



Software Compartmentalization Everywhere - *What Will it Take?*



Hugo Lefeuvre

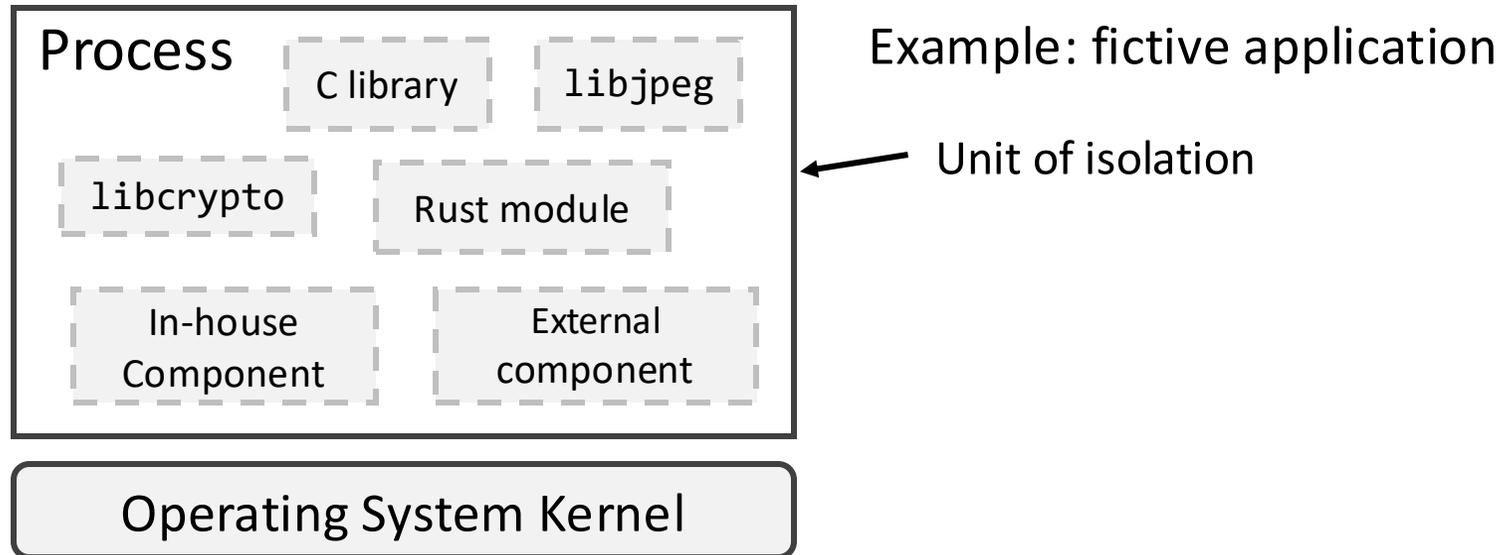
Postdoctoral Research Fellow

University of British Columbia

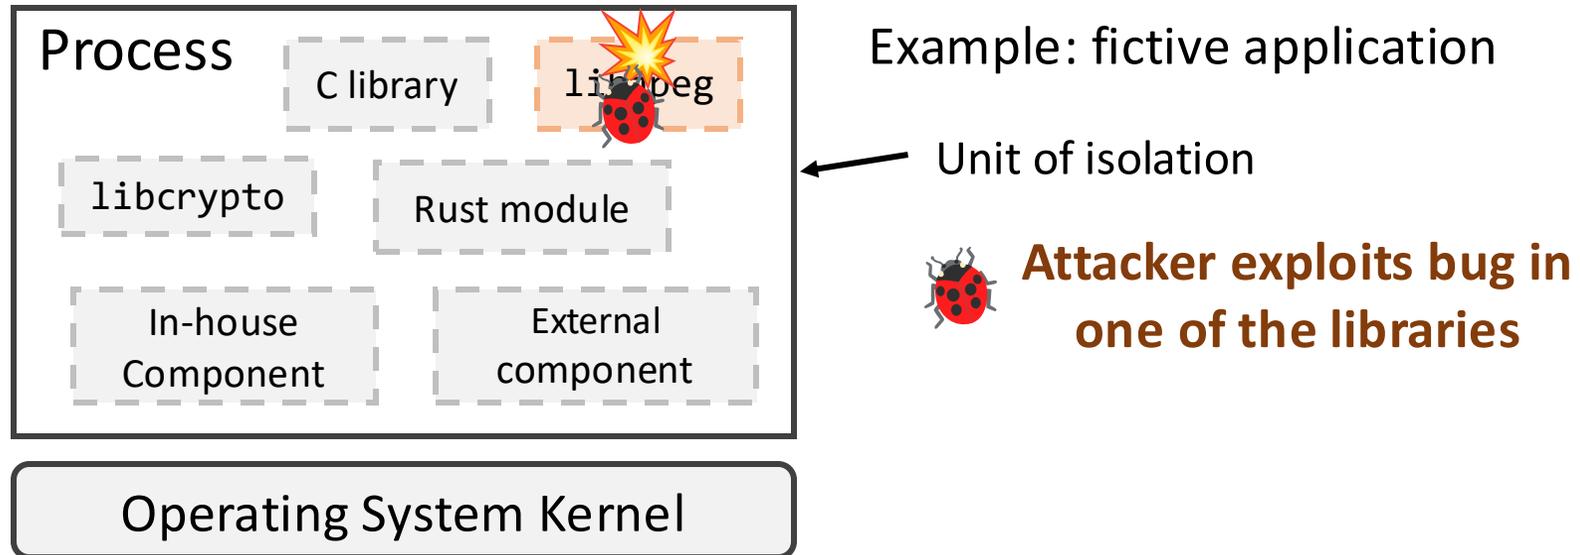
Vancouver (🇨🇦)

Historically, **program = unit of trust**

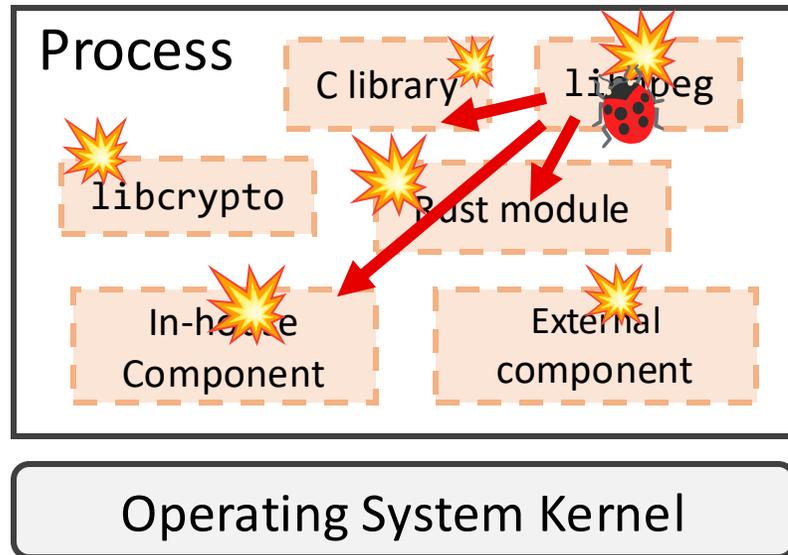
Historically, **program = unit of trust**



Historically, **program = unit of trust**



Historically, **program = unit of trust**



Example: fictive application

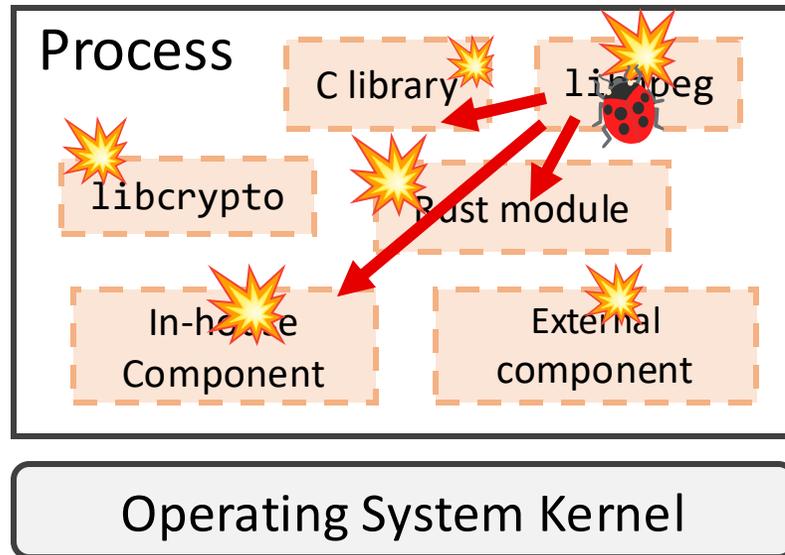
← Unit of isolation



Attacker exploits bug in one of the libraries

Spreads trivially because of **ambient trust**

Historically, **program = unit of trust**



Example: fictive application

← Unit of isolation

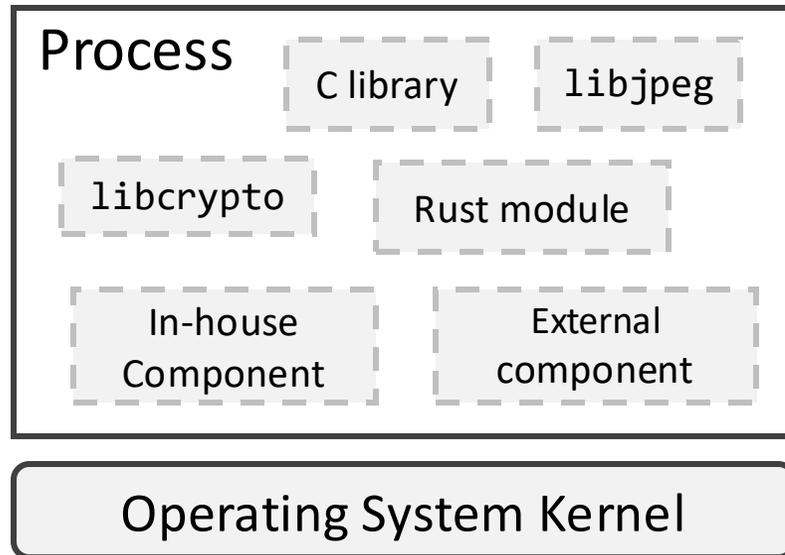


Attacker exploits bug in one of the libraries

Spreads trivially because of **ambient trust**

This historical approach is **increasingly problematic**

Historically, **program = unit of trust**

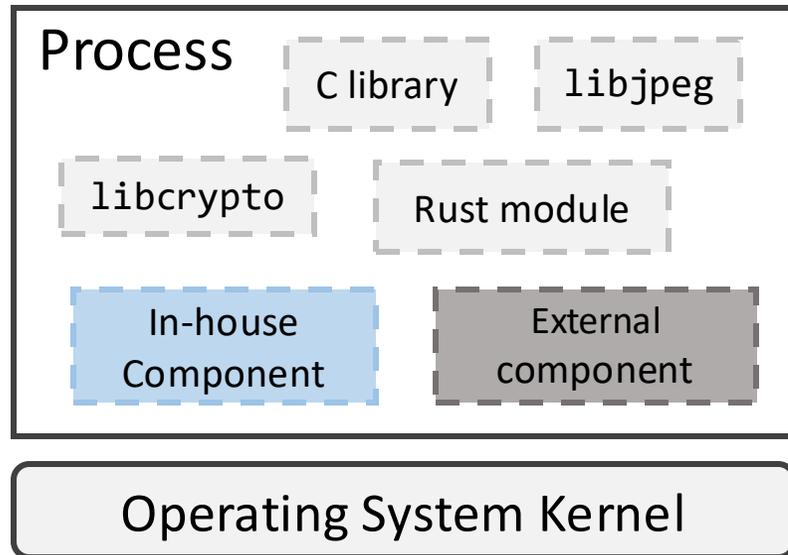


This historical approach is **increasingly problematic**

Among others...

Modern programs are **deeply heterogeneous**

Historically, **program = unit of trust**



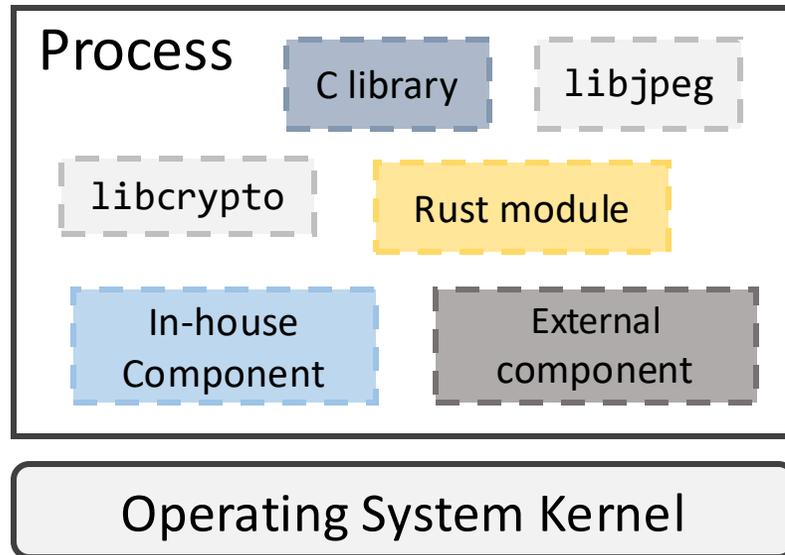
This historical approach is **increasingly problematic**

Among others...

Modern programs are **deeply heterogeneous**

Different origins...

Historically, **program = unit of trust**



This historical approach is **increasingly problematic**

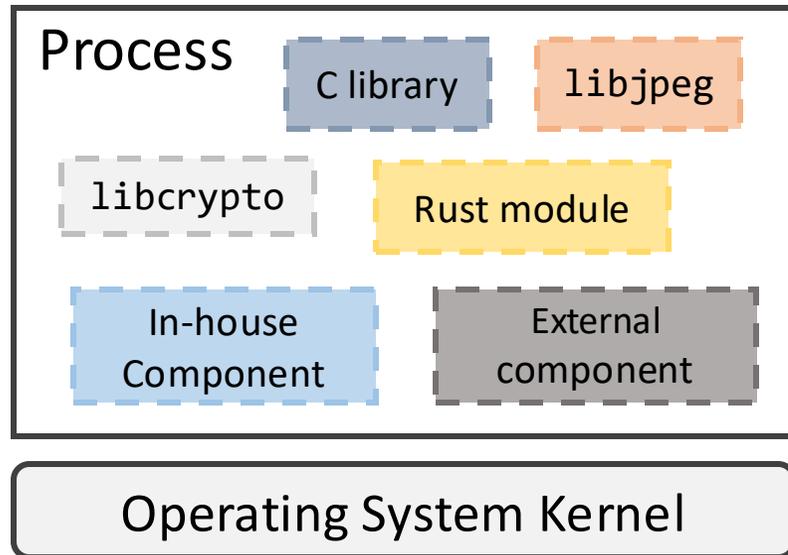
Among others...

Modern programs are **deeply heterogeneous**

Different programming languages...

Different origins...

Historically, **program = unit of trust**



This historical approach is **increasingly problematic**

Among others...

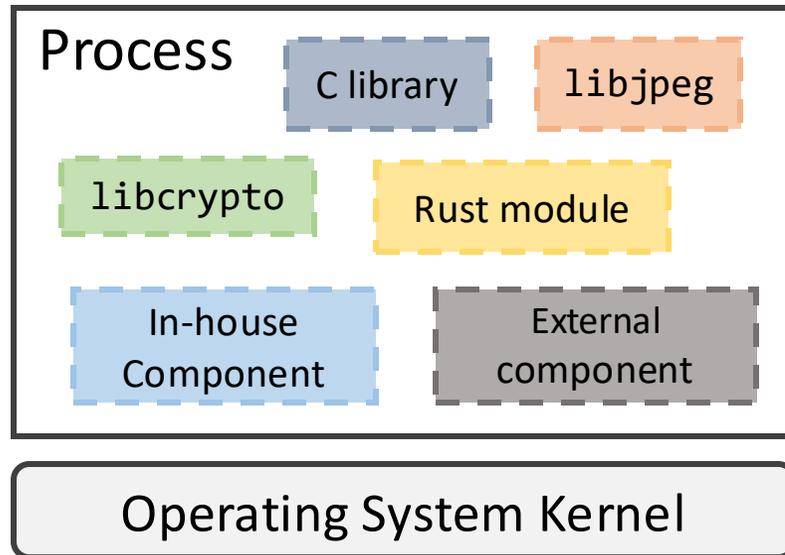
Modern programs are **deeply heterogeneous**

Different programming languages...

Different origins...

Different degrees of risk...

Historically, **program = unit of trust**



This historical approach is **increasingly problematic**

Among others...

Modern programs are **deeply heterogeneous**

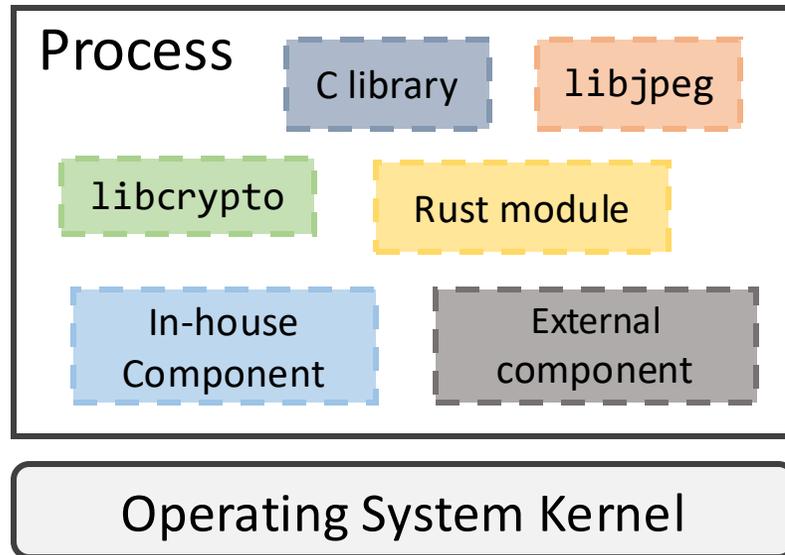
Different programming languages...

More or less valuable data...

Different origins...

Different degrees of risk...

Historically, **program = unit of trust**



This historical approach is **increasingly problematic**

Among others...

Modern programs are **deeply heterogeneous**

Different programming languages...

More or less valuable data...

Different origins...

Different degrees of risk...



*The average application has around **180 open-source components** — that's an increase from around 150 from which we found last year.*

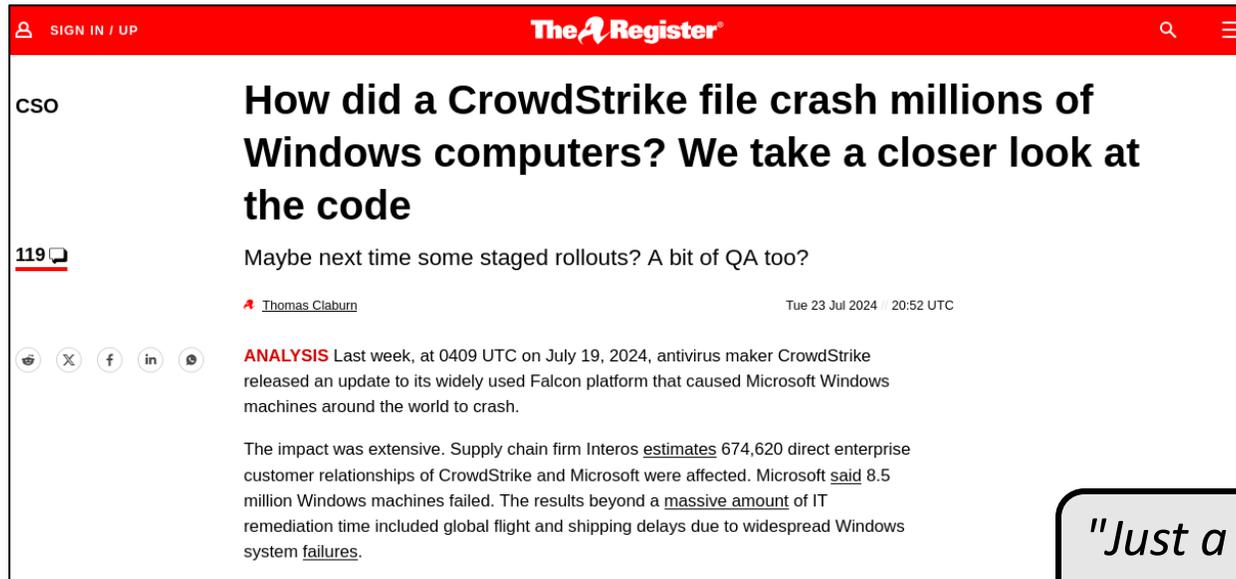
(1) 10th Annual State of the Software Supply Chain, Sonatype Inc.



(1)

Illustration of this problem:

Last year's Crowdstrike disaster where a **bug in an external antivirus plug-in crashes millions of computers**



(1)

"Just a crash": what if it had been a supply-chain attack?

Illustration of this problem:

Last year's Crowdstrike disaster where a **bug in an external antivirus plug-in crashes millions of computers**

(1) https://www.theregister.com/2024/07/23/crowdstrike_failure_shows_need_for/

SIGN IN / UP The Register

CSO

How did a CrowdStrike file crash millions of Windows computers? We take a closer look at the code

119

Thomas Claburn Tue 23 Jul 2024 20:52 UTC

ANALYSIS Last week, at 0409 UTC on July 19, 2024, antivirus maker CrowdStrike released an update to its widely used Falcon platform that caused Microsoft Windows machines around the world to crash.

The impact was extensive. Supply chain firm Interos estimates 674,620 direct enterprise customer relationships of CrowdStrike and Microsoft were affected. Microsoft said 8.5 million Windows machines failed. The results beyond a massive amount of IT remediation time included global flight and shipping delays due to widespread Windows system failures.

(1)

"Just a crash": what if it had been a supply-chain attack?

Illustration of this problem:

Last year's Crowdstrike disaster where a **bug in an external antivirus plug-in crashes millions of computers**



Program = single unit of trust

Compartmentalization

Compartmentalization

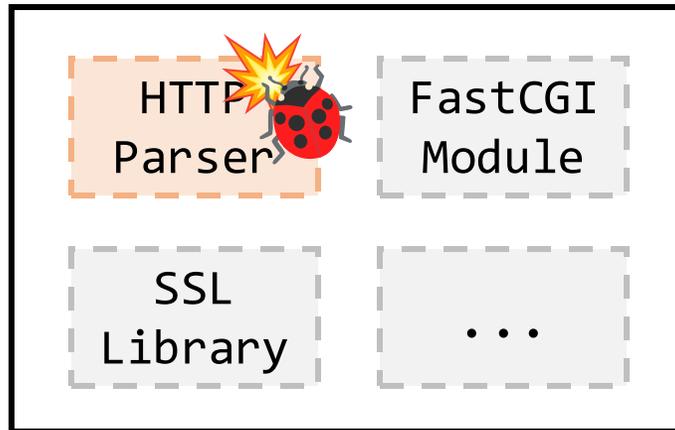
Software compartmentalization =

- design programs **split into distrusting and isolated components**
- such that if one of the components is breached, the attacker **does not compromise the whole program**

Compartmentalization

Software compartmentalization =

- design programs **split into distrusting and isolated components**
- such that if one of the components is breached, the attacker **does not compromise the whole program**



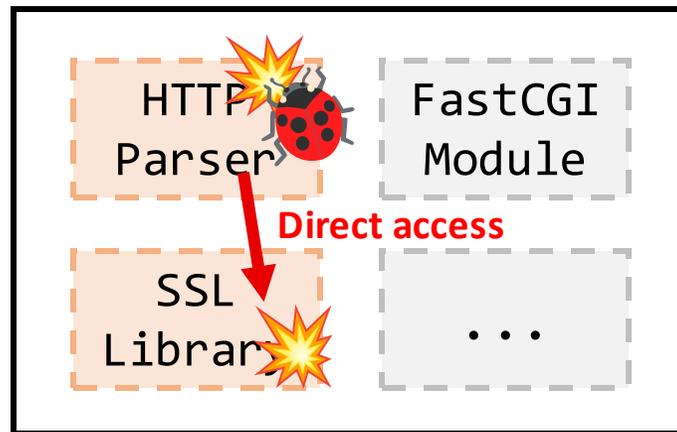
Monolithic

Example: a web server (Apache)

Compartmentalization

Software compartmentalization =

- design programs **split into distrusting and isolated components**
- such that if one of the components is breached, the attacker **does not compromise the whole program**



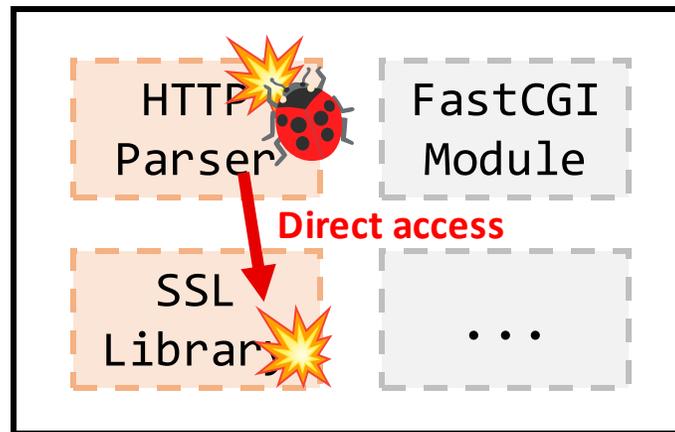
Monolithic

Example: a web server (Apache)

Compartmentalization

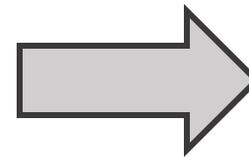
Software compartmentalization =

- design programs **split into distrusting and isolated components**
- such that if one of the components is breached, the attacker **does not compromise the whole program**

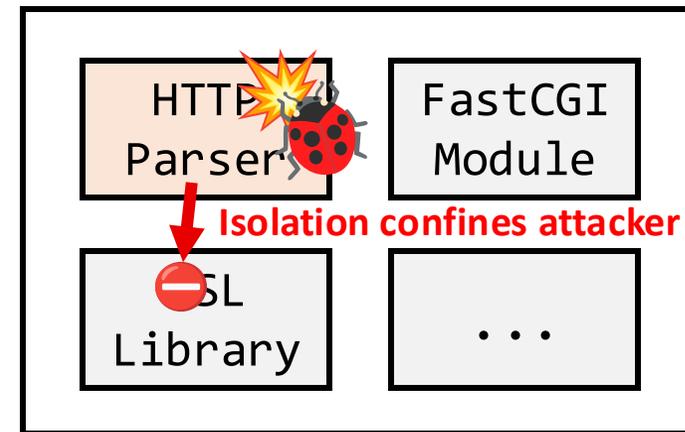


Monolithic

Example: a web server (Apache)



Compartmentalization



Compartmentalized

 Isolation boundary

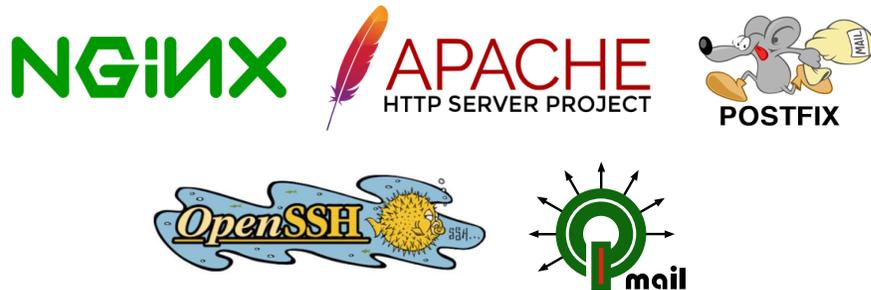
Compartmentalization

Software compartmentalization =

- design programs **split into distrusting and isolated components**
- such that if one of the components is breached, the attacker **does not compromise the whole program**

Compartmentalization is **applied to high-profile software**

Server Software



Web Browsers



But also...



Compartmentalization

£70M

Area of investment and support

Digital security by design

The digital security by design (DSbD) challenge funds business and researchers to update the foundation of the insecure digital computing infrastructure by creating a new, more secure hardware and software ecosystem.

Budget: £70 million
Duration: From 2020 to 2025

Partners involved: Innovate UK, Engineering and Physical Sciences Research Council (EPSRC)

Significant research funding in recent years

isolated components
When combined, the attacker does

£12M

Funding opportunity

Contracts for innovation: DSbD Advancing CHERI Tools and software

Opportunity status:	Closed
Funders:	Innovate UK
Co-funders:	Department for Science, Innovation and Technology
Funding type:	Other
Publication date:	19 May 2025
Opening date:	14 May 2025 9:30am UK time
Closing date:	18 June 2025 11:00am UK time

See the [full opportunity details on the Innovation Funding Service](#).
Organisations can apply for a share of up to £12 million, inclusive of VAT, to work on maturing and enabling the availability of CHERI Tools and Software components for RISC-V embedded devices that implement the CHERI architecture extensions.

\$50M

INDUSTRY EMERGING INNOVATION INDUSTRY NEWS

DARPA to Hold Compartmentalization and Privilege Management (CPM) Proposers Day

Full solutions are anticipated to require new hardware support, though software-only solutions that can meet performance expectations are in-scope for the program.

 By Homeland Security Today April 1, 2023

We know from this experience that **compartmentalization works**

We know from this experience that **compartmentalization works**

Re: double-free vulnerability in OpenSSH server 9.1 (CVE-2023-25136)

(1)

From: Qualys Security Advisory <qsa () qualys com>
Date: Mon, 13 Feb 2023 12:02:13 +0000

Hi all,

On Thu, Feb 02, 2023 at 01:02:04PM +0000, Qualys Security Advisory wrote:
Exploiting this vulnerability will not be easy: modern memory allocators provide protections against double frees, and the impacted sshd process is unprivileged and heavily sandboxed.

Quick update: we were able to gain arbitrary control of the "rip" register through this bug (i.e., we can jump wherever we want in sshd's address space) on an unpatched installation of OpenBSD 7.2 (which runs OpenSSH 9.1 by default). This is by no means the end of the story: this was only step 1, bypass the malloc and double-free protections. The next steps, which may or may not be feasible at all, are:

- step 2, execute arbitrary code despite the ASLR, NX, and ROP protections (this will probably require an information leak, either through the same bug or through a secondary bug);
- step 3, escape from sshd's sandbox (through a secondary bug, either in the privileged parent process or in the kernel's reduced attack surface).



- OpenSSH is **vulnerable to a bug with remote-code execution**
- It is exploitable: attackers can **break allocator-based mitigations**
- Compartmentalization ultimately **mitigates the exploit** by containing it in an unprivileged process

(1) <https://seclists.org/oss-sec/2023/q1/92>

We know from this experience that **compartmentalization works**

Re: double-free vulnerability in OpenSSH server 9.1 (CVE-2023-25136)

(1)

From: Qualys Security Advisory <qsa () qualys com>
Date: Mon, 13 Feb 2023 12:02:13 +0000

Hi all,

On Thu, Feb 02, 2023 at 01:02:04PM +0000, Qualys Security Advisory wrote:
Exploiting this vulnerability will not be easy: modern memory allocators provide protections against double frees, and the impacted sshd process is unprivileged and heavily sandboxed.

Quick update: we were able to gain arbitrary control of the "rip" register through this bug (i.e., we can jump wherever we want in sshd's address space) on an unpatched installation of OpenBSD 7.2 (which runs OpenSSH 9.1 by default). This is by no means the end of the story: this was only step 1, bypass the malloc and double-free protections. The next steps, which may or may not be feasible at all, are:

- step 2, execute arbitrary code despite the ASLR, NX, and ROP protections (this will probably require an information leak, either through the same bug or through a secondary bug);
- step 3, escape from sshd's sandbox (through a secondary bug, either in the privileged parent process or in the kernel's reduced attack surface).



- OpenSSH is **vulnerable to a bug with remote-code execution**
- It is exploitable: attackers can **break allocator-based mitigations**
- Compartmentalization ultimately **mitigates the exploit** by containing it in an unprivileged process

(1) <https://seclists.org/oss-sec/2023/q1/92>

We know from this experience that **compartmentalization works**

Re: double-free vulnerability in OpenSSH server 9.1 (CVE-2023-25136) (1)

From: Qualys Security Advisory <qsa () qualys com>
Date: Mon, 13 Feb 2023 12:02:13 +0000

Hi all,

On Thu, Feb 02, 2023 at 01:02:04PM +0000, Qualys Security Advisory wrote:
Exploiting this vulnerability will not be easy: modern memory allocators provide protections against double frees, and the impacted sshd process is unprivileged and heavily sandboxed.

Quick update: we were able to gain arbitrary control of the "rip" register through this bug (i.e., we can jump wherever we want in sshd's address space) on an unpatched installation of OpenBSD 7.2 (which runs OpenSSH 9.1 by default). This is by no means the end of the story: this was only step 1, bypass the malloc and double-free protections. The next steps, which may or may not be feasible at all, are:

- step 2, execute arbitrary code despite the ASLR, NX, and ROP protections (this will probably require an information leak, either through the same bug or through a secondary bug);
- step 3, escape from sshd's sandbox (through a secondary bug, either in the privileged parent process or in the kernel's reduced attack surface).



- OpenSSH is **vulnerable to a bug with remote-code execution**
- It is exploitable: attackers can **break allocator-based mitigations**
- Compartmentalization ultimately **mitigates the exploit** by containing it in an unprivileged process

(1) <https://seclists.org/oss-sec/2023/q1/92>

Yet, in 2025, **compartmentalization is not a widespread practice**

Yet, in 2025, compartmentalization is not a widespread practice

Compartmentalization

Software compartmentalization =

- design programs **split into distrustful and isolated components**
- such that if one of the components is compromised, the attacker **does not own the full program**

Compartmentalization is applied to high-profile software

Server Software	Web Browsers	But also...
  	  	 
 		

22

Beyond the popular software I mentioned earlier, **compartmentalization is rather rare**

Yet, in 2025, compartmentalization is not a widespread practice

Compartmentalization

Software compartmentalization =

- design programs split into distrustful and isolated components
- such that if one of the components is compromised, the attacker does not own the full program

Compartmentalization is applied to high-profile software

Server Software Web Browsers But also...



22

Beyond the popular software I mentioned earlier, compartmentalization is rather rare

"Less than 56 apps out of the 1,520 most popular Debian applications are compartmentalized"

(1) Lefeuvre et al., SoK: Software Compartmentalization, S&P 2025



Why?

Why?

SoK: Software Compartmentalization

Hugo Lefeuvre¹, Nathan Dautenhahn¹, David Chisnall², Pierre Olivier³
¹The University of British Columbia, ²Serenity, ³SCI Semiconductor, ⁴The University of Manchester

Abstract—Decomposing large systems into smaller components with limited privileges has long been recognized as an effective means to minimize the impact of exploits. Despite historical roots, demonstrated benefits, and a plethora of research efforts in academia and industry, the compartmentalization of software is still not a mainstream practice. This paper investigates why, and how this status quo can be improved. Noting that existing approaches are fraught with inconsistencies in terminology and analytical methods, we propose a unified model for the systematic analysis, comparison, and directing of compartmentalization approaches. We use this model to review 211 research efforts and analyze 61 mainstream compartmentalized systems, confronting them to understand the limitations of both research and production works. Among others, our findings reveal that mainstream efforts largely rely on manual methods, custom abstractions, and legacy mechanisms, poles apart from recent research. We conclude with recommendations: compartmentalization should be solved holistically; progress is needed towards simplifying the definition of compartmentalization; towards better challenging our threat models in the light of confused deputies and hardware limitations; as well as current bridging the gaps we not only maps the historical and mainstream needs. This paper not only maps the historical and current landscape of compartmentalization, but also sets forth a framework to foster their evolution and adoption.

1. Introduction

Despite decades of effort, vulnerabilities still plague software, and thwarting them remains a game of cat and mouse. The *principle of least privilege* (PoLP) [209] is the last line of defense when protections fail or when flaws are unknown. By granting each entity only the privileges needed, the PoLP ensures that a compromise of one part will not imply that the whole. *Software compartmentalization* is a prominent implementation of the PoLP, in which developers divide a large program into smaller, lesser privileged components, to reduce the impact of potential security breaches. Software compartmentalization inherits from a large body of work, starting with processes [93]; including OS models and separation kernels [56, 58, 205, 206, 247], or capability OSes [83, 125]; all the way to fine-grain application compartmentalization in the 2000s [147, 181, 232, 238]

¹This work was primarily done while Hugo Lefeuvre was affiliated with the University of Manchester and SCI Semiconductor, and Nathan Dautenhahn with Rice University (with minor contributions at Riverside Research).

following qmail [68], Postfix [121], or OpenSSH [198]. Its premises are plenty: containing memory safety issues [72, 202], untrusted third parties [54, 181], or unsafe parts of safe languages [62, 149, 174, 201, 228], isolating cryptographic secrets [167, 171, 232] or shadow stacks [74], thwarting supply-chain attacks [111, 235] and side-channels [137, 176, 182], or providing fault resilience [164, 184].

