

Operating Systems Security: concepts in security controls, vulnerabilities and attacks

A quantitative approach
2025-02-24

Robert Malmgren
rom@romab.com

1 minute presentation

- Consultant in IT, infosec and cybersecurity since 25+ years
- Working alot on with critical infrastrucutre protection, process control, SCADA security etc, but also in financial sector, government, etc
- Work covers everything from writing policies, requirement specs and steering documents to development, penetration testing, incident handling and forensics

Outline of talk

- Intro
- Background and basics
- Security problems & vulnerabilities
- Example of operating systems and security

Some short notes

- The focus is on general operating system used in general computers - COTS products
- *Embedded systems, code for micro controllers, etc often lack most fundamental security features*
- Some experimental OS's and domain specific solutions have better-than-average security concepts and security controls, e.g. military grade usage

Background and basics

Part I: protection, security controls

Intro - foundation

- Complex systems

- ...have multiple users,
- ...run multiple programs at once,
- ...store huge amounts of data,
- ...is interconnected via networks

Multiuser

Multitasking

Locally & remote

Multiple services
and clients

Intro - foundation: *Isolation*

- Modern software is normally formed into **components, parts** and **layers** in *systems*
- This will create a software stack
 - Layers in the stack provides abstraction
 - Layers in the stack provides supporting frameworks, functionality and support mechanisms
- Layers in the stack is one form of **isolation**

Intro - foundation: *Isolation*

- Layers and **isolation** is a way to provide **separation**, which can be:
 - **Logical/Virtual**: A way to make it appear that execution environment have exclusive access
 - **Physical**: Different computers, different CPUs/cores, different disks
 - **Time based**: Separation of execution time/Timeshare
 - Based on security technologies, i.e. cryptographic algorithms and crypto mechanisms can also be used to compartmentalise and isolate information.

Intro - foundation: IAM


- This there is to built-in security into the foundation of the systems - the operating system
 - To **identify** and **authorize** users of the system
 - To allow for an environment where necessary basic controls are in place
 - To prevent unauthorised access to OS resources

Capabilities and requirements


Need	Description	Example
Protect a system resource	<i>Prohibit malicious or unintentional access to system resources</i>	System tables, direct access to I/O-units, memory protection
Authorization checks for usage of system calls and system resources	<i>Provide controlled access to system, so that system maintain system integrity and provide continuous security to application and information</i>	reference monitor
Separation of resources	<i>Physical, Logical, temporal or cryptographical separation</i>	separation in running time

