



THE UNIVERSITY *of* EDINBURGH

Operating Systems Security

Security Engineering (Spring 2026)

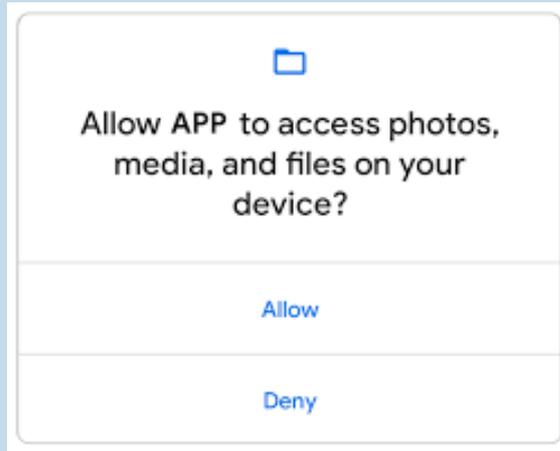
Lecturer: Jingjie Li & Daniel Woods



- Operating systems managing hardware and software resources – main entry point of enforcing security policies.
- How does an OS enforce security policy we defined?
- How do we isolate resources for security?



Who are interested in defining and setting up a security policy?





Access Control Matrix

| | | Objects (files) | | | | |
|-----------------------------|----------------|------------------------|----------|----------|----------|----------|
| | | a | b | c | d | e |
| Subjects (users) | jingjie | r,w | - | r,w, own | - | r |
| | bob | - | - | r | r | r,w |
| | alice | w, own | r | r | - | - |
| | eve | r | r,w | r,w | - | r |

Discretionary Access Control - Access Control List

| | | Objects (files) | | | | |
|---------------------|---------|-----------------|-----|----------|---|-----|
| | | a | b | c | d | e |
| Subjects (users) | jingjie | r,w | - | r,w, own | - | r |
| | bob | - | - | r | r | r,w |
| | alice | w, own | r | r | - | - |
| | eve | r | r,w | r,w | - | r |

- Access Control Lists: store permissions with file. May need different permissions for different programs, so actually a (user, file, program) triple
- ACLs scale badly without Role-based access control.
- Finding all the files a user has access to is a massive pain.

Discretionary Access Control - Capabilities

| | | Objects (files) | | | | |
|---------------------|---------|-----------------|-----|----------|---|-----|
| | | a | b | c | d | e |
| Subjects (users) | jingjie | r,w | - | r,w, own | - | r |
| | bob | - | - | r | r | r,w |
| | alice | w, own | r | r | - | - |
| | eve | r | r,w | r,w | - | r |

- Capabilities: store per user, not per file.
- Finding all the users who have access to a file is a pain.
- Hard to revoke access to a particular file, or produce evidence of who could have broken said file.
- Easily transferred
- Public key certificates are really capabilities.



Unix Access Control Lists

```

sam@sam-W65-67SC: ~/.ssh
sam@sam-W65-67SC:~$ cd .ssh
sam@sam-W65-67SC:~/.ssh$ ls -la
total 40
drwx----- 2 sam sam 4096 Mar 26 16:42 .
drwxr-xr-x 76 sam sam 4096 Aug 11 18:38 ..
-rw----- 1 sam sam 1671 May 14 2020 id_dsa
-rw----- 1 sam sam 1675 May 14 2020 id_rsa
-rw-r--r-- 1 sam sam 398 May 14 2020 id_rsa.pub
-rw----- 1 sam sam 5546 Mar 31 15:55 known_hosts

sam@sam-W65-67SC: ~
sam@sam-W65-67SC:~$ groups sam
sam : sam adm cdrom sudo dip plugdev lpadmin sambashare
sam@sam-W65-67SC:~$
```

- In Unix (and thus Linux, Android, iOS...): rwx attributes, with owner,group,world, per file. Can have richer Posix extended ACL extension.
- Sysadmin can do anything!



Mandatory access control

- Security policy is defined and controlled centrally, enforcing for all users
- DAC: Start in supervisor mode, and as admin, can make less privileged accounts available for less trusted tasks.
- MAC: sysadmin no longer the boss: ultimate control rests with the security policy (possibly set by remote govt authority in defence setting).
- Alternate view (Android): DAC requires **permission** of the user. MAC requires **consent** of user, developer AND platform -- 3-party consent.



Mandatory access control

Bell LaPadula

- Simple Rule (No Read Up): A subject at a given security level may not read an object at a higher security level.
- * Property: (No Write Down): A subject at a given security level may not write to any object at a lower security level.

