Open Confidential Computing Conference 2025 Recap

27 March 2025

Key insights from the premier event for confidential computing

Edgeless Systems

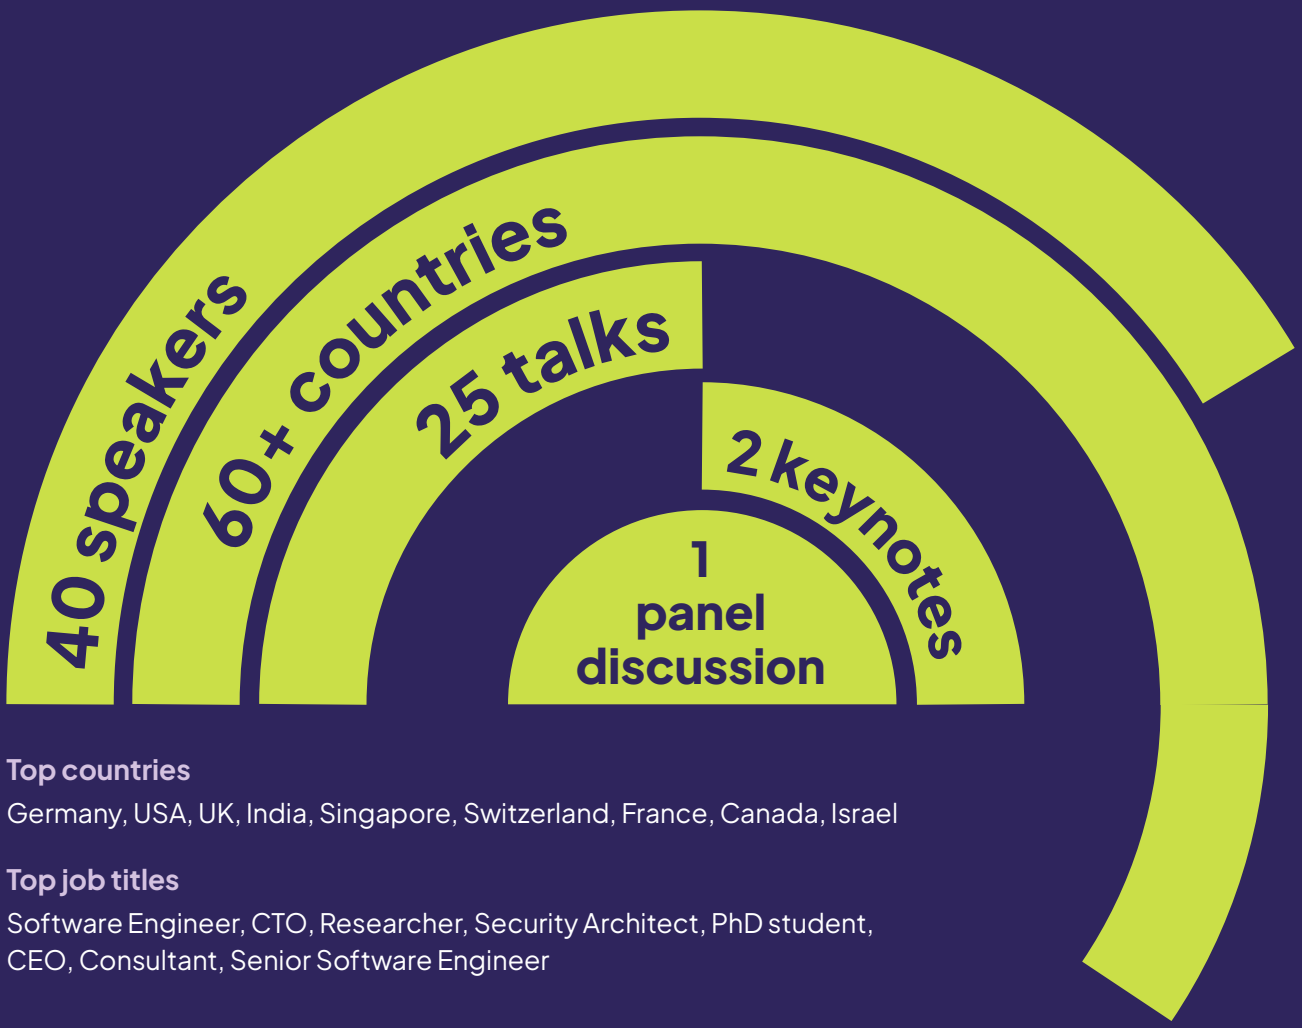

Top countries

Germany, USA, UK, India, Singapore, Switzerland, France, Canada, Israel

Top job titles

Software Engineer, CTO, Researcher, Security Architect, PhD student, CEO, Consultant, Senior Software Engineer

Growing strong.

OC3 2025 reached a significant milestone as our first hybrid conference. While continuing to serve our global audience through streaming, we gathered with the community in Berlin, creating a valuable blend of online access and personal interaction that strengthens our expanding network.

The event received outstanding feedback from both online and on-site attendees, achieving an average user score of 9/10. Participants showed strong engagement, with an average time spent of approximately 5 hours interacting with conference content.

176	9.0	5h 4min
Attendees on-site in Berlin	**Average user score**	**Average time spent per user**

It was a blast.

Held in Berlin and streamed globally, this year's OC3 conference marked a clear shift from theory to impact. 40 speakers—from major cloud providers to startups and academic researchers—shared how confidential computing is reshaping AI, infrastructure, attestation, and application design.

The conference featured live demos, real-world case studies, and panel discussions that moved beyond vision to practical implementation. Whether it was securing battlefield deployments, enabling privacy-preserving AI, or creating interoperable attestation standards, the talks showed that confidential computing is no longer experimental—it's becoming the trusted foundation for modern, secure computing.

