



WHITE PAPER • OCTOBER 2026

Secure Operational AI

Enabling AI and AI Agents for Any Mission

Build secure agents, tools, and operational capabilities without changing the systems and AI you already use.

From the tactical edge to the enterprise • Bring Your Own AI

**ZeroTrusted.ai . . . The AI/SI Force Multiplier for your enterprise,
department, and/or projects/programs**

“Our products allow you to enable your mission. We just do it with security, privacy, and transparency, every time you go to work.”

Waylon Krush, CEO

For DoW and federal mission owners, prime integrators, and enterprise CISOs

ZeroTrusted.ai LLC • Awardable on the CDAO Tradewinds Solutions Marketplace •
www.zerotrusted.ai

Contents

Executive Summary.....	2
Build Secure Agents and Tools in Simple Steps.....	3
The Problem: AI Adoption Is Blocked by Trust, Not Technology.....	4
Force Multiplier for Cleared and Certified Workforces.....	4
Bring Your Own AI (BYOAI).....	5
Architecture Overview.....	5
From the Edge to the Enterprise.....	6
Security, Privacy, and Transparency: Every Time You Go to Work.....	8
Mission Use Cases.....	9
Deployment Options.....	9
How to Start.....	10
The Bottom Line.....	10
About ZeroTrusted.ai.....	10

Executive Summary

ZeroTrusted.ai . . . The AI/SI Force Multiplier for your enterprise, department, and/or projects/programs

AI and AI agents are ready to do real mission work. Most missions are not yet allowed to use them. The blocker is rarely the model. It is trust: who sees the data, what the agent is allowed to do, whether anyone can prove what happened, and whether the security team can sign off without rebuilding the systems the mission already depends on.

ZeroTrusted.ai removes that blocker. We enable secure operational AI and AI agents, no matter the mission. Your teams can easily build secure agents, tools, and operational capabilities on top of the systems and AI they already use, without changing those systems or replacing their models.

Our products allow you to enable your mission. We just do it with security, privacy, and transparency, every time you go to work.

Waylon Krush, CEO

- **Build secure agents and tools in a few simple steps.** Bring your AI, connect your systems, build or clone an agent, set its guardrails, test it, and put it to work. No code changes to your applications and no rip-and-replace.
- **Bring Your Own AI (BYOAI).** Keep your own models, LLMs, agents, and tools: commercial, open-source, or built in-house. ZeroTrusted.ai secures, governs, and makes them transparent without changing them.
- **Run it from the edge to the enterprise.** The same secure layer runs on a forward-deployed node in a disconnected (DDIL) environment, at mission and regional nodes, and in enterprise data centers and private clouds. It is enterprise software that deploys inside your environment and authorization boundary. It is not SaaS.
- **Multiply your workforce.** Secure agents take on repeatable analysis, triage, evidence collection, and documentation. Your cleared and certified people keep decision authority and execute more operations and missions. Agents augment people; they do not replace them.

- **See everything.** Prompts, responses, tool calls, approvals, and outcomes are logged. Evidence packages are hashed and tamper-evident for assessor and authorizing-official review.

Build Secure Agents and Tools in Simple Steps

Securing AI is half the job. The other half is putting AI to work. ZeroTrusted.ai gives mission teams a simple, governed way to build their own agents, tools, and playbooks and field them with confidence. You don't need a new stack, a new model, or code changes to the systems you already run.

1. **Bring your AI.** Point ZeroTrusted.ai at the models you already use: commercial LLMs with your own keys (BYOK), or self-hosted and local models. Assign an approved model to each agent and tenant, and record its provider, version, token limits, cost controls, and data egress.
2. **Connect your systems.** Use connectors for the platforms you already operate, including SIEM, EDR, identity, email, cloud, ticketing, threat-intelligence, and OT tools, or add your own connectors and tools.
3. **Build or clone an agent.** Start from a platform-shipped agent and clone it, or register your own agent. Give it an identity, an owner, skills, tools, and a mission.
4. **Set the guardrails.** Define permissions, tool and function-call allow-lists, autonomy limits, and the approval policy. AI Firewall policy covers every prompt, response, and tool call the agent makes.
5. **Test before you trust.** A one-button workflow runs preflight, benchmark rubrics, adversarial testing, agent-vs-agent and escape testing, and human review, then exports a signed evidence package and mission trace.
6. **Put it to work, and keep it honest.** The agent operates inside its envelope. High-impact actions wait for a person, every action is traced, and findings drive approval-gated remediation and retest.




