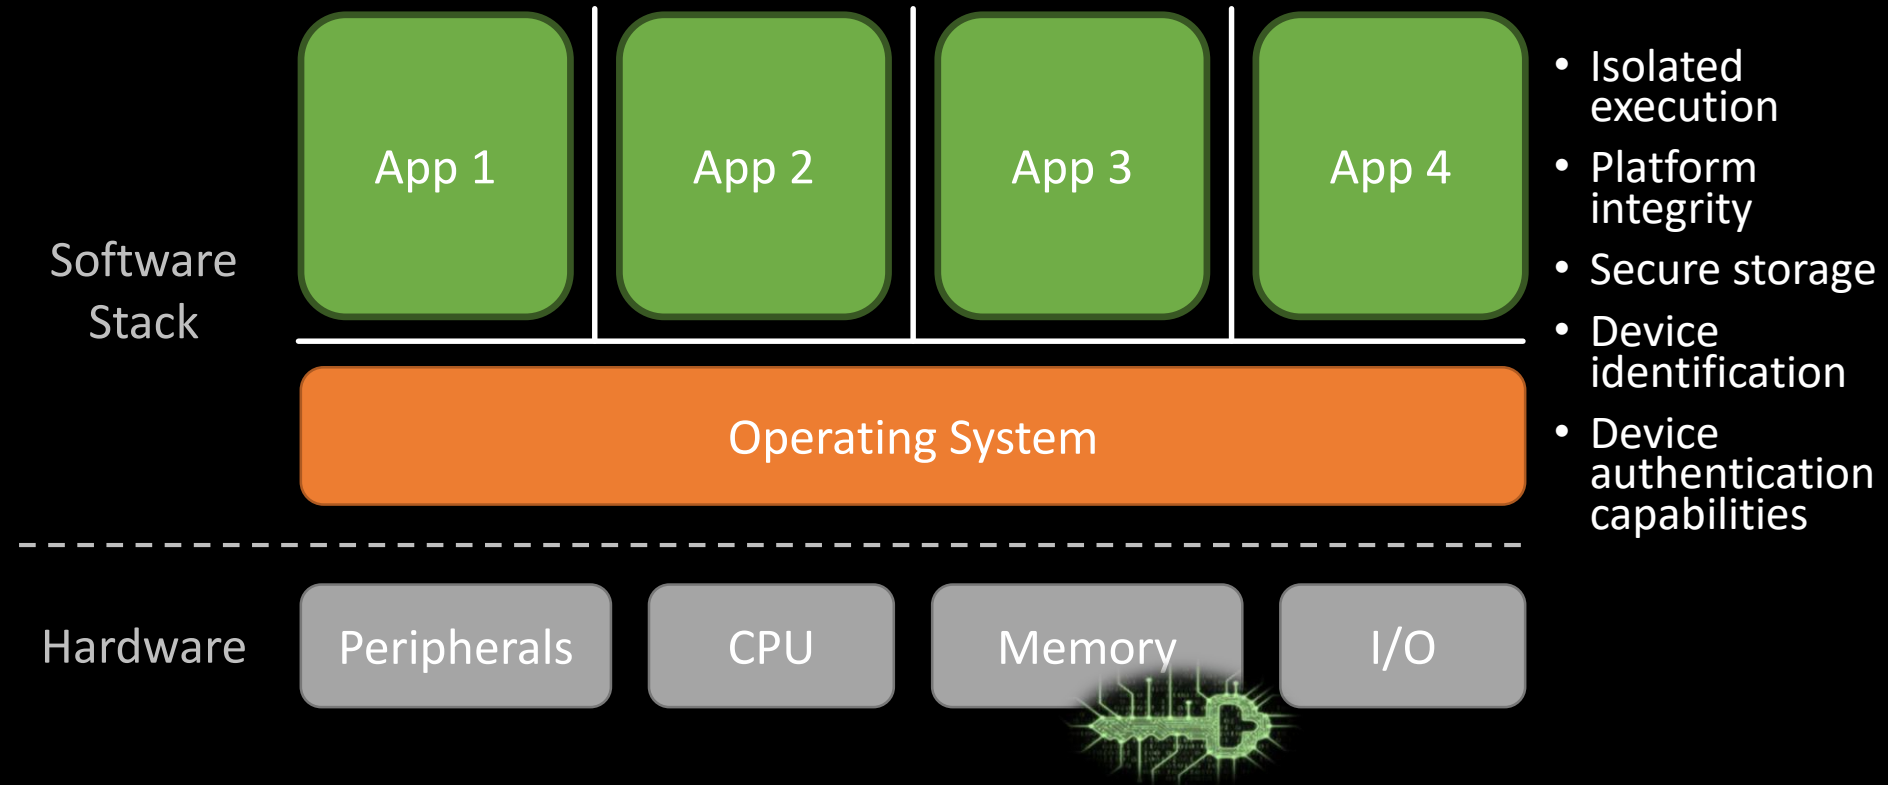


In Hardware We Trust: Enriching the World with Hardware Security

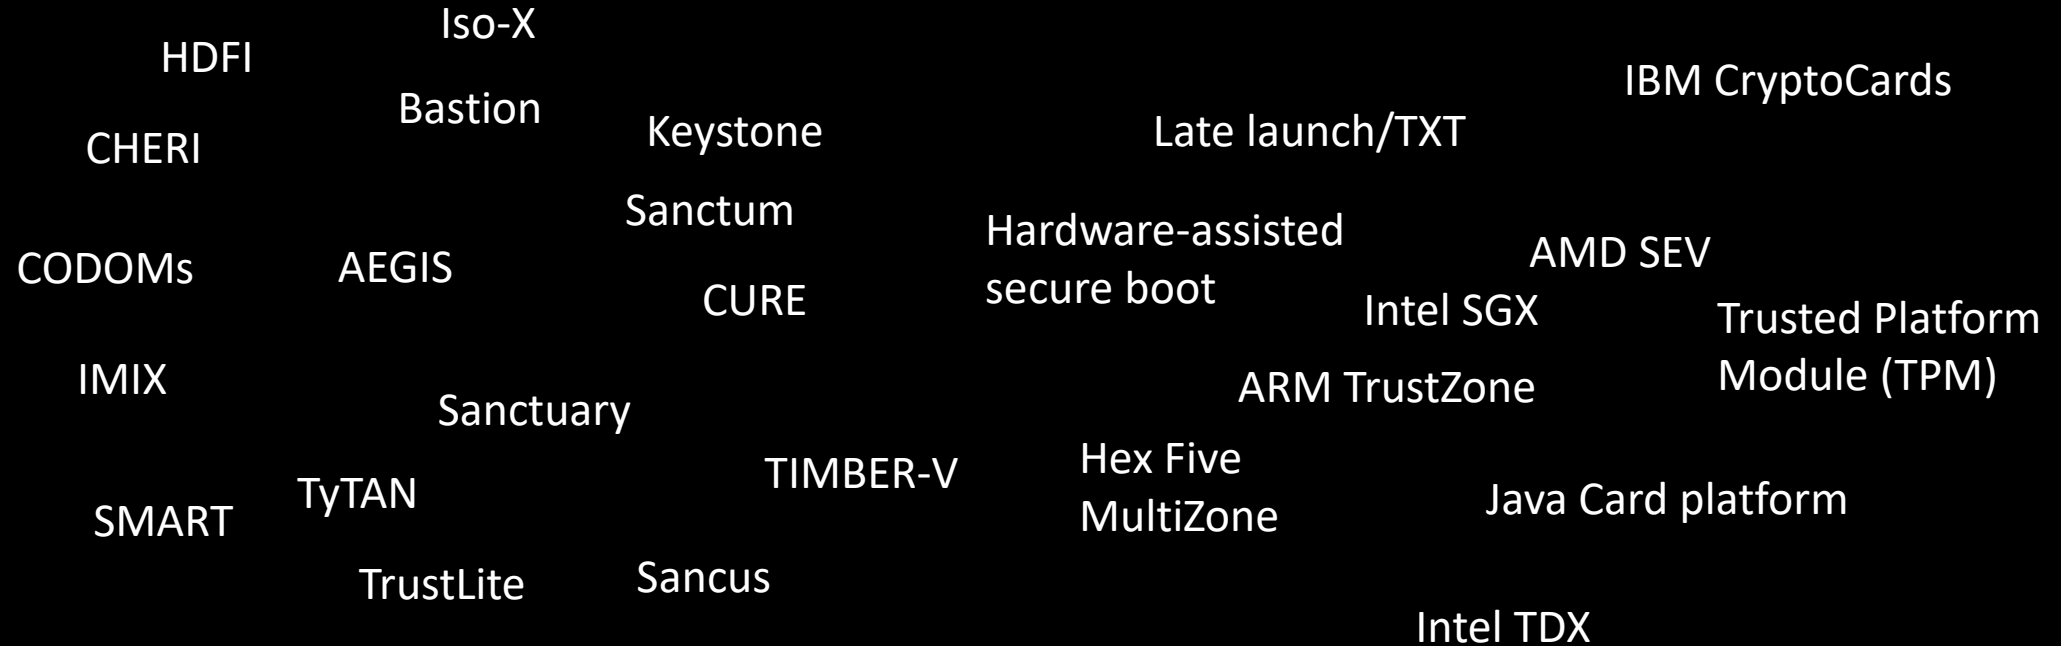
Ahmad-Reza Sadeghi
Technical University of Darmstadt



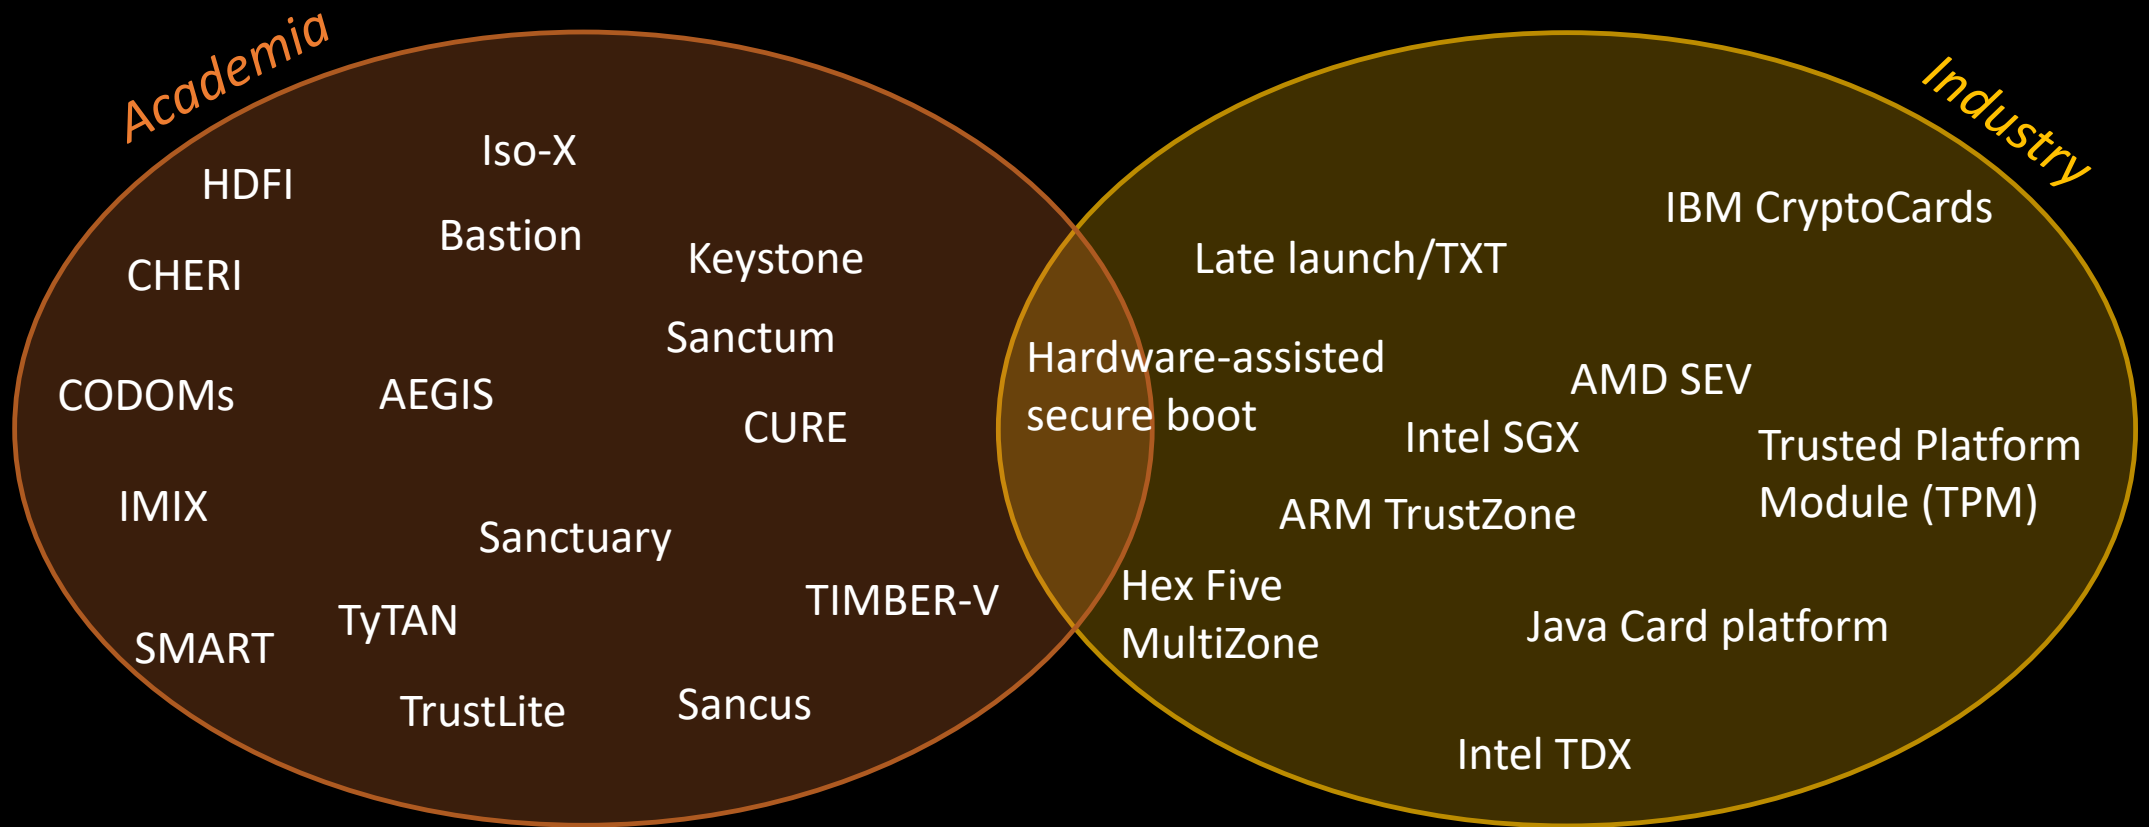
Trusted Computing Goal: Self-Contained Security



