankname/login/`

`http://149.32.40.1/bankname/login/`

- Free web hosting

`http://www.bank.com.freespacesitename.com/`

Fraud and Phishing Patterns


- Fraudsters do pretty well everything that normal marketers do
- The IT industry has abandoned manuals – people learn by doing, and marketers train them in unsafe behaviour (click on links...)
- Banks' approach is 'blame and train' – long known to not work in safety critical systems
- Their instructions 'look for the lock', 'parse the URL' are easily turned round, and discriminate against nongeeks


Online Safety & Personality Survey :: FF1 - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://www.tylermoore.com/limesurvey/index.php

Getting Started Latest Headlines

 UNIVERSITY OF CAMBRIDGE

Online Safety & Personality Survey
0%  100%

FF1

Is the following website legitimate or phishing?

Login - PayPal - Mozilla Firefox

File Edit View History Bookmarks Tools Help


http://httpspaypalcom/cgi-bin-webscr-sec-cmdlogin.secdatal.com/cgi-bin-webscr

Getting Started Latest Headlines

[Sign Up](#) | [Log In](#) | [Help](#) | [Security Center](#) Search

PayPal U.S. English

Home Personal Business Products & Services

Account login 

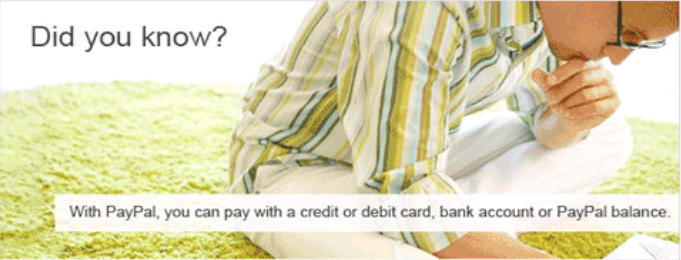
Email address

PayPal password

Forgot your [email address](#) or [password](#)?


New to PayPal? [Sign up](#)

Did you know?



With PayPal, you can pay with a credit or debit card, bank account or PayPal balance.

[About](#) | [Accounts](#) | [Fees](#) | [Privacy](#) | [Security Center](#) | [Contact Us](#) | [Legal Agreements](#) | [Developers](#) | [Jobs](#) | [Merchant Services](#) | [Mobile](#) | [Plus Card](#) | [Referrals](#) | [Shops](#) | [Mass Pay](#)


Copyright © 1999-2008 PayPal. All rights reserved.
[Information about FDIC pass-through insurance](#)

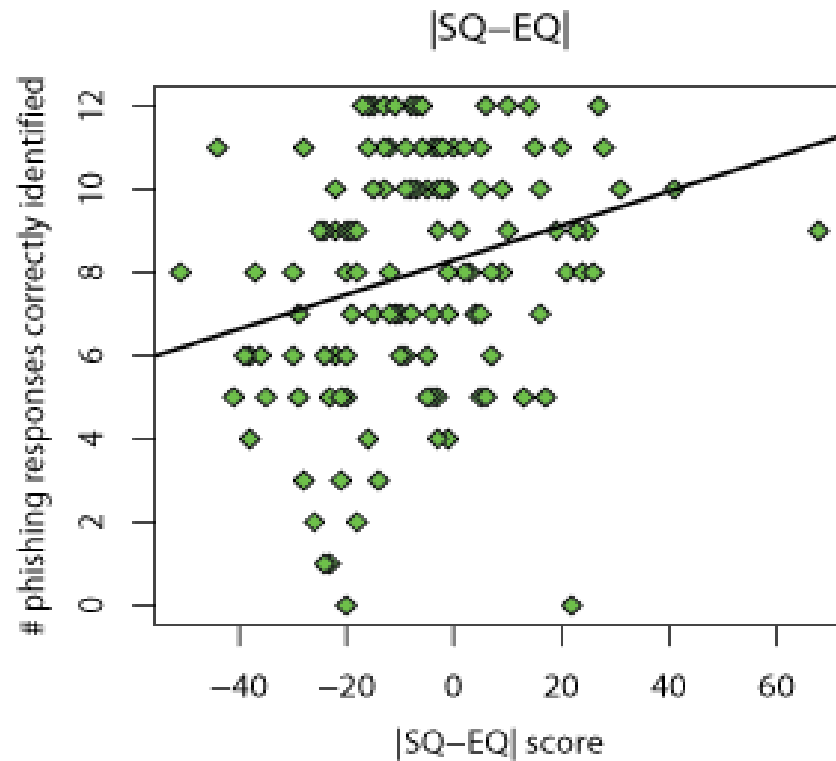
Done

Choose one of the following answers

Legitimate website
 Illegitimate phishing website
 I'm not sure

Done

Results



- Ability to detect phishing is (somewhat loosely) correlated with SQ-EQ (S=systemizer, E=empathizer)
- It is (independently) correlated with gender
- So the gender HCI issue applies to security too

Framing effects

- Imminent outbreak of some Asian Disease is expected to kill 600
- Two programs to combat the disease have been proposed
 - Program A:
 - 200 people saved
 - Program B:
 - 1/3 chance that 600 saved
 - 2/3 chance that 0 saved

Framing effects

- Imminent outbreak of some Asian Disease is expected to kill 600
- Two programs to combat the disease have been proposed
 - Program A':
 - 400 people die
 - Program B':
 - 1/3 chance that 0 die
 - 2/3 chance that 600 die

Framing effects

- Imminent outbreak of some Asian Disease is expected to kill 600
- Two programs to combat the disease have been proposed
 - Program A:
 - 200 people saved
 - Program B:
 - 1/3 chance that 600 saved
 - 2/3 chance that 0 saved
 - Program A':
 - 400 people die
 - Program B':
 - 1/3 chance that 0 die
 - 2/3 chance that 600 die

*Substantial majority favours A:
let's save those 200 (risk-averse)*

*Substantial majority favours B':
let's not kill those 400
(risk-seeking)*

Why the switch?