The classical *ring model*


UNIX



x86

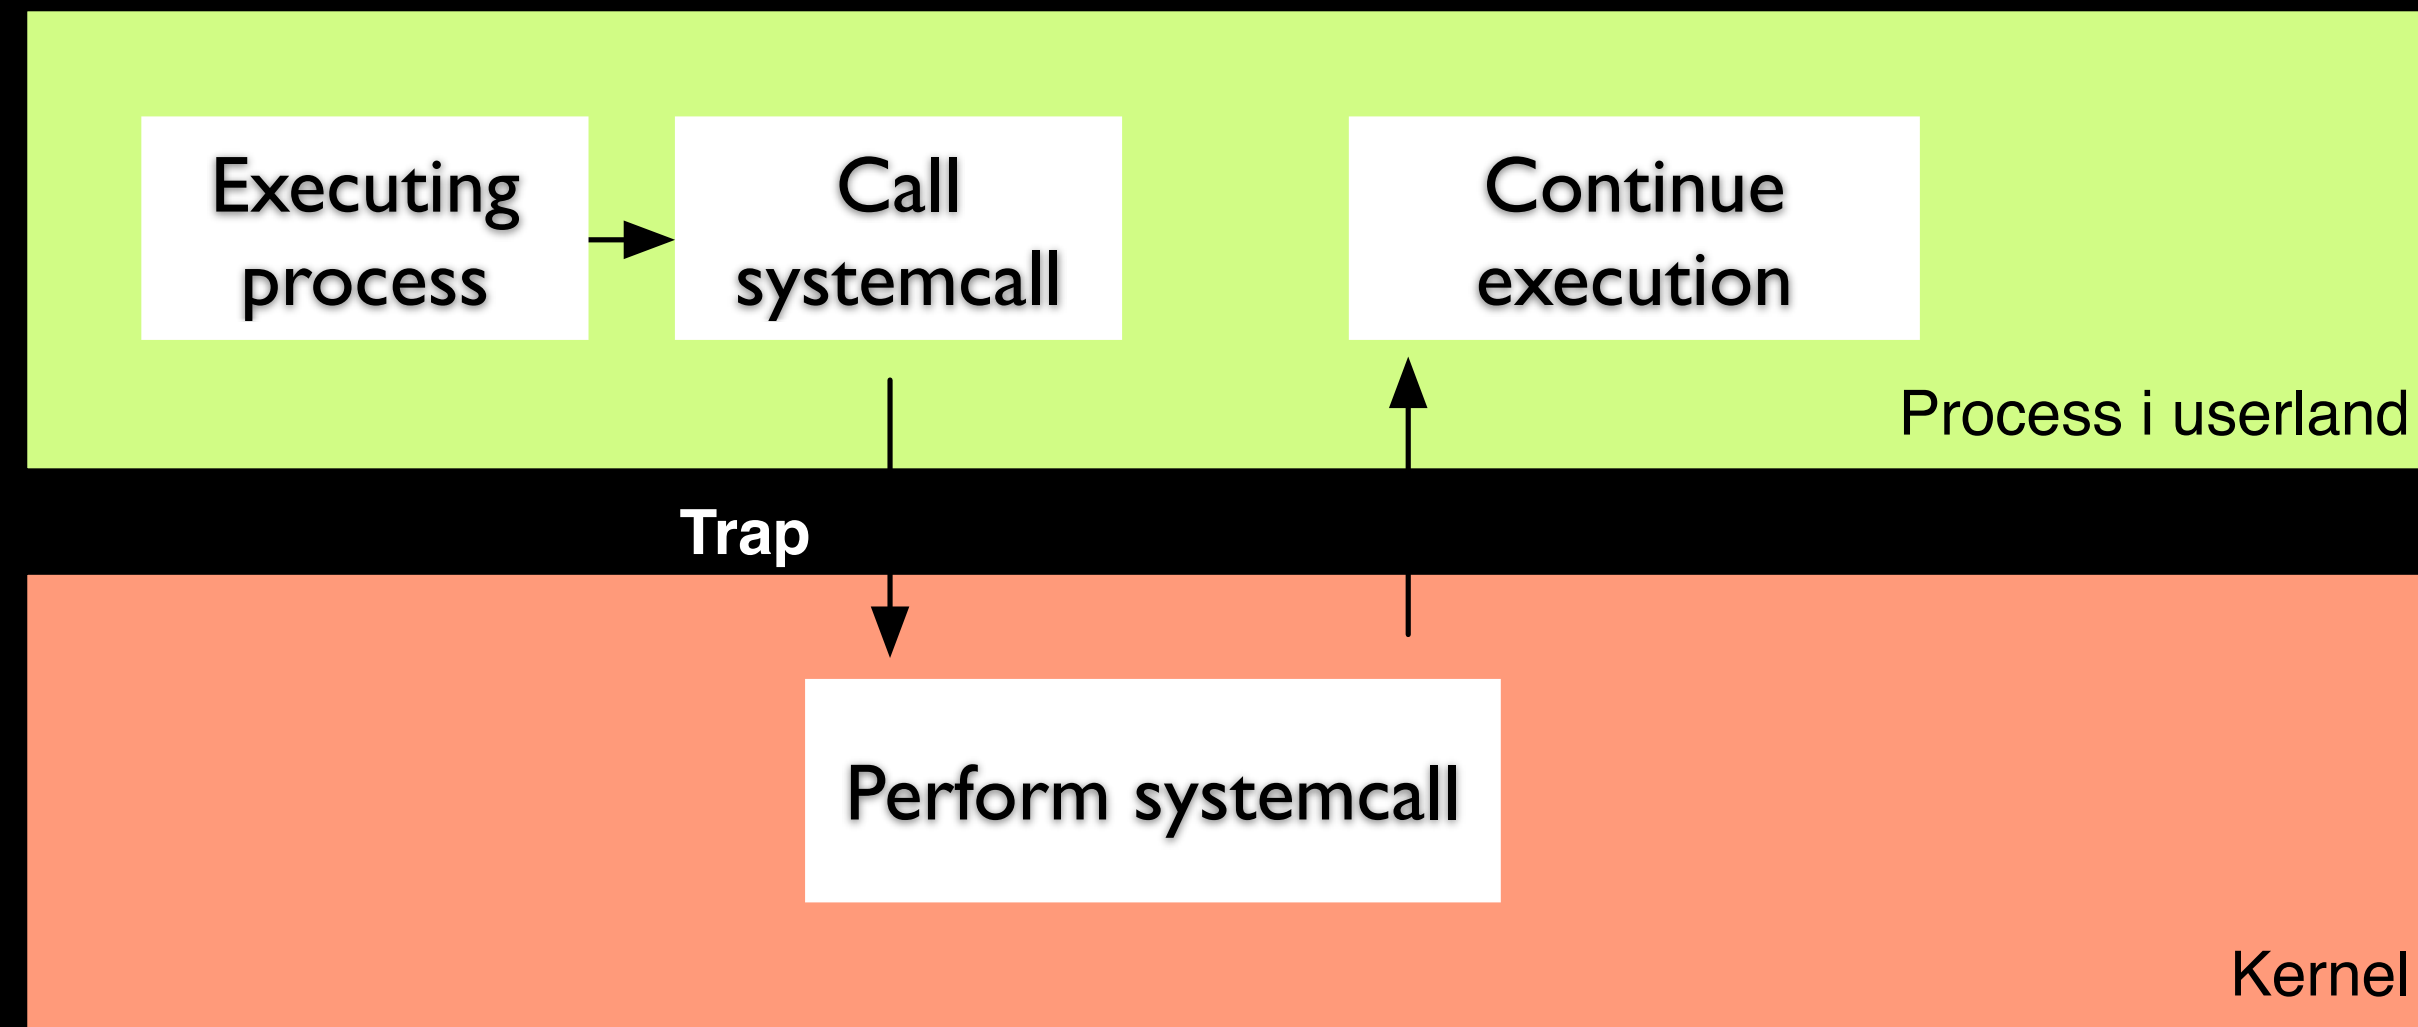


*Least
privileges*




*Highest
privileges*

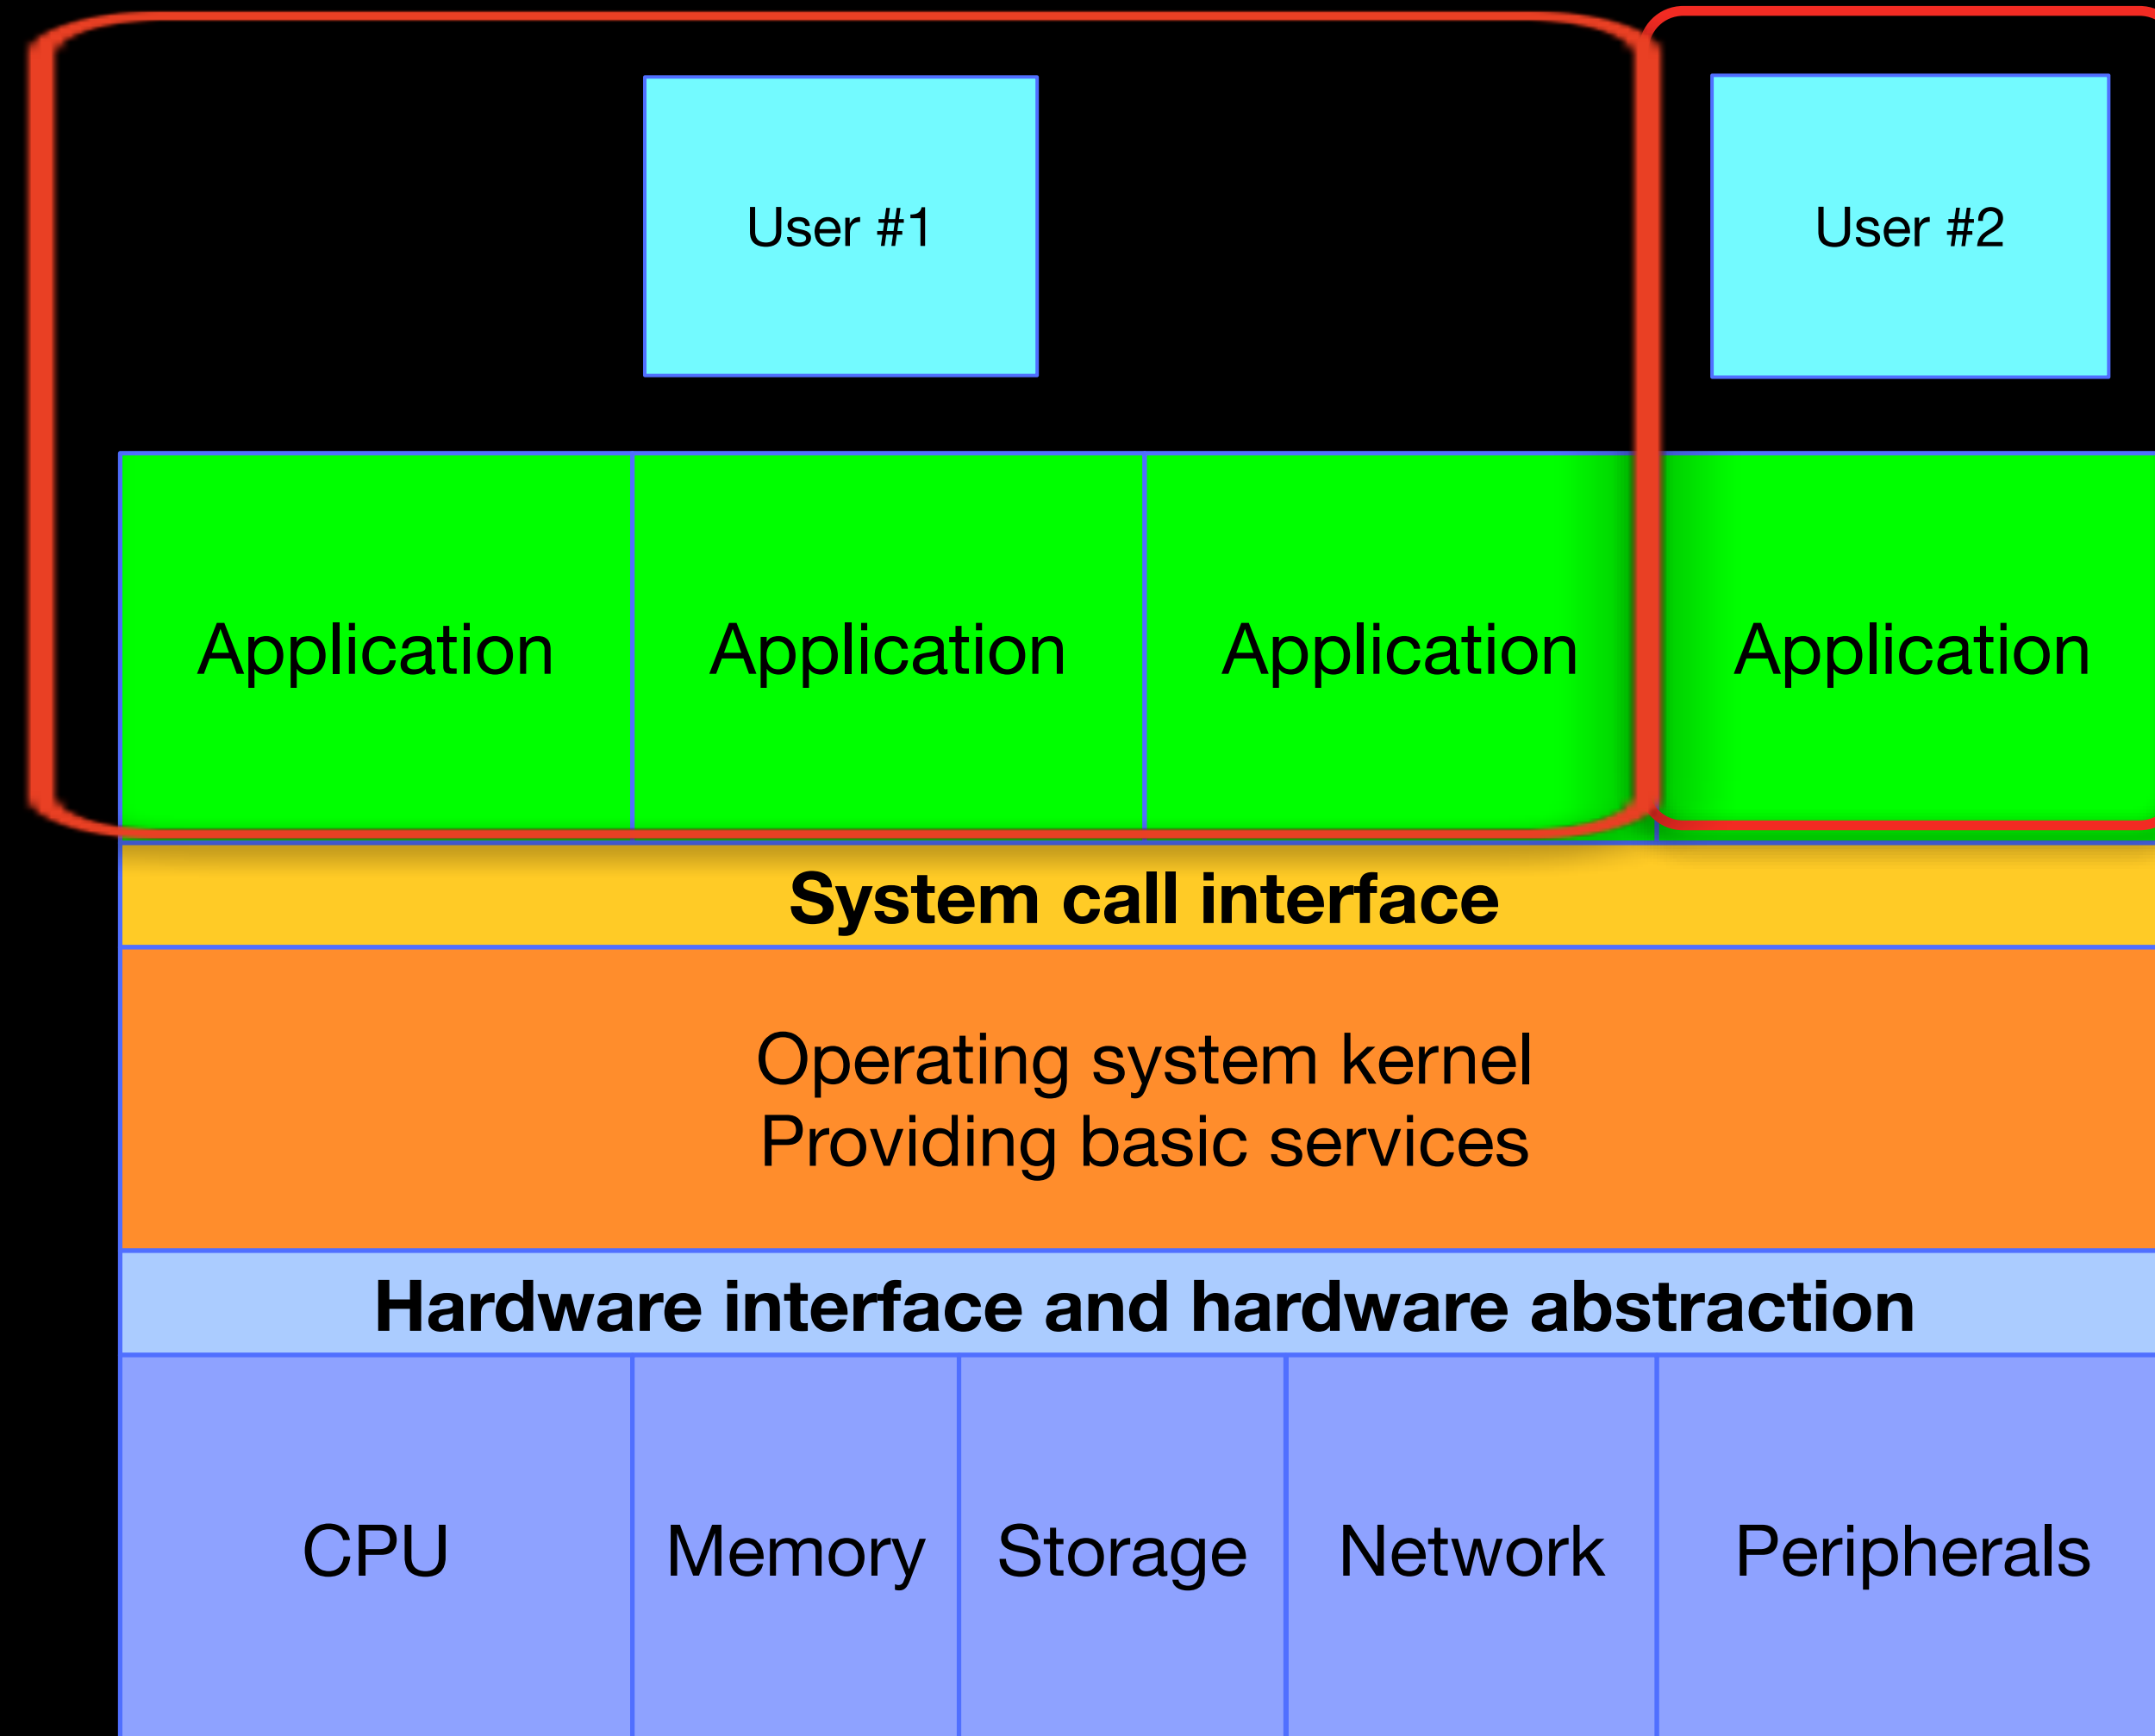
Interaction between application and OS



Overview of operating system (1/2)



Overview of operating system (2/2)

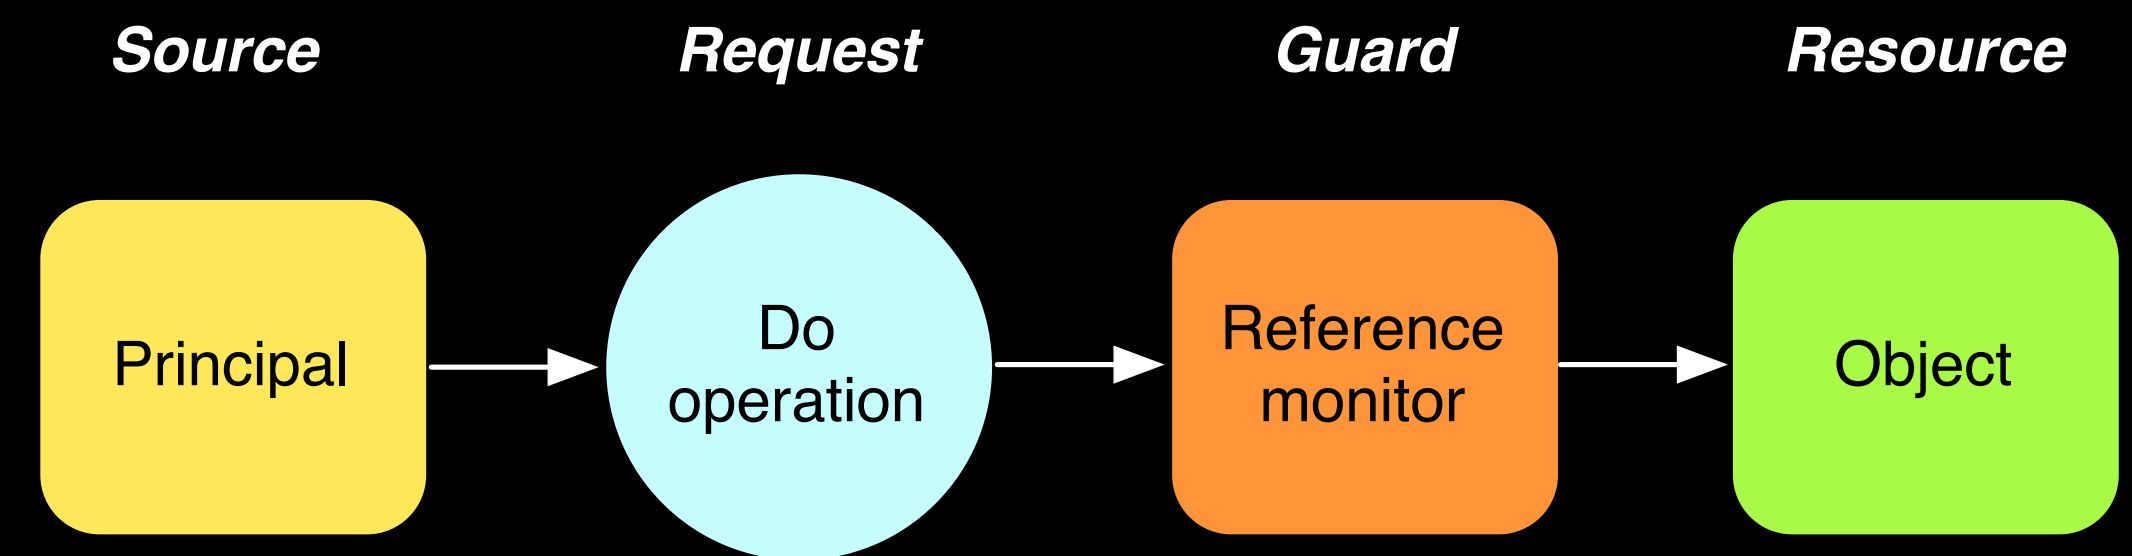


Some important concept

- A concept called *Trusted Computing Base*, or TCB
 - It contain *all things in the trusted part* of the OS necessary to enforce the security policy
 - Important that TCB is *small, clearly written, easy to see that it does not contain design or logical flaws, and that it is protected against alterations and tampering*

Some important concept

- Reference monitor




- A Reference Monitor is an abstract security component that enforces access control rules. It ensures that every access attempt to system resources complies with security policies.
- The TCB is the *entire collection of system components that enforce and maintain security, including the Reference Monitor.*

The classical *ring model*, *updated!*

Other rings

-1	Hypervisor	Allow guest OS "ring 0"
-2	System Management Mode (SMM)	APM/ACPI/TPM-support
-3	<i>Intel</i> Management Engine / <i>AMD</i> Platform Security Processor	Special software running in the Platform Controller Hub (PCH) processor



Problem with these pictures and concepts

- Layering violation

- some software might skip a layer and call an underlying layer directly and hence bypass controls

- In some scenarios attackers might *come an unexpected way*

- Attacking from *host operating system* against *guest operating systems* in a virtual machine environment

Problem with these concepts

- You have a “hidden” processor on your computer
- Its functionality has never been publicly documented
- It appears to have been customized for certain TLA government agencies
- It has unlimited access to the main processor
- It has unlimited access to all memory
- It has unlimited access to all peripherals
- It has its own MAC and IP addresses
- It runs a web server
- It is always running
- You can't turn it off
- You can't disable it
- It has had multiple known exploitable vulnerabilities
- It is the single most privileged known element of an Intel Architecture processor chipset

Memory handling

- RAM memory is a central resource that in a controlled way must be shared and handled *between operating system, applications and other components*
- Modern computer systems have hardware support for memory protection, e.g. **MMU**
- OS support is required to use the hardware supported memory protection
- Modern hardware support can enforce several security features related to isolation, non-executable memory areas, etc

File system

- A file system is often a central component in a computer system w.r.t. security and protection
- Besides the actual file content, there is meta data that is of importance
 - File owner, dates of creation/change/access, access information, security labels, etc
- Manipulation of meta data can in some cases be more serious security breach than the manipulation of the file content itself. Or a combo of both can be misleading and hide the fact that a file has been altered

Local filesystem

File system	Description	Comment
FAT	<i>No access control</i>	Classic MS-DOS
NTFS	<i>Discretionary Access Control via ACL</i>	Advanced possibilities to make controls
UFS	<i>Discretionary Access Control, writing & program execution for owner, group, “others”</i>	Simple access controls

Network file systems

File system	Description	Comment
NFSv3	<i>Hostbaserad accesskontroll, uid</i>	Trivial to circumvent
NFSv4	<i>Secure RPC, KRB5_a, KRB5_p, KRB5_i</i>	Require a Kerberos server, KDC a= authentication i=integrity = calculate MAC p= privacy = encrypt packet
SMB/CIFS	<i>KRB5_a</i>	

How do you create
security in the OS?

How do you create security in the OS?

- Follow well-known design principles
- Use well-known pattern
- Use ordinary developer best-practises
- Decide to use principles, e.g. secure-by-default
- Use programming languages that support secure practices
- Have the design and implementation evaluated and certified

Principles for secure design*

<i>Economy of mechanism</i>	Keep the design as <i>simple</i> and small as possible
<i>Fail-safe defaults</i>	Base access decisions on permission rather than exclusion
<i>Complete mediation</i>	<i>Every access to every object</i> must be checked for authority
<i>Open design</i>	The design should not be secret
<i>Separation of privilege</i>	technique in which a program is divided into parts which are limited to the specific privileges they require in order to perform a specific task
<i>Least privilege</i>	Every program and every user of the system should operate using the least set of privileges necessary to complete the job
<i>Least common mechanism</i>	Minimize the amount of mechanism common to more than one user and depended on by all users
<i>Psychological acceptability</i>	It is essential that the human interface be designed for ease of use, so that users routinely and automatically apply the protection mechanisms correctly

Comparing security in Operating systems (1/6)

- When in time was the system developed?
 - What was the state-of-the-art at that time?
 - What trends were currently in fashion?
 - What languages were available for creating the operating system?

OS	Developed	Released
Unix	1969	1971
Mach Kernel	1985	1986
Windows NT	1988	1993
Linux	1991	1991
MacOS	Late 1990's	2001
iOS	2005-07	2007
Android	2003	2008

Comparing security in Operating systems (2/6)

- In what language is an operating system developed?
 - **Unix**: assembler (1969), C (1973)
 - **Windows**: C/C++, C#&.NET
 - **Linux**: C/C++/asm, Python/bash/perl, Rust (2022-)
 - **MacOS X/iOS**: C/C++/Objective-C (1999–2010), Swift (2014-)
 - **Android**: C/C++/Java (2008–2016), Kotlin (2014-), Rust (2020-)

Comparing security in Operating systems (3/6)

*"Given enough eyeballs, all
bugs are shallow"
- Linus' Law*

- Development methodologies
 - Open Source or Closed Source?
 - What support do one use to ensure that security is *built into the product*?
 - How does one ensure that implementation is a correct representation of the design, that is a correct interpretation of the analysis?

But really, *what good is this comparison?*

Write more code = get higher salary?

Manage a 200K-SLOC project is *cooler* than a 5K-SLOC?

More code = more bugs?

Yes, more code is often more bugs

More code = more *security checks* and *advanced concepts* like
crypto, resilient failure checking built into everything?

But certainly, complexity is considered **bad** and **evil** in the context of security.

There is often a relation between complexity, size of program and bugs

Code Quality Level	Bugs per 1,000 SLOC (Defect Density)
Typical commercial software	10 - 50 bugs
Well-tested open-source software	1 - 5 bugs
Mission-critical software (NASA, avionics, medical, etc.)	0.1 - 1 bugs
High-reliability software (e.g., space systems, nuclear control)	< 0.1 bugs

There is often a relation between complexity, size of program and bugs

Comparing security in Operating systems (6/6)

- What can one gain by having formal certification of operating systems, subsystems or application
- Trusted Computer System Evaluation Criteria (TCSEC), Common Criteria (CC, ISO/IEC 15408), etc
- More a theoretical excersice than of any real value?

Background and basics

Part 2: bugs and vulnerabilities

Intro - *just the basic facts*

- All software *is prone to bugs*
- Some bugs will have an impact that can have **security implications**
 - data leaks,
 - destruction of data,
 - local privilege escalations (**LPE**),
 - execution of remotely uploaded malicious code (**RCE**),
 - etc

Intro - *just the basic facts*

- Some bugs help to circumvent security mechanisms
- Some **security designs** are flawed, or build on *flawed assumptions*



Operating system security

- Security problems in the operating system can affect the integrity of the system itself
 - Someone else can control the system to their own liking - **pwnd!**
 - Bugs in OS kernel can affect system integrity
- Security problems with the operating system can, as a result, affect the security in **applications** and **subsystems** (databases, middle ware, etc)

General examples of threats and attacks

Wrong file permissions

Sensitive plaintext in RAM

fork bombs

SYN flood

Confidentiality

Crashdumps with credentials or crypto keys

Bypassed security checks

malformed network packets

Availability

unintentional filling of disk space

intentional filling of disk space

Manipulated system configuration

System integrity

Manipulated application

program binaries

Manipulated system binaries

Manipulated user files

Data integrity

Zapped system logs

Manipulated database content

Some concepts and terms

Memory
corruption bugs

Stack smashing Stack overflow

Heap overflow

Time related bugs

Race conditions

*Time-of-Check to Time-of-Use
(TOCTTOU)*

Information
disclosure bugs

File inclusion

File/object permissions

Directory traversal

Some concepts and terms

Vulnerability

Exploit

0day exploit
1 day exploit/Nday
Foreverday exploit

vulnerability
unknown to vendor &
no patch *available*

vulnerability
disclosed but *not* widely
patched

unpatchable
or *unsupported*
systems

Intro - *the basics*

- Some bugs are undiscovered for some time, they lay latent
- Once discovered, they can be abused, if it is an **security vulnerability**, that can be **exploited**
- A discovered security bug, is sometime called a **0day**, until it is mitigated

Intro - *the basics*

- Nowadays bugs and vulnerabilities tend to get **names** (*heartbleed, ghost, shellshock, etc*) and **logos**
- Used by security companies for marketing their knowledge and brand




Some concepts and principles

- **Attack vector** - Different paths to reach an vulnerability. One path might be closed by a vendor patch, but another might still be there, if the root cause is not identified and fixed.
- **Attack surface** - exposed parts that an attacker can reach, i.e. all the different attack vectors
- **Reverse engineering (RE)** - To re-create the original design by observing the final result, in computer science - to re-create some source code by examining a binary.