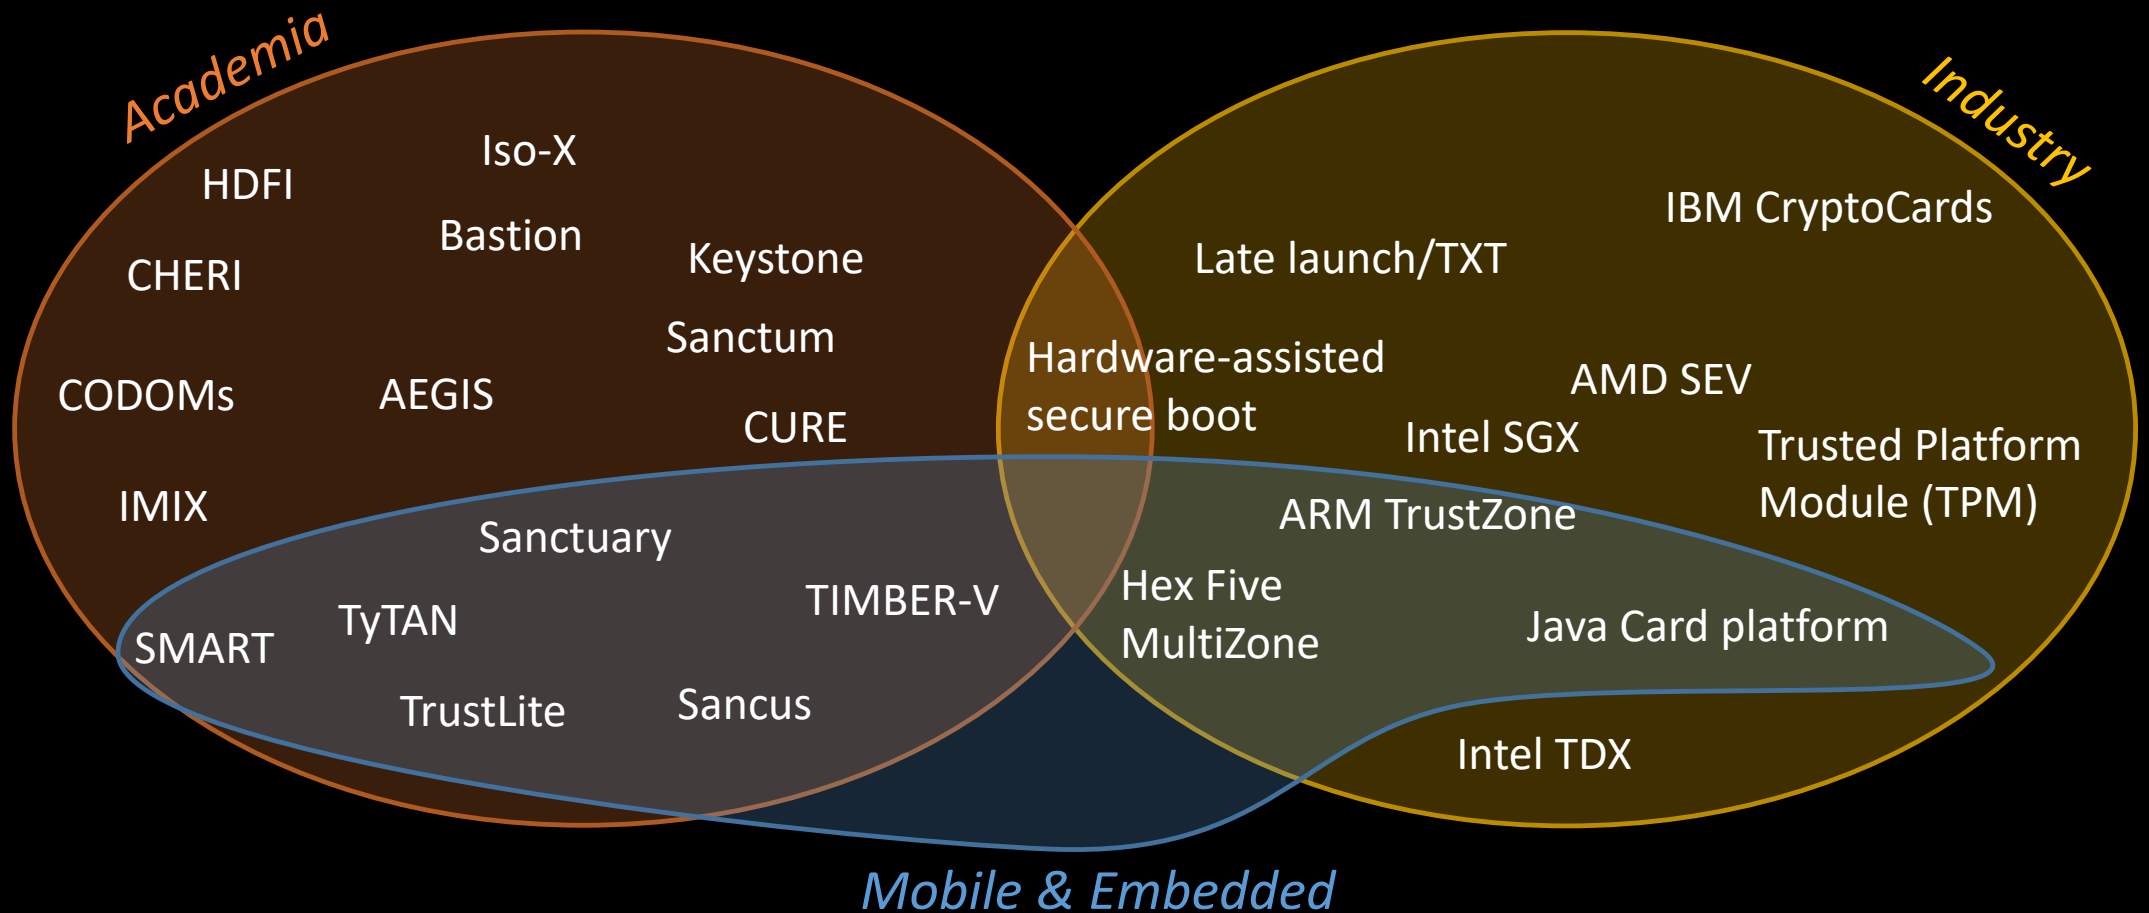
Trusted Computing Landscape



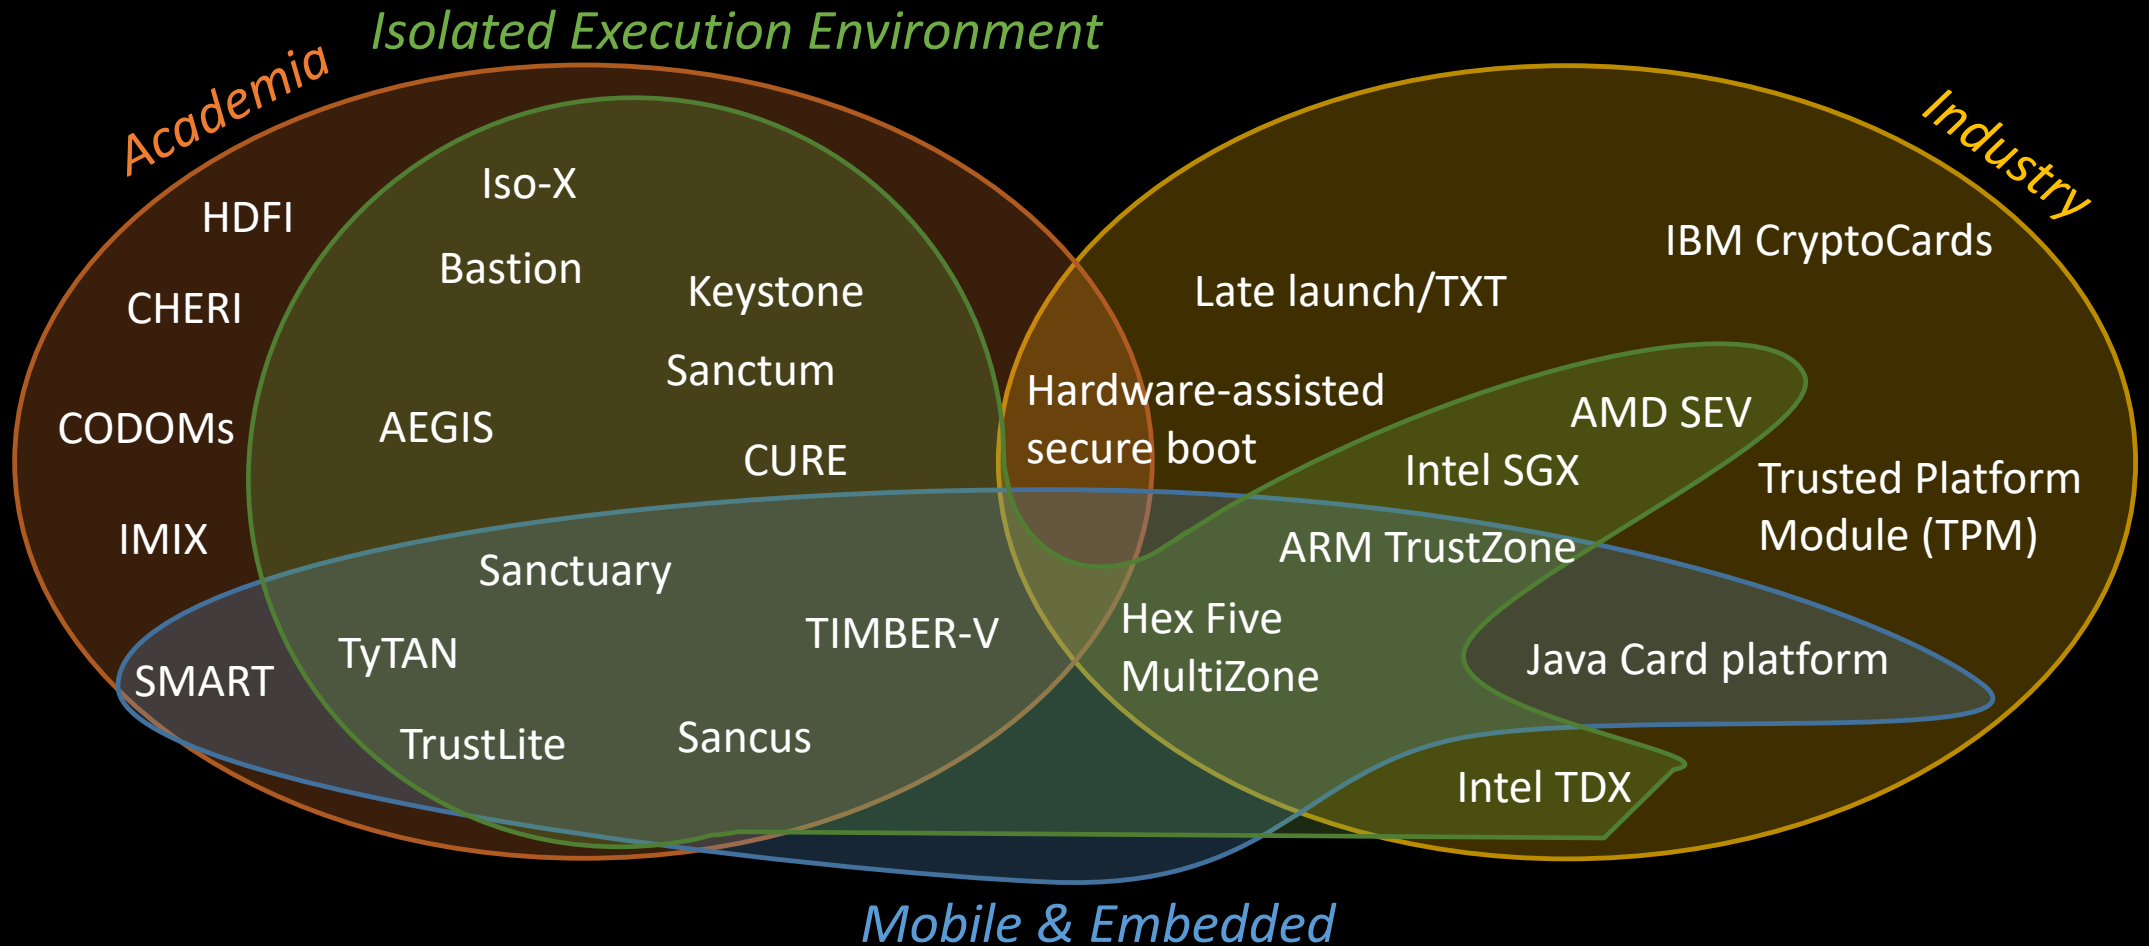
Trusted Computing Landscape



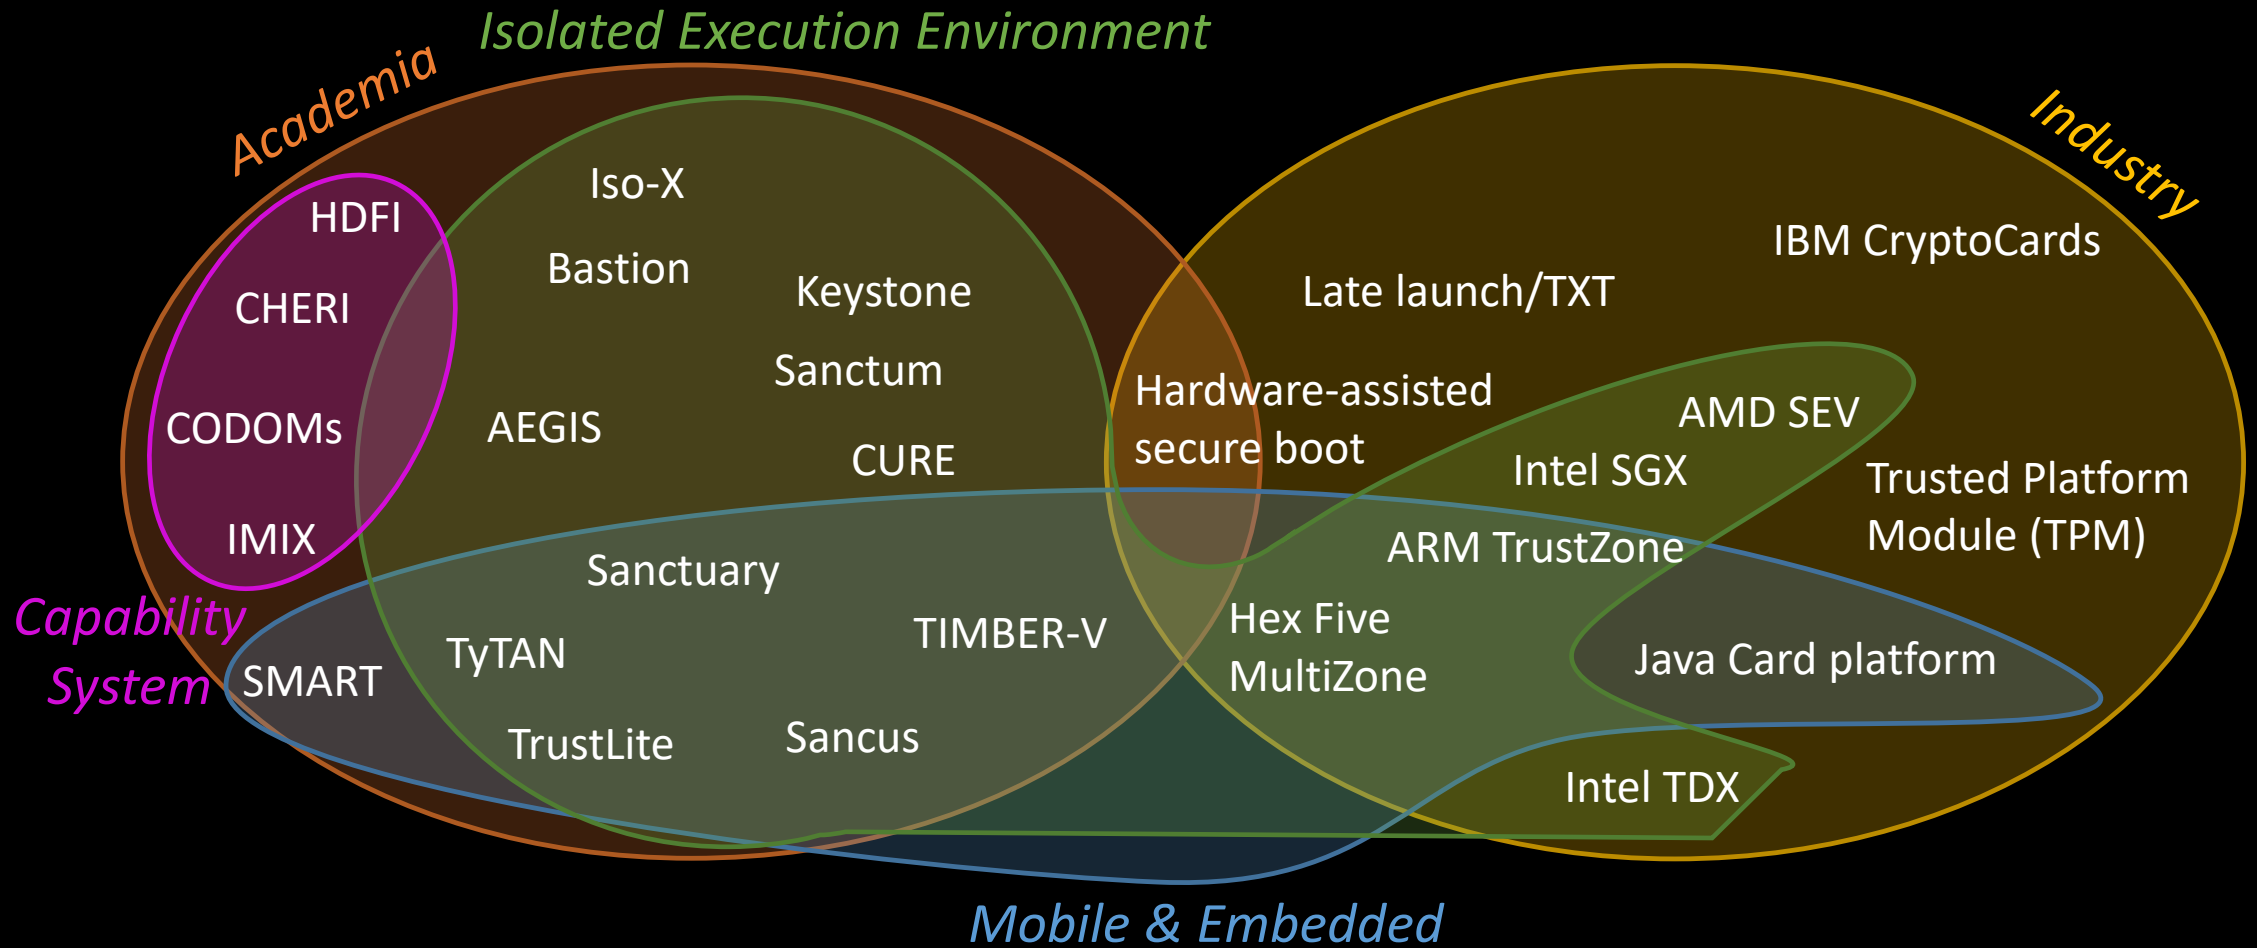
Trusted Computing Landscape



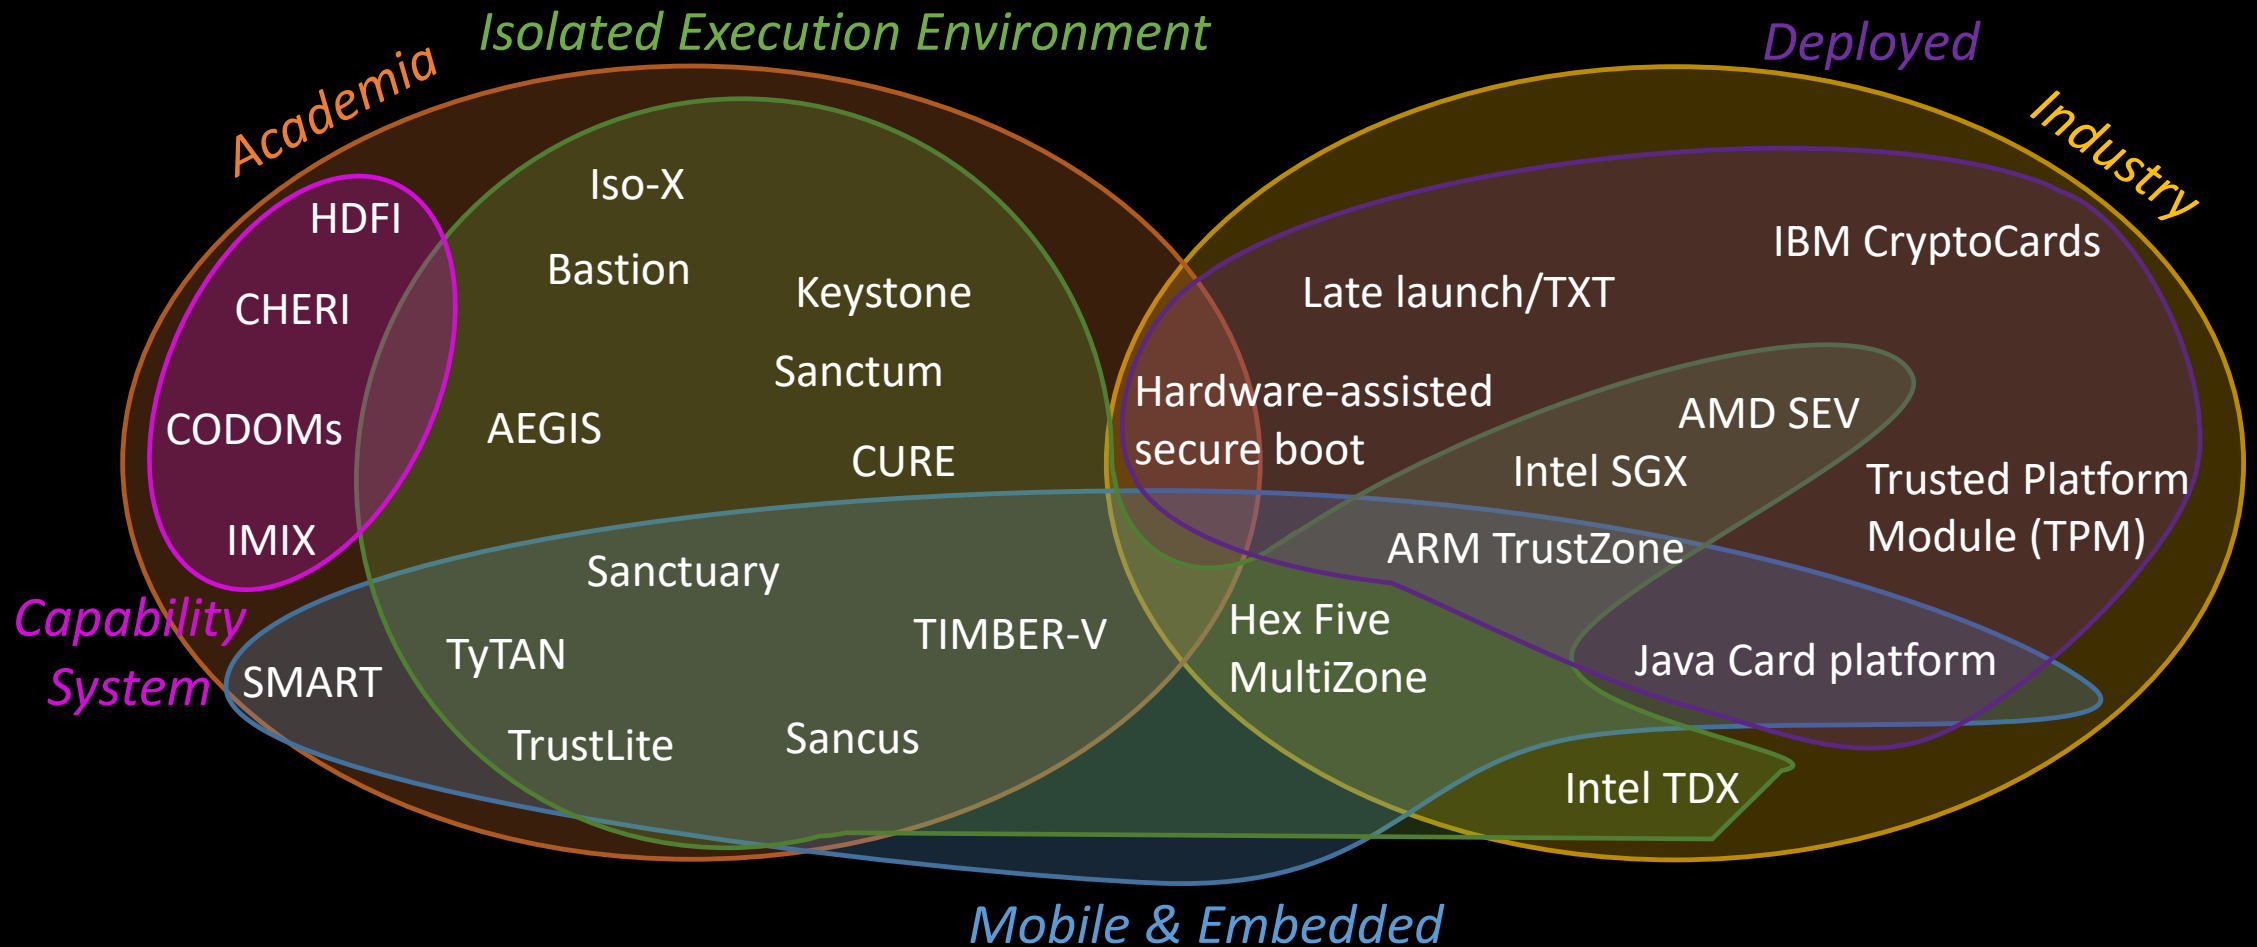
Trusted Computing Landscape



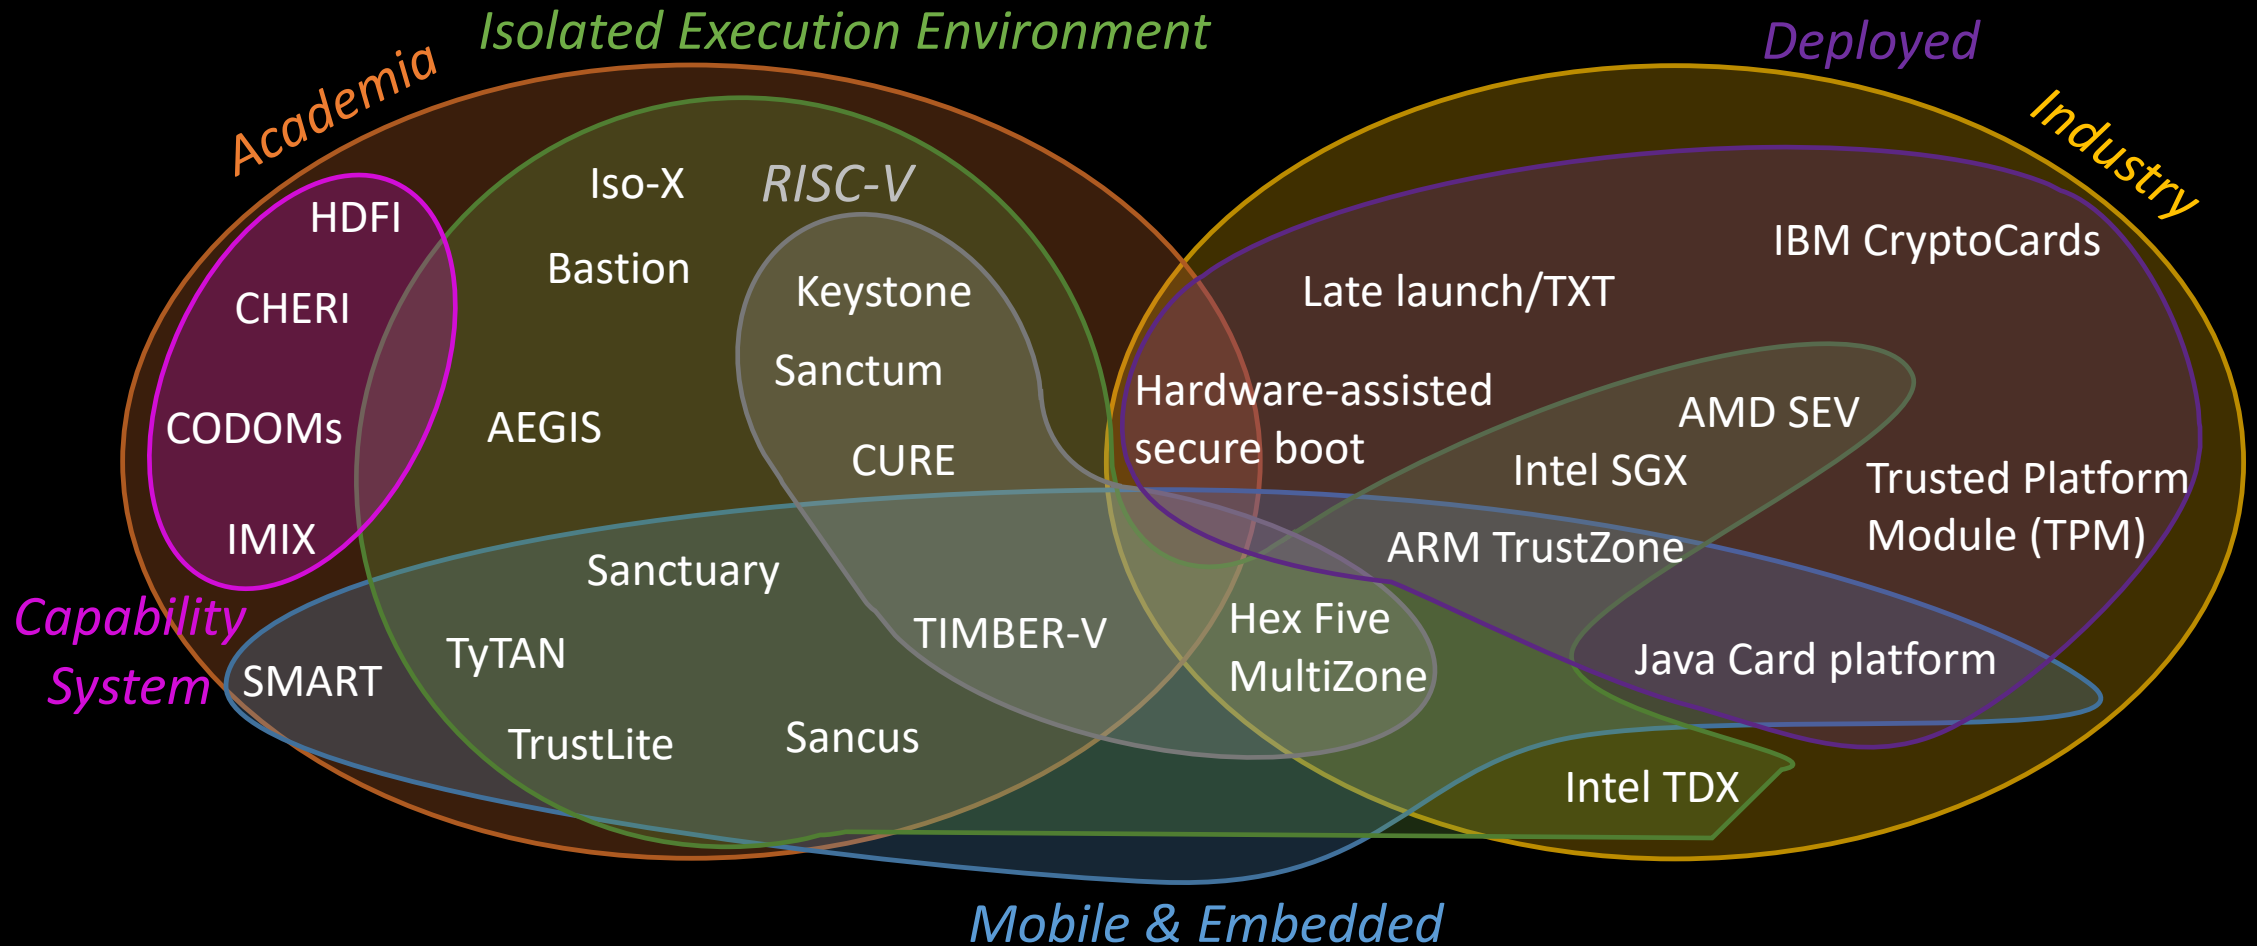
Trusted Computing Landscape



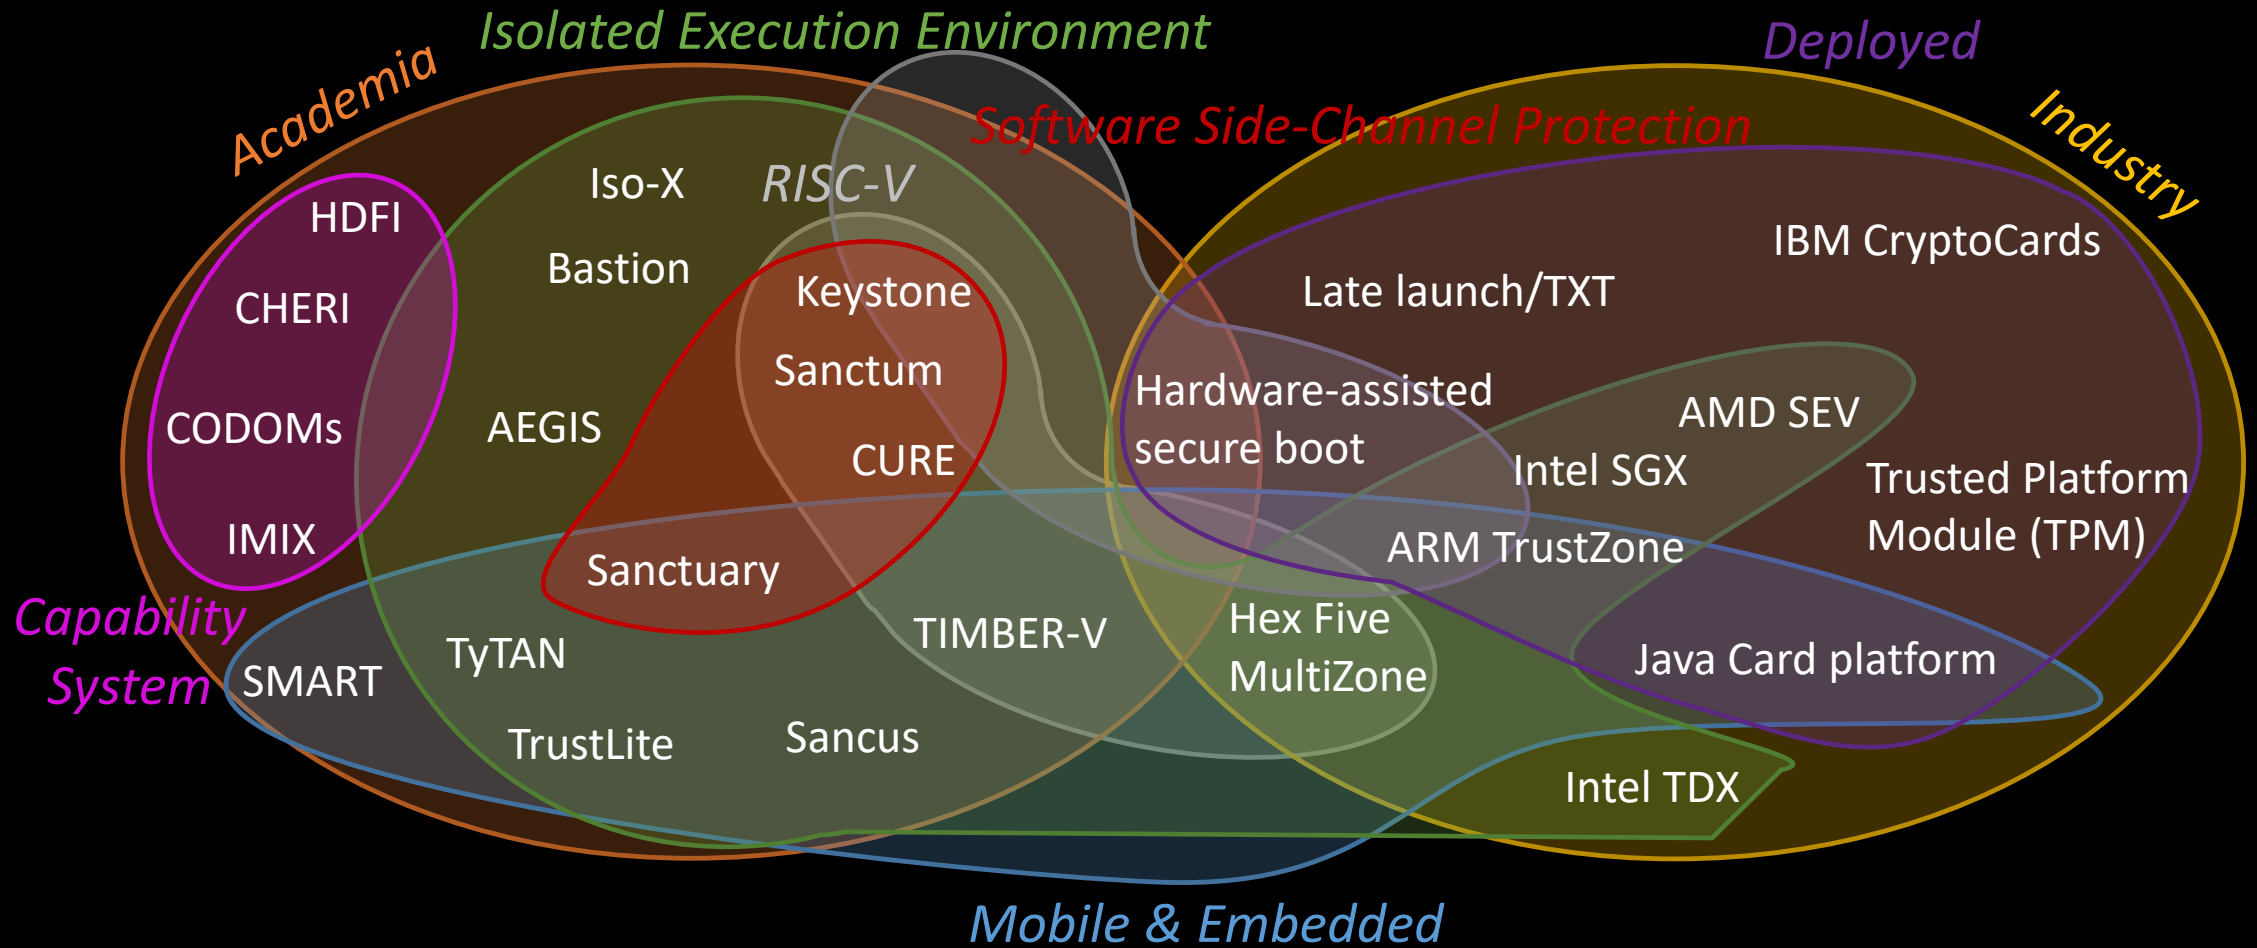
Trusted Computing Landscape



Trusted Computing Landscape



Trusted Computing Landscape

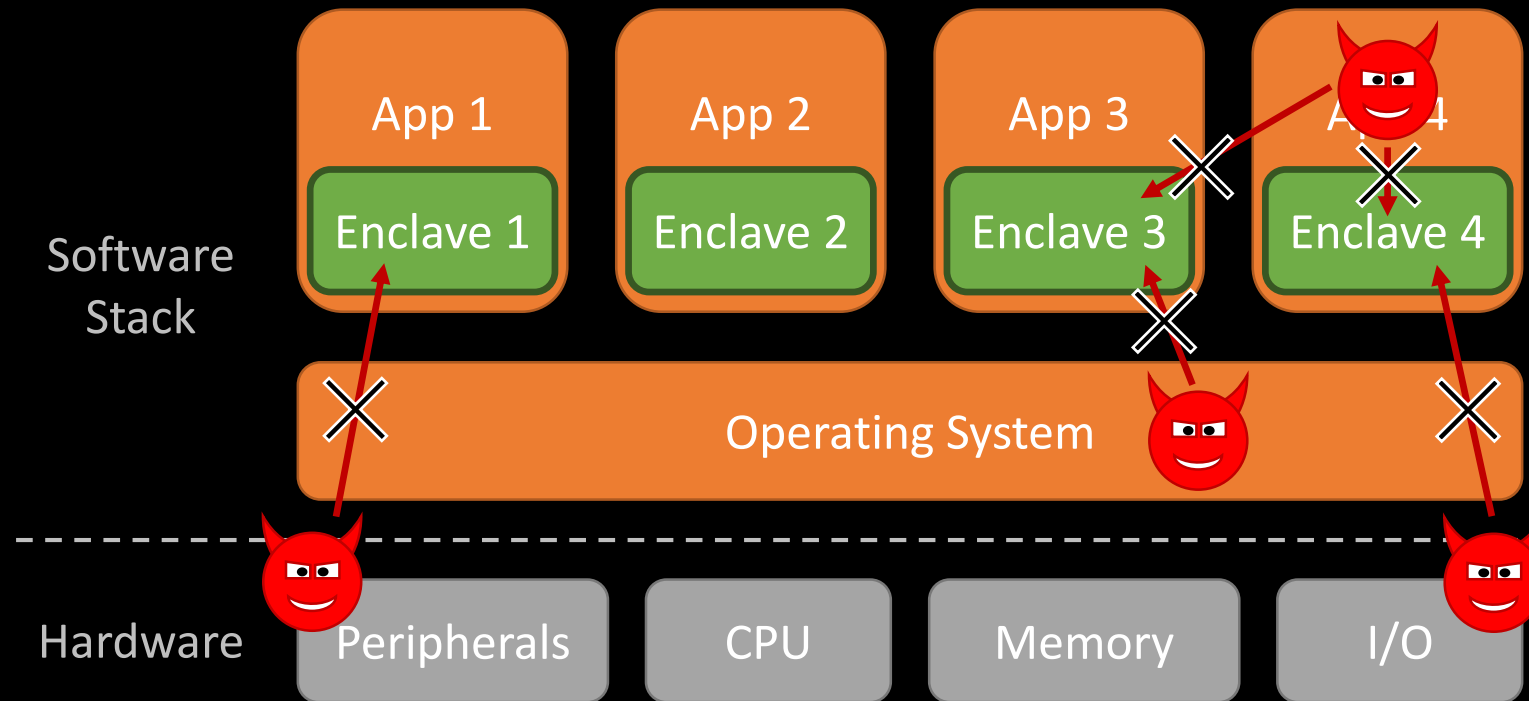


Example: Intel SGX




Intel Software Guard Extensions (SGX)

Assumptions: All software, and some hardware components, can be untrusted




Information Leakage



Real-World Consequences

Extracting 2048-bit RSA
decryption key from the enclave

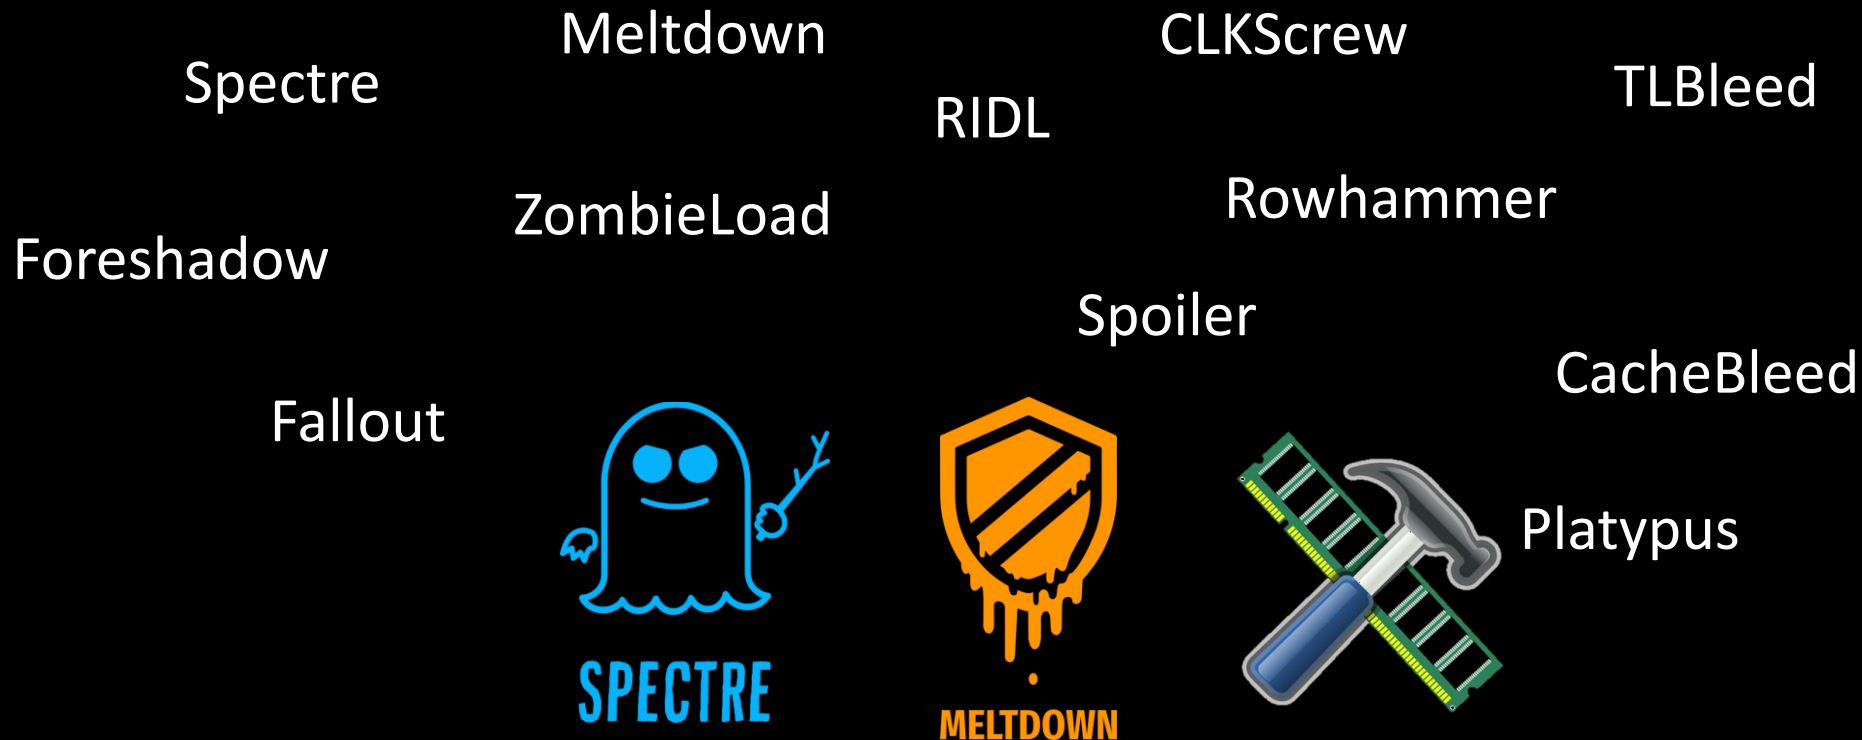


