OneTrust

Mike Wiseman

June 2026

AI Governance Demonstration

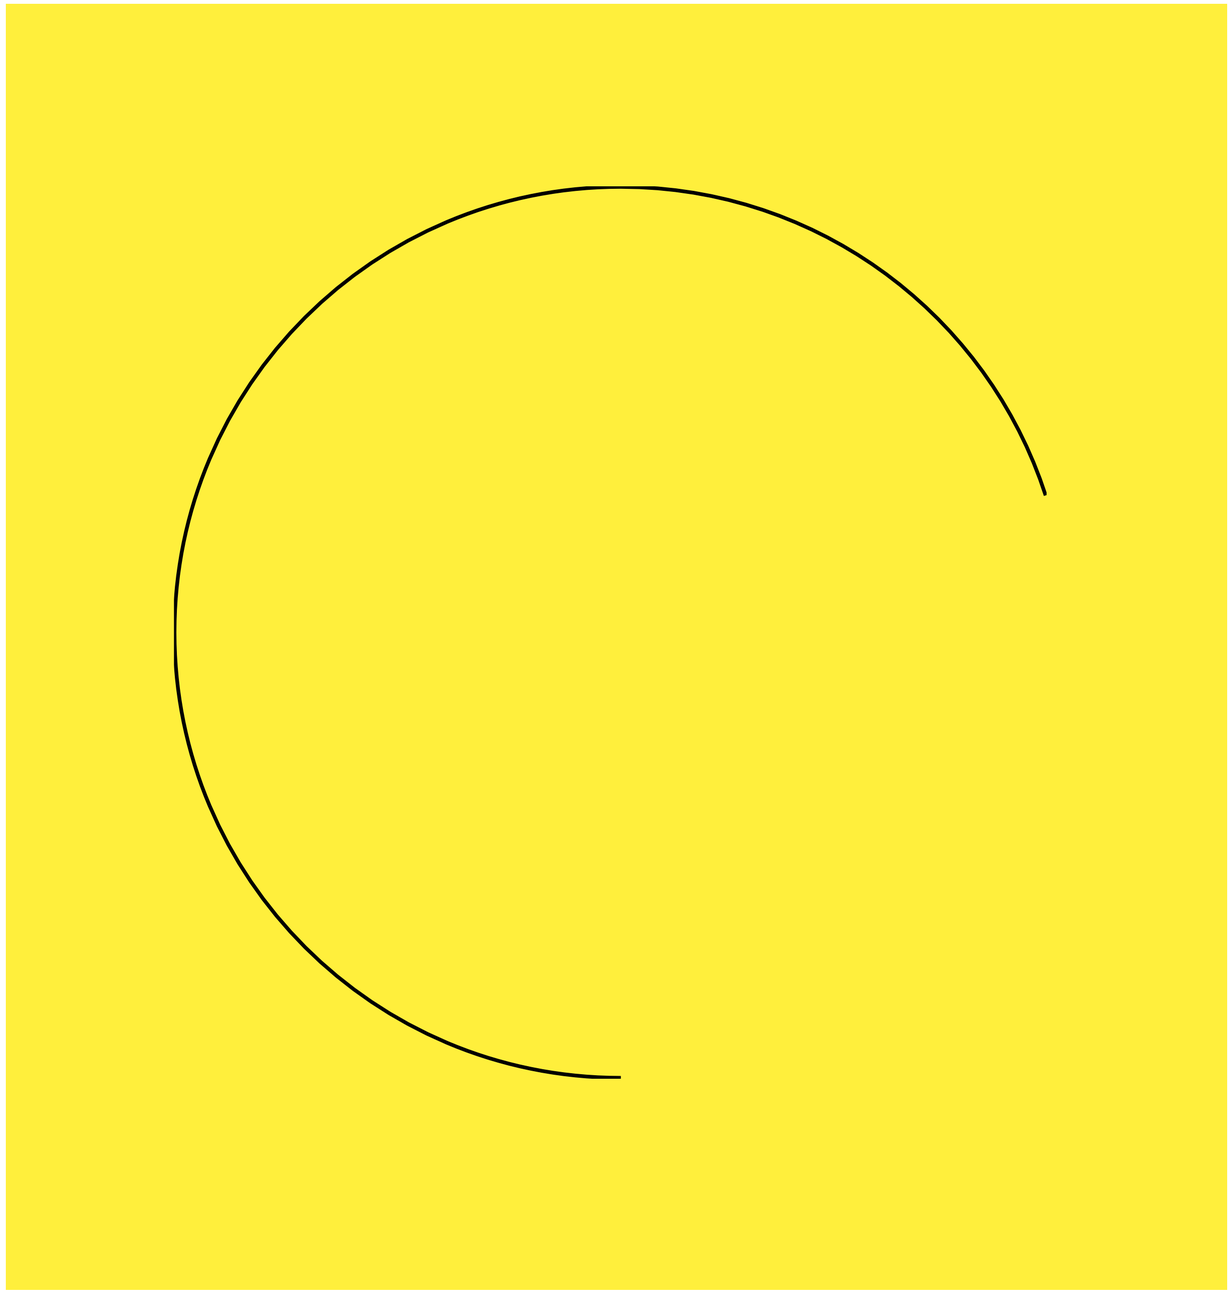

70%

say their ability to **govern AI** is at odds with the **speed** at which AI initiatives move

