AMD × OPAQUE

FROM RISK TO RESILIENCE: CONFIDENTIAL COMPUTING WITH AMD AND OPAQUE

WHITE PAPER

EXECUTIVE SUMMARY

As the use of artificial intelligence becomes more common in business, the need to protect confidential data also grows. Information travels constantly between on-premises data centers and the cloud. While data is often encrypted when stored or in transit, it must typically be decrypted for use, creating a significant security gap when deployed on shared resources. Confidential computing closes this gap.

What is confidential computing? According to the Confidential Computing Consortium, confidential computing (CC) is the protection of data in use by performing computation in a hardware-based, attested Trusted Execution Environment (TEE). The TEE is a secure area in the main processor that prevents unauthorized access to the data being processed. This technology helps organizations prevent data breaches, enables secure collaboration, and addresses strict regulatory requirements.

The OPAQUE Confidential AI Platform relies on AMD Secure Encrypted Virtualization (SEV) protected environment to secure data in use through hardware-based memory encryption and attestation. Building on this foundation, OPAQUE adds verifiable runtime governance to enforce access controls, data-use policies, and auditability for AI workflows. This enables enterprises to securely use valuable proprietary data to power AI agents and workflows, improve model accuracy, and support compliance.

THE MODERN DATA CHALLENGE IN AI

The lifecycle of an AI workload – from data acquisition to model deployment – involves nearly continuous data manipulation across different environments. Historical data might be moved from mainframes to the cloud for processing, or new data could stream in from remote sensors.

With traditional security methods, encryption protects data when it's stored or in transit. When it's in use, however, it must be decrypted, which introduces security challenges.

Organizations working with sensitive information face three fundamental barriers to AI adoption:

1. THE COMPLEXITY OF CROSS-BORDER DATA COMPLIANCE

Managing data across multiple systems is challenging, and data sovereignty laws add complexity by requiring certain types of data (e.g., personal, financial, medical) to comply with local regulations. These laws can hinder global collaboration and make it harder for teams to share information or leverage AI for fast, informed decisions.

2. THE PRIVACY-UTILITY TRADEOFF

AI performs best with rich, detailed data, but privacy techniques such as anonymization or masking often degrade data quality, reducing AI accuracy. Organizations face a tradeoff between protecting privacy and optimizing AI performance, while regulations such as GDPR and CCPA further restrict access to critical datasets for training.

3. GOVERNANCE GAPS THAT LEAD TO AUDITABILITY GAPS

Data used for AI workloads is often scattered across disconnected systems within an organization. As data from these silos is selected for AI use, a lack of visibility into the data and poor auditability within AI workflows can create significant risks. Without clear governance frameworks, organizations face challenges in tracking data lineage, monitoring policy compliance, and ensuring accountability. These gaps hinder the scalability of AI initiatives while undermining their reliability and ethical integrity.

SECURING DATA WITH TRUSTED EXECUTION ENVIRONMENTS

Confidential computing mitigates data risks by performing computations in a hardware-based TEE, helping ensure the integrity of code execution and data protection within the environment. Before sharing data for processing in a TEE, an attestation process verifies that the hardware and software environment is genuine, secure, and compliant with required standards.

By leveraging TEEs, organizations can securely process fragmented data across silos, without exposing it to unauthorized access, and help ensure compliance with data sovereignty laws such as GDPR and HIPAA. TEEs balance privacy and utility by enabling encrypted data analysis for AI workloads, preserving confidentiality while extracting value.

TEEs can also enhance governance and auditability by tracking data lineage, enforcing policies, and providing accountability through secure logging.

AMD SECURE ENCRYPTED VIRTUALIZATION (SEV) TECHNOLOGY

AMD Secure Encrypted Virtualization (SEV) technology, featured in the AMD EPYC™ Series of data center CPUs, is at the core of modern confidential computing. AMD SEV is a powerful feature for virtual machines (VMs) that protects data while it is being processed. Supported by major cloud providers, hardware manufacturers, and Linux® distributors, it helps organizations easily adopt confidential computing using existing VM images, tooling, and operations, typically with no application changes required. With each EPYC generation, AMD has enhanced SEV with stronger protection and scalability. The dedicated AMD Secure Processor (ASP), integrated in each EPYC CPU, manages encryption keys — never exposing them to the hypervisor — while the memory controller executes high-performance AES cryptography engines (now 256-bit AES-XTS). With current generation EPYC CPUs, SEV can protect up to 1000+ VMs with unique keys per VM and extends encryption to CXL™-attached memory, **an industry first**.

AMD SEV technology has evolved to close security gaps and provide comprehensive protection. It includes:

ENCRYPTED STATE (ES): Adds encryption to CPU register states, preventing the hypervisor or any component with similar privileges from viewing or altering the data actively processed by the CPU.

