

# Cognizant Secure AI Services based on Zscaler AI Protect

## Executive overview

AI is your enterprise's greatest competitive opportunity and its fastest-growing security liability.

As autonomous AI agents replace human-directed workflows, the attack surface has expanded beyond what traditional security can see or stop. Each unsanctioned model, unregistered agent or unguarded prompt is a potential breach, operating at machine speed with no perimeter to enforce or deterministic behavior to audit.

Legacy security was built for deterministic software with predictable inputs, auditable logic and gate-based releases. AI systems are probabilistic, context-rich, multiturn and increasingly autonomous: They learn, drift and act on enterprise data beyond what any perimeter control can govern, and one compromised prompt in an agentic workflow can trigger cascading, confidently incorrect actions at scale. Securing AI needs a new operating model: The enterprise must both secure AI systems and use AI to defend the enterprise, including autonomous agentic SOC capabilities.

## Challenges

### IP Leakage and Observability Blindspots

Lack of visibility and real-time observability over prompt interactions, sensitive data, PII, and proprietary IP exposure through public LLMS and APIs

### Data Poisoning and Ingestion Governance

Corrupted training pipelines and RAG knowledge bases produce biased outputs when integrity controls and data-input governance are not applied

### Shadow AI and Governance Gaps

Local AI models, IDE plugins, and browser-based tools create shadow AI that operates outside corporate compliance, governance, auditability, and security controls

### Adversarial Attacks and Runtime Control Gaps

Complex threats like prompt injection, jailbreaking, and agent impersonation bypass traditional security, requiring continuous runtime observability and governance

# Service offering

Securing AI-powered enterprises takes an inline, zero trust framework built for AI workflows: continuous posture management, strict boundary controls for model-to-model and agent-to-agent interactions, and active runtime protection across the full AI lifecycle, from training data ingestion to production deployment, aligned with NIST AI RMF, the EU AI Act, ISO/IEC 42001 and the OWASP Top 10 for LLMs.

Cognizant's offering is structured around a dual mandate, delivered as an integrated, partner-enabled managed service:

| Mandate | Purpose | Outcome |
|---|---|---|
| **Security for AI** | Provide deep observability, and policy governance for AI models, agents, training pipelines, data stores, APIs, endpoints and runtime surfaces | Foundational control layer with continuous runtime observability, automated governance guardrails, eliminating threats, data leakage, and compliance exposure |
| **AI for Security** | Deploy AI-powered detection, investigation and response capabilities within the SOC | Reduced MTTD and MTTR through autonomous agents that triage, correlate and escalate at machine speed |

The offering rests on three Cognizant-originated foundations: secure ADLC, which builds security and governance into every stage of agentic systems; Cognizant Neuro® Cybersecurity, which unifies AI and enterprise signals for faster response; and responsible AI, delivered through Cognizant Trust™ for traceable, audit-ready AI.

Cognizant and Zscaler have built a joint service offering to secure enterprise AI. Zscaler's Zero Trust Exchange™ enforces inline security at every AI interaction, inspecting every prompt, governing every agent and controlling every dataflow in transit. Cognizant Neuro Cybersecurity is the intelligence and orchestration layer, unifying AI security signals, runtime posture and automated response into a single control plane. Together, they deliver enforcement plus intelligence across the full AI lifecycle.

### What this means for your organization:

**Proactive Governance:** Inline, policy-based guardrails governing prompt inputs, data access, and agent workflows in real time.

**Full-Stack AI Observability:** Real-time visibility and deep transaction logging across every AI app, model, prompt, response, and autonomous agent interaction—eliminating shadow AI blind spots.

**Audit-Ready Assurance:** Continuous compliance reporting for EU AI Act, NIST AI RMF, ISO 42001, and enterprise acceptable-use policies.

## Secure model service

Gain full visibility and control of enterprise AI assets. Using Zscaler AI Asset Management and AI Security Posture Management (AI-SPM), Cognizant discovers LLMs, agents, pipelines and MCP servers, and eliminates shadow AI:

- AI asset discovery and AI-BOM generation, with continuous AI-SPM monitoring for policy violations, misconfigurations and supply-chain risks
- Model signing, integrity verification, supply-chain risk management and governance alignment to NIST AI RMF

## Data and privacy protection service

Sensitive data stays protected across every AI interaction. Cognizant deploys Zscaler AI Guard to inspect prompts and responses inline, blocking data leakage, prompt injection, jailbreaking and toxic content through ZIA proxy mode and API-based custom LLM integration:

- Inline DLP for prompts and responses; PII, PHI and IP redaction; data lineage tracking across multicloud AI
- EU AI Act and GDPR compliance controls for AI data handling

## AI runtime protection service

Zero trust enforcement extends to the device and runtime layer. With Zscaler Endpoint AI Security, Cognizant brings AI developer tools, IDE plugins, local models and browser extensions under enterprise control:

- Real-time monitoring of inference endpoints and model-serving infrastructure, with visibility into IDE plugins, local LLM runtimes and browser AI extensions
- Runtime prompt injection and jailbreak detection; anomaly detection for agent drift and unexpected tool calls

