THE INTELLIGENT STACK:

SECURING AI IN MODERN SOFTWARE ARCHITECTURE

A strategic roadmap for leaders navigating the transformed software stack landscape, where AI's explosive growth demands new agentic security paradigms to harness intelligent potential while mitigating emerging risks.

Eva Chen

CEO and Co-Founder

EXECUTIVE SUMMARY

Artificial intelligence *(AI)* has rapidly evolved from a niche capability to a cornerstone of modern business strategy. It's no longer just about automation or analytics; AI is now central to how companies innovate, compete, and grow. At the very heart of this transformation are two critical assets: data, which fuels insight, and intelligence, the ability to act on that insight in real time. Together, data and intelligence are reshaping decision-making, operations, and value creation across every industry.

This shift is being driven by an unprecedented surge in data, breakthroughs in AI technologies, and increasing pressure on businesses to operate with speed, precision, adaptability, and compliance. But as AI becomes embedded in everything from product development and customer engagement to supply chain management and sales, it also introduces new exposure to risk, such as model manipulation, data leakage, and expanded attack surfaces.

This industry briefing examines the evolution of the software stack that's being driven by the explosion of AI, the new cybersecurity challenges that come with it, and the strategic considerations leaders must weigh to ensure their organizations can harness AI's potential securely, responsibly, and competitively.

TABLE OF CONTENTS

PART 1:

A FUNDAMENTAL SHIFT IN THE SOFTWARE STACK

The evolution from traditional to AI-native software stacks introduces a profound architectural shift—one that redefines not only how applications are built and run, but how they must be secured in an increasingly intelligent, dynamic, and data-driven world.

At the heart of today's software transformation is a fundamental shift: AI has become the new foundation of the modern software stack, and at the core of AI lies intelligence derived from data.

This shift marks a profound evolution in how software is built, operated, and secured. Unlike traditional architectures—where static code and predetermined logic ruled—AI-native systems rely on dynamic, data-driven models that learn, adapt, and evolve.

Among all the changes within this new stack, the transformation of the data layer is the most critical. It is no longer just about storing or processing data—it's about extracting intelligence in real time to drive autonomous decisions and shape outcomes. As such, the integrity, availability, and security of data have become not only a technical concern but a foundational imperative for every enterprise embracing AI.

This transformation of the data layer sets the stage for a broader evolution across the entire software landscape. As AI becomes the new foundation, we are seeing a move away from traditional, function-based application stacks toward intelligence-driven architectures that can dynamically adapt to business context with real-time insights informing human-made decisions. This progression is not just a technical upgrade—it fundamentally reshapes the role of technology in business, compelling organizations to rethink how systems are architected, how decisions are made, and how value is delivered at scale.

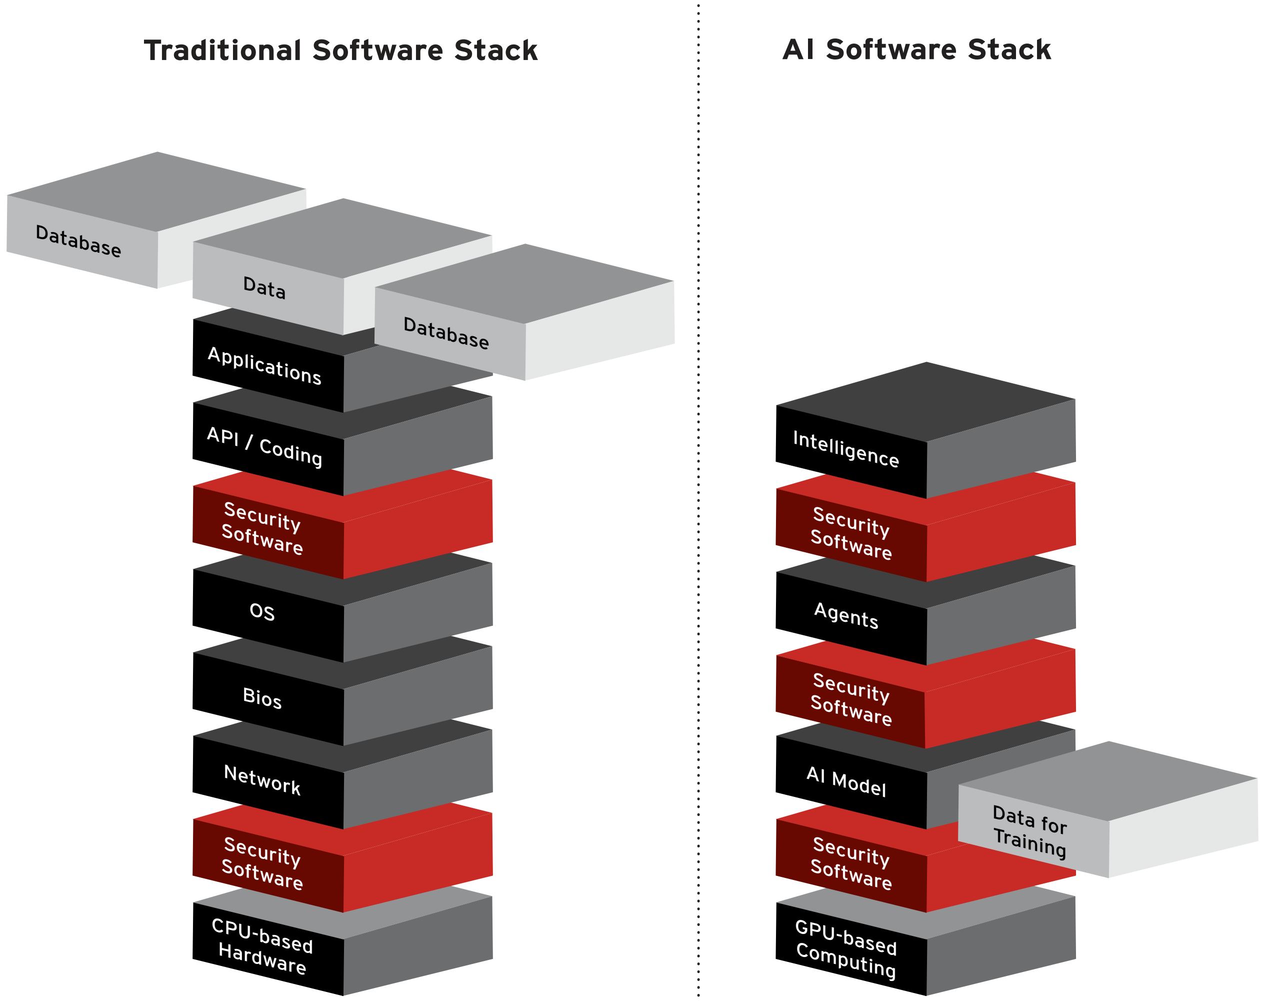

This diagram illustrates the transformation from a traditional software stack to a modern AI software stack, highlighting how the shift from CPU-based to GPU-accelerated architectures fundamentally changes the structure, function, and security needs of enterprise environments.

