ANCHORPOINT
ADVISORY

JULY 2026

INSIGHT REPORT

AI GOVERNANCE: WHERE STRATEGIC VALUE IS EMERGING

Identity, runtime enforcement and the consolidation of the enterprise AI control stack

INVESTMENT TRENDS AND M&A ANALYSIS

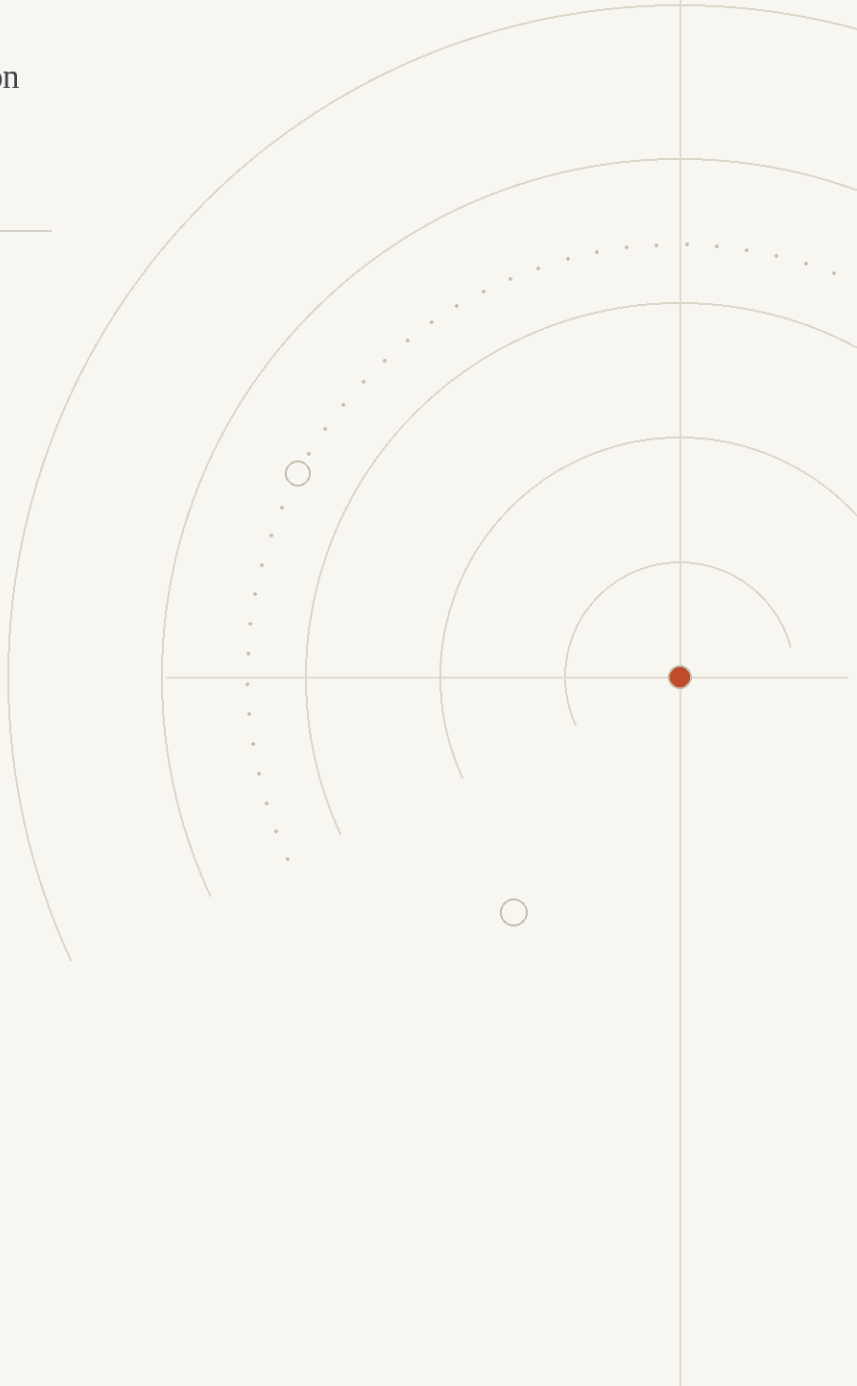

PREPARED BY ANCHORPOINT ADVISORY

Independent technology investment advisory

anchor-point.co.uk

Six takeaways

Artificial intelligence is reshaping enterprise governance as models and agents gain the ability to access systems, make decisions and execute actions with increasing independence. Yet AI deployment is advancing faster than the controls surrounding ownership, identity, permissions, runtime activity and accountability. Enterprises therefore face a growing gap between the capabilities they are introducing and their ability to understand, constrain and take responsibility for them. Governance is consequently moving beyond model quality and output accuracy towards the control of AI activity across the enterprise.