Year	Overflow	Memory Corruption	Sql Injection	XSS	Directory Traversal	File Inclusion	CSRF	XXE	SSRF	Open Redirect	Input Validation
2015	1037	1104	221	776	152	6	249	50	8	46	379
2016	1180	1173	97	497	99	12	87	41	16	33	519
2017	2478	1542	505	1500	282	155	334	109	57	97	936
2018	2084	1731	503	2042	569	112	479	188	118	85	1248
2019	1205	2030	544	2387	488	126	560	137	103	121	908
2020	1218	1879	465	2201	436	110	415	119	131	100	815
2021	1664	2530	742	2724	548	91	520	126	192	133	678
2022	1863	3369	1788	3404	729	96	769	126	230	146	779
2023	1673	2298	2121	5136	769	116	1392	127	243	181	627
2024	1781	2534	2650	7455	945	257	1435	112	378	121	133
2025	237	207	416	1683	134	46	426	16	58	13	0
Total	16420	20397	10052	29805	5151	1127	6666	1151	1534	1076	7022

Vulnerabilities by type



Vulnerabilities by impact types

Year	Code Execution	Bypass	Privilege Escalation	Denial of Service	Information Leak
2014	1041	165	186	1597	356
2015	1430	177	255	1793	602
2016	1239	469	608	2050	704
2017	1870	857	1027	3372	1394
2018	1728	666	850	2207	1418
2019	1534	670	916	1699	1326
2020	1691	816	1386	1677	1094
2021	2087	806	1121	2297	926
2022	2067	943	1527	2437	1144
2023	2578	1059	1525	2557	1545
2024	276	139	150	333	119
Total	17541	6767	9551	22019	10628

Intro - *the basics*


- Many vulnerabilities also gets "formal name", i.e. **CVE***, and a scoring **CVSS****
 - e.g. CVE-2024-21762 (A out-of-bounds write in Fortinet **FortiOS**) with **CVSS score of 9.8**
- A CVE is assigned by a CNA, a *CVE numbering authority*
- All issued CVE is stored in central database
- Not all vulnerabilities gets an CVE
- Not all issued CVE numbers ends up being used in public vulnerability info

* "Common Vulnerabilities and Exposures;" <https://cve.mitre.org/>

** <https://www.first.org/cvss/specification-document>**

Intro - *the basics*

Number of CVEs by year



Intro - *the basics*

CVSS Scores Between 2024-01-01 and 2024-12-31

CVSS score distribution for CVEs published between 2024-01-01 and 2024-12-31

Period

2024-01-01



2024-12-31



☐ Group By Year

Submit

CVSS Score Range

0-1



1-2



2-3



3-4



4-5



5-6



6-7



7-8



8-9



9+



Total

40296

Weighted Average CVSS Score: 7.2

Vulnerabilities

1644

15

227

589

2983

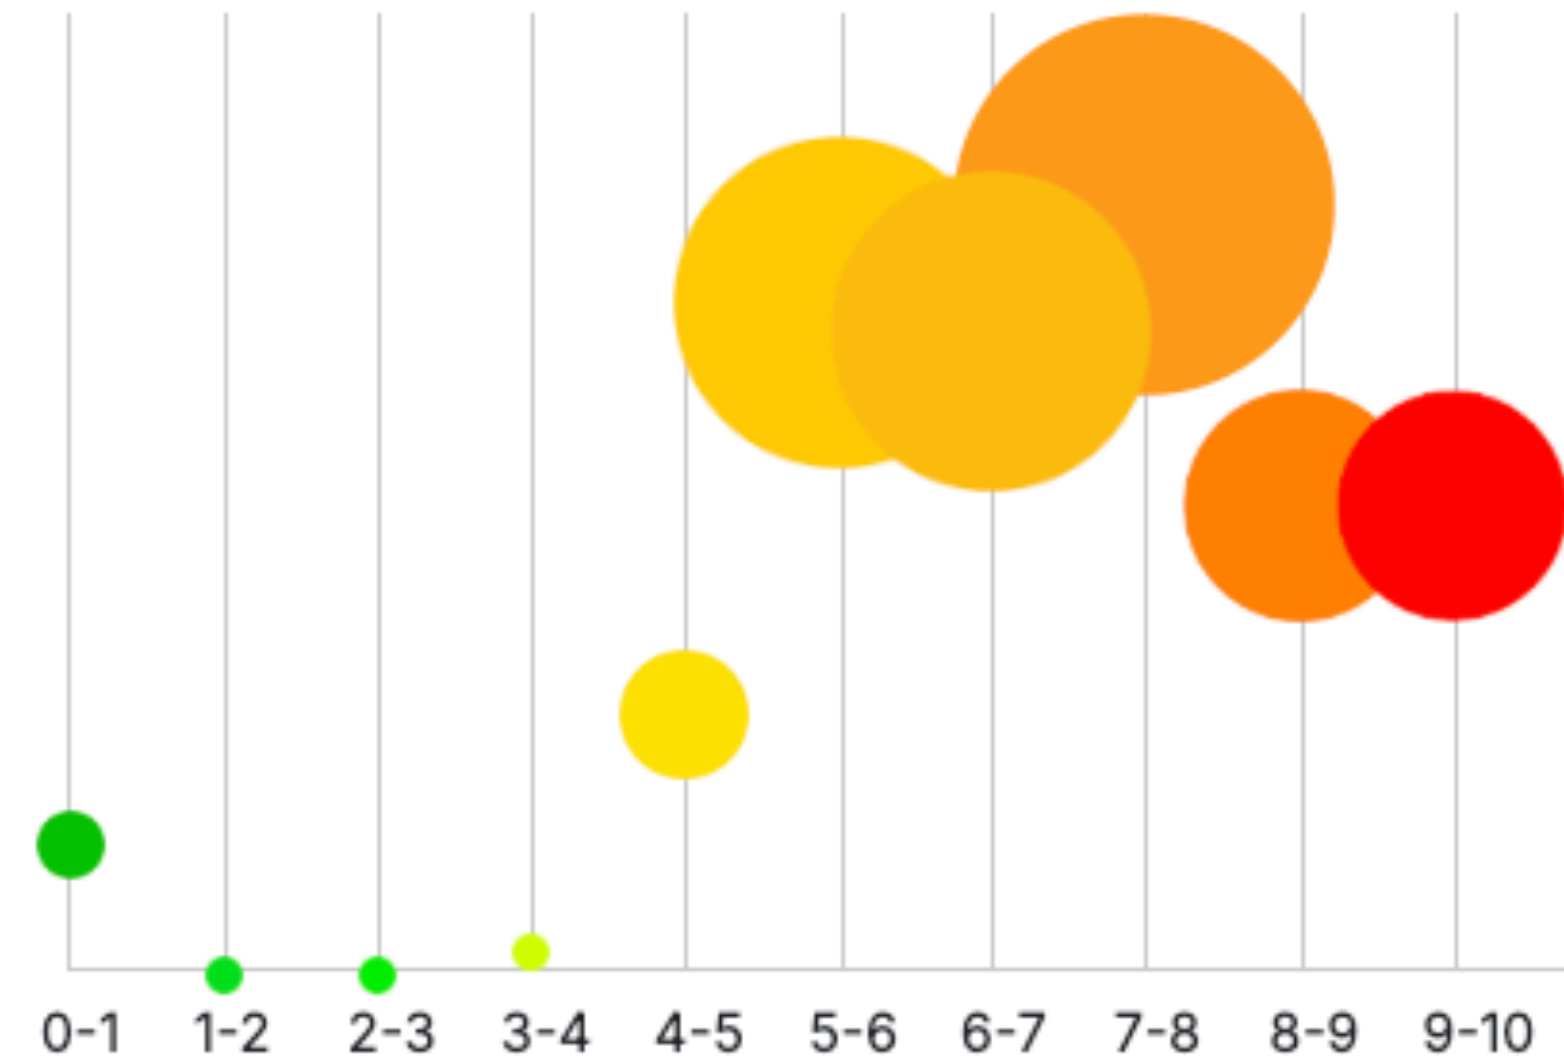
7735

7437

8890

5386

5390



Intro - *the basics*

CVSS Scores Between 2023-01-01 and 2023-12-31

Period

2023-01-01




2023-12-31



☐ Group By Year

Submit

CVSS Score Range




Total

29065


Weighted Average CVSS Score: 7.7

Vulnerabilities



Examples of vulnerabilities and attacks

Where do attacks occur?



The classical *ring model*, *updated!*

CVE-2022-40261
HIGH


InformationCPEsPlugins

Description

An attacker can exploit this vulnerability to elevate privileges from ring 0 to ring -2, execute arbitrary code in System Management Mode – an environment more privileged than operating system (OS) and completely isolated from it. Running arbitrary code in SMM additionally bypasses SMM-based SPI flash protections against modifications, which can help an attacker to install a firmware backdoor/implant into BIOS. Such a malicious firmware code in BIOS could persist across operating system re-installs. Additionally, this vulnerability potentially could be used by malicious actors to bypass security mechanisms provided by UEFI firmware (for example, Secure Boot and some types of memory isolation for hypervisors). This issue affects:
Module name: OverClockSmiHandler SHA256:
a204699576e1a48ce915d9d9423380c8e4c197003baf9d17e6504f0265f3039c Module GUID: 4698C2BD-A903-410E-AD1F-5EEF3A1AE422

Details

Source: [Mitre](#), [NVD](#)
Published: 2022-09-20
CVSS v3
Base Score: 8.2
Vector:
CVSS:3.0/AV:L/AC:L/PR:H/UI:N/S:C/C:H/I:H/A:H
Severity: High



Example of attacks

Attack method	Description	Synonyms and variants
Buffer overflow	<p>Attacks that allow an attacker to <u>deterministically alter the execution flow of a program by submitting crafted input to an application</u>. Executable code is written outside the boundaries of a memory buffer originally used for storing data. The executable parts is somehow made to execute, e.g. by manipulate return adress to be used when a function call is finished.</p> <p>Real world examples: OpenBSD IPv6 mbuf's* remote kernel buffer overflow[1], windows kernel pool</p>	<p>Synonyms: memory corruption attack, Buffer overrun, Stack smashing,</p> <p>Variants: Heap smashing, format string bugs,</p>

[1] <http://www.coresecurity.com/content/open-bsd-advisorie>

* An mbuf is a basic unit of memory management in the kernel IPC subsystem

Example of attacks

Attack method	Description	Examples
Backward compability	<p>Attacks that allow an attacker to use</p> <ul style="list-style-type: none">• <i>an older version of a service, or</i>• <i>an old protocol, or</i>• <i>an older mode, or</i>• <i>call legacy code</i> <p>Sometime triggered by downgrade attack, a negotiation to use older variant</p>	<p>Remote Desktop NTLMv1 XML encryption SSLv2, SSLv3, incl POODLE, FREAK Encryption modes Kerberos v4 in v5</p>

Most common attacks?

OWASP top-10 list

2021

A01:2021-Broken Access Control

A02:2021-Cryptographic Failures

A03:2021-Injection

A04:2021-Insecure Design

A05:2021-Security Misconfiguration

A06:2021-Vulnerable and Outdated Components

A07:2021-Identification and Authentication Failures

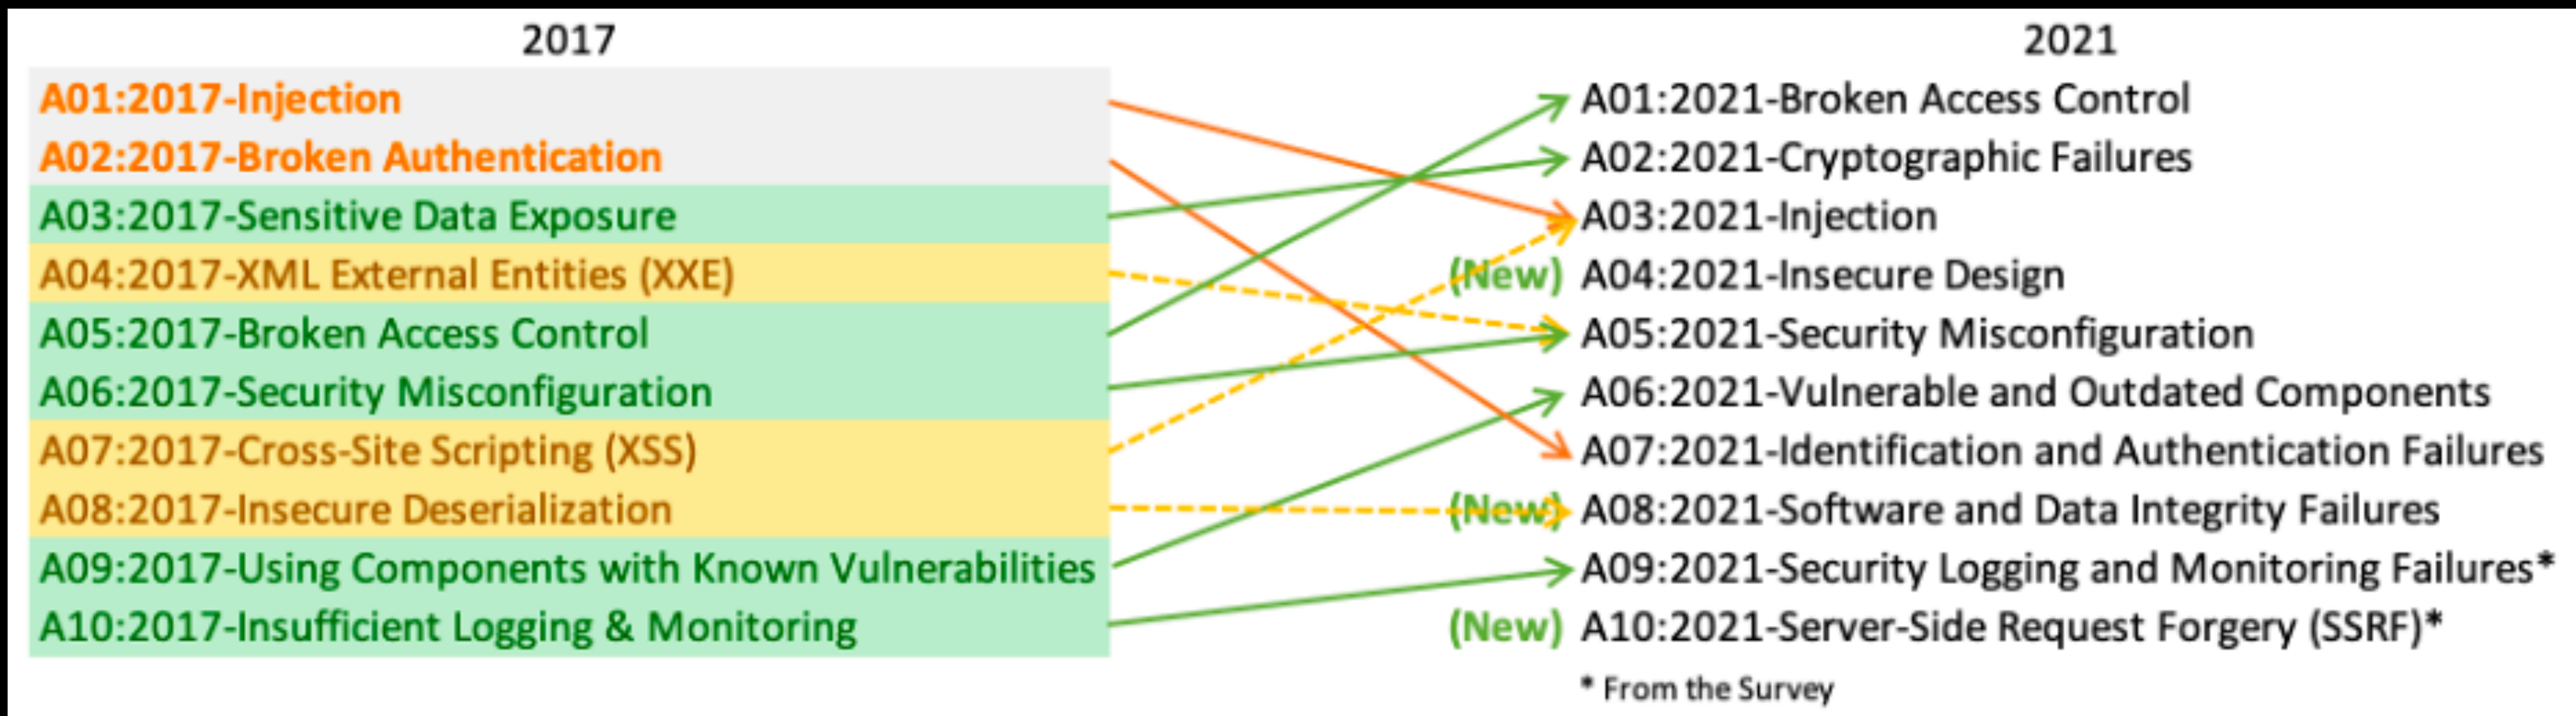
A08:2021-Software and Data Integrity Failures

A09:2021-Security Logging and Monitoring Failures*

A10:2021-Server-Side Request Forgery (SSRF)*

* From the Survey

Most common attacks?



OWASP top-10 list

Reconnaissance 10 techniques	Resource Development 7 techniques	Initial Access 9 techniques	Execution 12 techniques	Persistence 19 techniques	Privilege Escalation 13 techniques	Defense Evasion 40 techniques
<ul style="list-style-type: none"> Active Scanning (2) Gather Victim Host Information (4) Gather Victim Identity Information (3) Gather Victim Network Information (6) Gather Victim Org Information (4) Phishing for Information (3) Search Closed Sources (2) Search Open Technical Databases (5) Search Open Websites/Domains (2) Search Victim-Owned Websites 	<ul style="list-style-type: none"> Acquire Infrastructure (6) Compromise Accounts (2) Compromise Infrastructure (6) Develop Capabilities (4) Establish Accounts (2) Obtain Capabilities (6) Stage Capabilities (5) 	<ul style="list-style-type: none"> Drive-by Compromise Exploit Public-Facing Application External Remote Services Hardware Additions Phishing (3) Replication Through Removable Media Supply Chain Compromise (3) Trusted Relationship Valid Accounts (4) 	<ul style="list-style-type: none"> Command and Scripting Interpreter (8) Container Administration Command Deploy Container Exploitation for Client Execution Inter-Process Communication (2) Native API Scheduled Task/Job (5) Shared Modules Software Deployment Tools System Services (2) User Execution (3) Windows Management Instrumentation 	<ul style="list-style-type: none"> Account Manipulation (4) BITS Jobs Boot or Logon Autostart Execution (15) Boot or Logon Initialization Scripts (5) Browser Extensions Compromise Client Software Binary Create Account (3) Create or Modify System Process (4) Event Triggered Execution (15) External Remote Services Hijack Execution Flow (11) Implant Internal Image Modify Authentication Process (4) Office Application Startup (6) Pre-OS Boot (5) Scheduled Task/Job (5) Server Software Component (4) Traffic Signaling (1) Valid Accounts (4) 	<ul style="list-style-type: none"> Abuse Elevation Control Mechanism (4) Access Token Manipulation (5) Boot or Logon Autostart Execution (15) Boot or Logon Initialization Scripts (5) Create or Modify System Process (4) Domain Policy Modification (2) Escape to Host Event Triggered Execution (15) Exploitation for Privilege Escalation Hijack Execution Flow (11) Process Injection (11) Scheduled Task/Job (6) Valid Accounts (4) 	<ul style="list-style-type: none"> Abuse Elevation Control Mechanism (4) Access Token Manipulation (5) BITS Jobs Build Image on Host Deobfuscate/Decode Files or Information Deploy Container Direct Volume Access Domain Policy Modification (2) Execution Guardrails (1) Exploitation for Defense Evasion File and Directory Permissions Modification (2) Hide Artifacts (9) Hijack Execution Flow (11) Impair Defenses (9) Indicator Removal on Host (5) Indirect Command Execution Masquerading (2) Modify Authentication Process (4) Modify Cloud Compute Infrastructure (4) Modify Registry Modify System Image (2) Network Boundary Bridging (1) Obfuscated Files or Information (6) Pre-OS Boot (5) Process Injection (11) Reflective Code Loading Rogue Domain Controller Rootkit Signed Binary Proxy Execution (13) Signed Script Proxy Execution (1) Subvert Trust Controls (6) Template Injection Traffic Signaling (1) Trusted Developer Utilities Proxy Execution (1) Unused/Unsupported Cloud Regions Use Alternate Authentication Material (4) Valid Accounts (4) Virtualization/Sandbox Evasion (3) Weaken Encryption (2) XSL Script Processing