Outcomes that are certain are over-weighted

Attitudes towards risk

Which would you choose?

- Win £900 for sure
- 90% chance of winning £1000

*(For a rational agent, there should be no difference;
but what will you do if it's your own real money?)*

Attitudes towards risk

Which would you choose?

- Lose £100 for sure
- 10% chance of losing £1000

Again: in theory, no difference. But with your real money?

Attitudes towards risk

Which would you choose?

- Win £900 for sure
- 90% chance of winning £1000

Give me the certain money!
risk-averse

Attitudes towards risk

Which would you choose?

- Win £900 for sure
- 90% chance of winning £1000

Give me the certain money!
risk-averse

Bernoulli's Expected Utility theory explains why.

Utility of wealth is less than linear.

$£900 = 0.9 * £1000$; but $utility(£900) > 0.9 * utility(£1000)$

Attitudes towards risk

Which would you choose?

- Win £900 for sure
- 90% chance of winning £1000

Give me the certain money!
risk-averse

Bernoulli's Expected Utility theory explains why.

Utility of wealth is less than linear.

$£900 = 0.9 * £1000$; but $utility(£900) > 0.9 * utility(£1000)$

Which would you choose?

- Lose £100 for sure
- 10% chance of losing £1000


Let me try to get away without paying!
risk-seeking

If I give you £1000 after the loss, expected values are same as top problem.

Bernoulli's utility of wealth *can't explain* why you switch.

The difference is only *framing* (as a gain or a loss).

Prospect theory



- Kahneman & Tversky, 1979
 - Critique of Expected Utility theory
 - Descriptive analysis of decision under risk
 - Importance of *framing*
 - Cfr *perception* (hot or cold water)

- Utility of wealth? No, of changes from reference point (*gains and losses matter, not overall wealth*)
- Both curves lean towards the horizontal (*both higher gains and higher losses are progressively less relevant*)
- But steeper for losses than for gains (*we dislike a loss much more than we like a win; and by about 2x*)
- Concave for gains (*risk-averse*), convex for losses (*risk-seeking*)

Risk Misperception

- Why do we overreact to terrorism?
 - Risk aversion / status quo bias
 - ‘Availability heuristic’ – easily-recalled data used to frame assessments
 - Our behaviour evolved in small social groups, and we react against the out-group
 - Mortality salience greatly amplifies this
 - We are also sensitive to agency, hostile intentions
- Anderson chapters 2, 24
- But read the original Kahneman
(Nobel lecture; *Thinking fast and slow*)

Passwords

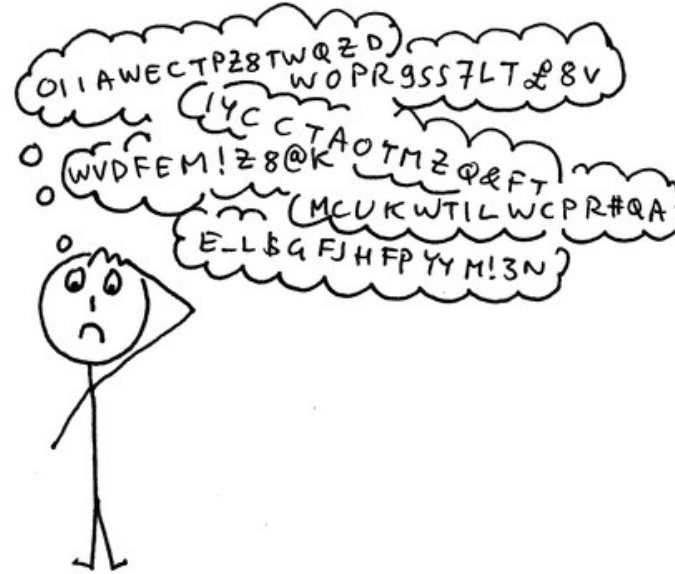
- Have many usability shortcomings
- Also have many security shortcomings
- But continue to be dominant

There is no doubt that over time, people are going to rely less and less on passwords. People use the same password on different systems, they write them down and they just don't meet the challenge for anything you really want to secure.

--Bill Gates, 2004.

What's wrong with passwords

- Must resist guessing
- Must resist brute-force
- Must be all different
- Must be regularly changed
- Must be memorable (cannot be written down)
- Must be typed correctly, despite *****



...but the intersection of these requirements is the empty set!

