March 8, 2024

Media Alert: Intel at Open Confidential Computing Conference 2024

Intel technologists to present advances in confidential computing and participate in a CTO panel discussion at this year's Open Confidential Computing Conference.

SANTA CLARA, Calif.--(BUSINESS WIRE)-- Join Intel experts for panel discussions and talks at this year's Open Confidential Computing Conference (OC3), a virtual event on March 13. Hosted by Edgeless Systems, OC3 is the premier event for security architects, cloud-native software engineers, IT security experts, CISOs, CTOs, security researchers and developers who want to learn about confidential computing.

This press release features multimedia. View the full release here: https://www.businesswire.com/news/home/20240308183003/en/

Intel technologists will present advances in confidential computing and participate in a CTO panel discussion at this year's Open Confidential Computing Conference. The virtual event is March 13. (Credit: Intel Corporation)

Protecting systems and sensitive, confidential or regulated data, especially while in use, has never been more critical. As computing moves to span multiple environments – from on-prem to public cloud to edge – organizations need protection controls that help safeguard sensitive intellectual property (IP) and workload data wherever that data resides. Learn how Intel, together with its partners and customers, builds the trusted foundation for protecting AI workloads and computing in a data-centric world.

Open Confidential Computing Conference (OC3) 2024
When: March 13, 2024
Where: Virtual Event
Registration: Free to attend

Confidential Cloud Native Attestation – Challenges and Opportunities

Confidential computing brings with it tamper-resistant registers to measure digital ingredients, akin to what the Trusted Computing Group's TPM 2.0 offers, such as BIOS, firmware, kernel and beyond. Clouds are varied in their infrastructure and multiple confidential computing vendors, each potentially with multiple product generations, offering confidential CPUs, GPUs and other special-purpose processing units. Further, there are at least three flavors of confidential virtual machine (CVM) use – whole confidential Kubernetes clusters, launching traditional virtual machine payloads as a CVM using KubeVirt or Virtual Kubelet, or running a confidential container, like CoCo. What should one measure, particularly with confidential clusters where workloads come and go? The trick lies in capturing invariants and keeping them separate to not have a combinatorial explosion of values to register in an attestation service as good values. Further, what is the essence that

we must keep invariant to protect the workloads in the various contexts?

In this talk, Mikko Ylinen, senior Linux software engineer at Intel, andMalini Bhandaru, senior principal engineer and cloud native architect at Intel, will share an overview of the landscape followed by a proposal to measure invariants in a typed data structure with a summary in the CVM tamper-resistant measurement registers and how it supports scalable attestation. It will be illustrated in the context of Intel® Trust Domain Extensions (Intel® TDX) using established techniques, such as CoCo, Linux IMA, dm-verity or CCNP.
When: Wednesday, March 13, 8-8:30 a.m. PDT
Where: Virtual through conference platform

Confidential Computing in 2024 – Innovating Secure and Scalable Solutions

We are on the cusp of a transformative era. Technical readiness and market momentum will converge in 2024 to accelerate growth and adoption of confidential computing. This session, presented by Anand Pashupathy, vice president and general manager of Security Software and Services at Intel, will offer a comprehensive assessment of the industry's progress as the industry aligns with imperatives described in Intel CTO Greg Lavender's 2023 keynote at OC3. Pashupathy will also provide an in-depth look at Intel's strategic initiatives to address remaining adoption barriers and elevate confidential computing to new levels of security, performance and user-friendly scalability.
When: Wednesday, March 13, 10:30-11 a.m. PDT
Where: Virtual through conference platform

Tightening Side Channel Protections with Intel SGX AEX-Notify

Intel® Software Guard Extensions (Intel® SGX) supports the creation of shielded enclaves within unprivileged processes. Code and data within an enclave cannot be read or modified by the operating system or hypervisor, nor by any other software. However, side-channel attacks can be challenging to comprehensively mitigate. This talk by Scott Constable, research scientist, Cybersecurity and Computer Security at Intel, will give an overview of AEX-Notify, a new flexible architecture extension that makes enclaves interrupt-aware: Enclaves can register a trusted software handler to be run after an interrupt or exception (such as a fault). AEX-Notify can be used as a building block for implementing countermeasures against different types of interrupt- and fault-based attacks. AEX-Notify is available on 4th Gen Intel® Xeon® Scalable processors and newer products with Intel SGX and is also backward-portable to all older server products via a microcode update. The Intel SGX SDK for Linux now supports a default trusted software handler that mitigates attacks that use interrupts or exceptions to exert fine-grained control over enclave execution, for example, by forcing a single enclave instruction to execute each time the enclave is entered.
When: Wednesday, March 13, 11-11:15 a.m. PDT
Where: Virtual through conference platform

Asterinas: A Safe and Efficient Rust-Based OS Kernel for TEE and Beyond

In the realm of OS kernels, particularly those within virtual machine (VM) trusted execution environments (TEEs), memory safety is a paramount concern. Rust, known for its safety features, aids in developing secure kernels but is not a panacea. Firstly, Rust's unsafe features, such as pointer dereferencing and inline assembly, are necessary for low-level, error-prone tasks, often permeating the codebase. Secondly, the guest kernel in a VM TEE

often processes untrusted inputs (over 1,500 instances in Linux, per Intel's estimation) from the host (through hypercalls, MMIO, etc.), posing a risk of exploitable memory safety vulnerabilities.

