PRODUCTBEACON

ProductBeacon — State of Cyber Security Markets 2026, Grand Unification: One Force, Two Wars

Yohay Etsion

Managing Director, ProductBeacon

2026-06-21

Disclosure: *This is ProductBeacon's independent, impartial market research, based solely on publicly verifiable signals, published on the open web. No vendor sponsors, no paywalled data, no analyst-firm reuse. It reflects ProductBeacon's independent view and is not investment advice. Full independence and sourcing methodology: Methodology →.*

By Yohay Etsion · Head of Product (Fractional), AXIA · Creator of Product Org OS · Author of Leading the Charge (2023) and Vision to Value (coming 2026)

AXIA is an AI-powered data-security (Insider Risk Management / DLP) company, which competes in the DLP segment covered in this report.

10.1 The Unifying Frame

The question that has to be answered before this report can call itself one report is the simplest and the most dangerous one in the whole project: is this one report, or is it two? Part 1 marched through the data war — Insider Risk Management, Data Loss Prevention, Data Security Posture Management — and resolved on the data layer. Part 2 marched through the wider platform theater — the SOC, the Secure Service Edge, AI Security, Identity — and resolved on agent identity. That is two arcs, two convergence chapters, two chosen theses that do not name the same core capability.

The honest default a disciplined reader should hold is that these are two market stories that happen to share a few large names. On that read, any attempt to braid them into a single capstone is the analyst manufacturing a thesis the evidence does not carry. I take that default seriously enough to put it on the table first, in its own words, before I say what I am betting against it.

Three theses were placed on the table before a word of this chapter was written, and the capstone discipline requires that I quote all three verbatim and attach my read to the one I am betting on. The chosen thesis:

> **The Shared-Incumbency Thesis.** The State of Cyber 2026 report is one report, not two, because the same handful of platform incumbents — **Microsoft, Palo Alto Networks, and CrowdStrike above all** — are winning *both* wars it documents: the Part-1 data war (Insider Risk, Data Loss Prevention, Data Security Posture Management) and the Part-2 platform war (the SOC, the Secure Service Edge, AI Security, and Identity). And they are winning both for *one* reason, which is why the two arcs are a single story: the autonomous AI agent is simultaneously a **data-access problem** (Part 1's question — where sensitive data sits, what is moving, who can reach it) and an **identity-and-detection problem** (Part 2's question — who the actor is, what it may authorize, and the alert it generates), and only a platform that already spans *both* the data layer and the identity/SOC layer can govern an agent end to end… The report is one report because the winners are the same, and the winners are the same because the problem is one problem.

[Editor's note — read before quoting the paragraph above: this is the Chosen Thesis exactly as pre-committed, quoted verbatim per the capstone discipline and then qualified below. The claim this chapter actually argues is the bounded one — the same FORCE drives both arcs (the autonomous agent and the 80-to-1 identity inversion), expressed as two DIFFERENT enforcement primitives: data-source enforcement in Part 1, agent-identity governance in Part 2. The report is one report because the force is one force, not because the two wars are one problem.]

The null, which I owe the reader at full strength:

> **The Two-Unrelated-Arcs Thesis.** There is no unifying force; Part 1 and Part 2 are two separate market stories that happen to share some large names because large names are everywhere… the incumbent overlap is a statistical artifact of company size, not evidence of one contest: Microsoft, Palo Alto, and CrowdStrike show up in both reports for the same reason they show up in every cybersecurity sub-market — they are the biggest acquirers with the broadest portfolios, and a big company appearing in two adjacent markets no more proves the markets are one war than Amazon being in both retail and cloud proves retail and cloud are one market.

And the alternative, which differs from the chosen only on the driver, not the outcome:

> **The Pricing-Power (Not Convergence) Thesis.** The same incumbents do win both wars — but the unifying force is *commercial*, not architectural. It is a margin-and-distribution story: Microsoft, Palo Alto, and CrowdStrike win both the data battlegrounds and the platform theater because they own the broadest bundles and the deepest distribution into the enterprise… the E5 estate that already ships Purview (Part 1) also ships Defender, Sentinel, and Entra (Part 2), so the marginal cost of adding each adjacent module trends toward zero and the incumbent wins by default-distribution on every front.