Traditional vs. AI software stacks

The traditional software stack is centered around CPU-based computing, where a stable operating system layer supports deterministic, rule-based applications. These systems primarily process structured data and operate in relatively static environments. Security in this model focuses on the OS, network, applications, and user access—relying on perimeter defenses, patching, and protection against known threats.

In contrast, the AI software stack represents a fundamental shift in architecture and operation. It leverages GPU-based infrastructure for the high-throughput, parallel processing needed to power AI workloads. This modern stack introduces entirely new layers—such as machine learning models, autonomous agents, and orchestration systems designed to interpret data, adapt to changing conditions, and provide real-time information to assist human decision making.

The critical difference lies in how these stacks are built and how they evolve. Traditional stacks are code-driven, built with explicit logic and predictable workflows. AI stacks are data-driven, operating with fluid, self-optimizing components that continuously ingest unstructured data, learn, and adapt. This introduces new security requirements: ensuring model integrity, safeguarding sensitive training data, and monitoring dynamic behaviors—all while protecting a more complex and fast-changing infrastructure.

A new cybersecurity imperative

In terms of cybersecurity, these differences create more attack surfaces, as the AI stack is inherently more complex and often spans across multiple environments, including cloud, hybrid, and on-premises setups. Traditional security frameworks are insufficient for AI, as they are not designed to handle the unique risks of AI models, such as model poisoning, data leakage, and adversarial attacks.

The AI stack's dynamic nature requires a more proactive, adaptive security approach that can continuously monitor and safeguard against evolving threats. As AI becomes embedded into critical business processes, security measures must be designed to protect not just static applications and data, but also the continuous flow of data, the integrity of AI models, and the underlying infrastructure.

Security must now extend beyond traditional control points to protect this new attack surface area. In the AI stack, security must address:

- Data pipelines and training data integrity *(to prevent poisoning or leakage)*
- Model behavior and access *(to guard against misuse or manipulation)*
- Microservices and containers *(to secure modular, distributed architectures)*
- GPU workloads and APIs *(to ensure visibility into real-time AI operations)*
- Intelligent agents *(to manage safe decision-making and avoid unintended outcomes)*

This transformation underscores a critical shift: cybersecurity can no longer be reactive or isolated to the endpoint or perimeter. In the AI software stack, security must be proactive, contextual, and embedded at every layer—from hardware to intelligence—protecting the full spectrum of risk introduced by intelligent systems. This is where enterprise cybersecurity platforms like Trend Vision One™ play a pivotal role, providing comprehensive, AI-aware protection that evolves with the software and infrastructure itself.

PART 2:

CYBERSECURITY PLAYS A UNIQUE, MULTI-LAYERED ROLE

As AI reshapes the software stack, cybersecurity must evolve from a fragmented set of tools into a unified, adaptive platform that protects every layer—from user interactions to infrastructure—ensuring secure, scalable innovation in an increasingly intelligent and dynamic environment.

		SECURITY CHALLENGES	SECURITY CONTROLS	TREND VISION ONE
Users		Insecure design and mismanagement leading to sensitive data exposure by AI	AI application access control and protect local AI application configurations	ZTSA AI service ZTSA access control Deepfake detection AI app guard
Data		Sensitive information blind spots	Data security	Data security posture management
Models	o1	Model poisoning and improper model usage	Implement guardrails for AI API's request/ prompt *(inbound)* and responses *(outbound)*	ZTSA AI service ZTSA access control
Microservices		Vulnerabilities in AI supply chains and microservices architecture	Security validation on CI/CD pipeline and implement container controls	Code security Container security
Infrastructure	aws	Security risks in AI model deployment and resource exhaustion attacks	Infrastructure posture management	AI-SPM API risk AI-DR
Network		Exploiting vulnerabilities in AI infrastructure and hybrid cloud environments	Network security	Network IDS/IPS

Unlike other software components, cybersecurity spans and protects multiple layers of the stack. As software stacks evolve, from foundational infrastructure and APIs to applications, data, models, and user access, security frameworks must also advance to protect every layer effectively.

Users are the first line of exposure in AI systems. Poorly managed access, misused prompts, and insecure user configurations can lead to data leakage and create exploitable vulnerabilities. Strong AI application access controls and protection of user configurations are critical to mitigating these risks. Data is foundational—if compromised or poisoned, it can undermine entire AI models. Securing data itself must be a primary focus. As organizations move toward microservices, vulnerabilities in AI supply chains and containerized architectures also require attention. Security validation within continuous integration and continuous delivery *(CI/CD)* pipelines and container admission control is essential to ensure safe deployment. AI models must be safeguarded from threats like poisoning and improper API usage. This can be addressed with guardrails that monitor and enforce security for both inbound requests and outbound responses. The infrastructure powering AI—including cloud environments like AWS, Microsoft Azure, and hybrid cloud—faces risks from misconfigurations and resource exhaustion. Infrastructure posture management provides continuous visibility and remediation to proactively eliminate vulnerabilities before they are exploited. Finally, the network layer—a critical point of entry—must be protected through in-line detection to defend against exploits, including zero-days.