Extracting genome sequences
from the enclave



The Rise & Hype of Cross-Layer Exploits

- Recent microarchitectural/hardware-based attacks are exploiting issues that originate in the underlying hardware/microarchitecture



The Rise & Hype of Cross-Layer Exploits

- Recent microarchitectural/hardware-based attacks are exploiting issues that originate in the underlying hardware/microarchitecture



Towards Resilient Enclave-based Security Architectures

CURE: A Security Architecture with CUstomizable and Resilient Enclaves

Raad Bahmani, Ferdinand Brasser, Ghada Dessouky,
Patrick Jauernig, Matthias Klimmek, Ahmad-Reza Sadeghi, Emmanuel Stapf at USENIX Security 2021

at different privilege levels can be supported.

FR.2: Enclave-to-peripheral binding. Secure communication between enclaves and selected system peripherals, e.g., when offloading sensitive machine learning tasks to hardware accelerators [84], must be explicitly supported.

FR.3: Minimal hardware changes. The hardware changes required to integrate the proposed security primitives into a commodity SoC (cf. Section 2) must be minimal, no invasive changes to CPU internals must be required to enable a higher adoption of CURE in future platforms.

FR.4: Reasonable performance overhead. The performance overhead incurred during enclave setup and run time must be minimized and must not render the computer system impractical for certain uses cases or degrade user experience.

FR.5: Configurable protection mechanisms. Protection mechanisms against cache side-channel attacks must be applicable dynamically at run time and on a per-enclave basis.