This shift is changing where strategic value sits within the AI governance market. Visibility remains important, but the strongest control positions are emerging around agent identity, permissions, policy enforcement and the ability to intervene before consequential actions are completed. Security, identity, cloud, data and workflow vendors are approaching the market from different starting points but are increasingly assembling capabilities across the same control chain. The resulting investment and acquisition activity provides an early indication of which technologies may become essential to enterprise AI adoption and where important gaps remain.

This report draws on Anchorpoint Advisory's continuing coverage of AI governance and security, spanning category development, ecosystem mapping, financing activity and M&A. The principal conclusions are:

01 **Agentic execution is shifting the centre of risk.** Control is moving from what AI says to what AI can access, decide and execute. Identity, permissions and runtime accountability are now critical as agents gain credentials and tools.

02 **Strategic value concentrates closest to identity and execution.** Runtime-enforcement transactions provide the clearest evidence; agent discovery and detection-and-response deals show buyers assembling the surrounding control architecture.

03 **Leading security vendors are buying complementary capabilities rather than relying on a single acquisition.** Palo Alto Networks has made three core AI-control acquisitions, supported by CyberArk in identity and Chronosphere in observability. Most data and workflow buyers have so far completed only their first acquisition in the category.

04 **Four buyer groups are moving towards the same control chain.** Security and identity vendors start with risk and permissions; cloud, data and workflow platforms start with information and execution. All want influence over identity, policy, enforcement and evidence.

05 **Two important needs remain under-served.** Agent identity received strong strategic endorsement through Palo Alto Networks' CyberArk acquisition, although its value reflects a broader identity-security business. Customer-controlled, private and on-premises governance also remained uncommon across the specialist set reviewed.

06 **Conditions favour continued investment and consolidation.** Agent adoption, clearer regulation and gaps in incumbent products should sustain financing and M&A. The market may eventually support a few broad suites and a smaller group of independent vendors that work across platforms, operate in customer-controlled environments or own hard-to-build technology.

BACKGROUND

The story so far

As enterprises move from using AI for assistance to allowing it to perform actions, they need stronger operating controls. Governance is not a tax on AI adoption; it is the infrastructure that allows adoption to scale.

Three waves, one control response

Wave I — Assistance Drafting, search, coding and customer support improved individual productivity. The control problem centred on policy, accuracy and data leakage.	**Wave II — Connectivity** Models connected to enterprise data, retrieval systems, APIs and applications. Every model call and data retrieval may need to be monitored or controlled.
Wave III — Delegated execution Agents can select tools, act under credentials, modify systems and coordinate workflows. Risk shifts from output quality to delegated authority.	**The control response** Enterprises need a real-time layer for identity, permissions, policy, evidence, containment and accountability, not another policy document.

Adoption and control

Adoption is broad, but scaled value and operating control remain much narrower. That gap is why a new market is forming. The survey figures below describe different populations and should not be combined into a single prevalence estimate.

88%
of organisations use AI in at least one function
McKinsey 2025 survey

39%
report any enterprise-level EBIT impact
McKinsey 2025 survey

97%
of AI-incident organisations lacked proper access controls
IBM breach research 2025

63%
of AI-incident organisations had no governance policy
IBM breach research 2025

From answer risk to action risk

A chatbot that produces a poor answer creates a quality, legal or reputational problem. An agent connected to credentials, databases, software repositories and payment systems creates an execution problem. The difference is delegated authority, and it is why “human in the loop” cannot be the entire control strategy once agents operate continuously, at machine speed, across chained systems.

Sources: Anthropic, Model Context Protocol; OWASP, Top 10 for Agentic Applications 2026; UK National Cyber Security Centre.

Nine layers, one control architecture

AI governance is better understood as a set of connected controls than as one product. The map sets out nine functional layers. For the M&A analysis, these are grouped into six broader transaction categories.