Despite longstanding recognition within the academic sphere and proven effectiveness in seminal industry projects, the adoption of compartmentalization techniques in mainstream software remains inconsistent: compartmentalizing is still far from being a common engineering practice. Take long advocated by the community [65, 70, 73, 87, 115–117, 140, 158, 165, 167, 211, 226, 232] but without viable adoption by leading cryptography libraries. At a time of growing threats, we are missing out on the security benefits compartmentalization can bring. This paper investigates the reasons behind this status quo, and how to improve it.

Research speculated that this situation is due to a lack of automation [65, 253], limitations of mechanisms [190, 259], excessive performance overheads [188, 232], a lack of strong security guarantees [129, 160], among others. All are likely part of the problem, and through decades of research progress was made on every front. Today the community lacks a global overview of this progress. Designing for or retrofitting compartmentalization is often hampered by inconsistencies in the understanding and application of its concepts. Existing models and terminologies are numerous, ad-hoc, and often contradictory, leading to confusion and a growing body of work that cannot be compared. The lack of a systematic perspective leads to a mismatch between what software compartmentalization needs to progress towards the mainstream, and the focus and framing of research efforts: most do not tackle compartmentalization's key aspects as a whole and thus produce solutions that cannot be relevant to making it a mainstream practice.

Recognizing these challenges, we propose a unified model providing a consistent framework for defining, understanding, and implementing compartmentalization. This model comes with a comprehensive taxonomy that allows to classify compartmentalization strategies based on their policy definition methods, abstractions, and mechanisms, providing a basis for systematic evaluation and comparison.

We validate our model and taxonomy by systematizing 211 research and 61 mainstream software systems implementing compartmentalization. Doing so, this study provides unique insights into where mainstream efforts have

(1)

Systematizing Compartmentalization

Here is how we approached the question:

- Started by making a fundamental observation: experts do not even agree on what compartmentalization is
 - Define a **fundamental model** and **consistent terminology** for compartmentalization
- Propose a **taxonomy** for evaluating compartmentalization approaches and **apply it systematically** to a wide set (200+) of *research* and *mainstream* efforts
- From this: **extract a set of core challenges**

Systematizing Compartmentalization

Here is how we approached the question:

- Started by making a fundamental observation: experts do not even agree on what compartmentalization is
 - Define a **fundamental model** and **consistent terminology** for compartmentalization
- Propose a **taxonomy** for evaluating compartmentalization approaches and **apply it systematically** to a wide set (200+) of *research* and *mainstream* efforts
- From this: **extract a set of core challenges**

Systematizing Compartmentalization

Here is how we approached the question:

- Started by making a fundamental observation: experts do not even agree on what compartmentalization is
 - Define a **fundamental model** and **consistent terminology** for compartmentalization
- Propose a **taxonomy** for evaluating compartmentalization approaches and **apply it systematically** to a wide set (200+) of *research* and *mainstream* efforts
- From this: **extract a set of core challenges**

Systematizing Compartmentalization

Here is how we approached the question:

- Started by making a fundamental observation: experts do not even agree on what compartmentalization is
 - Define a **fundamental model** and **consistent terminology** for compartmentalization
- Propose a **taxonomy** for evaluating compartmentalization approaches and **apply it systematically** to a wide set (200+) of *research* and *mainstream* efforts
- From this: **extract a set of core challenges**

My promise for this talk

*A journey through twenty years of
compartmentalization*

1. What is software compartmentalization? (slightly more formal)
2. A systematic perspective on compartmentalization
3. The *why*: compartmentalization everywhere, what will it take?

My promise for this talk

*A journey through twenty years of
compartmentalization*

1. What is software compartmentalization? (slightly more formal)
2. A systematic perspective on compartmentalization
3. The *why*: compartmentalization everywhere, what will it take?



What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice**

What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice** where developers **break down a program into groups of isolated/distrusting components**

What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice** where developers **break down a program into groups of isolated/distrusting components**, each controlling only the resources they precisely require

What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice** where developers **break down a program into groups of isolated/distrusting components**, each controlling only the resources they precisely require
- an embodiment of the **principle of least-privilege**

What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice** where developers **break down a program into groups of isolated/distrusting components**, each controlling only the resources they precisely require
- an embodiment of the **principle of least-privilege**

Fundamental idea:

- if a component is compromised, attackers are **restricted to the permissions of the compromised component**

What is Software Compartmentalization?

Software compartmentalization is...

- a **software engineering practice** where developers **break down a program into groups of isolated/distrusting components**, each controlling only the resources they precisely require
- an embodiment of the **principle of least-privilege**

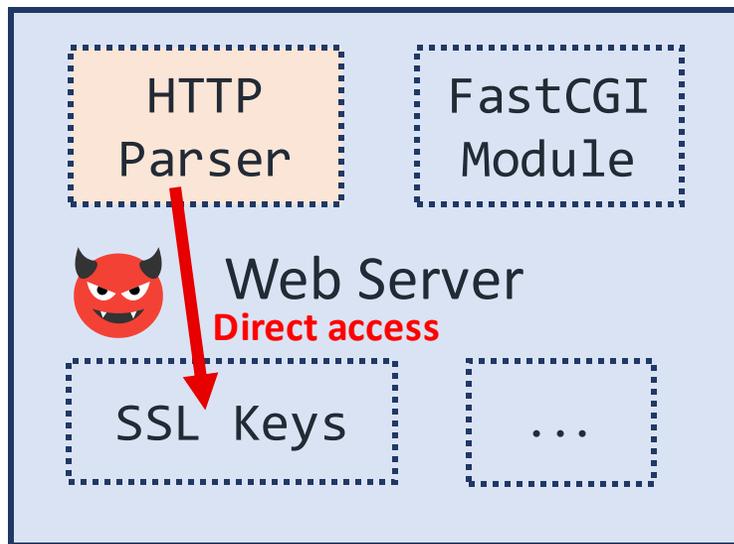
Fundamental idea:

- if a component is compromised, attackers are **restricted to the permissions of the compromised component**
- to escalate privileges and compromise the rest of the system, attackers must **find and exploit additional software vulnerabilities**

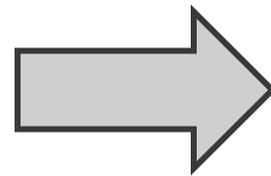
What is Software Compartmentalization?

Going back to our previous example...

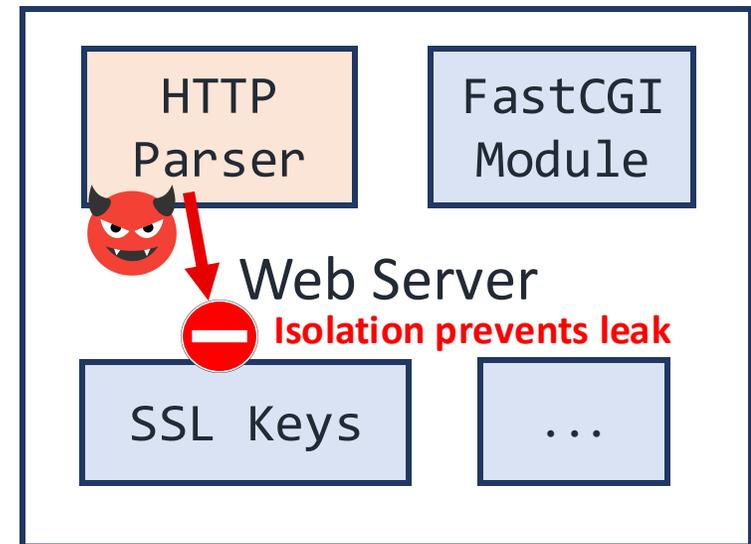
 Isolation Boundary



Monolithic



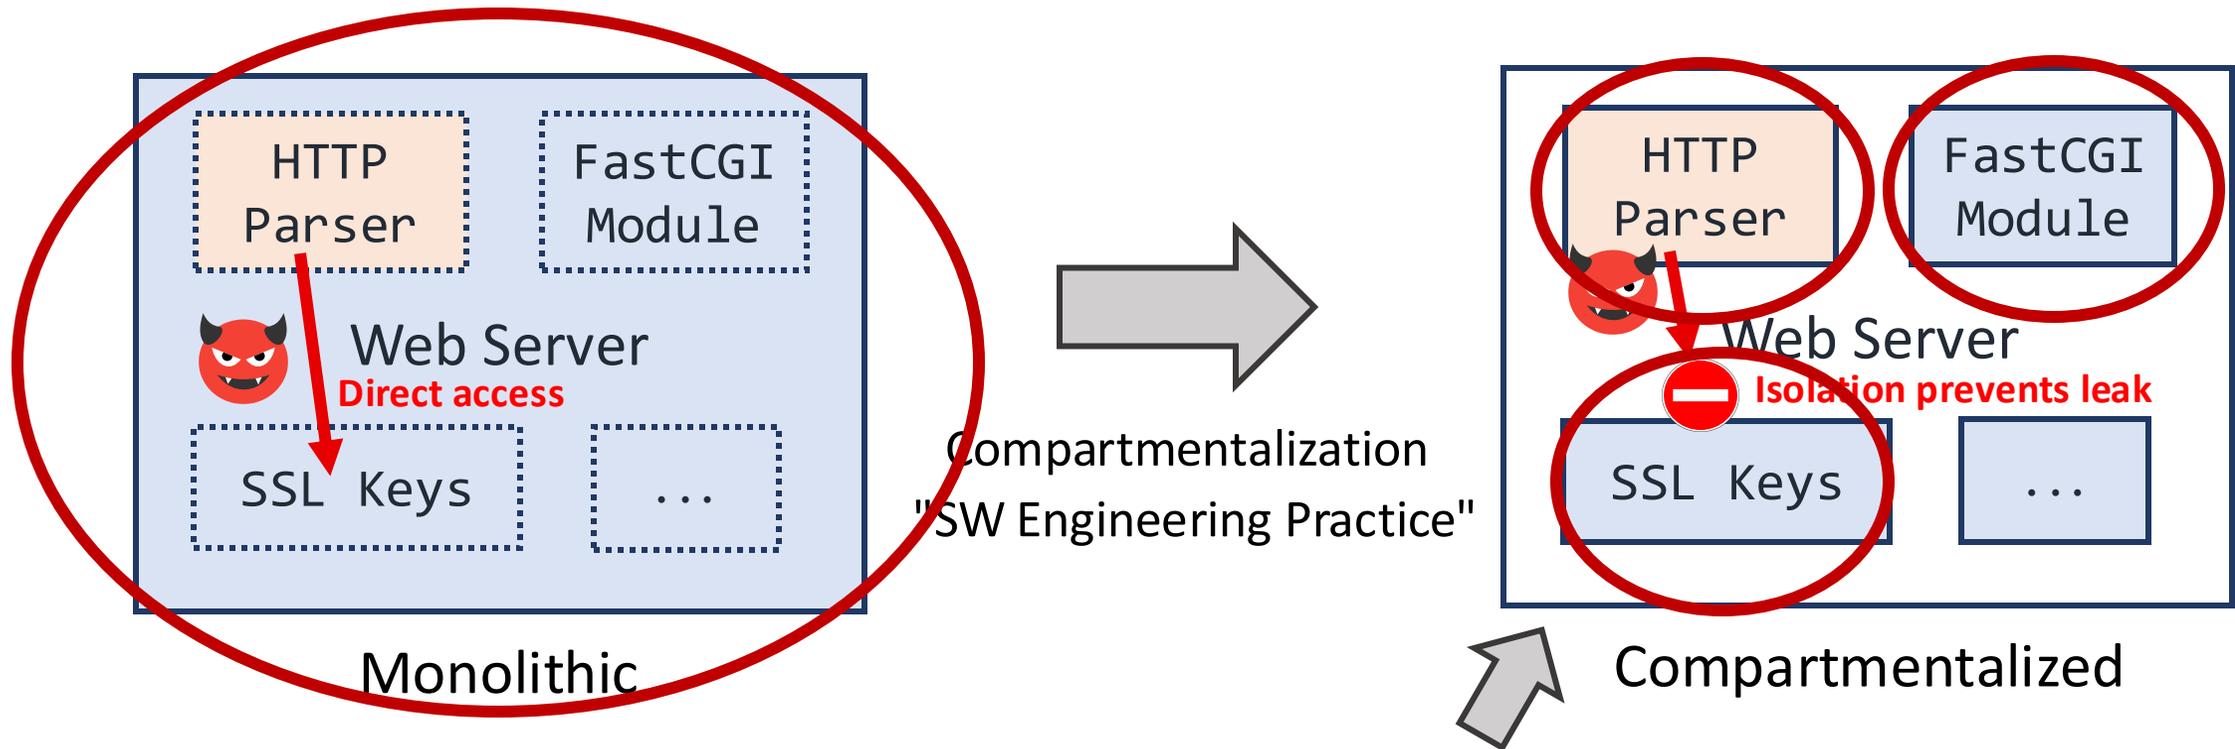
Compartmentalization
"SW Engineering Practice"



Compartmentalized

What is Software Compartmentalization?

Going back to our previous example...

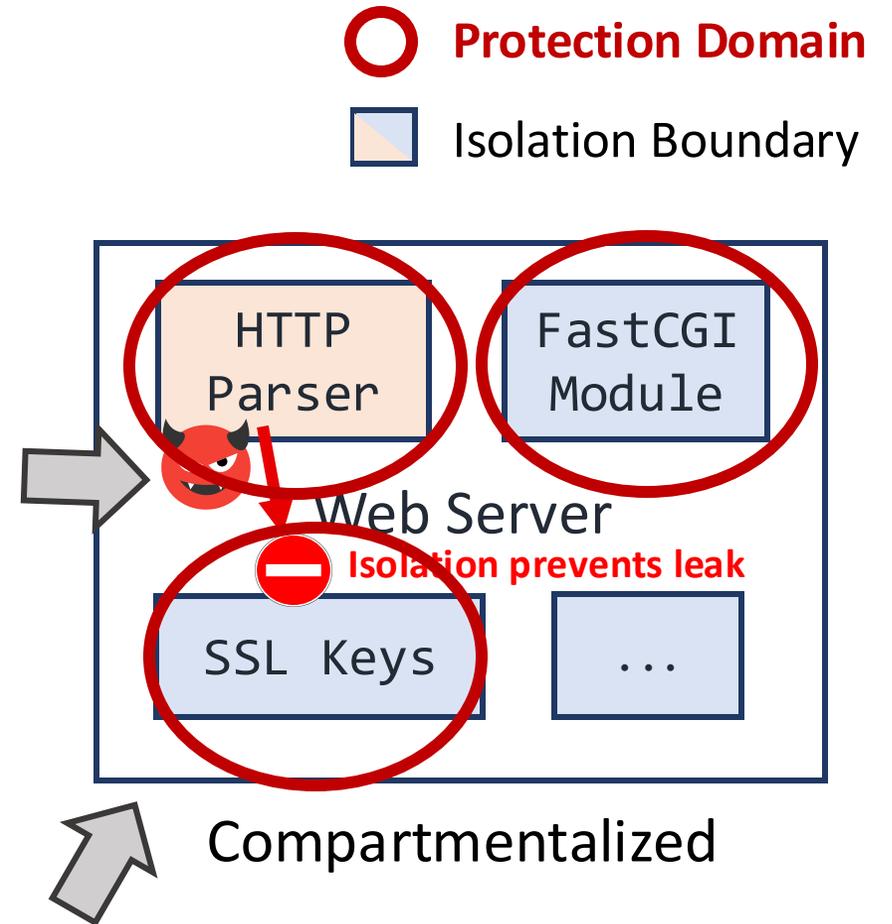


Components now **distrust** each other

What is Software Compartmentalization?

Going back to our previous example...

Compartmentalizing **adds** a (potentially very tricky) **step on the attacker's path** towards taking over the program



Components now **distrust each other**

What is Software Compartmentalization?

Software compartmentalization is applicable to any kind of program

Web Browsers



Server Software



NGINX



Microkernels



What is Software Compartmentalization?

Software compartmentalization is applicable to any kind of program

Web Browsers



Server Software



NGINX



Microkernels



Applications

What is Software Compartmentalization?

Software compartmentalization is applicable to any kind of program

Web Browsers



Server Software



NGINX



Microkernels



OS kernels

What is Software Compartmentalization?

Software compartmentalization is applicable to any kind of program

Web Browsers



Server Software



Microkernels



But also: hypervisors, firmware, etc.

What is Software Compartmentalization?

Software compartmentalization can be **applied to new as well as existing programs:**

- When applied to existing programs, we talk about **retrofitting**

What is Software Compartmentalization?

Software compartmentalization can be **applied to new as well as existing programs:**

- When applied to existing programs, we talk about **retrofitting**

Some programs which included compartmentalization in their design in the first place (non-exhaustive!):



*Mail transfer agent, one of the first compartmentalized programs to the best of my knowledge

What is Software Compartmentalization?

Software compartmentalization can be **applied to new as well as existing programs:**

- When applied to existing programs, we talk about **retrofitting**

Some programs which included compartmentalization in their design in the first place (non-exhaustive!):



...and some where it was retrofitted:



*Mail transfer agent, one of the first compartmentalized programs to the best of my knowledge

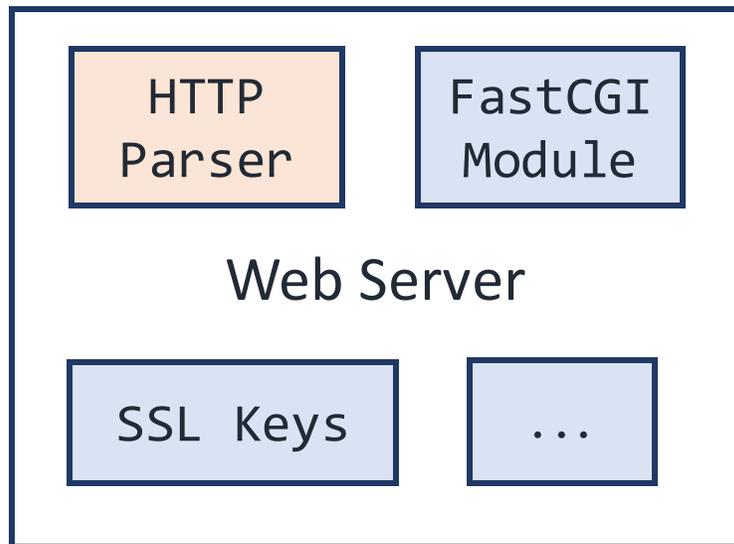
What is Software Compartmentalization?

Still, software compartmentalization as we refer to here applies **within one program, not across programs**

What is Software Compartmentalization?

Still, software compartmentalization as we refer to here applies **within one program, not across programs**

 Isolation Boundary

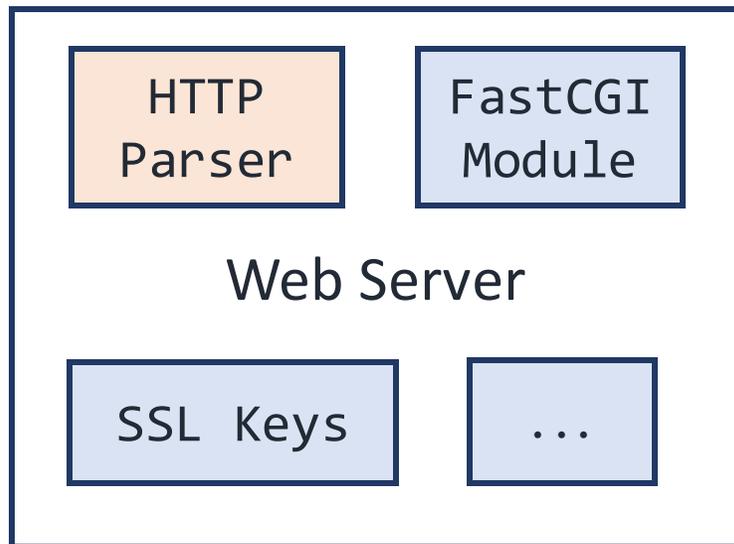


 Within one program
SW Compartmentalization

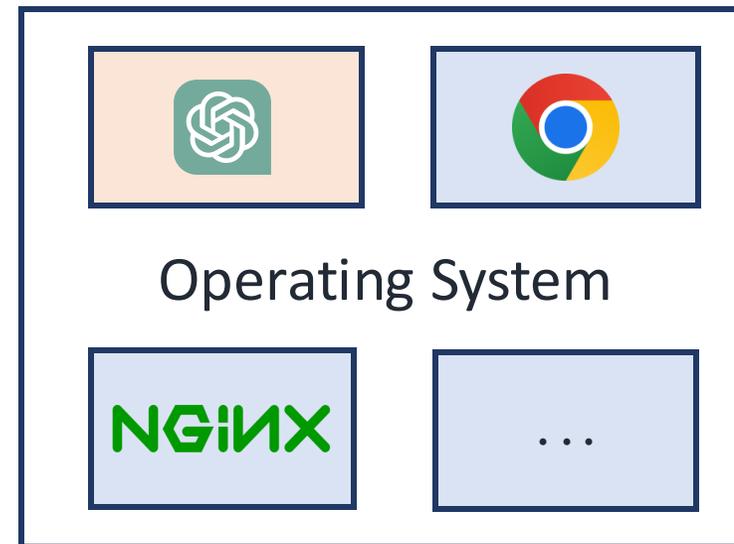
What is Software Compartmentalization?

Still, software compartmentalization as we refer to here applies **within one program, not across programs**

 Isolation Boundary



 Within one program
SW Compartmentalization

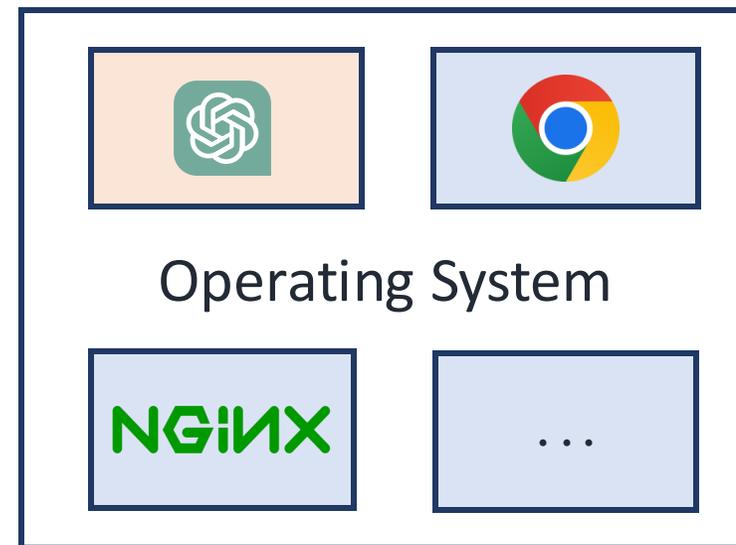
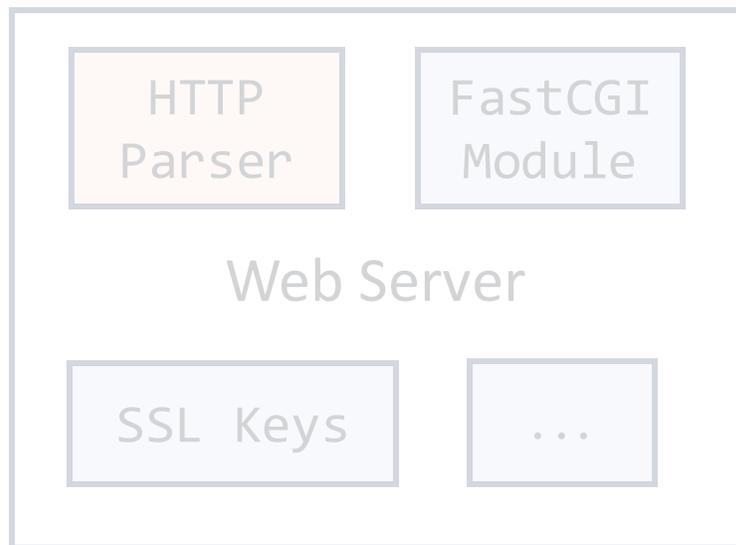


 Across programs
Not SW Compartmentalization

What is Software Compartmentalization?

Still, software compartmentalization as we refer to here applies **within one program**

More general isolation: share challenges and solutions with software compartmentalization as we define it



 Within one program
SW Compartmentalization

 Across programs
Not SW Compartmentalization

What is Software Compartmentalization?

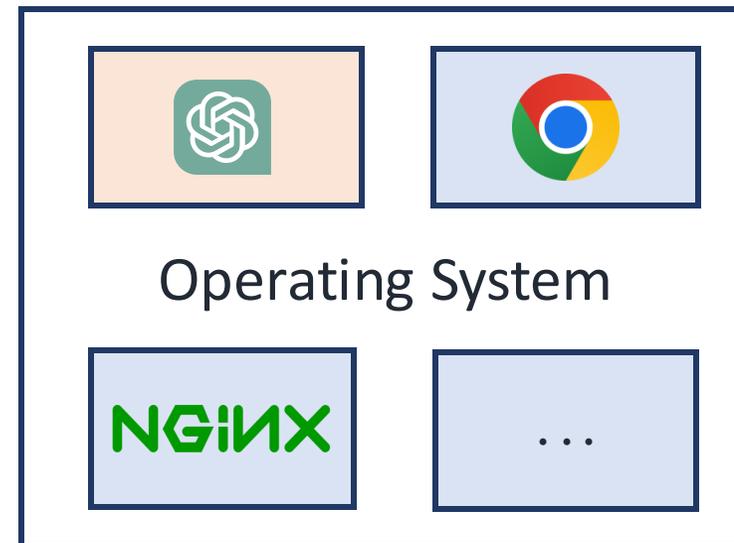
Still, software compartmentalization as we refer to here applies **within**

one program

More general isolation: share challenges and solutions with software compartmentalization as we define it



(1)



Across programs

Not SW Compartmentalization

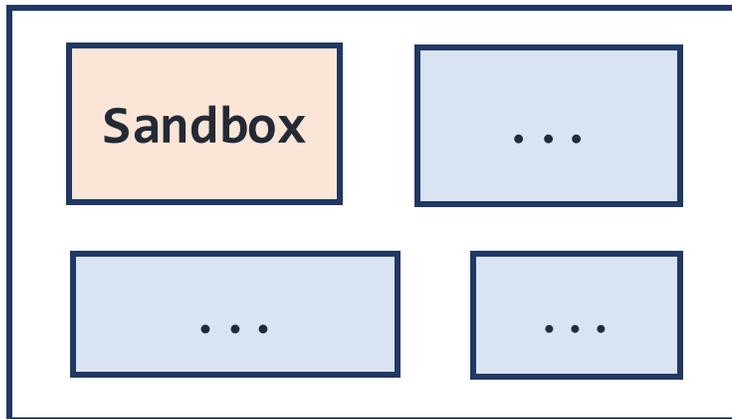
What is Software Compartmentalization?

Software compartmentalization can target **different trust models**

What is Software Compartmentalization?

Software compartmentalization can target **different trust models**

 Isolation Boundary

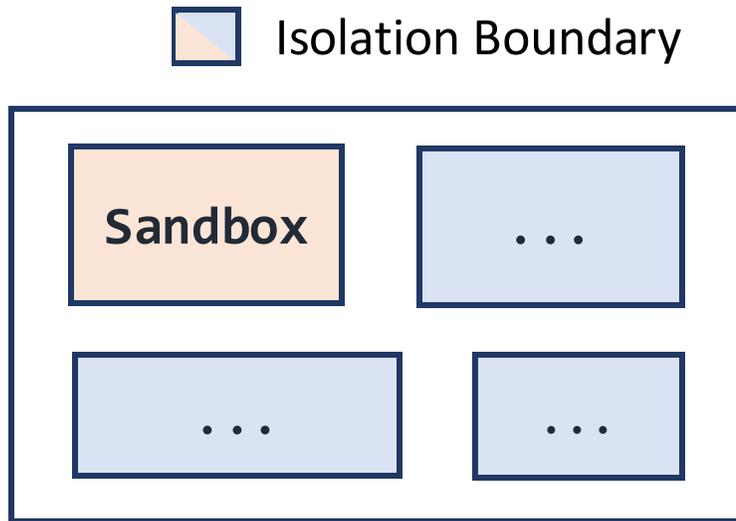


Sandbox: component isolated to protect the rest of the system

All but the sandbox is trusted

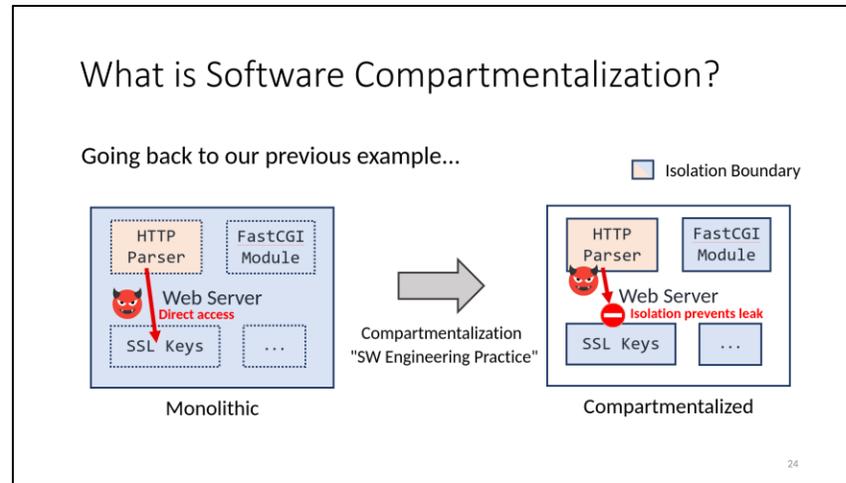
What is Software Compartmentalization?

Software compartmentalization can target **different trust models**



Sandbox: component isolated to protect the rest of the system

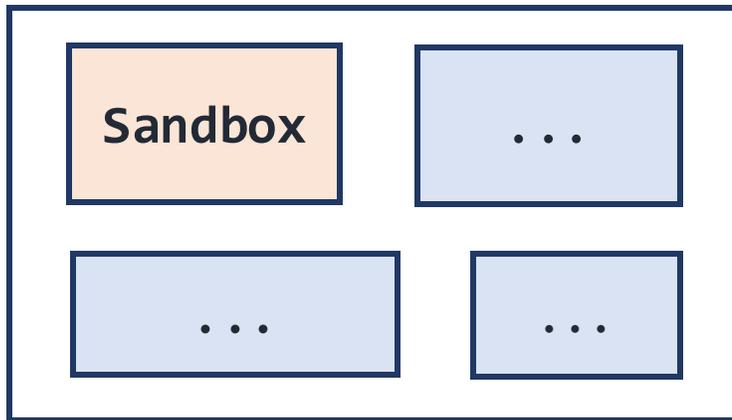
All but the sandbox is trusted



What is Software Compartmentalization?

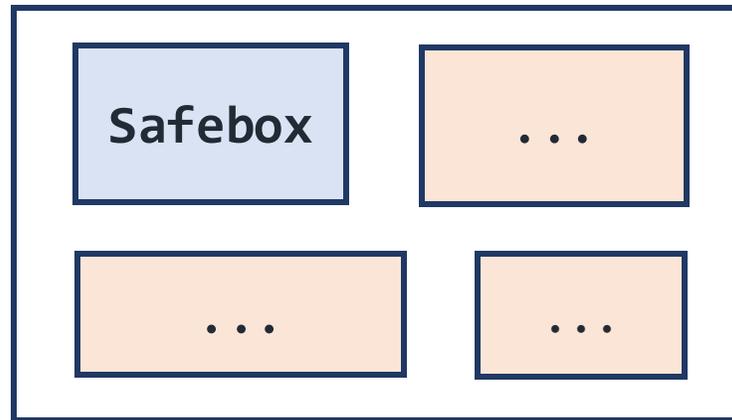
Software compartmentalization can target **different trust models**

 Isolation Boundary



Sandbox: component isolated to protect the rest of the system

All but the sandbox is trusted



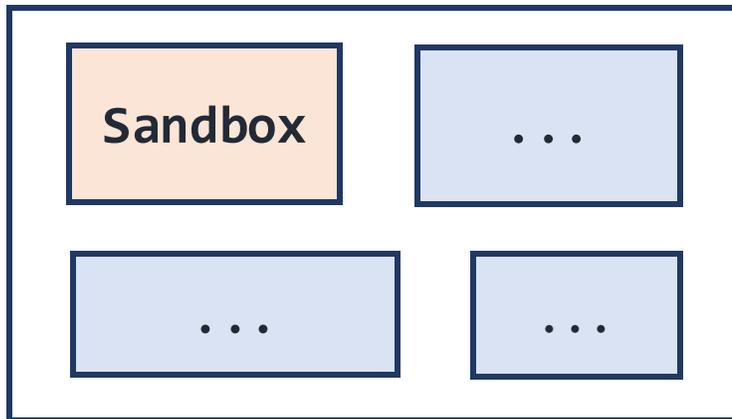
Safebox: component isolated to protect it from others

Only the safebox is trusted

What is Software Compartmentalization?

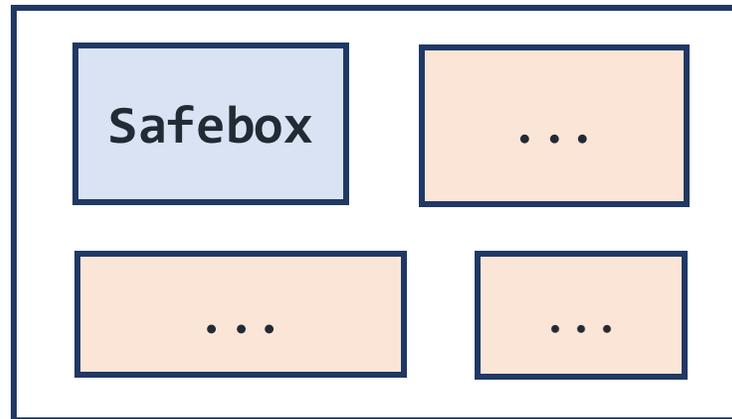
Software compartmentalization can target **different trust models**

 Isolation Boundary



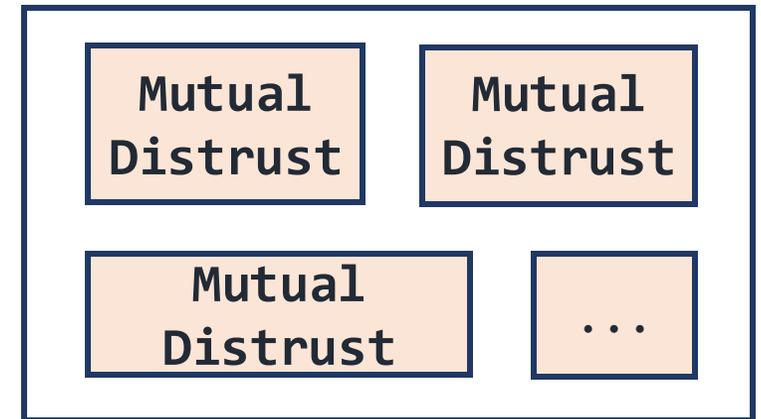
Sandbox: component isolated to protect the rest of the system

All but the sandbox is trusted



Safebox: component isolated to protect it from others

Only the safebox is trusted



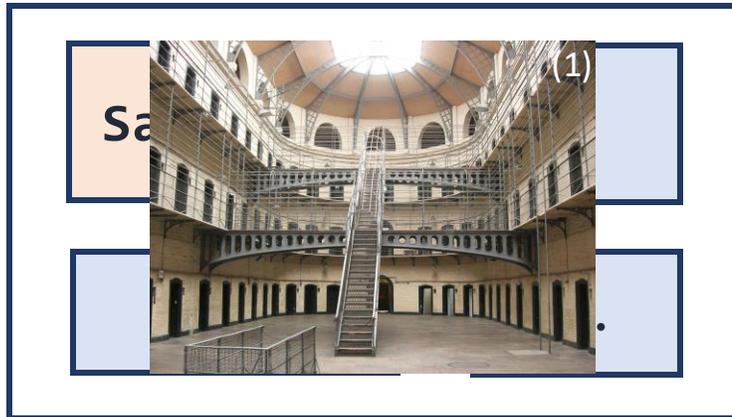
Mutual distrust: components distrust each other

What is Software Compartmentalization?