- Occasionally synonymous with MLS - Multi-Level Security (Unclassified, Confidential, Secret, Top Secret)
- Enforced by system policy, not by user discretion!
- Traditionally for military systems, e.g. Bell LaPadula



Mandatory access control

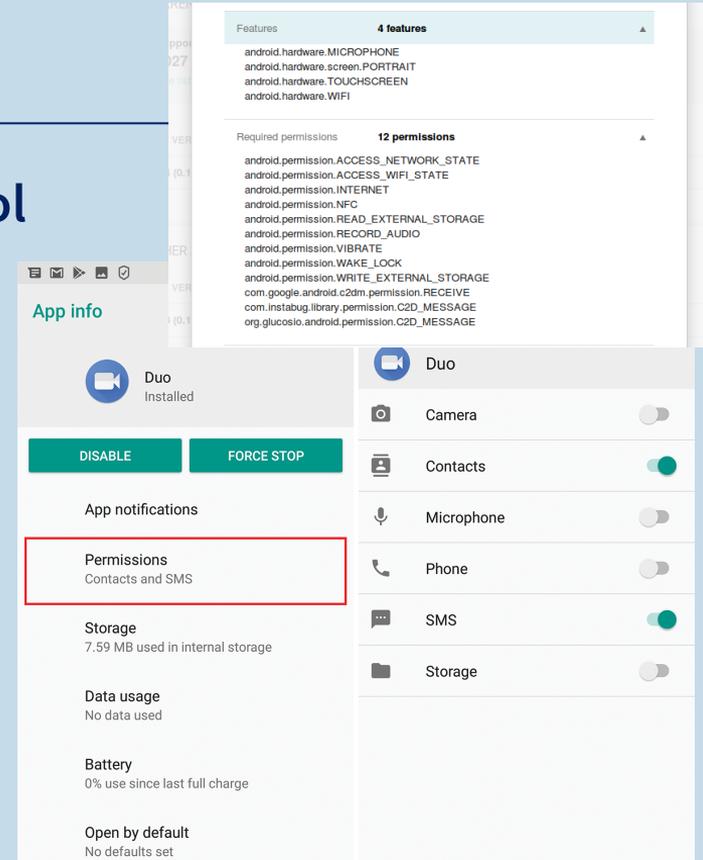
BIBA

- Simple Rule (No **Write** Up): A subject at a given security level may not write to an object at a higher security level.
- * Property: (No **Read** Down): A subject at a given security level may not read from any object at a lower security level.

- BIBA: Uses the opposite duality of confidentiality and integrity and thus reverses BLP -- low water mark: integrity of an object is the lowest level of all the objects that contributed to its creation

Case study: Android Discretionary Access Control

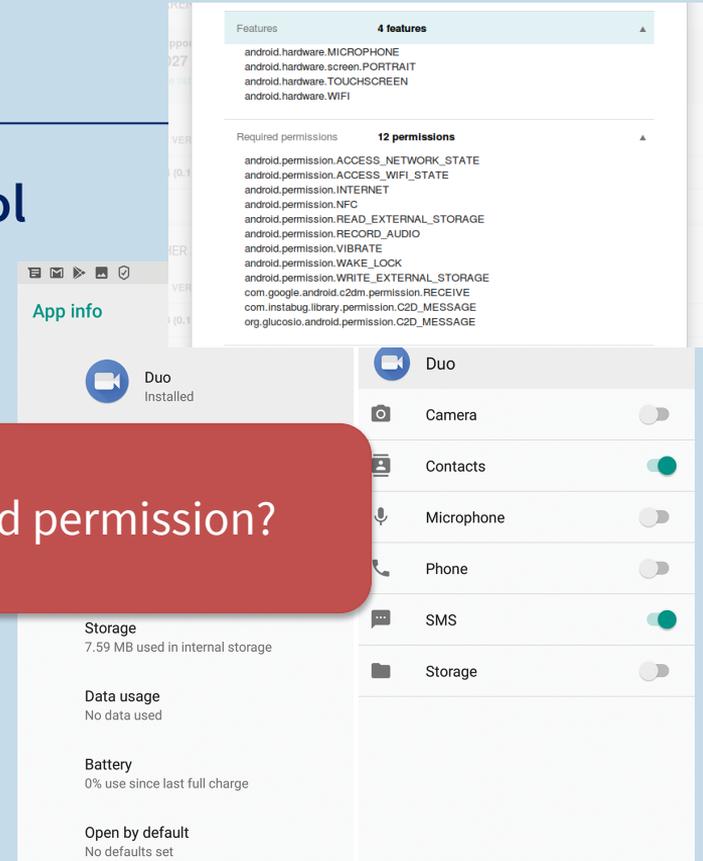
- App Isolation: Treat Apps by different companies as different users, using SETUID.
- Permissions also effectively capabilities, implemented by adding GIDs to the list of groups of the SETUID. “Permissions manifests” basically compile down to this.
- Early versions: all granted at install time. So flashlight apps started demanding your address book at install time so they could sell it.
- Since Android 6, Google moved to Apple model of TOFU, but earlier apps still demand on installation.



Case study: Android Discretionary Access Control

- App Isolation: Treat Apps by different companies as different users, using SETUID.
- Permissions also effectively capabilities.
implement
the SETUID
compile d
- Early versions: all granted at install time. So flashlight apps started demanding your address book at install time so they could sell it.
- Since Android 6, Google moved to Apple model of TOFU, but earlier apps still demand on installation.

Think: what can go wrong with Android permission?





