NUCLEAUS

Competitive Landscape Deep Dive

AI Safety & Security - (NUCL-C2S02M005)

Technologies that protect AI systems from adversarial attacks, ensure model safety and alignment, govern responsible deployment, and prevent misuse.

35

Vendors Mapped
Across 4 sub-segments

41%

SAM Uncaptured
$1.44B still greenfield

17 deals

~$3.7B+ M&A in `2025-`26 (current)
Unprecedented consolidation

50%

of Top 10 Acquired
Pure-play window closing

Table of Contents

AI Safety & Security - Competitive Landscape Deep Dive (NUCL-C2S02M005)

MARKET CONTEXT

01 Executive Summary
02 Why This Market Exists
03 The Market at a Glance
04 Sub-Market Landscape

COMPETITIVE ANALYSIS

05 Competitor Scorecard (35 Vendors)
06 Top 10 Power Rankings
07 Capability Heat Map
08 How the SAM is Being Divided

THREATS & OPPORTUNITIES

09 Indirect Competitors
10 Platform Convergence Risk
11 M&A Trajectory
12 White Space Opportunity

STRATEGIC OUTLOOK

13 Threat Assessment
14 Strategic Implications

Delivering

- 15-slide presentation deck
- Highlights from the Nucleaus dataset (35 vendors, sourced ARR, M&A, scorecard, heat map)
- All market statistics independently sourced and cited



Market Context

Executive Summary

AI Safety & Security is a $3.5B market undergoing rapid platform consolidation - with $~3.9B in untapped white space still ahead.

$31B → $86B
TAM by 2030 (22.8% CAGR)

35
Pure-play vendors mapped

8 Deals / ~$2.0B+
M&A in 2025 alone

$34.9M
HiddenLayer ARR - #1 by 4x

38%
Dedicated to AI Runtime & Security

$~3.9B
White space by 2027

Bottom line: The pure-play AI security window is closing fast. Platform vendors are acquiring at 30-50x ARR. The remaining opportunity lies in specialization - agentic AI, regulated-industry governance, and edge security represent the next $~3.9B wave. Act now or pay more later.

Why This Market Exists

From the perspective of the people who feel these problems every day

AI Runtime Security

CISO / VP Security

> *AI-enabled attacks surged 89% over the last year with systems now a top target for cyber criminals.*
>
> - CrowdStrike 2026 Global Threat Report

540%

surge in prompt injection reports

AI Governance

Chief AI Officer / GRC

> *69% of respondents report they lack a monitoring system for AI adoption. Only 13% rate their shadow AI management as highly effective.*
>
> - Cato Networks Global AI Survey, Dec 2025 (600+ IT leaders)

61%

found unauthorized AI in their org

Model Safety

ML Eng. / AppSec

> *Hackers embedded hidden prompt content within phishing emails to confuse AI-based email triage. 90+ organizations were compromised.*
>
> - CrowdStrike 2026 Global Threat Report

#1

OWASP risk for LLM applications

Deepfake Defense

VP Fraud / Trust & Safety

> *A finance worker was tricked into wiring $25 million due to a deepfake video conference call impersonating the CFO.*
>
> - Arup / Hong Kong incident, Feb 2024 (CNN)

2,137%

deepfake fraud increase (3yr)

All statistics independently sourced. Full citations: HackerOne 9th Annual Report (2025); CrowdStrike 2026 Global Threat Report; OWASP Top 10 for LLM Applications (2025); Cato Networks AI Survey (Dec 2025); Signicat Identity Fraud Report (2025); Arup incident via CNN (Feb 2024).

The Market at a Glance

AI Safety & Security - Competitive Landscape Overview

TAM	SAM	SOM	CAGR
$31B	**$3.5B**	**$800M**	**28.0%**
AI Cybersecurity	AI TRiSM + Governance	Pure-Play Vendors	SAM 2025–2030

Solutions that secure AI systems throughout their lifecycle - including runtime guardrails, prompt injection defense, model safety evaluation, AI governance, and deepfake detection.

AI Runtime Security	Guardrails, prompt injection defense, model shielding	*HiddenLayer, Protect AI (Palo Alto), Lakera (Check Point), CalypsoAI (F5), Aim Security (Cato), Prompt Security (SentinelOne), Pangea (CrowdStrike), SPLX (Zscaler), GuardrailsAI, Credal, Kindo, WitnessAI, PromptArmor, Nvidia NeMo* **14 vendors**
AI Governance & Compliance	Policy enforcement, bias auditing, EU AI Act, model risk mgmt	*Credo AI, Holistic AI, Monitaur, Giskard, Fairly (Asenion), Lumenova, Trustible, Saidot.ai, Quantpi, Konfer, Enzai* **11 vendors**
AI Model Safety & Alignment	Red-teaming, alignment research, safety evaluation	*Gray Swan, Mindgard, FAR.AI, CAIS, TrojAI (A10), Preamble, Redwood Research* **7 vendors**
AI Content Authenticity & Privacy	Deepfake detection, AI content ID, provenance, privacy	*Reality Defender, Zama, Apheris* **3 vendors**

Top Buyers by Sub-Segment: CISO / VP Security (Runtime), Chief AI Officer / GRC Lead (Governance), ML Eng / AppSec (Model Safety), VP Fraud / Trust & Safety (Content Auth.)