Software compartmentalization can target **different trust models**

 Isolation Boundary

These trust models are distinct and result in different software architectures



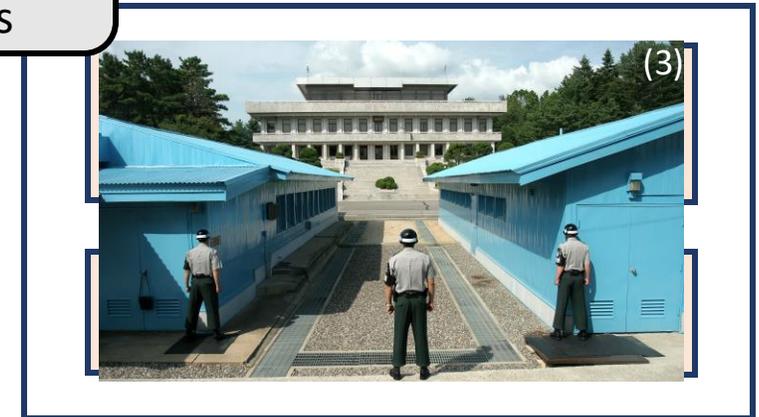
Sandbox: component isolated to protect the rest of the system

All but the sandbox is trusted



Safebox: component isolated to protect it from others

Only the safebox is trusted



Mutual distrust: components distrust each other

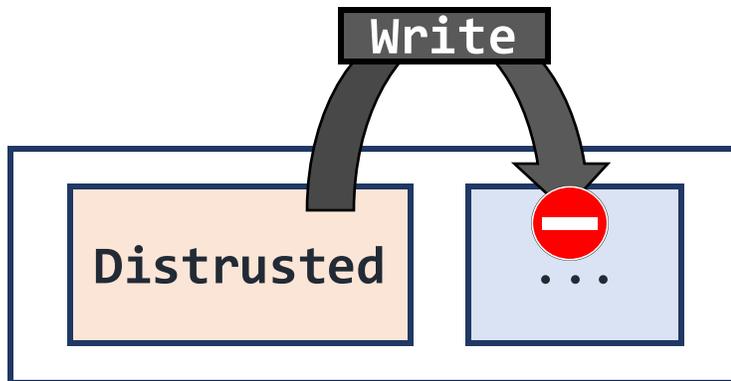
What is Software Compartmentalization?

Software compartmentalization can target **different properties**

What is Software Compartmentalization?

Software compartmentalization can target **different properties**

 Isolation Boundary

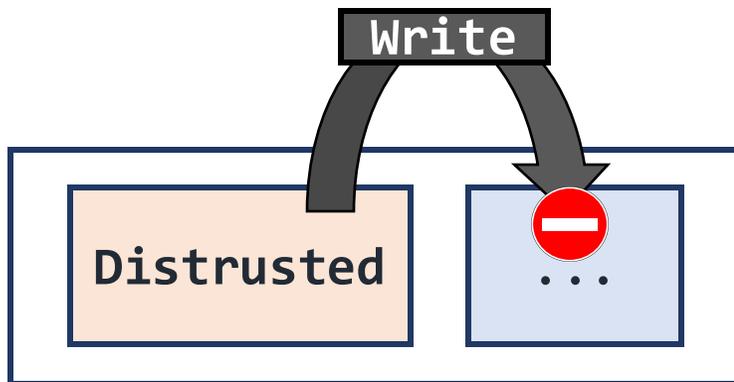


Integrity: compartment cannot alter other compartment's data

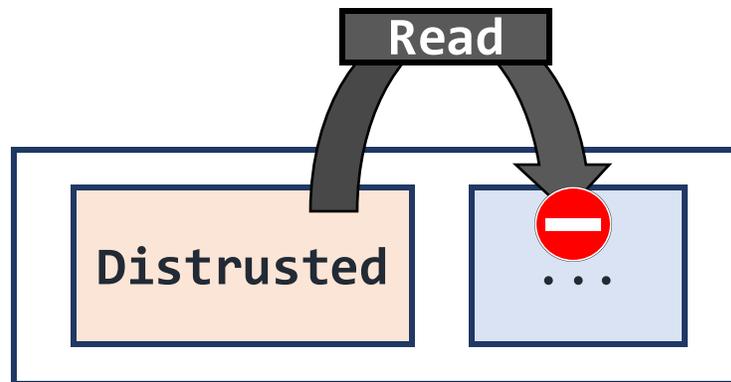
What is Software Compartmentalization?

Software compartmentalization can target **different properties**

 Isolation Boundary



Integrity: compartment cannot alter other compartment's data

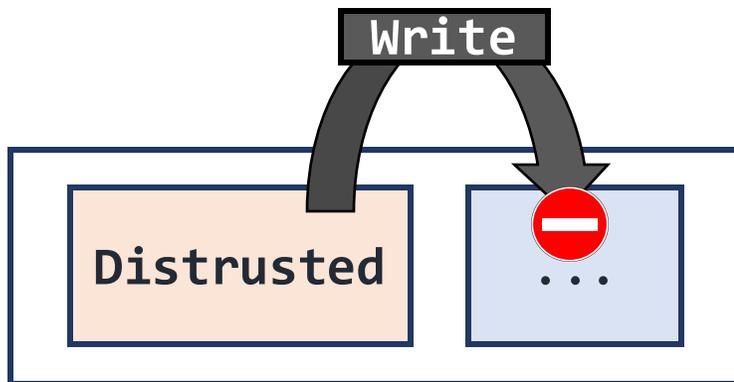


Confidentiality: compartment cannot read other compartment's data

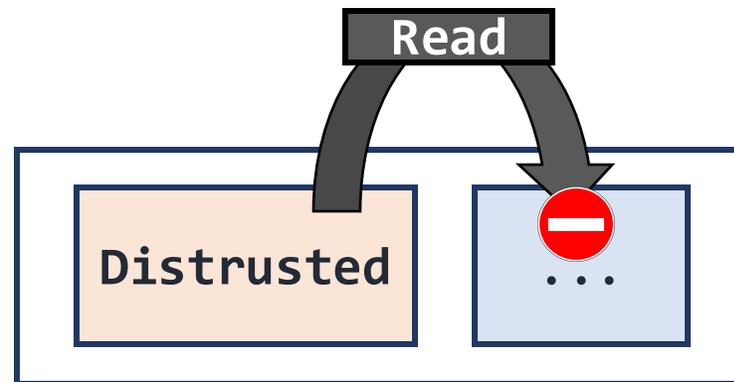
What is Software Compartmentalization?

Software compartmentalization can target **different properties**

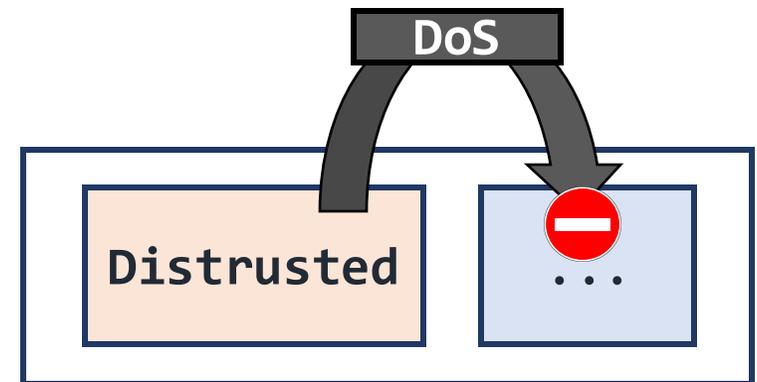
 Isolation Boundary



Integrity: compartment cannot alter other compartment's data



Confidentiality: compartment cannot read other compartment's data

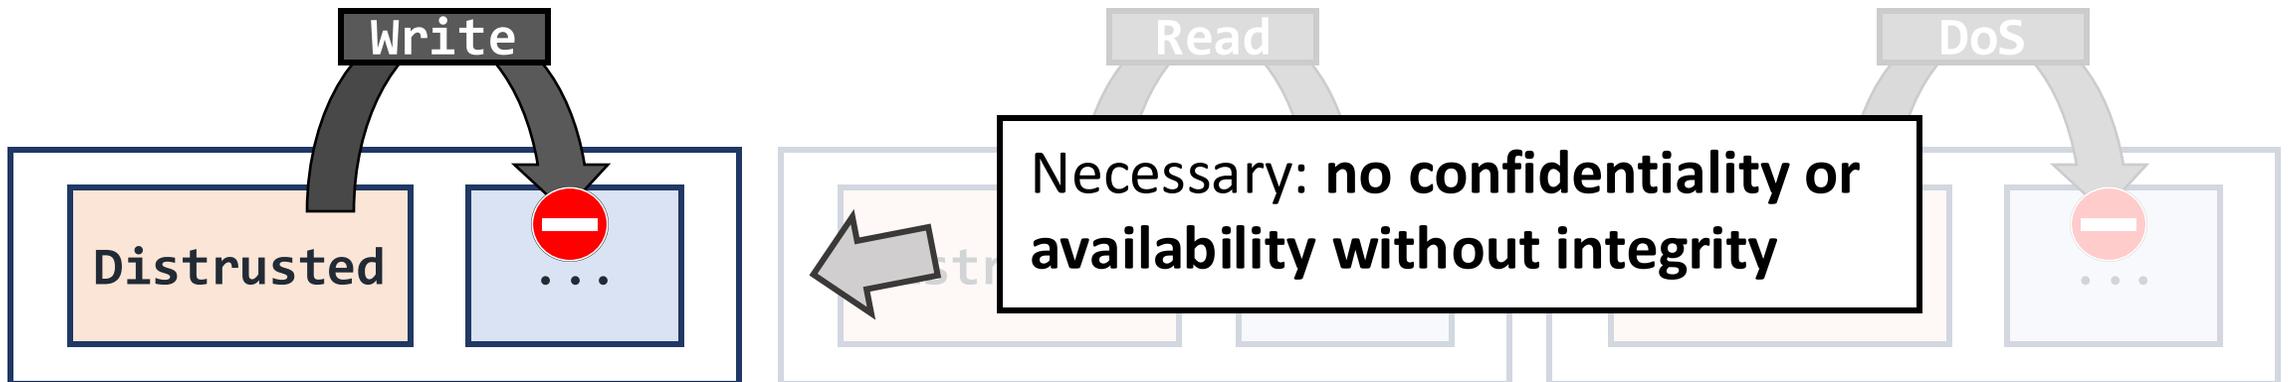


Availability: compartment cannot prevent other compartments from operating

What is Software Compartmentalization?

Software compartmentalization can target **different properties**

 Isolation Boundary



Integrity: compartment cannot alter other compartment's data

Confidentiality: compartment cannot read other compartment's data

Availability: compartment cannot prevent other compartments from operating

My promise for this talk

*A journey through twenty years of
compartmentalization*

- ~~1. What is software compartmentalization? (slightly more formal)~~
2. A systematic perspective on compartmentalization 
3. The *why*: compartmentalization everywhere, what will it take?

Take a step back: How do practitioners compartmentalize?
What does research say? What are open challenges?

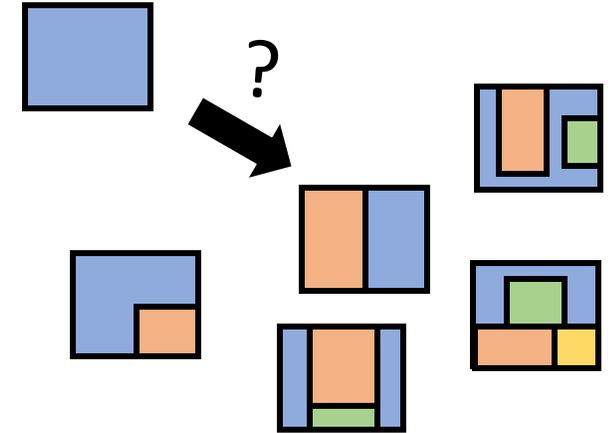
Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:

- 1. How to determine the right policy to enforce?**
 - Done with a *policy definition method*
- 2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?**
 - Done with a *compartmentalization abstraction*
- 3. How to enforce policies at runtime?**
 - Done with a *compartmentalization mechanism*

Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:



1. How to determine the right policy to enforce?

- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

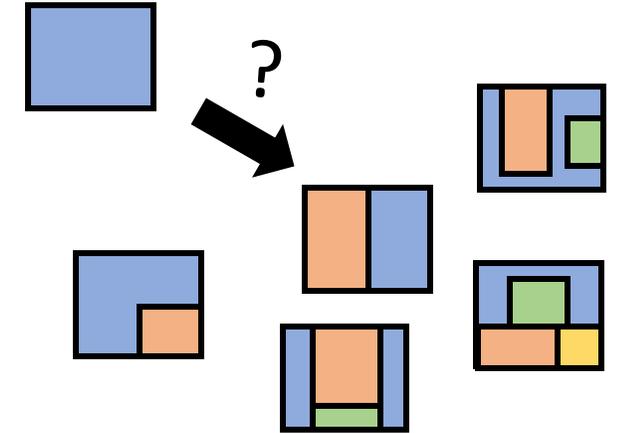
- Done with a *compartmentalization abstraction*

3. How to enforce policies at runtime?

- Done with a *compartmentalization mechanism*

Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:



1. How to determine the right policy to enforce?

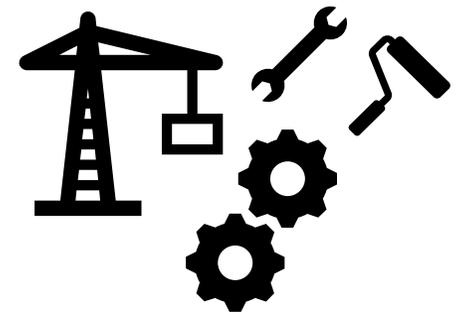
- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

- Done with a *compartmentalization abstraction*

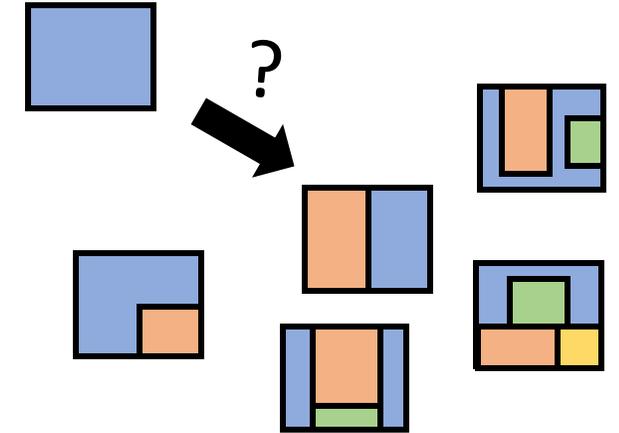
3. How to enforce policies at runtime?

- Done with a *compartmentalization mechanism*



Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:



1. How to determine the right policy to enforce?

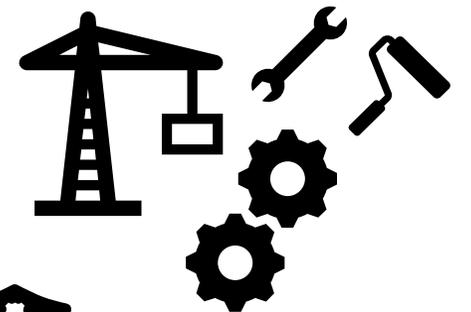
- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

- Done with a *compartmentalization abstraction*

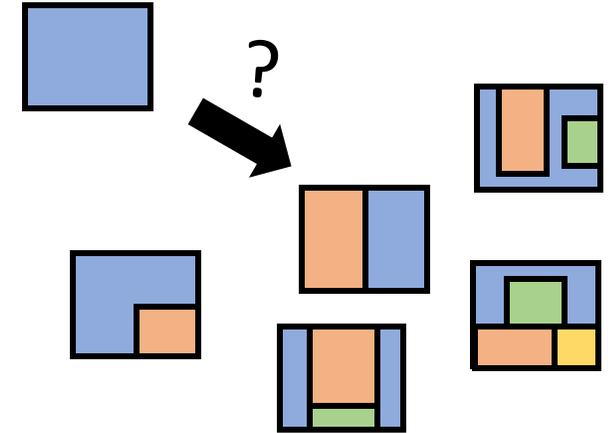
3. How to enforce policies at runtime?

- Done with a *compartmentalization mechanism*



Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:



1. How to determine the right policy to enforce?

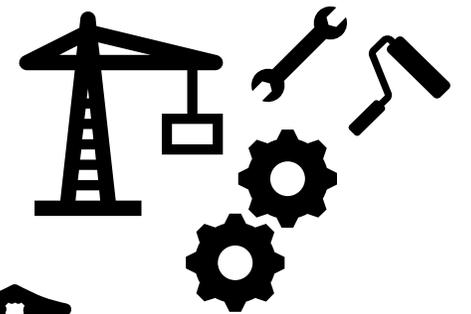
- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

- Done with a *compartmentalization abstraction*

3. How to enforce policies at runtime?

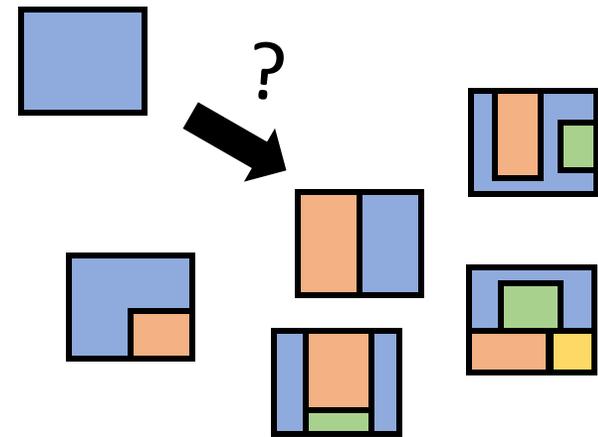
- Done with a *compartmentalization mechanism*



Problem #1

How to determine the right policy to enforce?

Developers must define **which components are to be separated** and **which properties should be enforced** on the resulting compartments

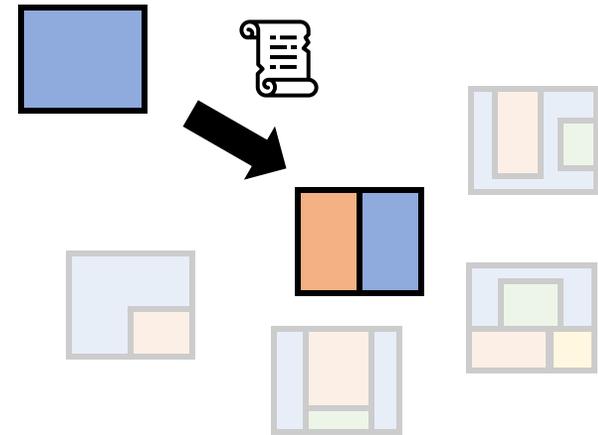


Problem #1

How to determine the right policy to enforce?

Developers must define **which components are to be separated** and **which properties should be enforced** on the resulting compartments

- The result is a **compartmentalization policy**



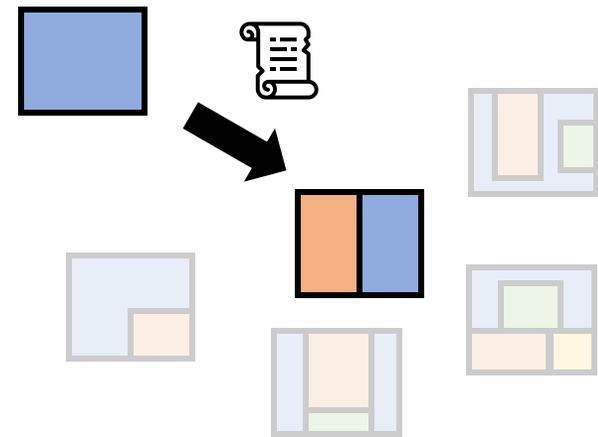
Problem #1

How to determine the right policy to enforce?

Developers must define **which components are to be separated** and **which properties should be enforced** on the resulting compartments

- The result is a **compartmentalization policy**

To design policies, developers employ a **Policy Definition Method**



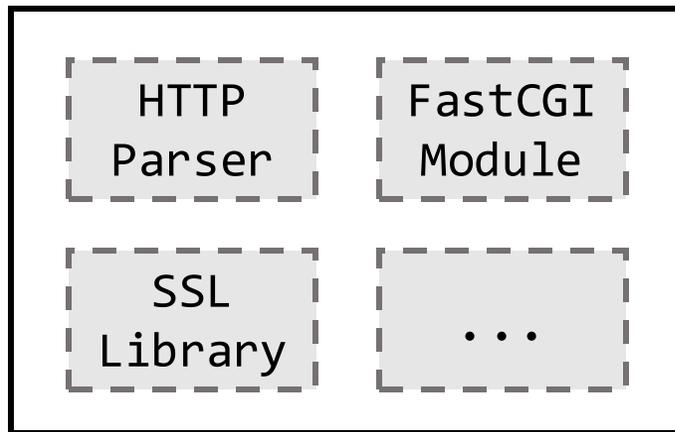
Policy Definition Methods

Historically, people have done this **manually**.

Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*

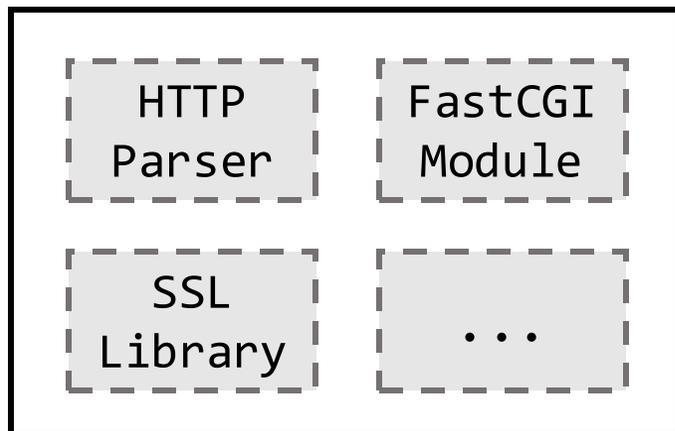


(Fictive monolithic program)

Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*



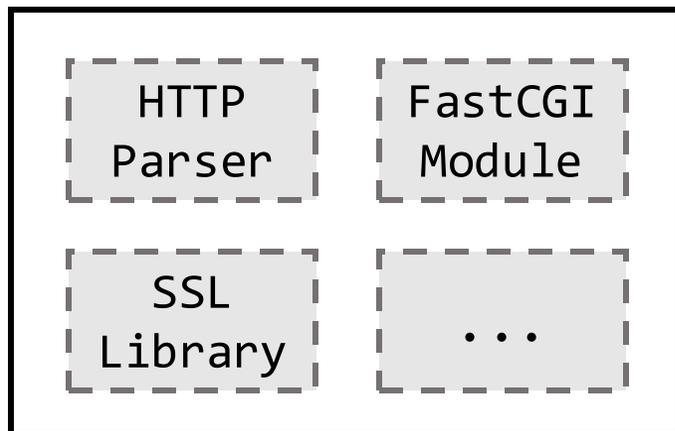
(Fictive monolithic program)

A developer looks attentively at the program...

Policy Definition Methods

Historically, people have done this **manually**.

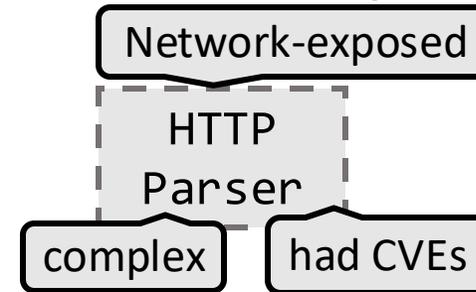
Example: *we want to split this C program.*



(Fictive monolithic program)

A developer looks attentively at the program...

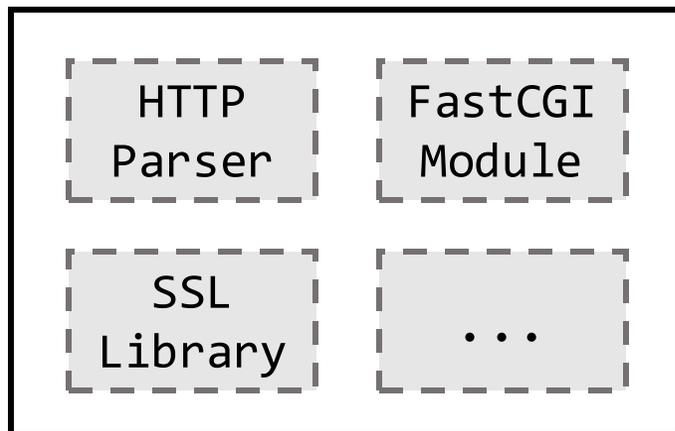
"I do not trust this component"



Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*



(Fictive monolithic program)

A developer looks attentively at the program...

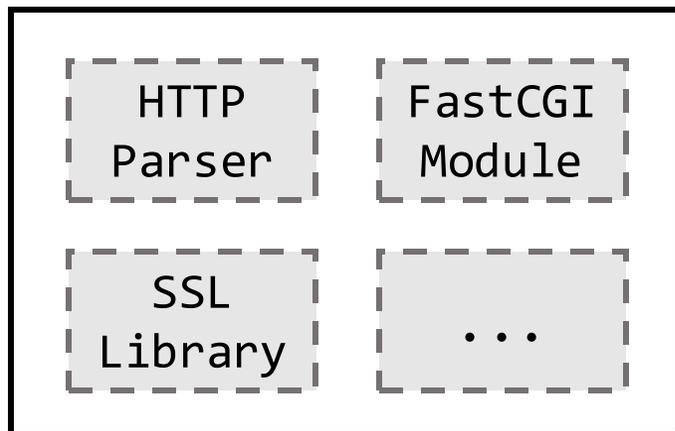
"I do not trust this component"



Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*

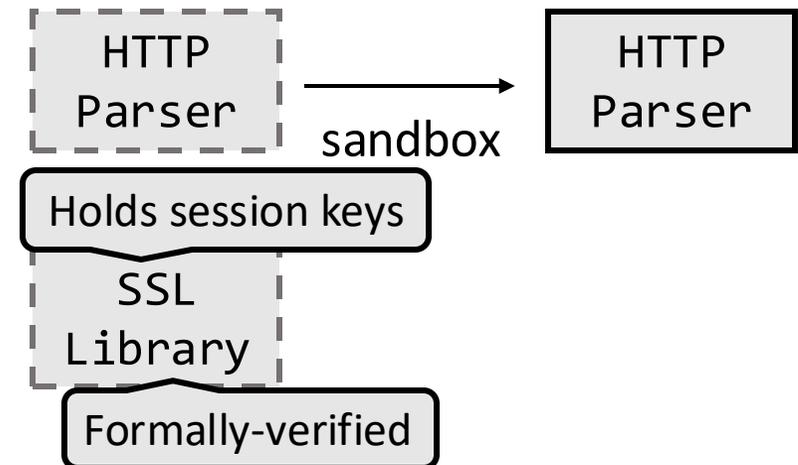


(Fictive monolithic program)

A developer looks attentively at the program...

"I do not trust this component"

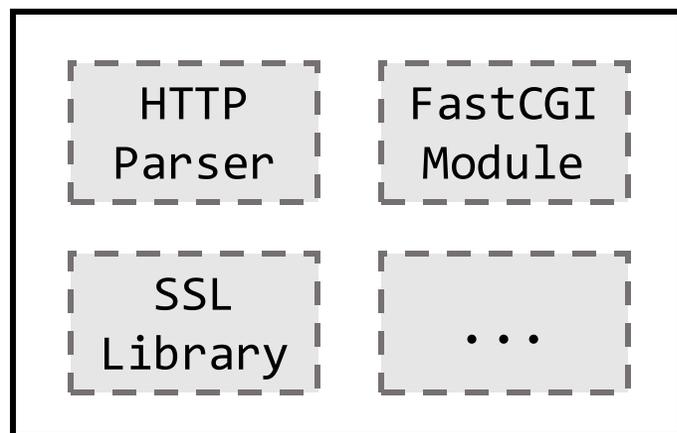
"I particularly value this component"



Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*



(Fictive monolithic program)

A developer looks attentively at the program...

"I do not trust this component"



"I particularly value this component"



This is still the most common way to do it today.

Systematizing Policy Definition Methods

Policy Definition Methods are a very active research area, with goals such as...

Minimize developer effort

Maximize soundness

Minimize over-privilege

Maximize performance

Maximize security of interfaces

We systematized these works in our SoK



(1)

TABLE 1: Taxonomy of Policy Definition Methods. PDMs that also propose an abstraction are marked with *. Manual (○) and fully automated (●) policies do not leverage a policy language, thus the column features Not/Applicable.

Separation Method	Automation ○ ●	Policy Language Type Annotations Placement Rules	Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation Performance Interface Safety
Manual [202, 1, ...]	○	N/A	N/A	Any	Manual	Any	○ N/A
Cowbar [76]	○	○	Function	Dynamic	Code-centric	○	○
MPD* [101, 102]	○	○	Component	Hybrid	Code-centric	○	○
CubicleOS* [111]	○	●	μLibrary	Static	Code-centric	●	○
Google SAPI* [23]	○	●	Function	Static	Code-centric	○	○
FlexOS* [143, 161]	○	●	μLibrary	Dynamic	Code-centric	●	○
ELIOS* [181]	○	○	Function	Static	Any	○	○
SOAAP* [111]	○	○	Any	Hybrid	Any	●	●
SeCaps* [171]	○	●	Function	Hybrid	Code-centric	●	○
PrivSplit* [169]	○	○	Function	Static	Code-centric	○	○
PrivTrans* [73]	○	○	Function	Hybrid	Code-centric	○	○
Glandring* [185]	○	○	Function	Static	Code-centric	○	○
Shield* [103], CAPSRY* [108]	○	○	Any	Static	Code-centric	○	○
DataShield* [77]	○	○	Any	Static	Code-centric	○	○
Swift* [101]	○	○	Any	Static	Code-centric	○	○
JP* [281]	○	○	Any	Static	Code-centric	○	○
PM [130]	○	○	Function	Hybrid	Code-centric	○	○
KSysP* [119]	○	○	Driver	Static	Code-centric	○	○
Call* [93]	○	○	Library	Static	Code-centric	○	○
CompartmentOS* [53]	○	○	Linkage Unit	Static	Code-centric	○	○
Enclousure* [111]	○	○	Package	Static	Code-centric	○	○
BreakApp* [135]	○	○	Package	Static	Code-centric	○	○
CompARTool* [113]	○	○	Library	Static	Code-centric	○	○
MCES* [101]	○	○	Function	Any	Code-centric	○	○
ProgramCutler [153]	○	N/A	N/A	Function	Dynamic	Code-centric	○
μSCOPE [202], SCALPOL [181]	○	N/A	N/A	Any	Dynamic	Code-centric	○

Systematizing Policy Definition Methods

Policy Definition Methods are a very active research area, with goals such as...

- Minimize developer effort
- Maximize soundness
- Minimize over-privilege

- Maximize performance
- Maximize security of interfaces

We systematized these works in our SoK



(1)

TABLE 1: Taxonomy of Policy Definition Methods. PDMs that also propose an abstraction are marked with *. Manual (○) and fully automated (●) policies do not leverage a policy language, thus the column features Not/Applicable.

Separation Method	Automation ○●	Policy Language Annotations	Type Placement Rules	Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation Performance	Interface Safety
Manual [28], [...]	○	○	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [76]	○	○	○	Function	Dynamic	Code-centric	○	○	○
MPD* [19], [92]	○	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [111]	○	●	●	μLibrary	Static	Code-centric	●	○	○
Google SAPI* [22]	○	●	●	Function	Static	Code-centric	○	○	○
FlexOS* [14], [81]	○	●	○	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	○	○	○	Function	Static	Any	○	○	○
SOAP* [117]	○	○	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	○	●	○	Function	Hybrid	Code-centric	●	○	○
PrivSplit* [169]	○	○	○	Function	Static	Code-centric	○	○	○
PrivTrans* [72]	○	○	○	Function	Hybrid	Code-centric	●	○	○
Glandring* [185]	○	○	○	Function	Static	Code-centric	●	○	○
Shred* [8], CAPACTY* [108]	○	○	○	Any	Static	Code-centric	○	○	○
DataShield* [77]	○	○	○	Any	Static	Code-centric	○	○	○
Swift* [89]	○	○	○	Any	Static	Code-centric	○	○	○
Jif* [261]	○	○	○	Any	Static	Code-centric	○	○	○
PM [130]	○	○	○	Function	Hybrid	Code-centric	○	○	○
ISys* [110]	○	○	○	Driver	Static	Code-centric	○	○	○
Call* [85]	○	○	○	Library	Static	Code-centric	○	○	○
CompartmentOS* [55]	○	○	○	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	○	○	○	Package	Static	Code-centric	○	○	○
BreakApp* [135]	○	○	○	Package	Static	Code-centric	○	○	○
CompARTIS* [112]	○	○	○	Library	Static	Code-centric	○	○	○
MCS* [89]	○	○	○	Function	Any	Code-centric	○	○	○
ProgramCutter [253]	○	○	○	Function	Dynamic	Code-centric	○	○	○
μSCOPE [202], SCALPEL [201]	○	○	○	Any	Dynamic	Code-centric	○	○	○

(1) Lefeuvre et al., SoK: Software Compartmentalization, S&P 2025

Systematizing Policy Definition Methods

Policy definition methods table from the paper

TABLE 1: *Taxonomy of Policy Definition Methods*. PDMs that also propose an abstraction are marked with *. Manual (○) and fully automated (●) policies do not leverage a policy language, thus the column features Not/Applicable.

Policy Definition Method	Automation ○●●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	●	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	●	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	●	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	●	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	●	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	●	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	●	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	●	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	●	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	●	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	●	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	●	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	●	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	●	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	●	●	●	Any	Static	Code-centric	●	○	●
PM [170]	●	●	●	Function	Hybrid	Code-centric	●	●	○
KSPLIT* [133]	●	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	●	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	●	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	●	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	●	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	●	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	●	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ● = guided manual, ● = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Names of the policy definition methods we consider



Characteristics we included in the taxonomy



Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◑	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◑	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◑	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◑	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◑	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◑	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◒	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◒	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◒	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◒	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◒	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◒	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◒	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◒	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◒	●	●	Function	Hybrid	Code-centric	●	●	○
KSPLIT* [133]	◒	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◒	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◒	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◒	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◒	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◒	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◒	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Names of the policy definition methods we consider



Characteristics we included in the taxonomy



Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [76]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◐	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◐	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◑	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◑	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◑	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◑	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◑	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◑	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◑	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◑	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◑	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◑	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◑	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◑	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◑	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◑	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Names of the policy definition methods we consider



Characteristics we included in the taxonomy



Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [76]	☉	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	☉	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	☉	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	☉	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	☉	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	☉	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	☉	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	☾	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	☾	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	☾	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	☾	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	☾	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	☾	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	☾	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	☾	●	●	Any	Static	Code-centric	●	○	●
PM [170]	☾	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	☾	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	☾	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	☾	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	☾	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	☾	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	☾	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	☾	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ☉ = guided manual, ☾ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

This talk: present a few interesting aspects, many more in the paper 😊

Automation is a key research topic in PDMs

Entirely manual

Fully automated

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◑	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◑	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◑	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◑	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◑	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◑	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◒	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◒	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◒	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◒	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◒	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◒	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◒	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◒	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◒	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◒	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◒	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◒	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◒	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◒	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◒	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◒	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Automation is a key research topic in PDMs

Entirely manual

Fully automated

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◑	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◒	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◓	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◔	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◕	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◖	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◗	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◘	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◙	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◚	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◛	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◜	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◝	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◞	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◟	●	●	Function	Hybrid	Code-centric	●	●	○
KSPLIT* [133]	◠	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◡	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◢	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◣	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◤	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◥	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◦	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Function	dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual

² Implemented with static analysis, dynamic analysis possible [90].

They achieve automation leveraging different kinds of inputs from developers.

Automation is a key research topic in PDMs

Entirely manual

Fully automated

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◑	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◑	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◑	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◑	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◑	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◑	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◑	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◑	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◑	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◑	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◑	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◑	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◑	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◑	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◑	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◑	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◑	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◑	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◑	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◑	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	●	○	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation

² dynamic analysis possible [90].