My read is that the report is one report. But I am qualifying the pre-commit in two places before I argue from it, exactly where the gate that cleared this capstone demanded. The chosen thesis as drafted said the two wars are "one problem." That is the over-reach I have to cut, and cutting it is the whole gate. The defensible claim is *one force, two enforcement primitives* — not one problem, one primitive.

The autonomous agent and the 80-to-1 machine-vs-human identity inversion are the **same upstream force** driving both arcs. But the two arcs resolve on two genuinely *different* enforcement primitives: data-source enforcement in Part 1, agent-identity governance in Part 2. Asserting they are the same problem hands the null its strongest counter, because Part 1's own convergence chapter chose the *data* primitive, not the identity one. The first qualification, then: same force, not same problem.

The second qualification: the win is bounded. The shared incumbents win the mid-market and greenfield decisively across both reports. The sophisticated enterprise tier is genuinely contested in both, and a capstone that claims the enterprise is settled is falsified by both convergence chapters' own buyer's-decision sections. So I am betting on shared incumbency as the unifying force of the whole report — but on the bounded, two-primitive, mid-market-decided form. The rest of this chapter argues from that bounded claim and names exactly what would break it.

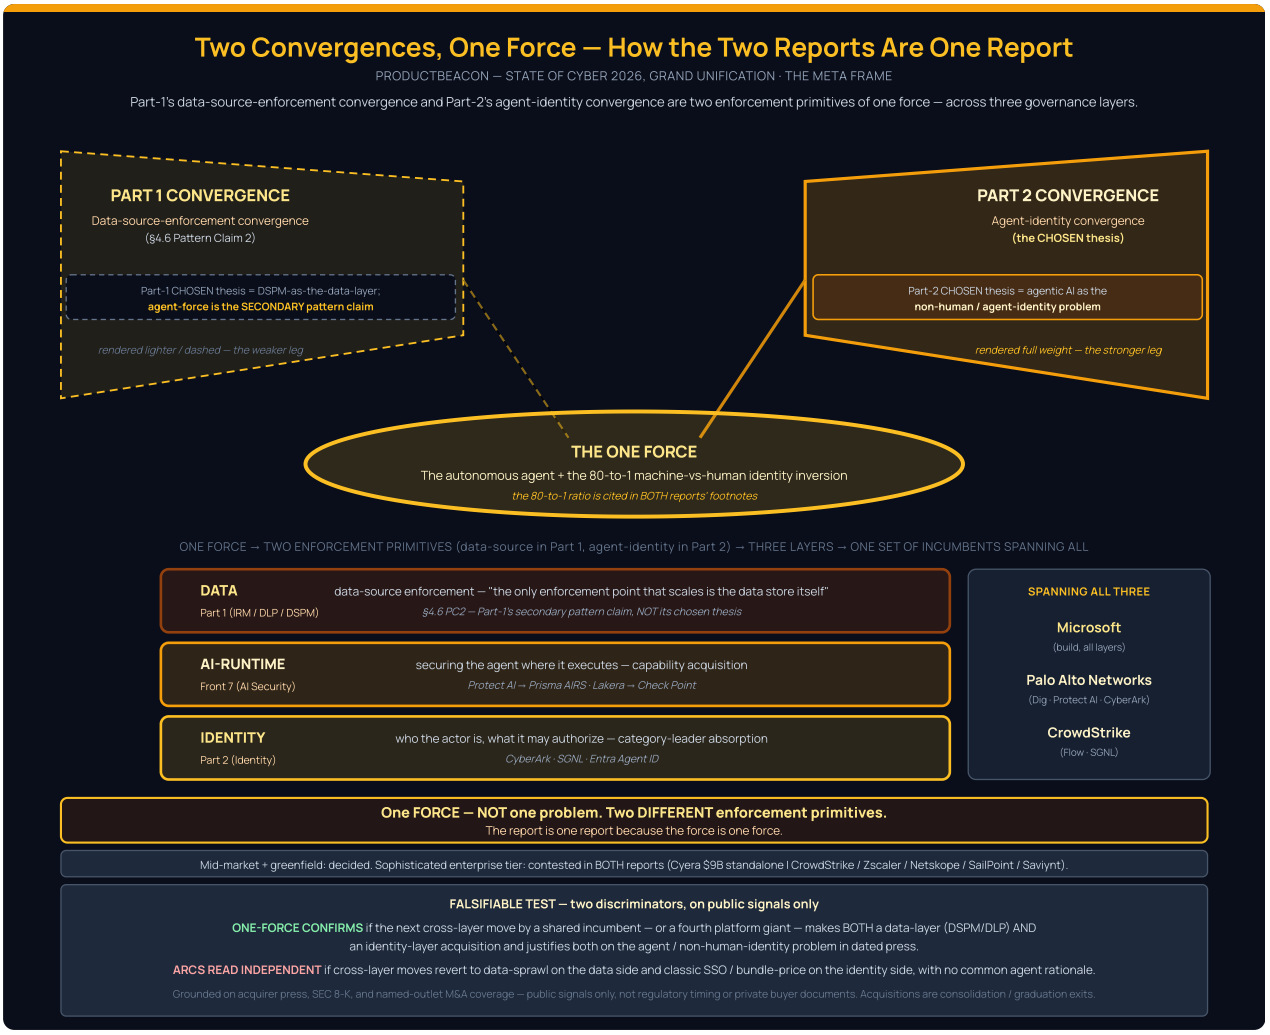

Caption: Part-1's data-source-enforcement convergence and Part-2's agent-identity convergence are two enforcement primitives of one force — the autonomous agent plus the 80-to-1 inversion — expressed across three governance layers. One force, not one problem; the report is one report because the force is one force.

10.2 One Force, Three Primitives

The architectural core of the capstone is a single claim about a single force. That force is the autonomous AI agent and the **80-to-1 machine-vs-human identity inversion**. The figure comes from Palo Alto Networks, cited in its CyberArk closing release: the enterprise's actor population is now overwhelmingly non-human, machine identities outnumbering human roughly eighty to one (developed in the Part-1 Force-1 record and the Part-2 identity convergence).[1] The precise multiple is a vendor estimate and is definition-sensitive: the identity chapter notes that no two firms agree where the machine-identity boundary falls. So the load-bearing claim here is the direction, not the exact ratio. Machine identities now vastly outnumber human ones, and the gap is widening; whether the figure is eighty to one or some other large multiple does not change the capstone.

