TRUSTAI RUNTIME SECURITY OVERVIEW

Secure Agentic AI at Runtime with TrustAI

TrustLogix TrustAI is a unified policy control plane built to secure enterprise AI and agentic workflows at runtime. It helps organizations govern how AI agents, copilots, MCP tools, applications, and machine identities access sensitive data, tools, and downstream systems while preserving auditability and operational speed.

Built on the TrustDSPM and TrustAccess foundation, TrustAI combines discovery, integration with data classification and semantic context, policy evaluation, and real-time enforcement in one architecture. This makes it especially relevant for organizations prioritizing agent security and runtime security rather than static governance alone.

WHY TRUSTAI

Runtime-first control for AI and agents

Runtime control for AI agents, copilots, MCP tools, and applications.

Context-aware policy evaluation using identity, sensitivity, requested action, semantic context, and runtime environment.

Advanced PDP controls for custom enterprise agents.

Continuous auditability and rapid-response kill-switch capabilities.

TWO DEPLOYMENT MODES

One control plane, two enforcement paths

Gateway mode

Gateway mode secures off-the-shelf MCP tools, SaaS applications, and copilots by intercepting and brokering requests before they reach protected systems.

Direct integration + PDP mode

Direct integration + PDP mode is designed for custom enterprise agents that require deeper runtime integration, advanced policy controls, and decisioning close to the agent workflow.

FOUNDATIONAL LAYERS

Built on TrustDSPM and TrustAccess

TrustDSPM

TrustDSPM provides discovery, exposure mapping, risk insight, and integration with data classification and semantic context across the enterprise data environment.

TrustAccess

TrustAccess provides least-privilege enforcement, masking, row-level controls, and dynamic policy enforcement across platforms such as Snowflake, Databricks, AWS, Azure, BI tools, and knowledge systems.

SOLUTION ARCHITECTURE

A unified control plane across runtime and data

The architecture illustrates two TrustAI enforcement paths: Gateway mode for MCP tools, SaaS, and copilots, and direct integration plus PDP mode for custom enterprise agents with advanced runtime controls and kill-switch protection.