The speed vs. control tradeoff is disconnect is **killing ROI**

$3.8B

Total enterprise spend on AI agents and assistants in 2026
(59% CAGR between 2025-2030)

40%

Will demote or decommission AI agents due to governance failures by 2027

Sources: * Gartner Sept 2025 / IoT Analytics Sept 2025 / OneTrust 2025 / Forrester Oct 2025

Every AI leader should be asking these questions

And most can’t answer them. OneTrust changes that.

What AI are we using?	• What AI tools, systems, models, and agents are running across your business? • What business processes are your AI actually driving? • Which AI uses have been reviewed and approved – which haven’t?
What risks is AI creating?	• Is your AI touching data it shouldn't? • Which AI is operating outside your internal policies or regulatory requirements? • Which AI systems are running over budget?
Can we actually govern AI?	• Can you respond quickly when AI use violates policies? • Which AI are producing unexpected results? • How will you demonstrate compliance to an auditor or regulator?

Govern with Confidence *and* Speed

Create a full picture of your AI based on runtime behavior

- Shift from attestation-based governance to signal-based governance
- Reduce reliance on ad-hoc conversations to collect information
- Evaluate runtime signal against AI Policies to automate AI governance
- Extend to capture deeper signals and activity tracking as AI governance needs evolve

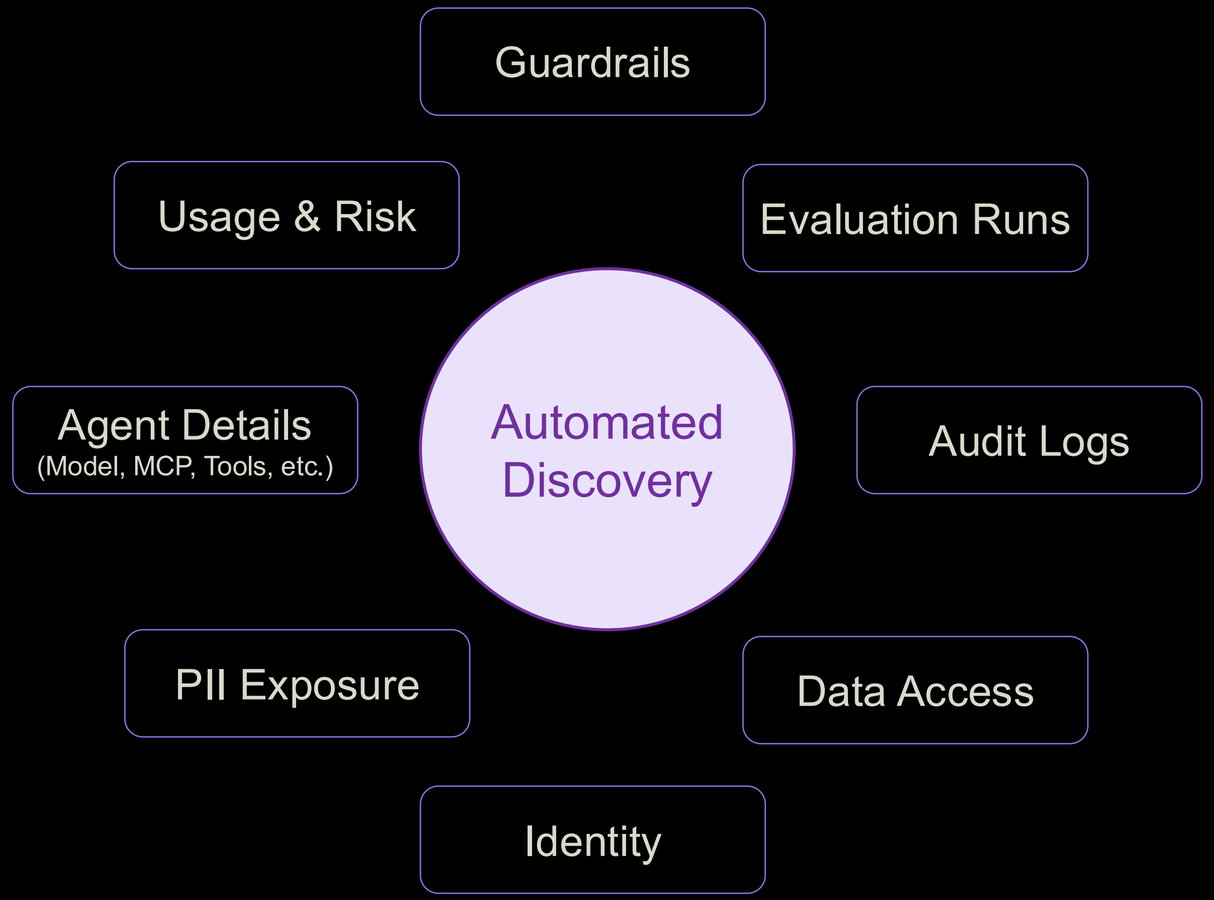

Always-On AI Governance

From one-time reviews to continuous discovery, evaluation, and enforcement.