Sub-Market Landscape

35 vendors across 4 sub-segments within AI Safety & Security SAM ($3.5B → $12B by 2030)

● Top 10 vendors by ARR in segment

AI Runtime Security — 14 vendors

1	● HiddenLayer	$34.9M ARR; RSAC winner; DoD SHIELD
2	● Protect AI (PANW)	Acquired Jul 2025; Prisma AIRS
3	● Aim Sec. (Cato)	Acquired Sep 2025; AI Firewall
4	● Lakera (Check Pt)	Acquired Sep 2025; Lakera Guard
5	● Prompt Sec. (S1)	Acquired Sep 2025;
6	● Pangea (CrowdStrike)	Acquired Sep 2025; AI guardrails
7	CalypsoAI (F5)	Acquired Sep 2025; red-teaming
8	SPLX (Zscaler)	Acquired Nov 2025; AI red-teaming
9	GuardrailsAI	Open-source guardrails framework
10	Credal	YC-backed; enterprise AI gateway
11	Kindo	AI security orchestration
12	WitnessAI	AI visibility platform
13	PromptArmor	Prompt injection defense
14	Nvidia NeMo	NVIDIA open-source guardrails

AI Governance & Compliance — 11 vendors

1	● Credo AI	Forrester Leader; Andrew Ng co-founded
2	● Holistic AI	$8M ARR; Gartner Cool Vendor
3	Monitaur	Model monitoring, governance
4	Giskard	Open-source ML testing
5	Fairly (Asenion)	Acquired; governance + monitoring
6	Lumenova	AI risk management
7	Trustible	AI assurance platform
8	Saidot.ai	AI transparency; Finnish
9	Quantpi	Explainability engine
10	Konfer	Model governance
11	Enzai	Responsible AI risk

AI Model Safety & Alignment — 8 vendors

1	● Gray Swan	CMU research; $40M Series A (2026)
2	● Mindgard	$41M raised; automated red-teaming
3	FAR.AI	Frontier AI safety research
4	CAIS	Center for AI Safety; advocacy
5	TrojAI (A10)	Model scanning, trojans
6	Preamble	LLM safety evaluation
7	Redwood Research	Alignment, interpretability

AI Content Auth. & Privacy — 3 vendors

1	● Zama	~$22M ARR; $150M raised; FHE; $1B val; 2026 blockchain pivot
2	● Reality Defender	$8.3M ARR; $52M raised; deepfakes
3	Apheris	Privacy-preserving AI compute

Vendors ranked by estimated ARR within each sub-segment. Sub-segment sizing: Nucleaus estimates.



Competitive Analysis

Competitor Scorecard

Strong | Good | Moderate | Limited

35 AI Safety and Security vendors assessed across product, pricing, go-to-market, and financial health

AI Runtime Security

		Product	Pricing	GTM	Fin.	Overall
HiddenLayer	$34.9M ARR; RSAC winner; DoD SHIELD					
Protect AI (PANW)	Acquired Jul 2025; $650-700M; Prisma AIRS					
Nvidia NeMo	Open-source guardrails; NVIDIA backed					
Lakera (Check Pt)	Acquired Sep 2025; Gandalf threat intel					
CalypsoAI (F5)	Acquired Sep 2025; $180M; RSAC finalist					
Aim Sec. (Cato)	Acquired Sep 2025; ~$350M; AI firewall					
Prompt Sec. (S1)	Acquired Sep 2025; ~$180M; GenAI DLP					
Pangea (CrowdStrike)	Acquired Sep 2025; ~$260M; AIDR focus					
GuardrailsAI	Open-source guardrails framework					
Credal	Enterprise AI gateway, access controls					
Kindo	AI security orchestration platform					
WitnessAI	AI visibility and governance platform					
SPLX (Zscaler)	Acquired Nov 2025; AI red-teaming					
PromptArmor	Prompt injection defense specialist					

AI Governance & Compliance

		Product	Pricing	GTM	Fin.	Overall
Credo AI	Forrester Leader; Gartner MQ 2025					
Holistic AI	EU AI Act compliance; bias auditing					
Monitaur	Model monitoring, performance gov.					
Giskard	Open-source ML testing and quality					
Fairly (Asenion)	Acquired; AI governance + monitoring					
Lumenova	AI risk management platform					
Trustible	AI assurance and compliance platform					
Saidot.ai	AI system transparency and registry					
Quantpi	Explainability and compliance engine					
Konfer	AI model governance and risk platform					
Enzai	Responsible AI risk management					