5 Design of the CURE Architecture

CURE provides a novel design that addresses the requirements described above and provides a TEE architecture with strongly-isolated and highly customizable enclaves, which can be adapted to the requirements of the services they protect. Unlike other TEE architectures, which only provide a single enclave-type, CURE allows to freely define enclave boundaries and thus, different enclaves can be constructed, as shown in Figure 2. First, in Section 5.1, we describe the ecosystem

Then, we elaborate on the diff

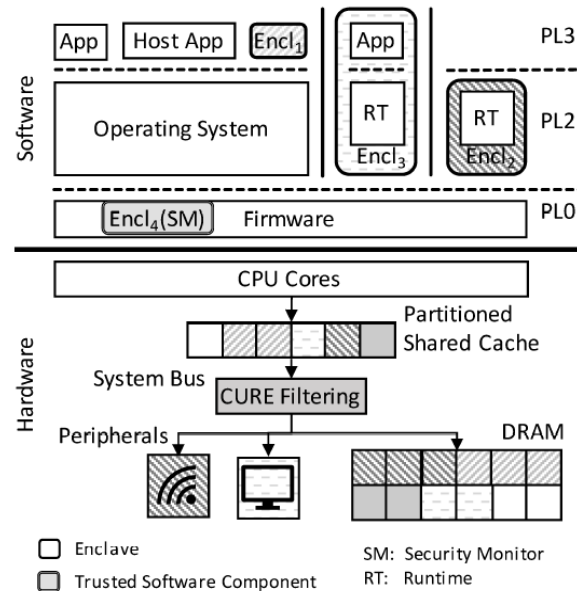


Figure 2: CURE privilege levels and enclave types, namely, user-space enclaves (Encl₁), kernel-space enclaves (Encl₂, Encl₃) and sub-space enclaves (Encl₄).

which is operated

L_{encl} is

- Multiple *types* of enclaves
- Security-critical tasks and services (e.g., remote attestation) managed by trusted software component (SM)
- Way-based cache partitioning on shared L2 cache
- Novel access control mechanism on system bus, minimal changes at processor
- Allows for secure binding between enclaves and peripherals

A person wearing a dark jacket and a headlamp is shown in a dark, industrial setting, possibly a server room or a hardware workshop. The person is leaning over a piece of equipment, illuminated by the beam of their headlamp. The background is dark with vertical lines, suggesting a server rack or a similar structure. The overall atmosphere is technical and focused.

Deep Dive into Hardware Vulnerabilities

International Hardware Security Capture-The-Flag Competition, HACK@ Franchise



Hack@DAC: Why and How?

- Deep dive into real-world hardware security vulnerabilities and detection methodologies
- Raise the bar for hardware security in the semiconductor industry
 - Lots of established techniques for software security assurance
 - But limited number of tools for hardware security assurance
- RISC-V SoC testbeds with injected bugs constructed in collaboration with Intel hardware security professionals
- Over the past 3 years: A total of 174 teams from academia and industry from all over the world
- Own investigation of the effectiveness of typical industry standard formal verification tools