Model & Agent Discovery	AI App Discovery (AI Guard SDK)	Guardrail Discovery	Evaluation Metrics	Operational Metrics	Log Analyzer (PII Detection)
• Vendor specific integrations which OneTrust manages with fast evolving AI landscape • Discover all agents across major platforms in few clicks	• AI Guard SDK and PII Classification • Discover in-progress projects early on • Test apps for PII exposure • Early course correction feedback	• Extend discovery to understand Guardrails applied to AI assets and how they are setup • Identify non-compliant settings from your Governance console	• Supports existing platform Evaluation Frameworks • Metrics for LLM-as-a-Judge • Measure Toxicity, Bias, Helpfulness	• Adds new dimensions to AI runtime • Track usage across large Agent estate • Connect dots to token consumption • Measures guardrail effectiveness	• 300+ classification profiles • PII detected, categorized, blocked, redacted • Extends support to real-time with metrics

Policy Evaluation & Enforcement

- Out-of-box policy rules for common standards and frameworks (NIST AI RMF, ISO 42001, EU AI Act)
- Custom policy rules, for your governance needs
- Model and agent guardrail enforcement
- Trigger model evaluation

Inventory AI assets

Discover and catalog AI systems, agents, models and data across the ecosystem and ensure accountability across the AI lifecycle.