They employ different kinds of analysis techniques (static, dynamic, hybrid)

Automation is a key research topic in PDMs

Specializing on a particular programming language is often necessary to simplify the problem

Entirely manual

Fully automated

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◑	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◑	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◑	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◑	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◑	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◑	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◑	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◑	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◑	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◑	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◑	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◑	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◑	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◑	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◑	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◑	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◑	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◑	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◑	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◑	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Pick one influential example from the literature: **PtrSplit**⁽¹⁾

Policy Definition Method	Automation ○●●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◐	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◐	●	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◑	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◑	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◑	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◑	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◑	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◑	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◑	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◑	●	●	Function	Hybrid	Code-centric	●	●	○
KSplit* [133]	◑	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◑	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◑	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◑	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◑	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◑	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◑	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Pick one influential example from the literature: **PtrSplit**⁽¹⁾

Policy Definition Method	Automation ○●●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◐	●	○	Function	Static	Any	●	○	●
SOAAP* [117]	◐	●	●	Any	Hybrid	Any	●	●	●
SeCage* [171]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◐	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◐	●	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [165]	◐	●	○	Function	Static	Code-centric	●	○	●
Shreds* [87], CAPACITY* [105]	◐	●	○	Any	Static	Code-centric	○	○	●
DataShield* [77]	◐	●	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◐	●	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◐	●	●	Any	Static	Code-centric	●	○	●
PM [170]	◐	●	●	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◐	●	●	Driver	Static	Code-centric	●	○	○
Cali* [65]	◐	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◐	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◐	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◐	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◐	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◐	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

Semi-automated tool
for *safeboxing* code

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Pick one influential example from the literature: **PtrSplit**⁽¹⁾

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	μLibrary	Static ²	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Static	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	μLibrary	Dynamic	Code-centric	●	●	○
RLBox* [181]	◐	○	○	Function	Static	Any	●	○	●
SOAAP* [117]	◐	○	○	Any	Hybrid	Any	●	●	●
SeCage* [171]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◐	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [73]	◐	○	○	Function	Static	Code-centric	●	○	●
Shreds* [87]	◐	○	○	Any	Static	Code-centric	○	○	●
DataShield* [87]	◐	○	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◐	○	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◐	○	○	Any	Static	Code-centric	●	○	●
PM [170]	◐	○	○	Function	Hybrid	Code-centric	●	●	○
KSplitt* [133]	◐	○	○	Driver	Static	Code-centric	●	○	○
Cali* [65]	◐	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◐	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◐	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◐	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◐	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◐	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

Semi-automated tool
for *safeboxing* code

Annotate valuable data,
PtrSplit then automatically
cuts the program to
safebox them

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Pick one influential example from the literature: **PtrSplit**⁽¹⁾

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	Function	Hybrid	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Hybrid	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	Function	Hybrid	Code-centric	●	●	○
RLBox* [181]	◐	○	○	Function	Hybrid	Any	●	○	●
SOAAP* [117]	◐	○	○	Function	Hybrid	Any	●	●	●
SeCage* [171]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◐	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
Glamdring* [169]	◐	○	○	Function	Static	Code-centric	●	○	●
Shreds* [87]	◐	○	○	Any	Static	Code-centric	○	○	●
DataShield* [169]	◐	○	○	Any	Static	Code-centric	○	○	●
Swift* [89]	◐	○	○	Any	Static	Code-centric	●	○	●
Jif* [261]	◐	○	○	Any	Static	Code-centric	●	○	●
PM [170]	◐	○	○	Function	Hybrid	Code-centric	●	●	○
KSplit* [133]	◐	○	○	Driver	Static	Code-centric	●	○	○
Cali* [65]	◐	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◐	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◐	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◐	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◐	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◐	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

Semi-automated tool for *safeboxing* code

The program is then split at *arbitrary function boundaries*

Annotate valuable data, PtrSplit then automatically cuts the program to safebox them

¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Pick one influential example from the literature: **PtrSplit**⁽¹⁾

Policy Definition Method	Automation ○○●● ¹	Policy Language Type		Separation Granularity	Analysis Approach	Subject Selection	Language Specific	Additional Goals of Automation	
		Annotations	Placement Rules					Performance	Interface Safety
Manual [128], [...]	○	N/A	N/A	Any	Manual	Any	○	N/A	N/A
Crowbar [70]	◐	○	○	Function	Dynamic	Code-centric	○	○	○
MPDs* [191, 192]	◐	○	○	Component	Hybrid	Code-centric	●	●	○
CubicleOS* [211]	◐	●	●	Function	Hybrid	Code-centric	●	○	○
Google SAPI* [22]	◐	●	●	Function	Hybrid	Code-centric	●	○	○
FlexOS* [143, 161]	◐	●	●	Function	Hybrid	Code-centric	●	●	○
RLBox* [181]	◐	○	○	Function	Dynamic	Any	●	○	●
SOAAP* [117]	◐	○	○	Function	Dynamic	Any	●	●	●
SeCage* [171]	◐	○	○	Function	Hybrid	Code-centric	●	○	○
PtrSplit* [169]	◐	●	○	Function	Static	Code-centric	●	○	○
PrivTrans* [73]	◐	○	○	Function	Dynamic	Code-centric	●	○	○
Glamdring* [169]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
Shreds* [87]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
DataShield* [169]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
Swift* [89]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
Jif* [261]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
PM [170]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
KSplit* [133]	◐	○	○	Function	Dynamic	Code-centric	○	○	●
Cali* [65]	◐	○	●	Library	Static	Code-centric	○	○	○
CompartOS* [55]	◐	○	●	Linkage Unit	Static	Code-centric	○	○	○
Enclosure* [111]	◐	○	●	Package	Static	Code-centric	●	○	○
BreakApp* [235]	◐	○	●	Package	Static	Code-centric	●	○	○
CompARTist* [132]	◐	○	●	Library	Static	Code-centric	●	○	○
ACES* [90]	◐	○	●	Function	Any ³	Code-centric	○	○	○
ProgramCutter [253]	●	N/A	N/A	Function	Dynamic	Code-centric	●	○	○
μSCOPE [202], SCALPEL [203]	●	N/A	N/A	Any	Dynamic	Code-centric	○	●	○

Semi-automated tool for *safeboxing* code

The program is then split at *arbitrary function boundaries*

Annotate valuable data, PtrSplit then automatically cuts the program to safebox them

The split is done *statically* by the tool, i.e., just by looking at the code (vs. a dynamic tool that would run the program)

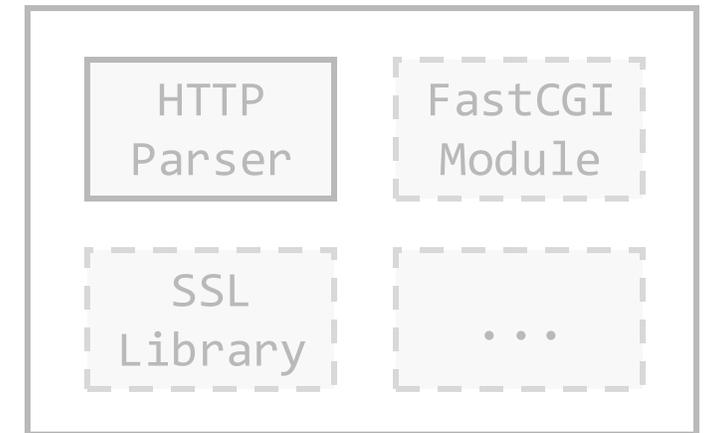
¹ ○ = manual, ◐ = guided manual, ◑ = policy refinement, ● = full automation. ² Loader-based. ³ Implemented with static analysis, dynamic analysis possible [90].

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path*. We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**



(Fictive monolithic program)



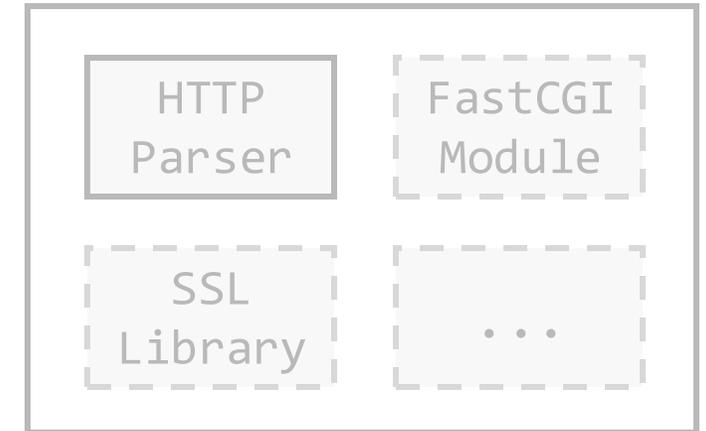
Which trade-offs are desirable in practice?

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path*. We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically**.



(Fictive monolithic program)



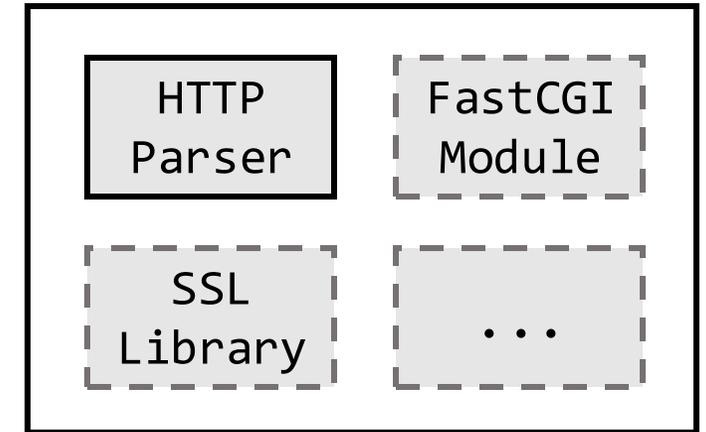
Which trade-offs are desirable in practice?

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**



(Fictive monolithic program)



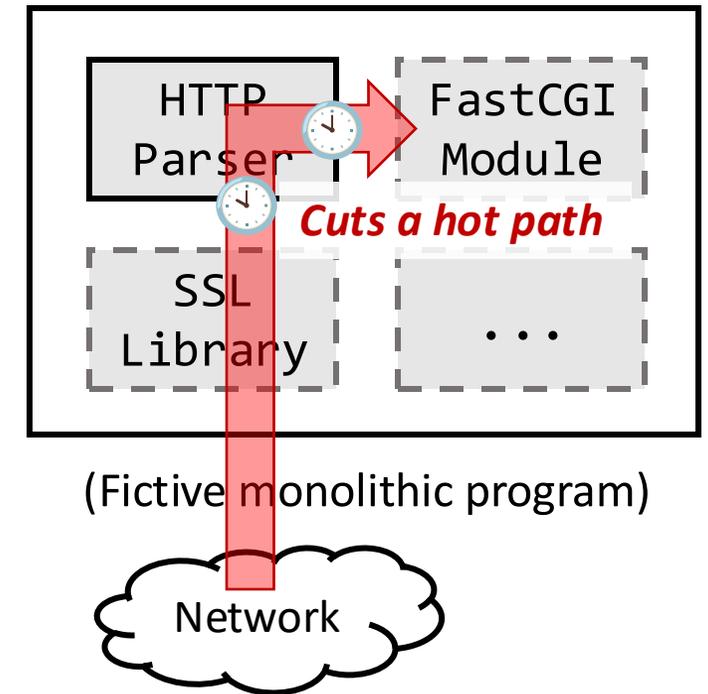
Which trade-offs are desirable in practice?

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**



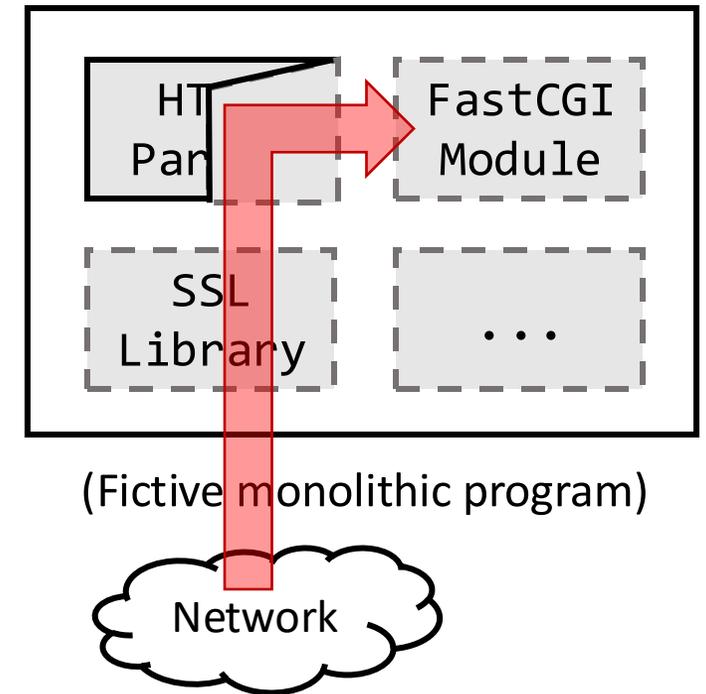
Which trade-offs are desirable in practice?

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**



Which trade-offs are desirable in practice?

Open Problem in Policy Definition Methods

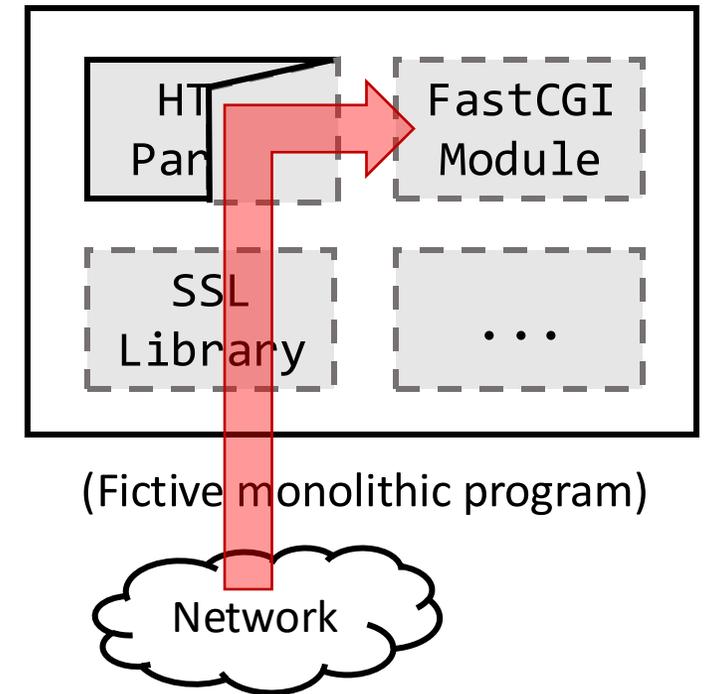
Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

Rice's theorem: "*non-trivial semantic properties of programs are undecidable*".

Which trade-offs are desirable in practice?



Open Problem in Policy Definition Methods

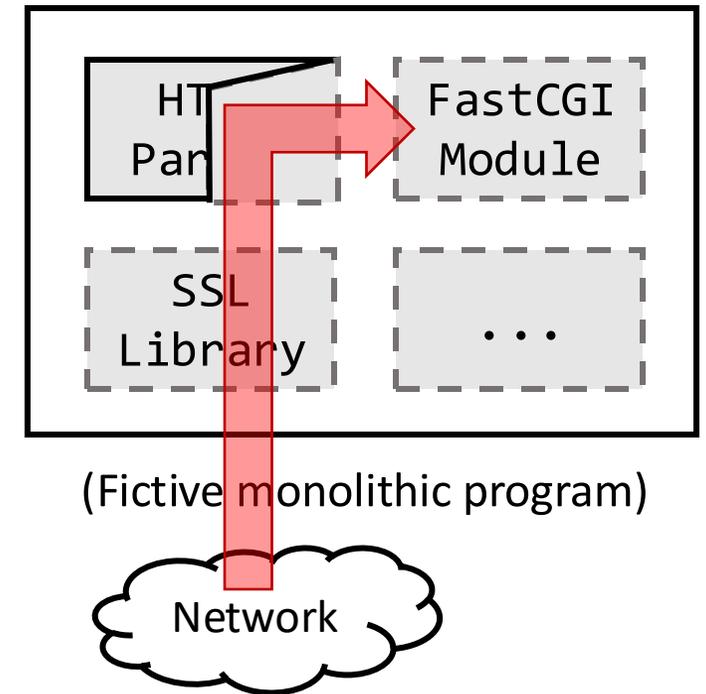
Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

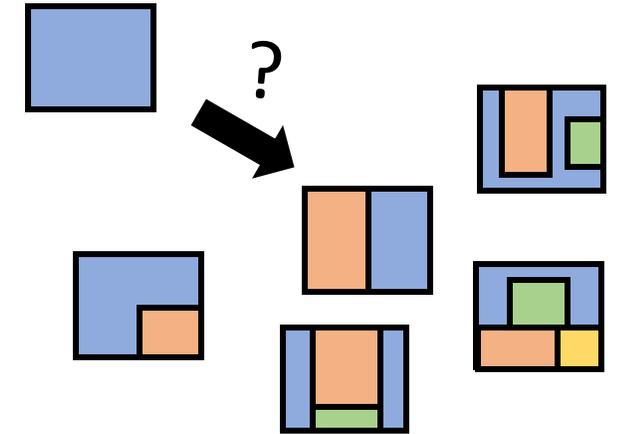
Rice's theorem: "*non-trivial semantic properties of programs are undecidable*".

Which trade-offs are desirable in practice?



Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:



1. How to determine the right policy to enforce?

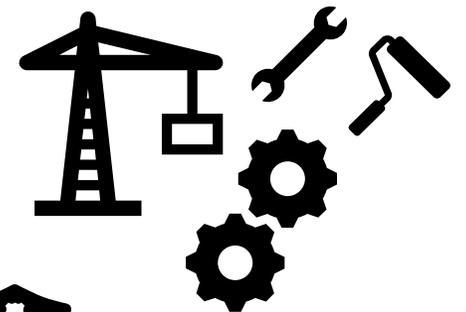
- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

- Done with a *compartmentalization abstraction*

3. How to enforce policies at runtime?

- Done with a *compartmentalization mechanism*



Problem #2

How to implement policies?

Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Problem #2

How to implement policies?

Having defined a policy, we need to **express it in the program**.

Developers perform this using **programming abstractions**.

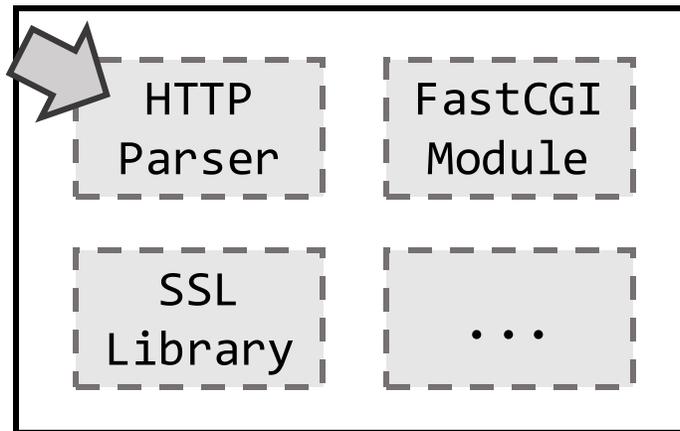
Historically, people have done this with **processes** (=the *process abstraction*).

Problem #2

How to implement policies?

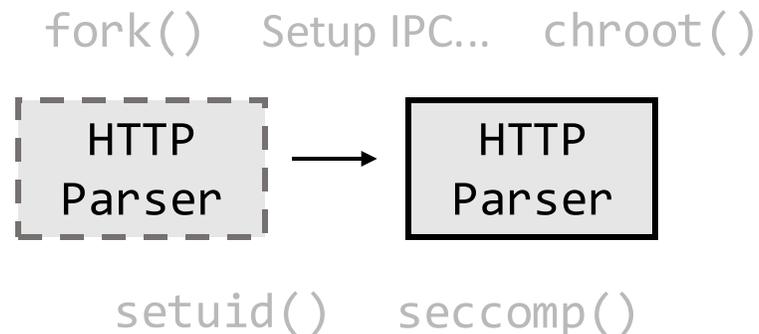
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



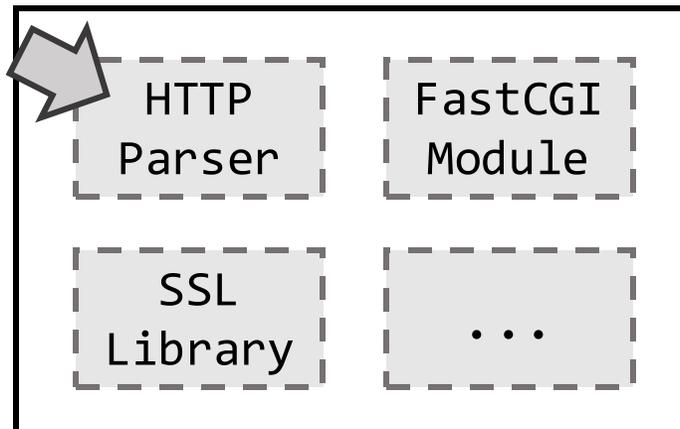
This is still the most common way to do it today.

Problem #2

How to implement policies?

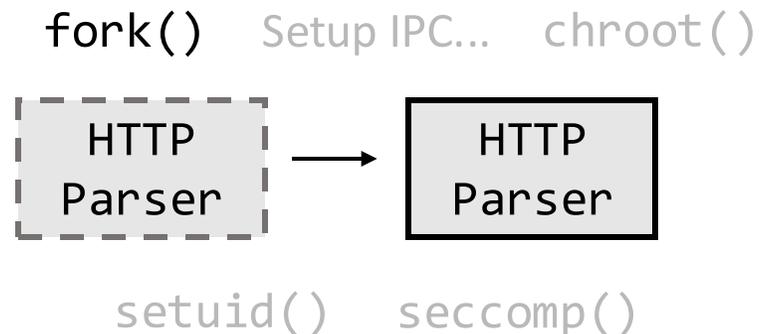
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



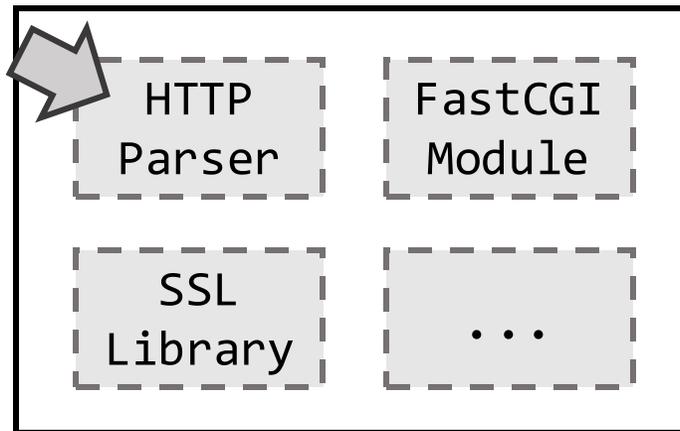
This is still the most common way to do it today.

Problem #2

How to implement policies?

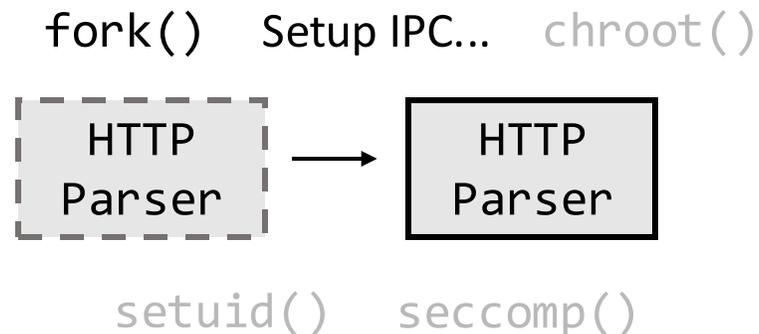
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



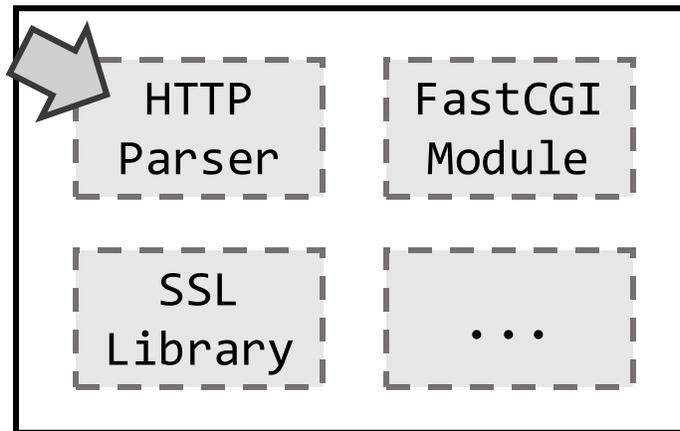
This is still the most common way to do it today.

Problem #2

How to implement policies?

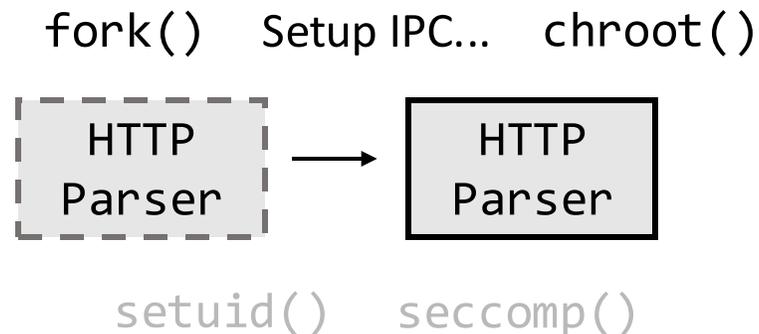
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



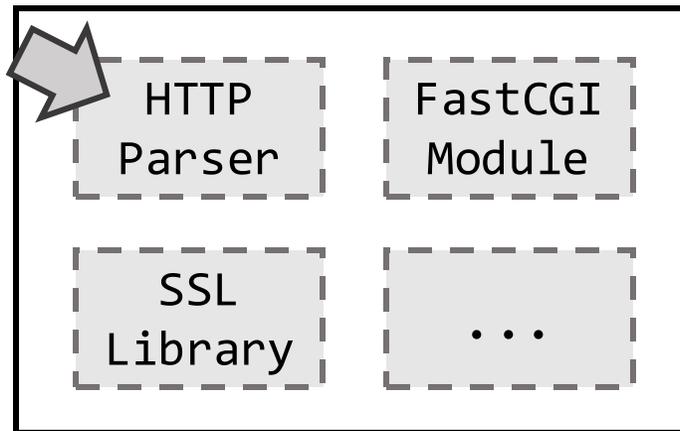
This is still the most common way to do it today.

Problem #2

How to implement policies?

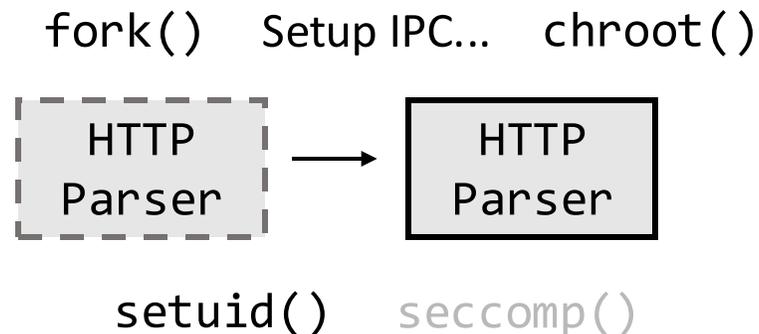
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



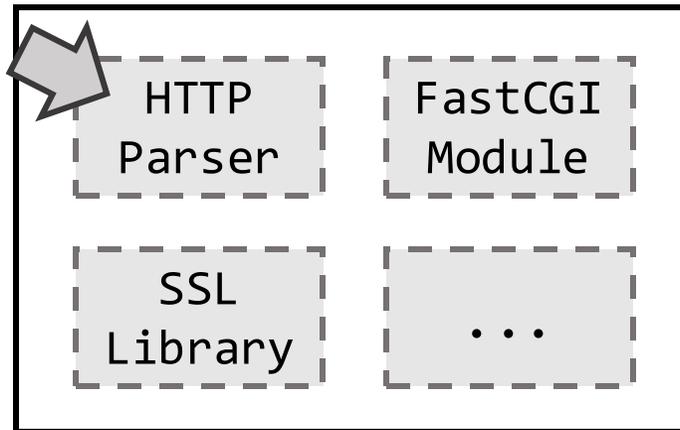
This is still the most common way to do it today.

Problem #2

How to implement policies?

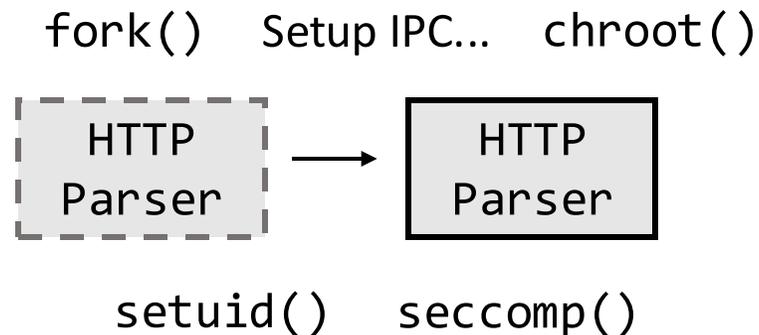
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



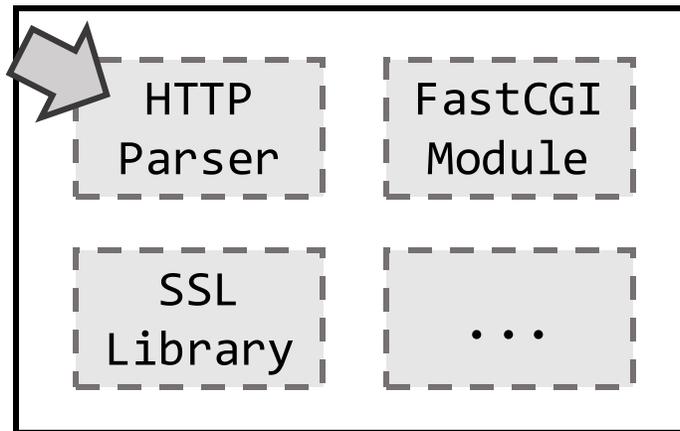
This is still the most common way to do it today.

Problem #2

How to implement policies?

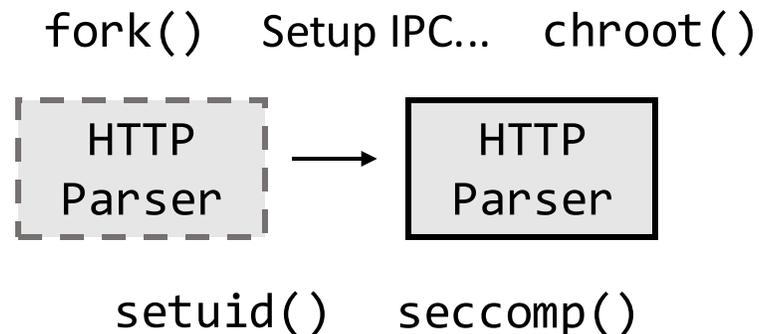
Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



(Fictive monolithic program)

Example: *assume we sandbox the HTTP parser.*



This is still the most common way to do it today.

Compartmentalization Abstractions

Programming abstractions for compartmentalization are also a very active research area:

Minimize developer effort (intuitive, easy to use)

Leverage domain-specific knowledge (threat model, deployment)

Maximize security properties that can be achieved

Making the most of a specific

enforcement mechanism

Support generic mechanisms

TABLE 2: Taxonomy of Compartmentalization Abstractions. Targets: User, Kernel, Hypervisor. Semantics: Synchronous, Asynchronous, Shared Memory, Message passing. S+A: the abstraction exhibits both S and A semantics. Properties: Confidentiality, Integrity, Availability, Recompartmentalization. Mechanism-independent abstractions are labeled with \emptyset .

Class	Target (U/K/HV)	Abstraction	Subject Selection	Semantics	Abstraction Granularity	Properties			Design Bound to Mechanism	
						C	I	A		
Manual/Domain	U	Virtues [202]	Code-centric	S	MES	Function	●	●	●	Virtual Machine (VPT)
		ACES [90]	Code-centric	S	SMM	Function	●	●	●	\emptyset^*
		SoCage [173]	Code-centric	S	SMM	Function	●	●	●	\emptyset^*
		HODOR [158]	Code-centric	S	SMM	Library	●	●	●	\emptyset^*
		CAPACITY [109]	Code-centric	S	SMM	Any	●	●	●	ARM PAC + MTE
		Jill [261]	Code-centric	S+A	MES	Any	●	●	●	\emptyset^*
		Arbiter [231]	Data-centric	S	SMM	Function ¹	●	●	●	\emptyset^*
		Secure Memory Views (SMVs) [128]	Data-centric	S	SMM	Function ¹	●	●	●	\emptyset^*
		Salus [226]	Data-centric	S	SMM	Function ¹	●	●	●	\emptyset^*
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SMM	Function ¹	●	●	●	Page Table ²
		TPMIX Processes (and earlier instances) [91]	Hybrid	Any	Any	Any	●	●	●	Page Table
		SOAAP [117]	Hybrid	S	SMM	Any	●	●	●	\emptyset^*
		IBMPK [199]	Hybrid	S	SMM	Any	●	●	●	Protection Keys
		CHEROS [108]	Code-centric	Any	Any	U/K-component	●	●	●	CHERI
		Sandbox	U+K	Microkernel Servers [106, 1, 1]	Code-centric	Any	MES	U/K-component	●	●
Mutable Protection Domains (MPDs) [191, 192]	Code-centric			S	SMM	U/K-component	●	●	●	\emptyset^*
RedLeaf [184]	Code-centric			S	SMM	U/K-component	●	●	●	Safe Languages
CubicleOS [111]	Code-centric			S	SMM	All-library	●	●	●	Protection Keys
FlexOS [161]	Code-centric			S	SMM	All-library	●	●	●	\emptyset^*
XMP [186]	Code-centric			S	SMM	Any	●	●	●	\emptyset^*
Monix [101]	Hybrid			A	SMM	Function ¹	●	●	●	\emptyset^*
VirtOS [180]	Code-centric			S+A	SMM	K-component	●	●	●	Virtual Machine (VPT)
THAKC [173]	Code-centric			S	SMM	Function	●	●	●	ARM PAC + MTE
LibertOS [187]	Code-centric			S	SMM	K-component	●	●	●	\emptyset^*
Call set	Code-centric			S	SMM	Library	●	●	●	\emptyset^*
CompARKit [103]	Code-centric			S	MES	Library	●	●	●	\emptyset^*
Enclosure [111]	Code-centric			S	SMM	Package	●	●	●	\emptyset^*
Google Sandbox API (SAPI) [111]	Code-centric			S	MES	Function	●	●	●	\emptyset^*
RELIX 2 (RELIX 00, 101)	Hybrid			S	SMM	Function	●	●	●	\emptyset^*
Wedge [191]	Hybrid	S	SMM	Function ¹	●	●	●	\emptyset^*		
Hypervisor/Domain	U+K	CompartOS [91]	Code-centric	S	SMM	Linkage Unit	●	●	●	CHERI
		LYNX 7 (SMM 110, 180)	Code-centric	S	MES	K-component	●	●	●	\emptyset^*
		XPLXPI [101, 120]	Hybrid	S	SMM	K-component	●	●	●	SMT
		Nexus [211]	Data-centric	S	SMM	Per-VM domain	●	●	●	Page Table ²
		Shreds [87]	Code-centric	S	SMM	Any	●	●	●	\emptyset^*
Sandbox/Domain	U	Pyramid [147]	Code-centric	S	MES	Function	●	●	●	Page Table ²
		Proxima [78]	Code-centric	S	MES	Function	●	●	●	Page Table ²
		Swift [80]	Code-centric	S+A	MES	Any	●	●	●	\emptyset^*
		Glandring [161]	Code-centric	S	MES	Function	●	●	●	\emptyset^*
		Proxifier / Program Mandering [106, 190]	Code-centric	S	MES	Function	●	●	●	\emptyset^*
		DataShield [71]	Hybrid	S	SMM	Any	●	●	●	Bounds Checking
		TERM [212]	Hybrid	S	SMM	Any	●	●	●	Protection Keys
		Nested Kernel [91]	Code-centric	S	SMM	Function	●	●	●	Page Table ²

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel. ⁴ The PDM does not have a certain extent.
⁵ The abstraction could plug into any into-AS mechanism, though the paper or documentation claims reliance on a particular one.



(1)

(1) Lefeuve et al., SoK: Software Compartmentalization, S&P 2025

Compartmentalization Abstractions

Programming abstractions for compartmentalization are also a very active research area:

Minimize developer effort (intuitive, easy to use)

Leverage domain-specific knowledge (threat model, deployment)

Maximize security properties that can be achieved

Making the most of a specific

enforcement mechanism

Support generic mechanisms

TABLE 2: Taxonomy of Compartmentalization Abstractions. Targets: User, Kernel, Hypervisor. Semantics: Synchronous, Asynchronous, Shared Memory, Message passing. S+A: the abstraction exhibits both S and A semantics. Properties: Confidentiality, Integrity, Availability, Recompartmentalization. Mechanism-independent abstractions are labeled with ∅.