That number is the hinge of the whole report, and it carries weight in *both* arcs. It is footnoted in Part 1's §4.2 Force-1 record and again in Part 2's §9.6 Pattern Claim 1 — the same dated release on both sides. An autonomous agent is, in the same instant, a thing that touches data and a thing that must be identified and authorized. That is Part 1's question (where sensitive data sits and who can reach it) and Part 2's question (who the actor is and what it may do). One force, asked as two governance questions on two layers.

That one force expresses as **three governance primitives**. The discipline of this chapter is to draw them cleanly and never collapse them into one undifferentiated "agent problem."

The first is **data-source enforcement** — Part 1's primitive. Part 1's convergence chapter put it in a single sentence I quote verbatim, because it is the bridge everything rests on between the two reports: "when the actor is a machine identity at machine speed, the only enforcement point that scales is the data store itself" (see the data chapter's second pattern claim).[2] When the agent population grows faster and operates further from human supervision than runtime controls can credibly govern, enforcement gets pulled back to where the data sits. That is the data primitive: govern the agent by governing the store it reaches into.

The second is the **AI-runtime** primitive — the frontier that Front 7 owns. This is the agent secured where it executes: is the model and its runtime safe, can the agent be governed at the point of its own behavior. Palo Alto bought into this layer when it folded Protect AI into its Prisma AIRS runtime line. Check Point bought into it when it acquired Lakera as the foundation of its AI-security center of excellence (see the Part-2 identity convergence).[3]

The third is **agent identity** — Part 2's primitive: who the actor is, what it may authorize, the alert it generates. This is the vector Part 2's convergence chose, the one CyberArk-into-Palo-Alto and SGNL-into-CrowdStrike resolve on. Data, runtime, identity — three primitives, one force, and only a platform spanning all three can govern an agent end to end.