And the more accounts we have, the worse it gets



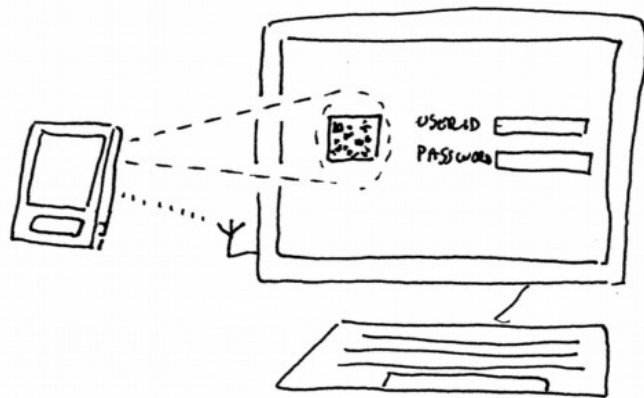
What else is wrong with passwords?

That was “just” usability. Then there's security.

- Thanks to the previous usability flaws...
 - weak passwords are guessed or brute-forced
 - (reused) passwords leaked at A cause a break-in at B
 - attacks on the recovery mechanism for forgotten passwords
 - passwords on post-its enable (casual or targeted) attacks
- Phishing
 - automated industrial-scale social engineering
- Keylogging
 - automated industrial-scale shoulder-surfing
 - or local, with hardware that can't be detected by security software
- Incompetent implementation at verifier end
 - Bonneau, Preibusch: “The password thicket”, 2010.

Pico: no more passwords!

Stajano, 2011.



+ BIOMETRIC



+ NETWORK SERVER

Want to join my Pico team for a master or PhD and make a dent in history?

The compliance budget

Beautement, Sasse, Wonham, 2008.

- Users assess cost/benefits of security measures and put up with some inconvenience for the good of the company
 - With virus scanner on, program takes longer to build
 - Encrypting USB stick might prevent me from giving talk
- But only up to a point!
 - User patience is finite: “the compliance budget”
 - Once exhausted, user quickly stops cooperating

Must be managed like any other budget

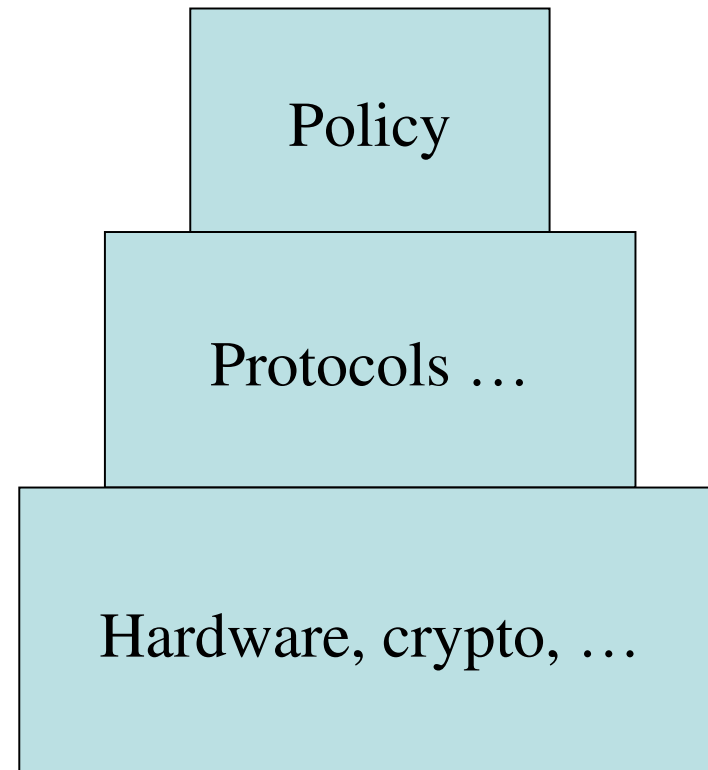
Security policies

Managing security

- Security awareness: measures must have, and be seen to have, full support of management
- Measuring security is hard
 - Measure security bugs, attack surface, attack cost...
- Risk analysis
 - Assets, vulnerabilities, threats, probabilities
 - That's quantitative, but inputs are usually guesswork
- Security policy: an instrument of communication

Design Hierarchy

- What are we trying to do?
- How?
- With what?



Security vs Dependability

- Dependability = reliability + security
- Reliability and security are often strongly correlated in practice
- But malice is different from error!
 - Reliability: “Bob will be able to read this file”
 - Security: “The Chinese Government won’t be able to read this file”
- Proving a negative can be much harder ...

Terminology

- A *subject* is a physical person
- A *person* can also be a legal person (firm)
- A *principal* can be
 - a person
 - equipment (PC, smartcard)
 - a role (the officer of the watch)
 - a complex role (Alice or Bob, Bob deputising for Alice)
- The level of precision is variable – sometimes you need to distinguish ‘Bob’s smartcard representing Bob who’s standing in for Alice’ from ‘Bob using Alice’s card in her absence’. Sometimes you don’t.

Terminology

- *Secrecy*: mechanisms limiting the number of principals who can access information
- *Privacy*: control of your own secrets
- *Confidentiality*: an obligation to protect someone else's secrets

Thus your medical privacy is protected by your doctors' obligation of confidentiality

Terminology

- *Anonymity* is about restricting access to metadata. It has various flavours, from not being able to identify subjects to not being able to link their actions
- An object's *integrity* lies in its not having been altered since the last authorized modification
- *Authenticity* has two common meanings –
 - an object has integrity plus freshness
 - you're speaking to the right principal

Terminology

Trust is the hard one! It has several meanings:

1. a warm fuzzy feeling
2. a trusted system or component is one that can break my security policy
3. a trusted system is one I can insure
4. a trusted system won't get me fired when it breaks

We'll use the NSA definition – number 2 above – by default.