Class	Target (U/K/HV)	Abstraction	Subject Selection	Semantics	Abstraction Granularity	Properties			Design Bound to Mechanism		
						C	I	A			
Manual Derived	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	Virtual Machine (EPT)
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	∅ ²
		SoCage [171]	Code-centric	S	SHM	Function	●	●	○	○	∅ ²
		HODOR [124]	Code-centric	S	SHM	Library	●	○	○	○	S
		CAPACITY [109]	Code-centric	S	SHM	Any	●	●	○	○	ARM PAC + MTE
		Jif [261]	Code-centric	S+A	MES	Any	●	○	○	●	∅
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	∅ ²
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	○	○	○	∅ ²
		Salus [284]	Data-centric	S	SHM	Function ¹	●	○	○	○	∅ ²
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	○	○	○	Page Table ²
		POXIX Processes (and earlier instances) [91]	Hybrid	Any	Any	Any	●	○	○	○	Page Table
		SOAAP [117]	Hybrid	S	SHM	Any	●	○	○	○	∅
		IBMPK [190]	Hybrid	S	SHM	Any	●	○	○	○	Protection Keys
		CHEROES [149]	Code-centric	Any	Any	U/K-component	●	○	○	○	CHERI
		Microkernel Servers [106, 1, 1]	Code-centric	Any	MES	U/K-component	●	○	○	○	∅ ²
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	S	SHM	U/K-component	●	○	○	○	∅ ²
		RedLeaf [184]	Code-centric	S	SHM	U/K-component	●	○	○	○	Safe Languages
		CubicleOS [111]	Code-centric	S	SHM	AllLibrary	●	○	○	○	Protection Keys
FlexOS [161]	Code-centric	S	SHM	AllLibrary	●	○	○	○	∅		
XMP [180]	Code-centric	S	SHM	Any	●	○	○	○	∅ ²		
Monza [101]	Hybrid	A	SHM	Function ¹	●	○	○	○	∅ ²		
Syntho	U	VirtOS [186]	Code-centric	S+A	SHM	k-component	●	●	○	○	Virtual Machine (EPT)
		THAKC [173]	Code-centric	S	SHM	Function	●	○	○	○	ARM PAC + MTE
		LibertOS [197]	Code-centric	S	SHM	k-component	●	○	○	○	∅ ²
		Call [101]	Code-centric	S	SHM	Library	●	○	○	○	∅ ²
		CompARtIt [132]	Code-centric	S	MES	Library	●	○	○	○	ARM PAC + MTE
		Enclosure [111]	Code-centric	S	SHM	Package	●	○	○	○	∅
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	○	○	○	∅
		REKEX [23] SWITCH [101, 102]	Hybrid	S	SHM	Function	●	○	○	○	∅
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	○	○	○	∅
		CompARtOS [131]	Code-centric	S	SHM	Linkage Unit	●	○	○	○	CHERI
Syntho	K	LYDR 7 (S/PM) [113, 189]	Code-centric	S	MES	k-component	●	○	○	○	∅
		XFPLXP [169, 172]	Hybrid	S	SHM	k-component	●	○	○	○	SPT
		Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	○	○	○	Page Table ²
		Shreds [87]	Code-centric	S	SHM	Any	●	○	○	○	∅ ²
Syntho	Dual World	Privman [147]	Code-centric	S	MES	Function	●	○	○	○	Page Table ²
		Proxman [78]	Code-centric	S	MES	Function	●	○	○	○	Page Table ²
		Swift [88]	Code-centric	S+A	MES	Any	●	○	○	○	∅
		Glandrind [165]	Code-centric	S	MES	Function	●	○	○	○	∅ ²
		Proxifier Program Mandering [169, 170]	Code-centric	S	MES	Function	●	○	○	○	∅ ²
		DataSet [71]	Hybrid	S	SHM	Any	●	○	○	○	Bounds Checking
		TERM [232]	Hybrid	S	SHM	Any	●	○	○	○	Protection Keys
		Nested Kernel [91]	Code-centric	S	SHM	Function	●	○	○	○	Page Table ²

¹ Inherited from thread-like semantics; ² from process-like semantics; ³ from the Nested Kernel. ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.



(1)

(1) Lefeuve et al., SoK: Software Compartmentalization, S&P 2025

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset	
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
	SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset		
	libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys		
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5	
		RedLeaf [184]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset	
		xMP [196]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset	
	Monza [35]	Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5		
K	VirtuOS [186]	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5		
Sandbox	U	Cali [65]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5	
		CompARTist [132]	Code-centric	\mathcal{S}	MES	Library	●	●	○	○	○	\emptyset^5	
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset	
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5	
		RLBox / μ SWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset	
		Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
		XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset^5
		U	Privman [147]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Privtrans [73]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Swift [89]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset
			Glamdring [165]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5
			PtrSplit / Program Mandering [169, 170]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5
			DataShield [77]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
		ERIM [232]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Names of the abstractions we consider

Characteristics we included in the taxonomy

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design to asm	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset	
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset	
		libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5	
		RedLeaf [184]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset	
		xMP [196]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset	
Monza [35]		Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5		
K	VirtuOS [186]	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5		
Sandbox	U	Cali [65]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5	
		CompARTist [132]	Code-centric	\mathcal{S}	MES	Library	●	●	○	○	○	\emptyset^5	
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset	
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5	
		RLBox / μ SWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset	
		Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
		XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset^5
		U	Privman [147]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Privtrans [73]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Swift [89]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset
			Glamdring [165]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5
			PtrSplit / Program Mandering [169, 170]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5
			DataShield [77]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
			ERIM [232]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys
		K	Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Names of the abstractions we consider

Characteristics we included in the taxonomy

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design to asm	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Light Weight Closures (LWCs)	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOXAN [177]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅	
	libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys		
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵	
		RedLeaf [184]	Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	∅	
		xMP [196]	Code-centric	S	SHM	Any	●	●	○	○	○	∅	
Monza [35]		Hybrid	A	SHM	Function ¹	○	●	○	○	○	∅ ⁵		
K	VirtuOS [186]	Code-centric	S+A	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	S	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	S	SHM	K-component	●	●	●	●	○	∅ ⁵		
Sandbox	U	Cali [65]	Code-centric	S	SHM	Library	●	●	○	○	○	∅ ⁵	
		CompARTist [132]	Code-centric	S	MES	Library	●	●	○	○	○	∅ ⁵	
		Enclosure [111]	Code-centric	S	SHM	Package	●	●	○	○	○	∅	
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	●	○	○	∅ ⁵	
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅	
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
	U+K	CompartOS [55]	Code-centric	S	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	S	MES	K-component	●	●	○	○	○	∅ ⁵	
		XFI/LXFI [107, 172]	Hybrid	S	SHM	K-component	●	●	○	○	●	SFI	
	HV	Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	Page Table ³	
Safebox	Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	∅ ⁵
		U	Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
			Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
			Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅
			Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵
			PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵
			DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
			ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

This talk: present a few interesting aspects, full discussion in the paper 😊

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset	
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
	SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset		
	libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys		
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5	
		RedLeaf [184]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset	
		xMP [196]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset	
	Monza [35]	Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5		
K	VirtuOS [186]	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5		
Sandbox	U	Cali [65]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5	
		CompARTist [132]	Code-centric	\mathcal{S}	MES	Library	●	●	○	○	○	\emptyset^5	
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset	
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5	
		RLBox / μ SWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset	
		Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
		XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset^5
		U	Privman [147]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Privtrans [73]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Swift [89]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset
			Glamdring [165]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5
			PtrSplit / Program Mandering [169, 170]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5
			DataShield [77]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
		ERIM [232]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Abstractions often specialize on a particular threat model

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Me	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	Virtual	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	Virtual	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	Virtual	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○	Virtual	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	Virtual	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOAAP [117]	Hybrid	S	SHM	Any	●	●	○	○	○	∅	
		libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵	
		RedLeaf [184]	Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	∅	
		xMP [196]	Code-centric	S	SHM	Any	●	●	○	○	○	∅	
K	Monza [35]	Hybrid	A	SHM	Function ¹	○	●	○	○	○	∅ ⁵		
	VirtuOS [186]	Code-centric	S+A	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	S	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
		LibrettOS [187]	Code-centric	S	SHM	K-component	●	●	●	○	∅ ⁵		
Sandbox	U	Cali [65]	Code-centric	S	SHM	Library	●	●	○	○	○	∅ ⁵	
		CompARTist [132]	Code-centric	S	MES	Library	●	●	○	○	○	∅ ⁵	
		Enclosure [111]	Code-centric	S	SHM	Package	●	●	○	○	○	∅	
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	●	○	○	∅ ⁵	
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅	
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
	U+K	CompartOS [55]	Code-centric	S	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	S	MES	K-component	●	●	○	○	○	∅ ⁵	
		XFI/LXFI [107, 172]	Hybrid	S	SHM	K-component	●	●	○	○	●	SFI	
	HV	Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	Page Table ³	
Safebox	Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	∅ ⁵	
		Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
		Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
		Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	○	∅	
		Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	○	∅ ⁵	
		PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○	∅ ⁵	
		DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○	Bounds Checking	
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table

vs. processes, which are quite generic but also trickier to use

For sandboxing (libraries, drivers)

For safeboxing (crypto keys, privileged code)



¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

A lot of interest in abstractions for kernel compartmentalization

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
U		Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		[128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		[167]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Instances) [93]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²	
			Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
			Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅	
			Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
			Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
		U+K		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○
Mutable Protection Domains (MPDs) [191, 192]	Code-centric			S	SHM	U/K-component	●	●	○	●	○	∅ ⁵	
RedLeaf [184]	Code-centric			S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
CubicleOS [211]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
FlexOS [161]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	∅	
xMP [196]	Code-centric			S	SHM	Any	●	●	○	○	○	∅	
Monza [35]	Hybrid			A	SHM	Function ¹	○	●	○	○	○	∅ ⁵	
VirtuOS [186]	Code-centric			S+A	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)	
HAKC [173]	Code-centric			S	SHM	Function	●	●	○	○	○	ARM PAC + MTE	
LibrettOS [187]	Code-centric			S	SHM	K-component	●	●	●	●	○	∅ ⁵	
Sandbox		[65]	Code-centric	S	SHM	Library	●	●	○	○	○	∅ ⁵	
		spARTist [132]	Code-centric	S	MES	Library	●	●	○	○	○	∅ ⁵	
		losure [111]	Code-centric	S	SHM	Package	●	●	○	○	○	∅ ⁵	
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	●	○	○	∅ ⁵	
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅ ⁵	
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		U+K	CompartOS [55]	Code-centric	S	SHM	Linkage Unit	●	●	●	○	○	CHERI
		K	LVDs / KSplit [133, 185]	Code-centric	S	MES	K-component	●	●	○	○	○	∅ ⁵
			XFI/LXFI [107, 172]	Hybrid	S	SHM	K-component	●	●	○	○	●	SFI
		HV	Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	Page Table ³
Safebox		[65]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	∅ ⁵	
			Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
			Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
			Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
			Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵	
			Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
			Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
			Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
			Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table

Compartmentalize userland and kernel together

Kernel only...

This came as a revival of research in microkernels

Or hypervisors

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

A lot of interest in rethinking **communication** and **sharing** between domains

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Salus [226]	Data-ce				●	●	○	○	○	∅ ⁵	
		Light-Weight Contexts (LwCs) [167]	Hybr									Page Table ²	
	POSIX Processes (and earlier instances) [93]	Hybr									Page Table		
	SOAAP [117]	Hybr									∅		
	libMPK [190]	Hybr									Protection Keys		
	U+K	CheriOS [108]	Code-ce									CHERI	
		Microkernel Servers [106], [...]	Code-ce									∅ ⁵	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	S		U/K-component	●	●	○	●	○	∅ ⁵	
		RedLeaf [184]	Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
FlexOS [161]		Code-centric	S	SHM	μLibrary	●	●	○	○	○	∅		
xMP [196]		Code-centric	S	SHM	Any	●	●	○	○	○	∅		
Monza [35]		Hybrid	A	SHM	Function ¹	○	●	○	○	○	∅ ⁵		
K	VirtuOS [186]	Code-centric	S+A	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	S	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	S	SHM	K-component	●	●	●	●	○	∅ ⁵		
Sandbox	U	Cali [65]	Code-centric	S	SHM	Library	●	●	○	○	○	∅ ⁵	
		CompARTist [132]	Code-centric	S	MES	Library	●	●	○	○	○	∅ ⁵	
		Enclosure [111]	Code-centric	S	SHM	Package	●	●	○	○	○	∅	
		Google Sandboxes (LPI, GADP)	Code-centric	S	SHM	Function	●	●	●	○	○	∅ ⁵	
		RLBox / μS	Code-centric	S	SHM	Function	●	●	○	○	○	∅	
		Wedge [70]	Code-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
	U+K	CompartOS	Code-centric	S	SHM	Process Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSp	Code-centric	S	SHM	K-component	●	●	○	○	○	∅ ⁵	
		XFI/LXFI [1]	Code-centric	S	SHM	K-component	●	●	○	○	●	SFI	
	HV	Nexen [221]	Code-centric	S	SHM	Domain	●	●	○	○	○	Page Table ³	
Safebox	Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	∅ ⁵	
		Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
		Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
		Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵	
		PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
		DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table

Leverage shared memory to speed up sharing between domains (vs. message passing)

Exposing communication between domains as function calls (vs. IPCs) to be more intuitive

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Interest in abstractions that **compose with different isolation mechanisms**

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset	
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	\emptyset^5	
	SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset^5		
	libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset^5		
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	\emptyset^5	
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5	
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5	
		RedLeaf [184]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages	
		CubicleOS [211]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys	
		FlexOS [161]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset	
		xMP [196]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset	
	Monza [35]	Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5		
K	VirtuOS [186]	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)		
	HAKC [173]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE		
	LibrettOS [187]	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5		
Sandbox	U	Cali [65]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5	
		CompARTist [132]	Code-centric	\mathcal{S}	MES	Library	●	●	○	○	○	\emptyset^5	
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset	
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5	
		RLBox / μ SWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset	
	Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5		
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
		XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset^5
		U	Privman [147]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Privtrans [73]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²
			Swift [89]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset
			Glamdring [165]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5
			PtrSplit / Program Mandering [169, 170]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5
			DataShield [77]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
		ERIM [232]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

For example: new abstractions that leverage the specifics of safe languages for isolation

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism		
				CALL	ASSIGN		C	I	A	R				
Mutual Distrust	U	Virtines [242]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)		
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5		
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5		
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset		
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE		
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset		
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5		
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5		
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5		
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	Page Table ²		
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table		
		SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset		
		libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys		
	U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI		
		Microkernel Servers [106], [...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5		
		Mutable Protection Domains (MPDs) [191, 192]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5		
		RedLeaf [184]	Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages		
		CubicleOS [211]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys		
		FlexOS [161]	Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset		
		xMP [196]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset		
K	Monza [35]	Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5			
	VirtuOS [186]	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)			
	HAKC [173]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE			
		LibrettOS [187]	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5		
Sandbox	U	Cali [65]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5		
		CompARTist [132]	Code-centric	\mathcal{S}	MES	Library	●	●	○	○	○	\emptyset^5		
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset		
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5		
		RLBox / μSWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset		
		Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5		
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI		
		K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
			XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
	HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset^5	
			Privman [147]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²	
			Privtrans [73]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²	
			Swift [89]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
			Glamdring [165]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5	
		K	PtrSplit / Program Mandering [169, 170]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5	
			DataShield [77]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
			ERIM [232]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
			K	Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		SeCage [171]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	\emptyset^5	
		HODOR [124]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset	
		CAPACITY [105]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
		Arbiter [241]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Secure Memory Views (SMVs) [128]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Salus [226]	Data-centric	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
		Light-Weight Contexts (LwCs) [167]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOAAP [117]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	\emptyset	
		libMPK [190]	Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
		U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI
	Microkernel Servers [106], [...]		Code-centric	Any	MES	U/K-component	●	●	○	○	○	\emptyset^5	
	Mutable Protection Domains (MPDs) [191, 192]		Code-centric	\mathcal{S}	SHM	U/K-component	●	●	○	●	○	\emptyset^5	
	RedLeaf [184]		Code-centric	\mathcal{S}	SHM	U/K-component	●	●	●	○	○	Safe Languages	
	CubicleOS [211]		Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	Protection Keys	
			Code-centric	\mathcal{S}	SHM	μ Library	●	●	○	○	○	\emptyset	
		Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○	\emptyset		
	Hybrid	\mathcal{A}	SHM	Function ¹	○	●	○	○	○	\emptyset^5			
	Code-centric	$\mathcal{S}+\mathcal{A}$	SHM	K-component	●	●	●	○	○	Virtual Machine (EPT)			
	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	ARM PAC + MTE			
	Code-centric	\mathcal{S}	SHM	K-component	●	●	●	●	○	\emptyset^5			
Sandbox	U	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Library	●	●	○	○	○	\emptyset^5	
		Enclosure [111]	Code-centric	\mathcal{S}	SHM	Package	●	●	○	○	○	\emptyset	
		Google Sandboxed API (SAPI) [22]	Code-centric	\mathcal{S}	MES	Function	●	●	●	○	○	\emptyset^5	
		RLBox / μSWITCH [181, 195]	Hybrid	\mathcal{S}	SHM	Function	●	●	○	○	●	\emptyset	
		Wedge [70]	Hybrid	\mathcal{S}	SHM	Function ¹	●	●	○	○	○	\emptyset^5	
	U+K	CompartOS [55]	Code-centric	\mathcal{S}	SHM	Linkage Unit	●	●	●	○	○	CHERI	
	K	LVDs / KSplit [133, 185]	Code-centric	\mathcal{S}	MES	K-component	●	●	○	○	○	\emptyset^5	
		XFI/LXFI [107, 172]	Hybrid	\mathcal{S}	SHM	K-component	●	●	○	○	●	SFI	
	HV	Nexen [221]	Data-centric	\mathcal{S}	SHM	Per-VM domain	●	●	○	○	○	Page Table ³	
	Safebox	Dual World	U	Shreds [87]	Code-centric	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴
Privman [147]			Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²	
Privtrans [73]			Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○	Page Table ²	
Swift [89]			Code-centric	$\mathcal{S}+\mathcal{A}$	MES	Any	●	●	○	○	●	\emptyset	
Glamdring [165]			Code-centric	\mathcal{S}	MES	Function	●	●	○	○	●	\emptyset^5	
PtrSplit / Program Mandering [169, 170]			Code-centric	\mathcal{S}	MES	Function	●	●	○	○	○ ⁴	\emptyset^5	
DataShield [77]			Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
ERIM [232]			Hybrid	\mathcal{S}	SHM	Any	●	●	○	○	○	Protection Keys	
K			Nested Kernel [94]	Code-centric	\mathcal{S}	SHM	Function	●	●	○	○	○	Page Table

Abstraction specialized for sandboxing untrusted code (such as libraries)

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism		
				CALL	ASSIGN		C	I	A	R				
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)		
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵		
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵		
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅		
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE		
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅		
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵		
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵		
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵		
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²		
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table		
		SOAAP [117]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅		
		libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys		
		U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI	
	Microkernel Servers [106], [...]		Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵		
	Mutable Protection Domains (MPDs) [191, 192]		Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵		
	RedLeaf [184]		Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages		
	CubicleOS [211]		Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys		
	Code-centric		S	SHM	μLibrary	●	●	○	○	○	∅			
Sanbox	U	Enclosure [111]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵		
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	○	○	○	∅ ⁵		
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅		
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵		
	U+K	CompartOS [55]	Code-centric	S	SHM	Linkage Unit	●	●	●	○	○	CHERI		
		K	LVDs / KSplit [133, 185]	Code-centric	S	MES	K-component	●	●	○	○	○	∅ ⁵	
	XFI/LXFI [107, 172]		Hybrid	S	SHM	K-component	●	●	○	○	●	SFI		
	HV	Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	Page Table ³		
	Safebox	Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	∅ ⁵
				Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
			U	Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
				Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅
Glamdring [165]				Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵	
PtrSplit / Program Mandering [169, 170]				Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
DataShield [77]				Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
ERIM [232]				Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
K			Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table	

Abstraction specialized for sandboxing untrusted code (such as libraries)

Implements cross-compartment communication with function calls, which is intuitive

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOAAP [117]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅	
		libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI
	Microkernel Servers [106], [...]		Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵	
	Mutable Protection Domains (MPDs) [191, 192]		Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵	
	RedLeaf [184]		Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
	CubicleOS [211]		Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
	Code-centric		S	SHM	μLibrary	●	●	○	○	○	∅		
Sandbox	U	Enclosure [111]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	○	○	○	ARM PAC + MTE	
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅ ⁵	
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		CompartOS [55]	Code-centric	S	SHM	Function	●	●	○	○	○	CHERI	
	U+K	LVDs / KSplit [133]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		XFI/LXFI [107, 172]	Code-centric	S	SHM	Function	●	●	○	○	●	SFI	
	K	Nexen [221]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table ³	
		HV	Shreds [87]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵
			Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
Privtrans [73]	Code-centric		S	MES	Function	●	●	○	○	○	Page Table ²		
Safebox	Dual World	Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵	
		PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
		DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table	

Abstraction specialized for sandboxing untrusted code (such as libraries)

Implements cross-compartment communication with function calls, which is intuitive

Builds on shared memory to speed up cross-compartment data sharing

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).

⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism	
				CALL	ASSIGN		C	I	A	R			
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅	
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table	
		SOAAP [117]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅	
		libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	CHERI
	Microkernel Servers [106], [...]		Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵	
	Mutable Protection Domains (MPDs) [191, 192]		Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵	
	RedLeaf [184]		Code-centric	S	SHM	U/K-component	●	●	●	○	○	Safe Languages	
	CubicleOS [211]		Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys	
	Code-centric		S	SHM	μLibrary	●	●	○	○	○	∅		
Sandbox	U	Enclosure [111]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)	
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	○	○	○	ARM PAC + MTE	
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅ ⁵	
		Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
		CompartOS [55]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵	
	U+K	LVDs / KSplit [133]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
		XFI/LXFI [107, 172]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵	
	K	Nexen [221]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table ³	
		Shreds [87]	Code-centric	S	SHM	Function	●	●	○	○	○ ⁴	∅ ⁵	
		Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
Safebox	Dual World	Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²	
		Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅	
		Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵	
		PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵	
		DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking	
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys	
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table

Abstraction specialized for sandboxing untrusted code (such as libraries)

Implements cross-compartment communication with function calls, which is intuitive

Builds on shared memory to enforce confidentiality and integrity up cross-compartment data sharing

¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Pick one influential example from the literature: **RLBox**⁽¹⁾

Class	Target U/K/HV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties				Interface Safety	Design Bound to Mechanism
				CALL	ASSIGN		C	I	A	R		
Mutual Distrust	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	Virtual Machine (EPT)
		ACES [90]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵
		HODOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	∅
		CAPACITY [105]	Code-centric	S	SHM	Any	●	●	○	○	○ ⁴	ARM PAC + MTE
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	Page Table ²
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	Page Table
		SOAAP [117]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	∅
		libMPK [190]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys
		U+K	CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○
	Microkernel Servers [106], [...]		Code-centric	Any	MES	U/K-component	●	●	○	○	○	∅ ⁵
	Mutable Protection Domains (MPDs) [191, 192]		Code-centric	S	SHM	U/K-component	●	●	○	●	○	∅ ⁵
	RedLeaf [184]		Code-centric	S	SHM	U/K-component	●	●	○	●	○	Safe Languages
	CubicleOS [211]		Code-centric	S	SHM	μLibrary	●	●	○	○	○	Protection Keys
	Code-centric		S	SHM	μLibrary	●	●	○	○	○	∅	
Sandbox	U	Enclosure [111]	Code-centric	S	MES	Function	●	●	○	○	○	∅ ⁵
		Google Sandboxed API (SAPI) [22]	Code-centric	S	MES	Function	●	●	○	○	○	∅ ⁵
		RLBox / μSWITCH [181, 195]	Hybrid	S	SHM	Function	●	●	○	○	●	∅
	U+K	Wedge [70]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	∅ ⁵
		CompartOS [55]	Code-centric	S	SHM	Function	●	●	○	○	○	∅ ⁵
		K	LVDs / KSplit [133]	Code-centric	S	SHM	Function	●	●	○	○	○
XFI/LXFI [107, 172]	Code-centric		S	SHM	Function	●	●	○	○	○	∅ ⁵	
Safebox	U	Nexen [221]	Code-centric	S	SHM	Function	●	●	○	○	○	Page Table ³
		Shreds [87]	Code-centric	S	SHM	Function	●	●	○	○	○ ⁴	∅ ⁵
	Dual World	Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
		Privtrans [73]	Code-centric	S	MES	Function	●	●	○	○	○	Page Table ²
		Swift [89]	Code-centric	S+A	MES	Any	●	●	○	○	●	∅
		Glamdring [165]	Code-centric	S	MES	Function	●	●	○	○	●	∅ ⁵
		PtrSplit / Program Mandering [169, 170]	Code-centric	S	MES	Function	●	●	○	○	○ ⁴	∅ ⁵
		DataShield [77]	Hybrid	S	SHM	Any	●	●	○	○	○ ⁴	Bounds Checking
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	Protection Keys
		K	Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○

Abstraction specialized for sandboxing untrusted code (such as libraries)

Implements cross-compartment communication with function calls, which is intuitive

Does not depend on a specific enforcement mechanism

Builds on shared memory to enforce confidentiality and integrity up cross-compartment data sharing

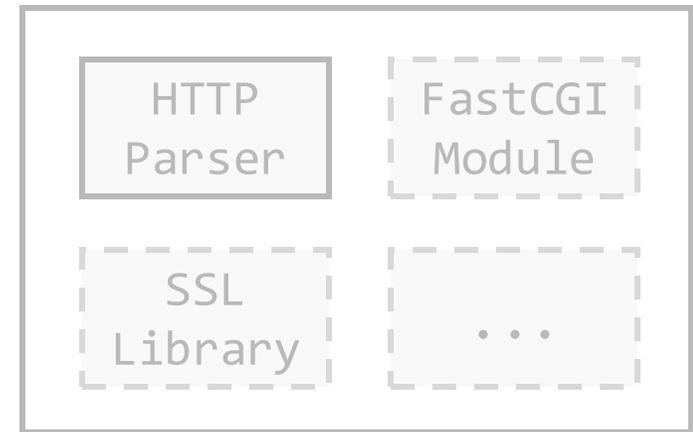
¹ Inherited from thread-like semantics, ² from process-like semantics, ³ from the Nested Kernel, ⁴ The PDM does (to a certain extent).
⁵ The abstraction could plug onto any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Open Problem in Compartmentalization Abstractions

Hardening compartment interfaces is key to obtaining strong security properties with compartmentalization

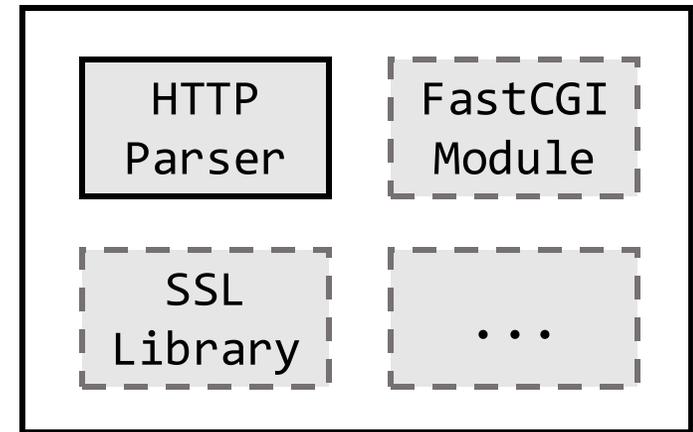
Example: assume the developer implemented our boundaries. Compartments *still need to communicate*, which opens an attack surface.



Open Problem in Compartmentalization Abstractions

Hardening compartment interfaces is key to obtaining strong security properties with compartmentalization

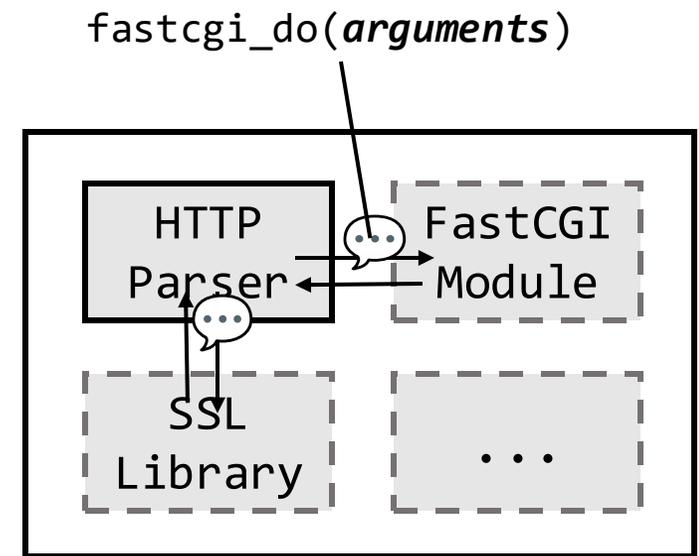
Example: assume the developer implemented our boundaries. Compartments *still need to communicate*, which opens an attack surface.



Open Problem in Compartmentalization Abstractions

Hardening compartment interfaces is key to obtaining strong security properties with compartmentalization

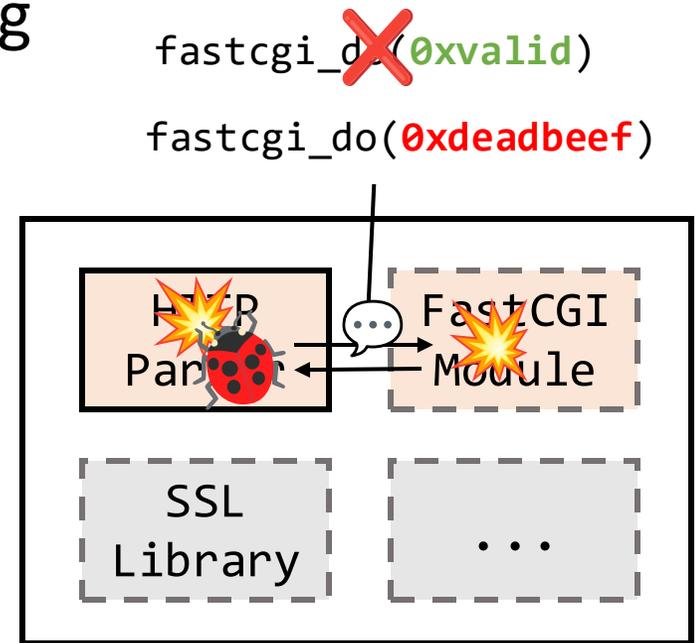
Example: assume the developer implemented our boundaries. Compartments *still need to communicate*, which opens an attack surface.



Open Problem in Compartmentalization Abstractions

Hardening compartment interfaces is key to obtaining strong security properties with compartmentalization

Example: assume the developer implemented our boundaries. Compartments *still need to communicate*, which opens an attack surface.



Open Problem in Compartmentalization Abstractions

Hardening compartment interfaces is key to obtaining strong security properties with compartmentalization

Example: assume the developer implemented our boundaries. Compartments *still need to communicate*, which opens an attack surface.

```

fastcgi_d(0xvalid)
fastcgi_do(0xdeadbeef)

```

(1)

Assessing the Impact of Interface Vulnerabilities in Compartmentalized Software

Hugo Lefeuvre¹, Vlad-Andrei Bădoiu², Yi Chien³, Felipe Hauck³, Nathan Dautenhahn³, Pierre Olivier¹

¹The University of Manchester, ²University Politehnica of Bucharest, ³Rice University, ⁴Unikraft

Abstract—Least-privilege separation decomposes applications into compartments limited to accessing only what they need. When compartmentalizing existing software, many approaches neglect securing the new inter-compartment interfaces, although what used to be a function call from a trusted component is now potentially a targeted attack from a malicious compartment. This results in an entire class of security bugs: **Compartment Interface Vulnerabilities (CIVs)**.

This paper provides an in-depth study of CIVs. We taxonomize these issues and show that they affect all known compartmentalization approaches. We propose **ConfFuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. We apply **ConfFuzz** to a set of 24 popular applications and 36 possible compartment APIs, to uncover a wide data-set of 629 vulnerabilities. We systematically study these issues, and extract numerous insights on the prevalence of CIVs, their causes, impact, and the complexity to address them. We stress the critical importance of CIVs in compartmentalization approaches, demonstrating an attack to extract isolated keys in OpenSSL, and uncovering a decade-old vulnerability in *uiskit*. We show, among others, that not all interfaces are affected in the same way, that API sites is uncorrelated with CIV prevalence, and that addressing interface vulnerabilities goes beyond writing simple checks. We conclude the paper with guidelines for CIV-aware compartment interface design, and appeal for more research towards systematic CIV detection and mitigation.

INTRODUCTION

The principle of least privilege has guided the design of safe computer systems for over half a century by ensuring that each unit of trust in a system can access only what it truly needs to fulfill its duties: in this way, system designers can proactively defend against unknown vulnerabilities [65]. Software compartmentalization is a prime example where unsafe, untrusted, or high-risk components are isolated to reduce the damage they would cause should they be compromised [50].

Recent years have seen the appearance of an increasingly large number of new isolation mechanisms [10], [4], [3], [65], [53], [45] that enable fine-grained compartmentalization. This resulted in compartmentalization works targeting finer and finer granularities, such as libraries [67], [60], [10], [42], [53], [35], [3], [51], [29], [2], modules [23], [3], [52], files [2], and even function-blocks of code [16], [64], [57], [1]. In that context, major attention was dedicated to compartmentalizing existing code, since rewriting software from scratch to work in a compartmentalized manner is costly and complex [16]. With

recent developments on compiler-based compartmentalization, frameworks offer to apply isolation at arbitrary interfaces for a low to non-existent porting cost [67], [5], [3], [5], [1].

Unfortunately, breaking down applications into compartments means that control and data dependencies through shared interfaces create new classes of vulnerabilities [61] in order to provide safe compartmentalization, it is not only necessary to ensure spatial memory isolation but also to design interfaces with distrust in mind. For example, objects passed through APIs can be corrupted to launch confused deputy attacks [39], [21], data structures can be manipulated to control execution or leak data through logic attacks [8], [11], called components can modify return values or indirectly access shared data structures to launch new forms of exploit, etc.

Even though interface-related vulnerabilities (denoted **Compartment-Interface Vulnerabilities (CIVs)** in this paper) were previously identified to various extents in the literature [39], [8], [21], [61], almost all modern compartmentalization frameworks [67], [60], [19], [53], [35], [23], [45], [5], [51], [57], [30], [29], [1] neglect the problem of securing interfaces, and rather focus on transparent and lightweight spatial separation. Since CIVs are already problematic for interfaces hardened from the ground up (e.g., the system call API [28], [8]) with well-defined trust models (kernel/users), their impact on safety is likely to be even greater when considering arbitrary interfaces and trust models that materialize when compartmentalizing existing software that was not designed with the assumption of hostile internal threats. Worse still, the complexity of safeguarding interfaces increases as more fine-grain components are targeted.

Beyond this lack of consideration, CIVs remain misunderstood: we ask the following research questions: *how widespread are CIVs when compartmentalizing unmodified applications? What are the API design patterns leading to them? What is the concrete impact of CIVs on the safety guarantees brought by compartmentalization, and what is the complexity of addressing them?* In order to achieve CIV mitigations that are generic and principled, we stress the need to formalize and quantify the problem.

This paper presents an in-depth study of CIVs. We taxonomize CIVs into a coherent framework, and systematize existing efforts to address them, highlighting categories that need attention in future research. In order to study existing CIVs in real-world scenarios, we propose **ConfFuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. **ConfFuzz** automatically explores the complexity of compartment interfaces by exposing data dependencies leading to vulnerabilities. Contrary to existing fuzzers, that inject malformed data in a single direction (e.g., a library),

Network and Distributed System Security (NDSS), Symposium 2023
27 February - 3 March 2023, San Diego, CA, USA
ISBN 1 891562 81 3
https://doi.org/10.4276/ndss.2023.26117
www.ndss-symposium.org

Open Problem in Compartmentalization

Abstraction

Hardening strong security boundaries

Example: a communication boundary

Most abstractions consider the hardening of compartment interfaces an orthogonal problem.