Systematic RTL Bugs Construction

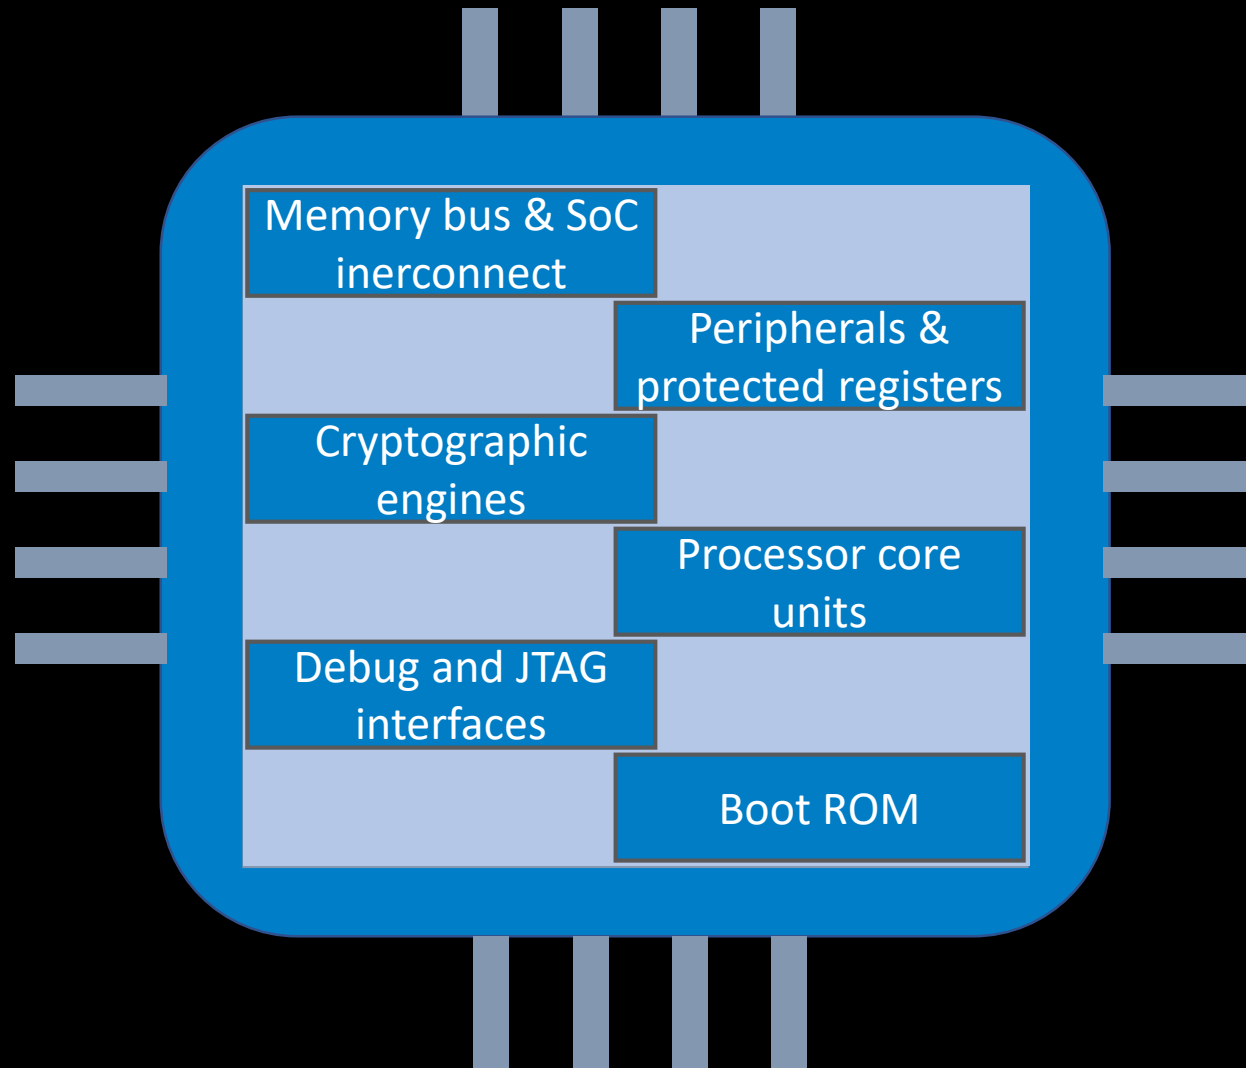


Spectrum of CVEs involving SW-exploitable hardware and firmware bugs

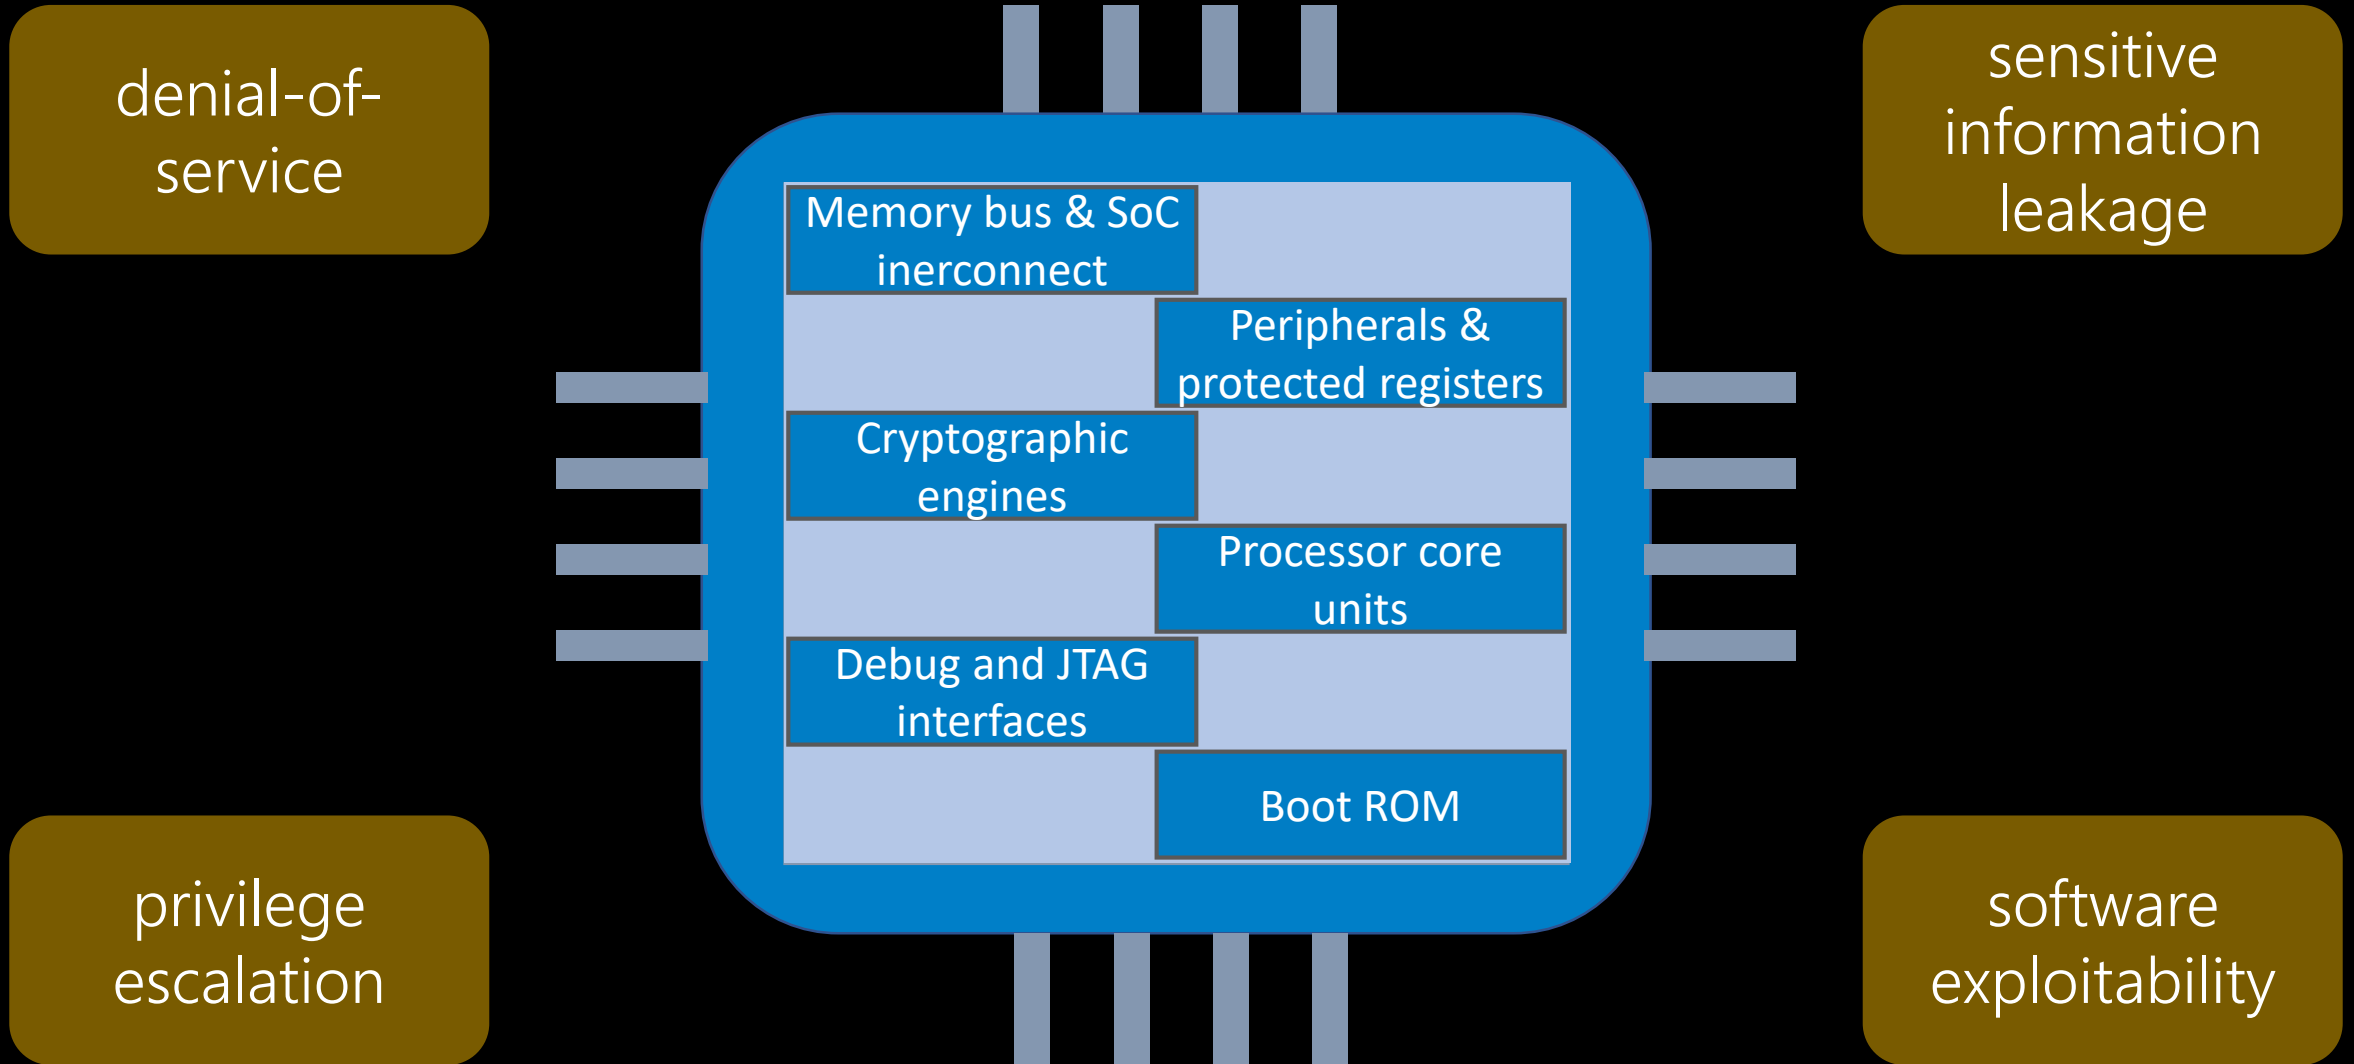


Real-world hardware bugs encountered by hardware and system security professionals at Intel

