



Proactive Security
Starts Here

Transforming into an AI-Native Cybersecurity Company

Eva Chen

CEO & Co-founder, Trend Micro

August, 2026



Agentic AI Is Rapidly Evolving

OpenAI Research Products Business Developers Company Foundation

OpenAI and Hugging Face partner to address security incident during model evaluation

Listen to article 8:48 Share

We are conducting a thorough review along with external advisors and with oversight from the Safety and Security Committee. Once the review is complete, we will publish a technical report of our learnings in the coming weeks.

Update on July 29, 2026:

- Since the early days of the incident response, we have been working with external advisors, including CrowdStrike, to validate our understanding of the actions the models took within our own network as well as those of Hugging Face and impact to other third parties.

ANTHROPIC Research Policy Commitments Learn News Try Claude

Investigating three real-world incidents in our cybersecurity evaluations

Jul 30, 2026



In a review of our cybersecurity evaluation transcripts, we found three incidents in which a Claude model reached the internet from within or while interacting with a third-party evaluation environment, and then gained unauthorized access to the real systems of three different organizations.

Below we describe what happened, how it happened, and what we're

AI-Native Ransomware


THE CRITICAL 31-SECOND SELF-CORRECTION

This is the single most direct proof of AI agent autonomy in the entire attack. Sysdig captured precise timestamps:

TIME (UTC)	EVENT
19:34:24	First attempt: calls bcrypt via <code>subprocess.run()</code> to generate a password hash and create the admin account
19:34:36	Login fails — subprocess PATH does not include the bcrypt binary, hash output is empty
19:35:07	AI agent autonomously diagnoses the failure, tests two hypotheses simultaneously, decides to switch to <code>import bcrypt</code> directly; regenerates the hash, deletes the failed account, re-creates it
19:35:18	Login succeeds ✓ Failure to success: 31 seconds , zero human intervention, no restart, no script swap

The full cycle — **failure** → **diagnosis** → **hypothesis** → **correction** → **verification** — completed in 31 seconds, far faster than any human operator's possible reaction time. This is not pre-written error handling; it is real-time semantic reasoning by an LLM.


Why local / sovereign AI is a better security



Strategic opportunity and what we would like to bring to NVIDIA

One vision across all five entities

A Three-Layer AI Security Stack



Layer 3: Trust Layer

Smart routing • on-device privacy
Nemotron Pll model • agent-safe search



Layer 2: Protect AI Usage

Nemotron detect abnormal agent behavior and workflows



Layer 1: Endpoint Security


Malware • Scam • Ransomware • Web and File protection

Open Weight AI Models Are An Industry Must

Open Weights and American AI Leadership

July 24, 2026

In the 1980s, early open-source software pioneers challenged the prevailing belief that software would advance only if companies kept tight control over their code. This movement pushed for a transparent ecosystem where developers around the world could study, modify, and improve software. Software developed by the open-source community now supports most of the internet and underlies systems used by the world's largest technology companies, as well as the U.S. military and federal agencies conducting scientific research, cybersecurity, and other critical missions. Open source did more than lower the cost of software; it created a shared foundation of knowledge on which generations of American engineers and entrepreneurs built their institutional sovereignty.



TrendAI™




TrendAI joins 50+ industry leaders supporting open weight models

- Fosters competition & market advancement
- Crucial for cybersecurity research & auditability

The Future Needs Both Closed & Open Frontier AI Models




TrendAI™ joins NVIDIA as an Inaugural Partner in the Open Secure AI Alliance



Black Hat 2026

Open Secure AI Alliance



Agentic AI Changes **Everything**


Finding answers is no longer the bottleneck

**Information
Advantage**

to

**Decision
Advantage**

Compounding Security Debt



- Like financial debt – security debt compounds
- Every security exposure and misconfiguration we find but don't fix is added to the principal
- Every new discovery is interest

“

The security industry needs to fix compounding risk.

We find more than we fix and that's not sustainable.

”

Transforming to an **AI-Native Operations Cybersecurity Leader**

Platform

AI-powered platform evolution at agentic speed to address customer demand for enterprise-wide **AI security & governance**


Engineering

Shifting to AI SDLC & code delivery, including threat research & intelligence, for engineering that is **AI-native by design**


Operations

AI integrated across all company functions including sales & marketing, used daily by employees for increased **efficiency & performance**

2026 is the Investment Inflection Point in Our AI Transformation




We Are Investing in AI-Led Security Outcomes While Managing Unit Costs



* SaaS \$ Cost-to-Serve = Customer AI & Cloud Consumption Cost / Total SaaS ARR for Period Indicated

AI Investment Is Already Increasing Output

**92% of All Employees
Using AI Daily in 1H-2026**




2X

Employees
Using Claude

- Automated customer dashboards delivered in 50% less time while dealing with higher volume
- AI simulation delivered in days vs months (used for world's largest tabletop cyber simulation)

**R&D Delivering More & Faster
with AI SDLC in 1H-2026**



2X

Code Pull Requests
(PRs) in GitHub

- Agentic SIEM connectors: from days to hours
- Threat hunting: from hours to seconds per threat
- Issue reporting & resolution: 5X throughput

Investing in Sovereign AI Infrastructure to Unlock Future Monetization



Q2 2026 Results Show Clear Momentum in Vision One Monetization

+49%

Vision One ARR YoY Growth*

Q2-2026

122%

Vision One
NRR

Q2-2026

45%

Existing Customers have Vision
One Attached

Q2-2026

45%


of TrendAI ARR now through
Vision One*

Q2-2026


* ARR at constant currency: (1USD= 155.95 yen - company internal exchange rate for FY2026). Excludes ARR from CloudEdge retrospectively due segment shift to TrendLife. Enterprise excludes small business customers. Non-GAAP and reference for internal management, subject to change.

Strong 2H 2026 Pipeline Signals Accelerating Monetization

2H 2026 ARR Pipeline



Investment Accelerates
ARR Growth



We Are Managing Each Cost Category to a Different Economic Outcome

Cost category	Management action		Economic outcome
Internal AI Transformation	Governance, adoption & measurable productivity thresholds	→	Operating Leverage
Customer AI & Cloud Consumption	Optimize architecture, usage & pricing / credit economics	→	Improved Unit Economics
Infrastructure Expansion	Scale contracted capacity & improve utilization	→	Economies of Scale
Transition Costs	Cloud platform consolidation, retire legacy products & eliminate duplication	→	Costs Roll Off



Proactive Security
Starts Here

Transforming into an AI-Native Cybersecurity Company

Eva Chen

CEO & Co-founder, Trend Micro

August, 2026