SECURE NESTED PAGING (SNP): Adds memory integrity protection, preventing the hypervisor from tampering with the VM's memory. It also deters firmware rollback attacks, helping ensure the environment remains secure and unmodified.

TRUSTED I/O (TIO): Extends the protection domain by protecting data moving between the VM and PCIe®-attached devices that support TDISP (TEE Device Interface Security Protocol), such as disks, network cards, or GPUs. TIO helps ensure that I/O data is encrypted and that only trusted devices can access the VM's memory.

TRANSPARENT SECURE MEMORY ENCRYPTION (TSME): Traditionally, data in RAM has been vulnerable to attacks. Transparent Secure Memory Encryption (TSME) was developed to provide system-wide memory encryption using a single AES-128 encryption key, managed by the processor.[1]

With these layers of protection, AMD SEV allows organizations to unlock value from sensitive data, enable secure collaboration, and run confidential AI workloads. By enabling confidentiality and isolation, even from privileged administrators or compromised components like hypervisors, AMD SEV enforces least privilege and minimizes system trust, aligning with zero-trust principles. It works with existing applications without requiring code changes, offering organizations in transition a path to a zero-trust security model.

A broad ecosystem supports AMD SEV, from CSPs and OEMs to hypervisors and orchestration tools, backed by open standards and trusted by major corporations. AMD has become a foundational technology for confidential computing.

OPAQUE SYSTEMS: ENABLING CONFIDENTIAL AI

While hardware provides a secure foundation, organizations need a flexible platform to manage and orchestrate these capabilities for complex AI workflows. **OPAQUE Systems** unblocks AI value by making trust verifiable, enabling enterprises to use their most sensitive data to power AI innovation while adhering to security, compliance, and privacy guidelines.

OPAQUE's platform delivers verifiable runtime policy enforcement, backed by cryptographic audit logs. It leverages transparent confidential computing and TEEs, enabled by AMD SEV, that protect data throughout processing.

The platform provides verifiable governance via cryptographic evidence that data and model weights are protected, and policies are enforced throughout workflows, before, during, and after runtime. By performing remote attestation, which cryptographically verifies that a TEE is running on genuine hardware and executing expected code, organizations can analyze and collaborate on encrypted data helping reduce exposure to cloud providers, system administrators, or unauthorized users.

Providing verifiable privacy through hardware-based remote attestation, cryptographic verification, and audit trails empowers organizations to unlock sensitive data. By integrating with AMD SEV infrastructure, OPAQUE extends hardware-level protections such as isolated execution and runtime encryption. OPAQUE customers may see faster time-to-production for selected workloads, depending on the use case and environment.

"Confidential AI allows enterprises to unlock the full value of their most sensitive data without ever exposing it. By combining AMD hardware-backed trust with OPAQUE's software-level enforcement, organizations can now run AI workloads securely on encrypted data at cloud scale. Every computation is cryptographically verifiable, every access governed by policy, and every outcome auditable. This turns data privacy from a compliance checkbox into a foundation for innovation."

– Rishabh Poddar, Co-Founder and CTO, OPAQUE Systems

REAL-WORLD DEPLOYMENT: PROTECTING SENSITIVE CONSUMER INFORMATION

For a practical look at confidential AI in action, consider a large U.S. credit management company. As a leading debt buyer, it routinely collaborates with hundreds of debt settlement companies (DSCs), which involves sharing sensitive consumer debt information, including Social Security numbers, often in CSV files.

The old manual workflow for analyzing these files was slow and inefficient. This introduced several business challenges:

- **Data was processed in the clear**, leaving sensitive information largely unencrypted and vulnerable.
- **Sensitive consumer data was decrypted for analysis**, creating risks of fraud, identity theft, and regulatory violations.
- **Traditional anonymization methods were no longer sufficient**. Patterns in AI could be re-identified, even if the data was anonymized.
- **Personal Identifiable Information (PII)** was at risk of exposure to unauthorized individuals or systems.
- **Enforcing data policies was difficult**, making it hard to determine how long data should be retained and to ensure it was properly deleted.

The organization needed a way to securely store, process, and share this confidential information while providing verifiable proof of data deletion. They chose the OPAQUE Confidential AI Platform, running on AMD Confidential VMs in Microsoft Azure, to address these challenges.

THE OPAQUE AND AMD SOLUTION IN ACTION

The implemented solution created an automated and compliant workflow with strong security capabilities that transformed the organization's operations. It is a cloud-based, transparent, confidential computing environment powered by OPAQUE and AMD confidential VMs (CVMs) within the company's Azure environment. Here's an overview of how it works:

1. **ATTESTATION:** Remote attestation is the foundation of OPAQUE's security model. Before any work begins, the OPAQUE platform cryptographically verifies that it has been deployed within a transparent confidential computing environment deployed on genuine AMD SEV confidential VMs. This attestation process helps ensure that the OPAQUE deployment is in a trusted state before it is allowed to accept or process any sensitive data.