E.g. an NSA man selling key material to the Chinese is trusted but not trustworthy (assuming his action unauthorised)

Terminology

- A *security policy* is a succinct statement of protection goals – typically less than a page of normal language
- A *protection profile* is a detailed statement of protection goals – typically dozens of pages of semi-formal language
- A *security target* is a detailed statement of protection goals applied to a particular system – and may be hundreds of pages of specification for both functionality and testing

What often passes as 'Policy'

1. This policy is approved by Management.
2. All staff shall obey this security policy.
3. Data shall be available only to those with a 'need-to-know'.
4. All breaches of this policy shall be reported at once to Security.

What's wrong with this?

Policy: Multi Level Security

- Multilevel Secure (MLS) systems are widely used in government / intelligence / military contexts
- Basic idea: a clerk with 'Secret' clearance can read documents at 'Confidential' and 'Secret' but not at 'Top Secret'
- 1960s/70s: problems with early mainframes
- First security policy to be worked out in detail following Anderson report (1973) for USAF which recommended keeping security policy and enforcement simple

Levels of Information

- Levels include:
 - **Top Secret:** compromise could cost many lives or do exceptionally grave damage to operations. E.g. intelligence sources and methods, battle plans
 - **Secret:** compromise could threaten life directly. E.g. weapon system performance, combat reports
 - **Confidential:** compromise could damage operations
 - **Restricted:** compromise might embarrass
else “unclassified”
- Resources have classifications
- Principals have clearances
- Information flows upwards only

Context of Multilevel Security

- Information mustn't leak from High to Low
- Enforcement must be independent of user actions
- Perpetual problem: careless staff
- 1970s worry: operating system insecurity
- 1990s worry: virus at Low copies itself to High and starts signalling down (e.g. covert channel)


Manning (2010) and Snowden (2013)
show us how things actually go wrong in practice...

Context of Multilevel Security

Nagaraja, Anderson 'The Snooping Dragon', 2009.

- September 2008: Dalai Lama's office realised there had been a security failure
- Initial break: targeted email with bad pdf
- Then: took over the mail server and spread it
- About 35 or their 50 PCs were infected
- Fix (Dharamsala): take 'Secret' stuff offline
- Fix (UKUSA agencies): use MLS mail guards and firewalls to prevent 'Secret' stuff getting out

Authorized Information Flow



Formalising the Policy

- Bell-LaPadula (1973):
 - *simple security property*: no read up
 - **-property*: no write down
- With these, one can prove that a system which starts in a secure state will remain in one
- Ideal: minimise the Trusted Computing Base (set of hardware, software and procedures that can break the security policy) so it's verifiable
- 1970s idea: use a reference monitor

The Lattice Model

- Intelligence agencies manage ‘compartmented’ data by adding categories. **Label = (level, {set of categories})**
- Basic idea: BLP requires only a partial order (*dominates*).

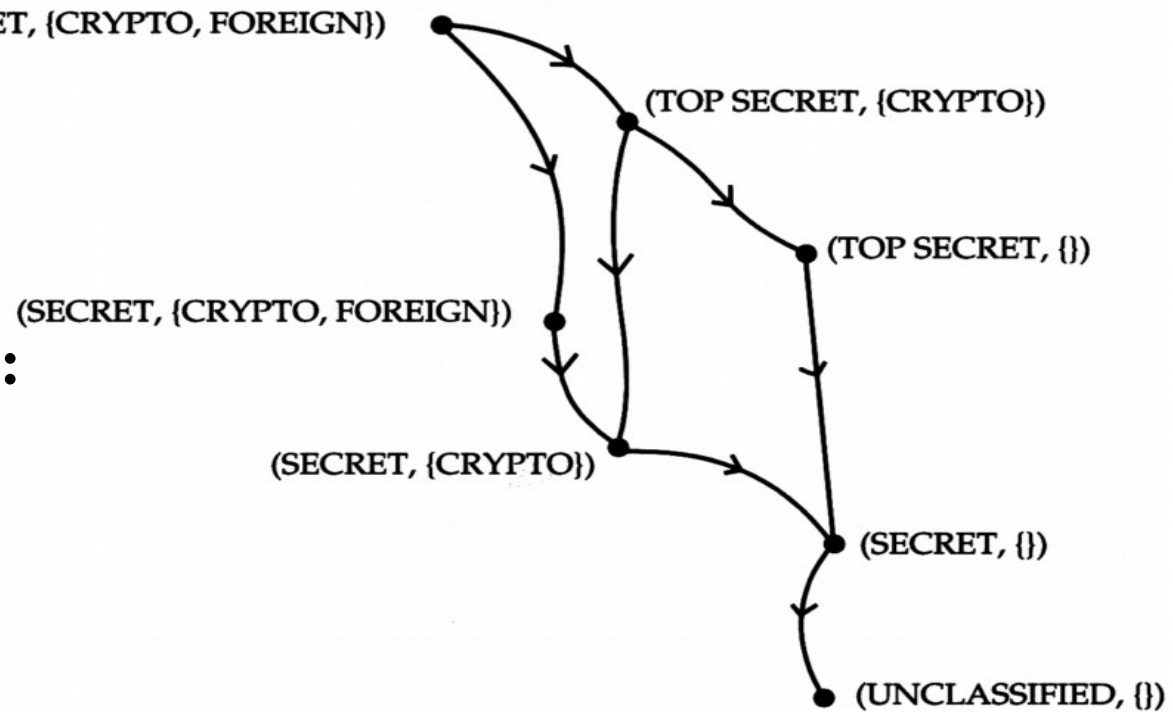
- X dominates Y iff
 $\text{level}(X) \geq \text{level}(Y)$ and
 $\text{cat}(X) \supseteq \text{cat}(Y)$