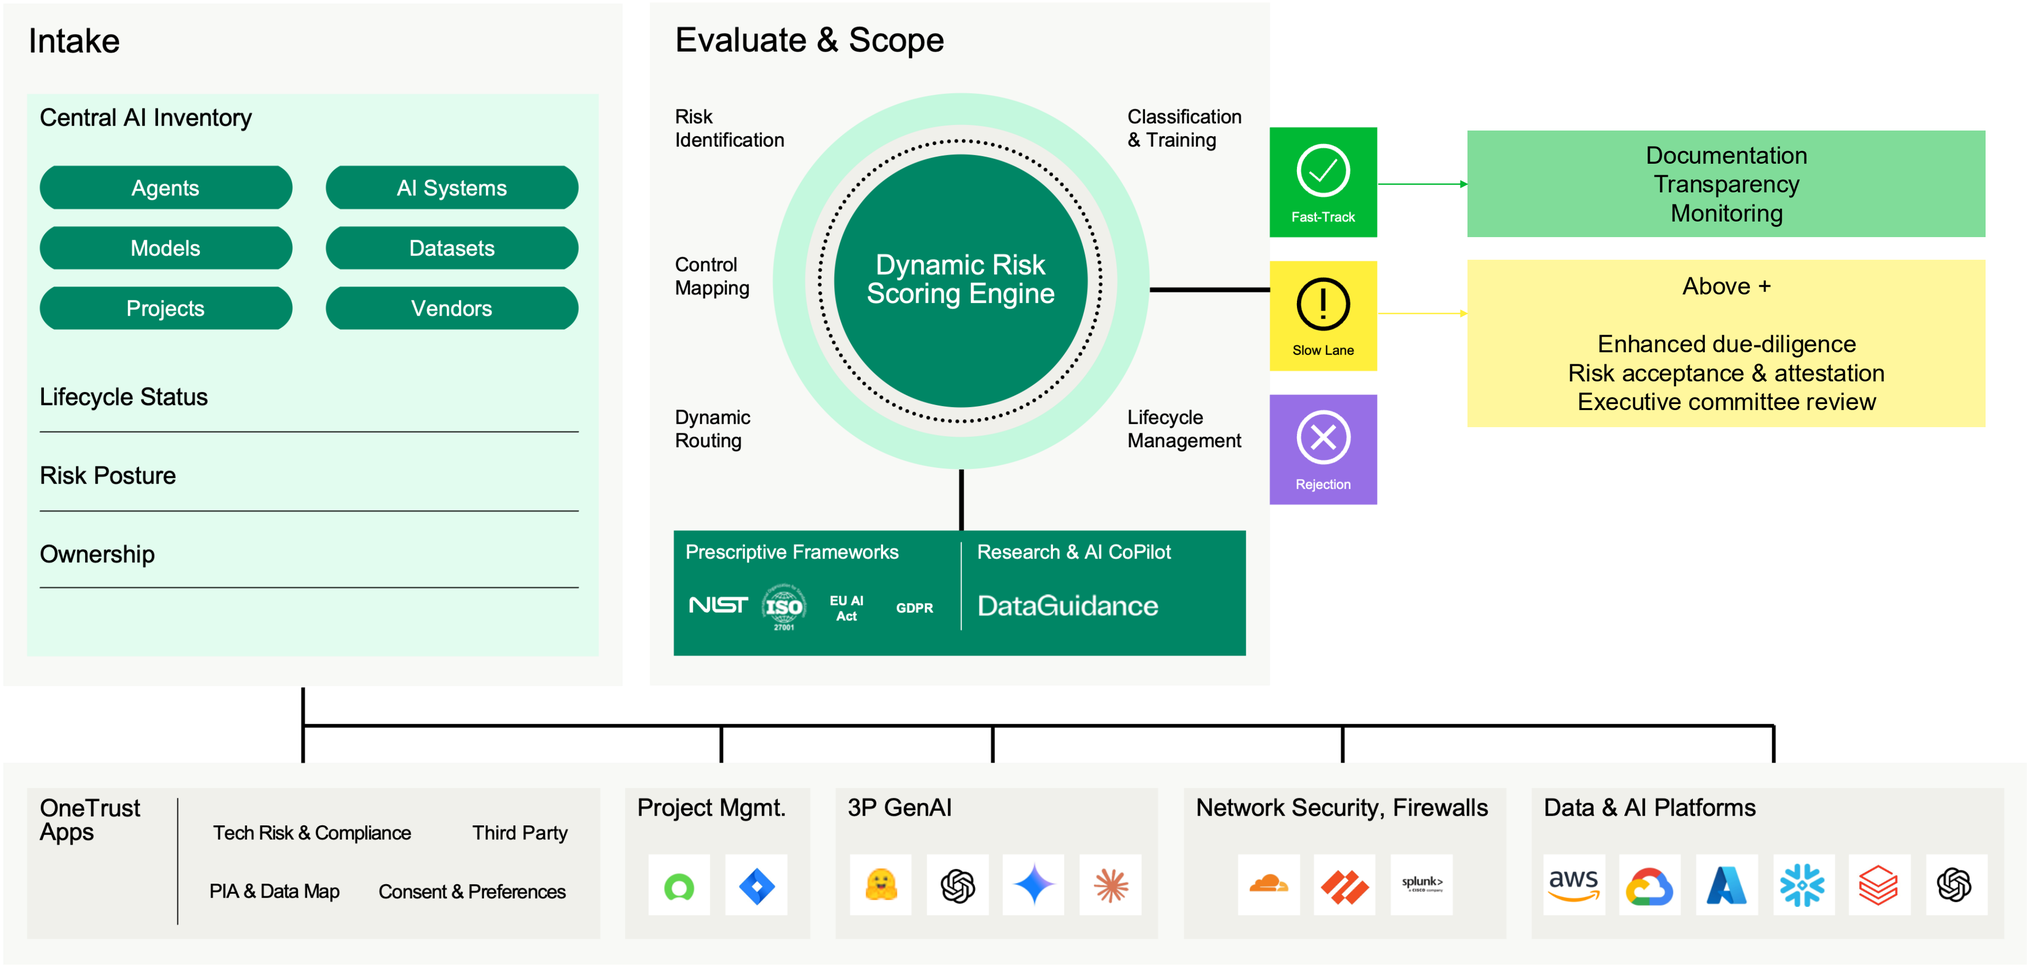

Activate Inventory

Most AI estates are ungoverned because nobody knows what's in them.