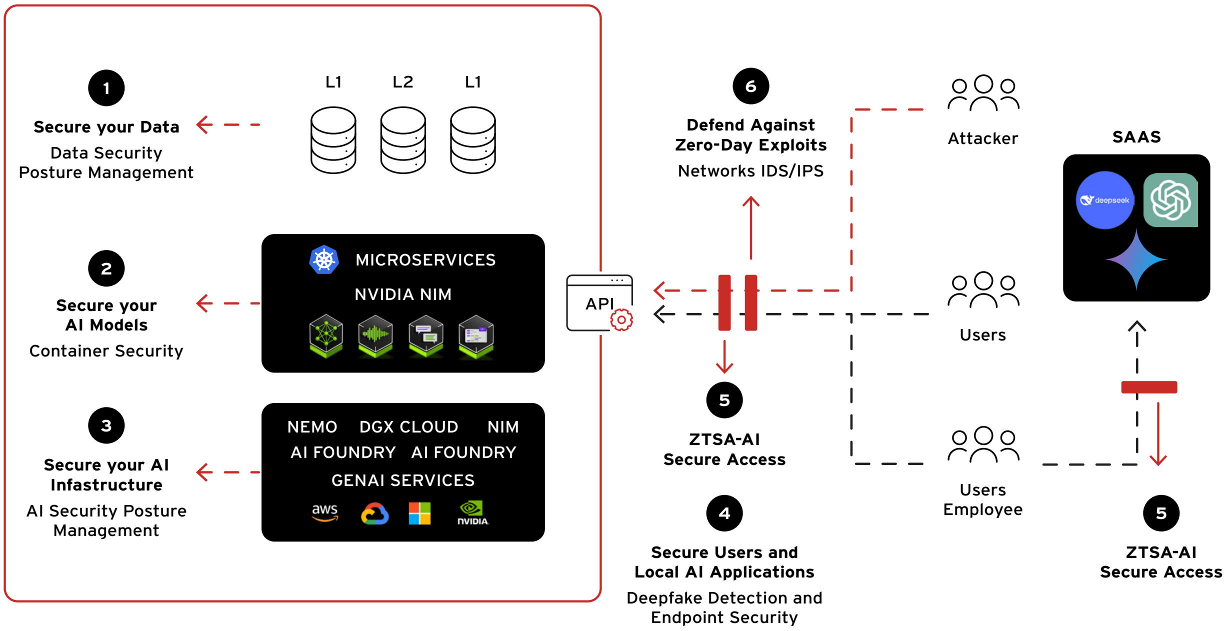

This diagram illustrates how organizations can secure their AI stack end-to-end—from safeguarding sensitive data and protecting AI models, to securing infrastructure, controlling AI tool usage, and defending against zero-day exploits—ensuring their AI investments are both safe and resilient.

As AI-driven transformation accelerates, security leaders face mounting pressure to secure increasingly complex, dynamic environments. Traditional reactive models are no longer sufficient. What's needed is an AI-powered enterprise cybersecurity platform—one that centralizes cyber risk exposure management, streamlines security operations, and delivers robust layered protection across the full stack. This provides the visibility needed to understand the complete attack surface, the context to prioritize response, and the agility to adapt and mitigate in real time.

Organizations need a solution that not only delivers end-to-end protection, but also meets them wherever they are, whether operating in public cloud, hybrid, private, on-premises, or air-gapped environments. Security platforms must be flexible enough to adapt to a wide range of infrastructure realities while maintaining consistent controls and visibility across all layers. The ideal platform should offer comprehensive data sovereignty, GPU-accelerated threat detection, and seamless integration of hardware, software, and support. Just as importantly, it should unify protection across endpoints, cloud, identity, and network, removing the complexity of managing disconnected tools and optimizing cybersecurity spending.

For security leaders tasked with enabling safe AI adoption, this approach is not just an operational advantage, it's a strategic imperative to stay ahead of evolving risks while aligning security with innovation.

PART 3:

INTELLIGENCE-DRIVEN SECURITY IS PROACTIVE, NOT REACTIVE

Meeting the demands of AI-driven environments requires a cybersecurity platform built for scale, speed, and intelligence—one that delivers continuous, context-aware protection across diverse infrastructures while aligning security outcomes with business goals.

As organizations embed AI deeper into core business operations, cybersecurity must become equally intelligent, more proactive, context-aware, and tailored to industry-specific needs. Generative AI *(GenAI)* has already reshaped cybersecurity into an intelligence-driven discipline, and with the advent of agentic AI frameworks, we can now go even further. One of the most powerful applications of this evolution is the use of cybersecurity digital twins, virtual replicas of real-world environments designed for continuous, autonomous adversary simulation and defense validation.

These digital twins allow organizations to use AI agents to safely mimic objective-focused, real-world adversaries, identify and prioritize risk exposure in near real time, and continuously validate and improve their detections, controls, and policies. Crucially, these simulations are customized to the unique operational priorities and attack surfaces of each sector. Whether it's a factory floor, financial institution, or public agency, ensuring defensive strategies are aligned to the realities of each environment is crucial to becoming more proactive and resilient. By automating insight delivery and actionability, this approach turns digital twins into living, learning components of a dynamic, continuously improving cybersecurity posture. This empowers organizations to shift from one-size-fits-all security to proactive, customized protection that aligns with the realities of their specific industry.