With 180 participants on-site and more than 600 online, OC3 2025 confirmed what many in the community have felt: confidential computing has momentum, and it's accelerating fast.


Welcome!

1 Panel Discussion

2 Keynotes

3 Ecosystem & Foundations

4 AI

5 Apps & Solutions

6 Attestation

Panel Discussion

Industry leaders' panel: The future is confidential.

A highlight of this year's OC3 was the panel discussion on the status quo and future of confidential computing. Daniel Rohrer (VP Software Security, NVIDIA), Greg Lavender (EVP & CTO, Intel Corporation), Mark Papermaster (EVP & CTO, AMD) and Mark Russinovich (CTO & Deputy CISO at Microsoft) covered the latest advancements in confidential computing, from use cases to industry collaborations and breakthroughs.

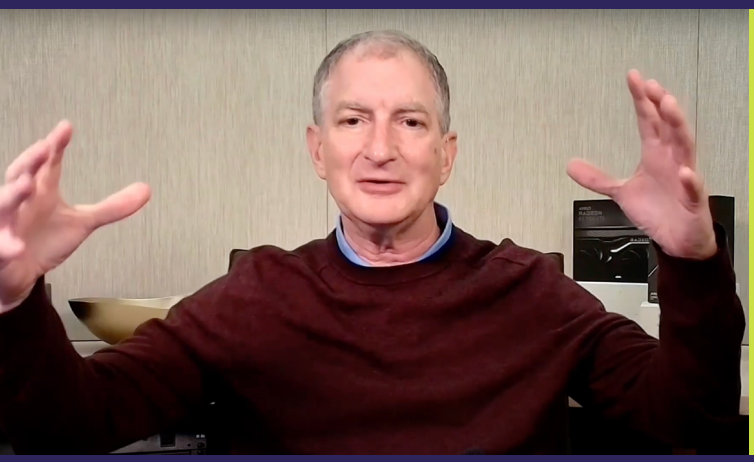

> **We're just at the beginning of this era of confidential computing. I think we're going to see a knee of the curve."**
>
> Mark Papermaster, CTO, AMD

AI is driving the next wave of confidential computing

The rise of AI, especially large language models and multimodal systems, is pushing confidential computing into the spotlight. Protecting sensitive training data, enabling federated learning, and securing inference at scale are becoming must-haves — and confidential computing is emerging as the default security foundation for these use cases.

