

CONFIDENTIAL hosted by
COMPUTING
SUMMIT™ 2025 **OPAQUE**



Building the Trust Layer For The Agentic Era

[OPAQUE.CO](https://opaque.co)

CC SUMMIT 2025



Confidential Computing Summit 2025: By the Numbers



145
Speakers



589
Participants



116
Sessions
& Posters



36%

C-Level / Executive /
Founder



Four Critical Transitions for the Agentic Era



02 Business



From compliance
checkbox to
competitive advantage



04 Ecosystem



From vendor
solutions to industry-
wide trust fabric

01 Technical



From securing data to
securing autonomous
agents and systems



03 Operational



From policy on paper
to policy enforced in
real time





Table of Contents



01 EXECUTIVE SUMMARY



VIEW

06 THE RISE



VIEW

02 THE STAKES



VIEW

07 REALITY CHECK



VIEW

03 AGENTS ARE HERE



VIEW

08 SECTOR INTELLIGENCE



VIEW

04 BUILDING THE FOUNDATION



VIEW

09 ENTERPRISE READINESS



VIEW

05 SCALING SAFETY



VIEW

10 WHAT'S NEXT



VIEW

OPAQUE

01

EXECUTIVE SUMMARY

The Agentic Inflection Point





At the 2025 Confidential Computing Summit in San Francisco, a global community of researchers, builders, and business leaders convened around a shared goal: shaping the trust layer for a rapidly evolving AI ecosystem. The momentum behind confidential AI was unmistakable. No longer a niche concern, it's now central to how enterprises are building for scale, safety, and speed.

Across keynotes and conversations, one thing was clear: we're entering a new phase of AI defined not just by models, but by interconnected, autonomous agents operating across systems, geographies, and institutions. The foundations of this global agentic web are being laid now—and confidential AI is at the heart of it, enabling the verifiable control, privacy, and policy enforcement these systems will require.

This moment marks more than technical progress. It signals a broader transformation in how we design, govern, and trust AI.

Yet, the biggest risk and opportunity for business is not cognitive. It's operational. These agents can already take human-like actions, make multi-step decisions, and adapt to new inputs at a pace no workforce can match. They automate, scale, and connect across APIs, contracts, org charts, and systems—unbounded by biology, shift schedules, or human alignment. That's what makes them both essential and, potentially, dangerous.

Our infrastructure, policies, and rate-limits were built for the “human internet,” where even the best-organized group could do only so much damage in a year. Now, an untrusted agent can do it in an hour—and doesn't even need to “think” to cause outsized impact.

This is why we need verifiable trust as a foundational layer—proof that an agent stayed within policy, respected data boundaries, and didn't go rogue. Without it, we're flying blind at machine speed.



These agents behave with human-like capabilities, but they operate at machine speed. Policy without proof is not trust—it's hope.”

— **AARON FULKERSON, CEO, OPAQUE**



[Learn more about turning policy into verifiable trust at scale](#)



We came together not for incremental progress, but to confront a fundamental challenge: Can we lay down the infrastructure for verifiable trust before our agentic systems outpace our ability to secure them?

This report maps out the current trajectory of AI, synthesizing exclusive insights from 70+ Summit speakers that explore the rise of agentic AI, why AI security has become so important in this new technological era, and where the future of AI security is headed.

It dives into how enterprises ranging from Big Tech to startups are embedding confidential AI into their tech stacks, followed by practical tips on how enterprises can begin their confidential AI journey.

Continue on to learn more.



KEY DATA POINTS:

10,000+

confidential VMs now deployed on Microsoft Azure, protecting live critical workloads

60 million

CrewAI: 60 million AI agents processed each month; usage doubling monthly; 94% efficiency gains

13%

Only 13% of enterprises achieving real ROI from AI, per Accenture



OPAQUE

02

THE STAKES

Hype Meets Hard Numbers



Every executive in the room knew the stakes:



AI adoption is surging, but trust and proof are lagging.



Usage is exponential (PwC deployed CrewAI's agents across its operations to automate code generation, for example), yet only a fraction of companies are seeing tangible value.