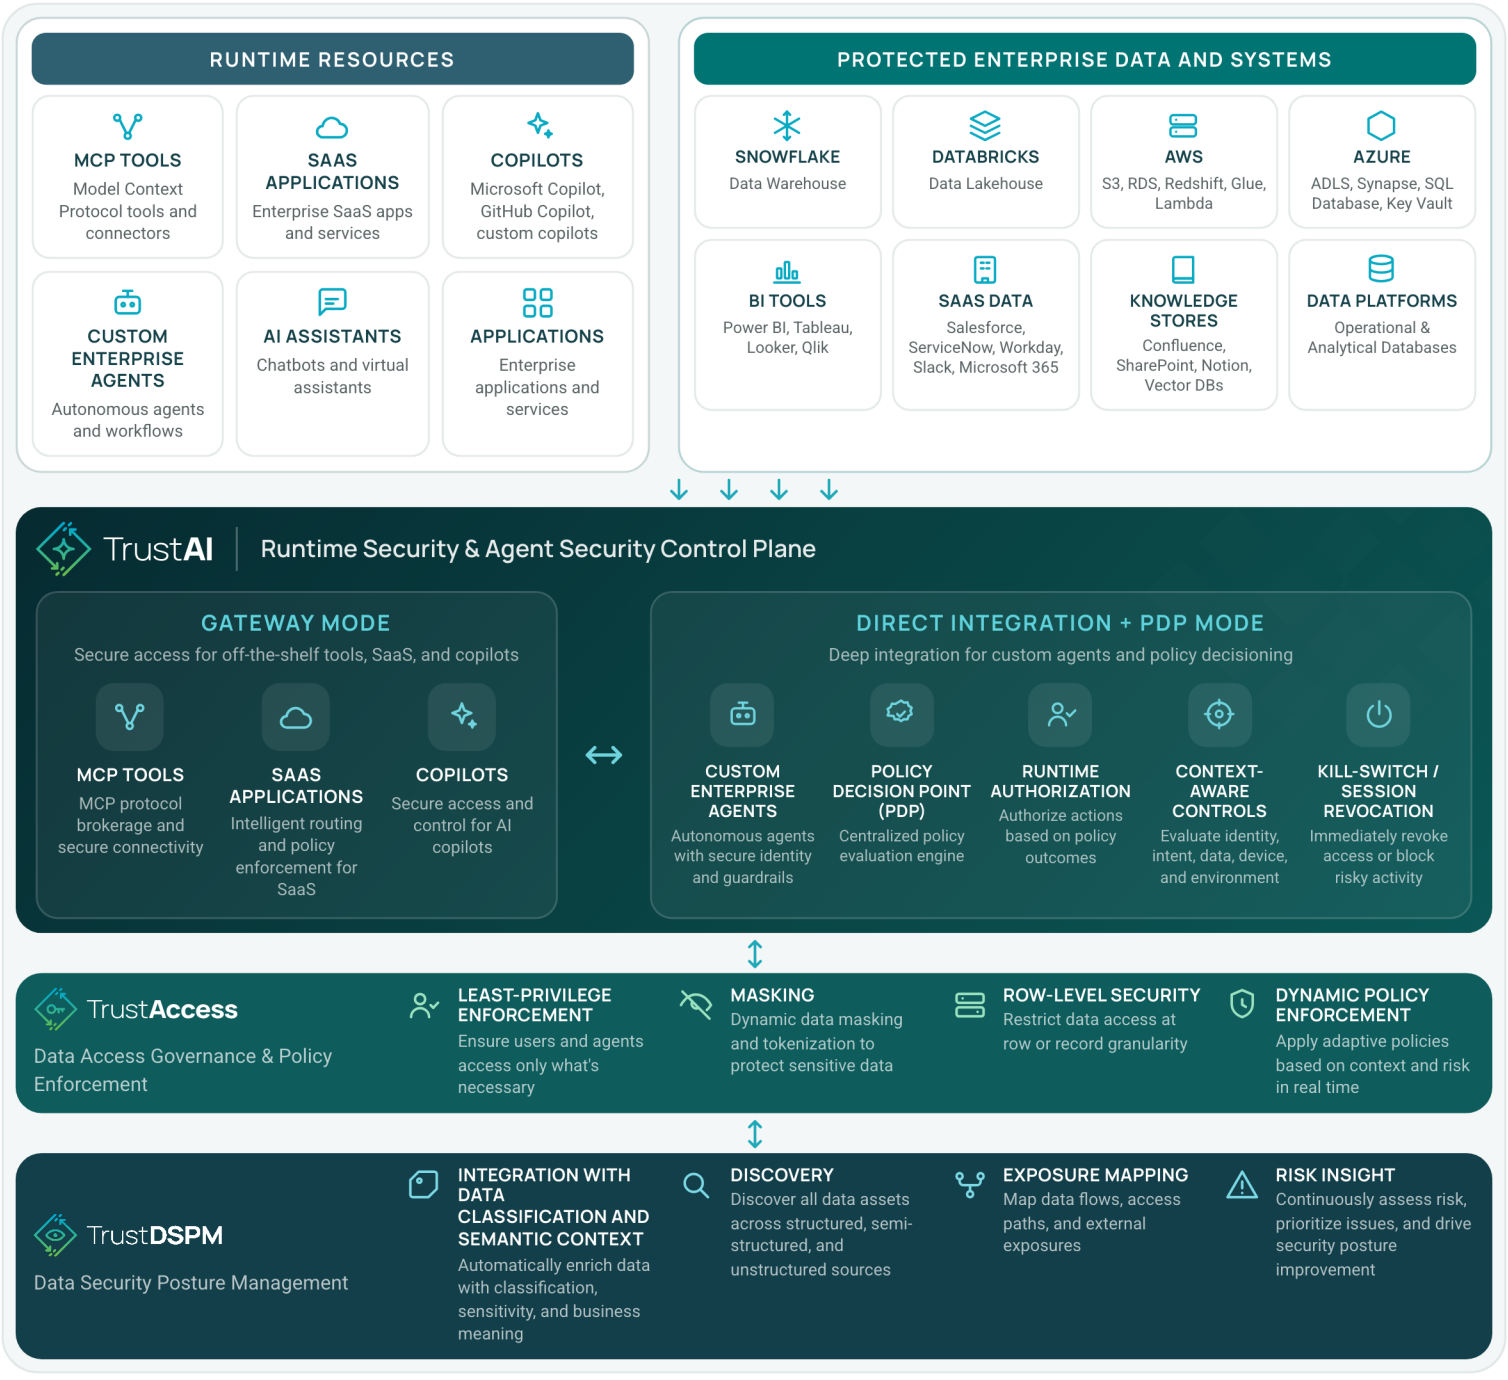

HOW IT WORKS

Evaluate intent, apply policy, authorize access

Runtime resources interact with TrustAI before reaching protected enterprise data and systems. The TrustAI layer evaluates intent and context, applies policy, and authorizes or blocks access in line with enterprise controls.

This architecture gives security and data teams a practical way to deploy AI faster while reducing exposure from autonomous behavior, over-permissioned access, and disconnected enforcement models.