AI Model Safety & Alignment

		Product	Pricing	GTM	Fin.	Overall
Gray Swan	CMU research; Cygnal + Shade products					
Mindgard	Automated AI red-teaming; $41M raised					
FAR.AI	Frontier AI safety research, alignment					
CAIS	Center for AI Safety; research + advocacy					
TrojAI (A10)	AI model scanning, trojan detection					
Preamble	LLM safety evaluation and testing					
Redwood Res.	Alignment research, interpretability					

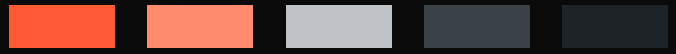
AI Content Authenticity & Privacy

		Product	Pricing	GTM	Fin.	Overall
Zama	FHE for secure AI; $150M raised; $1B val.					
Reality Defender	$52M raised; JPMorgan Hall of Innovation					
Apheris	Privacy-preserving AI computation					

Product = feature breadth + analyst mentions | Pricing = disclosed models, OSS, market accessibility
GTM = customer logos, partnerships, geographic reach, channel presence | Financial = verified/est. ARR, funding, acquirer backing

Capability Heat Map

AI Safety & Security - depth of capability across top 10 vendors

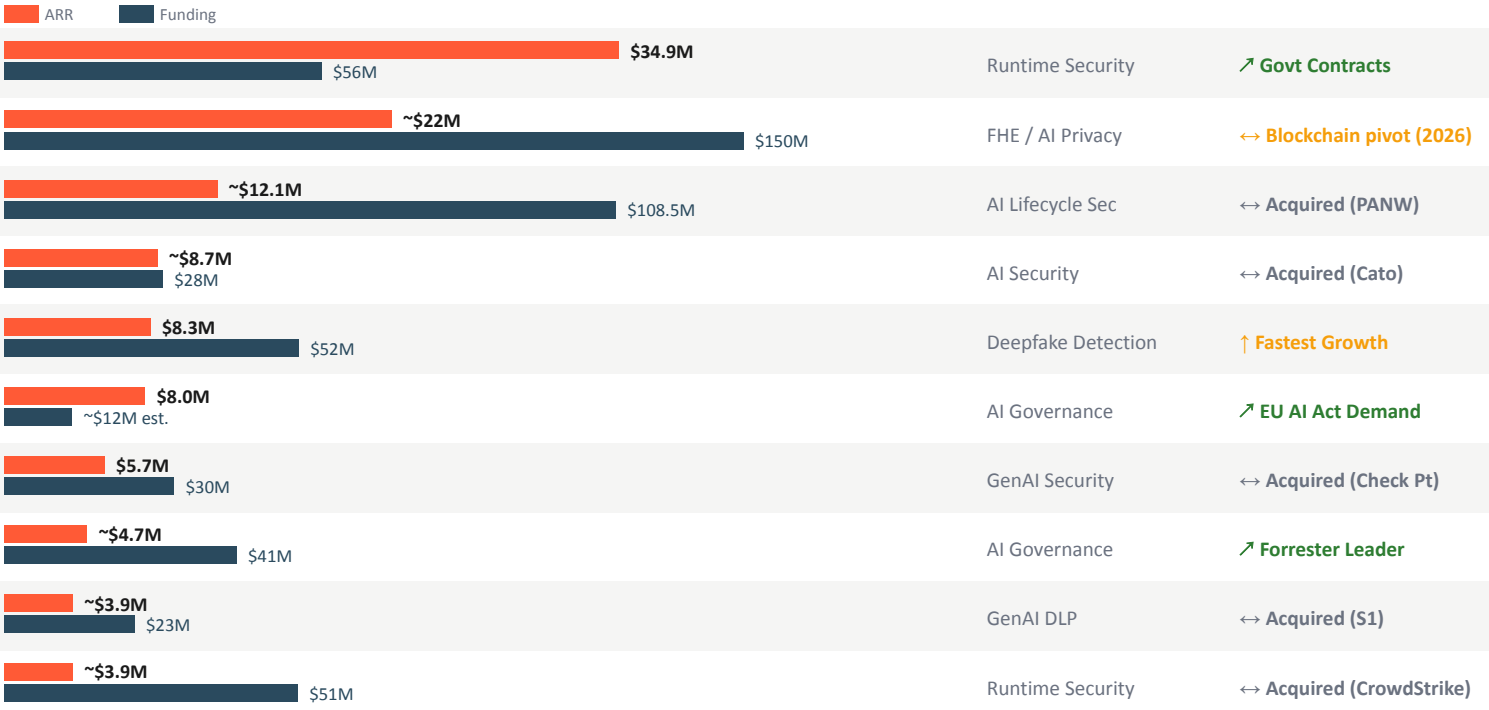


	Runtime Security	Supply Chain Security	Red-Team & Testing	AI Governance	Content Auth / Privacy	Agent Security
HiddenLayer	Strong	Strong	Strong	Good	Emerging	Good
Protect AI (PANW)	Strong	Strong	Strong	Good		Strong
Aim Sec. (Cato)	Strong	Moderate	Moderate	Good		Strong
Reality Defender				Emerging	Strong	
Zama	Good			Emerging	Good	
Holistic AI			Good	Strong		Emerging
Lakera (Check Pt)	Strong	Emerging	Good	Moderate		Good
Credo AI		Emerging	Moderate	Strong		Moderate
Prompt Sec. (S1)	Strong	Emerging	Moderate	Good	Emerging	Good
Pangea (CrowdStrike)	Strong	Emerging	Moderate	Good		Good

