

Cognizant's Secure AI Services powered by Palo Alto Networks

Overview

AI is rapidly becoming central to enterprise strategy, powering automation, decision intelligence and personalized experiences at scale. Adoption is accelerating at an unprecedented pace, global AI spending is projected to reach $632 billion by 2028 according to IDC[1], while enterprises are already putting 11x more AI models into production year-over-year as reported by Databricks[2]. Gartner further predicts that 80% of generative AI business applications will be built on enterprise data platforms by 2028, underscoring how deeply AI will be embedded into core systems and workflows[3].

As organizations scale AI, they increasingly depend on third-party LLMs, API-driven AI services and Model Context Protocol (MCP) servers. While these accelerate innovation, they also expand the attack surface introducing risks such as data leakage, model inversion, unauthorized retention, supply chain compromise and unpredictable model behaviour. Enterprises must also defend against AI-specific threats including prompt injection, data poisoning, adversarial manipulation and jailbreak attacks.

Key Challenges:

Securing AI now requires more than traditional infrastructure controls. Traditional security was built for deterministic software and release-gate controls. AI systems are probabilistic and context-driven they learn, drift, and depend on enterprise context that can be poisoned, manipulated, or exfiltrated. As agents move into core workflows, a single compromised context or prompt can trigger confidently wrong actions at scale—creating systemic risk that legacy, fragmented security wasn't designed to prevent.

- Rapid growth of enterprise AI applications is significantly expanding the organizational attack surface.
- New AI-specific threats such as prompt injection, data poisoning, model tampering, and agent impersonation are emerging.
- Autonomous and agentic AI systems increase risk through uncontrolled interactions with browsers, APIs, databases, and tools.
- AI is being actively weaponized for high-impact fraud, deepfake impersonation, phishing, and credential theft.
- Insecure APIs, credentials, and misconfigured agents enable hijacking of AI services and unauthorized actions.
- Lack of runtime visibility and monitoring makes stealthy AI-driven attacks difficult to detect and respond to.
- AI security incidents pose significant business risk, including financial loss, compliance penalties, and reputational damage.
- AI Security and Responsible AI converge into one mandate, how to design and upskill the new CISO office



If a model is secure but biased, it becomes a liability. If it is ethical but vulnerable to injection, it becomes a risk.

Securing the AI Powered Enterprise

Organizations must adopt AI-specific governance, Zero Trust for model interactions, secure integration patterns for external LLM ecosystems and continuous monitoring of model behaviour. Aligning with frameworks like NIST AI RMF and ISO/IEC 42001 helps ensure responsible and compliant adoption.

By embedding security across the entire AI lifecycle from data ingestion to deployment, and across the entire AI stack from AI Infrastructure to AI applications and services, enterprises can safeguard intellectual property, deploying responsible & AI governance practices, maintain regulatory confidence, and unlock the full strategic value of AI while mitigating emerging risks.

AI introduces a dual mandate for cybersecurity: accelerating defense with AI (AI for Security) while securing AI systems themselves (Secure AI Services). These mandates are inseparable, enterprises cannot protect the AI stack without the speed, consolidation, and orchestration that AI driven defence delivers. Strengthening both sides of this mandate is essential for realizing AI's full strategic value while safeguarding enterprise systems, data, and operations.

Key Foundation Principles

We actively align with leading global AI standards and embed their principles into our lifecycle controls and governance practices. Our adaptive governance model provides continuous oversight and real-time assurance, so systems remain accountable, consistent and aligned as technologies and expectations evolve.

Our Secure AI Services governance model is built on three foundations: Secure ADLC (build-time assurance), Neuro® Cybersecurity (consolidated control plane), and Responsible AI (trust & accountability).

Secure Agent Development Lifecycle (ADLC)

Given that Agentic systems are non-deterministic—they are autonomous, adaptive, goal-oriented and capable of proactive behaviour, we address these challenges by providing a repeatable, outcome-driven framework that enables agents to be purpose-built, compliant and optimized for enterprise success.