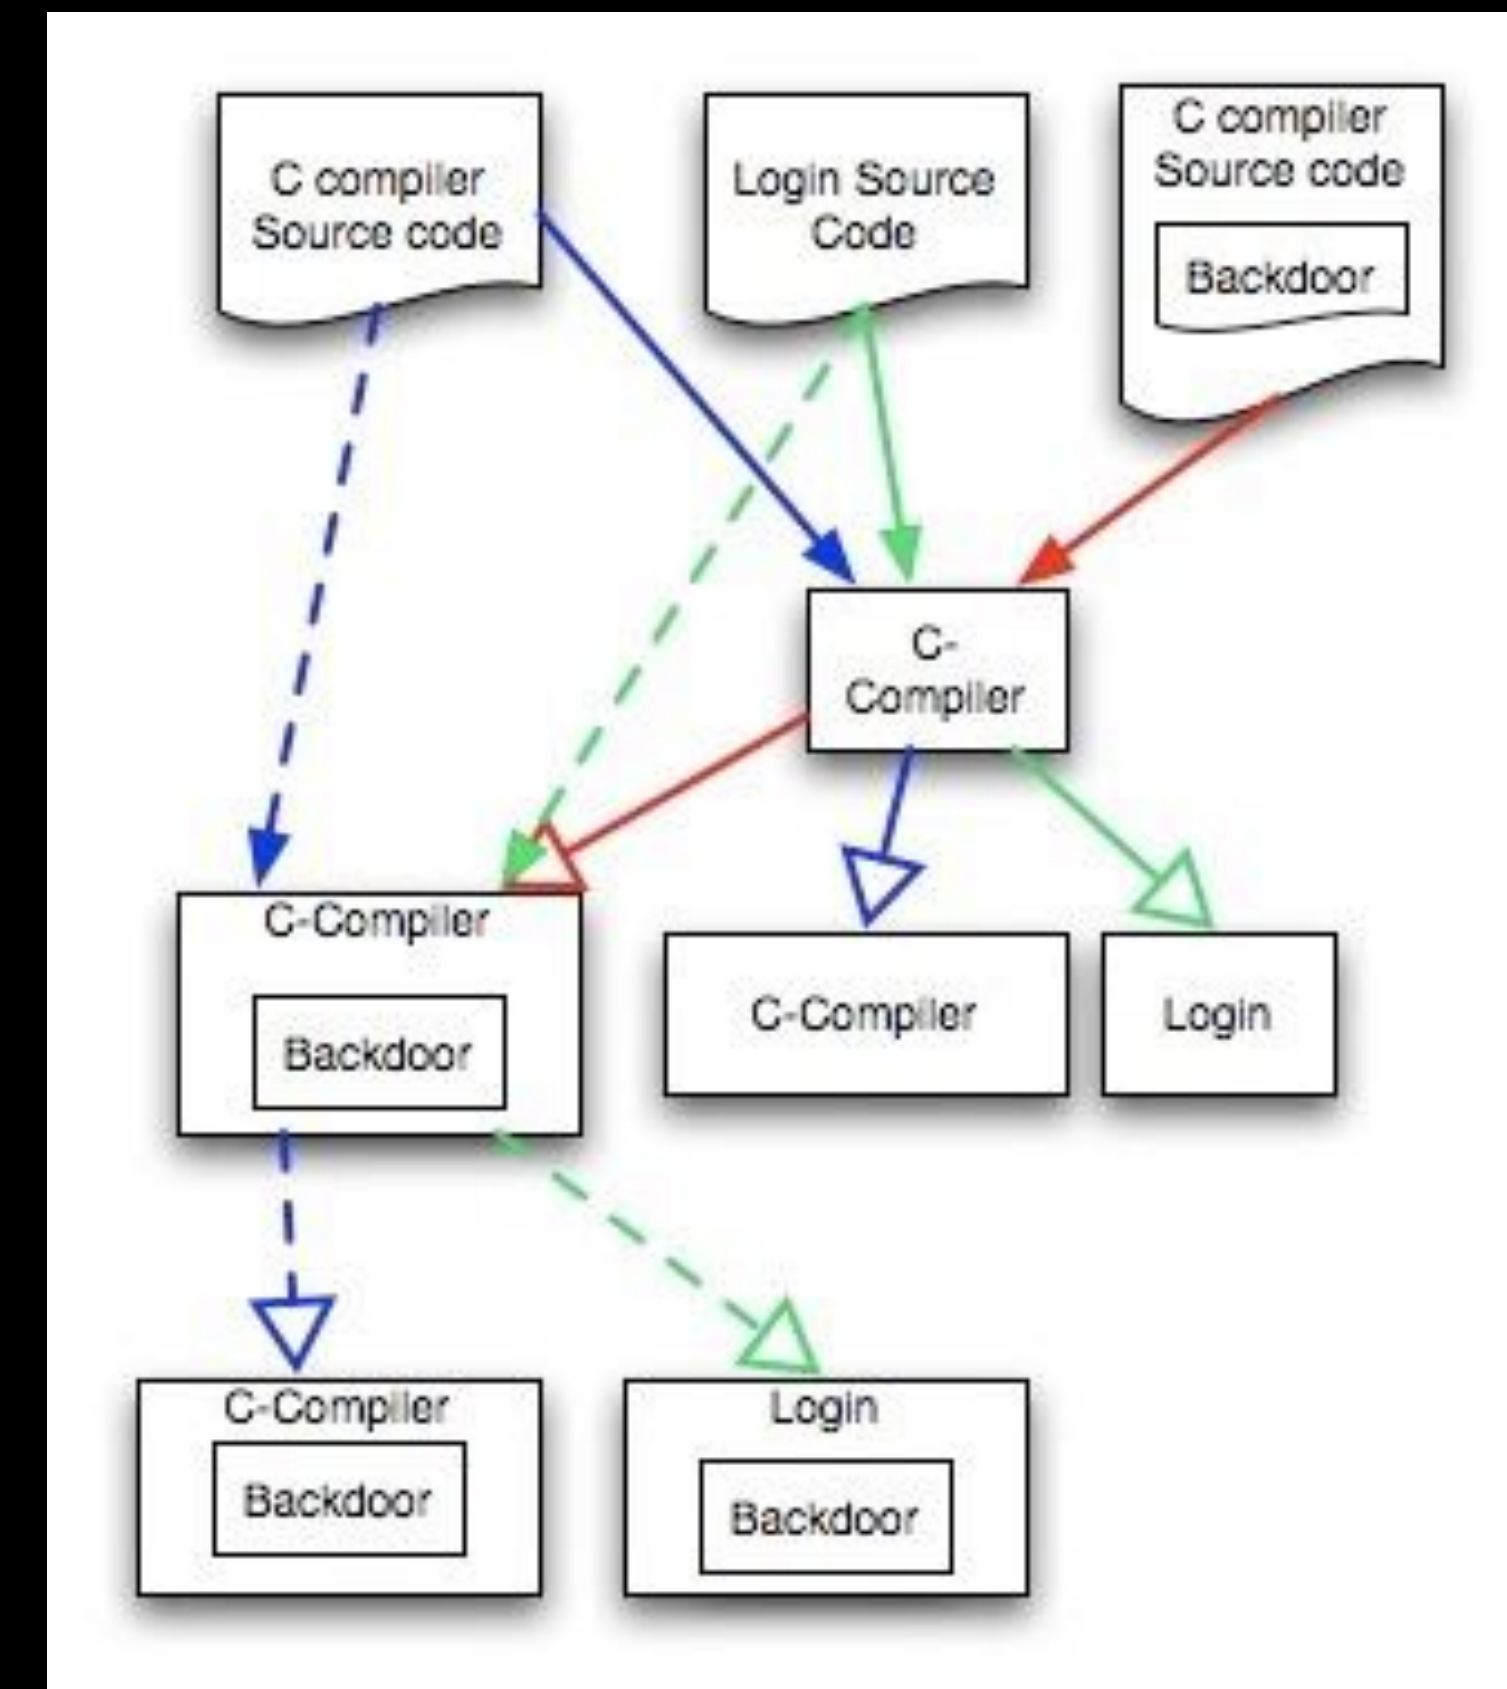
MITRE ATT&CK framework

Credential Access 15 techniques	Discovery 29 techniques	Lateral Movement 9 techniques	Collection 17 techniques	Command and Control 16 techniques	Exfiltration 9 techniques	Impact 13 techniques
Adversary-in-the-Middle (2)	Account Discovery (4)	Exploitation of Remote Services	Adversary-in-the-Middle (2)	Application Layer Protocol (4)	Automated Exfiltration (1)	Account Access Removal
Brute Force (3)	Application Window Discovery	Internal Spearphishing	Archive Collected Data (3)	Communication Through Removable Media	Data Transfer Size Limits	Data Destruction
Credentials from Password Stores (5)	Browser Bookmark Discovery	Lateral Tool Transfer	Audio Capture	Data Encoding (2)	Exfiltration Over Alternative Protocol (3)	Data Encrypted for Impact
Exploitation for Credential Access	Cloud Infrastructure Discovery	Remote Service Session Hijacking (2)	Automated Collection	Data Obfuscation (3)	Exfiltration Over C2 Channel	Data Manipulation (3)
Forced Authentication	Cloud Service Dashboard	Remote Services (6)	Browser Session Hijacking	Dynamic Resolution (3)	Exfiltration Over Other Network Medium (1)	Defacement (2)
Forge Web Credentials (2)	Cloud Service Discovery	Replication Through Removable Media	Clipboard Data	Encrypted Channel (2)	Exfiltration Over Physical Medium (1)	Disk Wipe (2)
Input Capture (4)	Cloud Storage Object Discovery	Software Deployment Tools	Data from Cloud Storage Object	Fallback Channels	Exfiltration Over Web Service (2)	Endpoint Denial of Service (4)
Modify Authentication Process (4)	Container and Resource Discovery	Taint Shared Content	Data from Configuration Repository (2)	Ingress Tool Transfer	Scheduled Transfer	Firmware Corruption
Network Sniffing	Domain Trust Discovery	Use Alternate Authentication Material (4)	Data from Information Repositories (3)	Multi-Stage Channels	Transfer Data to Cloud Account	Inhibit System Recovery
OS Credential Dumping (8)	File and Directory Discovery		Data from Local System	Non-Application Layer Protocol		Network Denial of Service (2)
Steal Application Access Token	Group Policy Discovery		Data from Network Shared Drive	Non-Standard Port		Resource Hijacking
Steal or Forge Kerberos Tickets (4)	Network Service Scanning		Data from Removable Media	Protocol Tunneling		Service Stop
Steal Web Session Cookie	Network Share Discovery		Data Staged (2)	Proxy (4)		System Shutdown/Reboot
Two-Factor Authentication Interception	Network Sniffing		Email Collection (3)	Remote Access Software		
Unsecured Credentials (7)	Password Policy Discovery		Input Capture (4)	Traffic Signaling (1)		
	Peripheral Device Discovery		Screen Capture	Web Service (3)		
	Permission Groups Discovery (3)		Video Capture			
	Process Discovery					
	Query Registry					
	Remote System Discovery					
	Software Discovery (1)					
	System Information Discovery					
	System Location Discovery (1)					
	System Network Configuration Discovery (1)					
	System Network Connections Discovery					
	System Owner/User Discovery					
	System Service Discovery					
	System Time Discovery					
	Virtualization/Sandbox Evasion (3)					

MITRE ATT&CK framework

A classic attack

- Ken Thompson's trojanized c compiler
 - Modify the source code to the compiler to recognize if it recompile itself or the login program - insert backdoor in login
 - recompile compiler
 - remove source code changes and recompile the compiler
 - recompile the login program with the modified compiler
- No visible signs for humans or tools to see the backdoor in source code. Calls for binary inspection or decompilation.



Attacks and counter measures

- *Chaining of attacks* - combining a number of exploits to achieve goal
 - finding and abusing a number of different vulnerabilities might allow an attacker to achieve goals not possible with just one potent exploit
 - *Code execution in gadgets (ROP) + sandbox escape + elevation of privileges + execution of privileged code*

Example of attacks

Remember that there is a number of ways that all OS security controls can be bypassed,
especially if the operating system is not running
- a very good side-channel attack ;-)

Examples of different protection solutions

General example of control principles

Security controls	Description	Example	Where?
Encryption	Protection against <i>eavesdropping</i> or unauthorized access	network traffic, file content, disk partitions, memory pages, swap files/ page area	OpenSSL, IPsec, SSH, OS kernel
Electronic signatures	Protection against <i>changes</i> or unauthorized modifications by third parties,	network traffic, file content, disk partitions	OpenSSL, IPsec, SSH, OS kernel
Cryptographically strong hash values	Protection against unauthorized changes, detect errors or changes	Saved passwords, file content,	Password file, user database, checksums on files

General example of control principles

Security controls	Description	Example	Where?
Random numbers	<i>Make a resource non-deterministic</i>	File names, process ID's, port numbers, session keys, session id's, transaction numbers, DNS query ID's, execution time & timing	getrandom() /dev/urandom
Constant numbers	<i>Make a resource non-deterministic</i>	execution time, timing of events	Crypto code to prevent side channel attacks

General example of control principles

Security controls	Description	Example
Compiler generated airbag - canary	<i>Make sure buffer overflows dont get undetected</i>	ProPolice, VisualStudio /GS
ASLR	<i>Randomize addresses used by applications. Make sure its hard to write code that knows of addresses. Where did that lib go?</i>	Android >4.0, iOS > 4.3, Windows >Vista, OpenBSD/NetBSD, Linux >2.6.12, MacOSX >10.5, Solaris >11.1, etc
KASLR	<i>Randomize addresses used by kernel</i>	Windows Vista, NetBSD, Linux >3.14, MacOSX 10.8, Android 11, etc

General example of control principles

Security controls	Description	Example
DEP, NX, W^X	<i>Make sure memory is not executable</i>	IE on Windows Vista, Android >2.3, FreeBSD > 5.3, OpenBSD, Linux >2.6.8, MacOSX >10.5, etc
MTE	<i>Memory Tagging Extension</i>	Using ARM architecture feature to better protect against memory safety violations

General example of control principles

Security controls	Description	Example
Secure boot chain / Verified boot	<i>Make system startup sequence is secure</i>	Make sure that each step of boot is cryptographically signed to ensure code integrity, e.g. BIOS vs UEFI
Secure pairing	<i>Make sure to connect to peripherals and resources in a secure way</i>	Using bluetooth to connect to headset,

General example of control principles

Security controls	Description	Example
Scrubbing, zeroing	<i>Make sure that old data areas are cleaned before usage or returned to system</i>	memory, file systems, VM system
Logs, audit trails	<i>Traces, error messages and dumps from systems and applications</i>	Windows Eventlog, Syslog, audit, BSM

Attacks and counter measures



Hijacking JIT compilers

ROP attacks

Address Space Layout
Randomization (ASLR)

No-executable
(NX, W^X) stacks

Data Execution
Prevention (DEP)

More advanced buffer
overflows, defeating canary

Stack canaries

Buffer overflow/memory
corruption attacks

Note - several of these counter measures does not work for protection **within** the kernel