Sources: HiddenLayer AISec 2.0; Prisma AIRS 2.0 (PANW, Oct 2025); Aim Security/Cato PR (Sep 2025); Reality Defender product page; Zama Concrete Framework (TechCrunch, Mar 2024); Holistic AI; Lakera Guard (TechCrunch); Credo AI; Prompt Security (SentinelOne PR, Sep 2025); CalypsoAI (RSAC finalist).

Competitive Power Rankings

Top 10 pure-play vendors by estimated ARR within AI Safety & Security - SAM ($3.5B)



Vendor	Estimated ARR / Funding	Segment	Momentum
HiddenLayer	$34.9M ARR / $56M Funding	Runtime Security	↗ Govt Contracts
Zama	~$22M ARR / $150M Funding	FHE / AI Privacy	↔ Blockchain pivot (2026)
Protect AI (PANW)	~$12.1M ARR / $108.5M Funding	AI Lifecycle Sec	↔ Acquired (PANW)
Aim Sec. (Cato)	~$8.7M ARR / $28M Funding	AI Security	↔ Acquired (Cato)
Reality Defender	$8.3M ARR / $52M Funding	Deepfake Detection	↑ Fastest Growth
Holistic AI	$8.0M ARR / ~$12M est. Funding	AI Governance	↗ EU AI Act Demand
Lakera (Check Pt)	$5.7M ARR / $30M Funding	GenAI Security	↔ Acquired (Check Pt)
Credo AI	~$4.7M ARR / $41M Funding	AI Governance	↗ Forrester Leader
Prompt Sec. (S1)	~$3.9M ARR / $23M Funding	GenAI DLP	↔ Acquired (S1)
Pangea (CrowdStrike)	~$3.9M ARR / $51M Funding	Runtime Security	↔ Acquired (CrowdStrike)

Verified: HiddenLayer $34.9M, Reality Defender $8.3M, Holistic AI $8.0M, Lakera $5.7M, Credo AI $3.7M (2024), CalypsoAI $3.8M (2024).

Estimated (~ prefix): Revenue/employee method using $110K median from verified cohort. Protect AI: 110 emp; Aim: 79 emp; Zama: ~200 emp; Prompt Sec: ~35 emp est.

Credo AI 2025 est. via 27% employee growth from 2024 base.

How the SAM Is Being Divided

AI Safety & Security - $3.5B SAM (2025) - who captures what, and what remains open



What's Captured

Top 5 vendors = ~$86M
(~2.5% of SAM)

HiddenLayer ($34.9M), Zama (~$22M), Protect AI (~$12.1M), Aim Security (~$8.7M), Reality Defender ($8.3M)

Other 30+ Pure-Play = ~$200M
(~8% of SAM)

~$200M est. across 30+ other pure-plays (range: $150–250M). Many pre-revenue or sub-$5M ARR. Total pure-play: ~$286M (10.5% of SAM).

Adjacent & Services

~$1.78B (51% of SAM)

Cybersecurity platforms adding AI modules (Palo Alto Prisma AIRS, CrowdStrike, Cisco), cloud provider built-in tools (Azure AI Content Safety, AWS Bedrock Guardrails), and consulting services (Deloitte, EY AI risk practices). Not pure-play but consuming SAM dollars.

What's Still Open

~$1.44B (41% of SAM)

Enterprises deploying AI without purpose-built safety tooling. Only 10-15% of AI-deploying orgs purchase dedicated safety tools in 2025. Massive greenfield concentrated in mid-market and regulated verticals with $12K-150K budgets.

Source: Nucleaus TAM (NUCL-C2S02M005). Top 5 ARR: verified figures (HiddenLayer, Reality Defender) and employee-method estimates (Protect AI, Aim Security, Zama ~$22M) Other pure-play est. $150–250M range. Adjacent = reconciliation residual.