Microsoft and AMD both emphasized production-scale deployments—this is no longer an experiment.

NELLY PORTER, DIRECTOR OF PRODUCT
MANAGEMENT FOR GCP CONFIDENTIAL
COMPUTING AND ENCRYPTION AT
GOOGLE

put it succinctly:



Don't wait too long
because it's taking
so much time for
us to migrate in
cryptography alone.”

The window to act is shrinking,
even as the risks and rewards grow.



OPAQUE

03

AGENTS ARE HERE.

Trust Isn't.



We entered the Summit with the reality that agents have already become our new virtual collaborators. They're already onboarding, automating, and performing human-like workflows—often at a scope and pace that's impossible for people to match. At Anthropic, 70% of code is now being written by their agent, Claude. By year's end, that number is expected to hit 95%.

Agents aren't hypothetical—they're here, they're networked, and their actions ripple across APIs, contracts, databases, and organizations.


But there's a hard truth emerging behind this acceleration. Despite headlines about larger context windows and "longer memory," the field is coming to terms with what power users already feel: more memory doesn't mean more thinking.

When you push today's models past a certain threshold of complexity or compositional reasoning, the illusion falls apart. LLMs don't simply get less accurate—they actually start to "think" less, even if there's plenty of context tokens to go around. This ceiling isn't a temporary bug, it's a fundamental limitation of the architecture.



MYTH BUSTED:

Longer memory $\neq$ deeper thinking. AI's ability to "remember" more isn't the same as developing genuine reasoning. When tasks get truly complex or require multiple conceptual steps, current models plateau or regress, even though they may be perfectly able to summarize, recall, or mimic for simpler queries.



So while the conversation often centers on “when does AI start thinking,” what matters for business—and society—is that agents don’t need cognition to have a seismic impact. Their composability, speed, and ability to trigger downstream consequences are what will drive both opportunity and risk.

And that’s where things get urgent. The core problem isn’t only whether an agent is “aligned,” but whether we can even prove what it did, who it acted as, or whether it respected policy boundaries. As **Jason Clinton, CISO at Anthropic, put it:**

“We don’t really have as an industry a good answer for identity. It is maybe the most challenging problem that we’re going to face as practitioners in the coming years.”


— JASON CLINTON, CISO, ANTHROPIC



For regulated industries, this is a dealbreaker. Guaranteeing identity, action trails, and code transparency isn’t optional—it’s the new baseline for trust. **Graham Mudd, Senior Vice President of Product Management at Mozilla and Founder of Anonym (now part of Mozilla)**, summed it up:

“We can provide guarantees through an audit trail and code transparency and attestation. If you’re in a regulated industry, this provides that comfort that you’re doing the right thing and you can back it up and prove it.”

— GRAHAM MUDD, SENIOR VICE PRESIDENT,
PRODUCT MANAGEMENT, MOZILLA



OPAQUE

04

BUILDING THE FOUNDATION

For Verifiable Trust



The focus of this year's Summit is clear: trust cannot be assumed. It must be implemented, measured, and continuously proven.

Companies were excited to share how they're deploying confidential AI across their tech stacks.

- Microsoft completed its global rollout of confidential VMs, securing both consumer token signing and payment systems in 70+ regions.
- AMD is operating Trusted I/O at global scale, with attestation pipelined all the way from hardware to software.
- OPAQUE launched the Confidential Agent Stack that embeds behavioral attestation, dynamic identity issuance, and real-time policy enforcement across agentic workflows.

The question is no longer if verifiable trust is required, but how quickly organizations can operationalize it—before the next breach or before regulatory pressure forces reactive change.



KEY DATA POINTS:

100,000+

confidential VMs deployed

94%

efficiency gains from agentic workflows

13%

of companies seeing ROI



OPAQUE

05

SCALING SAFETY

Trust Can't Be Proprietary





If there was one message that echoed throughout the summit, it was this: no single company, no matter how advanced, can build trust on its own. The agentic era demands open standards, real interoperability, and genuine industry collaboration.