Virtualization and isolation

sandboxes, containers, hypervisors, etc

Sandboxing

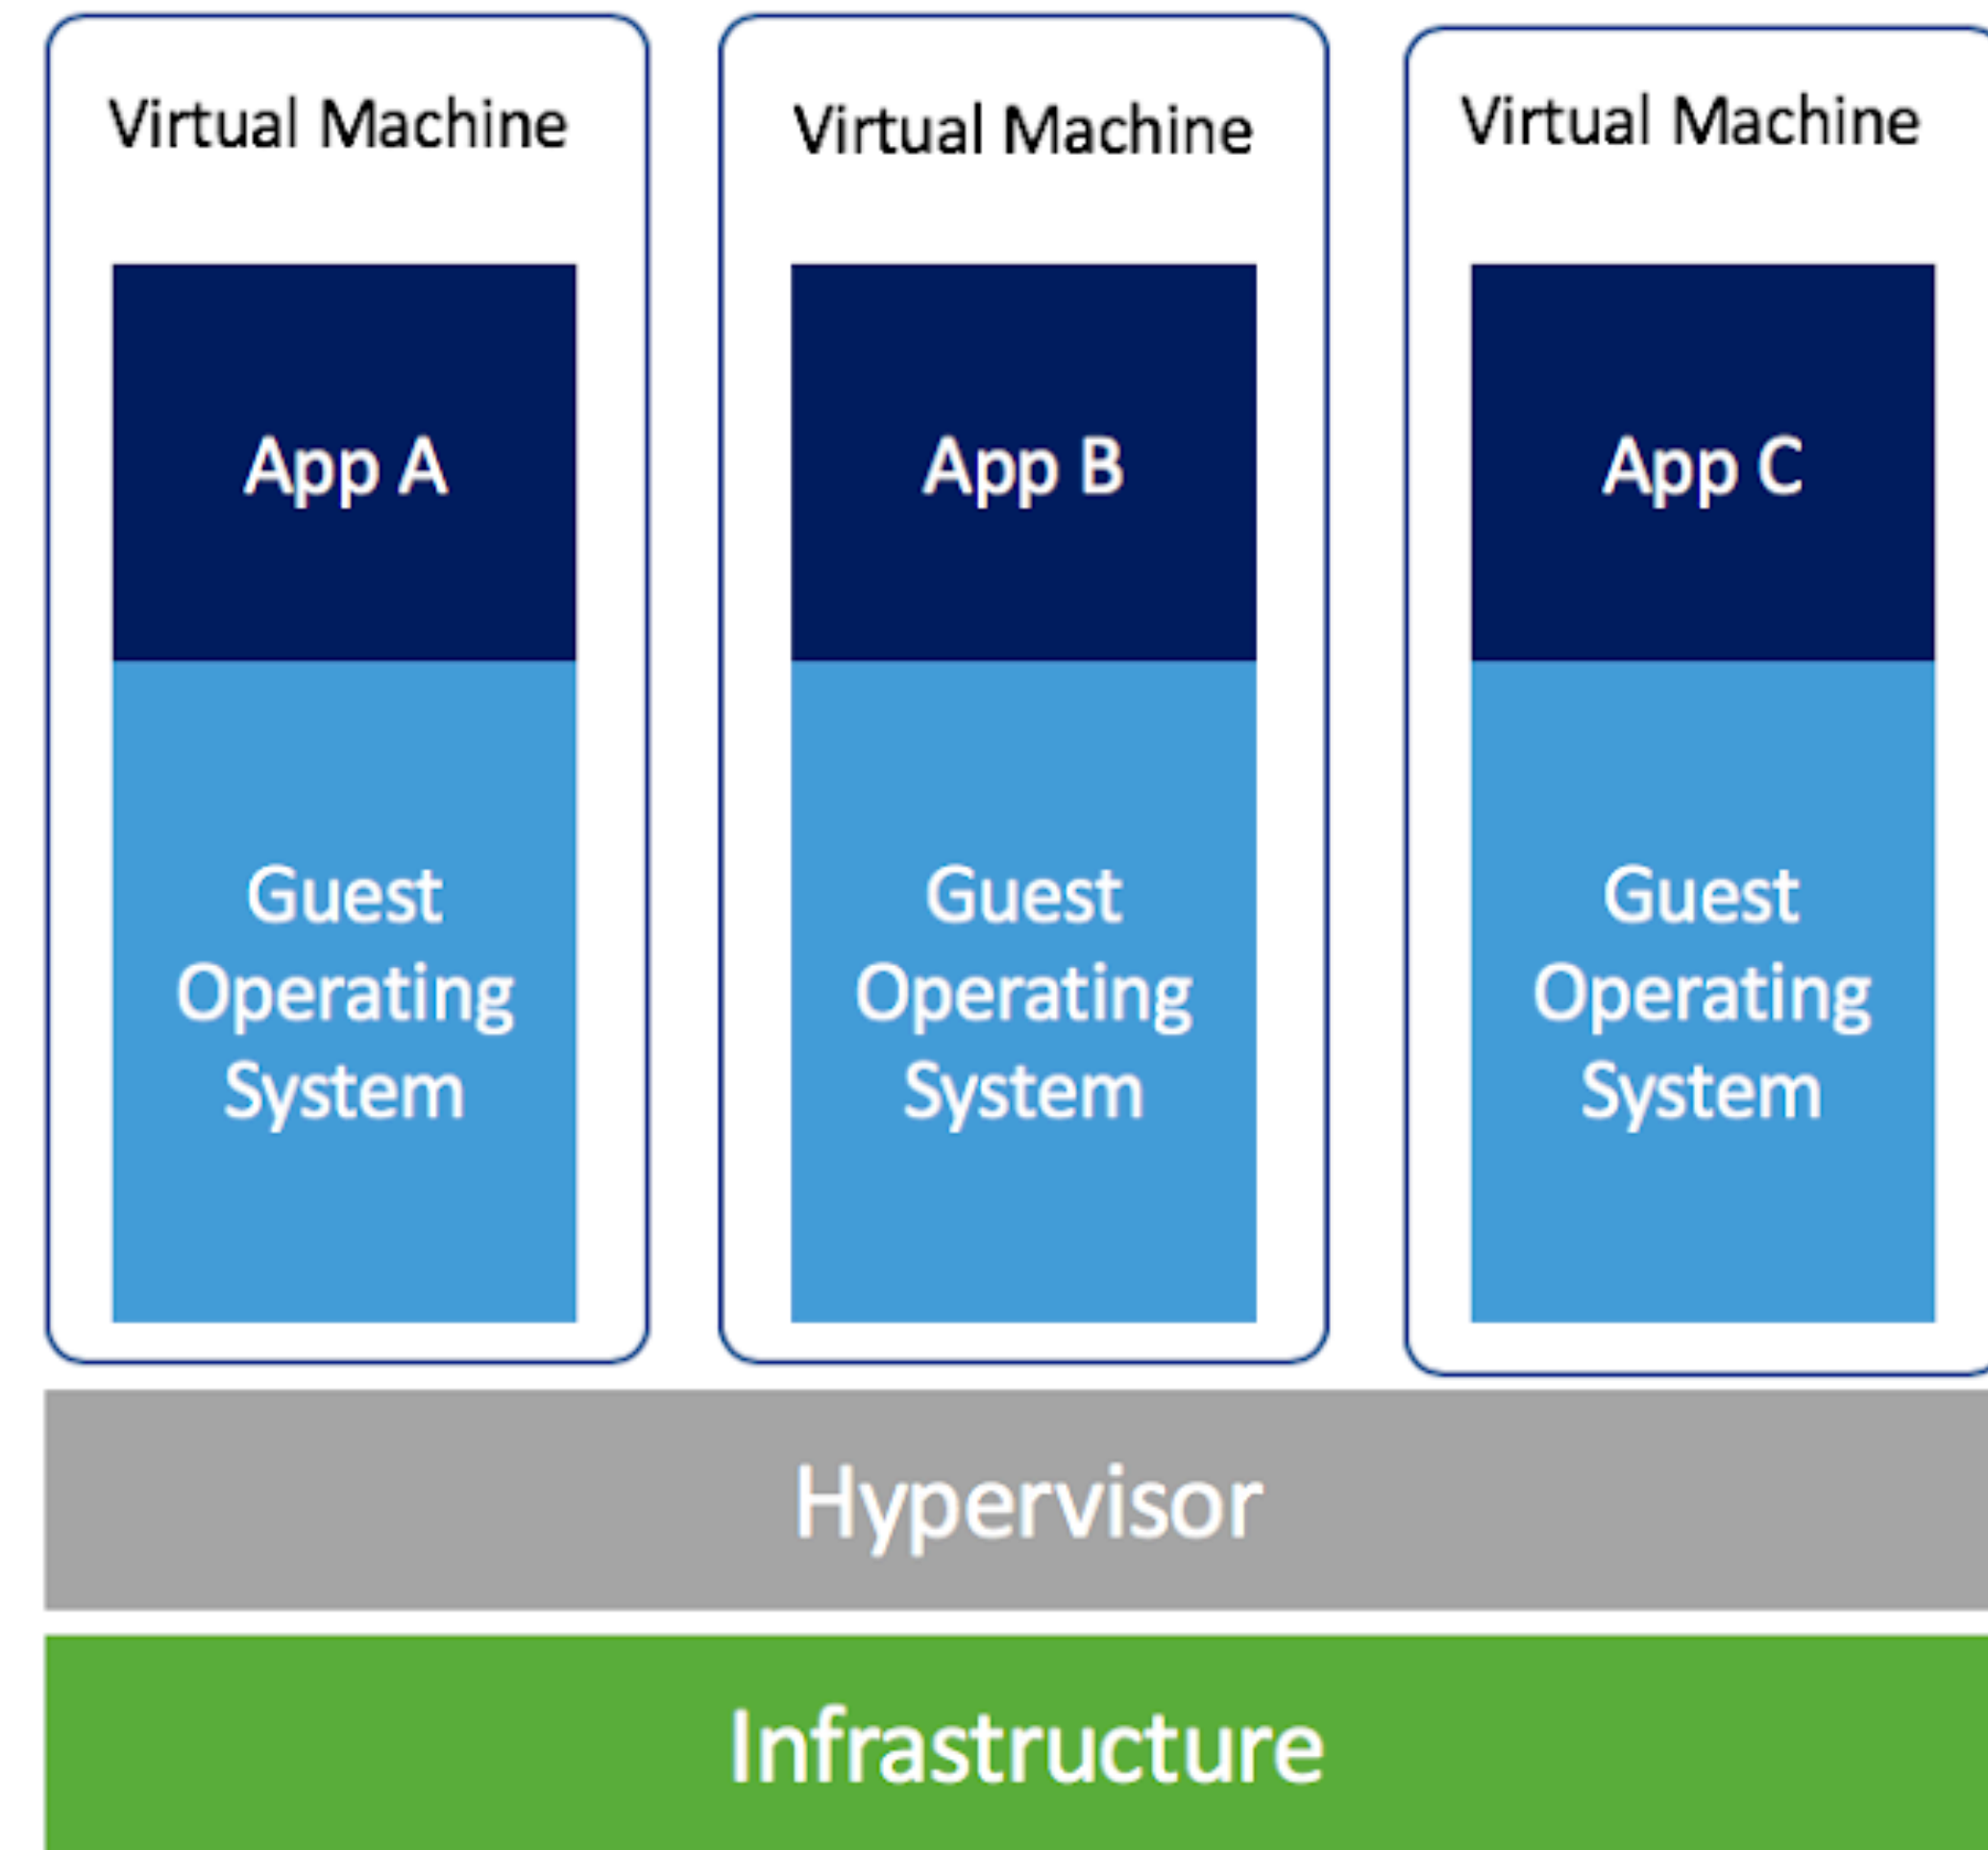

- Various types of OS supported or application supported **sandboxing** is good as a way to get defense-in-depth
- Create **temporary execution environments** for certain tasks
 - test of exe files to lure out malicious code execution
 - perform certain tasks that is more prone to attacks
 - perform certain tasks that is more sensitive
- Provide **isolation**, from other parts of system

Pro's and con's with virtualization

- Some sandbox and isolation technologies are not complete virtualization or separation
 - E.g. share *name space* (processes, file system, etc)
 - Share operating system kernel
 - Share drivers

Overview of virtualization

Containerized Applications



Pro's and con's with virtualization

- Isolation, and to have hardened and dedicated servers running specific services, are standard ways to minimize attack surface. Virtualization tools can help this
- Its easy to believe that virtualization will automatically make things secure, and that there is no way to jump between guest os', but exploits have shown this not hold true, e.g. cloudburst

VM's vs Containers vs WebAssembly

Feature	Virtual Machines	Containers	WebAssembly (Wasm)
Isolation	Full OS virtualization, strong	OS-level isolation	Sandboxed execution (virtual CPU)
Performance	Slower (heavy overhead)	Near-native	Near-native, optimized
Startup	Minutes	Seconds	Milliseconds
Size	GBs (full OS image)	MBs (includes OS dependencies)	KBs to MBs (minimal overhead)
System Access	Full OS access (kernel, drivers)	Shares OS kernel	No direct OS access (sandboxed)
Security	Strong (separate OS instances)	Moderate (kernel shared)	Strong (sandboxed, minimal attack surface)
Portability	Limited (OS-dependent)	Cross-platform (container runtimes)	Universal (Wasm runtimes)
Use Cases	Legacy applications, multi-OS environments	Cloud-native applications, microservices	Edge computing, serverless, high-security apps

Advanced attacks

Hardware attacks, etc

Attack tools



- Reverse Engineering Frameworks, such as Ghidra help debugging, disassemble, reverse engineer binaries
 - Give attackers powerful tools to introspect into firmware, drivers, kernels, applications

Example of attacks

- Attacks by attaching malicious hardware to buses and ports
 - Using debug interfaces to snoop & manipulate bus
 - **JTAG** (IEEE standard 1149.1-1990)
 - **SWD** (Serial Wire Debug)
- Firewire and other DMA based methods to access memory of a computer (*evil maid attacks, evil devices*)
- UEFI attacks via Thunderbolt (*thunderstruck attack*)

Example of attacks

- Removal of, or direct attachment to, physical memory chips (*cold boot attacks*)



Example of attacks: *cold boot attacks*



Example of attacks: *PCILeech*

Attacking
UEFI Runtime Services
and Linux


Example of attacks: *HW implants*



(TS//SI//NF) Left: Intercepted packages are opened carefully; Right: A “load station” implants a beacon


Advanced attacks

- *Rowhammer**
 - Flipping bits without accessing them
 - *Method of reading writing memory cells so that memory cells in adjacent rows become changed*
 - *Based on an unintended side effect in dynamic random-access memory (DRAM) that causes memory cells to leak their charges and interact electrically between themselves, possibly altering the contents of nearby memory rows that were not addressed in the original memory access*



Advanced attacks

- *Rowhammer**
 - *This circumvention of the isolation between DRAM memory cells*
 - *Memory leak == information leak*
 - *Have been used to **Gain Kernel Privileges, e.g. DRAMMER attack on Android***
 - *Can be used to attack **Virtual Machines***



Advanced attacks

- *Rowhammer*
 - *Have been implemented in JavaScript and runned in a browser*
 - *Modern variants* have been used to **defeat ECC memory***

* "Exploiting Correcting Codes: On the Effectiveness of ECC Memory Against Rowhammer Attacks" <https://cs.vu.nl/~lcr220/ecc/ecc-rh-paper-eccploit-press-preprint.pdf>

Advanced attacks

- *Rowhammer**
 - Initial research published 2014, but variants have been developed later
 - Rowhammer.js (2015)
 - Blacksmith (2022)
 - Half-double (2021)
 - Zenhammer (2024, AMD architecture)
 - RISC-Hammer (2024, RISC-Y architecture)
 - Hardware solutions to protect against it have been circumvented



Advanced attacks



- *Meltdown* & Spectre***
 - *Initial research published January 2018*
 - *Microarchitectural bugs in CPU*
 - *Meltdown breaks isolation between **user land and kernel***
 - *Spectre breaks isolation between **applications in user land***

<https://meltdownattack.com/>

* Lipp et al "Meltdown: Reading Kernel Memory from User Space" <https://meltdownattack.com/meltdown.pdf>

** Kocker et al "Spectre Attacks: Exploiting Speculative Execution" <https://spectreattack.com/spectre.pdf>



Advanced attacks



- *Meltdown & Spectre*

- *work on personal computers, mobile devices, and in the cloud*
- *Works on Windows, Linux, Android, etc*
- *Works on containers: docker, LXC, OpenVZ etc*



Advanced attacks



- *Meltdown & Spectre*

- *All modern CPUs are vulnerable (x86,AMD,ARM) in various degrees*

Attack		Spectre-PHT	Spectre-BTB	Spectre-RSB	Spectre-STL
Method					
Intel	same-address-space in-place	● [52, 50] ★		● [62]	● [32]
	out-of-place	★	● [18]	● [62, 54]	○
	cross-address-space in-place	★	● [52, 18]	● [62, 54]	○
	out-of-place	★	● [52]	● [54]	○
ARM	same-address-space in-place	● [52, 50] ★		● [6]	● [6]
	out-of-place	★	☆	● [6]	○
	cross-address-space in-place	★	● [6, 52]	☆	○
	out-of-place	★	☆	☆	○
AMD	same-address-space in-place	● [52]	★	★	● [32]
	out-of-place	★	☆	★	○
	cross-address-space in-place	★	● [52]	★	○
	out-of-place	★	☆	★	○