Figure 1. Secure agent lifecycle. Every agent is registered, verified, tested, and reviewed by people before it operates.

What you do not have to change

- **Your systems.** Agents work through connectors to the tools you already own, inside your existing boundary.
- **Your AI.** Models, agents, and tools stay as they are (see Bring Your Own AI). You can switch models without changing your policies, approvals, or evidence.
- **Your applications.** HealthCheck runs alongside your models and agents with no code changes to the application. The AI Firewall can run inline, out-of-band, or at the browser edge.
- **Where you operate.** The same build-and-test workflow applies on a forward-deployed laptop and in the enterprise data center (see From the Edge to the Enterprise).

What stays honest. Preflight verifies binaries, connectors, target reachability, authorization, model access, and evidence output before a mission starts. Missing tools stay unavailable and visible, and failed or unconfigured tools report as failures instead of fake results. Every agent carries an Agent Bill of Materials (Agent-BOM) and a signed version history, and model-integrity and hidden-weight scanning look for tampering before release.

The Problem: AI Adoption Is Blocked by Trust, Not Technology

Security, privacy, and trust stop AI at the boundary. Mission owners and CISOs ask the same questions every time: Will sensitive data (PII, PHI, CUI, secrets, or classified markings) leave our perimeter in a prompt? Can a prompt injection turn a helpful assistant into an insider? Which AI tools are people already using without approval? Will the model drift or hallucinate after an update? When an assessor asks what the AI did and why, can we show the evidence? Without good answers, the safe decision is "no," and the mission keeps doing the work by hand.

The cleared workforce gap. Cleared, certified subject-matter experts are the hardest resource to staff. Every hour an analyst spends assembling evidence, copying findings into a ticket, or writing documentation is an hour not spent on the operation.

AI agents create a new kind of insider risk. An agent can hold valid credentials and still act beyond its mission because of compromised dependencies, malicious instructions, excessive permissions, or unauthorized delegation. Valid credentials do not equal unlimited authority. Agents need the same Zero Trust discipline we apply to people.

Force Multiplier for Cleared and Certified Workforces

The fastest way to add capacity to a cleared team is not to hire faster. It is to stop spending cleared hours on work an agent can do safely. Secure agents help close the cleared-SME gap by taking on repeatable analysis, triage, evidence collection, and documentation, so operators can focus on judgment, decisions, and execution.




Figure 2. Force multiplier. Agents prepare the work; cleared and certified people decide and execute.

- **Augment, not replace.** Agents prepare the work; people decide. Every consequential action stays human-in-the-loop, with government or enterprise decision authority preserved.
- **Better operators, more missions.** Each analyst, engineer, and technician can execute more operations and missions, faster and more consistently, within existing accredited boundaries.

- **Workforce visibility.** The AI SOAR Workforce module maps governed agents to DoD Cyber Workforce Framework (DCWF) work roles under DoD 8140, and shows which roles are covered by people, by governed agents, or by neither.
- **For enterprises, departments, and programs.** Primes and integrators can deliver secure agents inside their own service delivery, and departments and program offices can start with a single project or program and grow from there.

Bring Your Own AI (BYOAI)

Your mission already has AI choices: the commercial LLMs you have contracted, the open-source models your engineers run, the agents your teams built, and the tools and MCP servers those agents call. You should not have to give any of that up to be secure. With ZeroTrusted.ai you bring your own AI, and we secure it, govern it, and make it transparent, without changing it.