Confidential computing is now core infrastructure

The technology has evolved from a niche solution into foundational infrastructure, especially for secure AI workloads. It's being adopted across industries, with companies like Microsoft moving critical systems (e.g., payment platforms) to confidential VMs, and NVIDIA pushing confidential containers for simplified deployment.

Enterprise adoption is accelerating fast

We're nearing feature parity with non-confidential systems, performance trade-offs are shrinking, and deployment is becoming user-friendly—soon to be as easy as ticking a checkbox. This growing ease of use is fueling broader enterprise adoption, beyond just cloud providers.

▶ **Watch industry's leaders panel discussion on demand**

Keynote

Keynote highlights

This year's OC3 featured two captivating keynotes from industry leaders.

Richard Grisenthwaite — Ensuring Confidentiality on Custom Silicon

> **Security is the greatest challenge computing needs to address to meet its full potential as computing becomes more and more central to our lives."**
>
> Richard Grisenthwaite, EVP & Chief Architect, Arm

Richard Grisenthwaite (EVP & Chief Architect, Arm) explored how the semiconductor industry is shifting from monolithic chips to modular chiplet architectures, enabling more scalable and customizable systems. This approach allows standard compute elements to be paired with specialized components like AI accelerators, supported by ARM's Chiplet System Architecture (CSA) and collaborations with partners such as Samsung and Rebellions.

As computing becomes more central to society, Grisenthwaite stressed that security is the defining challenge. ARM addresses this with its Confidential Computing Architecture, which uses hardware-isolated realms, memory encryption, and attestation to protect both data and AI models. A key example is Fujitsu's MONAKA SoC, which combines chiplets and confidential computing to deliver secure AI performance in data centers.

▶ **Watch the Richard Grisenthwaite's keynote on demand**

Daniel Rohrer — Securing AI's Third Dimension: Scaling Trust for Autonomous Intelligence

Daniel Rohrer (VP of Software Product Security, Architecture and Research at Nvidia) outlined the escalating demands of today's agentic AI systems, which require exponentially more compute and involve complex, cross-organizational data flows. Meeting this challenge requires both scale and trust—leading to the development of the Blackwell architecture, a 7x performance leap over Hopper, with full confidential computing support. It enables secure, high-performance AI with features like CUDA compatibility, advanced attestation using OPA, and multi-GPU scalability. Rohrer emphasized that industries from healthcare to finance now rely on confidential computing to safely unlock AI's full potential. **As AI becomes more autonomous, confidential computing is no longer optional—it's foundational.**

▶ **Watch Daniel Rohrer's keynote on demand**

> **"As we scale up compute, we also need to scale this notion of trust. And that's where confidential computing comes roaring to the front, as absolutely necessary, to achieve what we need in this new agentic space."**
>
> Daniel Rohrer, VP of Software Product Security, Nvidia

Ecosystem & Foundations

Ecosystem & Foundations

The goal we have is (...) making the amount of thought that needs to go into confidential computing as minimal as possible.”

Moritz Eckert, VP Product & Technology, Edgeless Systems

Moritz Eckert (Edgeless Systems) highlighted that the complexity of confidential containers lies not in the concept, but in the integration.

Despite offering strong isolation guarantees, confidential containers remain underused due to the operational burden they place on developers and platform teams. Eckert emphasized that the real barrier is usability: managing identity, attestation, and trust relationships in environments like Kubernetes that were never designed with untrusted infrastructure in mind. Making confidential computing accessible requires abstracting these challenges without sacrificing security, as shown in use cases from Airbus and Privatemode. The open-source framework Contrast was presented as one approach to addressing this integration gap.