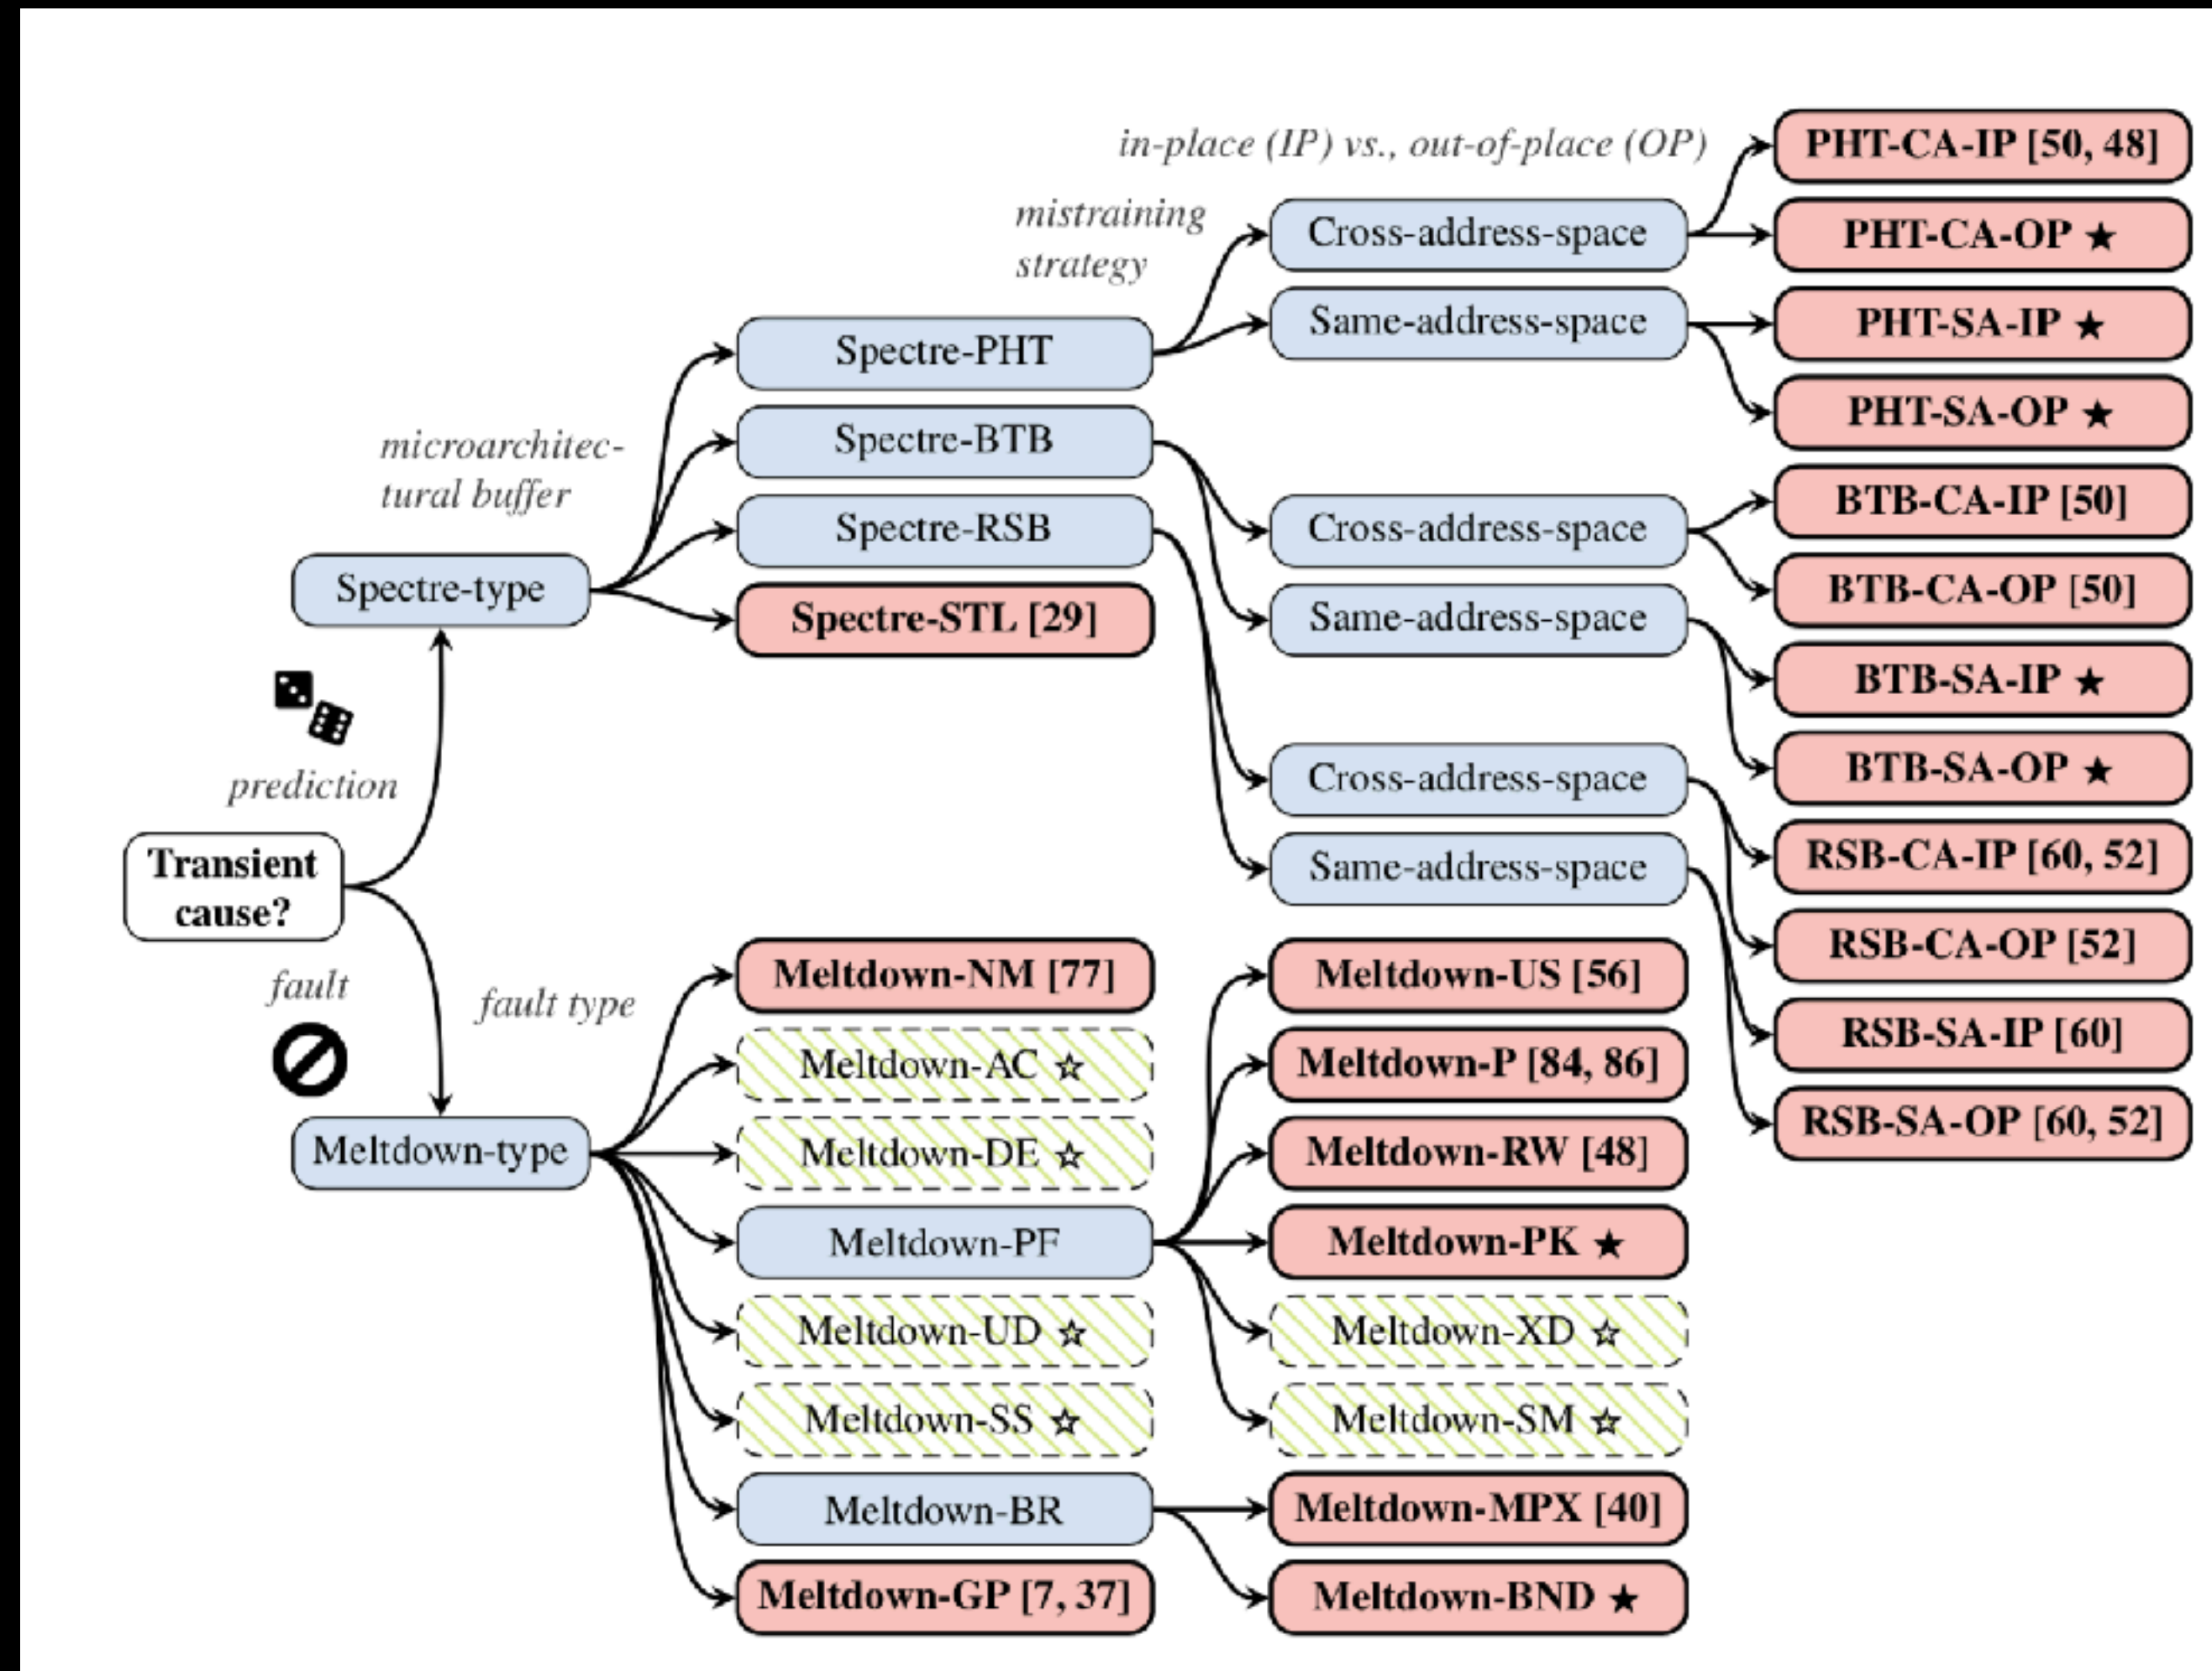
Symbols indicate whether an attack is possible and known (●), not possible and known (○), possible and previously unknown or not shown (★), or tested and did not work and previously unknown or not shown (☆). All tests performed with no defenses enabled.

Attack		Meltdown-US [59]	Meltdown-P [90, 93]	Meltdown-GP [10, 40]	Meltdown-NM [83]	Meltdown-RW [50]	Meltdown-PK	Meltdown-BR	Meltdown-DE	Meltdown-AC	Meltdown-UD	Meltdown-SS	Meltdown-XD	Meltdown-SM
Vendor														
Intel		●	●	●	●	●	★	★	☆	☆	☆	☆	☆	☆
ARM		●	○	●	—	●	—	—	☆	☆	☆	—	☆	☆
AMD		○	○	○	○	○	—	★	☆	☆	☆	☆	☆	☆

Symbols indicate whether at least one CPU model is vulnerable (filled) vs. no CPU is known to be vulnerable (empty). Glossary: reproduced (● vs. ○), first showed in this paper (★ vs. ☆), not applicable (—). All tests performed without defenses enabled.






Advanced attacks





Advanced attacks





Advanced attacks



- *Spectre class vulnerabilities will remain unfixed because otherwise CPU designers will have to disable speculative execution which will entail a **massive performance loss***

Advanced attacks

- *Microarchitecture Data Sampling attacks*
- *Side channel attacks*
- *Timing side channel attacks*
- *Power Analysis side channel attack*

Advanced attacks

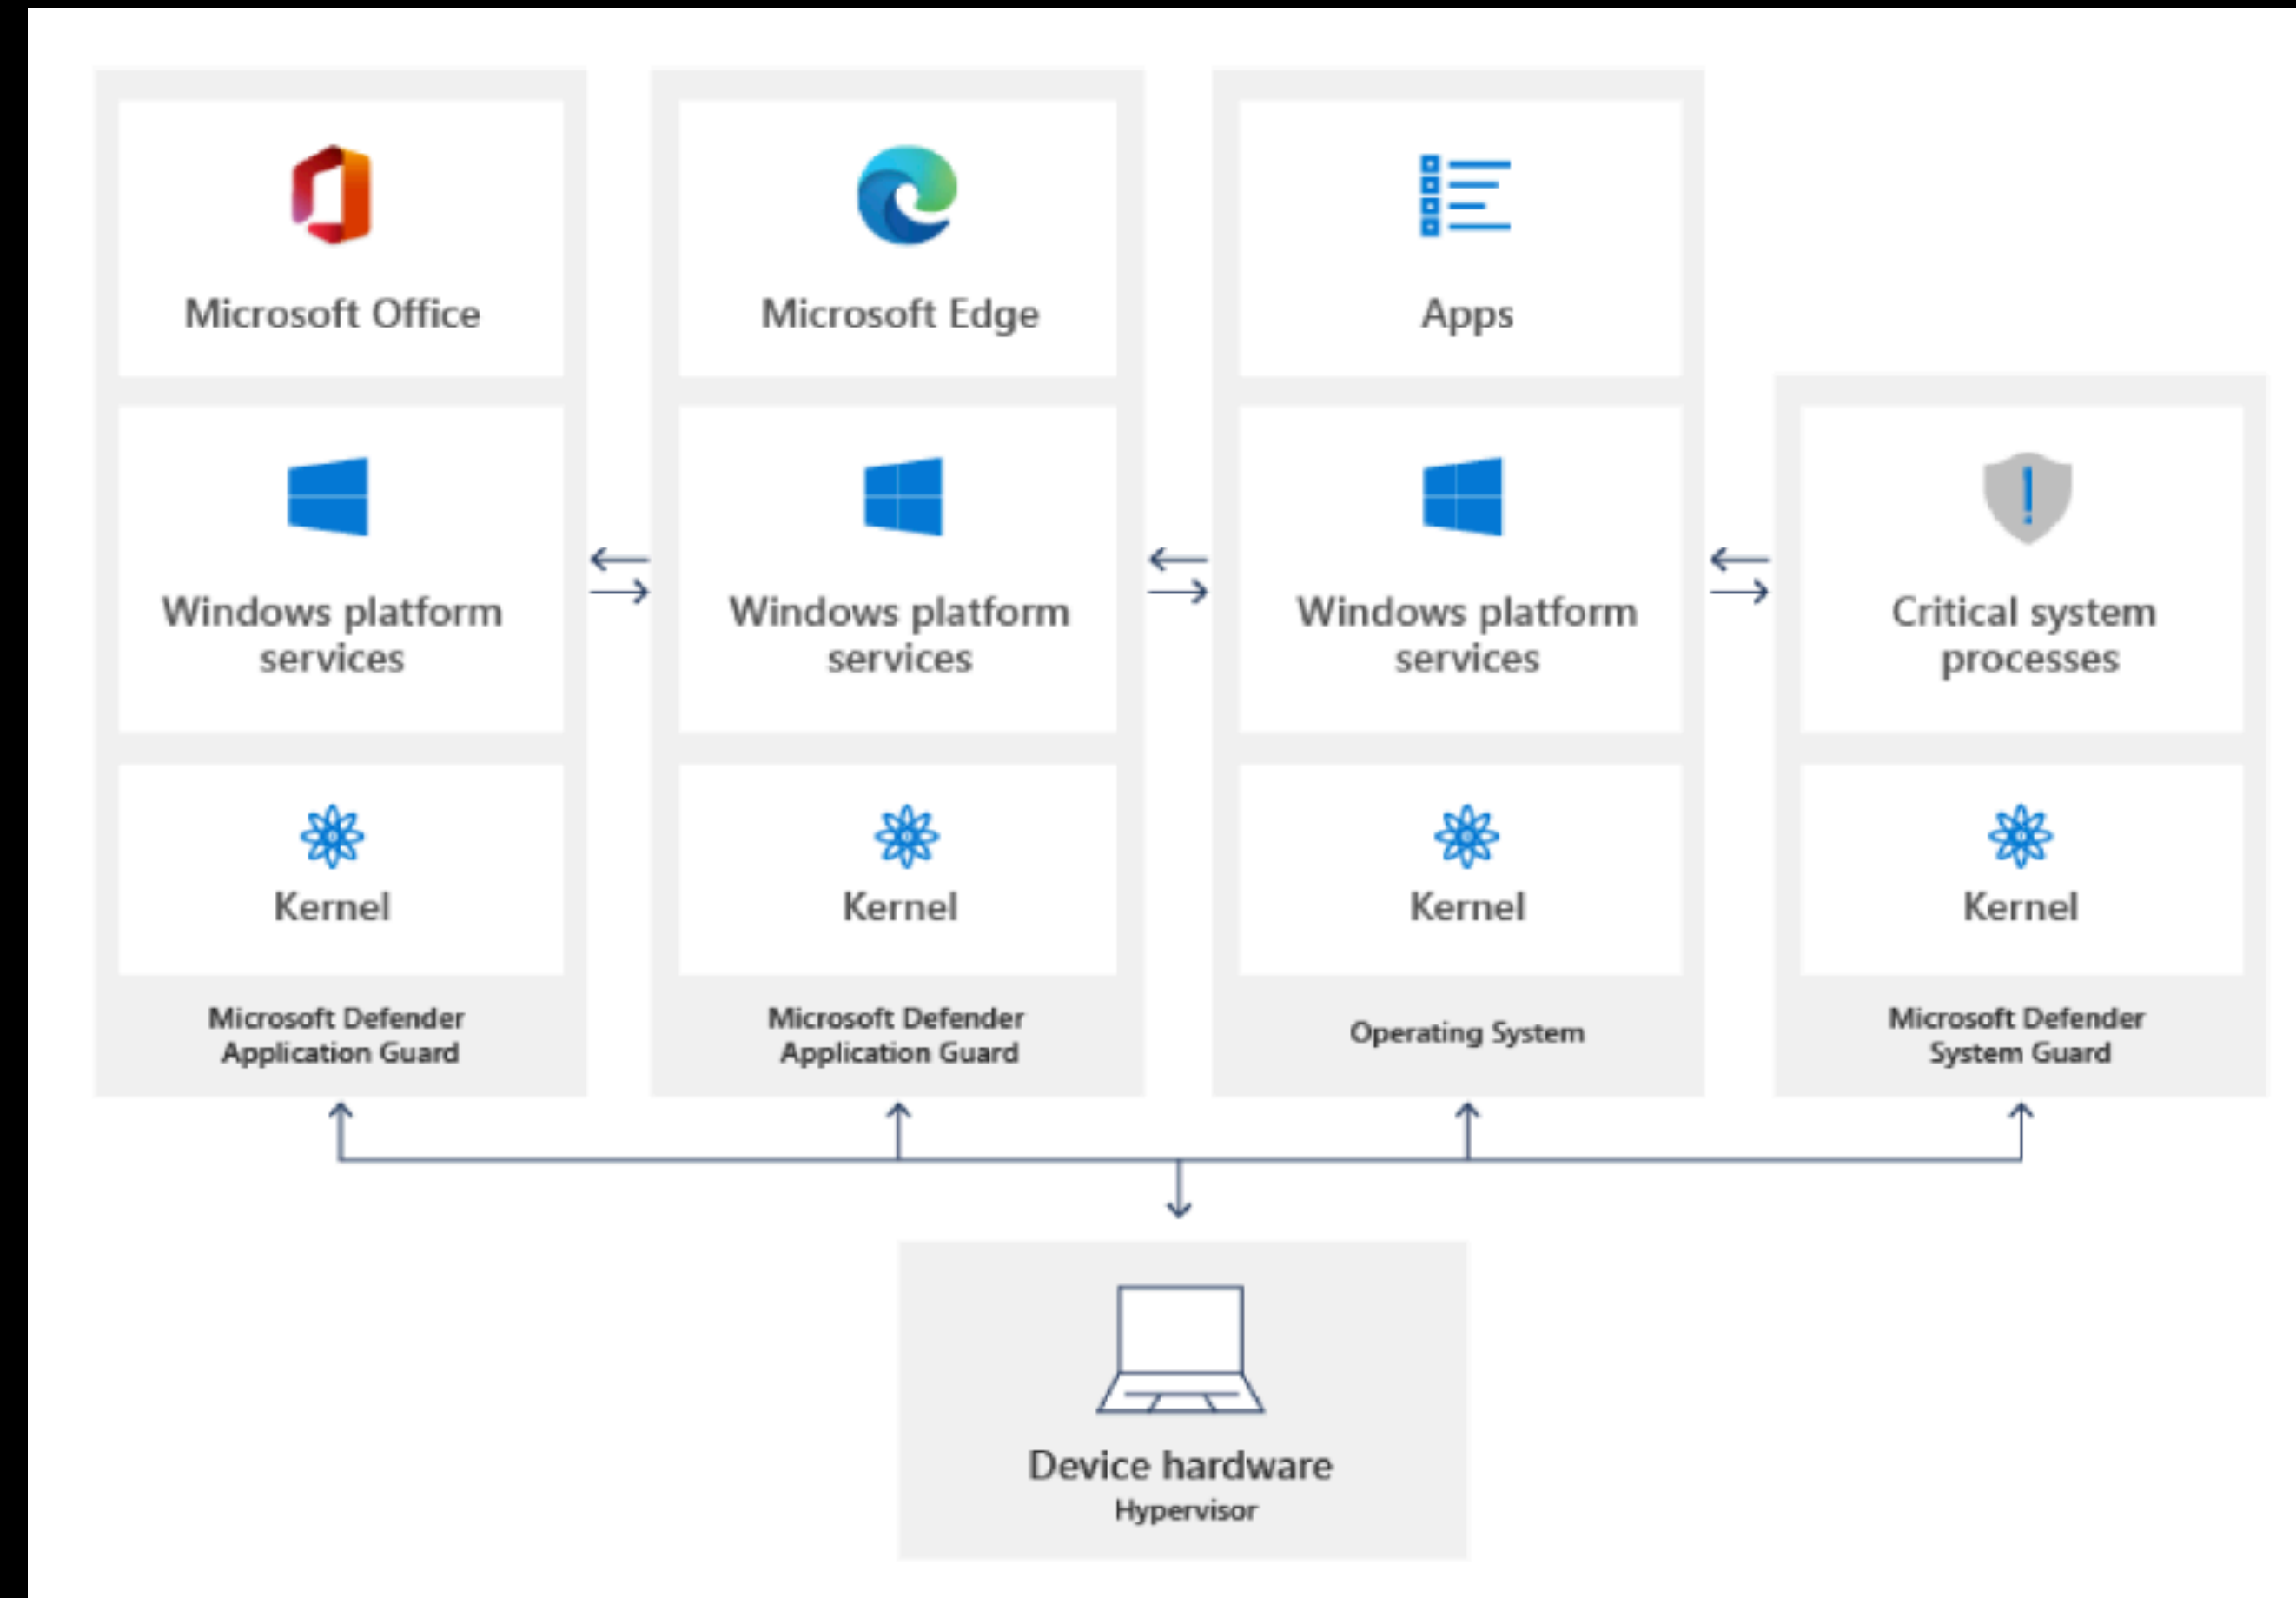
- *Attacks against Intel Management Engine*
 - *Proprietary and non-documented*
 - *Own OS (Minix!)*
 - *Reverse engineered and analysed by attackers*
 - *Found multiple vulnerabilities in Skylake & Kabylake architecture*