- **BLP simple property (NRU):**

X can read Y iff
 X dominates Y

- **BLP *property (NWD):**

X can write Y iff
 X is dominated by Y



Objections to BLP

- Some processes, such as memory management, need to read and write at all levels
- Fix: put them in the trusted computing base
- Consequence: once you put in all the stuff a real system needs (backup, recovery, comms...) the TCB is no longer small enough to be easily verifiable

Objections to BLP

- John MacLean's "System Z": as BLP but lets users request temporary declassification of any file
- Fix: add tranquility principles
 - Strong tranquility: labels never change
 - Weak tranquility: they don't change in such a way as to break the security policy
- Usual choice: weak tranquility using the "high watermark principle" – a process acquires the highest label of any resource it's touched
- Problem: have to rewrite apps (e.g. license server)

Objections to BLP

- High can't acknowledge receipt from Low
- This blind write-up is often inconvenient: information vanishes into a black hole
- Option 1: accept this and engineer for it (Morris theory) – CIA usenet feed
- Option 2: allow acks, but be aware that they might be used by High to signal to Low
- Use some combination of software trust and covert channel elimination


Covert Channels

- In 1973 Butler Lampson warned BLP might be impractical because of covert channels: “neither designed nor intended to carry information at all”
- A Trojan at High signals to a buddy at Low by modulating a shared system resource
 - Fills the disk (storage channel)
 - Loads the CPU (timing channel)
- Capacity depends on bandwidth and S/N. So: cut the bandwidth or increase the noise
- But it’s really hard to get below 1 bit/s or so...

Downgrading

- A related problem to the covert channel is how to downgrade information
- Analysts routinely produce Secret briefings based on Top Secret intelligence, by manual paraphrasing
- Also, some objects are downgraded as a matter of deliberate policy – an act by a trusted subject
- For example, a Top Secret satellite image is to be declassified and released to the press

Downgrading




...es, based on the observation
terms of the security perin
act. This translates generall

Category I – In this cate
The perpetrator must h
operate under th
standard ca

Text hidden in least significant bits of image

Downgrading



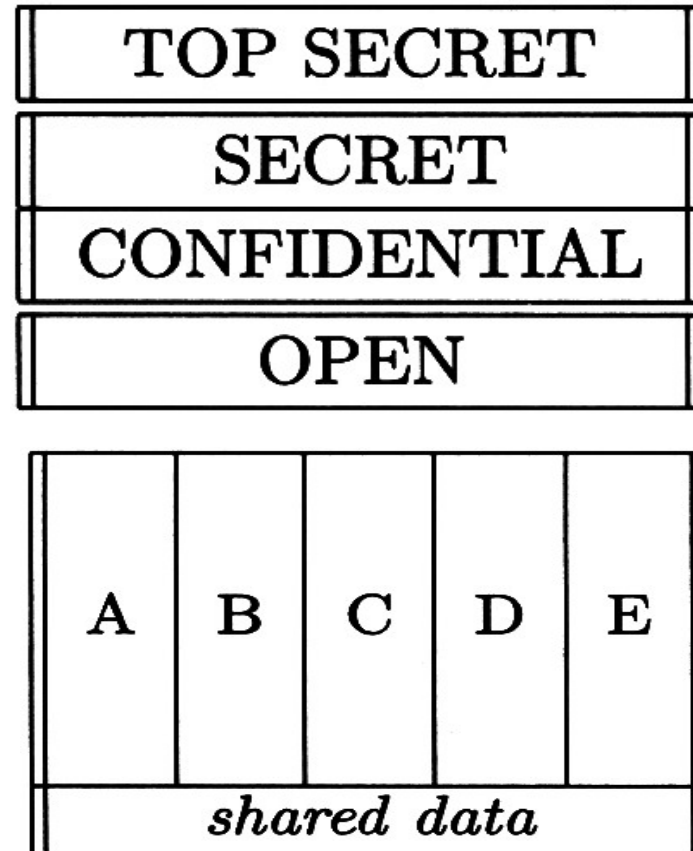
Picture hidden in three least significant bits of text

Multilevel Integrity

- The Biba model – data may flow only down from high-integrity to low-integrity
- Dual of BLP: (“BLP upside down”)
 - **Simple integrity property:** subject may write to object iff object has same or lower label as subject
 - ***-integrity property:** subject may read object iff object has same or higher label as subject
- So you have low watermark properties, etc
- Example: medical equipment with two levels, “calibrate” and “operate”

Multilateral Security

- Sometimes the aim is to stop data flowing down
- But other times, you want to stop lateral flows
- Examples:
 - Intelligence
 - Competing clients of an accounting firm
 - Medical records by practice or hospital



The Chinese Wall Model

- Industries such as investment banking, advertising and accounting have too few top firms for each big client to have its own
- So maybe you're auditing BP, and Shell too!
- Add a “Chinese Wall” that stops the two teams communicating
- *A subject can only access information that does not conflict with information they have already accessed*
- Idea: use a refinement of Bell-LaPadula