Take it from **Daniel Rohrer**, VP of Software Security for NVIDIA, who said that:



The ecosystem needs to help us build out that trust, adding that collaboration is required to scale trust, not just compute.”

— **DANIEL ROHRER, VP, SOFTWARE SECURITY, NVIDIA**

For us, this means moving beyond proprietary, bolt-on solutions. The future is about verification you can prove—not just assert—across clouds, vendors, and global workflows.

Leonardo Garcia, Principal Engineer at Linaro, takes trust a step further, saying that open source is the only way to verify and attest workloads haven’t been tampered with:



We believe that as an industry, if we agree on a very few set of standards, we can make the transition between environments easier.”

— **LEONARDO GARCIA, PRINCIPAL ENGINEER, LINARO**

Mike Bursell, Executive Director of the Confidential Computing Consortium, explained the real business value of having multiple stakeholders involved in the process:



We can establish a mechanism by which mutually distrusting parties can trust each other and thereby create value.”


— **MIKE BURSELL, EXECUTIVE DIRECTOR, CONFIDENTIAL COMPUTING CONSORTIUM**

OPAQUE

06

THE RISE

of the Confidential Model
Control Plane



What's emerging now isn't a collection of disconnected tools. It's a cross-industry control plane for confidential AI. The summit's working groups are aligning around three core pillars:

- **Shared APIs** for attestation and reporting
- **Portable identity** that moves securely with agents and workloads
- **Policy enforcement** that can flex with global AI law and standards

In practice, this shift means we're finally poised to deliver "verifiable-by-design" systems—not just for ourselves, but across the whole ecosystem. The story is moving from isolated innovation to shared infrastructure. That's how we raise the bar for everyone.



THE CONFIDENTIAL MODEL CONTROL PLANE

The Confidential Model Control Plane serves as the central management layer for AI—functioning as the “brain” that governs model behavior. It orchestrates workflows, enforces policies, and manages resources across the entire AI lifecycle, from training to deployment.

By establishing a common trust fabric, it ensures AI agents and workloads can move securely across clouds, vendors, and borders. This shared foundation enables interoperability, resilience, and compliance, making it possible for organizations to confidently deploy AI in multi-cloud, cross-enterprise environments.



At OPAQUE, our confidential agent stack is already built for this future: fully interoperable, standards-forward, and ready for a world where agentic systems cross enterprise and industry boundaries.



OPAQUE

07

REALITY CHECK

GenAI Must Work
in the Real World





As the excitement and speculation around GenAI reach a fever pitch, we kept returning to one grounding truth throughout the summit: proof, not promises, separates hype from real progress. This agentic era will be defined not by what's possible in a demo, but by what delivers in production.

Teresa Tung, Senior Managing Director at Accenture, drove this home with some surprising numbers from a recent survey.

While 83% of executives express enthusiasm about GenAI's potential, only 13% of companies are seeing actual value from their deployments.

Basic applications like Microsoft's Copilot are useful for boosting efficiency in mundane tasks like summarizing calls and scheduling. But they aren't moving the needle in terms of substantial ROI.

To achieve meaningful value, companies must shift from these "table stakes" solutions to making strategic bets on GenAI initiatives that directly target their core business functions.



As Tung put it, companies that are “reinvention-ready” have “invested in their data foundation,” where they’re likely to have “good governance and data product practices.”

— **TERESA TUNG, SENIOR MANAGING DIRECTOR, ACCENTURE**



James Kaplan, a Partner at McKinsey and CTO of McKinsey Technology, shares Tung's optimism for GenAI's business potential.

With GenAI tools, enterprises can extract insights from their untapped proprietary data, remediate technical debt, and free up IT resources once spent on system maintenance for innovation.

Kaplan frames it this way:



For the first time in decades, we may see a secular improvement in engineering or application development productivity as a result of generative AI. This is transformative.”

— **JAMES KAPLAN, PARTNER, MCKINSEY AND CTO, MCKINSEY TECHNOLOGY**



What does this mean on the ground? It means organizations are moving beyond experimentation and into scale, but only when trust and verification are integrated at every step.