Threats & Opportunities

Indirect Competitors

Market Threat: Major | Moderate | Emerging

Platform vendors and hyperscalers entering AI Safety & Security - market threat assessment

Company	Type	AI Security Product	Key Capabilities	Total Rev.	AI Sec. Rev. Est.	Market Threat	Source
Palo Alto Networks	Platform	**Prisma AIRS (incl. Protect AI acq.)**	Runtime security, model scanning, red-teaming, agent security, AI-SPM, posture mgmt	$9.2B	**~$100-200M**	Major	Prisma AIRS 2.0 PR Oct 2025; PANW earnings
CrowdStrike	Platform	**Charlotte AI + Falcon Shield + Pangea acq.**	Agentic SOC, shadow AI discovery, model scanning, agent identity governance	$4.9B ARR	**~$50-100M**	Major	Fal.Con 2025; Fall 2025 Release PR; VentureBeat
Microsoft	Hyperscaler	**Security Copilot + Azure AI Content Safety + Purview AI Hub**	NLP-driven SOC, content filtering, AI governance within Azure ecosystem	$236B	**Bundled**	Emerging	Microsoft Security Copilot product page
Google / DeepMind	Hyperscaler	**SAIF + Model Armor + SynthID + Mandiant AI**	Model Armor AI firewall (GA), AI red-teaming, SynthID watermarking (10B+ items), Vertex AI safety	$350B+	**Bundled**	Emerging	Model Armor product page, saif.google, Mandiant AI; DeepMind
AWS	Hyperscaler	**Bedrock Guardrails + SageMaker Model Monitor**	Content filtering, PII detection, model monitoring, guardrails for GenAI apps	$107B	**Bundled**	Emerging	AWS re:Invent 2025; Bedrock Guardrails docs
Cisco	Platform	**Cisco AI Defense (incl. Robust Intelligence, Astrix, WideField)**	AI model validation, runtime protection, guardrails	$56B	**~$10-30M**	Moderate	Cisco AI Defense launch 2025
Check Point	Platform	**Infinity AI Copilot + Lakera acq.**	GenAI security (via Lakera), threat prevention with AI, email security	$2.6B	**~$30-50M**	Moderate	Check Point / Lakera acq. Oct 2025
SentinelOne	Platform	**Purple AI + Prompt Security acq.**	GenAI DLP, runtime AI protection, agentic response, Singularity platform	$815M	**~$20-40M**	Moderate	SentinelOne / Prompt Sec. PR Sep 2025
Cato Networks	SASE	**Cato AISEC (incl. Aim Security acq.)**	AI Firewall, shadow AI discovery, AI-SPM, inline AI interaction monitoring	$350M ARR	**~$20-40M**	Moderate	Cato / Aim Security PR Sep 2025
Zscaler	Platform	**Zscaler AI Security (incl. SPLX acq.)**	AI asset discovery, auto red-teaming (5,000+ attack sim.), runtime guardrails, prompt hardening, AI governance	$2.4B	**~$5-15M**	Moderate	Zscaler/SPLX PR Nov 2025

Key insight: 8 pure-play AI security vendors acquired in 2025, and the wave continued in 2026 - shifting to AI-agent identity (Cisco/Astrix, CrowdStrike/SGNL, Okta/Permiso, SailPoint/Entro). Two acquirers (PANW, CrowdStrike) now pose Major market threats with dedicated AI security platforms and >$50M revenue. Three hyperscalers bundle AI security at no extra cost - an emerging long-term threat. Current pure-play offerings are consolidating. At the same time whitespace in AI Safety & Security is opening up, largely untapped by these competitors.

Revenue: public filings (PANW, CRWD, MSFT, GOOGL, AMZN, Cisco, CHKP, Zscaler); SentinelOne FY2025 shareholder letter; Cato Networks PR. AI Sec. Rev: Nucleaus estimates incl. acquired pure-play ARR. Market Threat: Major = acquired pure-play + >$50M AI rev

Platform Convergence Risk

How indirect competitors are expanding across AI Safety & Security capability areas



Company	Type	Coverage Risk	Runtime Security	AI Governance	Red-Team & Safety	Content Auth.	Agent Security	Coverage Score
Palo Alto Networks	Platform	Critical	Full	Strong	Full	None	Full	75%
Google / DeepMind	Hyperscaler	High	Strong	Strong	Strong	Strong	Emerging	65%
Cato Networks	SASE	High	Full	Strong	Moderate	None	Full	65%
CrowdStrike	Platform	Medium	Strong	Moderate	Strong	None	Strong	55%
Microsoft	Hyperscaler	Medium	Strong	Strong	Moderate	Emerging	Moderate	55%
Check Point	Platform	Medium	Strong	Moderate	Strong	None	Strong	55%
SentinelOne	Platform	Medium	Strong	Moderate	Moderate	Emerging	Strong	55%
Zscaler	Platform	Medium	Strong	Moderate	Strong	None	Emerging	45%
Cisco	Platform	Medium	Strong	Moderate	Moderate	None	Moderate	45%
AWS	Hyperscaler	Low-M	Moderate	Moderate	Emerging	Emerging	Moderate	40%