Case study: Android permissions issues

- API has poor documentation, and the permissions system is often the enemy of the developer, who ends up requesting more permission than they really need.
- Android still has malware! e.g. Pegasus via zero day, but costs \$1 million.
Alternative markets out of Google's control
- And lots of unpatched devices. The OS-update ecosystem is a disaster...
- Getting access control right intersects with lots of awkward edge cases, e.g. factory reset
- Over-privilege and coarse permissions
- Third-party SDKs inherits the app's granted permissions.

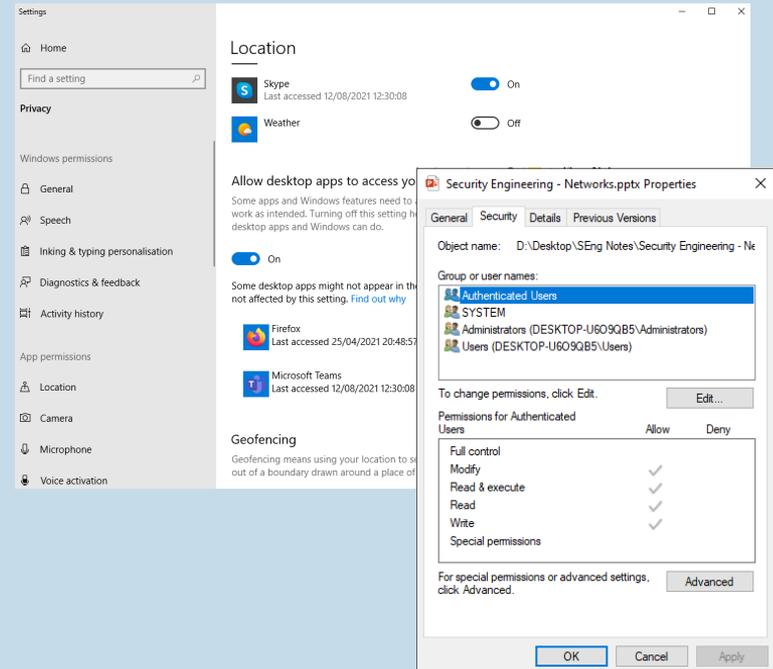


Case study: iOS

- Also a Unix derivative, via FreeBSD and Mach kernel.
- Domain and Type Enforcement for tamper-proof system components. App permissions are capabilities, granted on first use on consent.
- Signed ecosystem from the market, just as Android has its default supported Google Play. Allows screening and also revenue taking.
- On the App Store, Apple signs the binaries. On Google Play, the developer does.
- Biometrics stored via encryption by the secure enclave (SE). Neither iOS nor TrustZone are trusted with this data!
- Vertically integrated, closed ecosystem

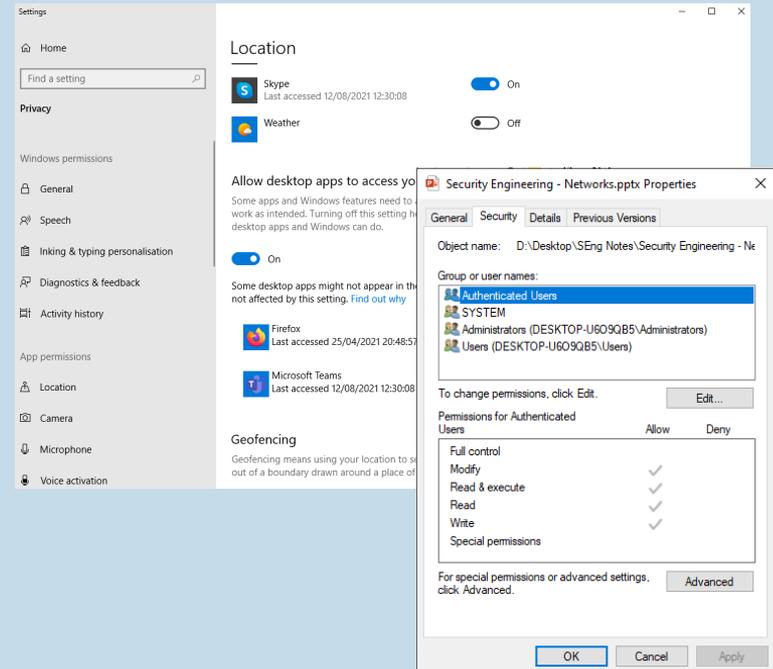
Case study: Windows

- Very complex Access Control, from Windows NT onwards. RWX, but also Take Ownership, Change Permissions and Delete.
- AccessDenied, AccessAllowed, SystemAudit: AD overrides AA if set multiple times.
- Not really integrated as a (user,program,file) triple – or even setuid.
- ACL just does (user,file), and a separate system handles a limited set of permissions, mostly for “apps”



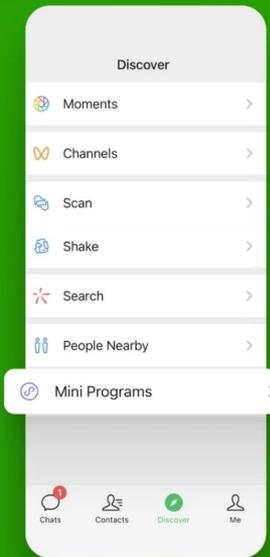
Case study: Windows – why so complicated?