THE AI GOVERNANCE & CONTROL ECOSYSTEM MAP

#	LAYER	WHAT IT ADDRESSES	REPRESENTATIVE PLAYERS	SIGNAL / READ-THROUGH
1	Data Security (Sovereignty / DLP / Privacy)	Protecting sensitive data at rest, in motion and in AI interfaces.	Microsoft Purview, OneTrust, RELIANCE AI, BigID, HARMONIC, Mine	Maturing AI-specific controls extending existing DLP and privacy stacks.
2	Shadow AI Discovery	Visibility into unauthorised or unmonitored AI tool use.	aranscape, HARMONIC, LayerX, zscaler, Pangea Security, LASSO	Fast-growing entry wedges; consolidating into control and DLP suites.
3	AI Lifecycle Governance	Policy, audit trail and regulatory adherence across the model lifecycle.	captain, enate, Artion, Holistic AI, ModelOp, trainIte, DEFPLOY, IBM watsonx governance	Regulatory pull; independent value layer often tied to model performance.
4	Non-Human Identity (NHID)	Identity and least-privilege governance for AI agents and tools.	zenity, Knostic, CYBERARK, Microsoft Entra, Cymfe, LASSO	Very early — critical emerging gap.
5	Runtime Enforcement / AI Gateway	Inline policy enforcement, routing, inspection and control of model, agent and tool interactions.	aidome (AI-native gateway), NOMA SECURITY, Writer, airia, PORTKEY, NEKOR.AI, PROXY.AI, CLOUDFLARE	Highest disclosed funding and M&A concentration.
6	Application / AI Security	Protect AI applications, jailbreak prevention, prompt guardrails for specific AI applications.	LAKERA, CALYPSO AI, Prompt Security, Pangea Security	Closest analog to legacy AppSec 2025–26 expansion.
7	AI Threat Detection & Response	Detecting risks, adversarial behaviour and anomalies in AI systems and interactions.	OBSERVEIT, CROWDSTRIKE /FALCON, paloalto, SentinelOne, HUNTERS, CORTEX	Early — category still being defined.
8	Model / MLOps Security	Model scanning, validation, and security across supply chain and runtime.	PROTECT AI, AI Defense (Forcewinst/Strider), HIDDENLAYER, preemptive, TrojAI, torch (Neuronenixe), deepestlabs	Early supply chain and gen/sec.
9	Unified Platform (UADP)	Integrating all prior layers into one intent-aware control plane.	NOMA SECURITY, Writer, LASSO, aidence (Agentic AI intelligence graph)	Pre-emergent — 2026–2030 formation window.

ADJACENT ECOSYSTEM

GRC, Risk, Compliance & Evidence: LOGICGATE, optro, TrustBits, DRATA, OneTrust, RSA ARCHER, service now, Atlassian Jira, safeactor, LITERA, Amazon Akirvian, splunk>, Google Cloud Assurance, Microsoft Entra Governance, aws Audit Manager, Snowflake, SAP, workday, GitHub, Jira, slack

Underlying Infrastructure: Microsoft Azure & Entra ID, aws, Google Cloud AI Studio, IBM watsonx, databricks, snowflake, salesforce, SAP, workday, GitHub

Foundation Models & Agentic Ecosystem: OpenAI, ANTHROPIC, Gemini, Meta Llama, MISTRAL AI_, cohere, Microsoft 365, salesforce, servicenow, workday

Unmet need — Identity controls for autonomous AI agents

Most identity-security products were designed for people and conventional software, not autonomous agents that may be created dynamically, exist briefly or act under delegated authority. Across the products reviewed, complete identity, least-privilege and revocation controls for such agents remained uncommon.

Unmet need — Private and on-premises AI governance

Many organisations need AI-governance software that can operate within their own infrastructure or private cloud. In Anchorpoint's April 2026 review of 13 funded specialists, none clearly offered a complete customer-controlled product. Few of the reviewed vendors met this need; the finding applies to the reviewed set, not the entire market.

Legend: incumbent / platform · AI-native / specialist · adjacency (GRC, privacy, observability) · UADP / control platform · open / under-served opportunity. Source: Anchorpoint Advisory ecosystem analysis, updated 2026. Company names are illustrative market participants, not an exhaustive or ranked list.

The funding dynamics

The AI TRiSM (Trust, Risk and Security Management) market is valued at $3.59B in 2025, projected to reach $46.8B by 2034 at a 35% CAGR; BFSI is the largest segment at 29.3% of revenue. A broader enterprise AI-governance estimate grows from $2.20B to $11.05B at a 15.8% CAGR through 2036. These estimates use different definitions and should not be combined. Venture funding demonstrates investor conviction, not customer adoption by itself. Investors have backed different starting products, including discovery, governance, agent security and broader platforms—before a standard product model has emerged.

