Qualys

SOLUTION BRIEF

FedRAMP Moderate Authorized

Qualys TotalAI

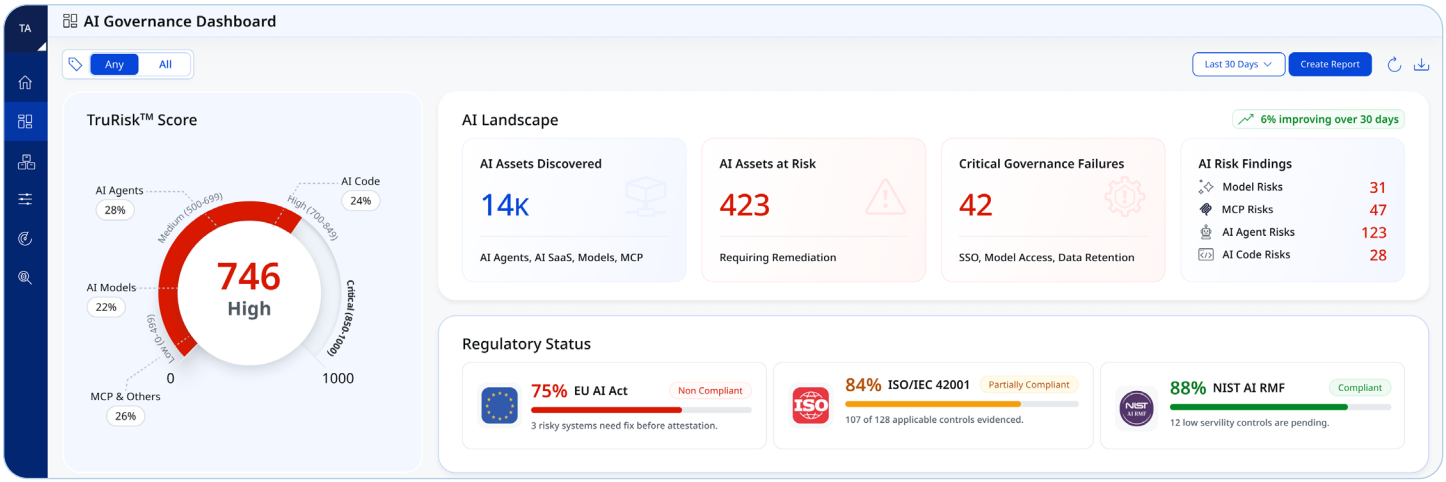
Know what AI exists. Test what can fail. Fix what matters. Prove your controls work.

AI has moved from isolated experiments into business-critical infrastructure and is adopted much faster than most organisations can govern them. Employees use generative AI through browsers and SaaS copilots. Developers embed models, agents, retrieval pipelines, and tool integrations into applications. Cloud teams deploy managed AI services and GPU workloads. Each path creates a different form of exposure, but most security programs still manage AI through self-reported inventories, questionnaires, periodic scans, and disconnected point tools. The business challenge is to prove AI is safe to their board, regulators, and customers. AI governance is the validation that security controls exist, operate continuously and reduce risk; and to prove this, it needs continuous technical evidence.

Turn AI Governance into Continuous Risk Operations

Qualys TotalAI makes AI risk measurable to bring fast-moving, ungoverned enterprise AI adoption under control with unified AI exposure management, security, and governance - covering both AI that the enterprise builds and AI that employees consume.

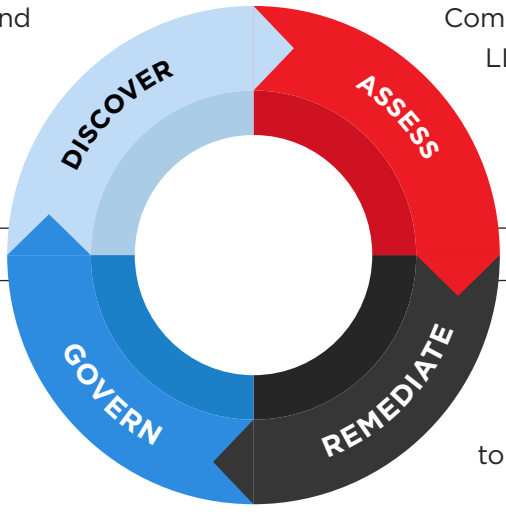
TotalAI treats AI as a first-class attack surface. It continuously discovers AI assets, evaluates infrastructure and AI-specific behavior, tests models and MCP servers, monitors AI workloads in production, prioritizes findings through TruRisk, guides remediation, re-tests controls, and produces evidence proving which controls work. Built on the Qualys Enterprise TruRisk Platform, it connects AI-specific behavior risks with infrastructure, cloud, container, application, and business context so teams can prioritize the AI risks capable of causing material harm.



One Place to Control Enterprise AI Risk

Continuously identify sanctioned and shadow AI across hosts, cloud, network, containers, code, SaaS, models, agents, and MCP servers.

Combine AI posture, vulnerability context, LLM and MCP adversarial testing, code checks, and runtime telemetry.