- Corporate customers need complicated access controls. MS made half its revenue from firms >25000 seats.
- Decades of backwards compatibility means testing at scale. And introducing features slowly, and complex compatibility layers e.g. Application Information Service

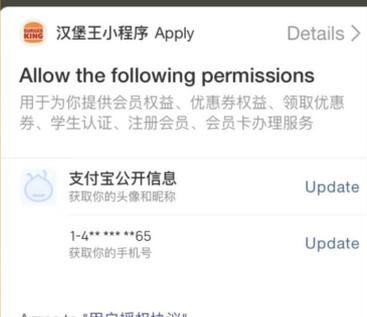


Discussion: Mini-app / app-in-app, what could go wrong in permission?

WeChat users in China complete their daily tasks within one app using mini apps




Alipay user name and photo →



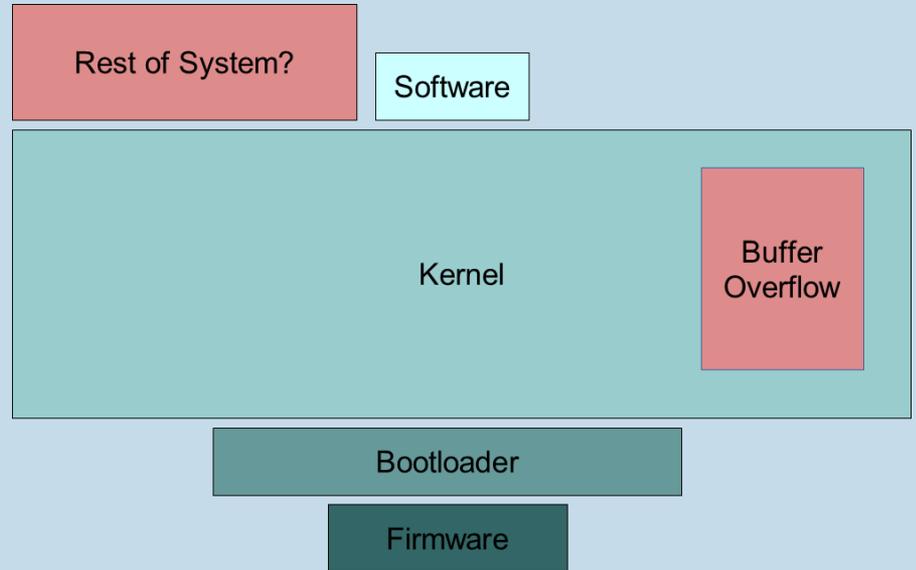
← Phone number

- What's the miniapp business model?
- What could go wrong in security?

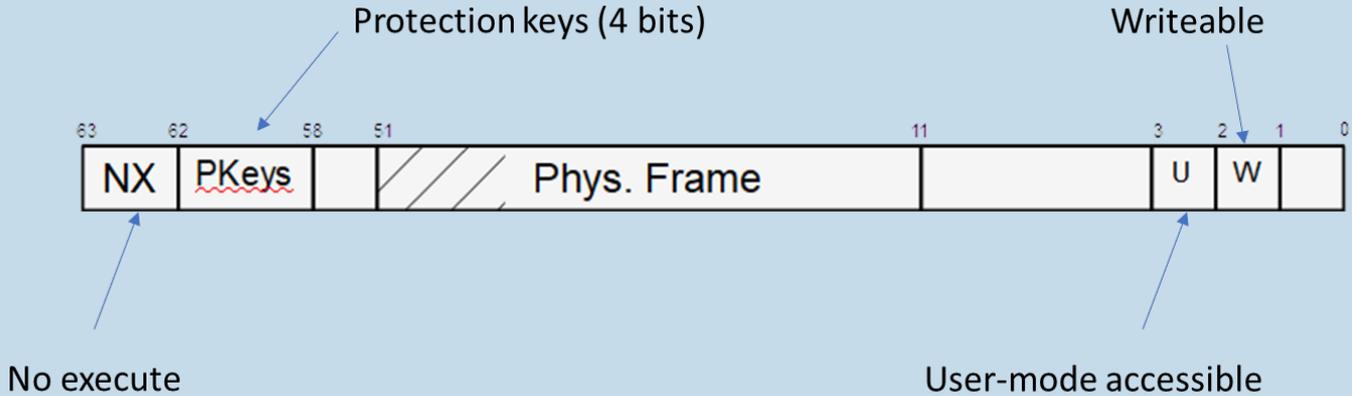


Isolation

- How do you stop others, using the same system(s), from being able to read your data / hack your software?
- I mean, really stop them (side channels, bugs in trust computing base)?