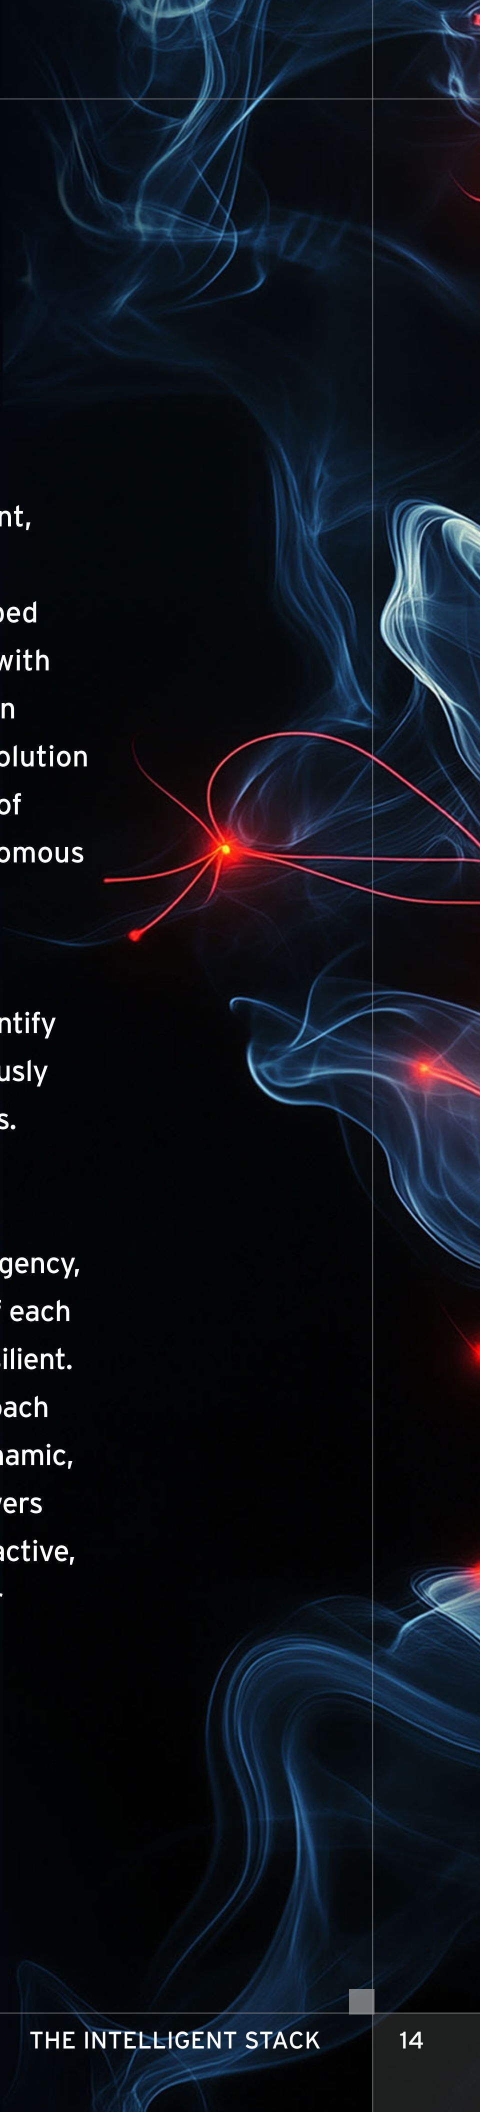

Agentic AI powers sector-specific defense

Using digital twins, domain-specific intelligence, internal knowledge, and standard operating procedures *(SOP)*, agentic AI can perform reasoning, planning, execution, and self-learning to continuously enhance cybersecurity unique to each industry. This allows for dynamic adaptation of controls and responses, raising the bar for resilience and response speed.

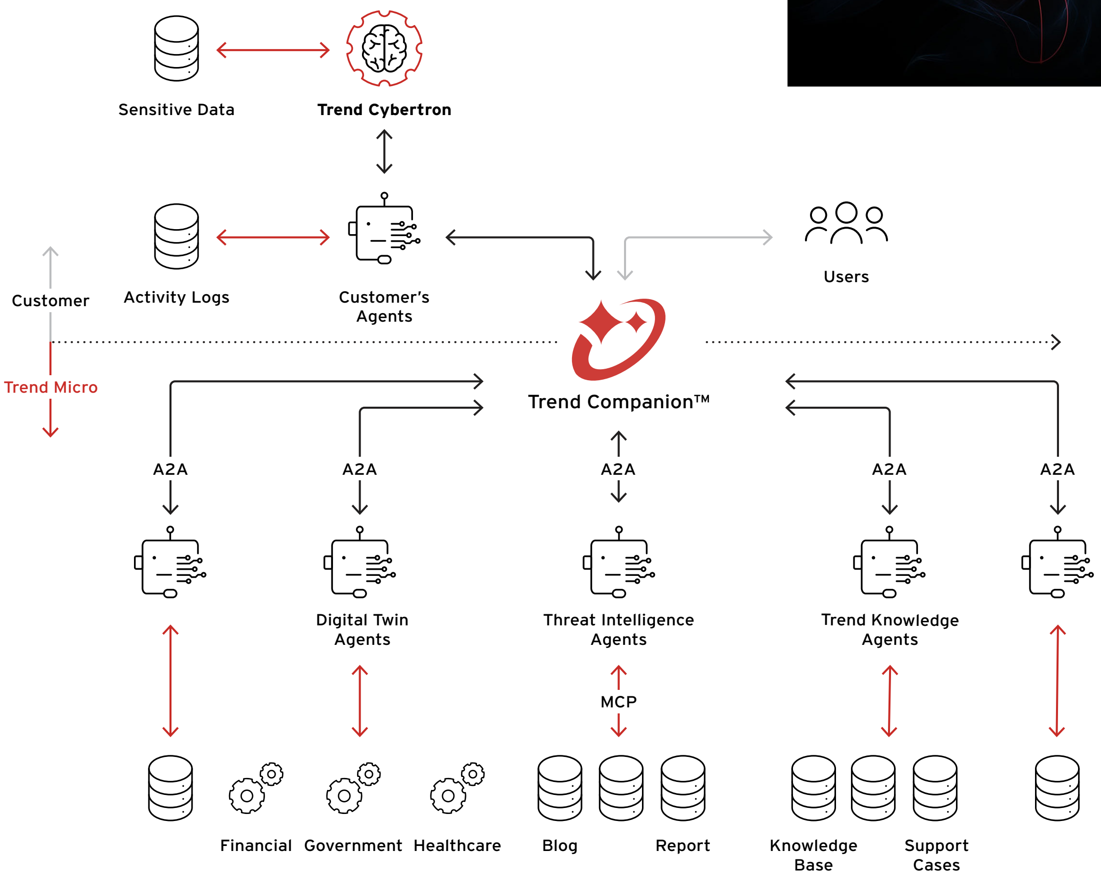

A2A: Agent-to-Agent protocol **MCP:** Model Context Protocol

The diagram above introduces three core types of AI agents, each representing a fundamental component within the agentic AI architecture of Trend Vision One. These components—threat intel, Trend Micro "know-how", and customer-specific insight—form the foundation of an intelligent, coordinated system designed to streamline operational workflows and elevate security outcomes.

Threat intelligence

This agent, which is powered by over 250 million global sensors, acts as a virtual threat intelligence analyst, responsible for continuously ingesting, organizing, and presenting detailed, actionable threat information relevant to the evolving cyber landscape. It specializes in curating high-fidelity threat insights and makes this knowledge available to other agents—such as digital twins—to support informed decisions.

Agentic AI for adaptive defense at the speed of threats

Trend knowledge

This agent acts as a virtual embodiment of our more than 35 years of cybersecurity expertise at Trend Micro. It consolidates and delivers deep product-specific knowledge, including defensive strategies, optimal configurations, tuning guidance, and best practices across Trend Vision One. This agent serves as the go-to authority on how to most effectively operationalize the platform for maximum security value. By distilling decades of expert insight into a structured, accessible AI agent, this agent ensures that security suggestion and improvement is rooted in proven, product-aligned defense intelligence which helps customers get the most out of their existing investments.