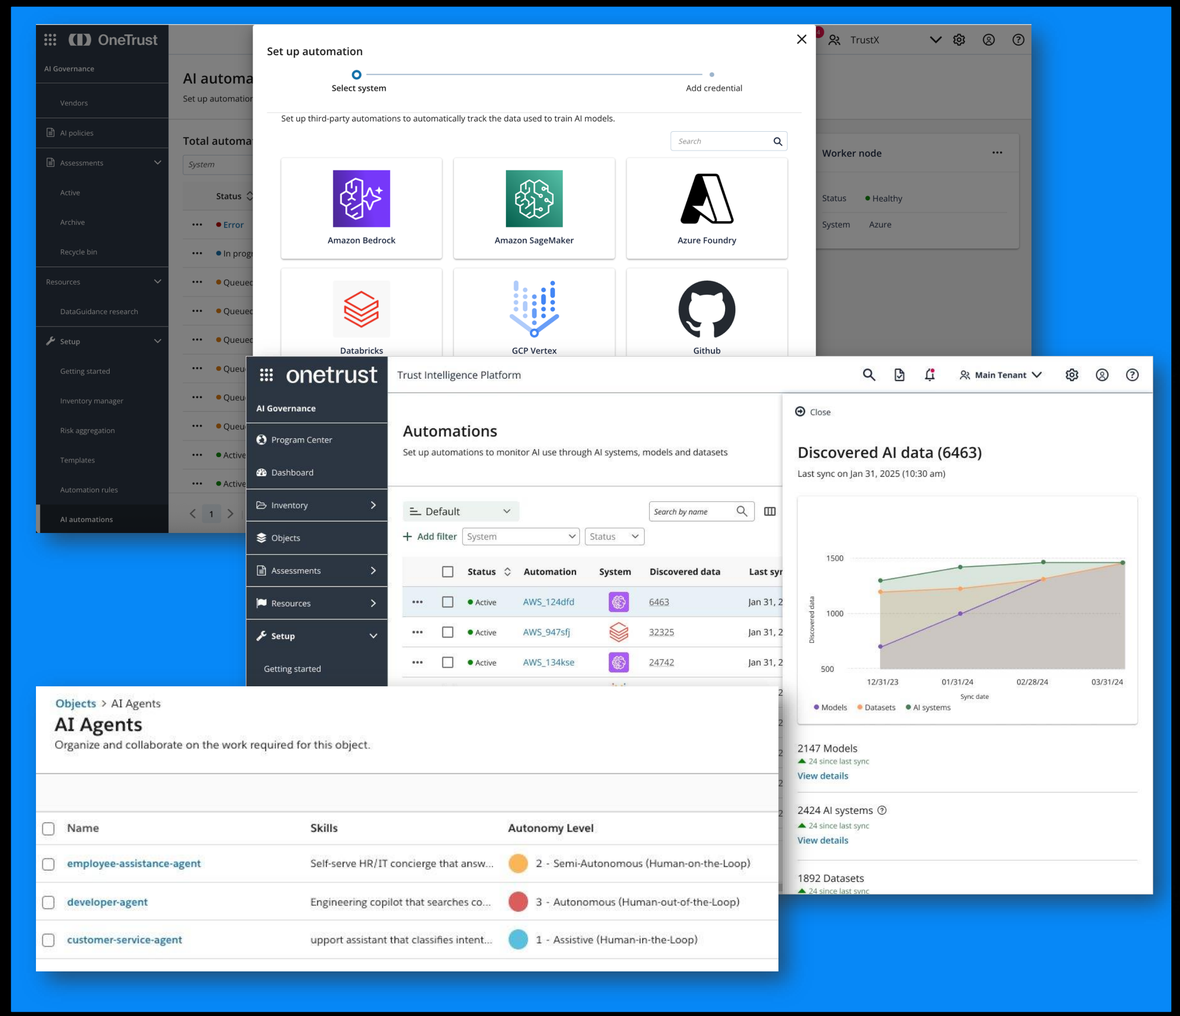

AI Automated Discovery
Surface every AI asset including shadow AI and models you didn't register. Automated scanning across the AI estate keeps inventory continuously current without manual effort.

Agent Inventory
Maintain a live, structured record of every agent in your environment, capturing ownership, purpose, data access, integrations, and lifecycle status.

Evaluation and Service Health Metrics
Track how AI performs in production, not just that it exists. Surface quality scores, latency, guardrail invocations, PII exposure, and uptime alongside your inventory.

Intelligence, Not Just Inventory

Knowing AI exists isn't the same as understanding it.