▶ **Watch Moritz Eckert’s talk on demand**

Jörg Rödel (SUSE) and Chris Oo (Microsoft) demonstrated how guest-side infrastructure is becoming essential for running unmodified operating systems in secure environments.

Rödel presented COCONUT-SVSM, which offers secure services like TPMs inside encrypted CVMs, independent of potentially compromised hypervisors. Oo presented OpenHCL, a Rust-based paravisor which provides core virtualization features from within the VM itself, enabling OS compatibility with technologies like AMD SEV-SNP and Intel TDX. Together, these talks show a shift toward redesigning runtime infrastructure to live entirely within the trusted boundary of the guest.

▶ **Watch the Jörg Rödel's talk on demand**

▶ **Watch the Chris Oo's talk on demand**

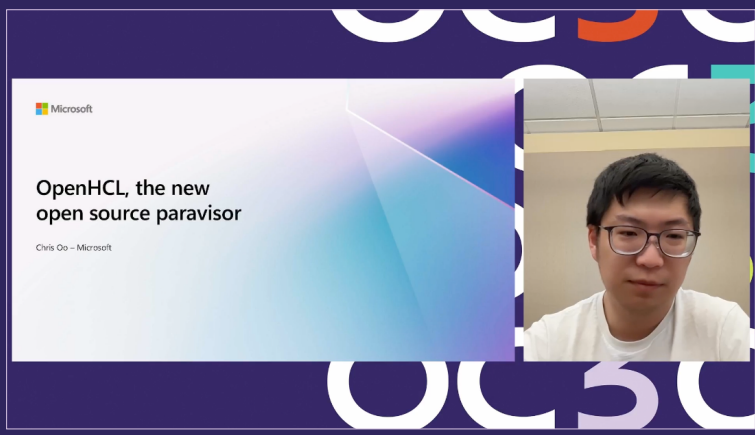

Vikas Bhatia (Microsoft) and Ijlal Loutfi (Canonical) showed how confidential computing is being embedded into cloud and edge platforms as a baseline security model.

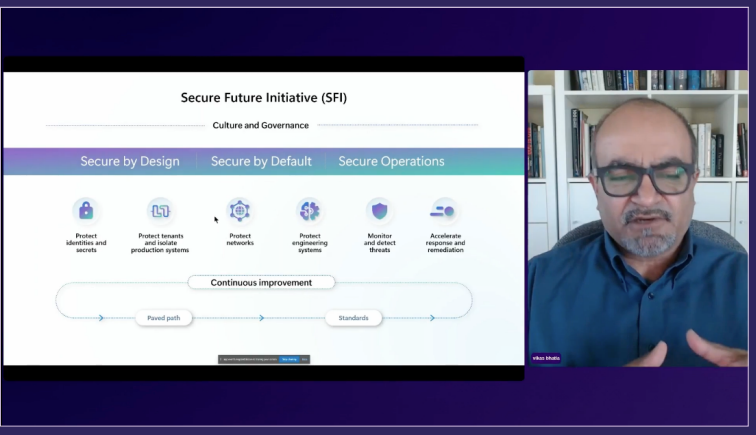

Bhatia outlined Microsoft's trajectory from early SGX adoption to Confidential Clean Rooms and Confidential AI, which now protect GPU memory and are used in production for identity, finance, and secure collaboration.

▶ Watch the Vikas Bhatia's talk on demand

Loutfi introduced Ubuntu Core as a minimal, snap-based OS that enforces strong separation between system components and workloads, with runtime integrity checking and long-term support. Their sessions made clear that confidential computing is evolving from optional feature to foundational capability.

▶ Watch Ijlal Loutfi's talk on demand

Mike Reed (Intel) and Darshan Patel (Fujitsu) illustrated how hardware and software co-design is enabling secure and scalable trusted execution.