Examples of modern security controls

Windows Defender security features in Win 10, Win 11

Windows

- Application Guard, **WDAG**
 - App & browser control
 - Isolation browsing
- WDAC (Windows Defender Application Control) give application & driver whitelisting
- VBS (Virtualization-Based Security)
- WDAC & VBS used to be Windows Device Guard



Windows

- Windows Device Guard, And Applocker, now called Windows Defender Application Control
 - Attributes of the **codesigning certificate**(s) used to sign an app and its binaries
 - Attributes of the app's binaries that come from the **signed metadata for the files**, such as Original Filename and version, or the hash of the file
 - The **path from which the app or file is launched**

Windows

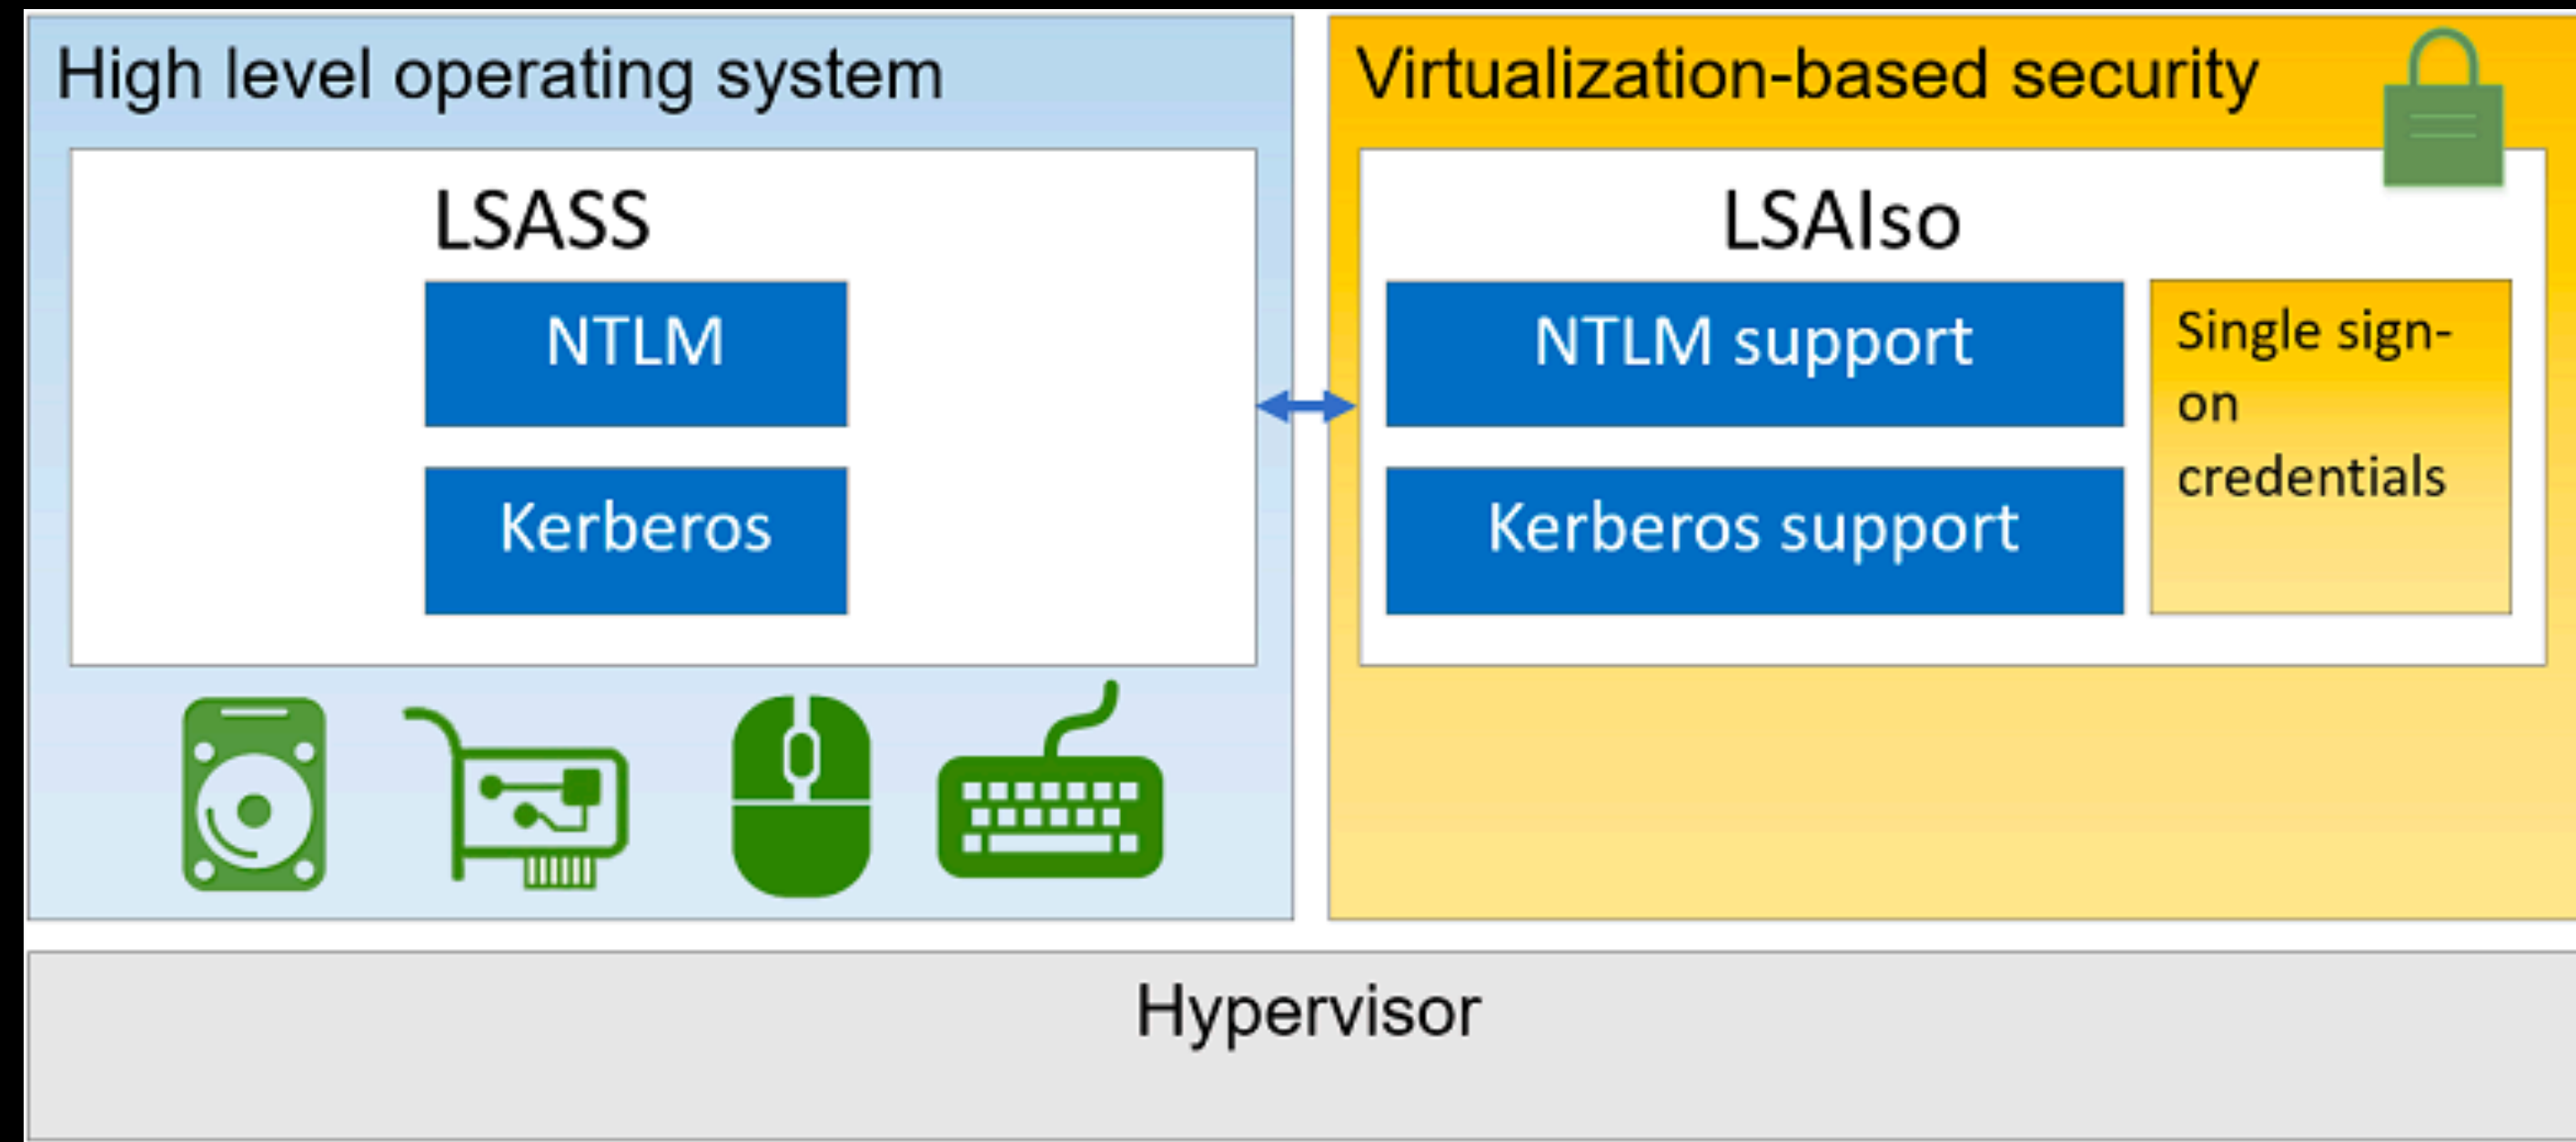
- *Core isolation with Memory integrity*, aka Hypervisor-protected Code Integrity (**HVCI**)
 - make it difficult for malicious programs to use low-level drivers to hijack your computer

Windows

- Windows Defender Exploit Guard, **WDEG**
- **Attack Surface Reduction (ASR)**: A set of controls that enterprises can enable to prevent malware from getting on the machine by blocking Office-, script-, and email-based threats
- **Network protection**: Protects the endpoint against web-based threats by blocking any outbound process on the device to untrusted hosts/IP through Windows Defender SmartScreen
- **Controlled folder access**: Protects sensitive data from ransomware by blocking untrusted processes from accessing your protected folders
- **Exploit protection**: A set of exploit mitigations (replacing EMET) that can be easily configured to protect your system and applications