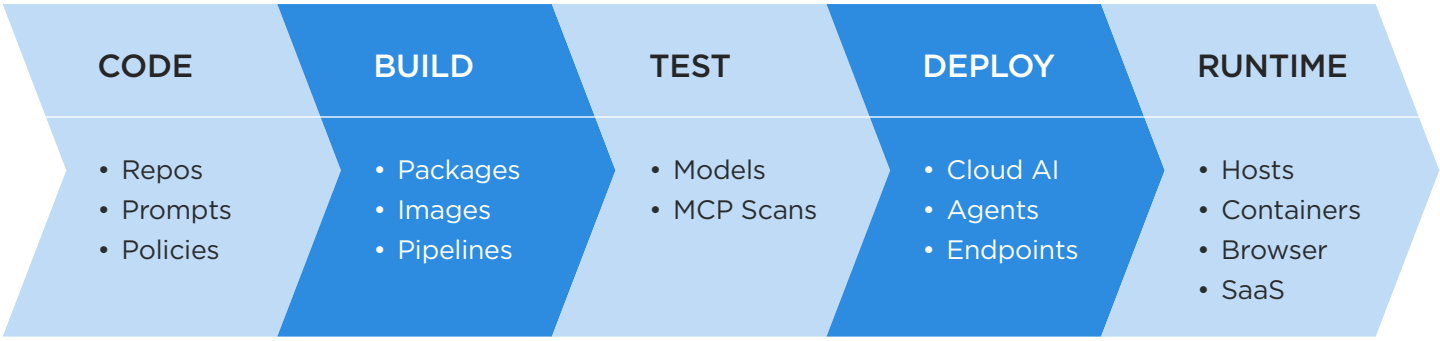


Produce repeatable evidence and reporting for boards, auditors, procurement, model review, GRC, ATO, and continuous monitoring.

Prioritize with TruRisk, route findings to the right owner, guide technical fixes and guardrails, then re-test.

Coverage Across the AI Lifecycle

TotalAI connects pre-production controls with production evidence.



Four-Pillars of AI Security Across the AI Risk Lifecycle

AI-SPM

AI Discovery & Posture Management

960+ AI-specific detections & 55+ AI posture controls.

- Continuously discover sanctioned and shadow AI across hosts, cloud services, models, MCP servers, AI software, images, containers, and workloads.
- Connect AI assets to ownership, environment, exposure, business criticality, and approval status.
- Inventory MCP servers, tools, resources, and endpoints, and classify agentic AI risk.
- Assess cloud AI posture across supported AWS, Azure, GCP, and OCI environments.
- Evaluate supported AI SaaS configurations, users, groups, and control posture.
- Assess failed controls, cloud misconfigurations, vulnerable AI infrastructure, excessive access, and exposed AI endpoints.
- Surface toxic combinations and prioritize AI risk using TruRisk™ and contextual risk insights.

AI-Runtime Security

Continuous Behavioral Monitoring

Monitor AI workloads, agents, models, MCP servers, browsers, and supported AI SaaS after deployment.

- Detect posture drift, abnormal model and MCP activity, suspicious workload behavior, and data-exposure indicators.
- Use eBPF telemetry to observe processes, files, APIs, network activity, and outbound connections inside AI agent sandboxes.
- Identify C2, malicious destinations, DNS tunneling, reconnaissance, command execution, and exposed credentials.
- Reveal sanctioned and shadow workforce AI usage through browser telemetry and SaaS connectors.
- Correlate code, policy, workload, and runtime evidence to show what AI was allowed to do and what it actually did.
- Preserve timestamped evidence for investigations, audits, and governance reporting.

Qualys

SOLUTION BRIEF

AI Adversarial Testing

Active LLM and MCP Behavioral Testing

300+ attack scenarios to test models for prompt injection, jailbreaks, sensitive-data disclosure, hallucination, denial of service, multilingual, encoded, and multimodal attacks.

- Scan supported model runtimes, including Chat Completion, Hugging Face, AWS Bedrock, etc.
- Scan public, private, cloud-hosted, and on-premises models using repeatable Option Profiles.
- 100+ detections to assess MCP servers for tool poisoning, injection, SSRF, spoofing, unsafe file access, command execution, and credential exposure.
- Preserve prompts, responses, severity, pass/fail/error results, successful jailbreaks, and technical justification.
- Map findings to OWASP LLM, MITRE ATLAS, and EU AI Act requirements.
- Recommend corrective guardrails and retest to prove whether risk was reduced.

AI Code Security

Shift-Left AI Security

Find vulnerabilities, secrets, malware, risky dependencies, and unsafe AI patterns before deployment.

- Assess prompts, RAG logic, model-output handling, agent tools, MCP configurations, and Cedar authorization policies.
- Scan AI repositories, images, packages, ML frameworks, and build artifacts.
- Apply pass/fail repository controls and policy-based CI/CD gates.
- Generate SBOM and build-to-deploy evidence for governance.
- Feed failed adversarial tests back into code, prompts, retrieval logic, policies, and guardrails.
- Correlate code findings with deployed AI assets, workloads, and runtime behavior.



Govern Enterprise AI from Discovery to Runtime.

Start a 30-day, no-cost trial on the Qualys Enterprise TruRisk™ Platform.

Book a Demo

About Qualys

Qualys, Inc. (NASDAQ: QLYS) is a pioneer and leading provider of disruptive cloud-based Security, Compliance and IT solutions with more than 10,000 subscription customers worldwide, including a majority of the Forbes Global 100 and Fortune 100. Qualys helps organizations streamline and automate their security and compliance solutions onto a single platform for greater agility, better business outcomes, and substantial cost savings. Qualys, Qualys VMDR® and the Qualys logo are proprietary trademarks of Qualys, Inc. All other products or names may be trademarks of their respective companies.

For more information, please visit **qualys.com**

v260824

©2026 Qualys, Inc. All Rights Reserved.