Reed showcased TDX Connect, a framework for direct, encrypted communication between VMs and PCIe devices like GPUs, removing the performance penalties of bounce buffers. Patel presented Fujitsu's MONAKA processor, which leverages Arm CCA to isolate workloads into cryptographically protected "realms," enforced by the CPU itself. Both emphasized that hardware innovation is no longer just about performance—it's about embedding attestation, isolation, and portability across heterogeneous environments.

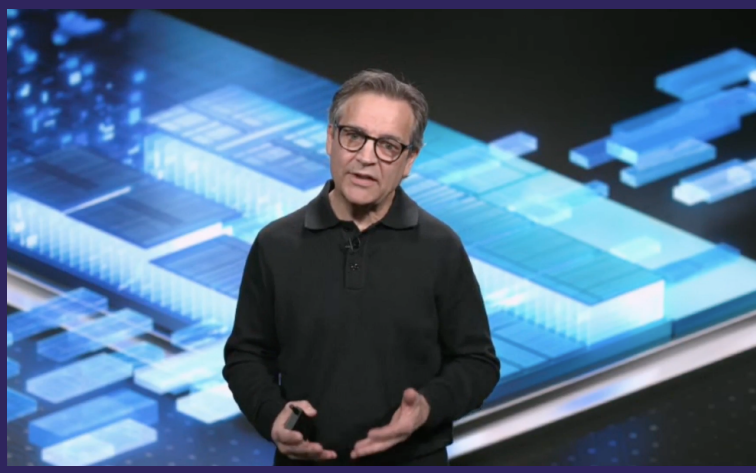

▶ **Watch Mike Reed's talk on demand**

▶ **Watch Darshan Patel & Tatsuya Kitamura's talk on demand**

Mike Bursell (Confidential Computing Consortium) reframed attestation as the foundational enabler of secure multi-party systems.

Bursell argued that attestation is not just a technical requirement, but the critical trust primitive for collaborative AI, data sharing, and regulatory compliance. He categorized emerging use cases by levels of trust and complexity—showing how attestation enables organizations to work together without revealing sensitive data or ceding control.

▶ **Watch the Mike Bursell's talk on demand**

AI

AI

Security and privacy is fundamental for AI future.”

Nelly Porter, Director of Product Management, GCP Confidential Computing, Google

The AI track explored how confidential computing is addressing the unique security demands of modern AI systems. From protecting model weights and prompts to enforcing runtime integrity at scale, the sessions revealed how trust boundaries are being redefined for AI development, deployment, and operation.

Wenhui Zhang (ByteDance) presented a comprehensive framework for securing model deployment across the AI lifecycle.

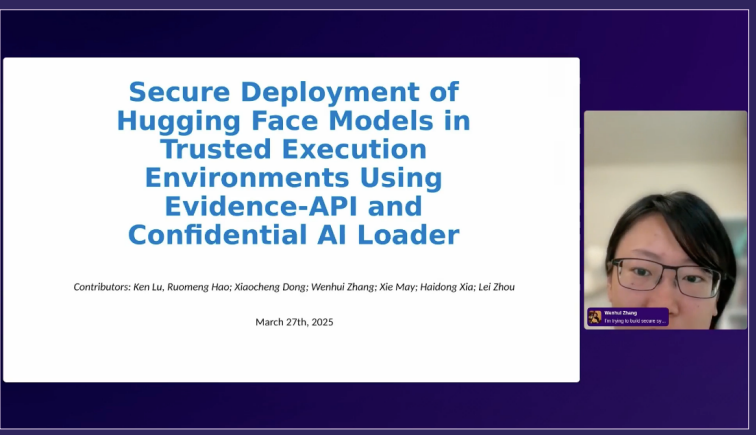

This solution ensures confidentiality not just during inference, but also for models at rest and in transit. It combines an encrypted model loader, a broker-based key release mechanism, and full-stack attestation—enabling secure deployment of LLMs across TEEs in VMs, containers, and clusters. The open-source approach promotes hardware-agnostic adoption and invites community contributions.