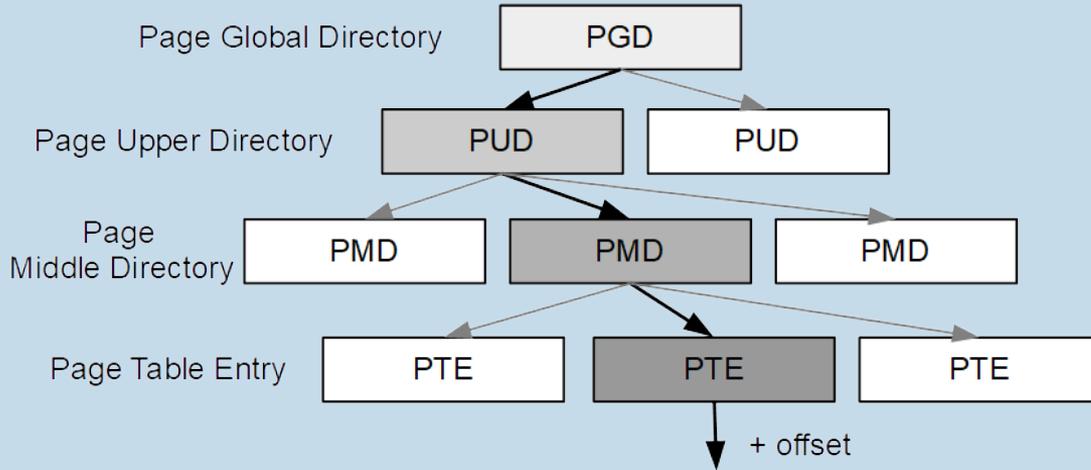


Isolation: Processes and Memory Management



You can get some isolation from individual **processes** in a single operating system. See Android last time for examples of precise **file-system** isolation, but it's also typical for OSes to support isolation within virtual memory.

Isolation: Page-Table Walk



- Up to four memory accesses, though many cached, to transform isolated virtual address to shared resource of physical frames.
- Data Execution Prevention (NX XOR W) – forces attackers to use Return-Oriented Programming.
- Address Space Layout Randomisation

Isolation granularities: Sandboxing, Virtualization and Containers



- Sandboxing e.g. Chrome, eBPF: often within-process
- Containers e.g. Docker: between-process with filtering
- Virtualization e.g. VirtualBox: whole separate “guest” OS on top of a shared “host” OS



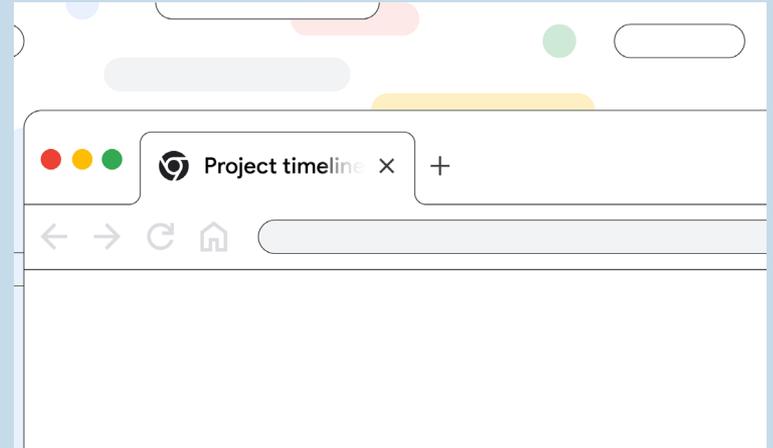
Isolation: Sandboxing

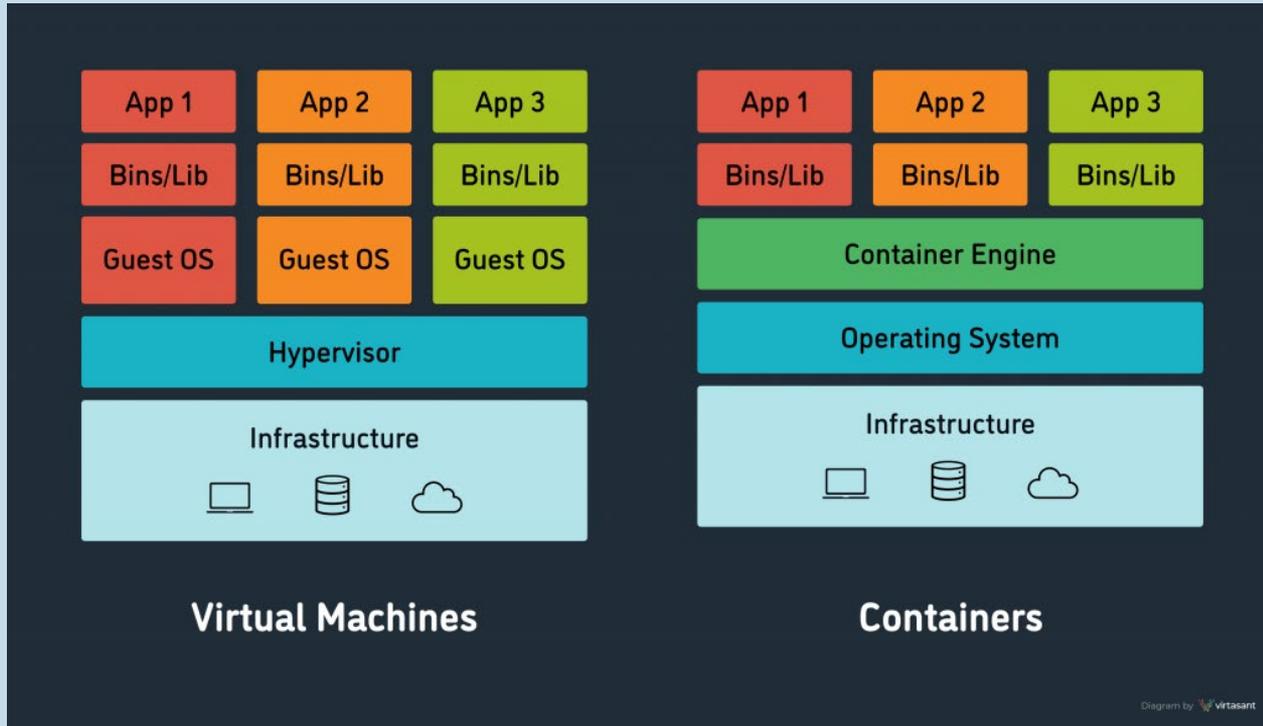
- Pioneered by Sun with the JVM
- Restrict the environment in various ways: temporary access to a single directory, communication via same-origin policy
- Prevent access of address space outside of a predefined region, e.g. the Javascript interpreter's memory in a browser (Spectre klaxon).