SELECTED PRIVATE FINANCINGS

Latest 2025-2026 raises of $20M+ | round size, not cumulative funding

NewCore 2026 **$66M** Launch / seed	**Straiker** 2026 **$64M** Series A	**Arcade.dev** 2026 **$60M** Series A	**WitnessAI** 2026 **$58M** Strategic
Patronus AI 2026 **$50M** Series B	**Gray Swan AI** 2026 **$40M** Series A	**Onyx Security** 2026 **$35M** Series A | $40M cumulative	**JetStream** 2026 **$34M** Seed
Reco 2026 **$30M** Series B	**Geordie AI** 2026 **$30M** Series A	**Runlayer** 2026 **$30M** Series A	**NeuralTrust** 2026 **$20M** Seed
Noma Security 2025 **$100M** Series B	**Aurascape** 2025 **$50M** Aggregate funding at launch	**Koi Security** 2025 **$38M** Series A | acquired	**Nudge Security** 2025 **$22.5M** Series A

16 additional companies completed sub-$20M financings; acquired companies are marked in the financing ledger.

Selected, not exhaustive. Company-reported funding and investor rationale from Anchorpoint's ecosystem and financing review, current through early July 2026.

What investors are backing

- **AI creates a new area that enterprises must control.** Agents, model interactions and probabilistic execution create activity that conventional security, identity and governance products were not designed to handle.
- **Enterprise adoption is moving faster than oversight.** The gap creates both security exposure and a barrier to wider production use.
- **Autonomous agents increase the need for control.** The focus has shifted from protecting employees who use generative AI to controlling software actors with identities, credentials, memory, tools and delegated authority.
- **Agent identity and authorisation are becoming essential.** Enterprise agents that access data, credentials or software tools will generally need a verifiable identity, limited permissions, action-level accountability and revocable authority.
- **Visibility comes before enforcement.** Enterprises must first identify shadow AI, unmanaged agents, model usage, tool connections, permissions and sensitive-data exposure before they can apply policy.
- **Products that can intervene are more valuable than reporting-only tools.** Platforms that can allow, deny, modify, route, interrupt or require approval sit closer to the point of economic, operational and security exposure than those that only observe and report.
- **Ongoing testing is becoming more important than one-off assessments.** Investors are backing systems that repeatedly evaluate, simulate, monitor and retest AI behaviour as models, prompts, tools and workflows change.
- **A shared control system across platforms may become a major category.** Enterprises use multiple clouds, models, applications, SaaS copilots and agent frameworks, creating demand for one independent governance layer that works across them.
- **Regulation will increase demand for technical controls.** The strongest platforms turn regulatory requirements into enforceable policy, traceable decisions, continuous monitoring and audit-ready evidence.
- **The category may support expansion within customers.** Companies can enter through an urgent use case—such as shadow-AI discovery, red teaming, agent authentication or prompt protection—and later add identity, runtime enforcement, policy and enterprise-wide governance.
- **Products may draw funding from several departments.** Security, identity, IT, data, model risk, compliance and AI teams may all contribute, although unclear ownership can lengthen sales cycles.
- **The ability to approve, block or change an action is a primary source of strategic value.** Investors are backing platforms that do not merely describe risk but determine what an AI system is permitted to access, decide and execute.

The M&A backdrop

From late 2024 through mid-2026, large cybersecurity and technology companies repeatedly acquired specialist capabilities they had not built internally or could not bring to market quickly enough. This shows that buyers need these capabilities quickly even though the market is still taking shape.

Anchorpoint identified 19 core AI-control acquisitions announced from 1 January 2024 to 14 July 2026. Known directly disclosed or filing-derived consideration is approximately $3.56 billion. Including the reported Aim Security estimate raises the total to approximately $3.91 billion; several additional transaction values remain undisclosed.

Separately, three broader acquisitions that support AI control carry much larger values: CyberArk in identity, Chronosphere in observability and Moveworks in workflow automation. These assets matter because they add identity, monitoring, distribution and workflow capabilities around AI, but their full transaction values should not be attributed to AI governance alone.

Ranking charts use announced headline values for cross-deal display; core consideration totals and detailed text use final or filing-derived values where available. This explains the difference between CyberArk's approximately $25.0 billion announced equity value and $21.1 billion final purchase price, Chronosphere's $3.35 billion announced and $2.95 billion final values, and Moveworks' $2.85 billion announced and $2.4 billion final values.

The core acquisitions span data governance, runtime enforcement, model and application lifecycle security, agentic discovery, detection and response, and evaluation or observability. Each deal is assigned to its primary acquired capability, although several span multiple layers.

The deals suggest that incumbents are buying complementary controls, combining them with data they already hold on identities, endpoints, applications and workflows, and selling them through established enterprise platforms.

The pattern does not prove that every acquired product will become a durable platform, nor that transaction value measures the value of one control point in isolation. Revenue, customers, talent, competitive tension and broader franchises also influence consideration.

Repeated acquisitions across several control layers show that major buyers consider these capabilities important and time-sensitive.

M&A by control point[1]

Policy enforcement
Closest to execution; strongest monetisation
Applies rules before an AI system acts. It can approve, block or redirect actions, control access and preserve evidence of what happened.

