

Security Engineering

INFR11208 (UG4) // NFR11228 (MSc)



Daniel W. Woods

Email: daniel.woods@ed.ac.uk

Reading for today's lecture

Security Policies

Ross Anderson <ross.anderson@cl.cam.ac.uk>
University of Cambridge Computer Laboratory

Frank Stajano <frank.stajano@cl.cam.ac.uk>, <fms@att.com>
University of Cambridge Computer Laboratory
AT&T Laboratories Cambridge

Jong-Hyeon Lee <jhlee@filonet.com>
Filonet Corporation

Abstract

A security policy is a high-level specification of the security properties that a given system should possess. It is a means for designers, domain experts and implementers to communicate with each other, and a blueprint that drives a project from design through implementation and validation.

We offer a survey of the most significant security policy models in the literature, showing how “security” may mean very different things in different contexts, and we review some of the mechanisms typically used to implement a given security policy.

<https://www.cl.cam.ac.uk/archive/rja14/Papers/security-policies.pdf>

Last week's lecture

Depending on the task, a security engineer might have to worry about:

1. Criminals (the crooks)

- Ransomware gangs, botnet operators, fraud gangs, malicious insiders

2. State actors (The spooks)–

- Five eyes; Russia; China; third-tier

3. Lawful operators (The geeks) –

- Employees, security researchers, competitors

4. The swamp

- hate crimes, bullying, family members etc


This week: What do we do with this information?

Example: Protecting Exam Scripts

- Process
 - Lecturers draft exams, sharing with reviewers
 - Scripts sent to graduate school to store and print
- Who are the threat actors?
- What should the security policy be?


Security policies as leadership/strategy

- Security policies are industry and sometimes even organization specific
 - Tailored to the threat landscape
 - Need to balance protection goals against other non-security goals
- By contrast, security mechanisms are mostly universal



Design Hierarchy

- What are we trying to stop?
- How are we trying to stop it?
- With what mechanisms?




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?

A simple security policy for a consulting firm



Problems

- Any off-site work
 - Client meetings
- Not all employees can be trusted to access **all** assets
- We want some assets to be accessed by external actors
 - Marketing
 - Email server

What's wrong with this?

A realistic security policy for a consulting firm

(and most of the firms you'll work for)

Threat model

Cybercriminals.

Security Policy

Put sensitive resources behind contextual access controls, with access determined by the owner.

Security Mechanisms

Non-sensitive resources (email, website) on public Internet.
Everything else on internal network, remote access by VPN.

Why this works better

- VPN solves offsite work
- Asset owners decide who should have access
 - HR can give access to investigations to only those who need it
- Public resources are public


Discretionary Access Control (DAC)

Discretionary Access Control (DAC)

Assumptions in a DAC security policy

- Users can be trusted to set policy for and defend each asset
 - Consequences of one user failing aren't that bad
- Usually have "crown jewels" controlled centrally
 - Finance, HR etc
 - But this can get v specific (exams in HE)
- Fine if you can accept the cost of discretion being misused

What about for an intelligence agency?



Discretionary Access Control (DAC)

Security policies that assume insider threat

- Insider threat could be a disloyal employee, **or** malware on their laptop
 - In an intelligence agency, tell the opponents or the press what's happening
 - In a health system, look at sensitive personal information such as celebrities' records
 - In a bank, steal money
- The following 3 policies are designed to limit the damage by removing discretion of asset owners

Mandatory Access Control (MAC)

Multilevel Secrecy

Access determined by position in hierarchy


First Policy Example – MLS

- Multilevel Secure (MLS) systems are widely used in government
- Goes back to President Roosevelt, 1940: a clerk with 'Secret' clearance can read documents at 'Confidential' and 'Secret' but not at 'Top Secret'
 - Easy to implement in a building with locks on doors
- 60s/70s: Anderson report (1973)
 - Following physical security, USAF used separate machines for top secret, secret etc but wanted to move to more efficient time-sharing model
 - **Problem** What if a general runs a virus that copies data to unclassified
 - **Solution** Reference Monitor and MAC
 - 'I don't care if you want to write this Top Secret data to a public drive; I won't let you'

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
 - Secret: compromise could threaten life directly. E.g. weapon system performance
 - Confidential: compromise could damage operations
 - Official: compromise might embarrass?
- Resources have classifications
- People (principals) have clearances
- Information flows upwards only
 - At what cost?

Computer Information Flows



How could you translate this into a policy for which employees communicate with each other? At what cost?

Formalising the Policy

- Initial attempt – WWMCCS – just said that no process could read a resource at a higher level. Not enough!
- Bell-LaPadula (1973):
 - *simple security policy*: no read up
 - **-policy*: no write down
- Theorem: a safe system stays safe
- Ideal: minimize the Trusted Computing Base (set of hardware, software and procedures that can break the security policy) in a *reference monitor*

Problems with Bell-LaPadula

- Processes such as memory management, need to read and write at all levels
- Fix: put them in the trusted computing base
 - Pointless if top secret info is copied to backup tapes along with unclassified
- 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)

Further problems with multilevel security

Pentagon leaks show difficulty of keeping secrets in a vast intelligence network

Britain issues more than 180,000 security clearances every year, as **Kim Sengupta** explains



Thursday 13 April 2023 21:40 BST

"The number of employees and contractors across the US administration with top-secret clearance is currently more than 1.25 million"

<https://www.independent.co.uk/news/uk/home-news/pentagon-leaks-security-clearance-employees-b2319307.html>

The Original Sin Is We Classify Too Much

Overclassification has a range of harms, from stifling democratic debate to harming national security itself.