Class	Target U/K/BV	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties					Interface Safety	Design Bound to Mechanism		
				CALL	ASSIGN		C	L	A	R					
Manual Domain	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	○	Virtual Machine (EPT)		
		ACES [96]	Code-centric	S	SHM	Function	●	●	○	○	○	○	∅ ²		
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	○	∅ ²		
		TRIPOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	○	∅		
		CAPACTV [105]	Code-centric	S	SHM	Any	●	●	○	○	○	○	ARM PAC + MTE		
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	∅		
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Salus [239]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	○	Page Table ²		
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	○	Page Table		
		SOAIP [117]	Hybrid	S	SHM	Any	●	●	○	○	○	○	∅		
		libMPK [199]	Hybrid	S	SHM	Any	●	●	○	○	○	○	∅		
		CheriOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	○	Protection Keys CHERI		
		U+K	U+K	Microkernel Servers [106], [1...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	○	∅ ²
Mutable Protection Domains (MPDs) [191], [92]	Code-centric			S	SHM	U/K-component	●	●	○	○	○	○	∅ ²		
ReelLeaf [184]	Code-centric			S	SHM	U/K-component	●	●	○	○	○	○	∅		
CubicleOS [211]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	○	Safe Languages Protection Keys		
FlexOS [161]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	○	∅		
SMP [196]	Code-centric			S	SHM	Any	●	●	○	○	○	○	∅		
Monza [197]	Hybrid			A	SHM	Function ¹	○	○	○	○	○	○	∅ ²		
VirtuOS [186]	Code-centric			S+A	SHM	K-component	●	●	○	○	○	○	Virtual Machine (EPT)		
HAKC [173]	Code-centric			S	SHM	Function	●	●	○	○	○	○	ARM PAC + MTE		
LibertOS [187]	Code-centric			S	SHM	K-component	●	●	○	○	○	○	∅ ²		
Sandbox	U	Call [61]	Code-centric	S	SHM	Library	●	●	○	○	○	○	∅ ²		
		CompaRTEK [132]	Code-centric	S	MES	Library	●	●	○	○	○	○	∅ ²		
		Enclosers [111]	Code-centric	S	SHM	Package	●	●	○	○	○	○	∅ ²		
		Google Sandboxed API (SAPi) [122]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²		
		RLBox / μSWITCH [181], [85]	Hybrid	S	SHM	Function	●	●	○	○	○	○	∅		
		Wedge [79]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		CompartOS [151]	Code-centric	S	SHM	Linkage Unit	●	●	○	○	○	○	∅ ²		
		K	K	LVIDs / KSplit [133], [185]	Code-centric	S	MES	K-component	●	●	○	○	○	○	∅ ²
				XPLXFI [167], [172]	Hybrid	S	SHM	K-component	●	●	○	○	○	○	SFI
		HV	Neven [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	○	Page Table ²	
Software Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	○	○	∅ ²		
		Privman [147]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²		
		Privtrans [72]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²		
		SwH [89]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	∅		
		Glandring [165]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²		
		PlusSplit / Program Mandering [169], [76]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²		
		DataShield [77]	Hybrid	Any	SHM	Any	●	●	○	○	○	○	∅ ²		
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	○	Protection Keys		
		Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	○	Page Table		

¹Inherited from thread-like semantics, ²from process-like semantics, ³from the Nested Kernel, ⁴The PDM does (to a certain extent).
⁵The abstraction could plug into any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

Abstraction is key to obtaining compartmentalization

Implemented our

needed to check surface.

```
fastcgi_do(0xvalid)
fastcgi_do(0xdeadbeef)
```

(1)

Assessing the Impact of Interface Vulnerabilities in Compartmentalized Software

Hugo Lefeuvre¹, Vlad-Andrei Bădoiu², Yi Chien³, Felipe Hauck⁴, Nathan Dautenhahn⁵, Pierre Olivier⁶

¹The University of Manchester, ²University Politehnica of Bucharest, ³Rice University, ⁴Unikraft

Abstract—Least privilege separation decomposes applications into compartments, limited to accessing only what they need. When compartmentalizing existing software, many approaches neglect securing the new inter-compartment interfaces, although what used to be a function call from a trusted component is now potentially a targeted attack from a malicious component. This results in an entire class of security bugs: **Compartment Interface Vulnerabilities (CIVs)**.

This paper provides an in-depth study of CIVs. We taxonomize these issues and show that they affect all known compartmentalization approaches. We propose **ConfPuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. We apply **ConfPuzz** to a set of 21 popular applications and 36 possible compartment APIs, to uncover a wide data-set of 629 vulnerabilities. We systematically study these issues, and extract numerous insights on the prevalence of CIVs, their causes, impact, and the complexity to address them. We stress the critical importance of CIVs in compartmentalization approaches, demonstrating an attack to various extents in the literature [195, 18, 211, 161], almost all modern compartmentalization frameworks [671, 660, 119, 153, 123, 143, 151, 151, 157], [30], [29], [11] neglect the problem of securing interfaces, and rather focus on transparent and lightweight spatial separation. Since CIVs are already problematic for interfaces hardened from the ground up (e.g., the system call API [28], [8]) with well-defined trust models (kernel/users), their impact on safety is likely to be even greater when considering arbitrary interfaces and trust models that materialize when compartmentalizing existing software that was not designed with the assumption of hostile internal threats. While still the complexity of safeguarding interfaces increases as more fine-grain components are targeted.

Beyond this lack of consideration, CIVs remain misunderstood: we ask the following research questions: *how widespread are CIVs when compartmentalizing unmodified applications? What are the API design patterns leading to them? What is the concrete impact of CIVs on the safety guarantees brought by compartmentalization, and what is the complexity of addressing them?* In order to achieve CIV mitigations that are generic and principled, we stress the need to formalize and quantify the problem.

This paper presents an in-depth study of CIVs. We taxonomize CIVs into a coherent framework, and systematically existing efforts to address them, highlighting categories that need attention in future research. In order to study existing CIVs in real-world scenarios, we propose **ConfPuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. **ConfPuzz** automatically explores the complexity of compartment interfaces by exposing data dependencies leading to vulnerabilities. Contrary to existing fuzzers, that inject malformed data in a single direction (e.g., a library,

INTRODUCTION

The principle of least privilege has guided the design of safe computer systems for over half a century by ensuring that each unit of trust in a system can access only what it truly needs to fulfill its duties: in this way, system designers can proactively defend against unknown vulnerabilities [65]. Software compartmentalization is a prime example where unsafe, untrusted, or high-risk components are isolated to reduce the damage they would cause should they be compromised [50].

Recent years have seen the appearance of an increasingly large number of new isolation mechanisms [101, 141, 131, 165], [53], [45], [43] that enable fine-grained compartmentalization. This resulted in compartmentalization works targeting finer and finer granularities, such as libraries [671, 660], [142], [151], [153], [151], [211], [21], modules [22], [21], [52], files [16], and even functions/blocks of code [146], [641], [571], [1]. In that context, major attention was dedicated to compartmentalizing existing code, since rewriting software from scratch to work in a compartmentalized manner is costly and complex [116]. With

Network and Distributed System Security (NDSS) Symposium 2023
 27 February - 3 March 2023, San Diego, CA, USA
 ISBN 1 891562 81 3
 https://doi.org/10.4272/ndss.2023.2017
 www.ndss-symposium.org

Open Problem in Compartmentalization

Abstraction

Hardening

strong security

Example: a

boundary

communicate

Class	Target U/K/B/V	Abstraction	Subject Selection	Semantics		Abstraction Granularity		Properties				Interface Safety	Design Bound to Mechanism		
				CALL	ASSIGN	C	L	A	R	U	K			B	V
Manual Domain	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	○	○	Virtual Machine (EPT)	
		ACES [96]	Code-centric	S	SHM	Function	●	●	○	○	○	○	○	∅ ²	
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	○	○	∅ ²	
		TRIPPOK [124]	Code-centric	S	SHM	Library	●	●	○	○	○	○	○	∅ ²	
		CAPACTIV [105]	Code-centric	S	SHM	Any	●	●	○	○	○	○	○	ARM PAC + MTE	
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	○	∅ ²	
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	○	∅ ²	
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	○	∅ ²	
		Salus [226]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	○	∅ ²	
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	○	○	Page Table ²	
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	○	○	Page Table	
		SOAIP [117]	Hybrid	S	SHM	Any	●	●	○	○	○	○	○	∅ ²	
libMPK [199]	Hybrid	S	SHM	Any	●	●	○	○	○	○	○	Protection Keys			
CherOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	○	○	CHERI			
U+K	U+K	Microkernel Servers [106], [1...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	○	∅ ²		
		Mutable Protection Domains (MPDs) [191], [92]	Code-centric	S	SHM	U/K-component	●	●	○	○	○	○	∅ ²		
		ReelLeaf [184]	Code-centric	S	SHM	U/K-component	●	●	○	○	○	○	∅ ²		
		CubicleOS [211]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	○	∅ ²		
		FlexOS [161]	Code-centric	S	SHM	μLibrary	●	●	○	○	○	○	∅ ²		
		SMP [196]	Code-centric	S	SHM	Any	●	●	○	○	○	○	∅ ²		
		Monza [197]	Hybrid	A	SHM	Function ¹	○	○	○	○	○	○	∅ ²		
		VirtuOS [186]	Code-centric	S+A	SHM	K-component	●	●	○	○	○	○	○	Virtual Machine (EPT)	
		HAKC [173]	Code-centric	S	SHM	Function	●	●	○	○	○	○	○	ARM PAC + MTE	
		LibertOS [187]	Code-centric	S	SHM	K-component	●	●	○	○	○	○	○	∅ ²	
		Sandbox	U	Call [61]	Code-centric	S	SHM	Library	●	●	○	○	○	○	∅ ²
				CompaRTK [132]	Code-centric	S	MES	Library	●	●	○	○	○	○	∅ ²
Enclosures [111]	Code-centric			S	SHM	Package	●	●	○	○	○	○	∅ ²		
Google Sandboxed API (SAPH) [122]	Code-centric			S	MES	Function	●	●	○	○	○	○	∅ ²		
RLBox / μSWITCH [181], [85]	Hybrid			S	SHM	Function	●	●	○	○	○	○	∅ ²		
Wedge [79]	Hybrid			S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
CompartOS [151]	Code-centric			S	SHM	Linkage Unit	●	●	○	○	○	○	○	CHERI	
K	K			LVIDs / KSplit [133], [185]	Code-centric	S	MES	K-component	●	●	○	○	○	○	∅ ²
				XPLXPL [167], [172]	Hybrid	S	SHM	K-component	●	●	○	○	○	○	∅ ²
HV	Neven [221]			Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	○	∅ ²	
Software Dual World	U			Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	○	○	∅ ²
				Privan [147]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²
		Privtrans [72]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²		
		SwH [89]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	∅ ²		
		Glandring [165]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²		
		PuSplit / Program Mandering [169], [76]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²		
		DataShield [77]	Hybrid	Any	Any	Any	●	●	○	○	○	○	∅ ²		
		ERIM [232]	Hybrid	S	SHM	Any	●	●	○	○	○	○	Protection Keys		
		Nested Kernel [94]	Code-centric	S	SHM	Function	●	●	○	○	○	○	Page Table		

¹Inherited from thread-like semantics, ²from process-like semantics, ³from the Nested Kernel, ⁴The PDM does (to a certain extent).
⁵The abstraction could plug into any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

is key to obtaining

compartmentalization

implemented our

ed to

back surface.

Most abstractions consider the hardening of compartment interfaces an orthogonal problem.

We stress that abstractions can and should facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?

```
fastcgi_do(0xvalid)
fastcgi_do(0xdeadbeef)
```

Assessing the Impact of Interface Vulnerabilities in Compartmentalized Software

Hugo Lefeuvre¹, Vlad-Andrei Bădoiu², Yi Chien³, Felipe Hauck⁴, Nathan Dautenhahn⁵, Pierre Olivier¹

¹The University of Manchester, ²University Politehnica of Bucharest, ³Rice University, ⁴Unikraft

Abstract—Last-privilege separation decomposes applications into compartments, limited to accessing only what they need. When compartmentalizing existing software, many approaches neglect securing the new inter-compartment interfaces, although what used to be a function call from a trusted component is now potentially a targeted attack from a malicious component. This results in an entire class of security bugs: **Compartment Interface Vulnerabilities (CIVs)**.

This paper provides an in-depth study of CIVs. We taxonomize these issues and show that they affect all known compartmentalization approaches. We propose **ConfPuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. We apply **ConfPuzz** to a set of 21 popular applications and 36 possible compartment APIs, to uncover a wide data-set of 629 vulnerabilities. We systematically study these issues, and extract numerous insights on the prevalence of CIVs, their causes, impact, and the complexity to address them. We stress the critical importance of CIVs in compartmentalization approaches, demonstrating an attack to extract isolated keys in OpenSSL, and uncovering a decade-old vulnerability in *uclibc*. We show, among others, that not all interfaces are affected in the same way, that API sites are uncorrelated with CIV prevalence, and that addressing interface vulnerabilities goes beyond writing simple checks. We conclude the paper with guidelines for CIV-aware compartment interface design, and appeal for more research towards systematic CIV detection and mitigation.

1. INTRODUCTION

The principle of least privilege has guided the design of safe computer systems for over half a century by ensuring that each unit of trust in a system can access only what it truly needs to fulfill its duties: in this way, system designers can proactively defend against unknown vulnerabilities [65]. Software compartmentalization is a prime example where unsafe, untrusted, or high-risk components are isolated to reduce the damage they would cause should they be compromised [50].

Recent years have seen the appearance of an increasingly large number of new isolation mechanisms [10], [4], [13], [65], [51], [57], [19], [29], [2], [21], [22], [23], [52], [5], [15], and even function-blocks of code [16], [64], [57], [1]. In that context, major attention was dedicated to compartmentalizing existing code, since rewriting software from scratch to work in a compartmentalized manner is costly and complex [16]. With

recent developments on compiler-based compartmentalization, frameworks offer to apply isolation at arbitrary interfaces for a low to non-existent porting cost [67], [5], [15], [35], [1].

Unfortunately, breaking down applications into compartments means that control and data dependencies through shared interfaces create new classes of vulnerabilities [61] in order to provide safe compartmentalization, it is not only necessary to ensure spatial memory isolation but also to design interfaces with distrust in mind. For example, objects passed through APIs can be corrupted to launch confused deputy attacks [39], [21], data structures can be manipulated to control execution or leak data through lags attacks [8], [11], called components can modify return values or indirectly access shared data structures to launch new forms of exploit, etc.

Even though interface-related vulnerabilities (denoted **Compartment-Interface Vulnerabilities (CIVs)** in this paper) were previously identified to various extents in the literature [39], [8], [21], [61], almost all modern compartmentalization frameworks [67], [60], [19], [51], [23], [45], [15], [51], [57], [30], [29], [11] neglect the problem of securing interfaces, and rather focus on transparent and lightweight spatial separation. Since CIVs are already problematic for interfaces hardened from the ground up (e.g., the system call API [28], [8]) with well-defined trust models (kernel/users), their impact on safety is likely to be even greater when considering arbitrary interfaces and trust models that materialize when compartmentalizing existing software that was not designed with the assumption of hostile internal threats. Worse still, the complexity of safeguarding interfaces increases as more fine-grain components are targeted.

Beyond this lack of consideration, CIVs remain misunderstood: we ask the following research questions: *how widespread are CIVs when compartmentalizing unmodified applications? What are the API design patterns leading to them? What is the concrete impact of CIVs on the safety guarantees brought by compartmentalization, and what is the complexity of addressing them?* In order to achieve CIV mitigations that are generic and principled, we stress the need to formalize and quantify the problem.

This paper presents an in-depth study of CIVs. We taxonomize CIVs into a coherent framework, and systematically existing efforts to address them, highlighting categories that need attention in future research. In order to study existing CIVs in real-world scenarios, we propose **ConfPuzz**, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. **ConfPuzz** automatically explores the complexity of compartment interfaces by exposing data dependencies leading to vulnerabilities. Contrary to existing fuzzers, that inject malformed data in a single direction (e.g., a library),

NDSS'23, March 27-29, 2023, San Diego, CA, USA
 ISBN 978-1-961661-81-3
 https://doi.org/10.47529/ndss.2023.24117
 www.ndss-symposium.org

Open Problem in Compartmentalization

Abstraction

Hardening

strong security

Example: a

boundary

communicate

is key to obtaining

compartmentalization

implemented our

ed to

black surface.

Most abstractions consider the hardening of

We stress that abstractions can and should

How can abstractions that facilitate the implementation of secure domain interfaces?

Class	Target U/K/B/V	Abstraction	Subject Selection	Semantics		Abstraction Granularity	Properties					Interface Safety	Design Bound to Mechanism		
				CALL	ASSIGN		C	L	A	R					
Manual Domain	U	Virtines [242]	Code-centric	S	MES	Function	●	●	○	○	○	○	Virtual Machine (EPT)		
		ACES [96]	Code-centric	S	SHM	Function	●	●	○	○	○	○	∅ ²		
		SeCage [171]	Code-centric	S	SHM	Function	●	●	○	○	○	○	∅ ²		
		TRIPOR [124]	Code-centric	S	SHM	Library	●	●	○	○	○	○	∅ ²		
		CAPACTIV [105]	Code-centric	S	SHM	Any	●	●	○	○	○	○	ARM PAC + MTE		
		Jif [261]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	∅ ²		
		Arbiter [241]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Secure Memory Views (SMVs) [128]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Salus [239]	Data-centric	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		Light-Weight Contexts (LwCs) [167]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	○	Page Table ²		
		POSIX Processes (and earlier instances) [93]	Hybrid	Any	Any	Any	●	●	○	○	○	○	Page Table		
		SOAIP [117]	Hybrid	S	SHM	Any	●	●	○	○	○	○	∅ ²		
		libMPK [199]	Hybrid	S	SHM	Any	●	●	○	○	○	○	Protection Keys		
		CherOS [108]	Code-centric	Any	Any	U/K-component	●	●	○	○	○	○	CHERI		
		Sandbox	U+K	Microkernel Servers [106], [1...]	Code-centric	Any	MES	U/K-component	●	●	○	○	○	○	∅ ²
Mutable Protection Domains (MPDs) [191], [192]	Code-centric			S	SHM	U/K-component	●	●	○	○	○	○	∅ ²		
ReelLeaf [184]	Code-centric			S	SHM	U/K-component	●	●	○	○	○	○	Safe Languages		
CubicleOS [211]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	○	Protection Keys		
FlexOS [161]	Code-centric			S	SHM	μLibrary	●	●	○	○	○	○	∅		
SMP [196]	Code-centric			S	SHM	Any	●	●	○	○	○	○	∅		
Monza [197]	Hybrid			A	SHM	Function ¹	○	○	○	○	○	○	∅ ²		
VirtuOS [186]	Code-centric			S+A	SHM	K-component	●	●	○	○	○	○	Virtual Machine (EPT)		
HAKC [175]	Code-centric			S	SHM	Function	●	●	○	○	○	○	ARM PAC + MTE		
LibertOS [187]	Code-centric			S	SHM	K-component	●	●	○	○	○	○	∅ ²		
Software Dual World	U			Call [61]	Code-centric	S	SHM	Library	●	●	○	○	○	○	∅ ²
				CompaRTK [132]	Code-centric	S	MES	Library	●	●	○	○	○	○	∅ ²
				Enclosures [111]	Code-centric	S	SHM	Package	●	●	○	○	○	○	∅ ²
				Google Sandboxed API (SAPH) [122]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²
				RLBox / μSWITCH [181], [185]	Hybrid	S	SHM	Function	●	●	○	○	○	○	∅ ²
		Wedge [79]	Hybrid	S	SHM	Function ¹	●	●	○	○	○	○	∅ ²		
		CompartOS [151]	Code-centric	S	SHM	Linkage Unit	●	●	○	○	○	○	CHERI		
		K	LVIDs / KSplit [133], [185]	Code-centric	S	MES	K-component	●	●	○	○	○	○	∅ ²	
			XPLXLFI [167], [172]	Hybrid	S	SHM	K-component	●	●	○	○	○	○	SFI	
			Nexen [221]	Data-centric	S	SHM	Per-VM domain	●	●	○	○	○	○	Page Table ²	
		Dual World	U	Shreds [87]	Code-centric	S	SHM	Any	●	●	○	○	○	○	∅ ²
				Privan [147]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²
				Privans [72]	Code-centric	S	MES	Function	●	●	○	○	○	○	Page Table ²
				SwH [89]	Code-centric	S+A	MES	Any	●	●	○	○	○	○	∅
				Glandring [165]	Code-centric	S	MES	Function	●	●	○	○	○	○	∅ ²
PuSplit / Program Mandering [169], [170]	Code-centric			S	MES	Function	●	●	○	○	○	○	∅ ²		
DataShield [77]	Hybrid			Any	Any	Any	●	●	○	○	○	○	Bounds Checking		
ERIM [232]	Hybrid			S	SHM	Any	●	●	○	○	○	○	Protection Keys		
Nested Kernel [94]	Code-centric			S	SHM	Function	●	●	○	○	○	○	Page Table		

¹Inherited from thread-like semantics, ²from process-like semantics, ³from the Nested Kernel, ⁴The PDM does (to a certain extent).
⁵The abstraction could plug into any intra-AS mechanism, though the paper or documentation claims reliance on a particular one.

```
fastcgi_do(0xvalid)
fastcgi_do(0xdeadbeef)
```

Assessing the Impact of Interface Vulnerabilities in Compartmentalized Software

Hugo Lefeuve¹, Vlad-Andrei Bădoiu², Yi Chien³, Felipe Hauck⁴, Nathan Dautenhahn⁵, Pierre Olivier¹

¹The University of Manchester, ²University Politehnica of Bucharest, ³Rice University, ⁴Unikraft

Abstract—Last-privilege separation decomposes applications into compartments, limited to accessing only what they need. When compartmentalizing existing software, many approaches neglect securing the new inter-compartment interfaces, although what used to be a function call from a trusted component is now potentially a targeted attack from a malicious component. This results in an entire class of security bugs: *Compartment Interface Vulnerabilities (CIVs)*.

This paper provides an in-depth study of CIVs. We taxonomize these issues and show that they affect all known compartmentalization approaches. We propose *Confuzz*, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. We apply *Confuzz* to a set of 21 popular applications and 36 possible compartment APIs, to uncover a wide data-set of 629 vulnerabilities. We systematically study these issues, and extract numerous insights on the prevalence of CIVs, their causes, impact, and the complexity to address them. We stress the critical importance of CIVs in compartmentalization approaches, demonstrating an attack to extract isolated keys in OpenSSL, and uncovering a decade-old vulnerability in *uak*. We show, among others, that not all interfaces are affected in the same way, that API sites are uncorrelated with CIV prevalence, and that addressing interface vulnerabilities goes beyond writing simple checks. We conclude the paper with guidelines for CIV-aware compartment interface design, and appeal for more research towards systematic CIV detection and mitigation.

1. INTRODUCTION

The principle of least privilege has guided the design of safe computer systems for over half a century by ensuring that each unit of trust in a system can access only what it truly needs to fulfill its duties: in this way, system designers can proactively defend against unknown vulnerabilities [65]. Software compartmentalization is a prime example where unsafe, untrusted, or high-risk components are isolated to reduce the damage they would cause should they be compromised [50].

Recent years have seen the appearance of an increasingly large number of new isolation mechanisms [10], [4], [3], [65], [53], [45], [14], [29], [72], [modules 22], [2], [52], [files 18], and even function-blocks of code [16], [64], [57], [1]. In that context, major attention was dedicated to compartmentalizing existing code, since rewriting software from scratch to work in a compartmentalized manner is costly and complex [16]. With

recent developments on compiler-based compartmentalization, frameworks offer to apply isolation at arbitrary interfaces for a low to non-existent porting cost [67], [5], [35], [1].

Unfortunately, breaking down applications into compartments means that control and data dependencies through shared interfaces create new classes of vulnerabilities [61] in order to provide safe compartmentalization, it is not only necessary to ensure spatial memory isolation but also to design interfaces with distrust in mind. For example, objects passed through APIs can be corrupted to launch confused deputy attacks [39], [21], data structures can be manipulated to control execution or leak data through lags attacks [8], [11], called components can modify return values or indirectly access shared data structures to launch new forms of exploit, etc.

Even though interface-related vulnerabilities (denoted *Compartment-Interface Vulnerabilities (CIVs)* in this paper) were previously identified to various extents in the literature [39], [8], [21], [61], almost all modern compartmentalization frameworks [67], [60], [19], [53], [25], [45], [5], [51], [57], [30], [29], [11] neglect the problem of securing interfaces, and rather focus on transparent and lightweight spatial separation. Since CIVs are already problematic for interfaces hardened from the ground up (e.g., the system call API [28], [8]) with well-defined trust models (kernel/users), their impact on safety is likely to be even greater when considering arbitrary interfaces and trust models that materialize when compartmentalizing existing software that was not designed with the assumption of hostile internal threats. Worse still, the complexity of safeguarding interfaces increases as more fine-grain components are targeted.

Beyond this lack of consideration, CIVs remain misunderstood: we ask the following research questions: *how widespread are CIVs when compartmentalizing unmodified applications? What are the API design patterns leading to them? What is the concrete impact of CIVs on the safety guarantees brought by compartmentalization, and what is the complexity of addressing them?* In order to achieve CIV mitigations that are generic and principled, we stress the need to formalize and quantify the problem.

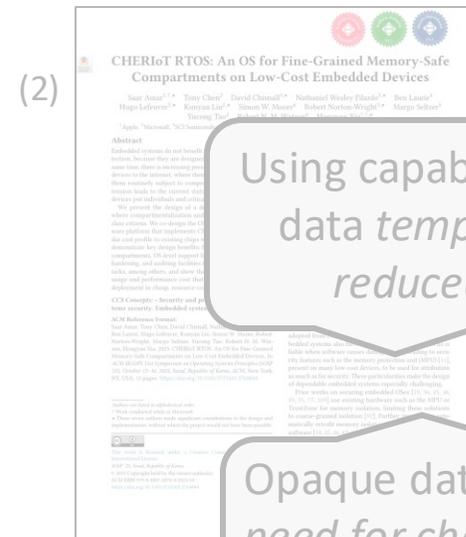
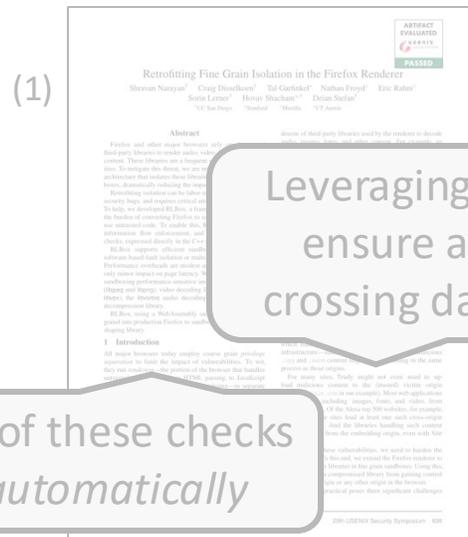
This paper presents an in-depth study of CIVs. We taxonomize CIVs into a coherent framework, and systematically existing efforts to address them, highlighting categories that need attention in future research. In order to study existing CIVs in real-world scenarios, we propose *Confuzz*, an in-memory fuzzer specialized to detect CIVs at possible compartment boundaries. *Confuzz* automatically explores the complexity of compartment interfaces by exposing data dependencies leading to vulnerabilities. Contrary to existing fuzzers, that inject malformed data in a single direction (e.g., a library),

NDSS'23, February 27 - March 2, 2023, San Diego, CA, USA
 ISBN 978-1-949875-81-3
 https://doi.org/10.4171/ndss.2023.2417
 www.ndss-symposium.org

Open Problem in Compartmentalization Abstractions

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.

Some do!



More work needed in that direction.

(2) Amar et al., CHERIOT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

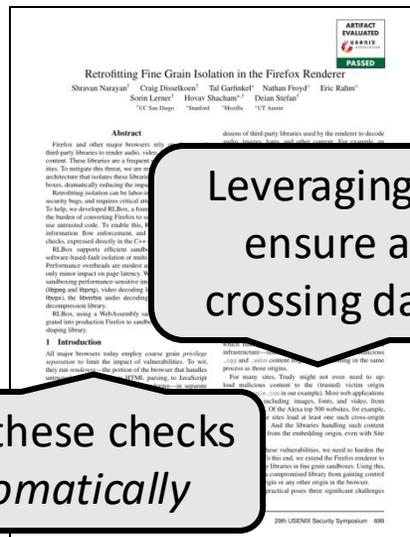
(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Open Problem in Compartmentalization Abstractions

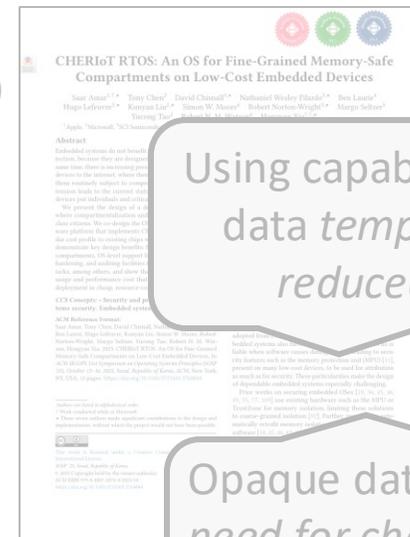
Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.

Some do!

(1)



(2)



More work needed in that direction.

(2) Amar et al., CHERIoT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

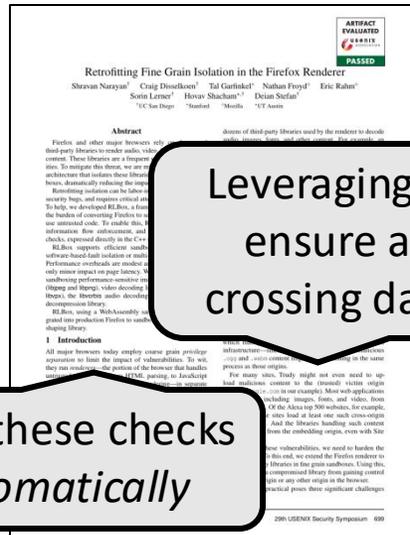
(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Open Problem in Compartmentalization Abstractions

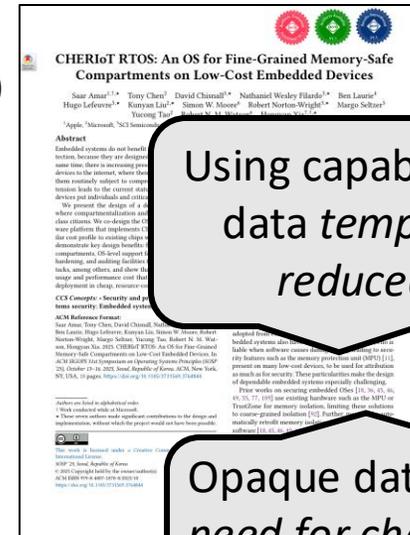
Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.

Some do!

(1)



(2)



More work needed in that direction.

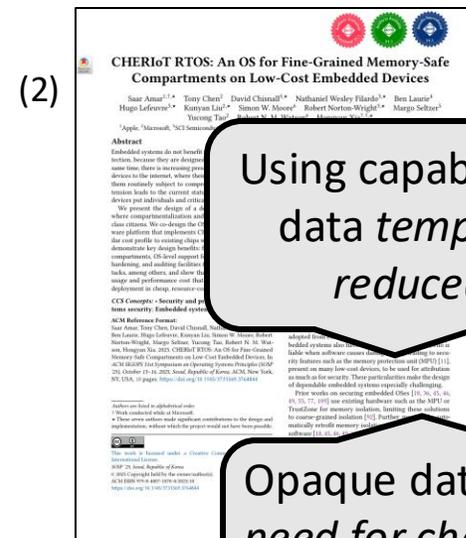
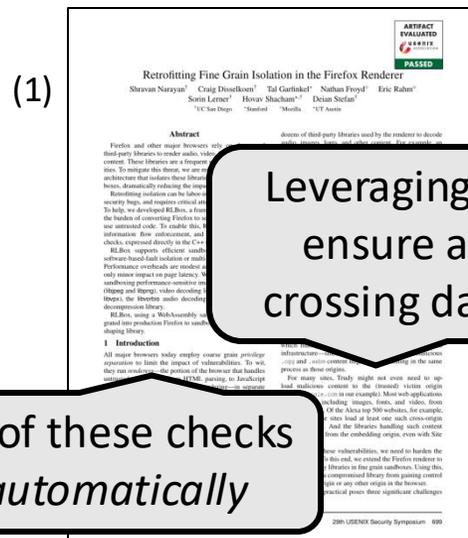
(2) Amar et al., CHERIOT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Open Problem in Compartmentalization Abstractions

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.

Some do!



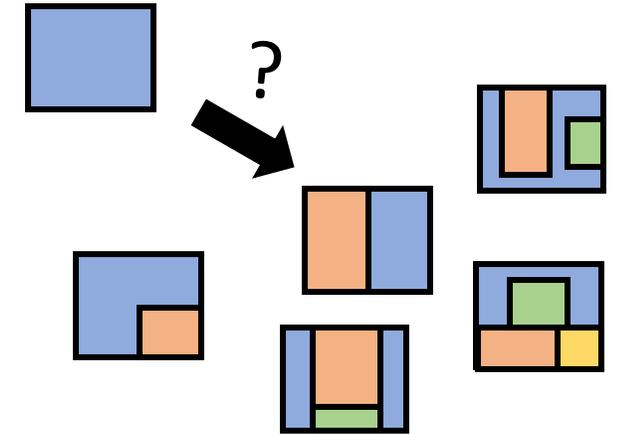
More work needed in that direction.

(2) Amar et al., CHERIOT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Bigger Picture of Compartmentalization

Propose to view compartmentalization as 3 problems:

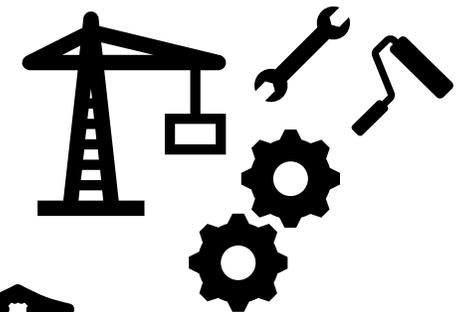


1. How to determine the right policy to enforce?

- Done with a *policy definition method*

2. How to integrate the notion of compartmentalization policies in software / programming models / idioms?

- Done with a *compartmentalization abstraction*



3. How to enforce policies at runtime?

- Done with a *compartmentalization mechanism*



Problem #3

How to enforce policies?

Finally, we must **enforce the policy at runtime.**

This is achieved with an **enforcement mechanism.**

Problem #3

How to enforce policies?

Finally, we must **enforce the policy at runtime**.

This is achieved with an **enforcement mechanism**.

Page tables (via the MMU) are the historical enforcement mechanism.

Problem #3

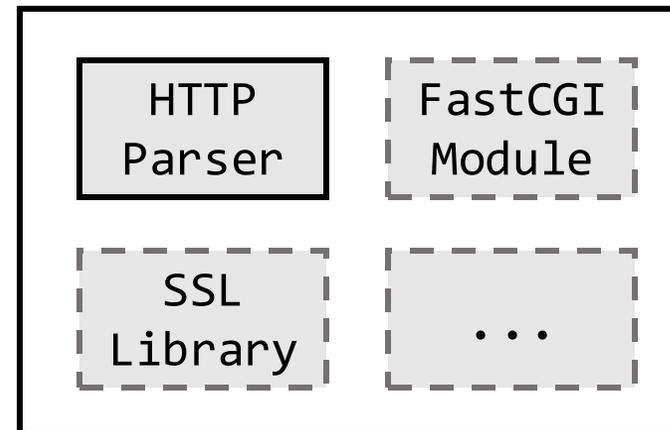
How to enforce policies?

Finally, we must **enforce the policy at runtime**.

This is achieved with an **enforcement mechanism**.

Page tables (via the MMU) are the historical enforcement mechanism.

Example from earlier: *assume we implemented the sandbox with processes.*



(Fictive monolithic program)

Problem #3

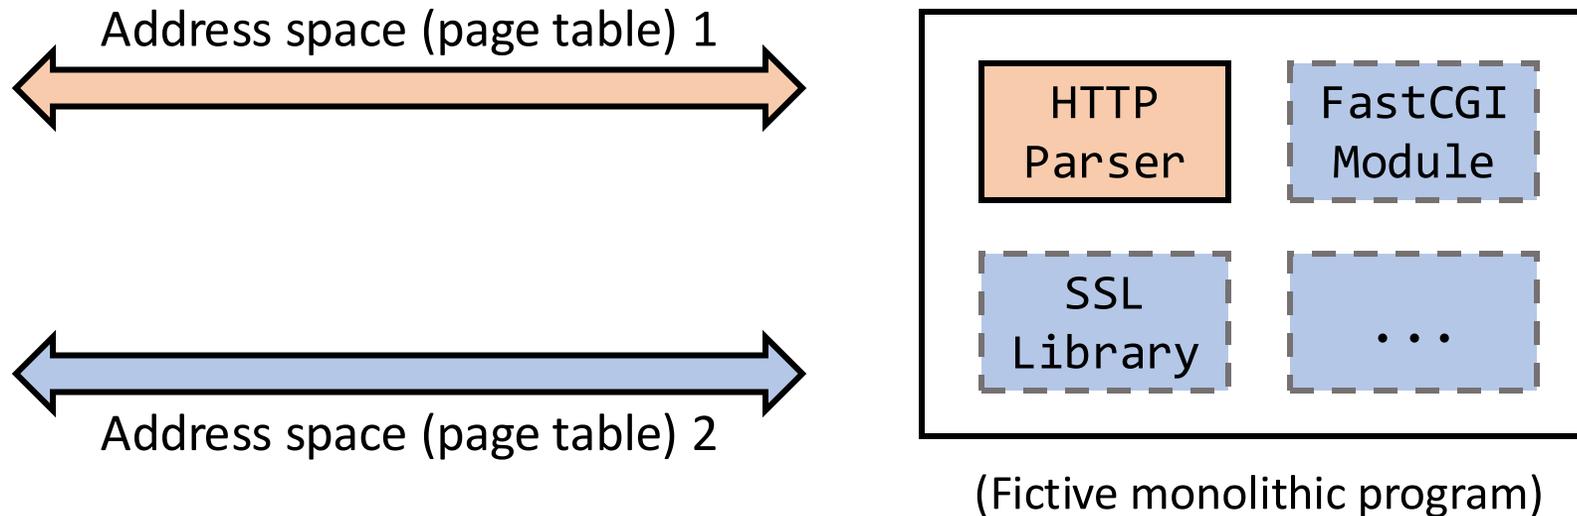
How to enforce policies?