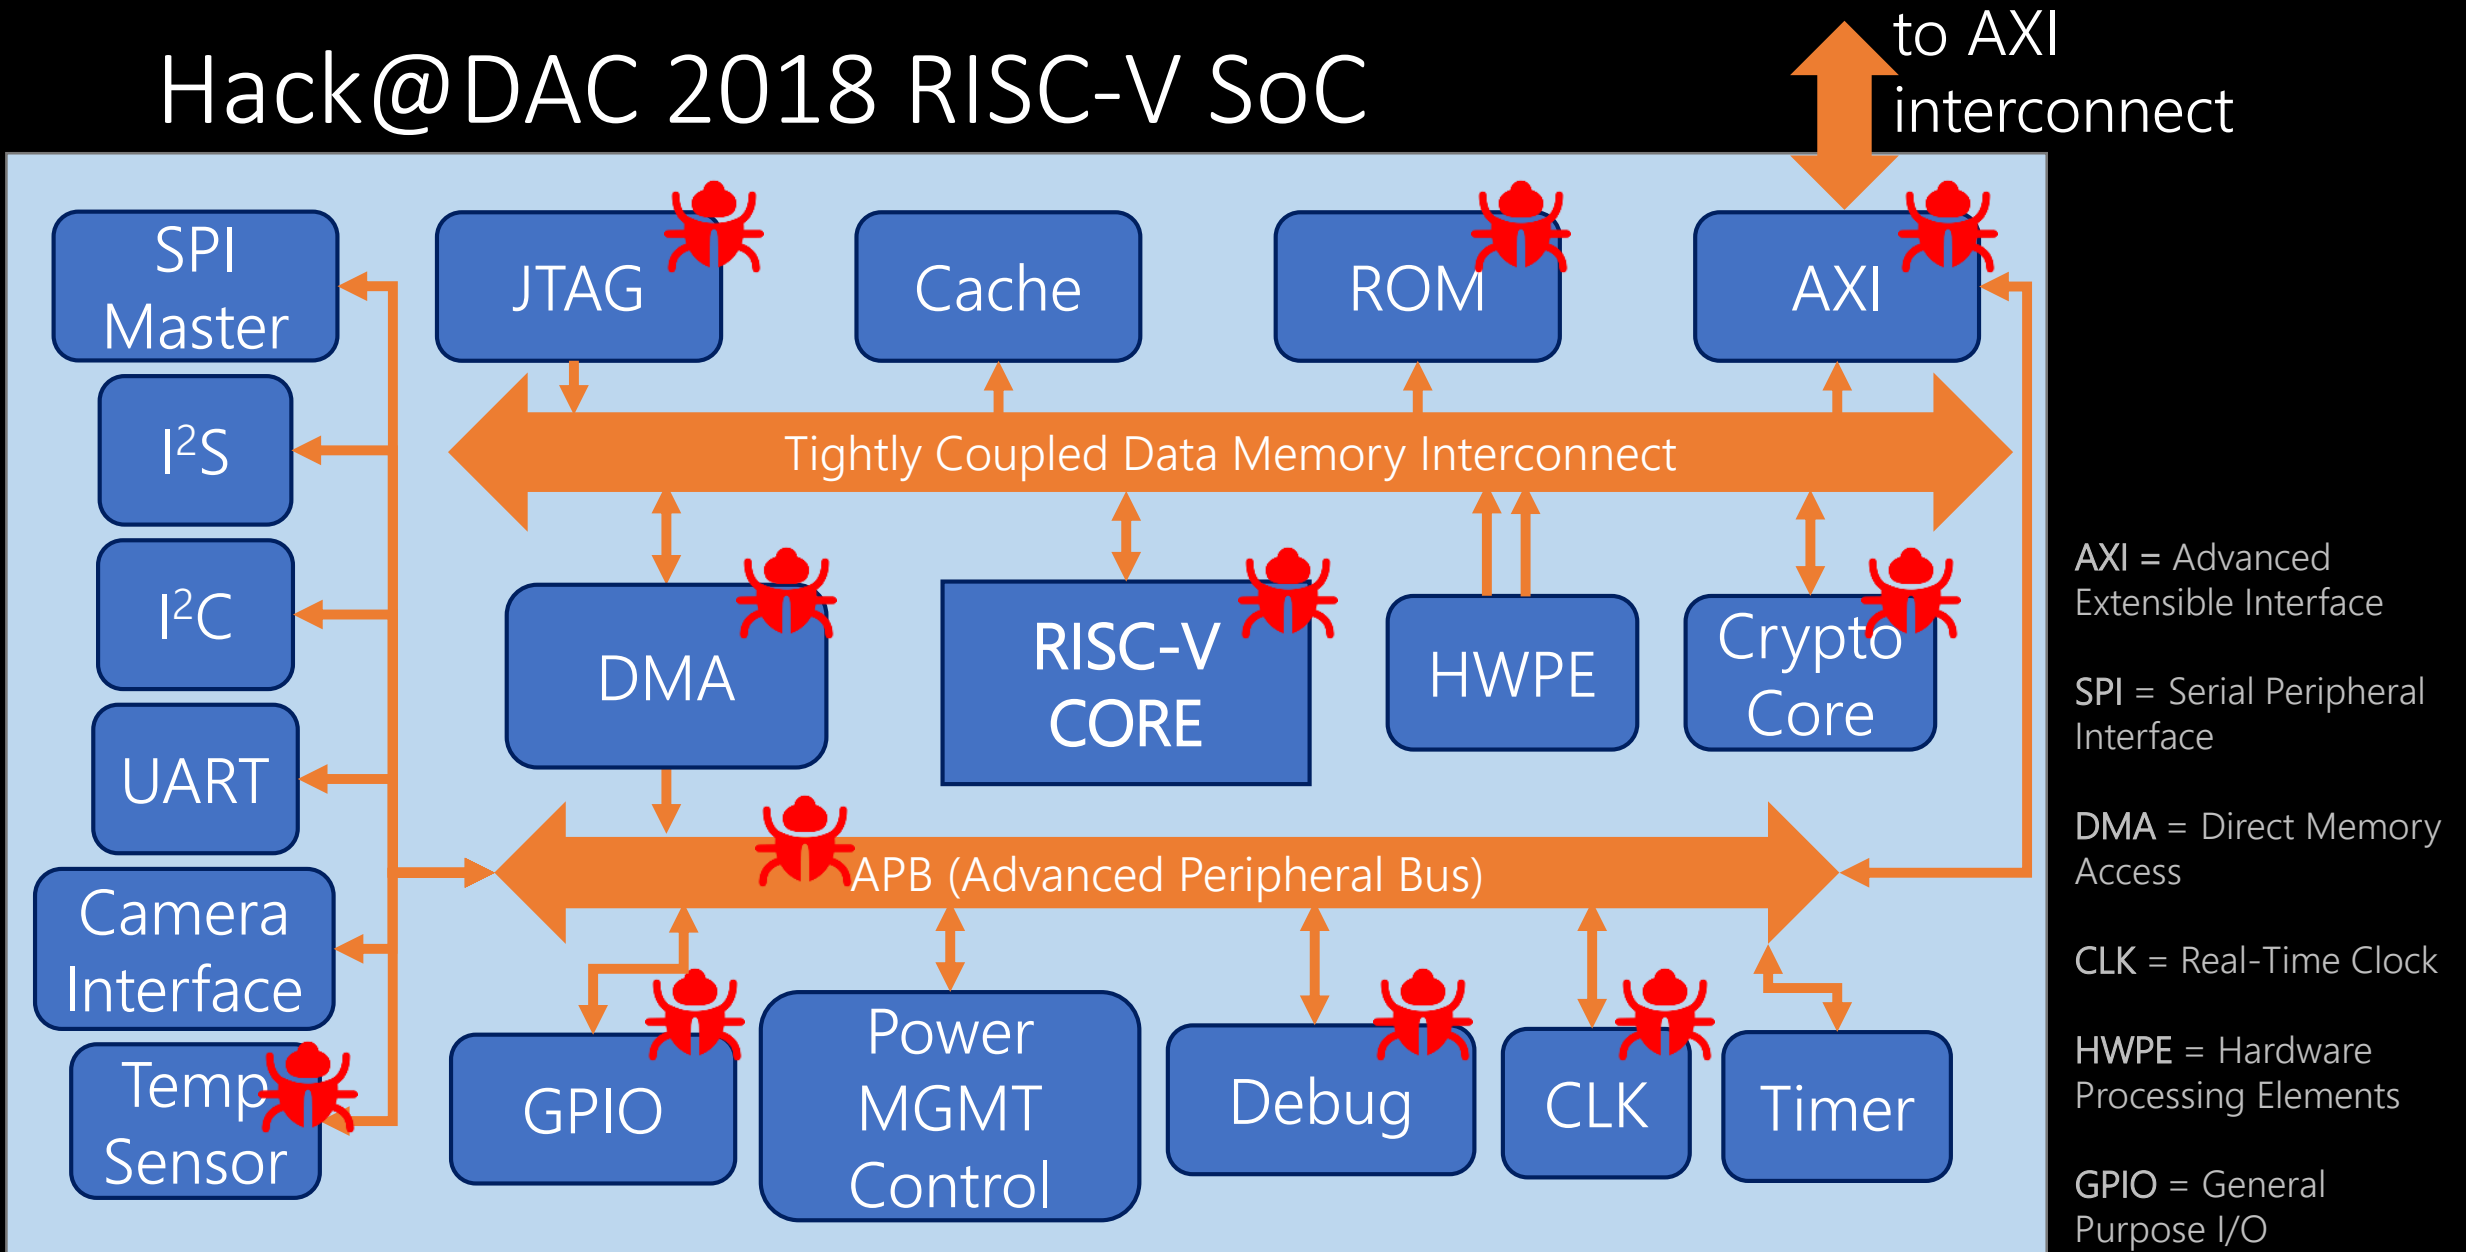
Systematic RTL Bugs Construction



Systematic RTL Bugs Construction

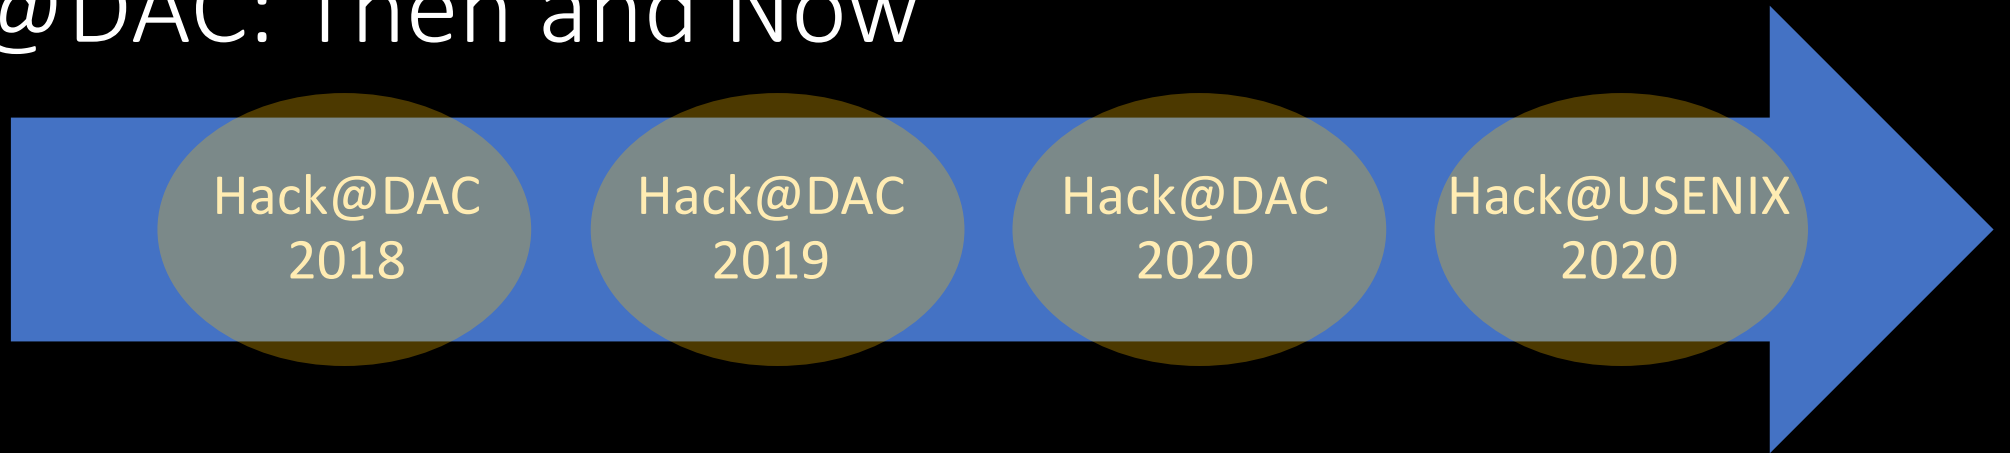


Hack@DAC 2018 RISC-V SoC



2 RISC-V SoCs used: PULPino & PULPissimo

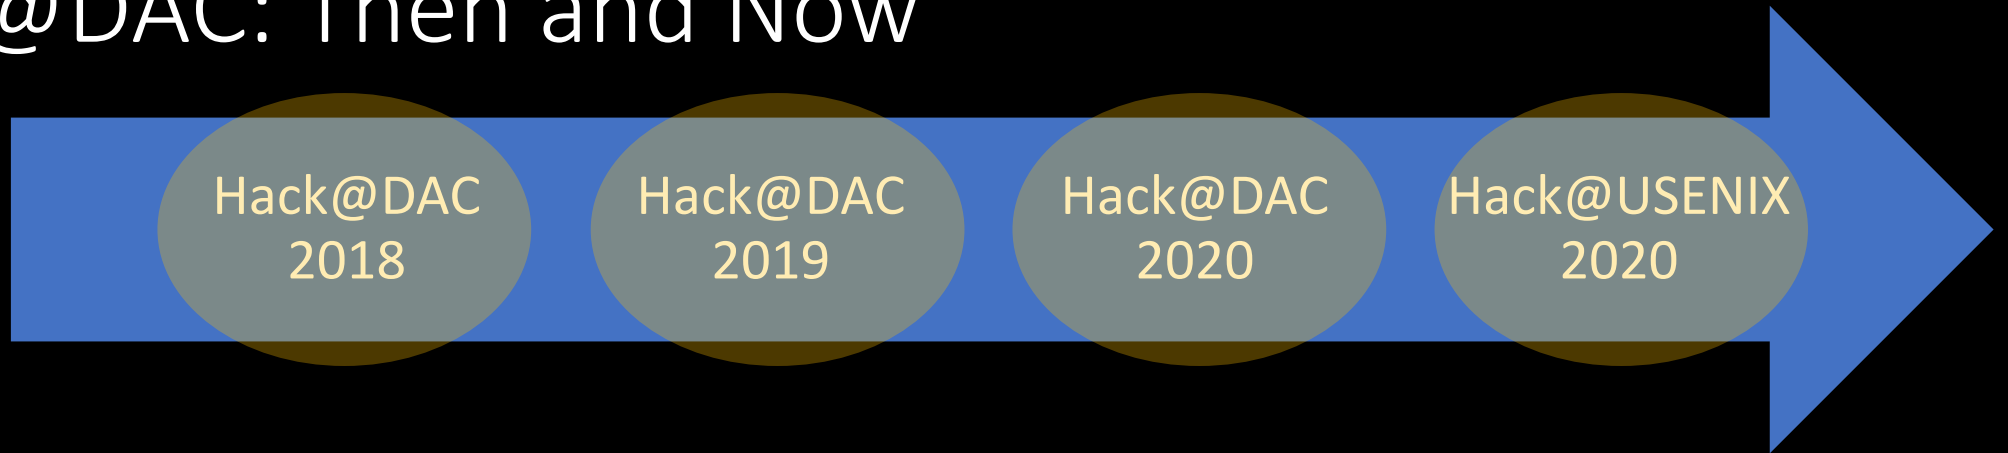
Hack@DAC: Then and Now



<https://hackat.events/dac18/>
<https://hackat.events/dac20/>

<https://hackat.events/dac19/>
<https://hackat.events/sec20/>

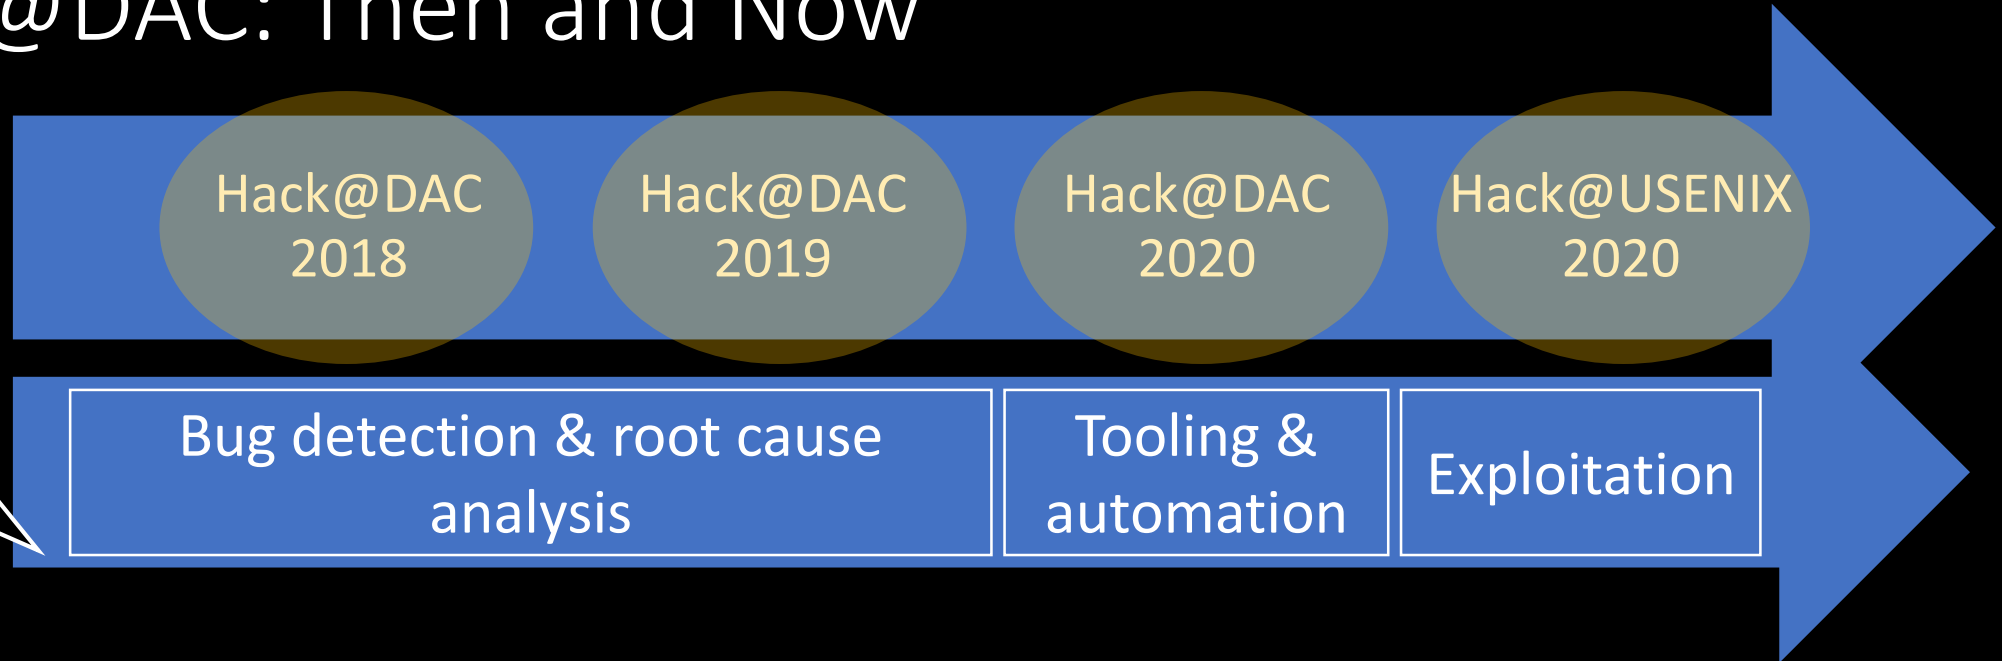
Hack@DAC: Then and Now



<https://hackat.events/dac18/>
<https://hackat.events/dac20/>

<https://hackat.events/dac19/>
<https://hackat.events/sec20/>

Hack@DAC: Then and Now



<https://hackat.events/dac18/>
<https://hackat.events/dac20/>

<https://hackat.events/dac19/>
<https://hackat.events/sec20/>

Hack@DAC: Then and Now

Hack@DAC
2018

Hack@DAC
2019

Hack@DAC
2020

Hack@USENIX
2020

Focus &
Scoring

Bug detection & root cause
analysis

Tooling &
automation

Exploitation

Complexity

Increasing SoC complexity, multiple privilege levels & MMU, new security features → more complex hardware-software vulnerabilities

<https://hackat.events/dac18/>
<https://hackat.events/dac20/>

<https://hackat.events/dac19/>
<https://hackat.events/sec20/>

Hack@DAC: Then and Now

Hack@DAC
2018

Hack@DAC
2019

Hack@DAC
2020

Hack@USENIX
2020

Focus &
Scoring

Bug detection & root cause
analysis

Tooling &
automation

Exploitation

Complexity

Increasing SoC complexity, multiple privilege levels & MMU, new security features → more complex hardware-software vulnerabilities

Insights &