Watch the Wenhui Zhang’s talk on demand

Felix Schuster (Edgeless Systems) introduced Privatemode, a confidential GenAI platform with full end-to-end protection.

Built with confidential containers and running on NVIDIA H100 GPUs, **Privatemode** encrypts all prompts and responses, making data invisible even to infrastructure operators. Its architecture combines a secure worker node (VLLM + Llama), a coordinator for attestation, and a key service—all verifiable by the user. Already in use in public-sector deployments, it aims to offer GPT-4-class performance with enterprise-grade assurances.

▶ **Watch the Felix Schuster's talk on demand**

Nelly Porter (Google) and Anand Pashupathy (Intel) emphasized that trust boundaries must evolve alongside AI complexity.

They showcased how Google Cloud's confidential AI stack uses Intel TDX, AMX, and the Tiber Trust Authority to isolate AI workloads and deliver up to 7× performance gains over previous hardware generations. As AI systems scale, cryptographic trust in execution environments is becoming as essential as model performance.

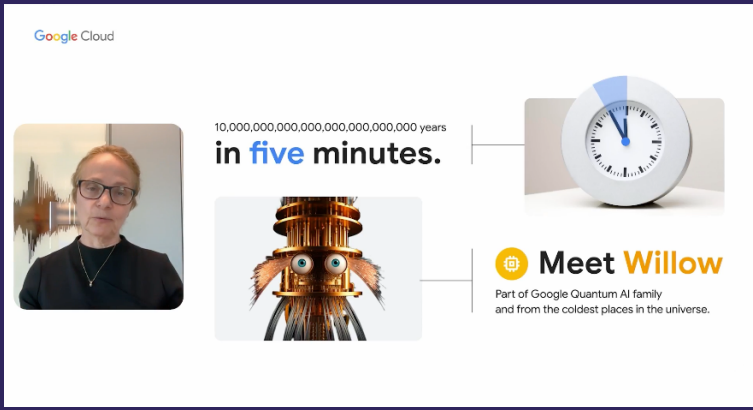

▶ **Watch the Nelly Porter & Anand Pashupathy's talk on demand**

Apps & Solutions

Apps & Solutions

The Apps & Solutions track demonstrated how **confidential computing is moving beyond pilots and into real-world deployments.** From battlefield communications and national healthcare systems to financial fraud detection, speakers shared how trusted execution technologies are solving high-stakes problems across sectors.

> **What I'm starting to see is that this technology is becoming mature to create security and privacy at scale."**
>
> Johan Bryssinck, Head of AI strategy, innovation & partnerships, SWIFT

Luc Gallay and Lucia Jeschonneck (Airbus) showed how confidential computing can support secure communications and edge AI in contested environments.

In its combat cloud proof-of-concept, Airbus migrated critical workloads from standard Kubernetes to confidential containers to protect against insider threats, network compromise, and physical capture. By combining secure mesh networking, confidential ingress, and hardware-enforced isolation, the system enables armored vehicles to act as trusted compute nodes for surveillance, communications, and edge inference—even in the field. The project reflects how container-based confidential computing can be deployed under extreme operational constraints.

▶ **Watch the Luc Gallay & Lucia Jeschonneck's talk on demand**

Rachel Wan and Ram Pai (IBM) shared how confidential computing is scaling to national-level healthcare systems.

In partnership with **Edgeless Systems** and Intel, IBM built a platform for securely processing millions of patient records using SGX-enforced enclaves managed by MarbleRun on OpenShift. The architecture enforces operator exclusion, end-to-end attestation, and encrypted data-in-use processing. Now extended to support TDX VMs for cloud-native deployment, the system serves as a blueprint not only for healthcare, but for any workload requiring both high assurance and regulatory alignment.