I have to state one asymmetry honestly, because it is the capstone's structural soft spot. Burying it would be the analyst manufacturing symmetry the evidence does not carry. **The agent-force is not the same kind of claim on both legs.** In Part 2, the agent-identity vector is the *chosen thesis* — the convergence chapter argues it directly, the headline of the whole arc. In Part 1, the chosen thesis is something else entirely: DSPM as the underlying data layer, a *data-economics* claim that the data layer becomes the anchor because it is the only capability all three categories ultimately need.

The agent-force reading in Part 1 is not the chosen thesis. It is a secondary cross-front pattern claim — §4.6 Pattern Claim 2, "Agentic AI Pulls Enforcement Back to the Data Source" — that runs *underneath* Part 1's three theses rather than sitting on top as the headline. So the capstone's bridge into Part 1 runs through that pattern claim, not through Part 1's chosen thesis. The honest framing: the Grand Unification **elevates a Part-1 cross-front pattern claim to capstone level, alongside Part-2's chosen thesis.**

Part 1's own author wrote, verbatim, that agentic AI is the force pulling enforcement to the data source. So the through-line is real, dated, and Part-1-authored. But it is the headline of one arc and the secondary current of the other, and I will not write that "Part 1 was about agent identity." It was not. Part 1 was about the data layer; the agent-force is the pattern claim running beneath it. That is precisely the bridge the capstone needs, and no more than the bridge it can defend.