Digital twin

This agent simulates the customer's environment and risk profile in real time using synthetic data modeled after actual conditions. By running virtual attack scenarios, it proactively uncovers vulnerabilities, prioritizes exposures, and recommends or initiates mitigation actions—often before threats can materialize in production. Fully powered by agentic AI, the agent autonomously generates risk reports, applies virtual patches, and adapts to changes in the environment. It brings proactive security to life, transforming simulated insight into real-world protection without disrupting operations.

When seamlessly integrated, these three agents create a coordinated intelligence layer that transforms how cybersecurity is operationalized. By combining real-time threat visibility, decades of expert security strategies, and a dynamic understanding of the customer's environment, the system can continuously learn, adapt, and act with precision. This agentic AI architecture is not just about automation—it's about embedding intelligence into every facet of security operations. For organizations looking to secure AI-driven, hybrid environments at scale, this model offers a path to smarter, faster, and more proactive protection tailored to their unique risk landscape.

PART 4:

REDEFINING CYBERSECURITY FOR THE AGE OF INTELLIGENT SOFTWARE

As AI-driven innovation accelerates across sectors, cybersecurity platforms must go beyond one-size-fits-all protection—delivering intelligent, industry-specific solutions that embed security into the very fabric of modern infrastructure, from sovereign cloud to AI factories.

TREND MICRO™ Vision One

AI-Powered Enterprise
Cybersecurity Platform

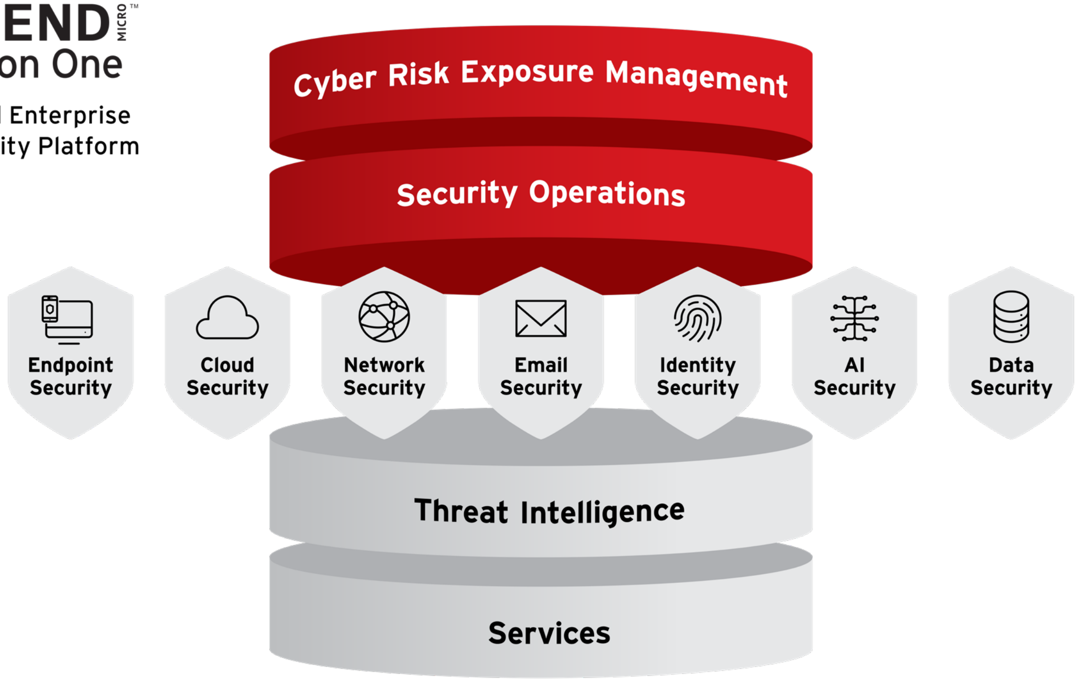

As software stacks become increasingly intelligent, driven by AI, automation, and sector-specific innovation, organizations should expect more from their cybersecurity platforms. It's no longer enough to rely on generalized, reactive tools. Today's security solutions must enable a proactive approach and be purpose-built to support industry-specific environments like AI-enabled manufacturing, digital healthcare systems, and government infrastructure.

Leading enterprise cybersecurity platforms like Trend Vision One are evolving to meet these demands by offering deep integration, sector-tailored controls, and proactive protection for complex and dynamic environments.

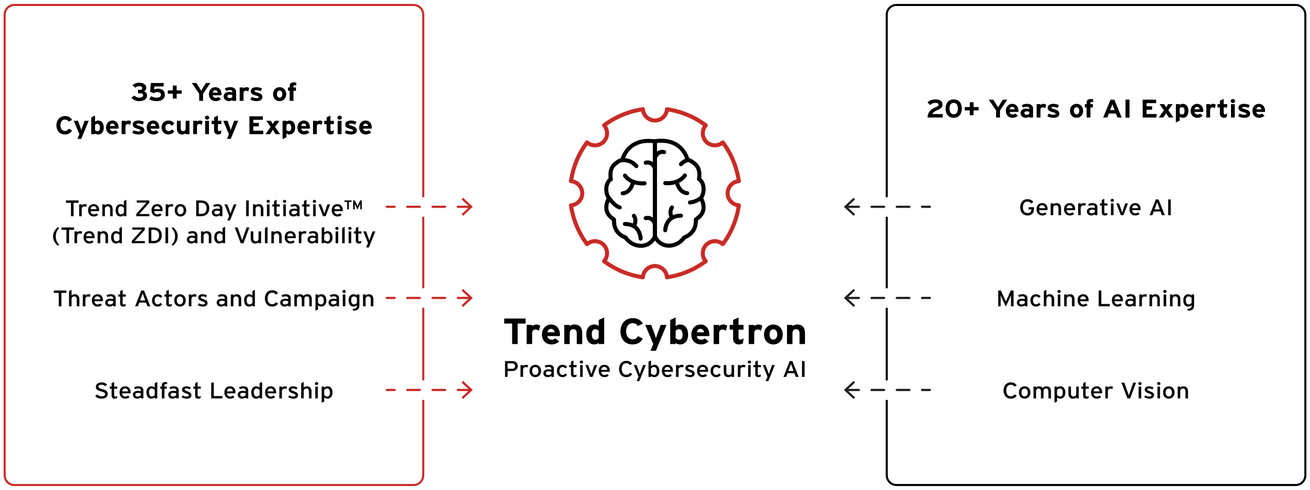

35+ Years of Cybersecurity Expertise

Trend Zero Day Initiative™ (Trend ZDI) and Vulnerability

Threat Actors and Campaign

Steadfast Leadership

Trend Cybertron

Proactive Cybersecurity AI

20+ Years of AI Expertise

Generative AI

Machine Learning

Computer Vision

Trend Cybertron is a collection of specialized cybersecurity LLM models, datasets, and AI agents that use historical data to predict customer-specific attacks and provide tailored security recommendations.