Finally, we must **enforce the policy at runtime**.

This is achieved with an **enforcement mechanism**.

Page tables (via the MMU) are the historical enforcement mechanism.

Example from earlier: *assume we implemented the sandbox with processes.*



Problem #3

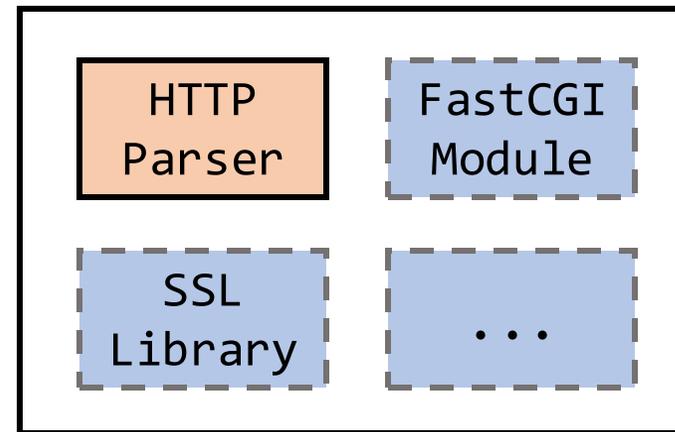
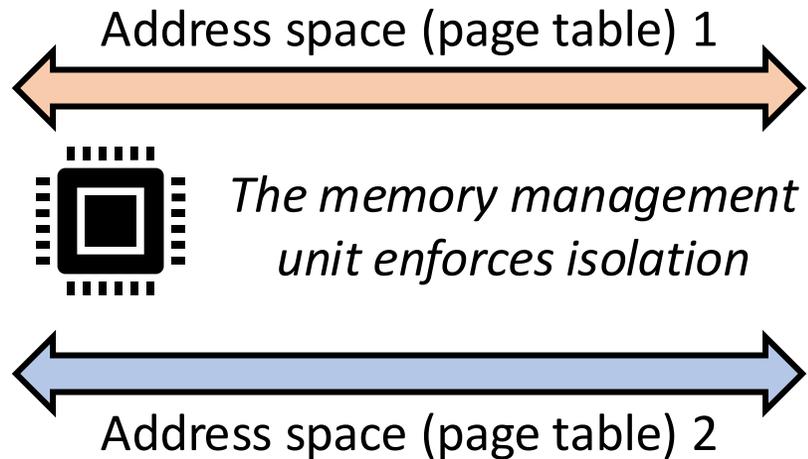
How to enforce policies?

Finally, we must **enforce the policy at runtime**.

This is achieved with an **enforcement mechanism**.

Page tables (via the MMU) are the historical enforcement mechanism.

Example from earlier: *assume we implemented the sandbox with processes.*



(Fictive monolithic program)

Problem #3

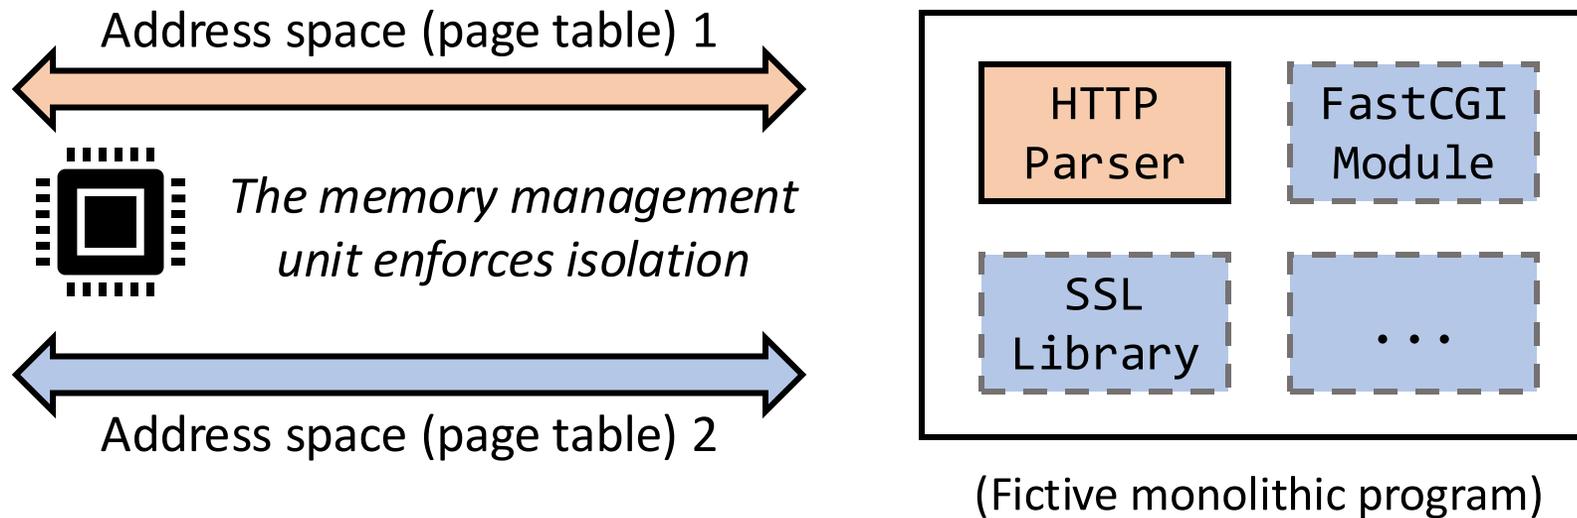
How to enforce policies?

Finally, we must **enforce the policy at runtime**.

This is achieved with an **enforcement mechanism**.

Page tables (via the MMU) are the historical enforcement mechanism.

Example from earlier: *assume we implemented the sandbox with processes.*



This is still the most common way to do it today.

Enforcement Mechanisms

Enforcement mechanisms (for compartmentalization) are also a hot area:

- Maximize performance
- Minimize hardware cost and complexity
- Enable for stronger security properties

TABLE 3: Taxonomy of Compartmentalization Mechanisms. Page-Table = PT; Permissions: Read, Write, Execute, Address (create pointers to), ● = supported, ◐ = supported by some, ○ = unsupported; Overhead: free = ○ < ◐ < ◑ < ◒ < ◓ = very costly.

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	◐ ¹	Mutual	Full	●	●	●	○	Page	∞	● (PT switch + ³)
		Supervisor Bit [25, 158, 159]	○	Single	Full	●	●	●	○	Page	2 (kernel/OS)	● (interrupt + ³)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	●	●	●	○	Word	∞	● (MMP hardware + ³)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	● (special register flip + ³)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	●	○	Byte - Page [178]	2 (safe/unsafe)	● (r ⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ³)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ³)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	○	Byte	∞	● (special instr. + ³)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	●	●	●	○	Byte	∞	● (bounds hardware + ³)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	●	●	●	○	Byte - Words ³	16 ³ - ∞ [138]	● (tagging hardware + ³)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ³)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ³)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	○	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	○	Byte	∞	● (r ⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ³)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.



(1)

(1) Lefevre et al., SoK: Software Compartmentalization, S&P 2025

Enforcement Mechanisms

Enforcement mechanisms (for compartmentalization) are also a hot area:

Maximize performance

Minimize hardware cost and complexity

Enable for stronger security properties

TABLE 3: Taxonomy of Compartmentalization Mechanisms. Page-Table = PT; Permissions: Read, Write, Execute, Address (create pointers to), ● = supported, ◐ = supported by some, ○ = unsupported; Overhead: free = ○ < ◐ < ◑ < ◒ < ◓ = very costly.

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N° of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N° of machines	● - ● (link latency)
	Access Bits [25], EPT / vmfunc [26]	◐ ¹	Mutual	Full	●	●	●	○	Page	∞	● (PT switch + ²)
	Supervisor Bit [25, 158, 159]	○	Single	Full	●	●	●	○	Page	2 (kernel/user)	● (interrupt + ²)
	Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	●	●	●	○	Word	∞	● (MMP hardware + ³)
	Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	◐ (special register flip + ³)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	● (link latency)
	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
	Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
	World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	◐ (special instr. + ³)
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	●	●	●	○	Byte	∞	◐ (bounds hardware + ³)	
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	●	●	○	○	Byte - Words ⁴	16 ⁴ - ∞ [138]	◐ (tagging hardware + ³)	
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	● - ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	●	Byte	∞	● - ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	◐ (link latency)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ³)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.



(1)

(1) Lefevre et al., SoK: Software Compartmentalization, S&P 2025

	Mechanism Class	Condi- tioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Names of the mechanisms we consider

Characteristics we included in the taxonomy

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Names of the mechanisms we consider

Characteristics we included in the taxonomy

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)
	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	PT Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)
	TEE Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
	TEE Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
	TEE World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)	
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)	
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Again, a few interesting aspects 😊

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Hardware and software are both popular research areas

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
		Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Protection / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² All combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Intel
MPK

CHERI

WebAssembly

A lot of interest in enforcing memory isolation at a **byte granularity** (vs. a page)

	Mechanism Class	Condi- tioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ 8-1024 domains. ⁵ Rubbing, stack switch.

Very widely varying – 128 bits!

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Nearly all works aim at cutting the **cost of domain switches**

- (1) Watson et al., CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization, S&P 2015
 (2) Amar et al., CHERIoT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

Pick one influential example from the literature: **CHERI** ⁽¹⁾⁽²⁾

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

- (1) Watson et al., CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization, S&P 2015
 (2) Amar et al., CHERIoT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

Pick one influential example from the literature: **CHERI** ⁽¹⁾⁽²⁾

Hardware mechanism that comes as an ISA extension

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	func [26]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)
	Monolithic Memory Protection (MMP) [249]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Segmentation-like Hardware [109, 178]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
	Enclaves [28, 92]	○	Mutual	TEE	●	●	○	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Confidential VMs [10, 29, 30]	○	Single	TEE	●	●	○	○	Page	2 (safe/unsafe)	○ (⁵)
	World Separation [9, 14]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
	Bounds-Checking Hardware [47, 98, 148, 155, 212]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	○	Mutual	Full	●	●	●	●	Byte	∞	● (special instr. + ⁵)
Software	Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)
	Software Capabilities [83, 125]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Mutual	Full	●	●	○	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Memory Encryption / AES-NI [155]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)
Memory Encryption / AES-NI [155]	○	Single	Full	●	●	○	○	128 bits	∞	○ (⁵)	
Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

- (1) Watson et al., CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization, S&P 2015
 (2) Amar et al., CHERIoT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

Pick one influential example from the literature: **CHERI** ⁽¹⁾⁽²⁾

Hardware mechanism that comes as an ISA extension

CHERI extends pointers with bounds and permission information. It is very expressive!

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	func [26]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)
	Mondrian Memory Protection (MMP) [249]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Segmentation-like Hardware [109, 178]	○	M								○ (MMP hardware + ⁵)
	Enclaves [28, 92]	○	M								○ (special register flip + ⁵)
	Confidential VMs [10, 29, 30]	○	M								○ (⁵)
	World Separation [9, 14]	○	Single	TEE	○	○	○	○	Page	2 (trusted/rest)	● (enclave call, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	● (> EPT switch)
	Bounds-Checking Hardware [47, 98, 148, 155, 212]	○	Mutual	Full	●	●	○	○	Byte	∞	● (world switch, incl. ⁵)
	(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (special instr. + ⁵)
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (bounds hardware + ⁵)
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	○	○	Byte	∞	○ (tagging hardware + ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ – ● (impl. dep., incl. ⁵)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (function call)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	○ (⁵)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Pick one influential example from the literature: **CHERI** ⁽¹⁾⁽²⁾

Hardware mechanism that comes as an ISA extension

CHERI extends pointers with bounds and permission information. It is very expressive!

It can be used to enforce isolation at a *byte granularity*, as opposed to entire pages

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	func [26]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)
	Mondrian Memory Protection (MMP) [249]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Segmentation-like Hardware [109, 178]	○	M								○ (MMP hardware + ⁵)
	Enclaves [28, 92]	○	M								○ (special register flip + ⁵)
	Confidential VMs [10, 29, 30]	○	M								○ (⁵)
	World Separation [9, 14]	○	Single	TEE	○	○	○	○	Page	2 (trusted/rest)	● (enclave call, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	● (> EPT switch)
	Bounds-Checking Hardware [47, 98, 148, 155, 212]	○	Mutual	Full	●	●	○	○	Byte	∞	● (world switch, incl. ⁵)
	(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	M							∞ [138]	○ (special instr. + ⁵)
Software	Software Capabilities [83, 125]	○	M							∞	○ – ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	M							∞	○ – ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Si							life/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Si							∞	○ (⁵)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

- (1) Watson et al., *CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization*, S&P 2015
- (2) Amar et al., *CHERIoT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices*, SOSP 2025

Pick one influential example from the literature: **CHERI**⁽¹⁾⁽²⁾

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	Nº of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	func [26]	○	Mutual	Full	●	●	●	○	Physical Mem.	Nº of machines	○ – ● (link latency)
	Mondrian Memory Protection (MMP) [249]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Segmentation-like Hardware [109, 178]	○	M								○ (MMP hardware + ⁵)
	Enclaves [28, 92]	○	M								○ (workload + ⁵)
	TEE Confidential VMs [10, 29, 30]	○	M								○ (workload + ⁵)
	World Separation [9, 14]	○	Single	TEE					Page	2 (trusted/rest)	○ (workload + ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	● (special instr. + ⁵)
	Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	●	●	○	○	Byte	∞	○ (bounds hardware + ⁵)
	(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	M							∞ [138]	○ (tagging hardware + ⁵)
Software	Software Capabilities [83, 125]	○	M							∞	○ – ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	M							∞	○ – ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Si							life/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Si							∞	○ (⁵)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)

Hardware mechanism that comes as an ISA extension

CHERI extends pointers with bounds and permission information. It is very expressive!

It comes with a very low domain-switching overhead.

It can be used to enforce isolation at a *byte granularity*, as opposed to entire pages

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N° of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N° of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	●	●	●	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	●	●	●	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	●	●	●	○	Word	∞	● (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
		Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (⁵)
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	●	●	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	●	●	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (⁵)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

We call them *conditioned*.

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)	
					R	W	X	A				
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)	
	PT	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
		Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
		Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
		Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (²)	
	TEE	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
		Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
		World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)	
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)		
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)		
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (²)	
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)	

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

We call them *conditioned*.

Intel
MPK

	Mechanism Class	Conditioned	Trust Model	TCB	Permissions				Granularity	N ^o of Domains	Domain Switch Cost (Versus Non-Separated)
					R	W	X	A			
Hardware	Physical Separation [205]	○	Mutual	Full	●	●	●	○	Physical Mem.	N ^o of machines	○ – ● (link latency)
	Access Bits [25], EPT / vmfunc [26]	● ¹	Mutual	Full	● ²	● ²	● ²	○	Page	∞	● (PT switch + ⁵)
	Supervisor Bit [25, 158, 159]	○	Single	Full	● ²	● ²	● ²	○	Page	2 (kernel/user)	● (interrupt + ⁵)
	Mondrian Memory Protection (MMP) [249]	○	Mutual	Full	● ²	● ²	● ²	○	Word	∞	○ (MMP hardware + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	○	Byte - Page [178]	2 (safe/unsafe)	○ (?)
	Enclaves [28, 92]	○	Mutual	TEE	●	●	●	○	Page	∞	● (enclave call, incl. ⁵)
	Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	●	○	Page	∞	● (> EPT switch)
	World Separation [9, 14]	○	Single	TEE	●	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	●	Byte	∞	○ (special instr. + ⁵)
	Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	● ²	● ²	○	○	Byte	∞	○ (bounds hardware + ⁵)
	(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	● ²	● ²	○	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	●	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	●	○	Byte	∞	○ – ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	●	Byte	2 (safe/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	●	Byte	∞	○ (?)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	○	128 bits	∞	● (copy key + encrypt + ⁵)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

We call them *conditioned*.

Intel MPK

Access permissions are stored in an unprotected register

								Granularity	Nº of Domains	Domain Switch Co. (Versus Non-Separated)
Hardware	Physical Separation							Physical Mem.	Nº of machines	○ - ● (link latency)
	Access Bits							Page	∞	● (PT switch + ⁵)
	Supervisor R							Page	2 (kernel/user)	● (interrupt + ⁵)
	Mondrian Memory Protec							Word	∞	○ (MMP hardware + ⁵)
	Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	●	●	○	Page	8-1024 [13, 217] ⁴	○ (special register flip + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	●	●	○	Byte - Page [178]	2 (safe/unsafe)	○ (?)
	Enclaves [28, 92]	○	Mutual	TEE	●	●	○	Page	∞	● (enclave call, incl. ⁵)
	Confidential VMs [10, 29, 30]	○	Mutual	TEE	●	●	○	Page	∞	● (> EPT switch)
	World Separation [9, 14]	○	Single	TEE	●	●	○	Page	2 (trusted/rest)	● (world switch, incl. ⁵)
	Hardware Capabilities [57, 78, 180, 236, 244]	○	Mutual	Full	●	●	●	Byte	∞	○ (special instr. + ⁵)
Bounds-Checking Hardware [47, 98, 148, 155, 212]	●	Mutual	Full	●	●	○	Byte	∞	○ (bounds hardware + ⁵)	
(Other) Tagged Architectures [12, 21, 99, 131, 138, 204, 224, 246]	●	Mutual	Full	●	●	○	Byte - Words ³	16 ⁴ - ∞ [138]	○ (tagging hardware + ⁵)	
Software	Software Capabilities [83, 125]	○	Mutual	Full	●	●	●	Byte	∞	○ - ● (impl. dep., incl. ⁵)
	Bounds-Checking Software [225]	●	Mutual	Full	●	●	○	Byte	∞	○ - ● (impl. dep., incl. ⁵)
	Safe Languages [50] / Software Verification [154, 163]	○	Single	Full	●	●	●	Byte	2 (safe/unsafe)	○ (function call)
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	●	●	●	Byte	∞	○ (?)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	●	●	○	128 bits	∞	● (copy key + encrypt + ⁵)

¹ In Ring 0. ² Not all combinations of R/W/X supported. ³ Covers many granularities [138]. ⁴ Some works [113, 173, 190] increase it. ⁵ Register saving/scrubbing, stack switch.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

Intel MPK

		Granularity	Nº of Domains	Domain Switch Co. (Versus Non-Separated)	
Hardware	Physical Separation	Physical Mem.	Nº of machines	○ – ● (link latency)	
	Access Bits			● (impl. dep., incl. ⁵)	
	Supervisor R/W/X			● (impl. dep., incl. ⁵)	
	Mondrian Memory Protection			○ (special instr. + ⁵)	
	Protection Keys [13, 27, 53, 217, 259]	○	Mutual	Full	○ (tagging hardware + ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	○ (tagging hardware + ⁵)
	Enclaves [28, 92]	○	Mutual	TEE	○ (incl. ⁵)
	TEE Confidential VM				○ (world switch, incl. ⁵)
	World Separation				○ (special instr. + ⁵)
	Hardware Capabilities				○ (bounds hardware + ⁵)
Software	Bounds-Checking Hardware (Other) Tagged Address			○ (tagging hardware + ⁵)	
	Software Capabilities			○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software			○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50]			○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	○ (copy key + encrypt + ⁵)
Memory Encryption / AES-NI [155]	○	Mutual	Full	○ (copy key + encrypt + ⁵)	

Access permissions are stored in an unprotected register

Skews performance and security comparisons

Instructions that can modify this register must be protected, but this comes at a cost and is often neglected

We call them *conditioned*.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

Open Problem in Enforcement Mechanisms

Mechanisms must fulfill several properties to be suitable to enforce compartmentalization (we formalize them in the SoK)

Some (popular!) mechanisms **do not fulfill all these properties**

Intel MPK

		Granularity	Nº of Domains	Domain Switch Co. (Versus Non-Separated)	
Hardware	Physical Separation	Physical Mem.	Nº of machines	○ – ● (link latency)	
	Access Bits			● (impl. dep. + ⁵)	
	Supervisor R/W/X			● (impl. dep. + ⁵)	
	Mondrian Memory Protection			○ (register flip + ⁵)	
	Protection Keys [13, 27, 53, 217, 259]	●	Mutual	Full	○ (incl. ⁵)
	Segmentation-like Hardware [109, 178]	○	Single	Full	○ (incl. ⁵)
	Enclaves [28, 92]	○	Mutual	TEE	○ (incl. ⁵)
	TEE Confidential VM				○ (incl. ⁵)
	World Separation				○ (incl. ⁵)
	Hardware Capabilities				○ (incl. ⁵)
Software	Bounds-Checking Hardware (Other) Tagged Address			○ (incl. ⁵)	
	Software Capabilities			○ – ● (impl. dep., incl. ⁵)	
	Bounds-Checking Software			○ – ● (impl. dep., incl. ⁵)	
	Safe Languages [50]			○ (function call)	
	Software Fault Isolation [79, 119, 141, 142, 179, 238, 257, 262]	○	Single	Full	○ (incl. ⁵)
	Memory Encryption / AES-NI [155]	○	Mutual	Full	○ (incl. ⁵)

Access permissions are stored in an unprotected register

Skews performance and security comparisons

Instructions that can modify this register must be protected, but this comes at a cost and is often neglected

We call them *conditioned*.

How to fairly evaluate the security and performance characteristics of isolation mechanisms?

My promise for this talk

*A journey through twenty years of
compartmentalization*

- ~~1. What is software compartmentalization? (slightly more formal)~~
- ~~2. A systematic perspective on compartmentalization~~
3. The *why*: compartmentalization everywhere, what will it take?



What are the remaining obstacles? Where to go from there?

Potential obstacle #1: **Lack of awareness?**

Potential obstacle #1: Lack of awareness?

"Out of the 13 groups of compartmentalized applications we constituted, 8 are exclusively authored by academics or security professionals"



(1)

(1) Lefeuve et al., SoK: Software Compartmentalization, S&P 2025

Potential obstacle #1: Lack of awareness?

"Out of the 13 groups of compartmentalized applications we constituted, 8 are exclusively authored by academics or security professionals"

We barely teach it at the university, not part of common "software design patterns"



(1)

(1) Lefeuve et al., SoK: Software Compartmentalization, S&P 2025

Potential obstacle #1: Lack of awareness?

"Out of the 13 groups of compartmentalized applications we constituted, 8 are exclusively authored by academics or security professionals"

We barely teach it at the university, not part of common "software design patterns"

Yes, but...

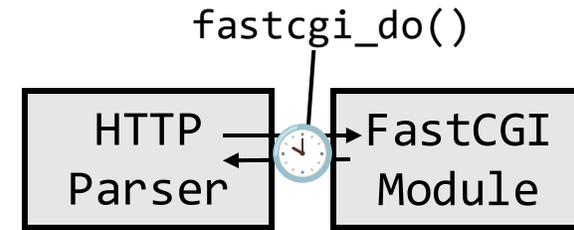


(1)

(1) Lefeuvre et al., SoK: Software Compartmentalization, S&P 2025

Potential obstacle #2: **Performance?**

Potential obstacle #2: Performance?



Compartmentalization has an **impact on performance**

- Crossing protection domains or IPC is not free (among others)
- This cost varies based on *where* crossings are in the control flow

This makes compartmentalization more complex and limits the extent to which one can compartmentalize

Critical stance: **in most cases, this is not a problem.**

- The community is overly focused on performance
- The cost is reasonable in most cases if the compartmentalization is done correctly
- In fact, many performance-sensitive programs are compartmentalized!

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

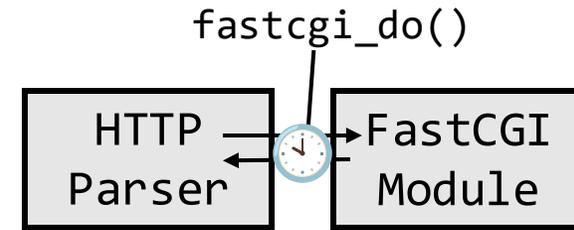
Rice's theorem: "non-trivial semantic properties of programs are undecidable".

Which trade-offs are desirable in practice?

A small diagram within the slide showing a 'Fictive monolithic program' box containing 'HTTP Parser', 'SS Library', and 'FastCGI Module'. A red arrow points from the 'FastCGI Module' to the 'Network' cloud below it.

Remember our policy definition challenge...

Potential obstacle #2: Performance?



Compartmentalization has an **impact on performance**

- Crossing protection domains or IPC is not free (among others)
- This cost varies based on *where* crossings are in the control flow

This makes compartmentalization more complex and limits the extent to which one can compartmentalize

Critical stance: **in most cases, this is not a problem.**

- The community is overly focused on performance
- The cost is reasonable in most cases if the compartmentalization is done correctly
- In fact, many performance-sensitive programs are compartmentalized!

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

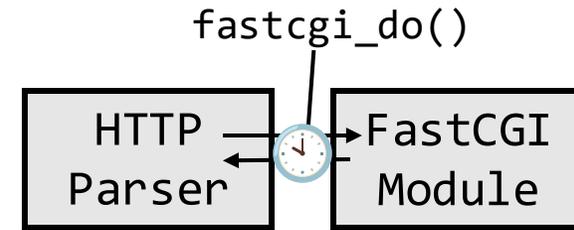
Rice's theorem: "non-trivial semantic properties of programs are undecidable".

Which trade-offs are desirable in practice?

A small diagram within the text block showing a box labeled 'HTTP Parser' on the left and a box labeled 'FastCGI Module' on the right. A red arrow points from the FastCGI Module to the HTTP Parser. Below these boxes is a cloud labeled 'Network'. A red arrow points from the Network to the FastCGI Module. The diagram is labeled '(Fictive monolithic program)'.

Remember our policy definition challenge...

Potential obstacle #2: Performance?



Compartmentalization has an **impact on performance**

- Crossing protection domains or IPC is not free (among others)
- This cost varies based on *where* crossings are in the control flow

This makes compartmentalization more complex and limits the extent to which one can compartmentalize

Critical stance: **in most cases, this is not a problem.**

- The community is overly focused on performance
- The cost is reasonable in most cases if the compartmentalization is done correctly
- In fact, many performance-sensitive programs are compartmentalized!

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

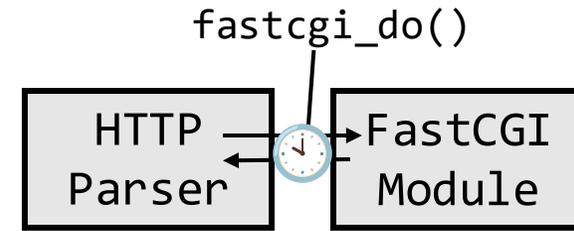
Rice's theorem: "non-trivial semantic properties of programs are undecidable".

Which trade-offs are desirable in practice?

The small diagram within the box shows a 'Fictive monolithic program' represented as a large rectangle. A red arrow points from the 'Network' (bottom) to the 'HTTP Parser' (left) and then to the 'FastCGI Module' (right). A red arrow also points from the 'FastCGI Module' back to the 'Network'. A red arrow points from the 'FastCGI Module' to the 'HTTP Parser'.

Remember our policy definition challenge...

Potential obstacle #2: Performance?



Compartmentalization has an **impact on performance**

- Crossing protection domains or IPC is not free (among others)
- This cost varies based on *where* crossings are in the control flow

This makes compartmentalization more complex and limits the extent to which one can compartmentalize

Critical stance: **in most cases, this is not a problem.**

- The community is overly focused on performance
- The cost is reasonable in most cases if the compartmentalization is done correctly
- In fact, many performance-sensitive programs are compartmentalized!

Open Problem in Policy Definition Methods

Automated Policy Definition Methods **trade off security and/or performance for developer effort.**

Take our example: *our split cuts a hot path.* We can spend more developer time to look attentively and refine our cut to avoid the critical code paths.

Automated methods struggle to do that since the semantics of programs fundamentally **cannot be captured automatically.**

Rice's theorem: "non-trivial semantic properties of programs are undecidable".

Which trade-offs are desirable in practice?

The small diagram within the text box shows a box labeled 'HTTP Parser' on the left and a box labeled 'FastCGI Module' on the right. A red arrow points from the FastCGI Module to the HTTP Parser. Below these boxes is a cloud labeled 'Network' with a red arrow pointing up towards the FastCGI Module. A red arrow also points from the FastCGI Module to the Network. The text '(Fictive monolithic program)' is written below the diagram.

Remember our policy definition challenge...

Yes, but...

Potential obstacle #3: **Complexity!**

Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Going back to our example:

Determining boundaries

Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*



This is still the most common way to do it today.

Implementing them with low-level process APIs

Problem #2

How to implement policies?

Having defined a policy, we need to **express it in the program**. Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



This is still the most common way to do it today.

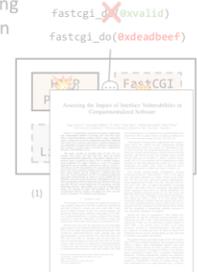
Securing them with ad-hoc audits

Open Problem in Compartmentalization

Abstraction is key to obtaining strong security in compartmentalization. Example: a developer implemented our boundaries to be exposed to communication on a back surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem. We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?



Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

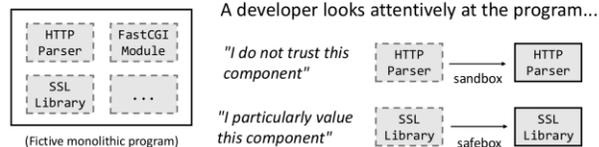
Going back to our example:

Determining boundaries

Policy Definition Methods

Historically, people have done this **manually**.

Example: we want to split this C program.



This is still the most common way to do it today.

85

Implementing them with low-level process APIs

Problem #2

How to implement policies?

Having defined a policy, we need to **express it in the program**. Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the process abstraction).



This is still the most common way to do it today.

117

Securing them with ad-hoc audits

Open Problem in Compartmentalization

Abstracting the boundaries between components is key to obtaining strong security in compartmentalized software.

Example: a developer implemented our boundaries to be exposed to communication over a network surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.

We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?



Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

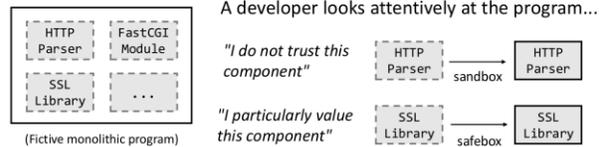
Going back to our example:

Determining boundaries

Policy Definition Methods

Historically, people have done this **manually**.

Example: *we want to split this C program.*



This is still the most common way to do it today.

85

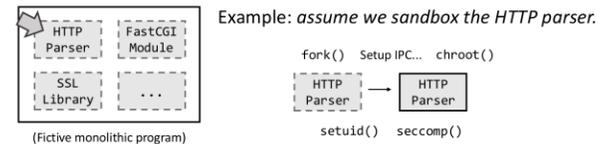
Implementing them with low-level process APIs

Problem #2

How to implement policies?

Having defined a policy, we need to **express it in the program**. Developers perform this using **programming abstractions**.

Historically, people have done this with **processes** (=the *process abstraction*).



This is still the most common way to do it today.

117

Securing them with ad-hoc audits

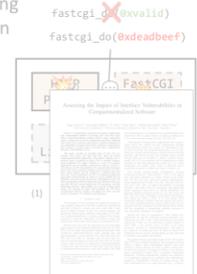
Open Problem in Compartmentalization

Abstracting boundaries is key to obtaining strong security in compartmentalization.

Example: a developer implemented our boundaries, but failed to communicate with the back surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem. We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?



Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Going back to our example:

Determining boundaries

Policy Definition Methods

Historically, people have done this **manually**.
Example: *we want to split this C program.*

A developer looks attentively at the program...

"I do not trust this component"

"I particularly value this component"

(Fictive monolithic program)

This is still the most common way to do it today.

Implementing them with low-level process APIs

Problem #2
How to implement policies?

Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.
Historically, people have done this with **processes** (=the *process abstraction*).

Example: *assume we sandbox the HTTP parser.*

This is still the most common way to do it today.

Securing them with ad-hoc audits

Open Problem in Compartmentalization

Abstraction is key to obtaining strong security in compartmentalization

Example: a developer implemented our boundaries to be exposed to communication back surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.
We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?

This approach is ad-hoc, costly in developer time, and requires a lot of expertise.

Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Going back to our example:

Determining boundaries

Policy Definition Methods

Historically, people have done this **manually**.
Example: *we want to split this C program.*

A developer looks attentively at the program...

"I do not trust this component"

"I particularly value this component"

(Fictive monolithic program)

This is still the most common way to do it today.

Implementing them with low-level process APIs

Problem #2
How to implement policies?

Having defined a policy, we need to **express it in the program**.
Developers perform this using **programming abstractions**.
Historically, people have done this with **processes** (=the process abstraction).

Example: *assume we sandbox the HTTP parser.*

This is still the most common way to do it today.

Securing them with ad-hoc audits

Open Problem in Compartmentalization

Abstraction is key to obtaining strong security in compartmentalization

Example: a developer implemented our boundaries to be exposed to communication back surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem.
We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?

This approach is ad-hoc, costly in developer time, and requires a lot of expertise.

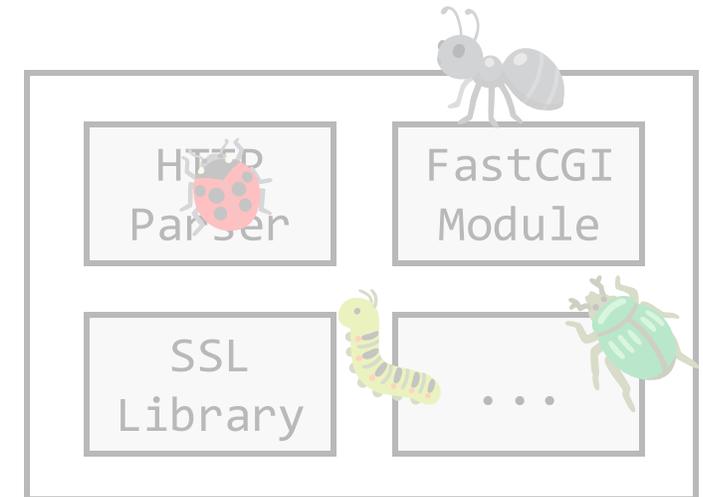
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider debugging:

- *There will be bugs to fix.*
- Bugs now occur across security boundaries.
- This impacts the debugging experience but the **debugging ecosystem is not there**.



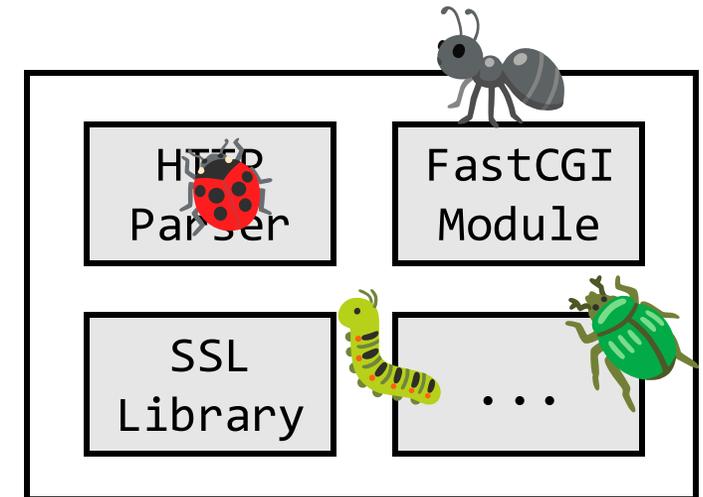
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider debugging:

- *There will be bugs to fix.*
- Bugs now occur across security boundaries.
- This impacts the debugging experience but the **debugging ecosystem is not there**.