You bring	Examples from our documentation	ZeroTrusted.ai adds
Commercial LLMs	ChatGPT, Claude, Gemini, and other major providers, with your own API keys (BYOK)	Inline inspection of every prompt and response, DLP and redaction, policy routing, and audit, while you keep the provider relationship, billing, and data routing
Open-source and local models	Self-hosted models through Ollama, vLLM, LM Studio, llama.cpp, or NVIDIA NIM	First-party, no-egress operation; a hardware-fit analyzer; and a curated open-source catalog that shows country of origin and supply-chain risk
In-house agents and agent frameworks	Internal agents, including agents built on common frameworks such as LangGraph, CrewAI, AutoGen, the OpenAI Agents SDK, and Google ADK	Agent registration, identity, owner, permissions, and autonomy limits; preflight and adversarial testing; and the Agent Mission Trace
Tools, MCP servers, and plugins	Your own tools and connectors, MCP and A2A connections, and the Developer API	Tool-use and function-call allow-lists, argument inspection, revocation on a policy violation, and supply-chain integrity checks

Without changing them. The AI Firewall runs inline, out-of-band, or at the browser edge, so you choose how traffic is inspected. HealthCheck runs alongside your models and agents with no code changes to the application. Agents are registered and tested as they are, and you can assign an approved local or external model to each agent and tenant. When you switch models, your policies, approvals, and evidence stay the same.

Architecture Overview

ZeroTrusted.ai deploys inside your boundary as a control layer between your people, your AI, and your existing systems. Figure 3 shows the logical view. The next section shows how that layer runs from the tactical edge to the enterprise.




Figure 3. Logical architecture. ZeroTrusted.ai sits between your people, your AI, and your existing systems, inside your boundary.

- **AI Firewall and AI Gateway.** Inline, out-of-band, or browser/edge enforcement for prompts, responses, APIs, MCP and A2A traffic, and tool calls. It covers prompt-injection and jailbreak blocking; redaction or blocking of PII, PHI, secrets, classified markings, and custom sensitive terms; tool-use allow-lists; file sanitization; and AI Gateway rules for shadow AI. Shadow AI Protection extends enforcement to the user's browser.
- **AI HealthCheck.** Continuous assurance for LLMs, SLMs, RAG pipelines, agents, and vector databases: drift baselining, bias and fairness, hallucination and factuality regression, prompt-injection posture, and policy-as-code gates in CI/CD. HealthCheck runs alongside your models and agents with no code changes to the application.
- **AI SOAR.** The governed agent workforce and operations control plane: Agent Hub, SWARM missions, Workforce mapping, Autonomous SOC, Deep AI System Testing, Supply Chain Risk, incident response, and compliance automation.
- **Governance and evidence.** Policy lifecycle, AI asset inventory, append-only audit trails, mission traces, and SHA-256 hash-chained, tamper-evident evidence packages.

From the Edge to the Enterprise

Missions do not happen only in the data center. The same AI that supports the enterprise SOC also has to work on a forward-deployed laptop with no reach-back, at a regional operations node with intermittent links, and in the enterprise data center. ZeroTrusted.ai uses one governance and evidence

model at every tier. AI SOAR runs on anything from a small laptop to a large enterprise cluster, within the limits of the selected models, tool workload, storage, network, and available compute.




Figure 4. Edge-to-enterprise architecture. One secure layer and one evidence model from forward-deployed DDIL nodes to enterprise data centers.

