# OPAQUE

Run sensitive AI workloads with verifiable trust, privacy, and compliance

OPAQUE.co
docs.OPAQUE.co

## The Challenge

Enterprises are blocked from using AI agents and LLMs with their most valuable, proprietary data due to data leakage and compliance risks, even with CEO mandates for productivity gains.

Yet technical leaders face a tradeoff between innovation and assurance:

- **Data exposure risk** inherent in AI workloads, exacerbated when they run in shared infrastructure or use external LLM services.
- **Limited control** over the ability to govern where workloads execute, which models are used, how systems update, and who can access them.
- **Weak auditability** with logs that can't verifiably prove where, when, or how a workload processed sensitive data.

As a result, most AI initiatives stay stuck in pilot mode because security, compliance, and engineering can't align on a verifiable path to production. Enterprises need a way to adopt AI quickly with confidence and control.

## The Solution

OPAQUE solves this by providing a confidential runtime and SDK, enabling enterprises to run sensitive AI workloads and agents on OPAQUE's confidential AI platform with end-to-end verifiable guarantees that your AI agents enforce data policies and keep sensitive data private before, during, and after execution, with full cryptographic auditability.

- **OPAQUE makes AI workflows verifiable** by cryptographically attesting AI pipelines, and ensuring that they run in a secure, verified environment
- **Customers maintain full control** of their data, models, workloads, and keys, and all executed inside their own confidential cloud environments.
- **Real-time policy enforcement** that prevents violations during execution, proving every run produces cryptographic evidence that the right code ran on the right hardware against the right data, under the right policies.

**The result:** Enterprises can plug AI into their most sensitive systems without weakening *security, compliance, or performance*.

## Key Outcomes

**End-to-end data protection**
Process your most sensitive proprietary data in TEEs with complete end-to-end encryption—at rest, in transit, and in use. No plaintext ever leaves your confidential environment.

**Verifiable runtime guarantees**
Generate hardware-backed attestation and cryptographically verifiable evidence showing exactly which models, data, configurations, and policies were in effect for every run.

**Immutable audit trails**
Automatically produce immutable, structured audit logs that record who ran each AI workflow, when, with which configurations, and under which policy, making both internal and external security and compliance reviews fast and easy.

**Flexible deployment**
Deploy the OPAQUE platform within your own cloud environment (Azure, AWS, GCP), while keeping sovereignty over workloads, keys, and data. Scale from department to global production without friction.

## System Architecture

OPAQUE separates orchestration and attestation from data and workloads. The diagram shows how the OPAQUE control plane, the customer's confidential environment, and the cloud provider work together.

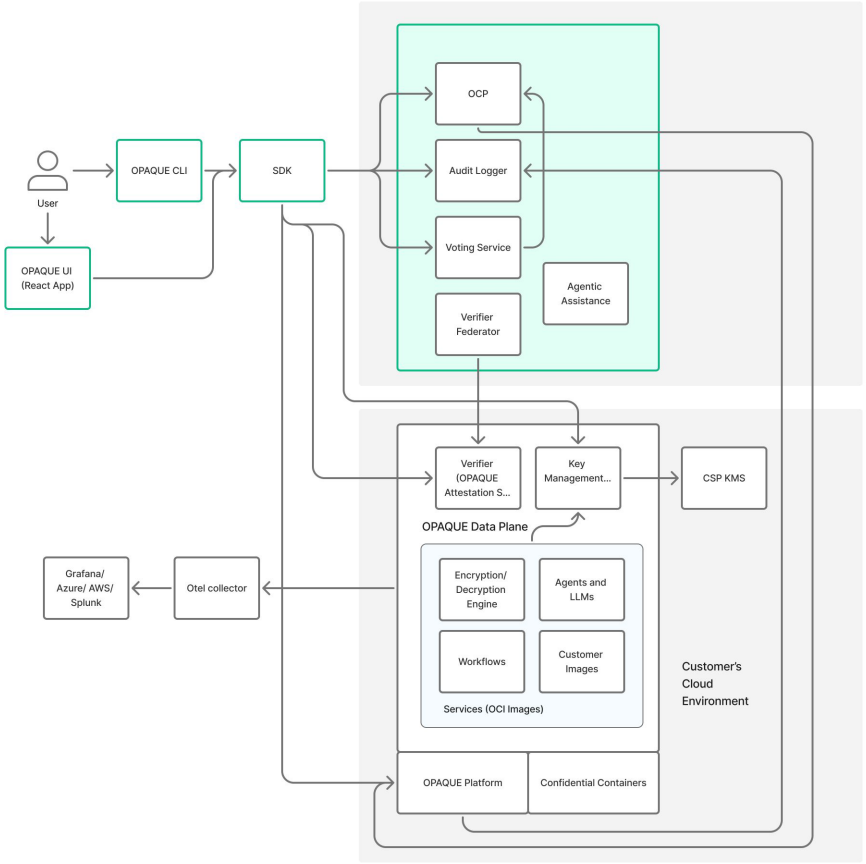


**OPAQUE Management Plane**

- OCP orchestrates workflow lifecycle and metadata.
- Audit Logger records every workflow event for compliance and forensics.
- Verifier Federator and Voting Service coordinate attestation results across verifiers.
- Agentic Assistance provides orchestration and tooling for agentic workflows.

**Customer Confidential Environment (OPAQUE Data Plane)**

- Customer workloads are packaged as signed OCI images.
- A Verifier (OPAQUE Attestation Service) validates that workloads are running in the expected TEE configuration.
- A Key Management System and CSP KMS manage customer keys used for encryption and signing.