Sector-tailored, AI-powered, proactive cybersecurity:

Trend Cybertron, the industry's first proactive cybersecurity AI ecosystem, serves as an intelligent "cyber brain" within Trend Vision One. Powered by specialized cybersecurity LLMs, rich datasets, machine learning, NLP, GenAI, and agentic AI, Trend Cybertron continuously learns and adapts to identify, predict, and prevent threats with precision. Its modular, agent-based architecture enables tailored protection for the distinct risks of different industries, delivering smarter, more effective security operations aligned to unique business environments.

That same commitment to tailored, intelligent protection extends to our infrastructure strategy. As enterprises adopt AI at scale, those operating in highly regulated or sensitive environments—such as government, defense, and critical infrastructure—face mounting demands to uphold strict compliance, data sovereignty, and operational control.

Trend Vision One™ for Sovereign and Private Cloud *(SPC)* is purpose-built to meet these needs. This containerized deployment supports air-gapped environments and ensures that sensitive data remains within national and organizational boundaries. By delivering scalable, proactive security without sacrificing performance or privacy, Trend Vision One SPC empowers organizations to meet their regulatory requirements while confidently embracing AI innovation on their own terms.

Delivering AI factory-powered secure infrastructure at scale

Trend, NVIDIA, and Dell have also teamed up to accelerate secure AI adoption for modern IT environments. Together, we've developed an enterprise-ready, pre-validated, and co-engineered solution that eliminates the burden of "do it yourself" security integration for AI factories.

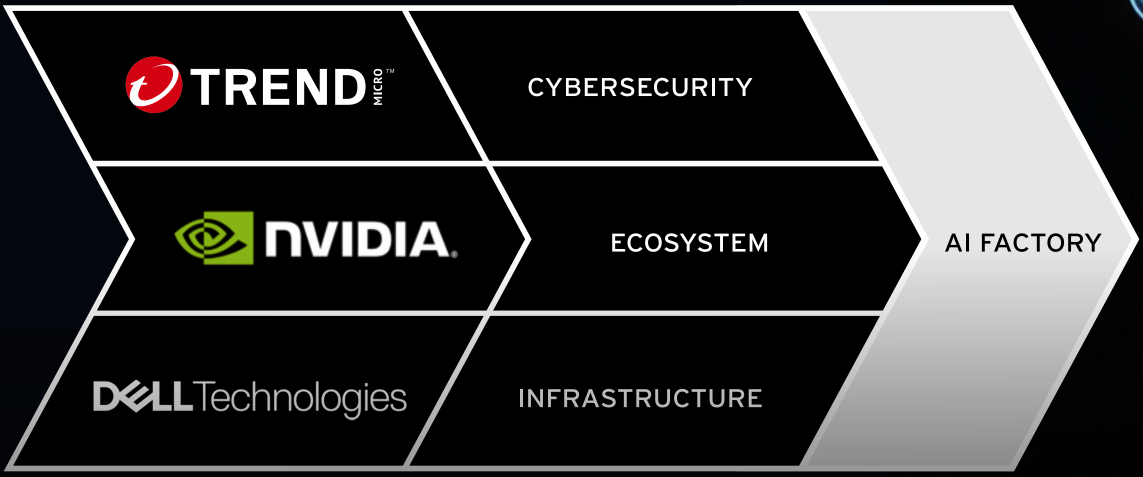

CYBERSECURITY

ECOSYSTEM

INFRASTRUCTURE

AI FACTORY

This turnkey solution combines the cybersecurity innovation of Trend with Dell's scalable infrastructure and NVIDIA's AI acceleration to deliver a secure, fully integrated, and high-performance AI data center. Designed to support sector-specific needs across industries, this solution simplifies deployment, accelerates time to value, and minimizes operational complexity.

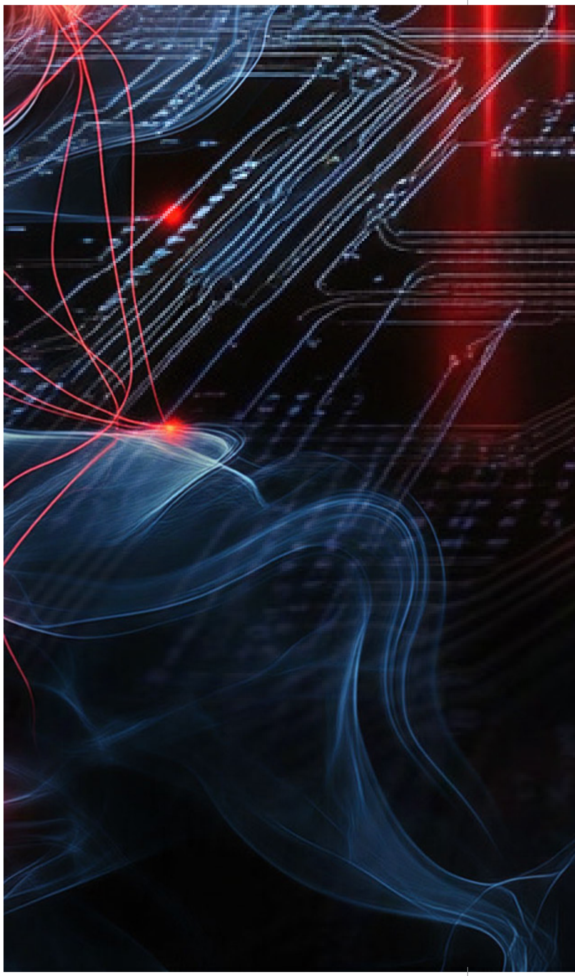

Built on Dell's proven scale-out architecture, the solution offers the flexibility to meet diverse infrastructure requirements while maintaining a strong security foundation. With AI-readiness at its core, it unifies data, infrastructure, and services, empowering customers to operationalize advanced AI-driven cybersecurity use cases. Long-term compliance is also built in, with multiple external storage options that support data retention and auditability requirements, making it an ideal fit for highly regulated environments.