- **Tactical and disconnected edge (DDIL, forward-deployed).** A small-footprint install, for example a laptop, a small server, or a lightweight Kubernetes distribution such as K3s, is loaded from an offline bundle. It runs local models and a focused set of agents. AI Firewall policy, human approvals, and the audit trail work locally with no outside network. Evidence is captured in signed, hash-chained packages for later review.
- **Mission and regional nodes.** Kubernetes nodes (RKE2, K3s, Talos, or another hardened distribution) host the mission's agents, tool runners, and models. Lightweight edge collectors in the networks being supported enroll with a one-time token that is exchanged for a scoped API key. Multi-tenant profiles keep units, missions, or customers separate.
- **Enterprise data centers and private clouds.** High-availability Kubernetes or Red Hat OpenShift on-premises, or your own AWS (including GovCloud), Azure, or Google Cloud tenancy, adds GPU capacity, more workers, enterprise models, the Autonomous SOC, and portfolio-level views across tenants.
- **Across every tier.** Identity (including CAC/PIV), policy, agent registration and testing, classification banners, and the evidence chain stay the same. Controlled synchronization moves approved content,

such as offline tool registries and signed evidence packages, between disconnected and connected environments when links allow. Adding capacity does not change the evidence and governance model.

Security, Privacy, and Transparency: Every Time You Go to Work

Security, privacy, and transparency are not features we add at the end. They are the operating conditions for every prompt, every agent, and every mission.

Pillar	What it means in ZeroTrusted.ai
Security	Zero Trust for users, models, and agents. Prompt-injection and jailbreak blocking; tool-use and function-call allow-lists; adversarial ML defense; Agent Escape testing; identity through local accounts, LDAP/AD, OIDC, SAML 2.0, and CAC/PIV; least privilege and autonomy limits for every agent.
Privacy and data protection	PII, PHI, secrets, classified markings, and custom sensitive terms are redacted or blocked before a prompt leaves your perimeter. Secure Chat swaps sensitive values for realistic stand-ins so no model sees the real data. Data stays inside your boundary; with local models, nothing needs to leave the network.
Transparency and audit	Activity History logs every prompt, response, and guardrail resolution. Agent Mission Trace shows every tool call, approval, and timing. Evidence packages are SHA-256 hash-chained and tamper-evident. Failed or unconfigured tools report as failures; missing evidence is never treated as a pass.
Policy and governance	A policy DSL with templates, keyword policies, AI Gateway rules for shadow AI, organization-wide safeguards, AI asset inventory, and an approval workflow for agents, models, and exceptions.
Model-agnostic	Works with commercial LLMs, self-hosted and local models, and agent frameworks, with BYOK for external providers. Change models without changing your security posture.




Figure 5. Agent Mission Trace: every tool call in an agent run is logged with timing, inputs, and outputs (demo environment; identifiers redacted).

Frameworks. AI SOAR includes packaged control mappings for frameworks such as NIST SP 800-53 Rev. 5, the NIST AI RMF, CMMC 2.0, DISA STIG, DoD 8140 (DCWF), ISO/IEC 42001, and the EU AI Act. Candidate NIST and OWASP mappings help organize reviewer-led analysis. Mappings support evidence;

they do not establish certification or regulatory compliance. An assessor, auditor, or authorizing official still makes the determination.

Mission Use Cases

The same secure layer supports very different missions. The workflows below are illustrative applications of documented capabilities, not customer case studies.

Mission	How ZeroTrusted.ai helps
Cyber operations and SOC	The Autonomous SOC correlates SIEM, EDR, identity, cloud, email, and threat-intelligence signals into governed investigation and response workflows, with human approval for high-impact actions. ZeroTrusted.ai does not run your SOC; your team runs its own operations with governed agents, 24x7x365.
Insider threat and UAM	Evidence-readiness assessments within Supply Chain Risk cover User Activity Monitoring (six control areas), Human Insider Threat (six), and AI Agent Insider Threat (ten), across 22 control prompts.
Supply chain risk	Supply Chain Risk Assessment with SBOM tracking per NIST SP 800-161, Agent-BOMs, model-integrity and hidden-weight scanning, and provenance checks for models, tools, MCP servers, plugins, and runtimes.
OT, UAS, maritime, and SOF	OT asset inventory and drift monitoring, and UAS assurance for firmware, C2 links, GPS resilience, Remote ID, and supply-chain provenance, evidence-only by default. Air-gapped installs with local models support disconnected and DDIL operations. The AI CNO Platform supports authorized counter-AI and information-operations missions.
Enterprise IT and AI adoption	AI Firewall and Shadow AI Protection govern employee use of public AI tools. Agents work inside existing service-desk workflows (for example, ServiceNow and Jira) with technician approval gates, and AI FinOps shows cost and risk by user, agent, and model.