This leads us to explore how effectively a Rust-based kernel can minimize its trusted computing base (TCB) against memory safety threats, including Iago attacks. The response is Asterinas: a safe and efficient OS kernel crafted in Rust, offering Linux ABI compatibility. Asterinas introduces a groundbreaking framekernel OS architecture. This design splits the kernel into two distinct halves within the same address space: the framework and services. The framework is the sole domain allowed to utilize unsafe Rust features, providing a high-level, safe and sound API for the services, which are exclusively developed in safe Rust. The services are responsible for providing most of the OS functionalities, including enabling all peripheral devices. As the entire kernel resides in the same address space, different parts of the kernel can communicate in the most efficient way.

In this talk, Chuan Song, principal engineer at Intel, and Hongliang Tian from Ant Group dive into the design and implementation of Asterinas. They will spotlight the pioneering framekernel OS architecture and show how the kernel is ported to and fortified for Intel TDX.
When: Wednesday, March 13, 11:15-11:45 a.m. PDT
Where: Virtual through conference platform

Seamless Attestation of Intel TDX and NVIDIA H100 TEEs for Confidential AI

AI is now the most significant workload in data centers and the cloud. It's being embedded into other workloads used for standalone deployments and distributed across hybrid clouds and the edge. Many of the demanding AI workloads require hardware acceleration with a GPU. Many AI models are considered priceless intellectual property – companies spend millions of dollars building them, and the parameters and model weights are closely guarded secrets. The datasets used to train these models are also considered highly confidential and can create a competitive advantage. As a result, data and model owners are looking for ways to protect these, not just at rest and in transit, but while in use as well.

Intel and Nvidia deliver confidential computing technologies that establish independent TEEs on the CPU and GPU, respectively. For a customer, this presents an attestation challenge, requiring attestation from two different services to gather the evidence needed to verify the trustworthiness of the CPU and GPU TEEs. Intel and Nvidia are collaborating to provide a unified attestation solution for customers to verify the trustworthiness of the CPU and GPU TEEs for confidential computing based on Intel® Xeon® processors with Intel® Trust Domain Extensions (Intel® TDX) and Nvidia Tensorcore H100 GPUs.

This session presented by Raghu Yeluri, senior principal engineer and lead security architect at Intel, and Michael O'Connor of Nvidia will look at the TEE architectures and how they are enabled for seamless attestation of the two TEEs using Intel® Trust Authority and Nvidia Remote Attestation Service (NRAS).
When: Wednesday, March 13, 12-12:30 p.m. PDT
Where: Virtual through conference platform

The Status Quo and Potential of Confidential AI

OC3 brings back this exciting panel with industry leaders, this time to discuss confidential AI.

The panelists will discuss what confidential AI is, use cases, technical challenges, regulatory incentives and limits. Panel members will also make predictions about the future of this technology. Will AI be the “killer app” for confidential computing? When will confidential computing be the standard for AI?

This panel will feature Greg Lavender, executive vice president, chief technology officer (CTO) and general manager of the Software and Advanced Technology Group (SATG) at Intel, alongside the CTOs of AMD and Microsoft Azure, and the vice president of Hyperscale and HPC from Nvidia.
When: Wednesday, March 13, 1-1:30 p.m. PDT
Where: Virtual through conference platform

Private Data Exchange – Leveraging Confidential Computing to Combat Human Trafficking and Modern Slavery

This session from Hope for Justice, Intel and Edgeless Systems will unpack the Private Data Exchange, an exciting and innovative project leveraging confidential computing as a powerful tool in the fight against human trafficking and modern slavery.

Organizations like Hope for Justice and Slave-Free Alliance have joined the effort to find victims, as well as perpetrators. The Private Data Exchange is an innovative project in partnership with Intel and Edgeless Systems to develop a platform that can encrypt data to protect sensitive information, knowing that behind it are the private lives of people who’ve been abused and traumatized and need protection.

Intel technology enables the Private Data Exchange to leverage confidential computing, which processes sensitive data out of view from unauthorized software or system administrators. The data is encrypted and processed in memory, lowering the risk of exposure to the rest of the system, which can compromise it. Confidential computing relies on hardware-based controls, enabled by Intel SGX enclaves.

This project will enable global organizations to collaborate and share analyses to prevent human trafficking, respond to situations of exploitation and ensure victims receive the support they need, while shielding their confidential information or regulated data.
When: Wednesday, March 13, 2:45-3 p.m. PDT
Where: Virtual through conference platform

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore’s Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers’ greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel’s innovations, go to newsroom.intel.com and intel.com.

© Intel Corporation. Intel, the Intel logo and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

View source version on businesswire.com:
https://www.businesswire.com/news/home/20240308183003/en/

Jennifer Foss
425-765-3485
jennifer.foss@intel.com

Source: Intel Corporation