Isolation: Chrome

- Cross-website theft increasingly important, so Site Isolation not just Renderer Isolation.
- Increasingly done by process-level isolation: also useful for Spectre.
- The “same-origin policy” gets complicated when you have untrusted ads in the same tabs...
- Vulnerabilities often exploited by Drive-by-Download attacks.



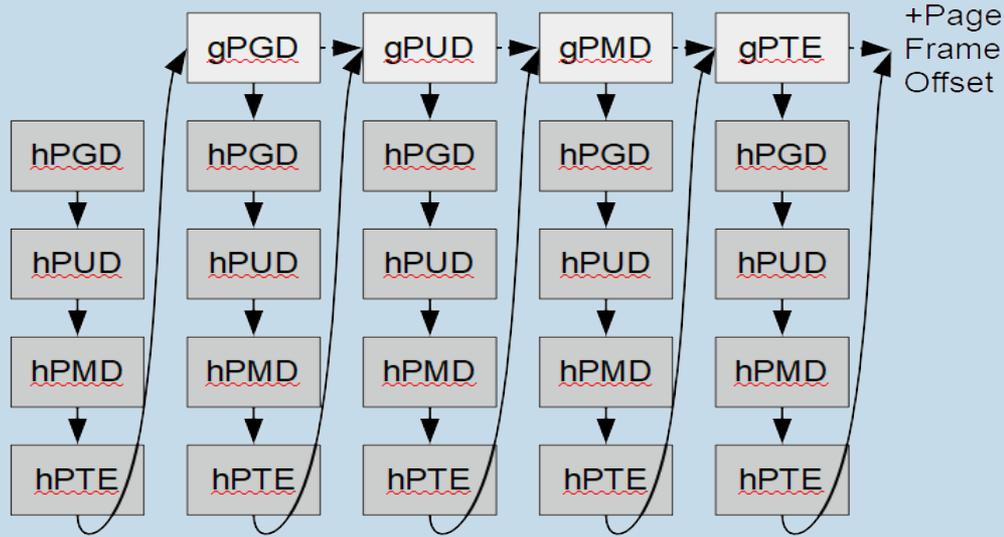




Isolation: Virtualization

- Replaces the entire operating system, and runs a “guest” operating system on top of a “host” via hypervisor.
- Powers cloud computing.
- HW support such as Intel VT-x makes things cleaner and faster.
- Why more secure? The hypervisor can be much smaller than a full OS and so easier to code-review and secure, right???
 - Why hardware compartmentalisation is still useful?
 - VM escape, cross-VM leaks...

Isolation: Virtualization



- Performance gain of nested page table adds to security costs



Isolation: Virtualization challenges

- Subtle issues around monitoring: If you're going to check all ACLs on your server, what about the containers or virtualized systems?
- Trouble at the interface: people still need to share data between VMs and ad-hoc mechanisms such as USB devices
- Bromium: VM per app, messy at the interface with untrusted files sent via host. Need specific exceptions and plugins, like Outlook being prevented from rendering files itself.



Isolation: Container

- Cheaper than a VM, but less secure. E.g. Docker
- Shared kernel, but isolates application code and libraries from the rest of the system -- not shared anymore.
- NSjail: trap processes in a "jail": restrict syscalls by seccomp, change root to a local directory.
- Virtualise some parts but not others e.g. PIDs, IPC and namespaces.
- Syscall filtering too, and sandboxing.



Isolation: Container challenges

- Not the same as VM, and not really meant for data isolation -- don't expect the same isolation as with a VM -- the trusted code base is still massive.
- Really for deployability, not security. Many subtle bugs, such as blank root passwords as defaults!
- On the flipside, deployability might **be** a security feature – why?