Insider risk, done responsibly. The insider-risk assessments evaluate whether monitoring is authorized and evidence exists. They are evidence-readiness assessments, not a live surveillance sensor or an automatic containment engine. Evidence collection and operational response depend on separately authorized tools and integrations. The emphasis is on observable actions and enforceable authority, and high-impact decisions about people remain with authorized human reviewers.

Deployment Options

ZeroTrusted.ai is on-premises first. It is enterprise software that you install and operate inside your environment and authorization boundary. **It is not SaaS.** You supply compute, storage, and AI; ZeroTrusted.ai provides the software.

Option	What it looks like
On-premises and Red Hat	Control plane and tool runners on RHEL, Rocky, Alma, or Red Hat OpenShift. Sensitive data, evidence, models, and credentials stay inside the boundary.
Kubernetes and private cloud	RKE2, K3s, Talos, or another hardened Kubernetes distribution, with separate workers for scanners, agents, evidence, and connectors.

Your own cloud tenancy	AWS (including GovCloud), Microsoft Azure, and Google Cloud on EKS, AKS, and GKE, with the control plane in your chosen boundary.
Tactical edge and DDIL	Small-footprint installs on a laptop or small server from an offline bundle, with local models, a focused agent set, and local approvals and evidence.
Hybrid and air-gapped	Offline bundles, local models, edge collectors, offline tool registries, signed evidence packages, and controlled synchronization for disconnected operations.
Multi-tenant operations	One platform for your organization and your customers, with tenant boundaries preserved, for enterprises, primes, and MSSPs.

The software is built to FedRAMP High security requirements with STIG-hardened Kubernetes configurations. ZeroTrusted.ai does not claim a FedRAMP authorization; it can be added to an existing authorization boundary, and we provide configuration and SSP information to support your own assessment. Classification banners support environments from unclassified and CUI through the highest national-security classification levels; accreditation of that environment remains the customer's.

How to Start

- 1. Start a free trial.** A free 15-day online trial (no credit card required) lets your team explore the products.
- 2. Request a demonstration** against the workflows you care about.
- 3. Run a scoped pilot.** Pick one mission, connect your systems, build and test the agents, and measure results from your own operations, in your environment.
- 4. Buy through Tradewinds or your prime.** ZeroTrusted.ai is Awardable on the CDAO Tradewinds Solutions Marketplace. Federal customers can also work through prime-contractor vehicles.
- 5. Train your team.** ZeroTrusted.ai Academy courses for operators, ISSOs, engineers, and leaders are available online or on-site.

The Bottom Line

AI is going to be part of every mission, either secured, private, and accountable, or as shadow AI and unmanaged agents. Keep the systems you trust. Bring your own AI. Run it from the tactical edge to the enterprise. Build the agents, tools, and operational capabilities your team needs, and keep your people in charge of the mission, with security, privacy, and transparency every time you go to work.

ZeroTrusted.ai . . . The AI/SI Force Multiplier for your enterprise, department, and/or projects/programs

About ZeroTrusted.ai

ZeroTrusted.ai LLC applies Zero Trust to AI users, models, and agents. Its products, AI SOAR, AI Firewall, and AI HealthCheck, let organizations secure the AI they already use and put governed AI agents to work for the mission, inside their own environment.

Contact: waylon@zerotrusted.ai • 703-200-8123 • **Web:** www.zerotrusted.ai

Free trial: identity.zerotrusted.ai/Account/SignUp

ZeroTrusted.ai LLC • 2198 Four Winds Blvd, Kissimmee, FL 34746 • UEI Y8SAH8UAYG27 • CAGE 9EEU1