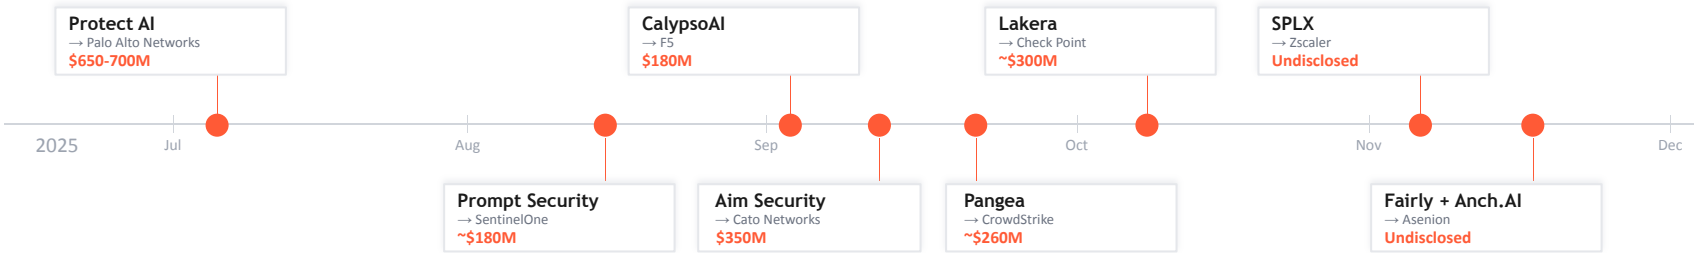
Convergence trend: Platform vendors are rapidly filling capability gaps through acquisition. Runtime Security is the most contested area (9 of 10 at Strong or Full). Agent Security is the fastest-moving - 2 vendors reached Full via acquisition (PANW, Cato). Content Auth. is the biggest gap - 6 of 10 have no capability, creating a defensive niche for pure-play specialists. No vendor covers all 5 dimensions at Strong or above.

Sources: Prisma AIRS 2.0 PR (PANW, Oct 2025); CrowdStrike Falcon AIDR / Pangea acq. (Sep 2025); Microsoft Security Copilot docs; Google Model Armor / SAIF 2.0 / Mandiant / SynthID; AWS Bedrock docs; Cisco AI Defense (Robust Intelligence acq.); Check Point / Lakera acq. (Oct 2025); SentinelOne / Prompt Security acq. (Sep 2025); Cato / Aim Security acq. (Sep 2025); F5 / CalypsoAI acq. (Sep 2025).

M&A Trajectory: 2025

2025: Unprecedented consolidation in AI Safety & Security - 8 acquisitions totaling ~$2.0B+

8	~$2.0B+	50%
Deals in 2025	Total Deal Value	of Top 10 Acquired



Deal Summary

Target	Acquirer	Date	Est. Price	Target Funding	Target Emp.	Strategic Rationale
Protect AI	Palo Alto Networks	Jul 2025	**$650-700M**	$108.5M	110	Largest AI security acquisition; forms basis of Prisma AIRS platform; full lifecycle coverage from model scanning to runtime defense
Prompt Security	SentinelOne	Aug 2025	**~$180M**	$23M	~35	Adds GenAI DLP and runtime protection to Singularity platform; 10x+ return for investors in <2 years
CalypsoAI	F5	Sep 2025	**$180M**	$43M	64	Extends F5's app delivery platform into AI security enforcement and red-teaming capabilities
Aim Security	Cato Networks	Sep 2025	**$350M**	$28M	79	First-ever Cato acquisition; embeds AI Firewall + AI-SPM into SASE platform; Fortune 500 customers
Lakera	Check Point	Sep 2025	**~$300M**	$30M	52	Adds Lakera Guard GenAI firewall; Gandalf-derived prompt injection taxonomy; Swiss engineering team
Fairly + Anch.AI	Asenion (merger)	2025	**Undisclosed**	Undiscl.	—	Fairly AI acquired Anch.AI and relaunched as Asenion; governance sub-segment consolidation
Pangea	CrowdStrike	Sep 2025	**~$260M**	$51M	~35	Pioneers AIDR category; prompt injection defense + AI governance for Falcon; $51M raised

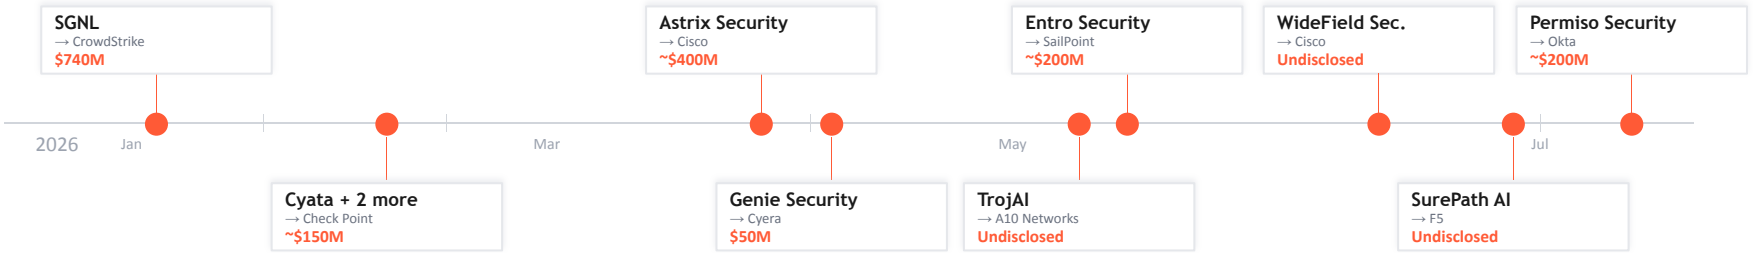
Next target: HiddenLayer ($34.9M ARR, $56M raised, 164 employees) is the largest remaining independent pure-play AI security vendor. Likely acquirers: Cisco, a hyperscaler, or a SIEM vendor seeking AI security credibility.