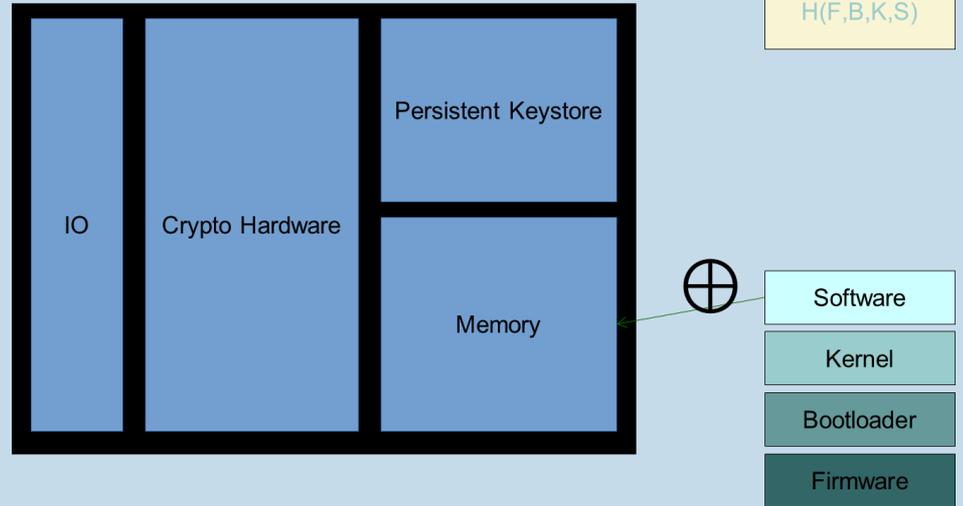


How do we trust remote systems, e.g., clouds?

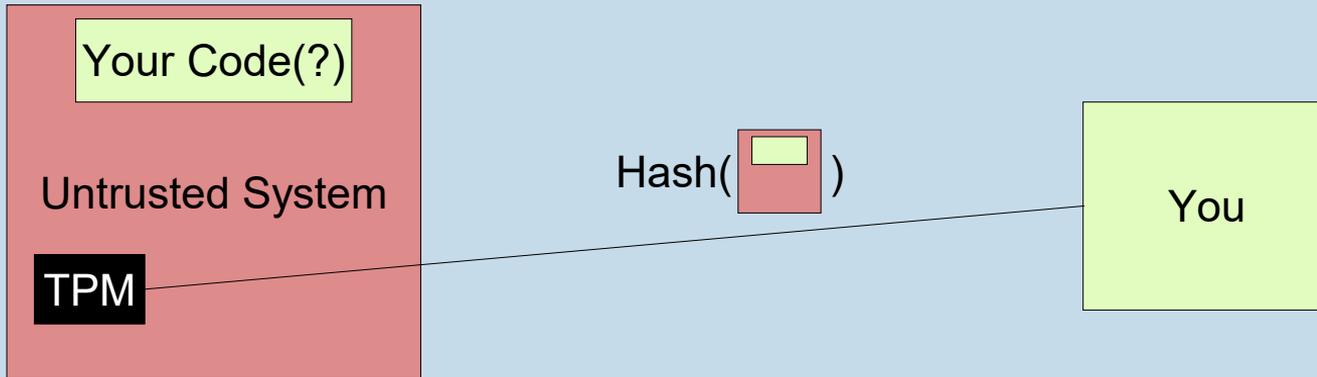
- Access control often specific to provider e.g. AWS Identity and Access Management – lock in
- Google App Engine: Provide only indirect access to networking/file system, by "safe" versions that integrate MAC and call other cloud infrastructure.
- Defence in depth: Integrate sanitization, mitigation/hardening, functionality reduction, indirection (Python -> NaCL) and ptrace filtering and logging.

Secure boot vs measured boot

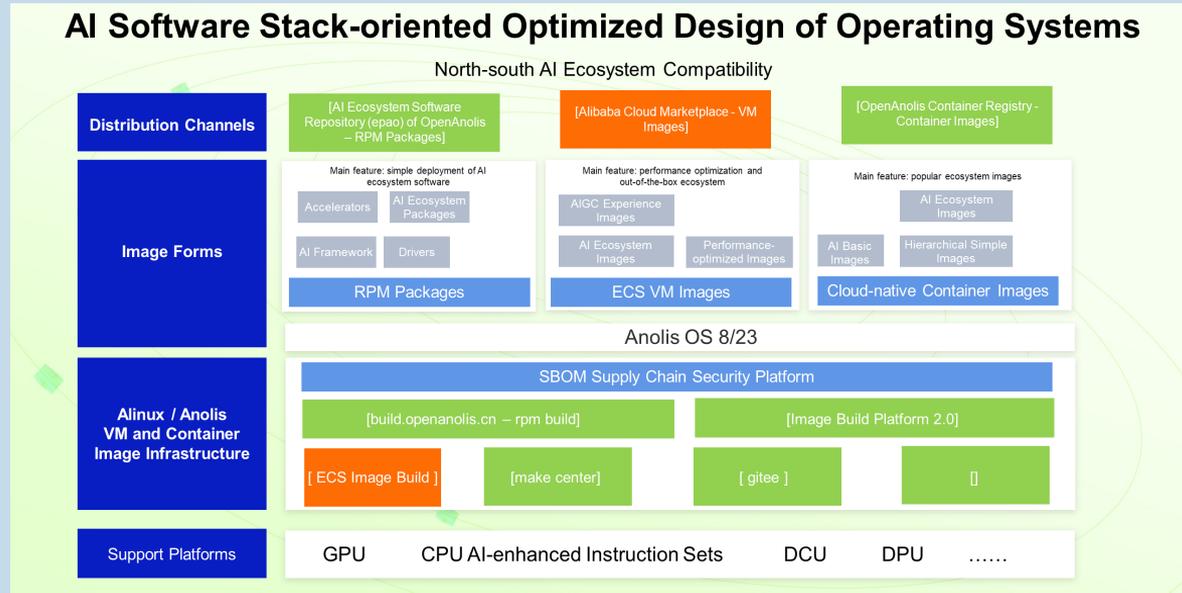
- Secure boot: verify then execute
 - Checks signatures against key
- Measure boot: measure then report
 - Creates and log hash
 - Report to a trust server
- Pros and cons?