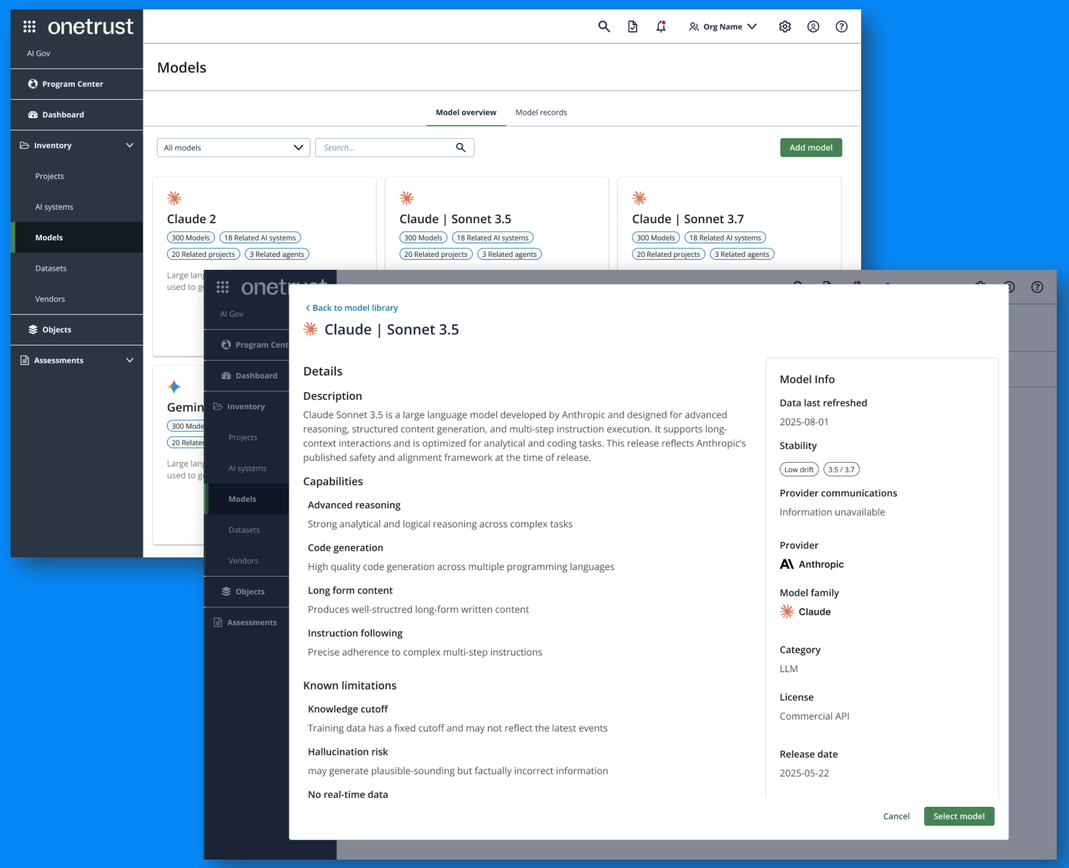

Foundation Model Intelligence
Curated, continuously refreshed intelligence on the models powering your AI inventory, embedded directly where governance decisions happen.

Provider Concentration Views
Understand at a glance how many records, systems, and projects depend on each model provider.

AI System Coverage (Next)
Extending intelligence beyond foundation models to the full ecosystem of vendor-packaged AI, so governance spans everything from the model layer to the SaaS products built on top of it.

Continuous Policy Enforcement

Most violations aren't caught until an audit. By then, the damage is done.

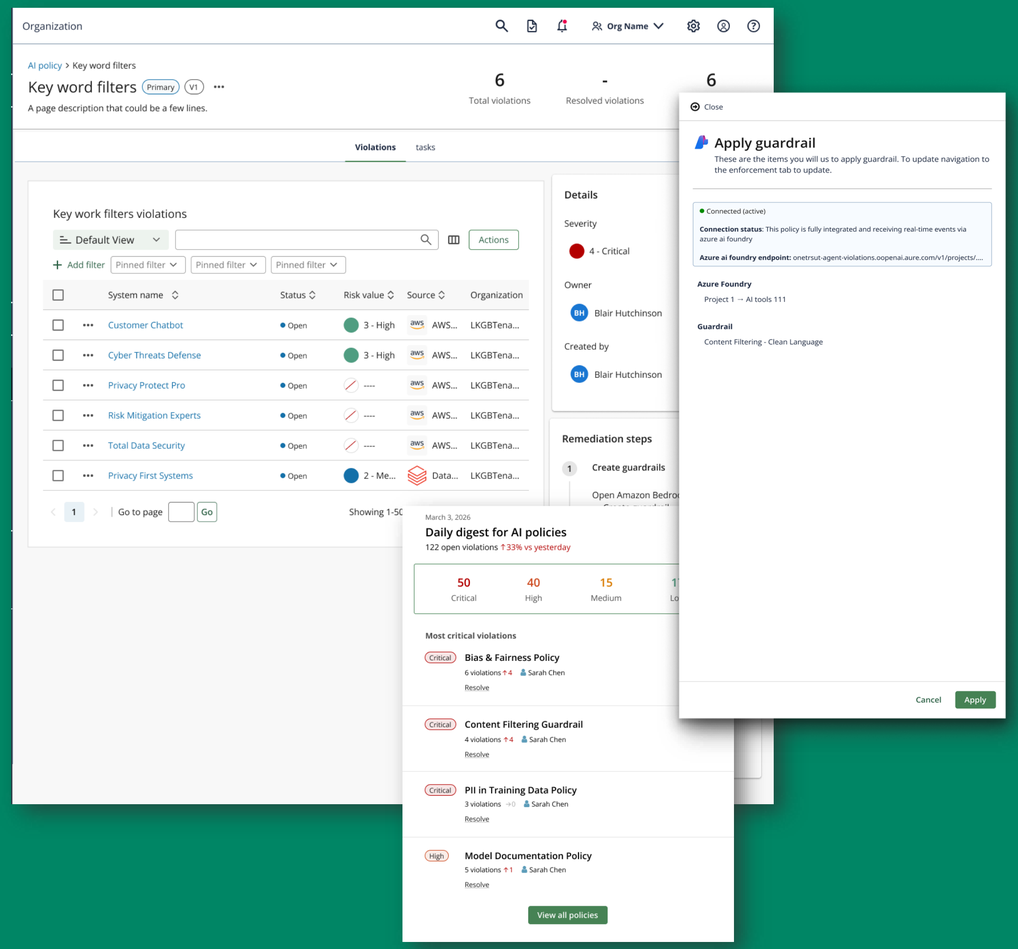

AI Policy Management

Define AI policies centrally, from scratch or adapted from technical standards and regulatory frameworks.

Automated Actions

Escalate risks, launch assessments, or adjust controls without manual intervention when policies are violated.

Email Digests

Alert policy owners to violations by severity with configurable digests so nothing falls through the cracks.

Guardrail Policy Enforcement

Push guardrails directly to connected AI platforms so policies hold at runtime, not just on paper.

Build AI Confidently

Agentic AI moves too fast for governance built around static systems.