Agent identity and permissions
Foundational and currently scarce
Controls which agents can act, which tools and data they may use, and how delegated authority is constrained or revoked.

AI discovery
Important control input; limited alone
Identifies where AI tools and agents are being used. Its value is limited unless the product also feeds policy, enforcement or approval.

Model monitoring and evaluation
Important control input; durable when tied to action
Measures performance and potential risk. It becomes more durable when its signals feed policy, release gates or runtime action.

[1] *A control point is a place in the AI system where a company can monitor activity, set rules or stop an unsafe action. It may sit around the AI model, the data it uses, the software tools it connects to, the identity of the user or agent, or the action it is about to take. Looking at M&A by control point shows which parts of the AI system buyers are trying to secure and control through acquisitions.*

M&A BY PRIMARY CONTROL POINT

Identified core transactions announced 1 Jan 2024 - 14 Jul 2026

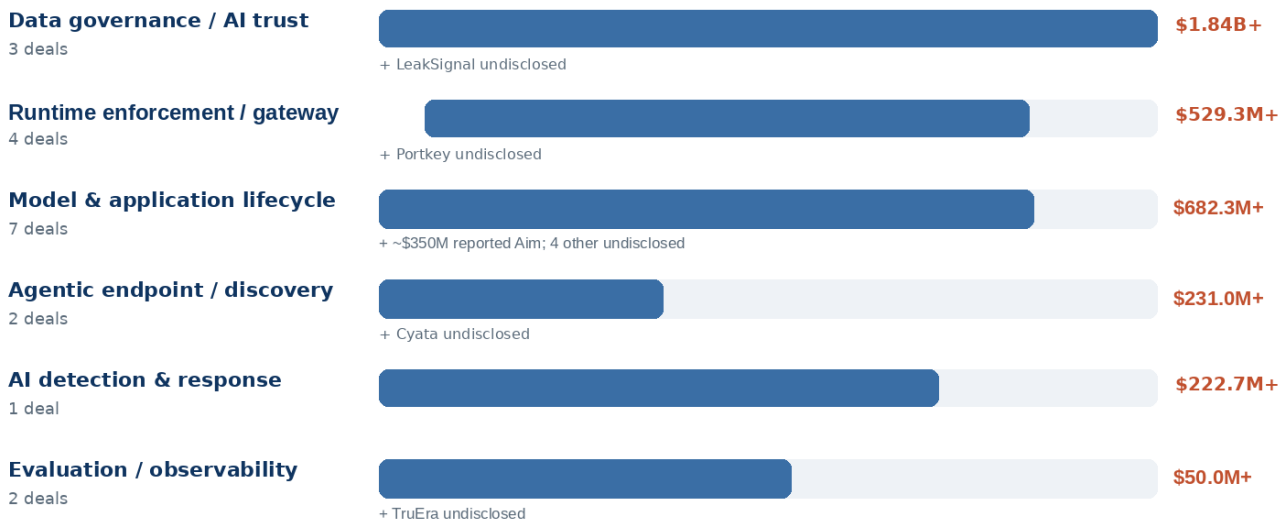

Known directly disclosed / filing-derived core consideration: ~$3.56B; including Reuters-reported Aim Security: ~$3.91B.

Key transactions

- **Cisco / Robust Intelligence (Sep 2024, undisclosed)** was one of the earliest transactions to show this strategy. Cisco later made the technology a foundation of AI Defense: acquire specialist AI-control technology, then distribute it through an established security platform.
- **Palo Alto Networks / Protect AI, Koi and Portkey** form the core AI-control acquisitions, moving from model lifecycle security to agentic endpoint control and an inline gateway for monitoring, routing and governing models and agents.
- **Palo Alto Networks / CyberArk ($25.0B announced; $21.1B final, completed Feb 2026)** supports identity as an AI-control layer. The value reflects CyberArk's broader identity business, not agent identity alone. Its relevance to AI is the potential extension of privileged-access controls from people and conventional software to autonomous agents.
- **Check Point / Lakera and Cyata** combine runtime guardrails and prompt-injection defence with discovery and governance for autonomous agents. Lakera controls AI interactions; Cyata discovers and analyses the agents operating inside an organisation.
- **CrowdStrike / Pangea ($222.7M total consideration, completed Sep 2025)** extends EDR logic into AI Detection and Response across data, models, agents, identities and infrastructure—from endpoint telemetry to the wider range of systems and actions an agent can reach.
- **SentinelOne / Prompt Security and F5 / CalypsoAI** show buyer interest in the interaction and runtime layers: discovery, data protection, prompt inspection, guardrails, red teaming and enforcement before an AI interaction reaches a sensitive business system.
- **ServiceNow / Moveworks ($2.85B announced; $2.4B final purchase consideration, completed 15 Dec 2025)** shows that workflow vendors also want to control how AI-generated actions are approved, executed and recorded.