Remote attestation

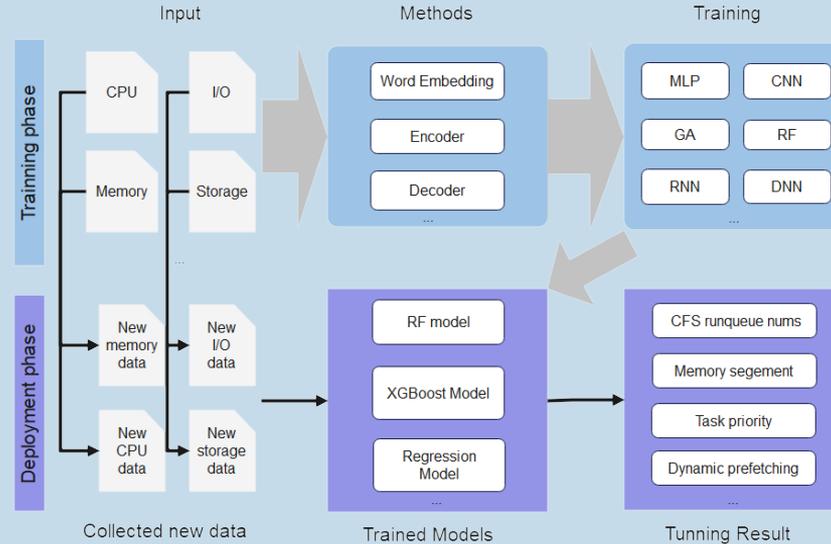


Think and discuss: security issues when AI meets OS?



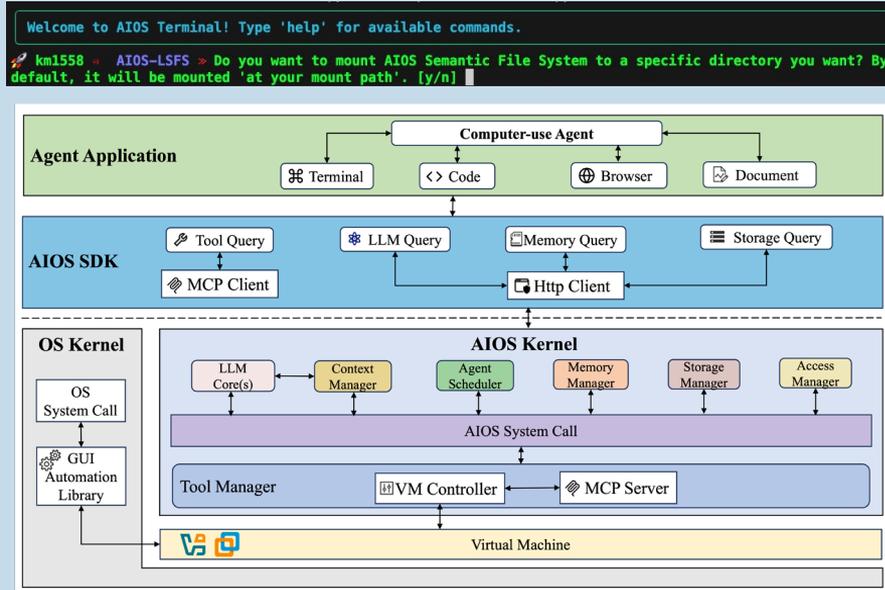
System for AI

Think and discuss: security issues when AI meets OS?



AI-integrated OS (scheduling)

Think and discuss: security issues when AI meets OS?



AI-copilot for end users



AI x OS challenges

- Probabilistic trust computing base
 - Same prompt yields different output due to sampling randomness
- Dynamic and task-specific security policy
 - No app, natural language task specification changing over time
- Fuzzy security boundary within an AI agent
- Dynamic instruction following = dynamic code loading?

comment on GitLab issue 745 that we're done

summarize issue 745 for me

Christodorescu, M., Fernandes, E., Hooda, A., Jha, S., Rehberger, J. and Shams, K., 2025. Systems Security Foundations for Agentic Computing. *arXiv preprint arXiv:2512.01295*.