2. **EXECUTION:** The data owners (DSCs) can now connect their encrypted data to OPAQUE. The system ingests the encrypted files and processes them via automated workflows on the platform.

   The data remains encrypted throughout its entire lifecycle, protected within AMD SEV confidential VMs with hardware-enforced memory encryption and integrity. Only approved queries and workflows are executed on the data. There is cryptographic proof of every action. Enterprise policies for data access and user controls are verifiably enforced at runtime, preventing any unauthorized access to PII (Personally Identifiable Information).

   Finally, the results of the data processing workflow are output to the company's teams to power their business operations.

3. **AUDITING:** After execution, the OPAQUE platform provides user-friendly, signed, tamper-proof audit trails that can be used to verify whether data used during processing was in compliance with the customer organization's policies. Users get cryptographic proof of the environment that executed the workload. These cryptographically signed reports show exactly how data was processed, which workflows were used, and what policies were applied. The exportable logs provide auditors and regulators with concrete proof of compliance.

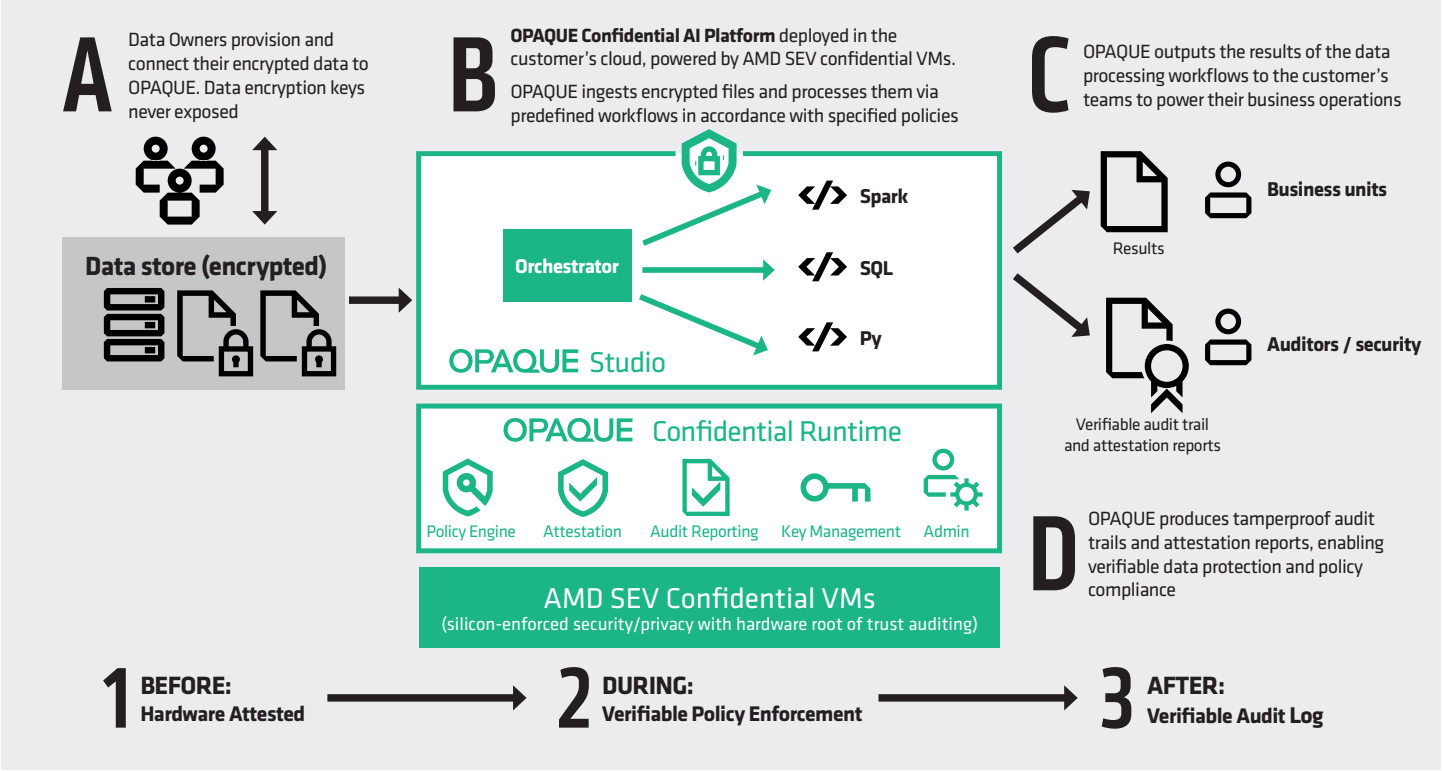

FIGURE 1: CONFIDENTIAL COMPUTE POWERED BY OPAQUE AND AMD

"Confidential AI is about turning sensitive data into advantage. AMD is thrilled to power partners like OPAQUE who help customers do that securely and at scale."