AI Governance & Control Transactions

Core transactions | Jan 2024 – Jul 2026

M&A by sub-sector

Sub-sector	Deals	Representative targets	Disclosed value
Data security, privacy & AI trust	3	securiti, alltrue.ai, leaksignal	~$1.85B+
Runtime enforcement / AI gateway	4	CalypsoAI, LAKERA, Prompt Security, portkey	~$529M+
Agentic endpoint & agent discovery	2	Koi Security, CYATA	$300M + undisclosed Cyata value
AI detection & response	1	PANGEA	~$223M
Model & application lifecycle security	7	PROTECT AI, ROBUST INTELLIGENCE, SYDELABS	Undisclosed
AI observability & evaluation	2	TRUERA, aporia	Undisclosed
Identity (adjacent)	1	CYBERARK	~$25.0B
Observability (adjacent)	1	chronosphere	~$3.35B
Workflow / agent assistant (adjacent)	1	moveworks	~$2.85B

Adjacent deals extend identity, observability or workflow platforms into AI governance rather than being purpose-built AI-security acquisitions. Primary classification reflects the principal acquired control point. Ranking uses announced headline values; final consideration is shown in detailed text where available.

Selected transactions

CyberArk (PANW)	CYBERARK → paloalto	$25.0B
Chronosphere (PANW)	chronosphere → paloalto	$3.35B
Moveworks (ServiceNow)	moveworks → servicenow	$2.85B
Securiti AI (Veeam)	securiti → veeam	$1.725B
Koi Security (PANW)	Koi Security → paloalto	$300M
Pangea (CrowdStrike)	PANGEA → CROWDSTRIKE	$222.7M
Lakera (Check Point)	LAKERA → CHECK POINT	~$190M
CalypsoAI (F5)	→ F5	$180M
Prompt Security (S1)	Prompt Security → SentinelOne	$159.3M
AllTrue.ai (Varonis)	alltrue.ai → VARONIS	$114.5M
Cyata + Cyclops (Check Point)	CYATA + CYCLOPS → CHECK POINT	$92M

M&A transaction trends

- **Buyers are acquiring new areas of AI control.** Deals add visibility and authority over models, prompts, agents, identities, data access and tool execution—areas conventional security platforms were not built to govern.
- **Products that intervene during an interaction provide the clearest evidence of strategic value.** Lakera, CalypsoAI, Prompt Security, Portkey and LeakSignal enable buyers to block, modify, route, redact or approve AI interactions rather than merely report risk.
- **Agentic AI is changing acquisition priorities.** Koi, Cyata, Portkey and AllTrue provide discovery, identity, permission and runtime controls for autonomous software actors operating across enterprise systems.
- **Incumbents are beginning to build broader platforms.** Acquisitions fill gaps across discovery, testing, deployment, monitoring, enforcement, incident response and audit evidence—reducing dependence on narrow tools.
- **Data the buyer already holds improves acquired capabilities.** Buyers can combine AI controls with existing information on identities, permissions, endpoints, applications and sensitive data to make enforcement more accurate.
- **Monitoring and policy tools become more valuable when they can trigger action.** Securiti AI, TruEra, Aporia and AllTrue contribute policy, monitoring and evidence; their value increases when those signals feed decisions and enforcement.
- **M&A speeds product entry and customer access.** Large vendors acquire specialist teams, technology and customers faster than they can build internally, then sell the capability through established enterprise channels.
- **Platform buyers want to control the path from AI discovery to action and audit.** They are positioning themselves as the main layer through which enterprises discover AI, govern access, evaluate behaviour, control execution and prove compliance.

Four buyer groups, one control chain

Security vendors bring telemetry and threat response; identity vendors determine who or what may act; data and cloud platforms govern information access; workflow platforms govern approvals and execution. The four groups are moving towards overlapping parts of the same chain: inputs, decisions, enforcement and evidence.

Security platforms

Palo Alto Networks, Check Point, F5, CrowdStrike, Cisco and SentinelOne are extending established security platforms into model security, prompt inspection, runtime enforcement, agent discovery, endpoint exposure and AI gateways. Their advantage is the ability to combine acquired AI controls with existing endpoint, network, cloud and threat telemetry.