The Chinese Wall Model

- Subtlety: it's not enough to stop a Shell analyst reading BP data; must also stop a BP analyst writing data to a Barclays file that the Shell analyst can also read.
- For each object O , let $y(O)$ be the firm it relates to.
 $O = \text{bp-payroll.xls}$
 $y(O) = \text{BP}$
- Let $x(O)$ be that firm's conflict-of-interest class.
 $x(O) = \{\text{set of all oil companies}\}$
- Let $x(O) = x_0$ (a special COI class) if the information has been sanitized, so anyone is allowed to see it

The Chinese Wall Model

(Brewer and Nash, 1989)

Reminder: $x(\text{bp-file}) = \text{all oil companies}$; $y(\text{bp-file}) = \text{BP}$.

- **Simple property:** access (for reading or writing) is granted only if the object belongs to a firm the subject has already accessed, or to a firm in a different conflict-of-interest class

S can access O only if, for all O' that S has accessed, either $y(O)=y(O')$ or $y(O) \notin x(O')$

Boolean matrix holds state: $N_{s,o} = \text{“did S ever access O?”}$

- ***-property:** Writing is only allowed if simple property satisfied and the subject cannot read any unsanitized object from other firms

S can write O only if (simple property grants access) and (S is not able to read any O' with $y(O) \neq y(O')$ and $x(O') \neq x_0$)

Flow of unsanitized info is kept to within the firm's dataset.

Can you spot the problem? (Was an exam question in 2013.)

Bookkeeping, c. 3300 BC



Bookkeeping c. 1100 AD

- How do you manage a business that's become too large to staff with your own family members?
- Double-entry bookkeeping – each entry in one ledger is matched by opposite entry in another
 - E.g. firm sells £100 of goods on credit – credit the sales account, debit the receivables account
 - Customer pays – credit the receivables account, debit the cash account

(Some of these may sound backwards but make sense to accountants)
- So bookkeepers have to collude to commit fraud

Banking Security Policy

- Threat model:
 - 1% of staff go bad each year
 - Mistakes happen – 1 in 500 paper transactions
 - There are clever fraudsters too
 - Loss of confidence means ruin
- Protection goals:
 - Deter/prevent the obvious frauds
 - Detect the rest as soon as possible
 - Be able to defend the bank's actions in court

The Clark-Wilson Policy Model

Work by David Clark (MIT) and David Wilson (Ernst & Whinney) in 1986 to model double-entry bookkeeping

- In addition to the normal objects in your system, which we call unconstrained data items (UDIs), you add constrained data items (CDIs)
- CDIs are acted on by special programs called transformation procedures (TPs) that preserve the invariants
- IVPs (integrity verification procedures) verify the validity of CDIs (eg that the books balance)
- Mental model: a TP in a bank must increase the balance in one CDI (account) by the same amount that it decrements another

Clark-Wilson rules

1. There's an IVP to validate integrity of each CDI
2. Applying a TP to a CDI maintains integrity
3. A CDI can only be changed by a TP
4. Subjects can use only certain TPs on certain CDIs
5. Triples (subject, TP, CDI) enforce separation of duty
6. Special TPs on UDIs can produce CDIs
7. Each TP application must be logged to special append-only CDI
8. System must authenticate subjects that attempt to launch a TP
9. Only special subjects (admins) can change auth lists

Clark-Wilson importance

- First influential security policy model not based on BLP
- Application-level security state
 - The audit log (with enough info to reconstruct each TP)
 - The triples
- Separation of duties
 - In parallel (require 2 signatures, e.g. for large and irreversible transactions)
 - In series (different people for raising an order, accepting delivery, paying invoice, balancing budget)

Ubiquitous computing

- Authentication and device pairing without infrastructure
 - Two devices meet for the first time
 - No online servers available
 - Can't do the key distribution protocols studied later in this course
 - Can't use PKI, because you can't check for revoked keys
 - One device wants the other to “do something”
 - Authentication as temporary master-slave pairing:
Secure Transient Association
e.g. Smart Home devices + Universal Controller (phone)

Big Stick policy

(Stajano, 2000)

Whoever has physical control of the device
is allowed to take it over

- E.g. you can press the hard reset button on your router and the admin password is restored to factory default
- A trivial policy; but effective because cynically realistic
- Works fine for your lawnmower, pocket calculator (stateless) or for your fridge or router (inside the home)
- But not good for devices with valuable state that may be left unattended (e.g. a vending machine, a wireless sensor)

Imprinting

- Inspired by Konrad Lorenz (1973 Nobel prize)
- First moving subject seen by duckling becomes its mother
- Duckling stays faithful to mother until death



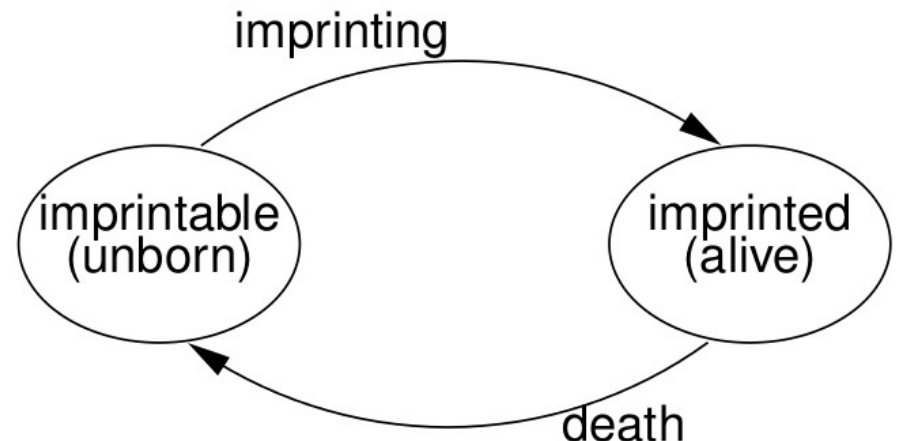
Out of metaphor:

- Slave device starts as imprintable
- First device that gives it a key becomes its master
- Bootstrap with unmediated physical channel
...but what if you then want to sell your Blu-Ray player?

Resurrecting Duckling policy

Stajano, Anderson 1999

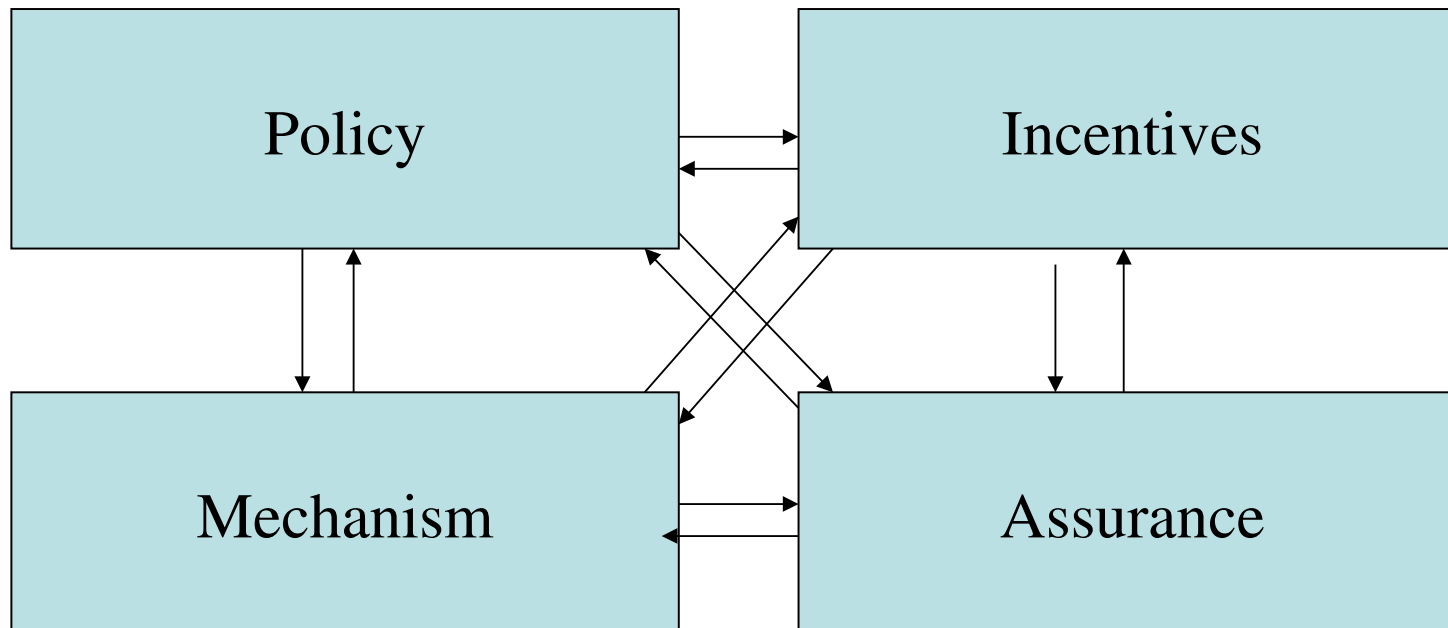
- **Two-state:** *duckling* can be imprintable or imprinted
- **Imprinting:** transition to imprinted when someone (henceforth *mother duck*) sends imprinting key over secure channel
- **Death:** transition to imprintable when mother duck orders it (like seppuku of the samurai)
- **Assassination:** uneconomical for attacker artificially to cause transition to imprintable (implies tamper resistance)



Policy

- Bell-LaPadula, Biba and Clark-Wilson are only three early examples of policy
- Many industries develop their own, and may get Protection Profiles evaluated
- Many things go wrong – people protect the things they can, not the things they should
- We often see deception at the policy level!
- For now, here's a useful framework

A Framework



Note how tinkering with mechanisms often feels easier than fixing incentives ...

Availability Policies

- Until recently, researchers ignored availability, but it's where the money goes

	research	industry
confidentiality	90%	1%
integrity/authenticity	9%	9%
availability	1%	90%

- Good example where availability matters: alarms!
Which bring us to...