Sources: GeekWire/PANW PR (Protect AI); SentinelOne 8-K/SEC (Prompt Sec. ~$180M); F5/CalypsoAI PR; Globes/Calcalist (Aim Sec.); SecurityWeek/WSJ (Pangea ~$260M); SecurityWeek/Globes (Lakera ~$300M); CrowdStrike Fal.Con 2025. Dates = announcement (Protect AI = close).

M&A Trajectory: 2026

2026: Consolidation shifts to AI-agent identity - 9 deals YTD, ~$1.74B disclosed

9	~$1.74B	$57B
Deals YTD (11 targets)	Disclosed Deal Value	CyberArk + Wiz Deals Closed



Deal Summary

Target	Acquirer	Date	Est. Price	Target Funding	Target Emp.	Strategic Rationale
SGNL	CrowdStrike	Jan 8, 2026	$740M	—	—	Largest 2026 AI-security deal; agentic identity security governing AI-agent access across the Falcon platform
Cyata + 2 more	Check Point	Feb 12, 2026	~$150M comb.	Private	—	Three-startup sweep (Cyata, Cyclops, Rotate); Cyata adds AI-agent identity to the Infinity platform
Astrix Security	Cisco	May 4, 2026	~$400M	$85M	121	Non-human and AI-agent identity security; closed Jun 29 into Cisco Identity Intelligence
Genie Security	Cyera	May 20, 2026	$50M	Undiscl.	5	Endpoint GenAI data-leak prevention added to Cyera's data-security platform
Entro Security	SailPoint	Jun 15, 2026	~$200M	Private	—	Machine-identity and AI-agent governance extends the SailPoint platform; closed Jun 29
TrojAI	A10 Networks	Jun 15, 2026	Undisclosed	Private	—	AI red-teaming and model/runtime protection joins A10's application-delivery security
WideField Security	Cisco	Jun 19, 2026	Undisclosed	$11M+	—	Agentic-SOC identity threat detection for Splunk; Cisco's second AI-security buy of 2026
SurePath AI	F5	Jun 26, 2026	Undisclosed	~$6M	~19	Shadow-AI discovery and control; anchors F5's new AI Security Platform
Permiso Security	Okta	Jul 30, 2026	~$200M	Private	—	AI-agent identity threat detection; Okta's push into securing non-human identities

Context: 2026's identity mega-deals closed - Palo Alto/CyberArk ($25B, Feb 11) and Google/Wiz ($32B, Mar 11). Cisco bought twice (Astrix, WideField); F5 twice across 2025-26 (CalypsoAI, SurePath). HiddenLayer still independent.

Sources: CNBC (SGNL); SecurityWeek/Calcalist (Check Point trio); Cisco PR (Astrix, WideField); Calcalist (Genie); SailPoint PR (Entro); BusinessWire (TrojAI); GeekWire (SurePath); Okta PR/TechCrunch (Permiso). Dates = announcement; ~ = reported.

White Space Opportunity

Underserved opportunity areas within AI Safety & Security - estimated addressable value by 2027

High | Mod.-High | Moderate | Emerging

~$3.9B
total white space

Opportunity	Opportunity	Market Potential vs. Captured	Gap Description	Potential Entrants
Agentic AI Security	High	$1.15B ~25% captured (~$290M)	Agent-to-agent, tool-use, and MCP security gaps. WitnessAI ($58M Series B) and multiple startups entering; still early but moving fast.	**Palo Alto (Prisma AIRS)** **CrowdStrike (Pangea)** **Cato / Zenity / Noma / WitnessAI**
Deepfake / Synthetic Media	Mod.-High	$875M ~20% captured (~$175M)	Voice deepfake detection, video KYC, synthetic media watermarking. $15.7B detection market projected by 2026. Reality Defender only scaled pure-play.	**Reality Defender** **Microsoft (Video Auth.)** **Pindrop / Hive**
AI Supply Chain Security	Moderate	$700M ~30% captured (~$210M)	70% of AI/ML repos have critical vulnerabilities (Mitiga). AIBOM standards nascent. No vendor covers full training-to-deployment chain.	**HiddenLayer (AIBOM)** **PANW (model scanning)** **Wiz (Google)**
AI Gov. for Regulated Ind.	Moderate	$575M ~35% captured (~$200M)	Sector-specific governance needed (FDA AI/ML, OCC model risk, EU AI Act). IBM and consulting firms active, but pure-play vertical tools remain scarce.	**Credo AI** **Holistic AI** **IBM watsonx.governance**
Encrypted AI / Privacy-Pres. ML	Emerging	$300M ~10% captured (~$30M)	FHE for AI is viable but slow. Secure enclaves, differential privacy, federated learning security adjacent. Awaiting performance breakthroughs.	**Zama** **Apheris** **Intel (SGX) / NVIDIA**
AI Security for Edge / IoT	Emerging	$250M ~10% captured (~$25M)	Edge AI market is $25B but AI model security for edge is nascent. HiddenLayer shipped airgapped platform (Jan 2026). Mass deployment is a 2027-2028 story.	**HiddenLayer (airgapped)** **CrowdStrike (XIoT)** **New entrants**