Cloud and data platforms

Veeam, Varonis, Snowflake and Coralogix show that buyers now include data, resilience and observability companies as well as cybersecurity vendors. Examples include Veeam/Securiti AI, Varonis/AllTrue.ai, Snowflake's announced acquisition of Natoma (agent identity and MCP governance, ~$110M, May 2026) and Coralogix's acquisition of Aporia (AI observability and guardrails, ~$50M, Dec 2024). They are combining AI governance with enterprise data, permissions, resilience, observability and recovery. Microsoft, AWS, Google Cloud and Databricks can pursue the same control layer organically, although native controls may favour their own infrastructure.

Workflow platforms

ServiceNow and adjacent enterprise application platforms can govern how AI-generated actions are approved, executed and recorded. They sit closest to the business process: what an agent requested, who approved it and what happened. ServiceNow's acquisition of Moveworks shows that workflow vendors also want to control this process.

Identity and access platforms

Human, machine, workload and agent identities increasingly require unified credentials, permissions, monitoring and revocation. Identity could become the foundational AI-control layer because every governed action must be attributable to an authorised actor. Palo Alto Networks' CyberArk acquisition explicitly extends identity security across humans, machines and agents.

Most active strategic acquirers

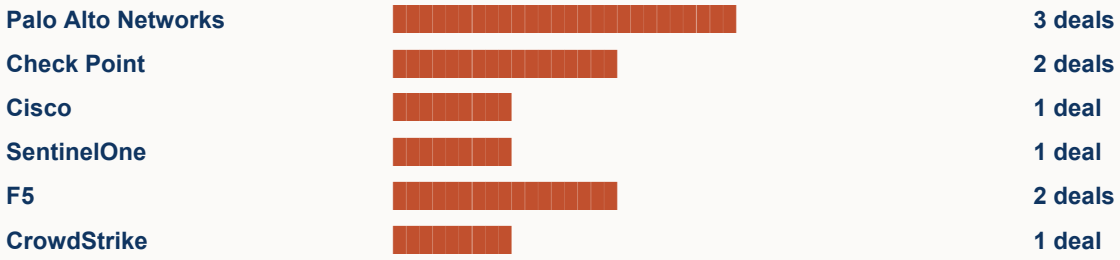

Core AI-control transactions only (Sep 2024–May 2026); adjacent enabling acquisitions—Palo Alto Networks / CyberArk and Chronosphere, and ServiceNow / Moveworks—are excluded from this count and discussed separately.

The market is moving beyond single-purpose acquisitions. Palo Alto Networks, Check Point and F5 are beginning to assemble multi-layer AI-control platforms: Palo Alto spans lifecycle security, agentic endpoints and gateways; Check Point combines runtime protection with agent discovery and governance; F5 combines data-in-transit protection with inference-time guardrails and red teaming.

Most other buyers have completed one entry transaction. The market remains unconsolidated, but competition is shifting towards which vendor can control the greatest share of the path from identity and data access to action and audit.

CASE STUDY, PALO ALTO NETWORKS

Palo Alto Networks' emerging AI-control platform

Palo Alto Networks provides the clearest example of a security company building a broader AI-control platform. It has made three acquisitions directly focused on AI control—Protect AI, Koi and Portkey—and two broader acquisitions that add identity and observability: CyberArk and Chronosphere. Together they outline, but do not yet prove, a five-layer platform.

Control point	Target	Capability
Model lifecycle	Protect AI ($634.5M, completed Jul 2025)	Scanning, posture, supply chain, red teaming
Agentic discovery / endpoint	Koi Security ($300M announced; $231M final, completed Apr 2026)	Coding agents, plugins, packages and local credentials
Runtime enforcement / gateway	Portkey (undisclosed, May 2026)	Routing, orchestration, monitoring and agent governance
Identity (adjacent enabling)	CyberArk (~$25.0B announced; $21.1B final, Feb 2026)	Human, machine and agent permissions and privileged access
Observability (adjacent enabling)	Chronosphere ($3.35B announced; $2.95B final, Jan 2026)	Telemetry, monitoring and evidence for AI-native systems

If successfully integrated, the value would come less from owning five separate products and more from coordinating context and enforcement across them. A model vulnerability detected by Protect AI could inform gateway policy; CyberArk could determine which agent is entitled to act; Portkey could enforce that entitlement; Koi could identify risky local tools or credentials; and Chronosphere could provide the telemetry required to investigate performance, failure and policy outcomes.

Palo Alto Networks has described these capabilities as complementary: Portkey provides centralised agent governance, Koi controls the agentic endpoint, CyberArk underpins agent identity and authorisation, and Chronosphere supplies high-fidelity observability. Its installed base creates a cross-selling opportunity, but product integration and common policy execution remain the key tests.