Our Secure ADLC process is designed for creating, testing, and monitoring non-deterministic, autonomous AI agents that reason and act.

Providing security and governance integrated across design, build, test, deploy, and change of agentic systems.

Neuro® Cybersecurity

As threats become increasingly sophisticated and pervasive, organizations need to stay ahead of bad actors. Siloed operations and risk mitigation decisions often lag the need to act in real time, leading to higher risk of exposure.

Need of the hour is AI-enabled, easy to consume UI that enables decisions across the enterprise.

Cognizant Neuro® Cybersecurity revolutionizes risk management, threat and vulnerability management and compliance assurance—by integrating and orchestrating point solutions to ensure comprehensive security coverage using AI.

It is the "unification layer" for AI security posture + runtime signals + response actions

- Consolidates AI signals (model, prompt, RAG, agent actions) with enterprise signals (identity, cloud, network, endpoint)
- Enables faster correlation and response orchestration across tools
- Produces audit-ready evidence and posture dashboards for AI risk

Responsible AI

Our focus is on operationalizing trust in AI systems through continuous assurance, control, and accountability across the lifecycle.

In addition to offer a compliance layer, Responsible AI functions as a continuous trust layer that:

- Ensures AI systems behave as intended in real-world conditions
- Provides traceability of decisions, prompts, and actions
- Enables policy enforcement and behavioural guardrails for AI agents
- Continuously validates performance, bias, drift, and safety
- Generates audit-ready evidence aligned to ISO 42001, NIST AI RMF, and regulatory requirements
- Through **Cognizant Trust™**, our **Responsible AI Assurance** framework, organizations gain:
- AI system inventory and risk classification
- Continuous compliance monitoring and policy validation
- Incident tracking and remediation workflows
- Real-time dashboards and metrics for AI system performance and trust

This enables organizations to move from static compliance to continuous, provable trust in AI systems.

Service Offering

Cognizant's Secure AI Services offering leverages Prisma AIRS enabling AI discovery, inventory, and continuous risk scoring for enterprise-built LLM applications, RAG pipelines, autonomous agents, and MCP servers. Cortex platform provides AI-specific threat detection, automated testing, adversarial evaluation, and runtime protection to secure models, data flows, and agent behaviour. Prisma Access, NGFW, Strata Cloud Manager and Prisma Browser secures AI usage for third-party AI services.

Cognizant's Secure AI Services approach provides comprehensive, end-to-end protection by strengthening the core pillars across the AI lifecycle: Model Security, Data Security, Runtime Security, AI DevOps Security, Identity Protection, Threat Defence and GRC Security integrated with Palo Alto Networks' Prisma AIRS, Cortex Cloud, Prisma Access, NGFW, and Strata Cloud Manager.

1. ***Secure Model Service*** - Safeguards enterprise AI models against extraction, inversion, poisoning, and adversarial attacks leveraging:

- Prisma AIRS for AI posture management and model risk scoring
- Scans for malicious scripts, unsafe code, deserialization attacks, and neural backdoors embedded within model files.
- Validation of third-party and open-source models (e.g., from Hugging Face) to prevent model tampering or poisoned data from entering your environment.
- enabling checks for hardcoded secrets or credentials accidentally left inside model artifacts.

2. ***Data & Privacy Protection Service*** - Protects sensitive data used across AI workloads through:

- Leveraging Prisma Access & NGFW for AI-aware DLP and safe outbound/prompt inspection
- Prisma Browser enables securing browser-based AI interactions and preventing data leakage

3. ***AI Runtime Protection Service*** - Defends AI systems from prompt injection, jailbreaks, unsafe tool execution, and malicious agent behaviour

- Cortex Cloud facilitates real-time AI threat detection
- Prisma AIRS enables guardrail-based runtime policy enforcement



4. ***AI Apps DevOps Security Service:*** Secures AI development lifecycle via APIs and AI Red teaming.