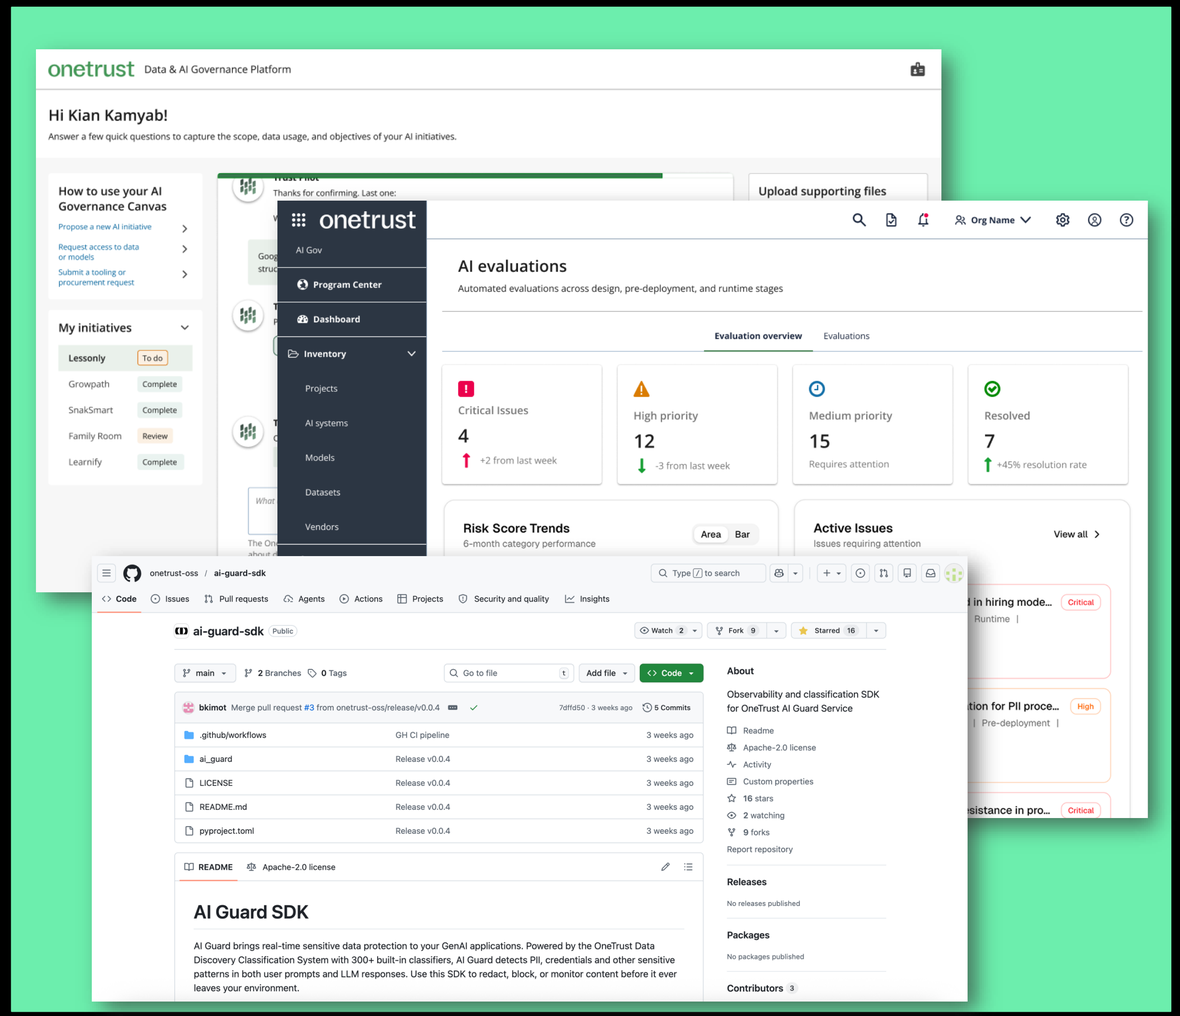

AI Guard SDK
Detect and block PII in your GenAI apps before leaks reach users.

AI Governance Intake Agent (Next)
Classify risk, apply policies, and route to reviewers automatically on intake.

AI Gov Agent Skills (Next)
Embed policy checks and compliance logic into the AI development workflow.

Compliance Evaluators (Next)
Test against safety, bias, and regulatory requirements to produce audit-ready evidence.

Scale AI Governance Operations

Programs that don't scale become bottlenecks for every new AI initiative.