Physical security

The Physical Security Revolution

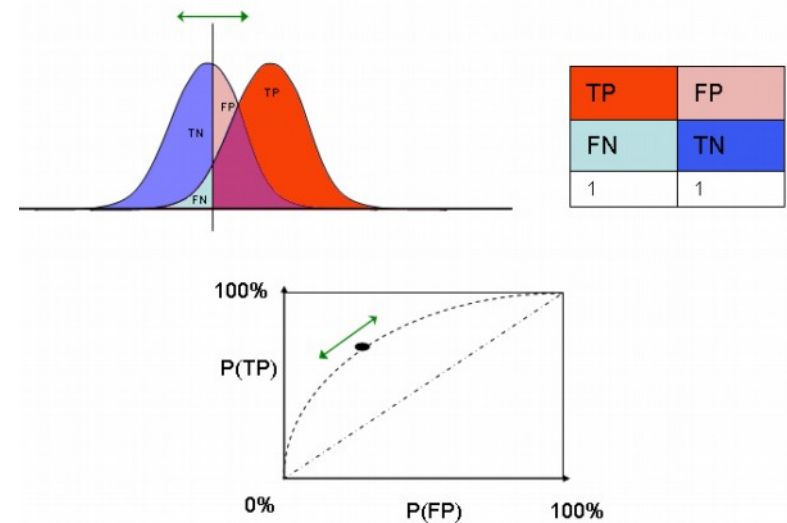
- IT security is rooted in physical security (for server rooms, crypto boxes etc)
- Old model: firms like Chubb with proprietary fire / burglar alarm systems; locks with master-keying systems
- New model: sensors run off ethernet like everything else
- We should be able to do better than metal systems
- Should be much easier to manage too – but many tensions (manageability, dependability)

Burglar Alarms

- Good example of a service where availability matters!
 - If there's a burglar in my vault I want the alarm company to be told!
 - Not bothered about confidentiality: you can tell other people too
 - Nor about authenticity: I don't care who tells them
- Wide range of systems: homes – supermarkets – jewellery stores – banks – nuclear facilities
- Wide range of standards, from Underwriters' Labs to the IAEA

How to Steal a Painting (1)

- Hollywood idea of art theft: cut through roof, climb down rope, grab painting without stepping on pressure mat (i.e. sensor defeat), get girl...
 - Response to this perceived threat is: more, fancier sensors
 - There are limits: set by false alarm rates and environmental conditions
 - Critical science: the Receiver Operating Characteristic (ROC) curve
 - Multisensor data fusion is really hard!
- But most high-grade attacks don't defeat sensors



How to Steal a Painting (2)

- More common type of art theft: hide in broom cupboard, come out at midnight, grab the Rembrandt and head for the fire exit
 - Understand the service you're supplying: deter – detect – alarm – delay – response
 - Don't rely on tech too much: 'Titanic effect'
- Or just toss in a smoke grenade. The fire alarm turns off the burglar alarm. Dash in and grab the Rembrandt
 - If caught, claim you were passing by and dashed in to save the national heritage

How to Steal a Painting (3)

- Wait for a dark and stormy night, when false alarms will be common. Create several (fence rattling). Wait till guards stop responding
 - Typical police force blacklists a property after 3–4 false alarms
 - Fix: multiple sensors, e.g. CCTV inside
 - Problem: we want best sensors on the outside for delay, but on the inside for low false alarm rate
- This is the standard way for professionals to do a bank vault! (Attack trust in the system)

How to Steal a Painting (4)



- Cut the wire from the sensor to the controller
- Connect a bogus controller to the phone line
- Cut the communications to the controller
- Cut the communications to many controllers
 - 2 independent channels for risks over £20m
 - Armed response force on premises for plutonium
- Insurance companies would like resilient anonymous communications to make service denial hard

Alarms – Lessons Learned

- Dealing with service denial is becoming more important, and harder
- Trade-off between false alarm rate and missed alarm rate is central
- You need to be clear what the service you're supplying (or buying) is – is it about sounding the alarm, or more?
- Critically, we need to design the system around the limitations of the human response. E.g. in airport screening, you insert deliberate false alarms. But what more can be said?

Really physical security

Basics of pin locks




When all pins align at the shear line, the cylinder can rotate and the lock opens.



Basic lockpicking

- While applying torque, pick one pin at a time, in order of stickiness
- Cost is linear in number of pins (not exponential!)
- Requires a modest (but non-zero) amount of skill




Basic lockpicking

- When all pins align at the shear line, cylinder can rotate and lock opens.
- While applying torque, pick one pin at a time, in order of stickiness
- Cost is linear in number of pins (not exponential!)
- Requires a modest (but non-zero) amount of skill




Lock Bumping

With lock bumping, that's not necessary




- Cut a key down to the (0,0,0,0,0) bitting
- Put it in the keyway, apply torque, and tap
- Pins bounce up to shear line and cylinder turns

Lock Bumping



- With fancy locks, like sidebar here: break that separately first
- pick it, or steal or photograph a key
- Enthusiasts have now defeated most mechanical locks

Master-key locks



- Intent: allow a single *master key* to open many locks that have different *change* keys
- Implementation: some or all pins are cut in more than one place

Example:

- Lock A opened by 11111
- Lock B opened by 22222
- Both opened by 44444

(but note unintentional cross-keying:
14114 opens A; 22442 opens B...)

Blaze attack on master key systems

Matt Blaze, “Rights Amplification in Master-Keyed Mechanical Locks”. *IEEE Security & Privacy* 1(2): 24-32 (2003). <http://www.crypto.com/papers/mk.pdf>

- **Preconditions:** Attacker (insider?) has one change key, several blanks, access to (own?) lock
- **Outcome:** Attacker recovers master key for whole system
- **Strategy:** For each pin, try all possible cuts, but copy the known key on the other pins (easy to find the “other” cut for this pin!)