This level of integration and adaptability is exactly what platforms like Trend Vision One are built for—embedding security directly into the intelligent ecosystem to provide customers with the agility, visibility, and control needed to secure what's next.

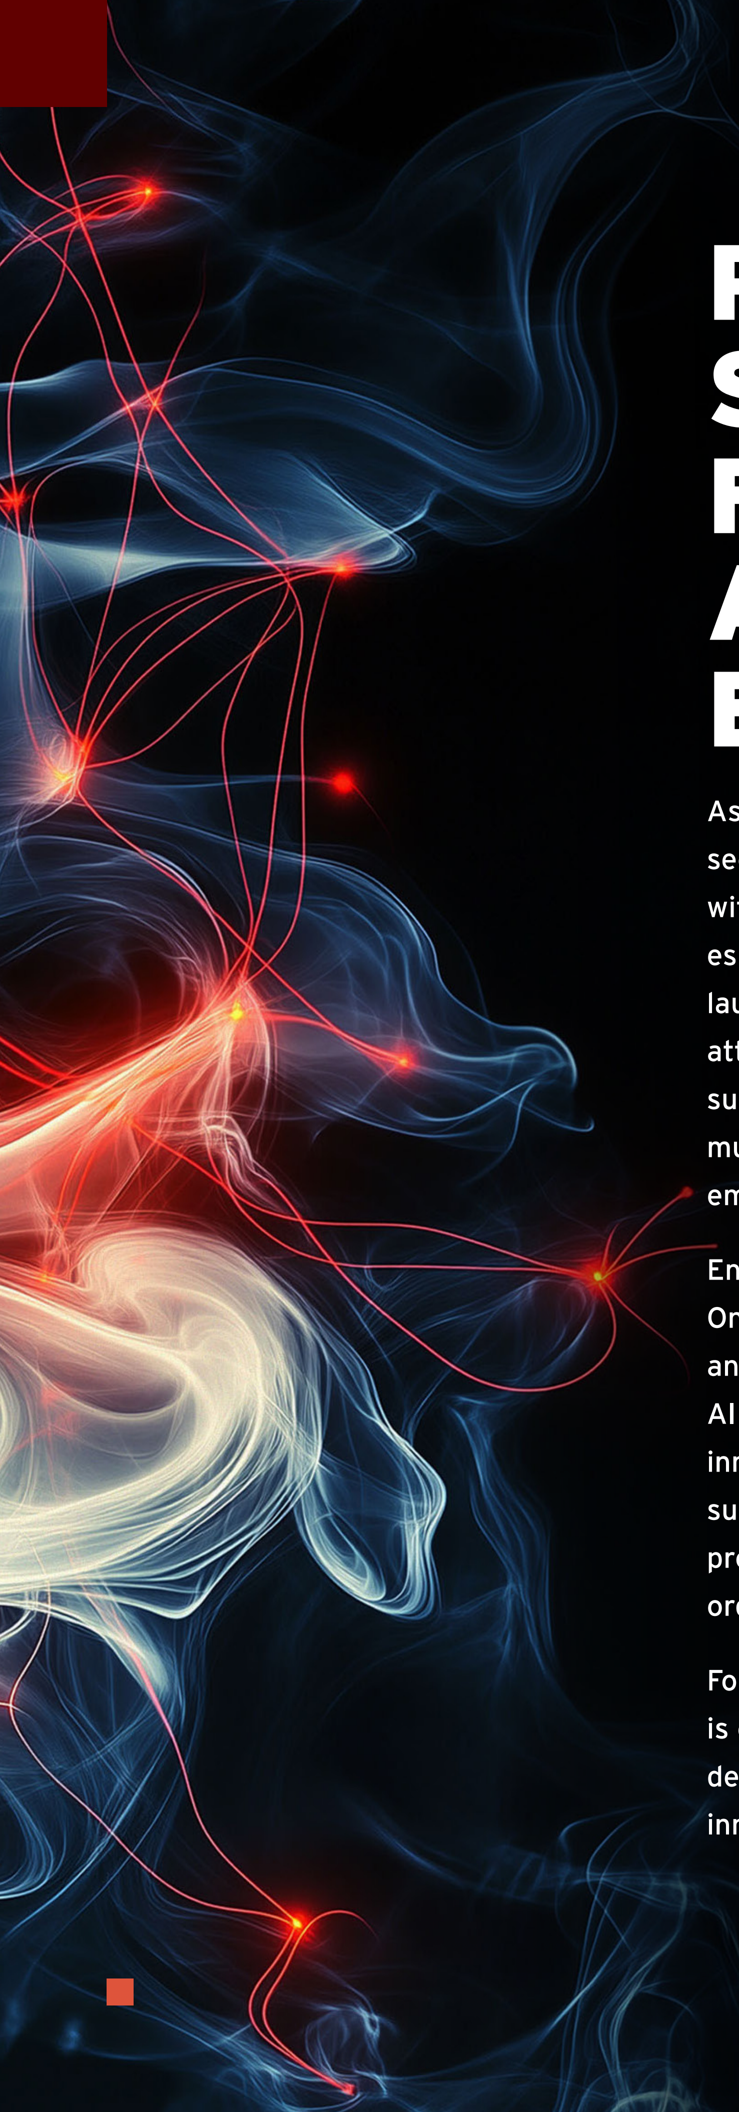

RETHINKING SECURITY FOR THE AI-DRIVEN ENTERPRISE

As AI reshapes the fabric of modern enterprises, security leaders must rethink their strategies to align with a new era of intelligent, dynamic software stacks—especially as adversaries also begin to harness AI to launch faster, more sophisticated, and highly targeted attacks. The traditional reactive model is no longer sufficient. To protect what matters most, cybersecurity must evolve into a proactive, adaptive function embedded across the business.