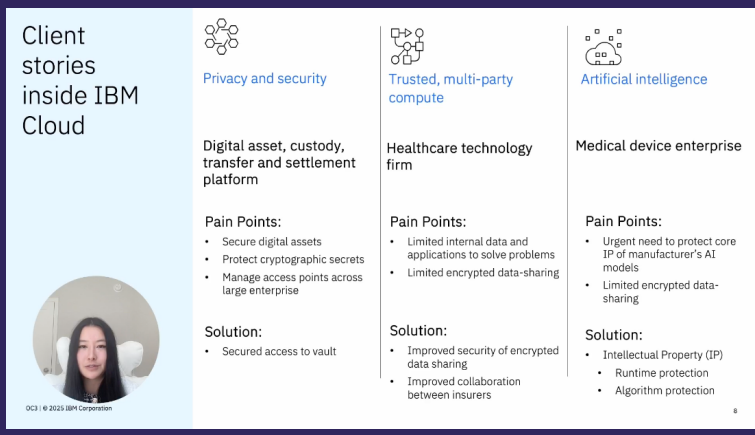

Watch the Rachel Wan & Ram Pai's talk on demand

Johan Bryssinck (Swift) demonstrated how confidential computing enables trust in cross-bank fraud detection.

Swift's industry pilot tackled one of the most challenging problems in finance: how to collaborate securely across organizational boundaries. Through a combination of federated learning and a shared secure environment, banks are able to detect suspicious accounts and share intelligence—without exposing sensitive internal data. Confidential computing plays a key role in enabling shared analytics and provenance guarantees while maintaining legal and regulatory boundaries.

Watch the Johan Bryssinck's talk on demand

Attestation

OC3OC
OC3OC3

Attestation

You are making sure that before you start your trust, that you have done your due diligence, and that is what attestation is today.”

Rob Nertney, Senior Software Architect, NVIDIA

The attestation tack showcased how this once highly specialized process is becoming more accessible, flexible, and relevant to real-world deployments. From enabling verification directly in web browsers to embedding attestation into GPUs and encrypted protocols, speakers demonstrated the expanding reach of attestation across platforms and trust boundaries.

Kosei Akama showed how attestation can be extended to web users—without plugins or proprietary software.

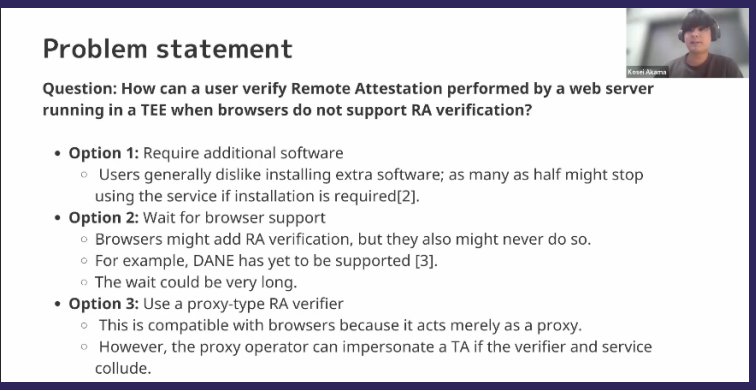

RA-WEBs leverages Certificate Transparency logs and existing browser infrastructure to enable verifiable attestations in typical web sessions. Akama’s approach challenges the assumption that attestation is only feasible in backend systems, demonstrating how open standards and minimal user interaction can bring stronger trust guarantees to browser-based applications.

▶ **Watch the Kosei Akama’s talk on demand**

Paul Howard (Arm) and Kevin Zhao (Linaro) highlighted the growing maturity and modularity of attestation tooling.

Their Docker-based learning platform makes attestation accessible for developers without special hardware, while production-ready integrations in Confidential Containers now support both Veraison and Rust-CA-Token. Their work emphasizes that simplifying attestation workflows and aligning with standards like IETF RATS is key to adoption—especially in complex, multi-vendor supply chains.