Proof in production

We heard from leaders who shared exactly how confidential computing is powering mission-critical workloads today:



Mark Russinovich, CTO, Deputy CISO and Technical Fellow of Microsoft Azure, talked about Microsoft's deployment of over 100,000 confidential virtual machines (VMs) to curb cyber threats.

Under Microsoft's Secure Future Initiative, the program is securing payment systems that cover over one billion credit card transactions a year and protecting 60 billion global licensing transactions per month.

Ravi Kuppuswamy, senior vice president of the Server Solutions Group at AMD, discussed AMD's Trusted I/O, a foundational technology designed to expand the trusted boundary of confidential computing beyond just the CPU to any of the devices attached to it.

Trusted I/O is now rolling out to new generations of processors—crucial for modern AI workloads that heavily rely on accelerators like GPUs.



João Moura, Chief Executive Officer of CrewAI, has led efforts to deploy 60 million CrewAI agents a month, a figure that is “almost doubling” every month.

Fortune 500 and Global 2000 companies use CrewAI's agents to generate code, analyze contracts, and automate back-office operations. One consumer packaged goods achieved 94% efficiency gains by automating coupon and discount approvals.

To deploy agents at scale, clients emphasize the need for secure environments, data protection, and verifiable trust.



Explore how leading organizations
are scaling agent deployments while
preparing for post-quantum security



The lesson is clear: when trust is built in from the start, infrastructure and business value scale together. Hype fades, but proof endures.



KEY DATA POINTS:

100,000+

Real-World Proof: 100,000+ confidential VMs
in deployment

1B+

credit card transactions protected annually

94%

efficiency gains from production agentic workflows



OPAQUE

08

SECTOR INTELLIGENCE

Where Confidential AI Is
Delivering Business Value



Confidential AI isn't just a future concept anymore—it's already making a measurable impact across high-stakes sectors. At the summit, leaders shared live deployments and real ROI. Here's what's working right now.



FINANCIAL SERVICES

Banks today face unrelenting cyber threats and high regulatory pressure, making AI security a key priority.

- **Microsoft** now protects over one billion credit card transactions every year on confidential VMs, preventing exposure even for the most sensitive data.
- **Swift's multi-bank collaborations** use confidential AI to spot money laundering as it happens—in real time and across institutions, all while maintaining strict privacy boundaries. That way, banks can collaborate and derive insights from sensitive data without directly exposing their proprietary data to each other.
- **Money to GO** uses confidential AI to create digital fingerprints of financial data, making its contents unknown to the company or cloud provider.



HEALTHCARE

Healthcare is about collaboration, privacy, and speed. In a space where patient data is vast and sensitive, confidential AI is proving to be essential in unlocking the full potential of AI, enabling secure data processing and ensuring privacy compliance for medical data.

- **AstraZeneca** uses Google Cloud to securely determine the right medication for patients, processing large, anonymized data sets through confidential computing.
- **Beekeeper AI**, in partnership with AMD, offers a solution that removes barriers to using healthcare data. The solution ensures patient data and AI models remain encrypted at all times, operate within secure environments, and generate results only published to authorized researchers and model owners.
- **AlGnomix is using confidential genomics workloads** to help researchers detect emerging pathogens faster, using runtime attestation to keep patient data secure.
- **Karma Health** is using confidential AI to power the fight against sepsis—a \$62 billion challenge affecting 1.7 million lives in the U.S. each year.
- **Multi-institution cancer research projects** are pooling data without risking leaks.



ADVERTISING

In digital advertising, the privacy arms race is on. **Anonym's** confidential ad targeting enables cross-DSP model training—brands and platforms can collaborate on high-performing campaigns without leaking audience data. Technologies like Oblivious HTTP, differential privacy, and ephemeral TEEs keep both identity and performance metrics protected.





ENTERPRISE OPERATIONS

AI agents are beginning to automate core enterprise workflows at scale. OPAQUE and ServiceNow demoed confidential agents for everything from contract review to customer support, all with full attestation trails.