Sizing = Nucleaus estimates. Captured = est. vendor revenue.



Strategic Outlook

Threat Assessment

Summary of competitive risks and strategic implications for AI Safety & Security

CRITICAL — Platform Consolidation

- 8 pure-play AI security vendors acquired in 2025; 9 more deals in 2026
- Palo Alto, CrowdStrike, Microsoft all building end-to-end AI security platforms
- Pure-play window closing - remaining independents face build-or-be-bought pressure
- HiddenLayer ($34.9M ARR) is the last scaled independent

Pure-play vendors must reach $50M+ ARR or find a strategic acquirer within 12-18 months

HIGH — Hyperscaler Bundling

- Microsoft, Google, AWS embedding AI security into cloud platforms at no incremental cost
- Security Copilot, Bedrock Guardrails, and SAIF framework commoditize basic runtime protection
- Governance vendors (Credo AI, Holistic AI) less affected - regulatory complexity protects them

Differentiation must come from depth, not breadth - specialized capabilities the hyperscalers won't build

MODERATE — Capability Gaps as Opportunity

- ~$3.9B in white space across 6 underserved areas - larger than current SAM
- Agentic AI Security ($1.15B) is the single biggest gap - no vendor has full coverage
- Edge/IoT AI security is nascent ($250M, ~10% captured) with real demand in 2028+
- Deepfake defense has only 1 scaled pure-play (Reality Defender) for an $875M opportunity

First movers in agentic and edge AI security will define the next wave of $500M+ acquisitions

Key Market Signals

$34.9M	HiddenLayer ARR - 4x larger than #2 vendor. Market has one clear revenue leader.	*Nucleaus dataset*
~$2.0B+	Total M&A deal value in 2025 across 8 deals. Acquirers paying 30-50x ARR for strategic AI security assets.	*GeekWire, Globes, Calcalist*
38.5%	SAM uncaptured ($1.44B). The majority of the $3.5B addressable market has no dedicated AI security vendor.	*Nucleaus estimate*
2 Major	indirect competitors (PANW, CrowdStrike) with acquired pure-plays and >$50M AI security revenue.	*Product launches, acquisitions*
~$3.9B	white space by 2027 across 6 opportunity areas - exceeding the entire current SAM ($3.5B).	*Nucleaus estimate*

Assessment based on: Power Rankings (verified ARR/employee method), Indirect Competitor analysis, M&A Trajectory (public filings/PRs), White Space sizing (Nucleaus estimates).

Strategic Implications

Actionable takeaways from the AI Safety & Security competitive landscape analysis

1 Invest Now - Window is Closing

5 of 10	top pure-plays acquired in 2025
30-50x ARR	average acquisition multiples
$34.9M	HiddenLayer ARR - last scaled independent

Evaluate acquisition targets now. Asset prices will only increase as platform vendors compete for the shrinking pool.

2 Specialize, Don't Compete on Breadth

~$3.9B	in white space exceeds current SAM
$1.15B	agentic AI security - biggest single gap
$250M	edge/IoT - ~10% captured, nascent but emerging

Back vendors solving problems hyperscalers can't bundle - sector-specific compliance, agent sec., or airgapped dep.

3 Watch the Governance Layer

EU AI Act	transparency live Aug 2026; high-risk Dec 2027
<$10M ARR	both leaders (Credo AI, Holistic AI)
Most defensible	regulatory moats outlast technical ones

Governance is the most defensible segment against platform bundling. This category will produce outsized returns.

Market Snapshot

TAM (2025 → 2030)	$31B → $86B	22.8% CAGR
SAM (2025 → 2030)	$3.5B → $12B	28% CAGR
Pure-Play Share	~10.5% of SAM	$286M / 35 vendors
Uncaptured SAM	38.5%	$1.44B unaddressed
White Space (2027)	~$3.9B	6 opportunity areas
M&A (2025)	8 / ~$2.0B+	Avg ~$250M at 30-50x

What to Watch in 2026

- HiddenLayer acquisition - who and at what multiple sets the benchmark
- EU AI Act Omnibus deferral (high-risk → Dec 2027) - does governance demand slip?
- First major agentic AI breach catalyzes the $1.15B opportunity
- CrowdStrike vs. Palo Alto - platform war defines the decade