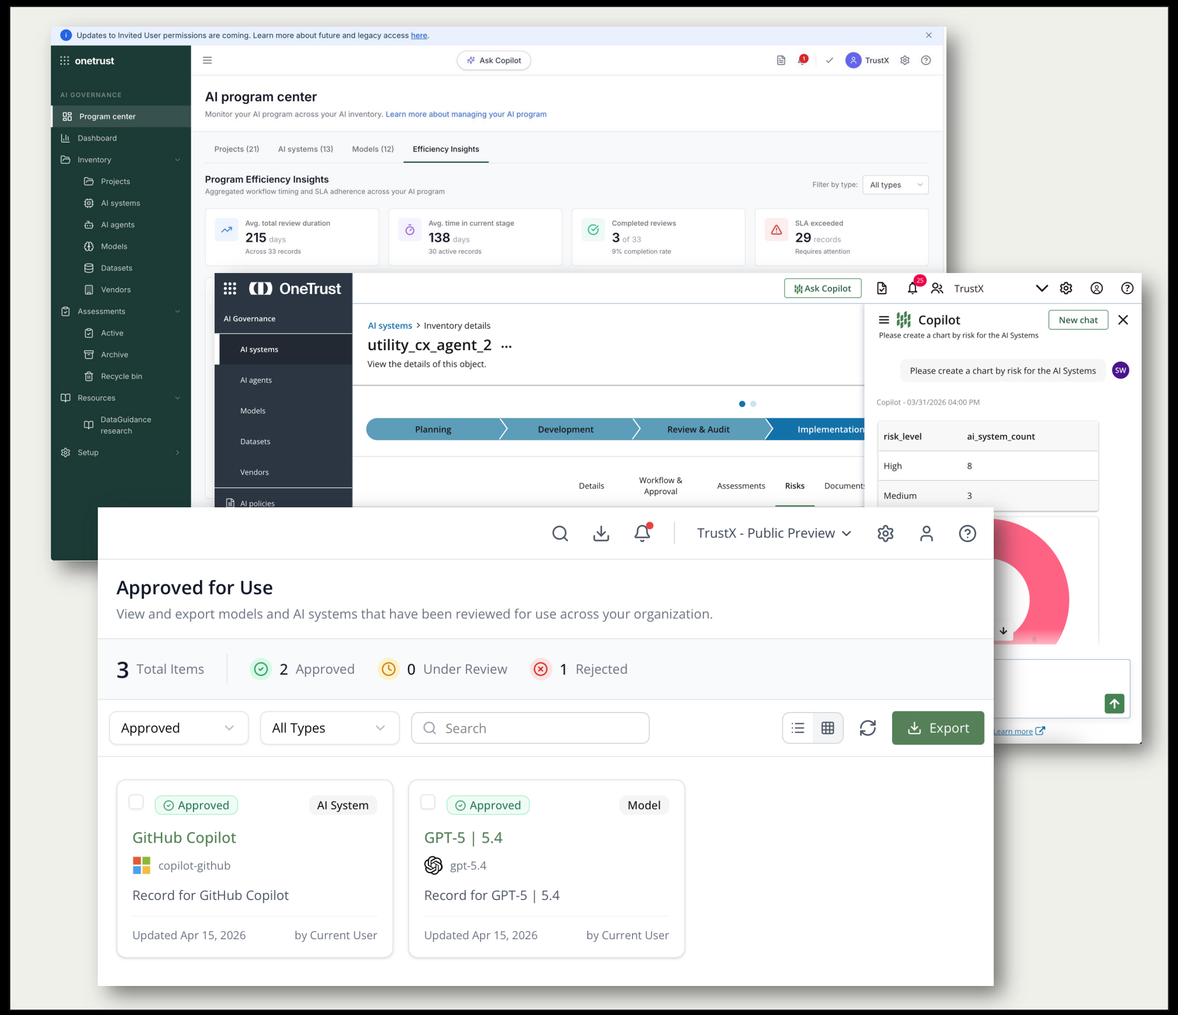

Acceptable Use Catalog (Next)
Pre-approved use cases business teams can adopt without triggering a new review.

Program Efficiency Insights
See where approvals stall and which teams are behind SLA.

Conversational Analytics
Ask governance questions in plain English. No reports or analyst support needed.

Each AI era adds new risk without removing the old

WE ARE HERE

Wave 1

Predictive AI

Characteristics
- Traditional ML
- Predictive outputs
- Specialized teams
- Dedicated infrastructure

Key Governance Challenges
- Quality data
- Accuracy

Wave 2

Generative AI

Characteristics
- Subset of deep learning
- Adaptable & non-deterministic
- Accessible to anyone
- Embedded in applications

Key Governance Challenges
- Appropriate use
- Sprawl

Wave 3

Agentic AI

Characteristics
- Autonomous problem solving
- Human out of the loop

Key Governance Challenges
- Access
- Autonomy

How the analyst firms define AI Governance

Gartner®

AI Governance Platforms

"…ensure organizations comply with responsible AI practices/policy/regulations and risk frameworks"

"…act as a **central repository** linking trust/risk/security runtime controls, **automate workflow approvals** for new AI use cases/apps/agents, and support **risk-based, real-time guardrails"**

FORRESTER®

AI Governance Solutions

"…**inventory of AI assets** and performing **risk assessments** plus the ability to scale deployments with standards, workflows, classification of risk, audits, and remediation"

IDC

Unified AI Governance Platforms

"…integrated suite to oversee the **entire AI lifecycle** (ML + GenAI + agentic AI), including **centralized model registry**, **model discovery**, **continuous monitoring** (bias/drift/security/performance), **policy management**, **risk assessment**, and **audit trails**, plus reporting and integration with enterprise systems"