Awareness


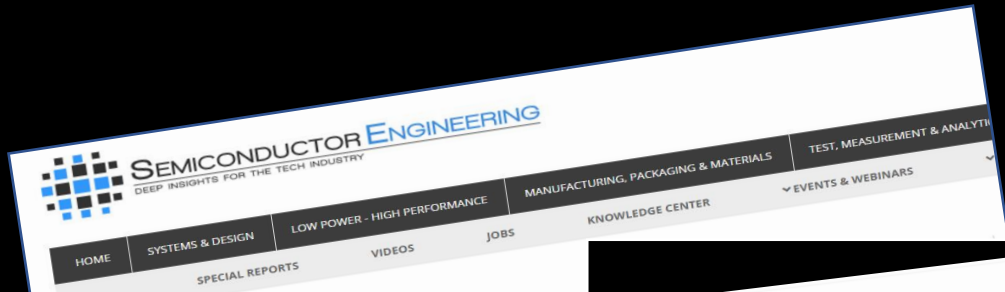
More interesting insights and results from the teams
Increasing awareness and attention in academia and industry

174 teams in
total over the
3 years

<https://hackat.events/dac18/>
<https://hackat.events/dac20/>

<https://hackat.events/dac19/>
<https://hackat.events/sec20/>

In the Press



Conclusion

- Trusted computing could not keep its promises
 - Still suffering from legacy issues
 - The impact of side-channels and transient execution were “ignored” ?
- Hardware security validation is still its infancy
- Current tools and metrics for security are limited
 - No fuzzing or symbolic execution for RTL code
 - Significant expertise, explicit definition of security property assertions and manual inspection required
 - Scalability and consolidated hardware/firmware verification are open challenges
- Real-world systems are highly proprietary
 - False sense of security („by-obscurity“)
 - Open-source hardware such as RISC-V may help, like Linux in software
- Customizable security architecture may be more successful towards more resilient computing platforms (see CURE, *Bahmani et al USENIX SEC 2021*)

What's Next for Hack@DAC?

- Hack@DAC goes to the cloud – Amazon's AWS FPGA instances to host the SoCs and open-source security assurance tools



- A learning hub on hardware security assurance for academia & industry
- Developing new effective & efficient hardware security assurance techniques, e.g. hardware fuzzing

Contact us:
ahmad.sadeghi@trust.tu-darmstadt.de