- facilitating Automated Penetration Testing against AI Systems using adaptive Agents

5. ***AI Threat Defence Services*** – Delivers continuous detection, monitoring, and response for AI-specific threats through:

- Cortex Cloud enables AI-aware analytics and behavioural threat models
- Strata Cloud Manager facilitates unified threat visibility and policy governance

6. ***AI GRC Services*** - Supports Responsible AI governance, continuous compliance, and audit readiness with:

- Strata Cloud Manager for centralized policy and governance controls
- Prisma AIRS for AI risk scoring, compliance mappings, and control insights

Together, these six pillars deliver secure-by-design AI, continuous risk visibility, and regulatory confidence across enterprise AI initiatives. Cognizant's Secure AI Services services enable organizations to scale trusted AI systems while protecting models, data, identities, and agentic workflows using Microsoft's industry-leading security platforms.

Engagement Methodology:



Use Cases:

- Prevent Deepfake-Driven Fraud and Impersonation
- Protect AI Models from Tampering and Theft
- Defend Against Prompt Injection and Jailbreak Attacks
- Secure AI-Driven Data Pipelines
- Safeguard AI Agents and Autonomous Workflows
- Secure AI Supply Chain and Third-Party Dependencies
- Enable AI-Specific Threat Detection and Incident Response
- Ensure Regulatory Compliance and Audit Readiness for AI
- Validate AI Systems Through Red Teaming and Adversarial Testing
- Establish Continuous AI Security Posture Management
- AI runtime protection

Why Choose Cognizant

Cognizant helps enterprises turn AI ambition into real, scalable business outcomes. Our approach goes beyond experimentation to embed AI deeply into core operations, decision-making and customer experiences. By combining deep industry expertise, proprietary AI platforms and responsible engineering, we deliver AI solutions that are contextual, secure and built for enterprise scale.

At the heart of this approach is Cognizant Neuro®, our enterprise AI platform suite that acts as a unification layer for AI security posture, runtime signals and automated response actions accelerating the journey from identifying high-value use cases to deploying and governing AI and agentic solutions in production. Our partnership with Palo Alto Networks strengthens this capability by bringing industry-leading AI-aware security platforms Prisma AIRS, Cortex Cloud, Prisma Access, NGFW and Strata Cloud Manager into our architecture, enabling enterprises to benefit from seamless enforcement, continuous posture monitoring and end-to-end protection across both first-party and third-party AI applications.

Cognizant's human-centred AI philosophy ensures technology augments human intelligence, drives productivity and enables new ways of working. Supported by a strong partner ecosystem including Palo Alto Networks and a platform-agnostic strategy, we deliver future-ready AI solutions without vendor lock-in. With a relentless focus on outcomes, trust and scale, Cognizant enables organizations to unlock sustained value from AI responsibly, securely and at speed. To know more, please contact- cybersecuritypulse@cognizant.com

Source

1. Worldwide Spending on Artificial Intelligence Forecast to Reach $632 Billion in 2028, According to a New IDC Spending Guide
2. State of AI: Enterprise Adoption & Growth Trends | Databricks Blog
3. Gartner Predicts by 2028, 80% of GenAI Business Apps Will Be Developed on Existing Data Management Platforms

Cognizant helps engineer modern businesses by helping to modernize technology, reimagine processes and transform experiences so they can stay ahead in our fast-changing world. To see how Cognizant is improving everyday life, visit them at **www.cognizant.com** or across their socials **@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
Toll Free: +1 888 937 3277

European Headquarters
280 Bishopsgate
London
EC2M 4RB England
Tel: +44 (01) 020 7297 7600

India Operations Headquarters
5/535, Okkiam Thoraipakkam,
Old Mahabalipuram Road,
Chennai 600 096
Tel: 1-800-208-6999
Fax: +91 (01) 44 4209 6060

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

© Copyright 2025–2027, 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 here in are the property of their respective owners.

cognizant®