OUTLOOK

The conditions favour continued activity

The market still consists mainly of specialist products, but larger vendors are beginning to combine them into broader suites. Regulatory requirements are becoming clearer, AI agents are gaining access to more systems, and incumbents still lack important controls.

Build, partner or buy

Build Best where the control is native to an existing platform, tightly coupled to proprietary telemetry, and not strategically urgent.	**Partner** Best where neutrality and interoperability matter, but the capability is not sufficiently scarce to justify acquisition.
Buy Best where the missing control must reach market quickly, sits in the execution path, and requires architecture or talent the incumbent cannot reproduce in time.	**Independent platform** Plausible where the vendor coordinates multiple control systems, remains cross-platform, owns evidence and serves deployment needs incumbents underserve.

A plausible market sequence

1. **Separate products.** Discovery, prompt security, model scanning, agent inventory and gateways emerge separately.
2. **Customers ask for integrated controls.** Enterprises begin to coordinate policy, permissions and evidence in the execution path.
3. **Consolidation.** Security, identity, cloud, data and workflow platforms acquire or bundle specialist controls.
4. **Possible end state.** A few broad platforms coexist with specialists that serve regulated markets, work across vendors or own hard-to-build technology.

Recent financings and acquisitions suggest a possible shift from narrow specialist products towards broader platforms, but the outcome is not yet settled. Signs of progress would include agent-specific identity products, gateways adding approvals and revocation, and platforms applying common policy and evidence across environments. The strongest commercial opportunities are likely to be in runtime enforcement, delegated permissions and controls that work across multiple platforms.

Bottom line

Better models are necessary. Enterprise value appears when AI can be embedded into workflows, trusted with limited and clearly defined authority, measured for productivity and audited for accountability. The most valuable capability in the next phase is not intelligence alone. It is the ability to convert intelligence into controlled action.

ABOUT ANCHORPOINT ADVISORY

Clarity at critical inflection points

Anchorpoint is an independent technology investment advisory firm. We advise founders, boards, shareholders, investors and family offices on strategic advisory, M&A and strategic financing, and family office services across the UK, Europe and the Middle East.

15+
years advising technology companies

$6bn+
transaction value across M&A and strategic financing

50+
completed transactions

17
countries, ~60% cross-border execution

The Anchorpoint Method — Orient. Align. Act.

Orient Establish the true position: where the company stands, what creates strategic advantage, and which choices are genuinely available.	**Align** Create a common basis for decision: the objective, positioning, decision criteria and valuation expectations stakeholders must agree before execution.
Act Translate strategy into execution: reposition the company, approach capital partners, and run the financing or M&A process while preserving leverage and optionality.	**Where we focus** AI, Software & Cybersecurity; Deep Technology & Advanced Computing; Industrial Automation & Intelligent Systems; Energy, Infrastructure & Climate Systems; Defence, Security & Resilience.

Methodology and disclosure

This report has been prepared by Anchorpoint Advisory Ltd solely for general information and discussion purposes. It does not constitute investment, financial, legal, tax, accounting or other professional advice, nor does it constitute an offer, invitation, recommendation or solicitation to buy, sell or subscribe for any security, financial instrument, business or asset.

The report does not take account of the investment objectives, financial circumstances, risk tolerance or particular needs of any recipient. Recipients should form their own independent judgement and obtain appropriate professional advice before making any investment, financing, acquisition or strategic decision.

The information contained in this report has been obtained from public sources, company disclosures and other materials believed to be reliable. However, Anchorpoint Advisory has not independently verified all such information. No representation or warranty, express or implied, is made as to its accuracy, completeness, timeliness, consistency or fitness for any particular purpose. Estimates, market data, transaction values and product capabilities may be based on company-reported, third-party or publicly available information and may be incomplete or subject to revision.

This report may contain opinions, estimates, projections and forward-looking statements. These reflect Anchorpoint Advisory's judgement as at the date of publication and are subject to assumptions, risks and uncertainties that may cause actual outcomes to differ materially. Anchorpoint Advisory is under no obligation to update or revise this report following its publication, except where required by applicable law.

Nothing in this report should be relied upon as a substitute for the recipient's own investigation, due diligence or professional judgement. To the fullest extent permitted by law, Anchorpoint Advisory and its directors, employees, advisers and affiliates accept no liability for any loss arising from the use of, or reliance on, this report or its contents.

Companies, transactions and financings included in this report are selected for illustrative purposes and do not constitute an exhaustive market review, ranking or endorsement. Anchorpoint Advisory may have advised, invested in, or maintained commercial relationships with companies or participants operating in the broader market discussed in this report. Any such relationships do not alter the analytical basis of the report.

anchor-point.co.uk | Prepared by Anchorpoint Advisory Ltd | July 2026