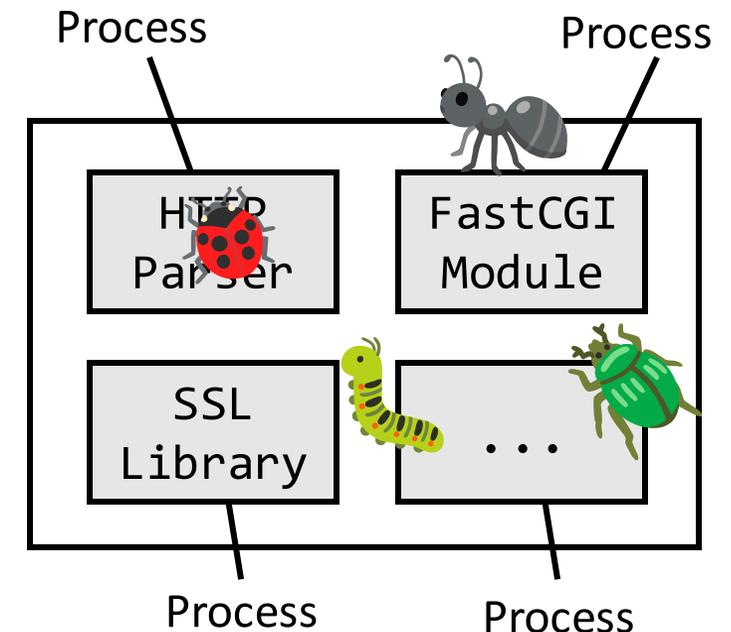
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider debugging:

- *There will be bugs to fix.*
- Bugs now occur across security boundaries.
- This impacts the debugging experience but the **debugging ecosystem is not there**.



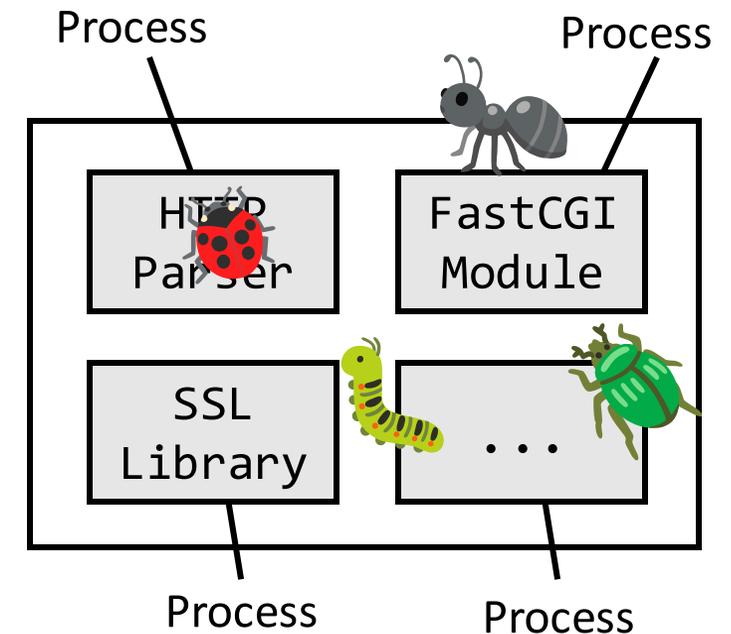
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider debugging:

- *There will be bugs to fix.*
- Bugs now occur across security boundaries.
- This impacts the debugging experience but the **debugging ecosystem is not there**.



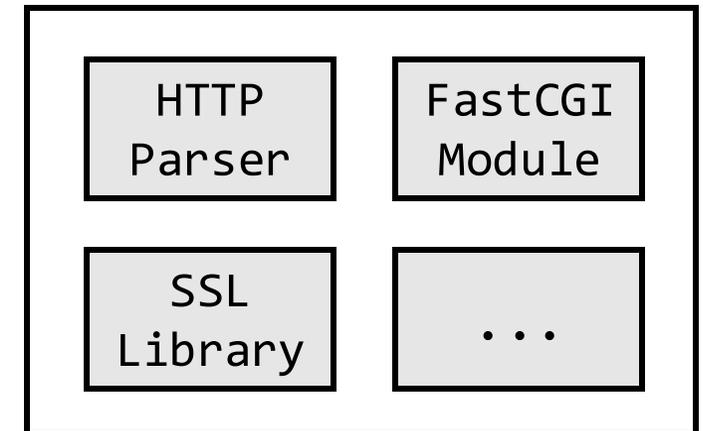
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider maintenance!

- Developers will need to maintain the partitioned program.
- Over time, the **partitioning strategy may need to change**.
- Again, there is historically **no tooling** to handle that safely and avoid regressions.



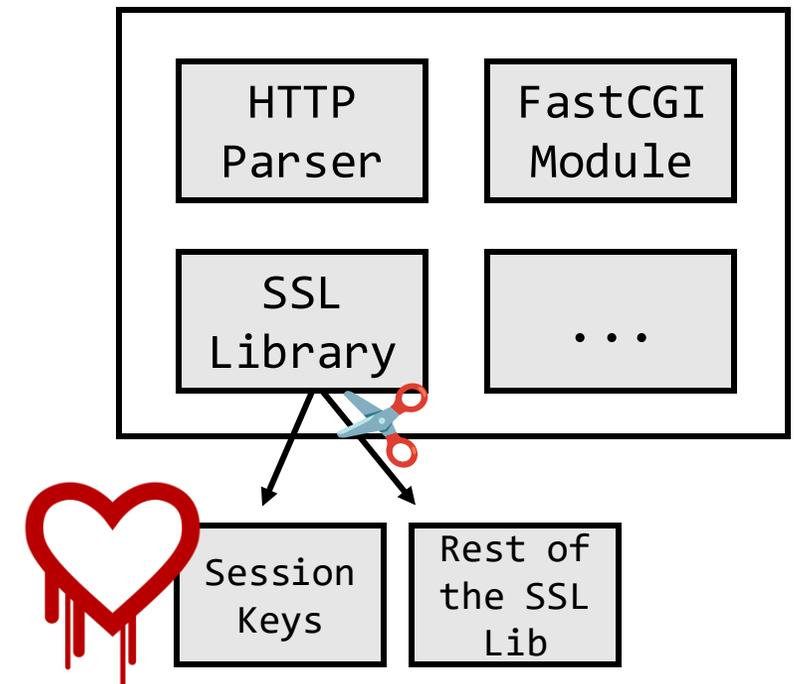
Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software is **too complex**.

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider maintenance!

- Developers will need to maintain the partitioned program.
- Over time, the **partitioning strategy may need to change**.
- Again, there is historically **no tooling** to handle that safely and avoid regressions.



Potential obstacle #3: Complexity!

The historical approach to compartmentalizing software

Worse yet, it considers compartmentalization as a one-off operation, **which it isn't**.

Consider maintenance!

- Developers will need to maintain the partitioned program.
- Over time, the **partitioning strategy may need to change**.
- Again, there is historically **no tooling** to handle that safely and avoid regressions.

(1) Lefeuve et al., Assessing the Impact of Interface Vulnerabilities in Compartmentalized Software, NDSS'23

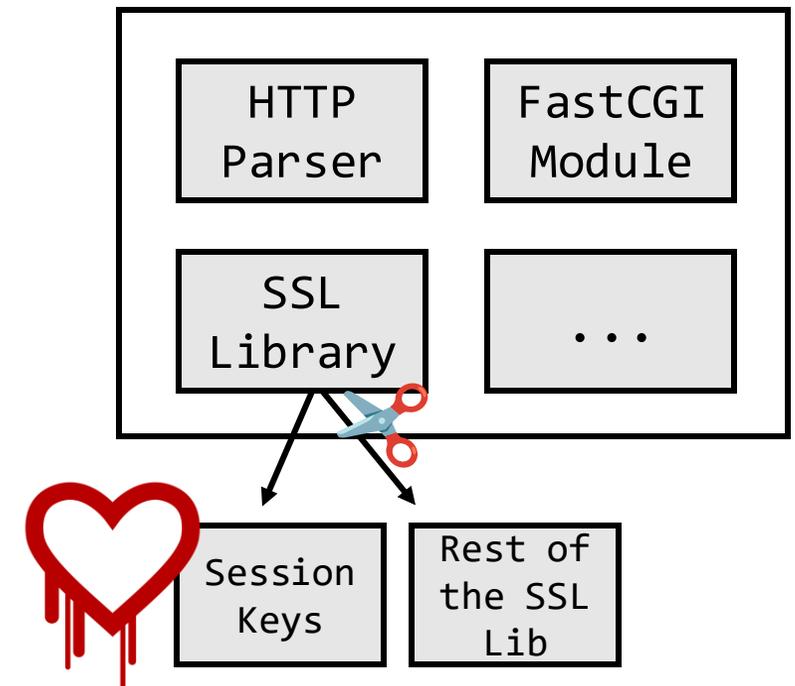
Open Problem in Compartmentalization

Abstraction is key to obtaining strong security in compartmentalization. Example: a boundary implemented our communication to the network surface.

Most abstractions consider the **hardening of compartment interfaces** an orthogonal problem. We stress that abstractions **can and should** facilitate the hardening of compartment interfaces.

How can abstractions that facilitate the implementation of secure domain interfaces?

```
fastcgi_do(0xvalid)
fastcgi_do(0xdeadbeef)
```



Potential obstacle #4: **The lack of a systematic approach?**

Potential obstacle #4: The lack of a systematic approach?



The historical approach to compartmentalization does not scale:

- Historically, compartmentalization is deeply **custom to each program**
- Compartmentalize program *A*, *then start all over again with B*

Can we make compartmentalization efforts reusable?

- Can we compartmentalize component *C*, and reuse that component compartmentalized everywhere?
- Analogy to the "software crisis" from the 1960s:
 - *The answer to "we are not producing enough software" is to build software from reusable components⁽¹⁾*

(1) <https://www.cs.dartmouth.edu/~doug/components.txt>

Potential obstacle #4: The lack of a systematic approach?



The historical approach to compartmentalization does not scale:

- Historically, compartmentalization is deeply **custom to each program**
- Compartmentalize program *A*, *then start all over again with B*

Can we make compartmentalization efforts reusable?

- Can we compartmentalize component *C*, and reuse that component compartmentalized everywhere?
- Analogy to the "software crisis" from the 1960s:
 - *The answer to "we are not producing enough software" is to build software from reusable components⁽¹⁾*

(1) <https://www.cs.dartmouth.edu/~doug/components.txt>

Potential obstacle #4: The lack of a systematic approach?



The historical approach to compartmentalization does not scale:

- Historically, compartmentalization is deeply **custom to each program**
- Compartmentalize program *A*, *then start all over again with B*

Can we make compartmentalization efforts reusable?

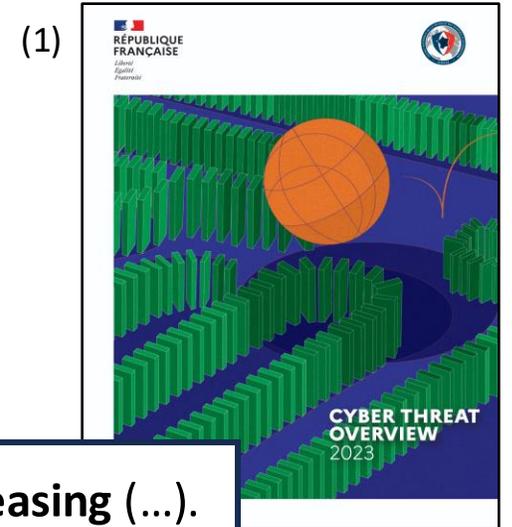
- Can we compartmentalize component *C*, and reuse that component compartmentalized everywhere?
- Analogy to the "software crisis" from the 1960s:
 - *The answer to "we are not producing enough software" is to build software from reusable components*⁽¹⁾

(1) <https://www.cs.dartmouth.edu/~doug/components.txt>

A golden age

It is the right *time* to push compartmentalization everywhere

- More than ever **we need more secure software**

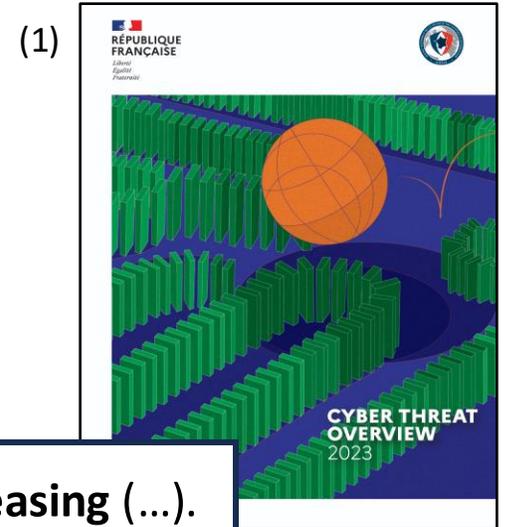


The level of **cyber threat keeps on increasing (...)**. Today, ANSSI assesses that attackers linked to China, Russia and cybercrime pose the **greatest threat to the most critical networks and to the French ecosystem in a systemic way.**

*emphasis mine

It is the right *time* to push compartmentalization everywhere

- More than ever **we need more secure software**
- Compartmentalization **absolutely fits the needs**



Secure Hardware Foundation: Incorporate architectural features that **enable fine-grained memory protection**, such as those described by Capability Hardware Enhanced RISC Instructions (CHERI)

The level of **cyber threat keeps on increasing (...)**. Today, ANSSI assesses that attackers linked to **cybercrime pose the greatest threat to the most critical networks and to the economy in a systemic way.**

*emphasis mine

Co-signed by, among others, UK, NL, DE, NO

(2) https://www.cisa.gov/sites/default/files/2023-10/SecureByDesign_1025_508c.pdf

(1) <https://www.cert.ssi.gouv.fr/uploads/CERTFR-2024-CTI-002.pdf>

It is the right *time* to push compartmentalization everywhere

- The challenges to push compartmentalization at scale are **not new**:
 - Research has **explored them in many ways** and come up with **solutions that work** (second part of this talk!)
 - These advances are just asking to be **developed and mainstreamed**



My promise for this talk

A journey through twenty years of compartmentalization

1. What is software compartmentalization? (slightly more formal)
2. A systematic perspective on compartmentalization
3. The why: compartmentalization everywhere, what will it take?

Take a step back: How do practitioners compartmentalize? What does research say? What are open challenges?

It is the right *time* to push compartmentalization everywhere

- The challenges to push compartmentalization at scale are **not new**:
 - Research has **explored them in many ways** and come up with **solutions that work** (second part of this talk!)
 - These advances are just asking to be **developed and mainstreamed**

(1)



My promise for this talk

A journey through twenty years of compartmentalization

1. What is software compartmentalization? (slightly more formal)
2. A systematic perspective on compartmentalization
3. The why: compartmentalization everywhere, what will it take?

Take a step back: How do practitioners compartmentalize? What does research say? What are open challenges?

It is the right *time* to push compartmentalization everywhere

- The challenges to push compartmentalization at scale are **not new**:
 - Research has **explored them in many ways** and come up with **solutions that work** (second part of this talk!)
 - These advances are just asking to be **developed and mainstreamed**

(1)



My promise for this talk

A journey through twenty years of compartmentalization

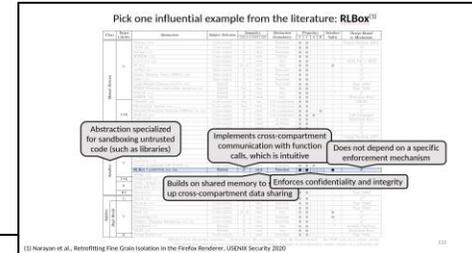
1. What is software compartmentalization? (slightly more formal)
2. A systematic perspective on compartmentalization
3. The why: compartmentalization everywhere, what will it take?

Take a step back: How do practitioners compartmentalize? What does research say? What are open challenges?

The missing frameworks are emerging



(1)



Take RLBox, discussed earlier in the talk

- Productized in Firefox for library isolation
- You can use it: <https://github.com/PLSysSec/rlbox>

Recently: increasing interest in coming together to create industry standards for compartmentalization

- The *Open Robust Compartmentalization Alliance* was just accepted as Linux Foundation project: <https://github.com/ORCA-LF/governance>
- Aiming to foster the adoption and standardization of compartmentalization practices

Retrofitted Fine Grain Isolation in the Firefox Renderer

Shravan Narayan¹ Craig Disselkosen² Tal Garfinkel² Nathan Froyd³ Eric Rahm⁴ Soren Lerner¹ Hovav Shacham^{2,3} Deian Stefan² UC San Diego¹ Stanford² Mozilla³ UT Austin⁴

Abstract

Firefox and other major browsers rely on dozens of third-party libraries to render audio, video, images, and other content. These libraries are a frequent source of vulnerabilities. To mitigate this threat, we are migrating Firefox to an architecture that isolates these libraries in lightweight sandboxes, dramatically reducing the impact of a compromise. Retrofitting isolation can be labor-intensive, very prone to security bugs, and requires critical attention to performance. To help, we developed RLBox, a framework that minimizes the burden of converting Firefox to securely and efficiently use untrusted code. To enable this, RLBox employs static information flow enforcement, and lightweight dynamic checks, exposed directly in the C++ type system. RLBox supports efficient sandboxing through either software-based fault isolation or multi-core process isolation. Performance overheads are modest and transient, and have only minor impact on page latency. We demonstrate this by sandboxing performance-sensitive image decoding libraries (libjpeg and libpng), video decoding libraries (libvorbis and libx264), the theora audio decoding library, and the cifs decompression library. RLBox, using a WebAssembly sandbox, has been integrated into production Firefox to sandbox the B2G-native font shaping library.

1 Introduction

All major browsers today employ coarse grain privilege separation to limit the impact of vulnerabilities. To wit, they run renderers—the portion of the browser that handles untrusted user content from HTML, parsing to JavaScript execution to image decoding and rendering—in separate sandboxed processes [1, 33, 40]. This stops web attackers that manage to compromise the renderer from abusing local OS resources to, say, install malware.

Unfortunately, this is no longer enough: nearly everything we care about today is done through a website. By compromising the renderer, an attacker gets total control of the current site and, often, any other sites the browser has credentials for [14]. With services like Dropbox and Google Drive, privilege separation is insufficient even to protect local files that sync with the cloud [24].

Browser vendors spend a huge amount of engineering effort trying to find renderer vulnerabilities in their own code [25]. Unfortunately, many remain—frequently in the

dozens of third-party libraries used by the renderer to decode audio, images, fonts, and other content. For example, an out-of-bounds write in libvorbis was used to exploit Firefox on PowerShell 2018 [7]. Both Chrome and Firefox were vulnerable to an integer-overflow bug in the theora video-decoding library [19]. Both also rely on the libx264 graphics library, which had four remote code execution bugs and severity [12, 17].

To appreciate the impact of these vulnerabilities and the difficulty of mitigating them, consider a typical web user, Alice, that uses Gmail to read email in her browser. Suppose an invisible, Trudy, sends Alice an email that contains a link to her malicious site, `http://evil.com/evil.com`. If Alice clicks on the link, her browser will navigate her to Trudy's site, which can embed an `` value to exploit vulnerabilities in libvorbis and libpng and compromise the renderer of Alice's browser. Trudy now has total control of Alice's Gmail account. Trudy can read and send emails as Alice. For example, to respond to password reset requests from other sites Alice belongs to. In most cases, Trudy can also attack cross site [14], i.e., she can access any other site that Alice is logged into (e.g., Alice's `facebook.com` account).

Recent versions of Chrome and upcoming versions of Firefox support Site Isolation [1], which isolates different sites from each other (e.g., `google.com` from `facebook.com`) to prevent such cross-site attacks. Unfortunately, Trudy might still be able to access `http://evil.com/evil.com`, which manage Alice's files, online payments, and cloud infrastructure—since the renderer that loads the malicious `` and `<video>` content might still be running in the same process as their origins.

For many sites, Trudy might not even need to upgrade malicious content to the (trusty) victim, `origin(1)://evil.com` in our example). Most web applications load content, including images, fonts, and videos, from different origins. Of the Alexa top 500 websites, for example, over 80% of the sites load at least one such cross-origin resource [8, 7]. And the libraries handling such content are not isolated from the embedding origin, even with Site Isolation [1].

To mitigate these vulnerabilities, we need to harden the renderer itself. To this end, we extend the Firefox renderer to isolate third party libraries in fine grain sandboxes. Using this, we can prevent compromised library from gaining control of the current origin or any other origin in the browser.

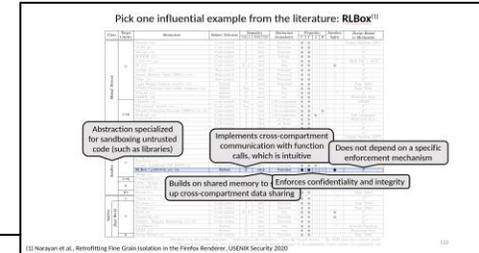
Making this practical poses three significant challenges:

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Render, USENIX Security 2020

The missing frameworks are emerging



(1)



Take RLBox, discussed earlier in the talk

- Productized in Firefox for library isolation
- You can use it: <https://github.com/PLSysSec/rlbox>

Recently: increasing interest in coming together to create industry standards for compartmentalization

- The *Open Robust Compartmentalization Alliance* was just accepted as Linux Foundation project: <https://github.com/ORCA-LF/governance>
- Aiming to foster the adoption and standardization of compartmentalization practices

Retrofitted Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

Shravan Narayan¹ Craig Disselkosen² Tal Garfinkel² Nathan Froyd³ Eric Rahm⁴ Sorn Lerner¹ Hovav Shacham^{2*} Deian Stefan² UC San Diego¹ Stanford² Mozilla³ UT Austin⁴

Abstract

Firefox and other major browsers rely on dozens of third-party libraries to render audio, video, images, and other content. These libraries are a frequent source of vulnerabilities. To mitigate this threat, we are migrating Firefox to an architecture that isolates these libraries in lightweight sandboxes, dramatically reducing the impact of a compromise. Retrofitting isolation can be labor-intensive, very prone to security bugs, and requires critical attention to performance. To help, we developed RLBox, a framework that automates the burden of converting Firefox to securely and efficiently use untrusted code. To enable this, RLBox employs static information flow enforcement, and lightweight dynamic checks, exposed directly in the C++ type system. RLBox supports efficient sandboxing through either software-based fault isolation or multi-core process isolation. Performance overheads are modest and transient, and have only minor impact on page latency. We demonstrate this by sandboxing performance-sensitive image decoding libraries (libjpeg and libpng), video decoding libraries (libvorbis and libx264), the theoretical audio decoding library, and the cifs decompression library. RLBox, using a WebAssembly sandbox, has been integrated into production Firefox to sandbox the B2G-optimized shipping library.

1 Introduction

All major browsers today employ coarse grain privilege separation to limit the impact of vulnerabilities. To wit, they run renderers—the portion of the browser that handles untrusted user content from HTML, parsing to JavaScript execution to image decoding and rendering—in separate sandboxed processes [1, 33, 40]. This stops web attackers that manage to compromise the renderer from abusing local OS resources to, say, install malware.

Unfortunately, this is no longer enough: nearly everything we care about today is done through a website. By compromising the renderer, an attacker gets total control of the current site and, often, any other sites the browser has credentials for [14]. With services like Dropbox and Google Drive, privilege separation is insufficient even to protect local files that sync with the cloud [24].

Browser vendors spend a huge amount of engineering effort trying to find renderer vulnerabilities in their own code [25]. Unfortunately, many remain—frequently in the

dozens of third-party libraries used by the renderer to decode audio, images, fonts, and other content. For example, an out-of-bounds write in libvorbis was used to exploit Firefox at Pwn2Own 2018 [7]. Both Chrome and Firefox were vulnerable to an integer overflow bug in the libpng video-decoding library [19]. Both also rely on the libx264 graphics library, which had four remote code execution bugs and severity [12, 17].

To appreciate the impact of these vulnerabilities, and the difficulty of mitigating them, consider a typical web site. Alice, that uses Gmail to read email in her browser. Suppose an invisible, Trudy, sends Alice an email that contains a link to her malicious site, hosted on `evil.com`. If Alice clicks on the link, her browser will navigate her to Trudy's site, which can embed an `img` or `video` tag to exploit vulnerabilities in libvorbis and libpng and compromise the renderer of Alice's browser. Trudy now has total control of Alice's Gmail account. Trudy can read and send emails as Alice. For example, to respond to password reset requests from other sites Alice belongs to. In most cases, Trudy can also attack cross site [14], i.e., she can access any other site that Alice is logged into (e.g., Alice's `amazon.com` account).

Recent versions of Chrome (and upcoming versions of Firefox) support Site Isolation [1], which isolates different sites from each other (e.g., `google.com` from `amazon.com`) to prevent such cross-site attacks. Unfortunately, Trudy might still be able to access `libvorbis.cloud.google.com`, which manage Alice's files, online payments, and cloud infrastructure—since the renderer that loads the malicious `img` and `video` content might still be running in the same process as their origins.

For many sites, Trudy might not even need to upgrade malicious content to the (trusty) victim origin (`google.com` in our example). Most web applications load content, including images, fonts, and videos, from different origins. Of the Alexa top 500 websites, for example, over 80% of the sites load at least one such cross-origin resource [8, 7]. And the libraries handling such content are not isolated from the embedding origin, even with Site Isolation [1].

To mitigate these vulnerabilities, we need to harden the renderer itself. To this end, we extend the Firefox renderer to isolate third party libraries in fine grain sandboxes. Using this, we can prevent a compromised library from gaining control of the current origin or any other origin in the browser.

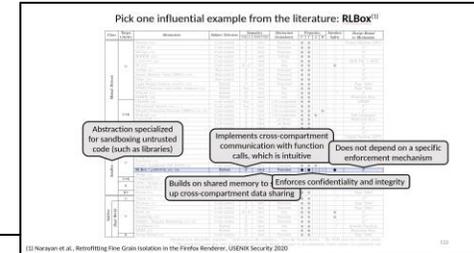
Making this practical poses three significant challenges:

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Render, USENIX Security 2020

The missing frameworks are emerging



(1)



Take RLBox, discussed earlier in the talk

- Productized in Firefox for library isolation
- You can use it: <https://github.com/PLSysSec/rlbox>

Recently: increasing interest in coming together to create industry standards for compartmentalization

- The *Open Robust Compartmentalization Alliance* was just accepted as Linux Foundation project: <https://github.com/ORCA-LF/governance>
- Aiming to foster the adoption and standardization of compartmentalization practices

Join us! Still in early stage, you can send me a message.

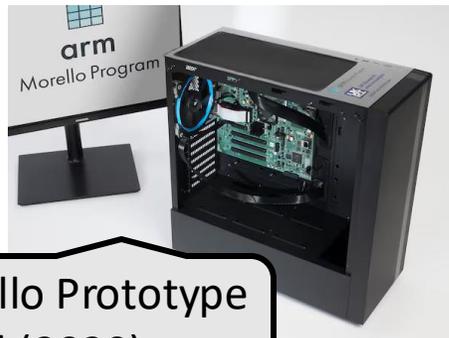
(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Renderer, USENIX Security 2020

(1) Narayan et al., Retrofitting Fine Grain Isolation in the Firefox Render, USENIX Security 2020

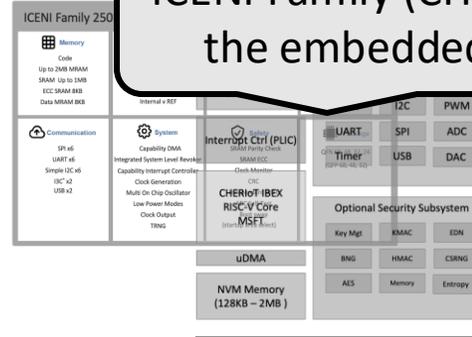
The hardware ecosystem is emerging

Take CHERI, discussed earlier as well.

It is not a research prototype. You can use it!



ARM Morello Prototype Board (2020)



ICENI Family (CHERIOT): CHERI for the embedded world (2025)

(1) Watson et al., CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization, S&P 2015
(2) Amar et al., CHERIOT RTOS: An OS for Fine-Grained Memory-Safe Compartments on Low-Cost Embedded Devices, SOSP 2025

Pick one influential example from the literature: CHERI⁽¹⁾⁽²⁾

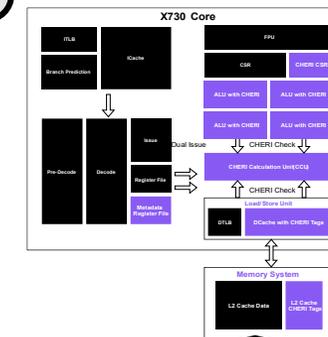
Mechanism Class	Config. Model	Time Model	TCR	Permissions	Granularity	# of Domains	Domain Switch Cost (Arms Non-Separated)
Hardware	Full	Full	Full	Full	Byte	10	100ns
Software	Full	Full	Full	Full	Page	10	100ns
Hardware	Full	Full	Full	Full	Page	10	100ns
Software	Full	Full	Full	Full	Page	10	100ns

Hardware mechanism that comes as an ISA extension

CHERI extends pointers with bounds and permission information. It is very expressive!

It comes with a very low domain-switching overhead.

It can be used to enforce isolation at a byte granularity, as opposed to entire pages.



Codasip X730 (2024)

Recently: creation of the *CHERI Alliance* to support industry standards and foster adoption

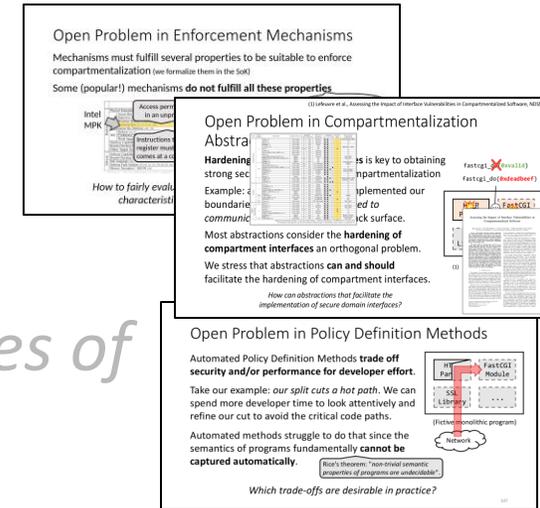
- *Join the CHERI Alliance!* <https://cheri-alliance.org/>

Software compartmentalization
everywhere: *what will it take?*

Yes, there are many interesting research challenges left.

But for the better part, it is about *transitioning two decades of research into the industry and open-source communities*

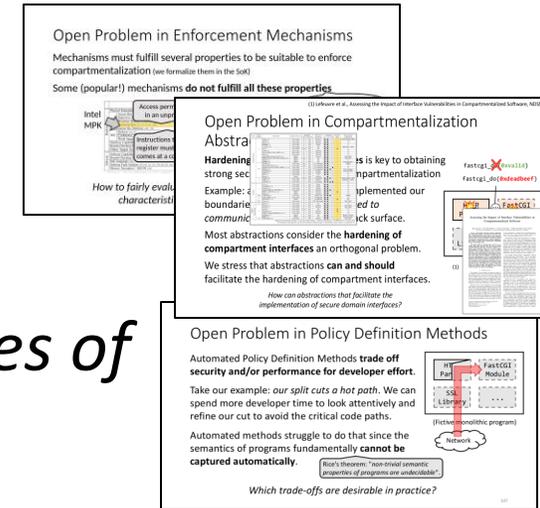
1. An effort on **industry-grade compartmentalization frameworks** to lower the bar to compartmentalize software
2. An effort on using these frameworks to **compartmentalize vulnerable components in a re-usable fashion**: can we do this under the umbrella of Linux distributions?
3. Supporting the push towards **emerging security-oriented hardware**: integrate CHERI in your hardware products!



Yes, there are many interesting research challenges left.

But for the better part, it is about *transitioning two decades of research into the industry and open-source communities*

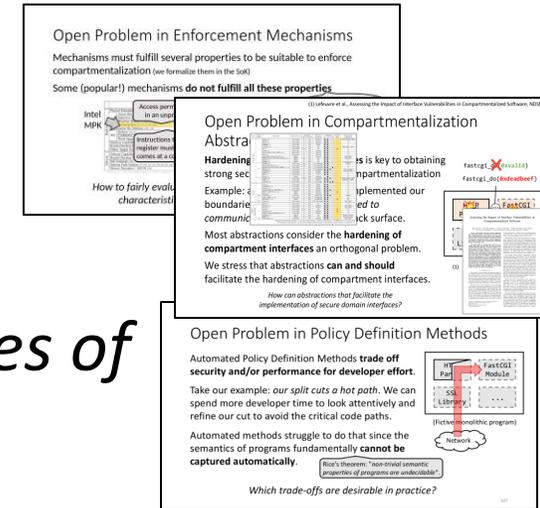
1. An effort on **industry-grade compartmentalization frameworks** to lower the bar to compartmentalize software
2. An effort on using these frameworks to **compartmentalize vulnerable components in a re-usable fashion**: can we do this under the umbrella of Linux distributions?
3. Supporting the push towards **emerging security-oriented hardware**: integrate CHERI in your hardware products!



Yes, there are many interesting research challenges left.

But for the better part, it is about *transitioning two decades of research into the industry and open-source communities*

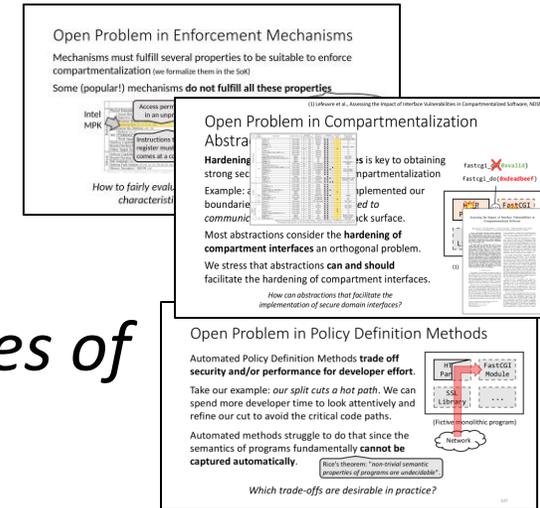
1. An effort on **industry-grade compartmentalization frameworks** to lower the bar to compartmentalize software
2. An effort on using these frameworks to **compartmentalize vulnerable components in a re-usable fashion**: can we do this under the umbrella of Linux distributions?
3. Supporting the push towards **emerging security-oriented hardware**: integrate CHERI in your hardware products!



Yes, there are many interesting research challenges left.

But for the better part, it is about *transitioning two decades of research into the industry and open-source communities*

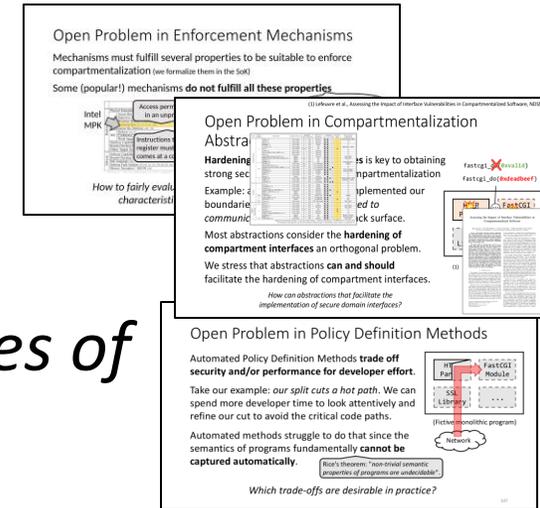
1. An effort on **industry-grade compartmentalization frameworks** to lower the bar to compartmentalize software
2. An effort on using these frameworks to **compartmentalize vulnerable components in a re-usable fashion**: can we do this under the umbrella of Linux distributions?
3. Supporting the push towards **emerging security-oriented hardware**: integrate CHERI in your hardware products!



Yes, there are many interesting research challenges left.

But for the better part, it is about *transitioning two decades of research into the industry and open-source communities*

1. An effort on **industry-grade compartmentalization frameworks** to lower the bar to compartmentalize software
2. An effort on using these frameworks to **compartmentalize vulnerable components in a re-usable fashion**: can we do this under the umbrella of Linux distributions?
3. Supporting the push towards **emerging security-oriented hardware**: integrate CHERI in your hardware products!





Key Takeaways

1. Compartmentalization is a key practice to **improve the trustworthiness of software**
2. Compartmentalization, in its historical form, is **hindered by its complexity**
3. We live at a **golden age to solve this problem**. A call to the community: we can make software more trustworthy!



Key Takeaways

1. Compartmentalization is a key practice to **improve the trustworthiness of software**
2. Compartmentalization, in its historical form, is **hindered by its complexity**
3. We live at a **golden age to solve this problem**. A call to the community: we can make software more trustworthy!



Key Takeaways

1. Compartmentalization is a key practice to **improve the trustworthiness of software**
2. Compartmentalization, in its historical form, is **hindered by its complexity**
3. We live at a **golden age to solve this problem**. A call to the community: we can make software more trustworthy!



Key Takeaways

1. Compartmentalization is a key practice to **improve the trustworthiness of software**
2. Compartmentalization, in its historical form, is **hindered by its complexity**
3. We live at a **golden age to solve this problem**. A call to the community: we can make software more trustworthy!



THE UNIVERSITY
OF BRITISH COLUMBIA

Sponsors:



Co-funded by
the European Union



UK Research
and Innovation

Collaborators:



Key Takeaways

1. Compartmentalization is a key practice to **improve the trustworthiness of software**
2. Compartmentalization, in its historical form, is **hindered by its complexity**
3. We live at a **golden age to solve this problem**. A call to the community: we can make software more trustworthy!

Questions?

Reach out: hugo.lefeuvre@ubc.ca