OneTrust

Data Sheet

# OneTrust AI Governance

## Centralize Governance from Policy to Runtime

AI's next frontier is multi-modal, adaptive, and autonomous, but governance foundations weren't built for it. Without an integrated, at-scale approach to risk modeling, monitoring, and governing AI systems, organizations face inefficient compliance, stalled innovation, lost competitive advantage, and significant reputational and financial risk from broken trust.

**OneTrust AI Governance translates AI risk into enforceable controls so teams can govern well and move fast.**

- **Intake:** Centralize and contextualize AI initiatives in one governance hub.
- **Evaluate & Scope:** Assess risk against regulations to right-size oversight.
- **Implement & Monitor:** Enforce policies with automated controls and guardrails.
- **Report:** Track compliance, program health, and AI Investment status in one view.

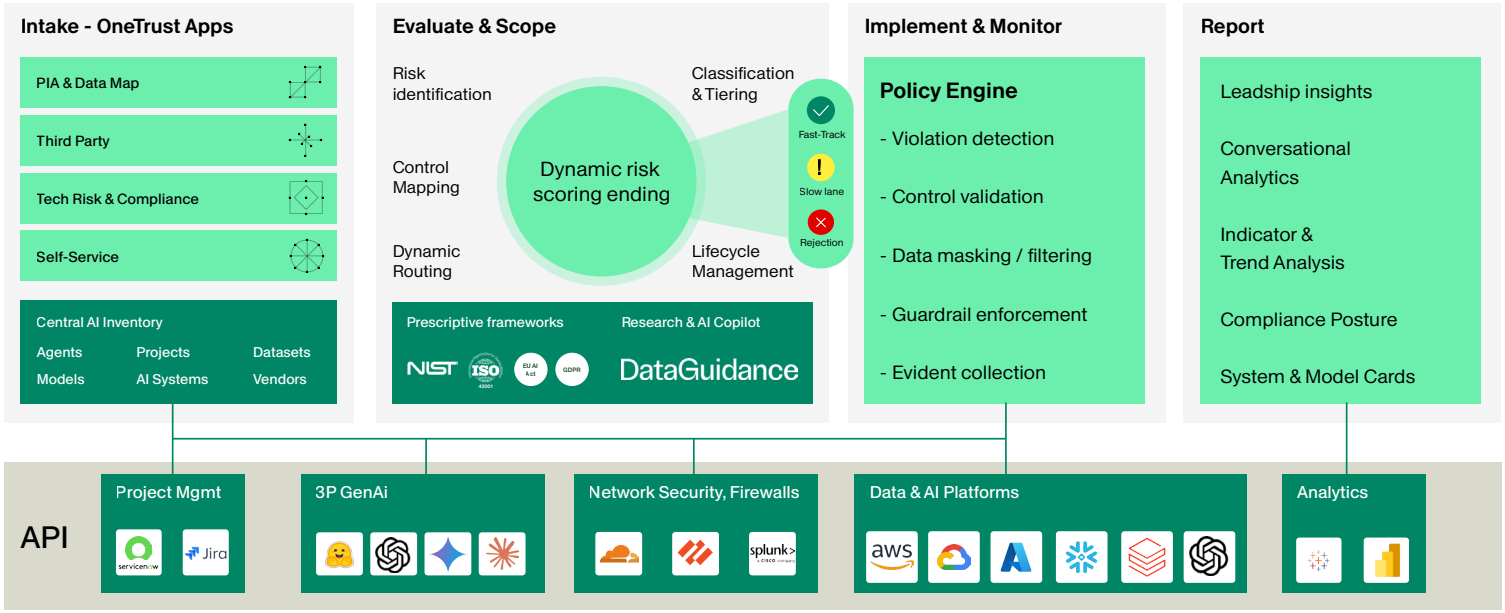


### Accelerate AI Delivery

Ship AI faster by removing governance friction safely.

### Scale Trusted Innovation

Confidently scale AI adoption with continuous system trust.

### Protect AI ROI

Demonstrate responsible AI by design to preserve trust and value.

OneTrust

Data Sheet

## Compliance & Risk Management

Establish a scalable foundation for AI compliance and risk readiness.

**The AI Governance Program center delivers a unified record for AI systems and their components, standardizing risk identification, classification, and control enforcement across environments. By establishing clear ownership and risk alignment upfront, it streamlines intake and compliance, so high-risk initiatives get the scrutiny they require while low-risk work moves faster.**

- Central AI Inventory: Track all models, datasets, agents, vendors and projects in one place.
- Regulatory Intelligence: EU AI Act, NIST, and ISO 42001 assessment templates with automated risk tiering and risk and control frameworks.
- Workflow Approvals & Sign off: Configurable stage gates and attestation tracking
- Compliance Reporting: Automated evidence, model and system cards, and audit ready outputs.

✓ **Faster decisions** ✓ **Consistent outcomes** ✓ **Risk-aware governance**

## Cross-Platform Monitoring

Contextualize AI telemetry with frameworks and policies to guide action.

**AI observability serves as the runtime risk and compliance context layer, interpreting AI telemetry against policy, purpose, and regulatory requirements. By enriching signals like drift, hallucinations, and anomalous behavior with governance context, it helps teams determine what matters and enforce action across AI systems before issues become incidents or audit findings.**

- Model & Agent Telemetry: Continuously ingest drift, quality, safety, and performance signals from AI platforms and providers.
- Context-Aware Monitoring: Correlate runtime signals with regulatory obligations, data sensitivity, intended purpose, and allowed uses.
- AI Policy Monitoring: Detect, log, and surface violations of approved AI policies in production environments.
- Sensitive Data Detection: Automatically identify PII and sensitive attributes in AI inputs and outputs.

✓ **Earlier risk detection** ✓ **Improved oversight without slowing teams**

## Programmatic Control Enforcement

Apply policy-driven guardrails to AI workflows.

**Programmatic enforcement embeds policy and data controls directly into development, deployment, and runtime environments—from data pipelines to agent orchestration. By enforcing guardrails automatically and consistently across platforms, teams move faster while ensuring every model, dataset, and agent operates within approved, compliant boundaries.**

- Runtime Guardrails: Enforce prompt and output filtering, block or allow actions, and constrain unsafe behavior in production.
- Data Policy Enforcement: Apply RAG rules, data use constraints, and access controls to govern how models and agents interact with data.
- CI/CD Gating: Require evaluations, approvals, and policy checks before models or agents are promoted.
- Agent Governance & Protocol Contracts: Register agents and enforce policy contracts governing purpose, permissions, and allowed actions.
- MCP Server Enforcement Layer: Enforce policies across MCP-enabled assistants and tools by filtering tool access, classifying requests by purpose and regulatory context, applying response controls, and generating auditable decision logs.

✓ **Responsible AI by design** ✓ **Reduced product incidents** ✓ **Automated compliance controls**