PwC is automating SAP and Salesforce workflows with improved accuracy and stronger privacy guarantees, while CrewAI's back-office automations are setting new standards for both efficiency and security.

Across sectors, one consistent takeaway emerged: businesses are now grappling with what **Michael Reed**, Intel's Senior Director for Confidential Computing, called a "diverse data problem." As Reed put it at the summit:



Many organizations today struggle with what we call a diverse data problem. They have multiple data sets with different sensitivities that include varying ownership, restrictions on viewership, and even regulatory controls."

— MIKE REED, SENIOR DIRECTOR FOR
CONFIDENTIAL COMPUTING, INTEL



MANUFACTURING & SUPPLY CHAIN

Confidential computing's appeal isn't just technical. It's about making collaboration possible without sacrificing regulatory compliance or ownership boundaries. This is fuelling growth in areas like joint healthcare research, multi-bank financial analytics, and privacy-first advertising—where the business value depends on securely handling diverse, sensitive information.

The supply chain is only as strong as its weakest link—which is often security. AMD's Trusted I/O now underpins silicon provenance tracking, ensuring every component's integrity from fab to finished product. Secure cross-vendor compute (NVIDIA and AMD together) is finally making it possible to build trusted boundaries for manufacturing at scale.

Gelato is using confidential agents to optimize logistics across Europe while protecting sensitive operational data.

**FINANCIAL
SERVICES**

1B +transactions protected

HEALTHCARE

multi-institutional
research secured

**MANUFACTURING
& SUPPLY CHAIN**

secure silicon
supply chain

ADVERTISING

cross-platform targeting
with privacy

**ENTERPRISE
OPERATIONS**

attested automation

**What do all these use cases have in common?**

They treat trust not as a compliance checkbox, but as a design principle—and they're seeing both business impact and future readiness as a result.

OPAQUE

09

ENTERPRISE READINESS

How to Put Confidential AI to
Work—For Real



Most companies are asking the same question right now: “How do we actually make confidential AI work in our business, without getting stuck in endless planning or abstract concepts?”

Here’s what came up again and again at the summit—
less checklist, more reality.

Start with security, but embed it everywhere. Waiting to “add security later” is how you fall behind. Companies seeing results are baking agent identity and attestation right into their development pipelines. That means every new feature, every deployment, gets verified as a matter of course—not as a crisis response.

“

Confidential computing is not the option. It's absolutely must capability for everything that we do in our product and platform”

— **NELLY PORTER, GOOGLE’S DIRECTOR
OF PRODUCT MANAGEMENT FOR GCP
CONFIDENTIAL COMPUTING AND ENCRYPTION**



“

Policy without proof
is not trust, it’s hope

— **AARON FULKERSON,
CEO OF OPAQUE**





Upskill at every level, not just the C-Suite. This isn't just a CTO conversation anymore. Companies that excel are training engineers, risk managers, even product leads in the practical side of confidential computing. Hosting simulations and incident drills has become the norm, so if something goes wrong, everyone knows their role—not just the experts.

“ For individual contributors who have never managed before, [we should also be] starting to lean into what the next career progression for them looks like where they're managing AI as a part of their work.”

— JASON CLINTON, CISO, ANTHROPIC



Procurement is proof-based. There's no more “just trust us.” RFPs now require vendors to demonstrate attestability and reproducibility before a contract is signed. Many teams are piloting new solutions in a controlled environment, measuring business impact before investing big. Sourcing is less about slide decks, more about shared evidence.



“

Nikhil Gulati, Global Head of Engineering and AI at Johnson & Johnson, put it plainly: The only way compliance teams are now comfortable engaging with AI initiatives is by being brought in as equal partners. That wasn't always the case. These teams used to approach AI projects with hesitation, raising basic questions like “Do you have rights to the data?” and “How is that data secured?” only after the fact.

— **NIKHIL GULATI, GLOBAL HEAD OF
ENGINEERING AND AI, JOHNSON & JOHNSON**



Now, those same questions are built into the process from the start. Security and data handling proofs are no longer an afterthought, but rather a prerequisite baked into the company's design thinking and planning cycles.