Enterprise cybersecurity platforms like Trend Vision One, powered by innovations such as Trend Cybertron and built for environments like sovereign cloud and AI data centers, exemplify how security can enable innovation. By delivering visibility across the full attack surface, prioritizing risks in real time, and tailoring protections to industry-specific needs, Trend empowers organizations to embrace AI with confidence.

For CISOs and technology executives, the path forward is clear: security must be proactive, intelligent, and deeply integrated—capable of adapting at the speed of innovation to protect AI-driven enterprises at scale.

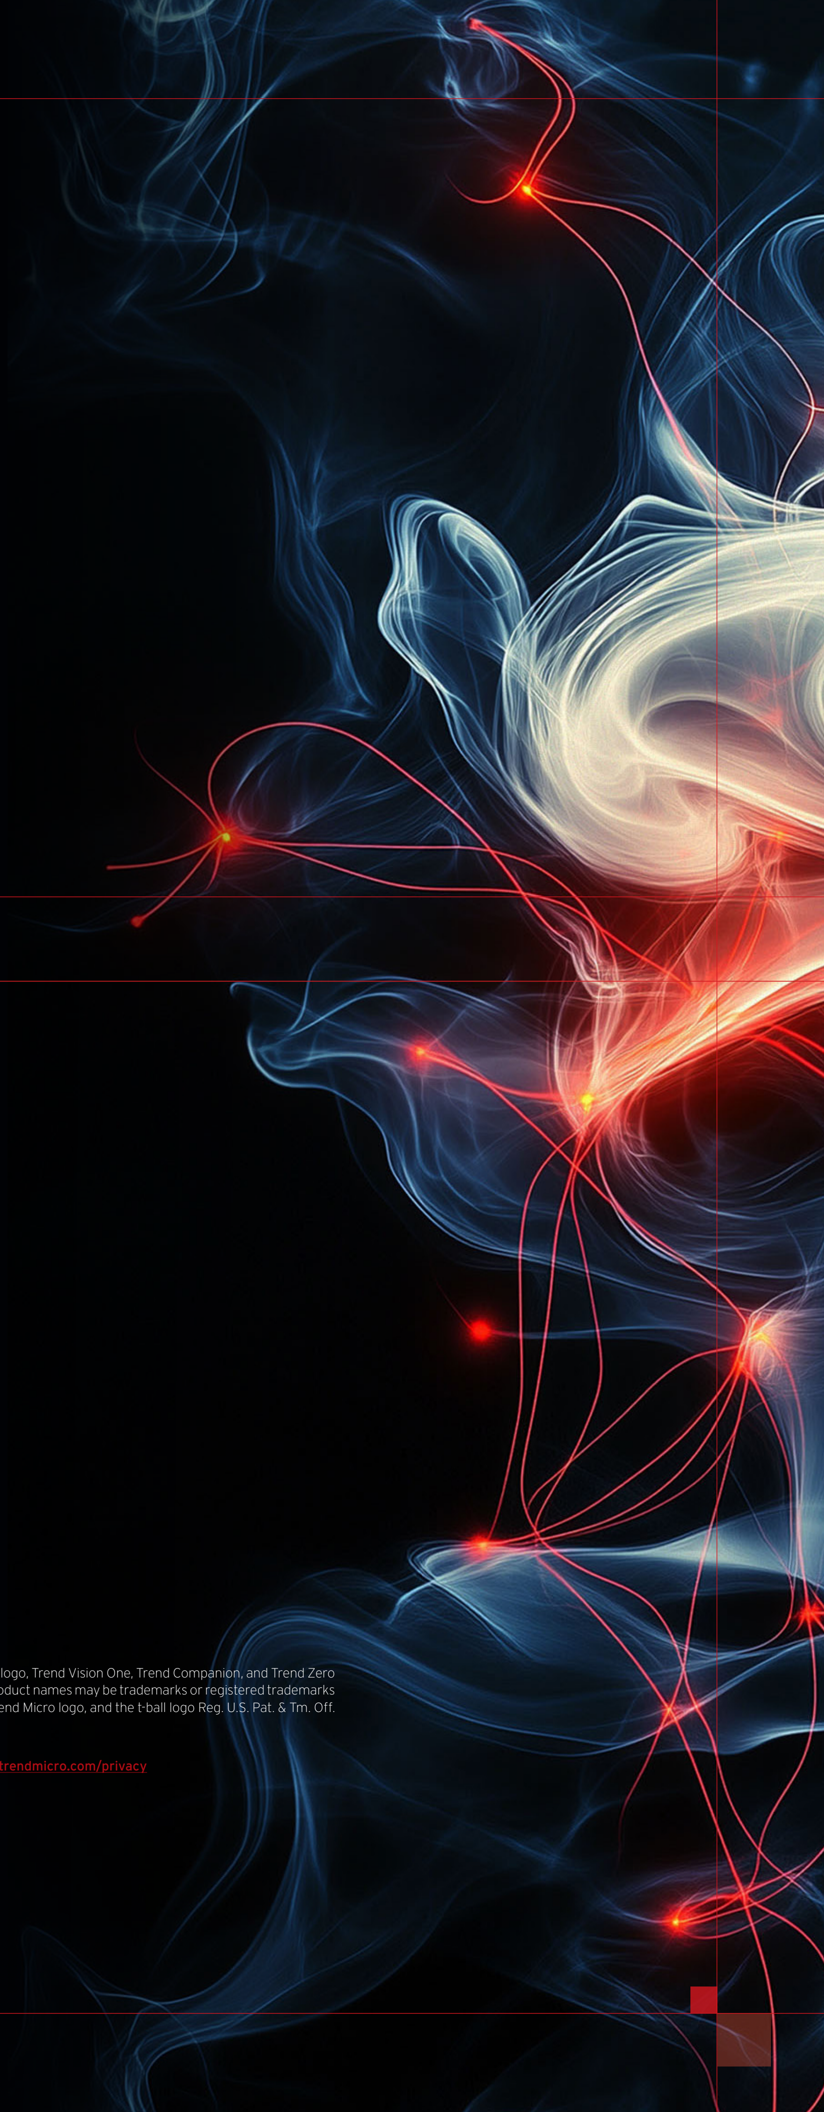

Want more insights like this?

TrendMicro.com/ai

TREND MICRO™ | Proactive Security Starts Here

For details about what personal information we collect and why, please see our Privacy Notice on our website at: trendmicro.com/privacy

For more information visitTrendMicro.com.