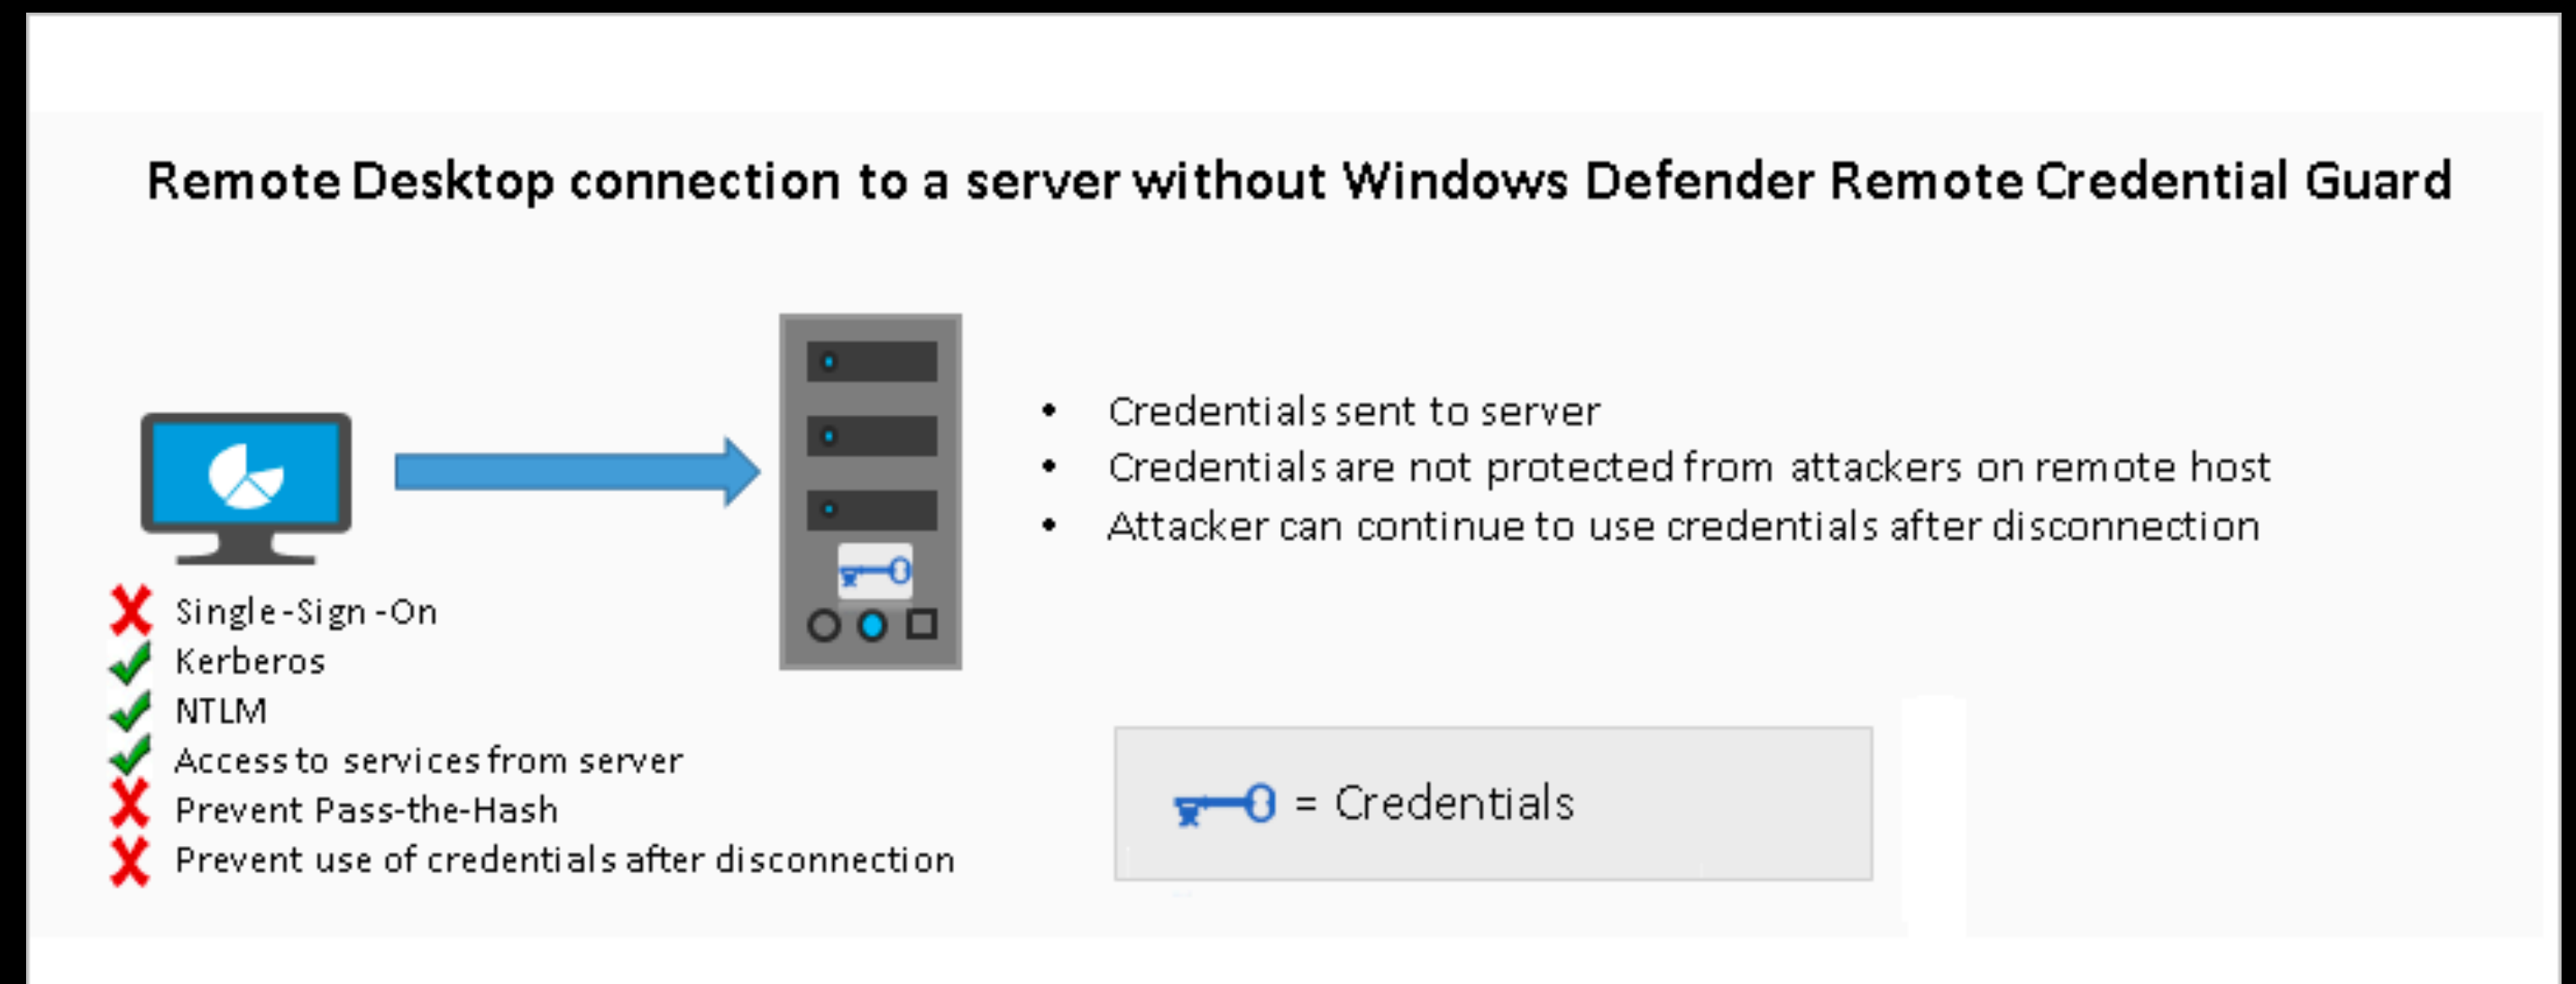
Windows

- Windows Credential Guard
 - To protect *Local Security Authority Server Service* (LSASS) by moving it into *LSAIso*
- Build on top of
 - Virtualization Based Security (VBS)
 - Secure boot
 - Trusted Platform Module (TPM)
 - UEFI lock



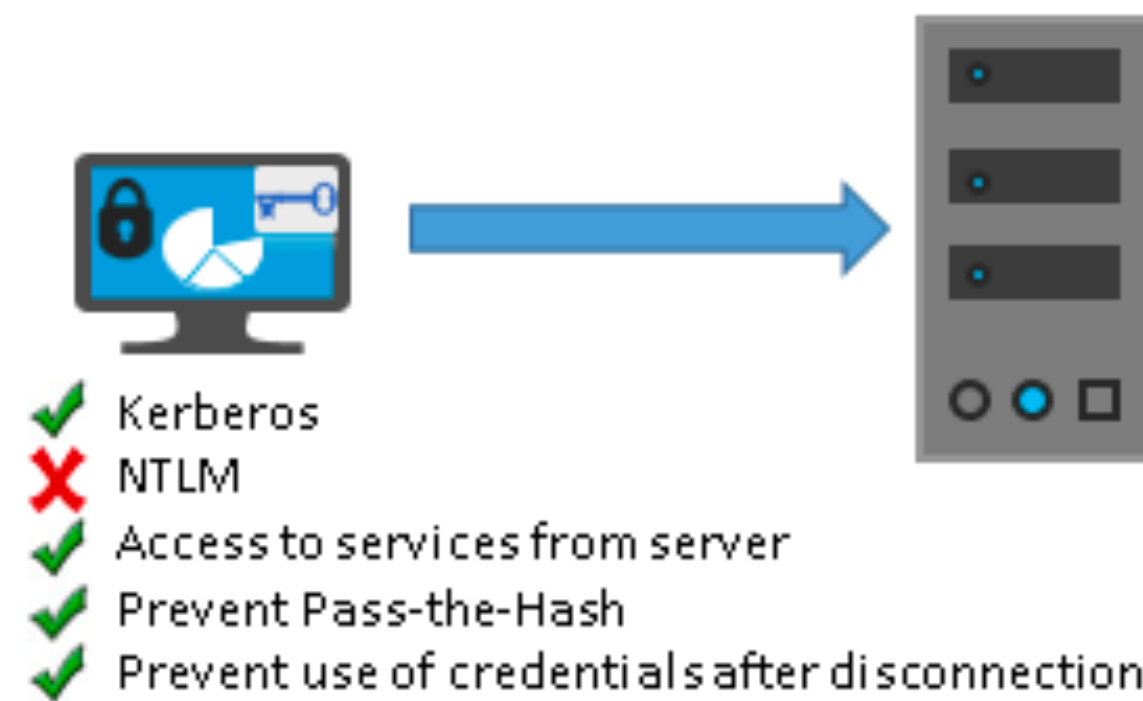
Windows

- Windows Remote Credential Guard
- To protect *against theft of credentials sent to server side*
 - Others that have admin access to the server
- Especially important on jump hosts




Windows

Windows Defender Remote Credential Guard



- Credentials protected by Windows Defender Remote Credential Guard
- Connect to other systems using SSO
- Host must support Windows Defender Remote Credential Guard

Restricted Admin Mode



- Credentials used are remote server local admin credentials
- Connect to other systems using the host's identity
- Host must support Restricted Admin mode
- Highest protection level
- Requires user account administrator rights

🔒 = Credential protection
🔑 = Credentials

MacOS X

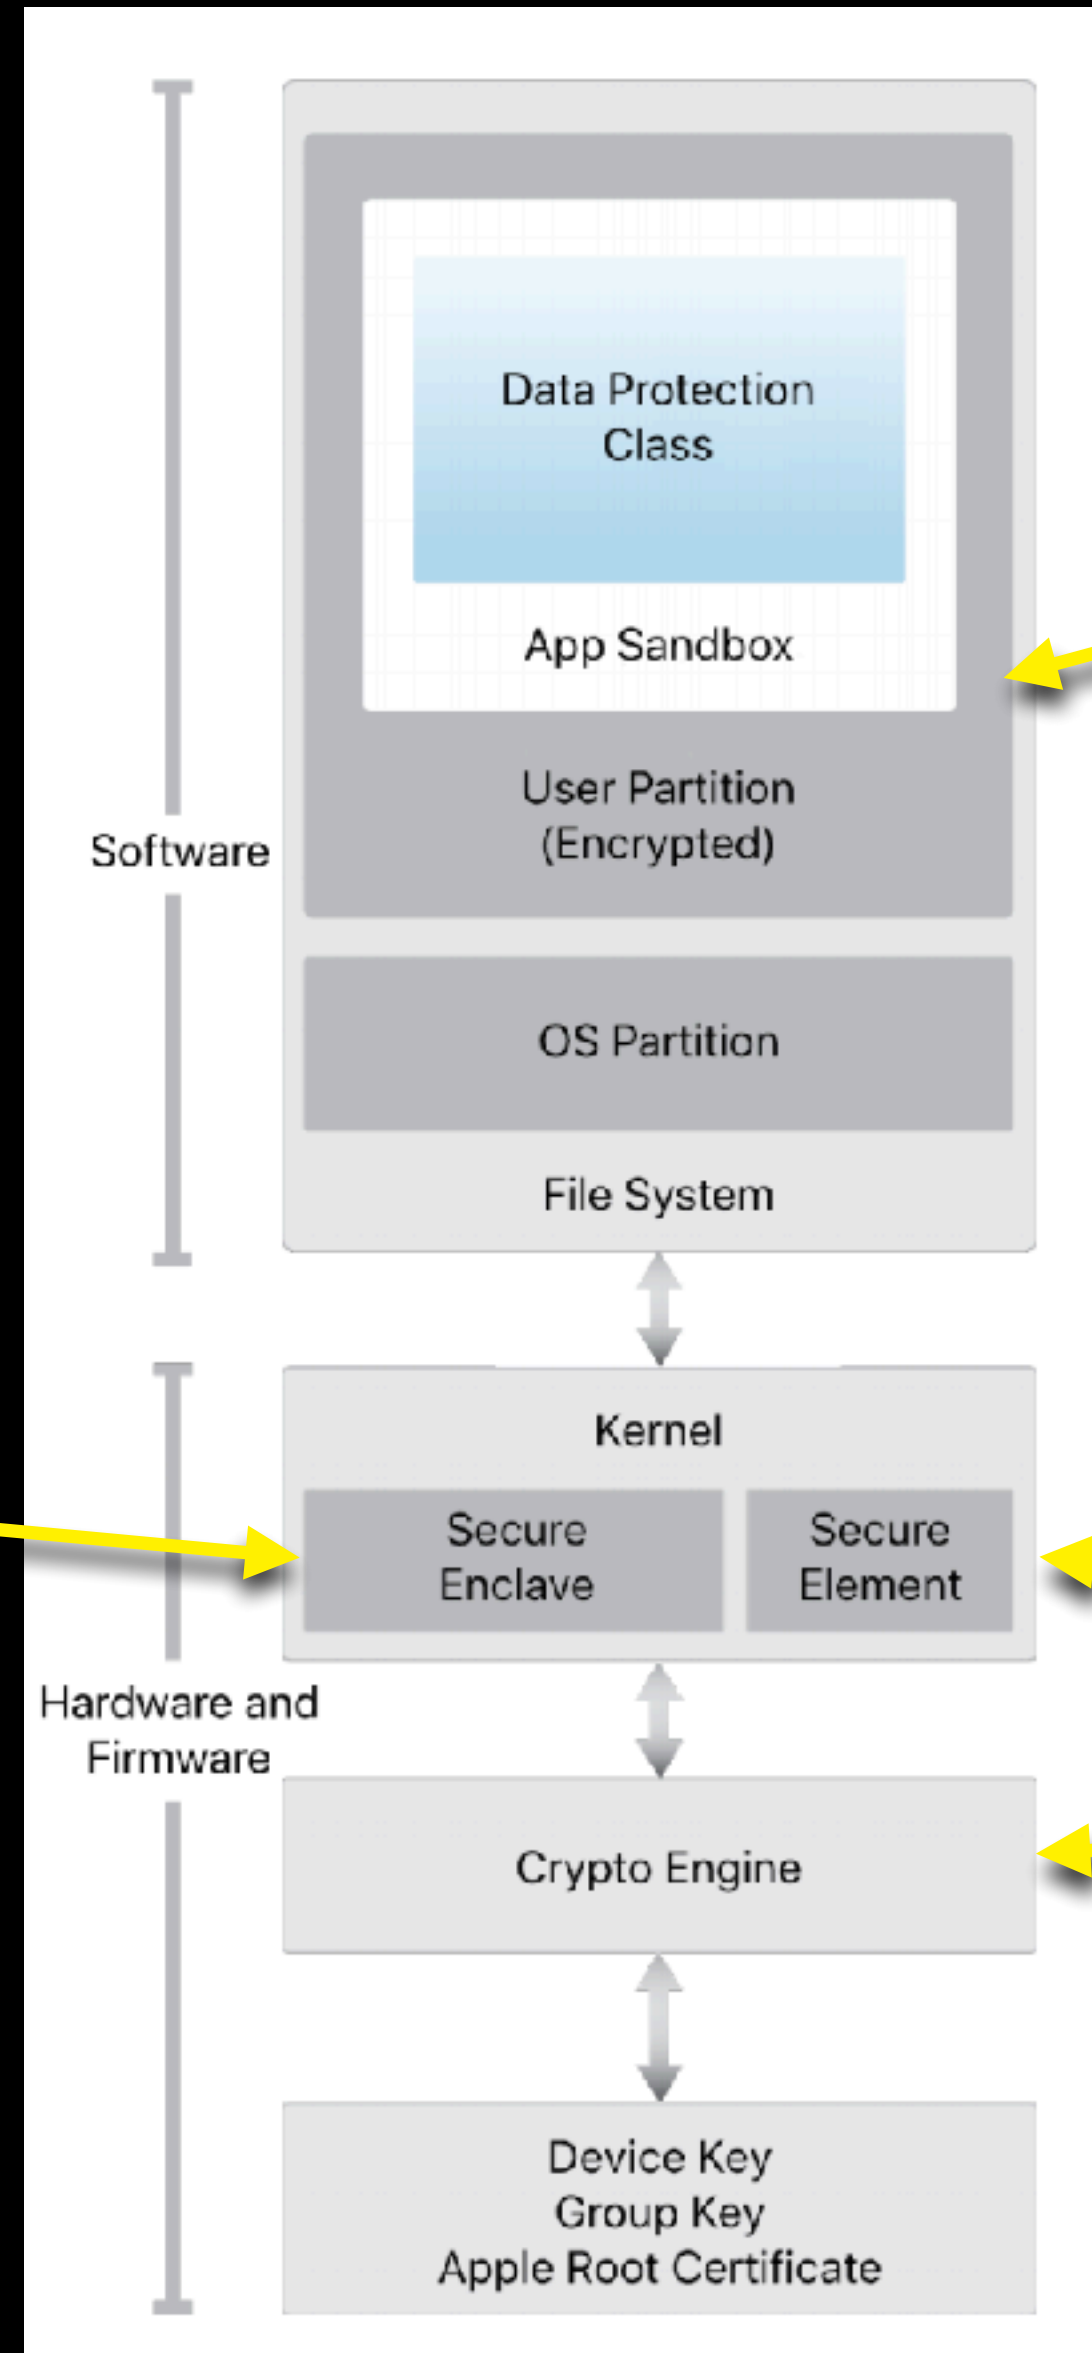
- **Hardware based security**
 - Secure Enclave
 - Memory Tagging & Pointer Authentication
 - Hardware-Accelerated Encryption
- **Isolation**
 - App sandbox
- **FileVault**
 - Full disk encryption
- **GateKeeper**
 - Checks code signing
- **XProtect**
 - Malware protection

iOS/iPadOS

- Share many security features with MacOS X
- Additional ones include
 - **BlastDoor**. A way to take a look at all incoming messages and inspect their content in a secure environment, which prevents any malicious code inside of a message from interacting with iOS or accessing user data.
 - **LockDown mode**: add many restrictions to applications, e.g. web browsing, messaging, FaceTime, photos, etc

Apple iOS device security

Secure Enclave




App Sandbox

Secure Element

Crypto Engine

Apple device security



Secure Enclave

Tools mentioned during the class

- Ghidra - Reverse Engineering Framework
- IDA pro - Disassembler
- Hexray - Decompiler
- Ollydbg, windbg - Other disassemblers
- Bindiff - Advanced tool from zynamics to compare binaries, with call graphs etc. Not same as built-in windows tool with same name.

References used during the class

- <https://www.commoncriteriaportal.org/>
- <https://www.cs.virginia.edu/~av6ds/papers/isca2021a.pdf>
- <https://www.cvedetails.com/top-50-products.php>
- <https://owasp.org/www-project-top-ten/>

References used during the class

- http://en.wikipedia.org/wiki/Source_lines_of_code
- <https://sources.debian.org/stats/>
- <https://informationisbeautiful.net/visualizations/million-lines-of-code/>
-

References used during the class

- <https://docs.microsoft.com/en-us/windows/security/identity-protection/credential-guard/credential-guard-how-it-works>
- <https://docs.microsoft.com/en-us/deployedge/microsoft-edge-security-windows-defender-application-guard>
- <https://docs.microsoft.com/en-us/windows/security/threat-protection/microsoft-defender-application-guard/md-app-guard-overview>
- <https://docs.microsoft.com/en-us/windows/security/identity-protection/remote-credential-guard>

References used during the class

- <https://support.microsoft.com/en-us/windows/core-isolation-e30ed737-17d8-42f3-a2a9-87521df09b78>
- https://en.wikipedia.org/wiki/Local_Security_Authority_Subsystem_Service
- <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.91.5728&rep=rep1&type=pdf>
-

References used during the class

- https://www.cisa.gov/sites/default/files/2024-04/CSRB_Review_of_the_Summer_2023_MEO_Intrusion_Final_508c.pdf
- https://media.ccc.de/v/37c3-12142-breaking_drm_in_polish_trains
- <https://www.ccc.de/en/updates/2024/das-ist-vollig-entgleist>
- <https://attack.mitre.org/>
- https://www.cs.cmu.edu/~rdriley/487/papers/Thompson_1984_ReflectionsonTrustingTrust.pdf
-

References used during the class

- <https://umatechnology.org/the-truth-about-the-intels-hidden-minix-os-and-security-concerns/>
- <https://www.bleepingcomputer.com/news/hardware/intels-secret-cpu-on-chip-management-engine-me-runs-on-minix-os/>
-