## AI-apps DevOps security service

Adversarial testing becomes part of AI application delivery. Cognizant integrates Zscaler AI Red Teaming into CI/CD pipelines, running 5,000+ purpose-built attack simulations across 25+ risk categories with dynamic prompt hardening:

- Automated AI red teaming in CI/CD pipelines, with dynamic prompt and model hardening recommendations
- Secure AI code review, SAST/DAST, secrets scanning and AI/ML dependency analysis; OWASP Top 10 for LLMs coverage

## AI threat defense service

Agent-to-agent and MCP communications are governed through an agent registry. Cognizant uses Zscaler AI Gateway and AI Access Graph to map relationships among human identities, AI agents, applications and data sources:

- Agentic governance (agent registry, A2A and MCP controls); AI identity graph for overprivileged agents and toxic permission pairings
- AI-powered SOC triage, correlation and escalation at machine speed, with ThreatLabz threat intelligence and incident response playbooks for AI-native attacks

## AI GRC service

Continuous AI governance reporting and compliance assurance across major AI and privacy frameworks, delivered by Cognizant and supported by Zscaler's data residency and governance capabilities:

- NIST AI RMF (Govern, Map, Measure, Manage) and OWASP Top 10 for LLMs alignment
- EU AI Act risk-tier classification and controls, ISO/IEC 42001 certification support, and GDPR and CCPA privacy controls
- Continuous Governance & Observability Audits: Real-time policy validation, prompt-level telemetry tracking, and automated reporting against EU AI Act, NIST AI RMF, and internal acceptable-use policies.

# Enterprise AI security use cases

| Purpose | Outcome |
|---|---|
| Shadow AI discovery and governance | AI-BOM creation, unauthorized AI asset discovery and access policy enforcement |
| LLM data leakage prevention | Blocking PII/IP leakage in prompts and enforcing DLP controls |
| Agentic workflow governance | Agent registry, cloud data stores, APIs, internal databases and privilege-escalation prevention |
| AI red teaming | Automated attack simulations, risk categorization and preproduction validation |
| Developer AI tool endpoint control | Visibility and control for IDE plugins, coding assistants, local AI runtimes and browser tools |
| Prompt injection attack detection | Detection and blocking of malicious documents, RAG pipeline abuse, inference-time attacks and data fetching risks |
| AI regulatory compliance reporting | EU AI Act, NIST AI RMF, dashboards and audit-ready evidence |
| Agentic SOC autonomous threat response | Autonomous triage, correlation, escalation and reduced MTTR |

# Engagement methodology

Cognizant delivers the offering in four phases with defined entry and exit criteria, continuous KPI measurement and control-framework alignment:

### Assess (identify):

Establish clarity for informed decisions—AI security posture assessment, risk identification and logging, regulatory and standards mapping, AI/ML threat modeling and red teaming preparation

### Architect (prevent and protect):

Build scalable, resilient security—secure AI architecture design, technology and platform selection, governance and RACI structure, red teaming integration and ethical AI compliance

### Act (detect and respond):

Enable rapid threat and incident response—deployment and onboarding, control testing and validation, incident response and breach readiness, red teaming execution and recovery planning

### Assure (govern):

Ensure compliance and resilience—continuous compliance and audit, security awareness and training, performance and capacity management, periodic red teaming and KPI/SLA tracking and reporting

# Why choose Cognizant × Zscaler?

Enterprise AI is already operating in your business, making decisions, accessing sensitive data and acting for your users. The question is whether you can secure it at the speed and scale you need.

Cognizant and Zscaler have launched a joint Center of Excellence (CoE) under the **ZScaler Project AI-Guardian** initiative. Cognizant Neuro provides the intelligence layer, consolidating AI security posture, runtime signals and automated response into one governed control plane, with Cognizant Trust keeping every AI system accountable, traceable and audit-ready.

Zscaler provides the enforcement layer: The Zero Trust Exchange inspects and governs every AI interaction inline—with no lateral movement and no implicit trust—backed by ThreatLabz intelligence from over 500 billion daily transactions.

Together, they help global enterprises move fast on AI without creating the risks that undo its value. Learn more at cognizant.com/zscaler.

For inquiries: **cybersecuritypulse@cognizant.com**



Cognizant Neuro® and Cognizant Trust™ are proprietary platforms of Cognizant. Zscaler, Zero Trust Exchange™ and all Zscaler product names are trademarks of Zscaler, Inc.

cognizant®

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

**World Headquarters**
300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243

**European Headquarters**
280 Bishopsgate
London
EC2M 4AG, England
Phone: +44 (01) 020 7297 7600

**India Corporate Office**
Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park - Chengalpattu District
Chennai 603103, Tamil Nadu
Phone: 1800 208 6999

**APAC Headquarters**
1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower
Singapore 138542
Phone: +65 6812 4000

© Copyright 2026, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice.
All other trademarks mentioned herein are the property of their respective owners.

WF5471550