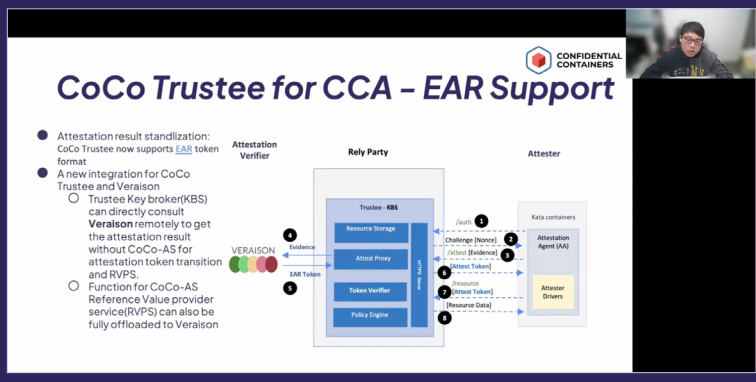

▶ **Watch the Paul Howard & Kevin Zhao's talk on demand**

Rob Nertney (NVIDIA), Dionna Glaze (Google), and Ivan Petrov and Katsiaryna Naliuka (Google DeepMind) illustrated how attestation is being woven into hardware, protocols, and supply chains.

Nertney presented NVIDIA's attestation models for GPUs, supporting both cloud and air-gapped environments. Glaze introduced CoRIM as a structured format for verifying trusted components across multi-vendor systems.

▶ **Watch the Rob Nertney's talk on demand**

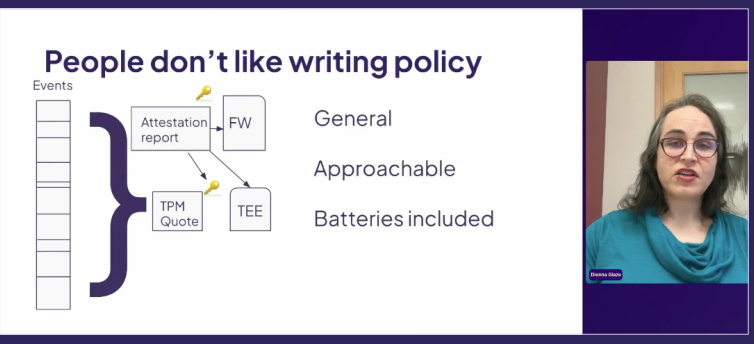

▶ **Watch the Dionna Glaze's talk on demand**

DeepMind researchers extended this thinking to encrypted communication, binding attestation to Noise protocol handshakes for lightweight but verifiable channels in constrained environments.

▶ **Watch the Ivan Petrov & Katsiaryna Naliuka's talk on demand**

Sponsors

Google intel Microsoft Azure

AIRBUS arm CONFIDENTIAL COMPUTING CONSORTIUM

Media Partners

ALL CONFERENCE ALERT
ACADEMIC CONFERENCES AT A GLANCE

GT:>_
GermanTechJobs.de

AIPressRoom

SICHERHEIT.
DAS FACHMAGAZIN

STARTUPS
MAGAZINE

www. it-sicherheit.de
Der Marktplatz IT-Sicherheit

AINEWS

INNOVATION
NEWSNETWORK

IT WELT.at

IMPACT
AI NEWS

IMPACT
newswire

INTERNATIONAL
CONFERENCE
ALERTS

DX
DX TALKS

About Edgeless Systems

Edgeless Systems is a German cybersecurity company that sets new standards in cloud and AI security. With its leading software for confidential computing, Edgeless Systems makes the cloud the safest place for sensitive data and offers the first end-to-end encrypted AI service, Privatemode AI. Edgeless Systems cooperates with leading organizations such as Airbus, Nvidia, IBM, and the Schwarz Group. The company also hosts the OC3, the world's leading specialist conference for confidential computing, with high-profile speakers like the CTOs of Microsoft Azure, Intel, and AMD.

OC3

hosted by Edgeless Systems