– Madhusudhan Rangarajan, Corporate VP, AMD

BUSINESS IMPACT AND BENEFITS

The joint solution delivered significant improvements:

- **STREAMLINED WORKFLOWS THROUGH AUTOMATION:** Automating manual, multi-step workflows can reduce operational friction and shorten turnaround time. This helps shift staff effort from routine review of raw outputs (e.g., CSV files) toward higher-value analysis and strategic work.
- **ENHANCED SECURITY POSTURE:** By accepting encrypted files from the start and keeping data encrypted even while in use, the process slashes the risk of exposing sensitive information. This end-to-end protection is crucial in a regulated environment.
- **SCALABILITY FOR GROWTH:** The OPAQUE platform is designed to scale horizontally, enabling the company to handle a massive and growing volume of files from hundreds of partners without compromising performance or security.

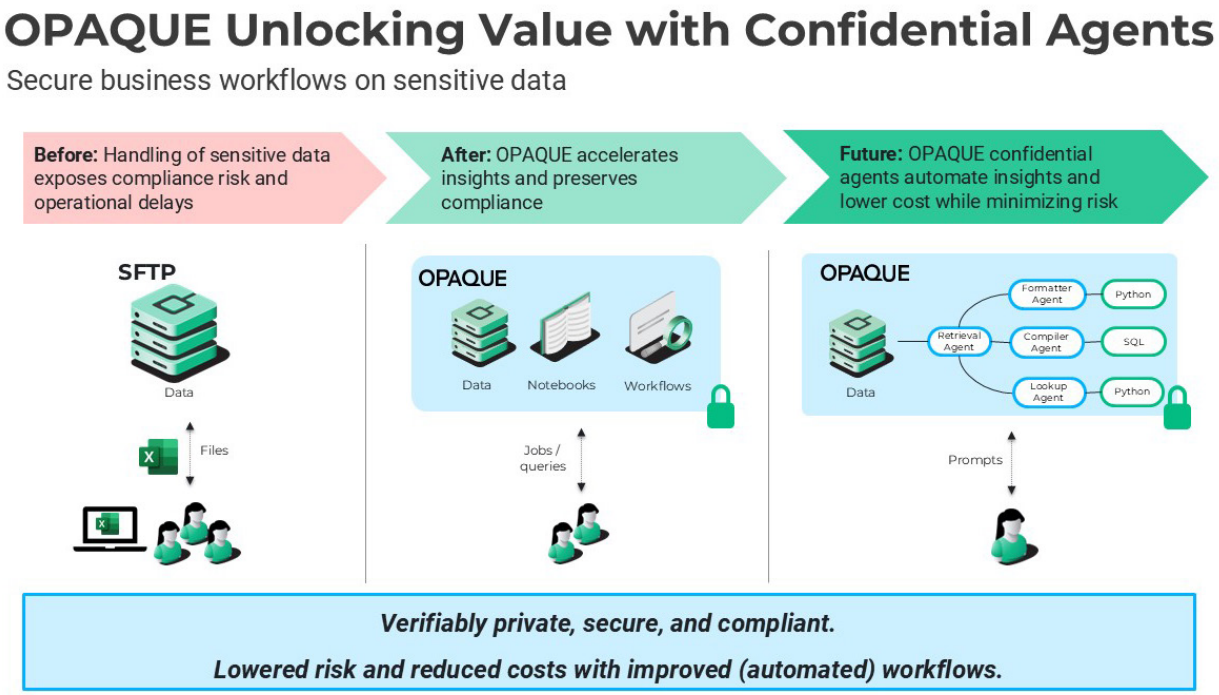

Illustration by OPAQUE

FIGURE 2: OPAQUE CONFIDENTIAL AGENTS

THE FUTURE OF CONFIDENTIAL COMPUTING

This collaboration among a leading U.S. credit management company, OPAQUE, and AMD shows what's possible when innovative hardware and software solutions come together to tackle real-world problems. By leveraging AMD SEV technology within the OPAQUE platform, the company has transformed its data handling processes, implementing robust security capabilities, enhancing efficiency and compliance. As AI continues to evolve, confidential computing will become an essential foundation for any organization looking to innovate responsibly with sensitive data.

Learn more about confidential computing and AMD SEV

1 VPSG, "Evolution of AMD SEV", March 27, 2025, https://www.vpsbg.eu/blog/evolution-of-amd-sev