**Cloud Services & Observability**
Telemetry flows via Otel collector into Grafana / Azure / AWS / Splunk so platform and security teams can reuse existing monitoring and compliance tooling.

## Workflow Lifecycle (Before, During and After)

**Before (Hardware Attested)**

- A user or platform team designs a workflow in the OPAQUE UI, CLI, or SDK, selecting data connections, tools, and guardrails.
- The workflow definition and its tools are packaged as OCI images and signed.
- Only approved, signed images are eligible to run; unsigned or tampered images are rejected.

**During (Verifiable Policy Enforcement)**

- A user or application launches a workflow. The request is authenticated, signed, and sent to the OPAQUE Control Plane (OCP).
- OCP validates the request and coordinates with the customer's confidential environment to start the right images using the customer's KMS / CSP KMS and configuration.
- Policies around who can run what, with which assets and parameters, are enforced during runtime before any code runs.

**After (Verifiable Audit Log)**

- The Verifier produces an attestation report describing the hardware, firmware, images, and policy in effect.
- OPAQUE's Verifier Federator and Voting Service evaluate attestation and policy results. If they pass, the workflow result is returned.
- The Audit Logger records the full run, and telemetry is exported via Otel. Customers can independently re-verify attestation evidence or view attestation documents in OPAQUE UI. A user can leverage their tools of choice to view our Otel compatible logs.

## Workflow Invocation

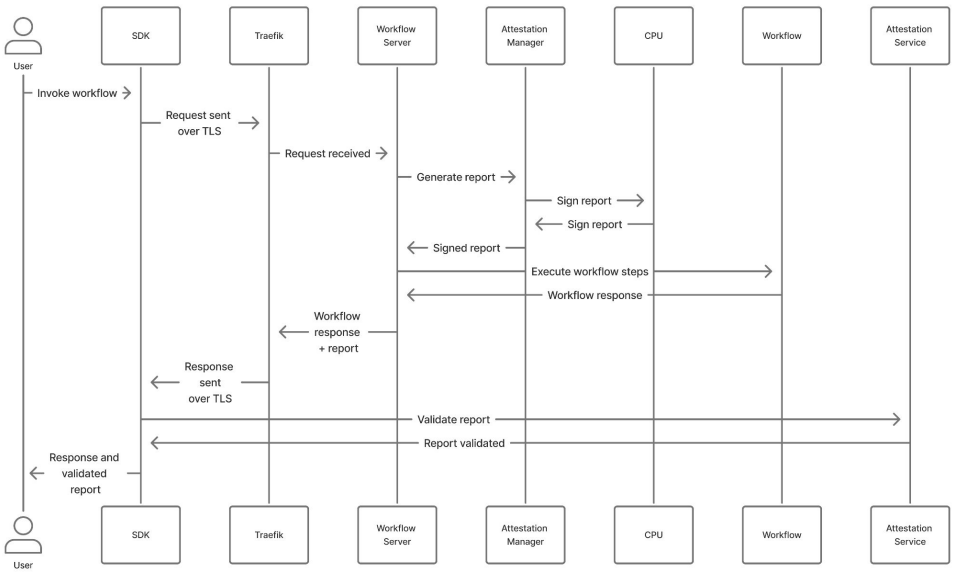


**When a client invokes a workflow:**

1. The client app uses the OPAQUE SDK or API to send a signed request.
2. The request is routed to the customer's confidential environment, where the appropriate OCI images are started.
3. The Verifier collects measurements from the TEE and generates an attestation report.
4. The report is validated against the cloud provider's attestation service and evaluated by the Verifier Federator / Voting Service.
5. Only if attestation and policy checks succeed does the workflow run and return results.
6. The client receives both the response and the attestation evidence, which can be validated independently of OPAQUE's SDK.

## FAQs

**How do you ensure only approved workflows can run?**
Each workflow and tool is packaged as a signed OCI image. OPAQUE and the customer Verifier both check signatures, configuration, and policy. Unsigned or mismatched images are rejected and never started.

**Who controls the keys and trust anchors?**
Customers keep control of their own keys via their Key Management System, backed by the cloud provider's KMS. OPAQUE coordinates attestation and orchestration but does not hold customer data-encryption keys.

**How can we independently verify attestation?**
Every run produces an attestation report and cryptographic signatures from the CSP attestation service. Customers can verify this evidence themselves, outside of OPAQUE's SDK or UI, using their own tooling and trust policies.

**How do we guarantee confidential execution of a workload running in OPAQUE?**
For every request submitted to the platform to process via a given workflow, the platform produces a tree of documents that vouch for the environment in which that request was processed, and which tie the processing of the request and the generation of the response, to the hardware on which said processing and generation took place. Those documents are produced if and only if the platform is hosted in an environment whose measurements are known to represent a trustworthy TEE that not only guarantees confidentiality and integrity of code and data at runtime, but also whose contents (i.e., the code running within it) are known, produced by OPAQUE, and thus known to behave correctly.

**Does OPAQUE see sensitive data related to agentic workflows running inside the environment?**
No. Workloads run inside the customer's confidential compute environment. Data is encrypted at rest, in transit, and in use; plaintext is only processed inside customer-owned TEEs with customer-controlled keys.

**How does this fit into our existing monitoring and compliance stack?**
All workflow activity and telemetry is exported as Otel compatible so you can use Otel compatible tools such as Grafana, Azure Monitor, CloudWatch, or Splunk. Platform and security teams can use tools of choice to display data from OPAQUE.