Elizabeth Goitein

January 26, 2023

Government Power

<https://www.brennancenter.org/our-work/analysis-opinion/original-sin-we-classify-too-much>

Terminology


- 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

Multilateral Secrecy

Access determined by relationship to data


Multilateral Security

- Sometimes aim is to stop data flowing down
 - From Top Secret to Secret
 - Misconduct investigations from HR down
 - Exam answers from lecturer to students
- More often, you want to stop lateral flows
 - Intelligence
 - Competing clients of an accounting firm
 - Medical records by practice or hospital



The Lattice Model

- This is how intelligence agencies manage ‘compartmented’ data – by adding labels
- Basic idea: BLP requires only a partial order




What didn't work so well (NHS)

- 1996: medical records in 11,000 surgeries
 - Prevents mega breaches, at the cost of micro breaches?
 - Hard to manage changing GPs
- 2021: now on three cloud services
 - Give patients access to own records, test results, and prescriptions
 - Multi level would fail : doctors have "Top Secret" access to entire country's medical records
- Idea: access by role and relationship
 - Need to be a doctor AND the patient needs to be getting treatment from you

Alternative lateral flow controls

- Chinese Wall Model
 - Accountancy firm: if you've worked for an oil company, you can't work for a competing oil company for (e.g.) two years
 - Bank: If IB works funding for a merger, the bank's traders can't know
 - Requires tracking state
- How reliably is this enforced?
 - "Cigarette on the Pavement" (ASIC v. Citigroup, 2005)



Alternative lateral flow controls

- Delegation
 - in a retail bank, you only get to see a customer's account details once they've passed authentication **for you**

Terminology matters

- A *subject* is a physical person
- A principal can be
 - a person
 - equipment (device receiving SMS)
 - a role (the executor of the will)
 - a complex role
 - Bank employee deputising for customer to check balance
 - Bank employee deputising for the executor who is deputising for the original customer




Authenticating the call centre employee as acting on your behalf

Multilevel Integrity

Modification determined by hierarchy

Multilevel Integrity


- The Biba model – data may flow only down from high-integrity to low-integrity
 - Dual of BLP!
- Example 1: electricity / gas / oil distribution
 - Safety: highest integrity level
 - Prevents harmful incidents
 - Must never be influenced by untrustworthy data.
 - Monitoring and control: next level
 - Monitors + stops usage (e.g. if no payment)
 - Enterprise apps (e.g. billing): third level
 - Collect payment from customer
- Colonial pipeline hack: operator turned off the pipeline when ransomware killed the billing system!



Small Group Exercise

Imagine you're writing a security policy for a small family.

- What kind of policy would you write?
 - What is sensitive?
 - Security levels or compartments?
 - Who gets to read what?
 - Who gets to write what?
 - Delegation?



2024 CRIME TYPES *continued*

BY COMPLAINT LOSS

Crime Type	Loss
Investment	\$6,570,639,864
Business Email Compromise	\$2,770,151,146
Tech Support	\$1,464,755,976
Personal Data Breach	\$1,453,296,303
Non-Payment/Non-Delivery	\$785,436,888
Confidence/Romance	\$672,009,052
Government Impersonation	\$405,624,084
Data Breach	\$364,855,818
Other	\$280,278,325
Employment	\$264,223,271
Credit Card/Check Fraud	\$199,889,841
Identity Theft	\$174,354,745
Real Estate	\$173,586,820

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 entries 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
- Why?

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 can mean 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 (accountant) in 1986 to model real bookkeeping systems
- 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)
- **Mental model:** a TP in a bank must increase the balance in one CDI (account) by the same amount that it decrements another
 - Double entry book-keeping

Clark-Wilson Framework

Unconstrained data items (UDIs)

Constrained data items (CDIs)

Transformation procedures (TPs)

- There's an IVP to validate CDI integrity
 - Validate double entry book keeping upheld
- Applying a TP to a CDI maintains integrity, and only TP can change CDI
 - TP = transfer money between accounts designed to uphold double entry
 - Your bank account balance can only be changed by a transfer/deposit action
- Subjects can use only certain TPs on certain CDIs
 - You can only action balance transfers on your own account
- Triples (subject, TP, CDI) enforce separation of duty
 - Person who can manipulate sales account cannot also manipulate receivables
- Each application of a TP writes enough for an audit-trail CDI to reconstruct its action
- Only special subjects (security officers) can set up and alter triples

Lessons learned from security policies

- No single solution to the insider threat!
- Multilevel security policies first to be explored, thanks to the military
 - Used for safety/integrity as well as secrecy
- Multilateral policies mitigate effects of scale
 - Patient records, Chinese walls in finance
- Often need to integrate roles/relationships/dependencies
 - Call center employee acting on behalf of customer X
 - Dr Foster acting as the GP of patient X
 - Can't divert sales to your personal account without colluding with receivables not to debit their account
- Academics make careers on specifying formal models (Bell-LaPadua, Wilson Cox etc), but actual failures often outside model
 - Implementation problems or side channels

Questions to think about

- Which systems in your life deploy mandatory access control?
- What about discretionary?
- Why isn't MAC more widely deployed in systems?