Organizations now see auditability—not just as a technical checkbox—but as essential to building trust internally and with external partners.

This shift means that contracts, vendor assessments, and even compliance reviews increasingly require demonstrable proof of these controls in action.

So, what does it mean when **trust is built into how teams work** each day? The companies leading this space have made trust tangible—woven into daily routines, not left to annual reviews. Is your company ready to take the same steps?

[Assess your organization's confidential AI readiness with our 5-level maturity framework](#)



OPAQUE

10

WHAT'S NEXT

Shared Trust,
Shared Responsibility





The big takeaway by the summit's close was simple but radical.

The question isn't "when will AI think like us," but how fast can we build a trust layer for a world where agents can already act, scale, and trigger real outcomes at machine speed?

AI doesn't need to "understand" to create opportunity or cause harm. The very qualities that make agentic AI so valuable—24/7 composability, interconnection, and autonomy—are what make verifiable proof and policy enforcement absolutely non-negotiable.

We're entering an era where trust isn't about intuition or brand, but proof: policies that are enforced in real-time, verifiable audit trails, and standards that cross org and sector boundaries. Collaboration and transparency are now the true differentiators. The organizations that help build this trust fabric—across vendors, users, regulators, and the open ecosystem—are the ones that will shape the agentic internet for everyone.

The bottom line is this: In the age of AI, traditional methods of security and governance are insufficient for the current and future technological landscape. Inherent, verifiable trust is essential at every level of the system.

As **OPAQUE CEO Fulkerson** put it: "Our future will not be secured with more policies or regulations or legal agreements or retroactive governance that doesn't scale in this new world. It must be built on verifiable proof, baked into the architecture, baked into our systems, baked into the network, enforced at runtime."





OPAQUE

APPENDIX & RESOURCES



Keynote Speakers



Mark Russinovich

CTO, Deputy CISO and Technical Fellow, Microsoft Azure



Jason Clinton

Chief Information Security Officer, Anthropic



Nelly Porter

Director of Product Management, GCP Confidential Computing and Encryption, Google



James Kaplan

Partner, McKinsey and CTO, McKinsey Technology, McKinsey



Ravi Kuppuswamy

Senior Vice President, Server Solutions Group Server Business Unit, AMD



Ion Stoica

Co-Founder and Executive Chairman, Anyscale and Databricks | Professor, UC Berkeley | Co-Founder Board Member, OPAQUE



Aaron Fulkerson

Chief Executive Officer, OPAQUE



Daniel Rohrer

VP of Software Product Security, Architecture and Research, NVIDIA



João Moura

Chief Executive Officer, CrewAI



Teresa Tung

Senior Managing Director, Accenture



Raluca Ada Popa

Computer Security Professor at UC Berkeley, Leading Frontier Security at Google DeepMind, Co-Founder of OPAQUE, PreVeil and IMUA



Sam Lugani

Product Lead, Confidential Computing, Google



Graham Mudd

Senior Vice President, Product Management, Mozilla



Michael Reed

Senior Director, Confidential Computing, Intel



Mike Bursell

Executive Director, Confidential Computing Consortium



Event Sponsors



HOSTED BY

OPAQUE

DIAMOND



GOLD



SUPPORTING SYSTEMS



PLATINUM



SILVER




Interested in sponsoring the Confidential Computing Summit?

Contact sponsorships@opaque.co to learn more about sponsorship opportunities!



Watch Summit highlights



Access the full collection of keynotes, breakout sessions, and poster presentations from the leading event on trusted AI and confidential computing.

[EXPLORE ON-DEMAND SESSIONS](#)

CONFIDENTIAL hosted by
COMPUTING
SUMMIT 2025 **OPAQUE**



CONFIDENTIAL COMPUTING SUMMIT 2025 hosted by OPAQUE



Thank you

Interested in learning more about the latest trends and advancements in confidential AI?

Check out [OPAQUE's Confidential AI platform](#) and subscribe to our [AI Confidential Newsletter](#), in your inbox every two weeks.

AUGUST 2025