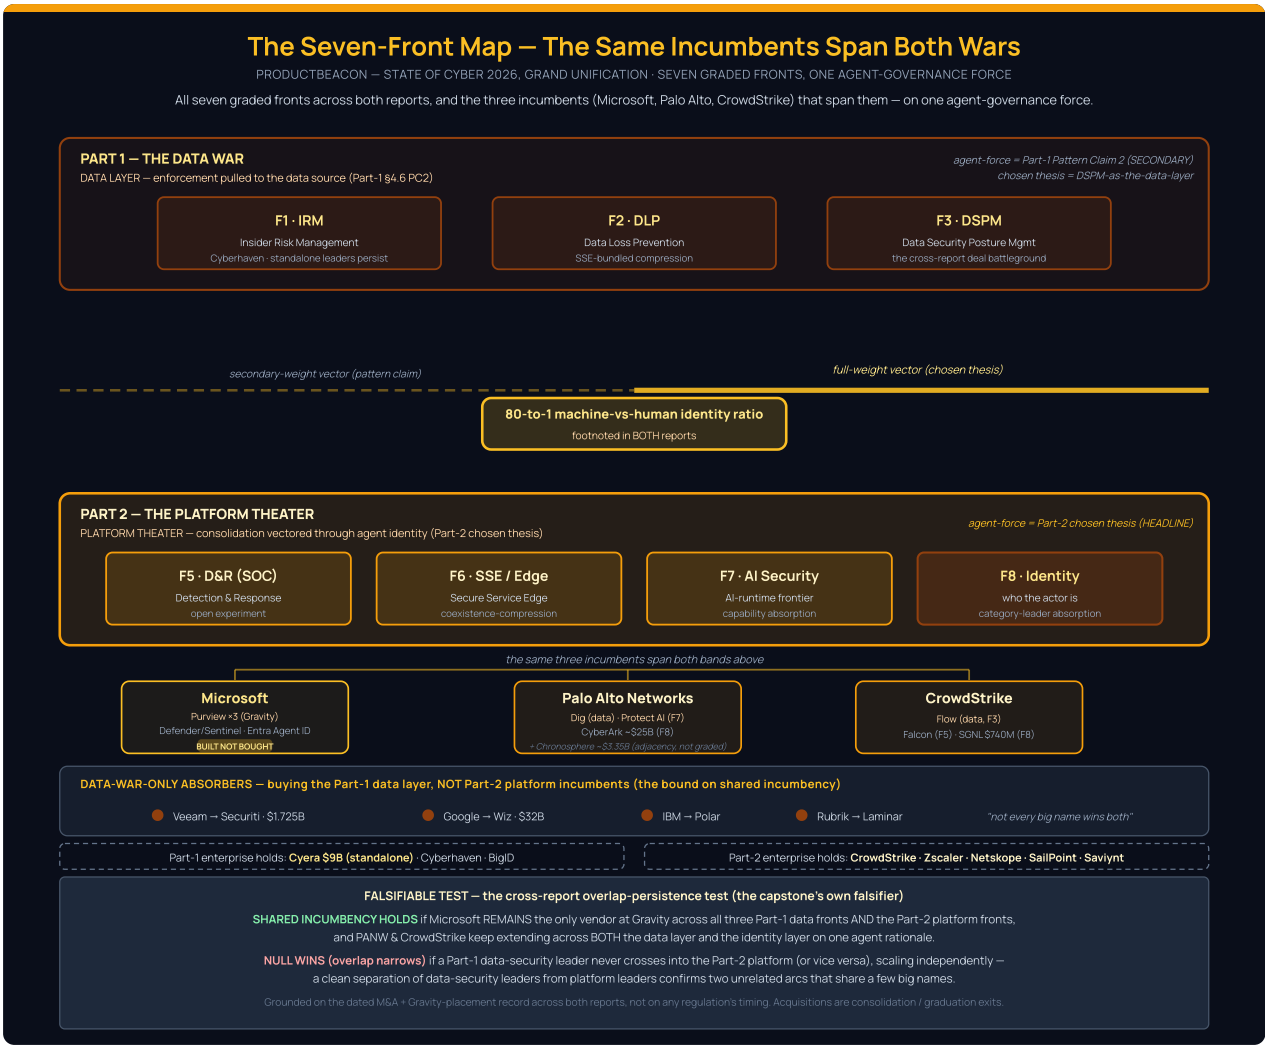

Caption: all seven graded fronts across both reports, the three shared incumbents rendered as pillars piercing both bands, and the 80-to-1 hinge pinned where the spine crosses between the data layer and the platform theater. The Part-1 agent-vector is drawn at secondary weight (pattern claim) and the Part-2 vector at full weight (chosen thesis) — the leg-asymmetry made visible.

10.3 The Shared Incumbents

The evidence that turns two reports into one report is a single exhibit: three names, each with a Part-1 data-layer move *and* a Part-2 platform move, justified on one rationale.

Microsoft holds Gravity-tier placement across all three Part-1 data fronts — Purview ships IRM, DLP, and DSPM as modules of the same M365 E5 estate. It anchors the Part-2 platform fronts too: Defender and Sentinel in the SOC, Entra Agent ID in Identity, Defender-for-Cloud in AI. Microsoft reaches the agent-identity destination by building. Entra Agent ID makes agent identities a first-class directory construct inside Agent 365 (see the Part-1 and Part-2 convergence chapters).[4]

Palo Alto Networks bought into the Part-1 data layer with **Dig Security (DSPM)**, one of Part 1's six DSPM-relevant M&A events. It then moved across the Part-2 theater, closing **CyberArk (~$25B, 2026-02-11)** and folding **Protect AI** into Prisma AIRS (see the Part-1 and Part-2 convergence chapters).[5] **CrowdStrike** bought into the Part-1 data layer with **Flow Security (DSPM)** and then the Part-2 platform with **SGNL (~$740M reported, 2026-01-08)** (see the Part-1 convergence and the Part-2 identity convergence).[6] Three names, both wars.

The skeptic's objection is the hardest thing in this analysis, and it deserves its strongest form before I answer it. It is the report's own "battery-firm" argument lifted one level up:

> *Microsoft, Palo Alto, and CrowdStrike show up in both reports for the same reason they show up in every cybersecurity sub-market — they are the biggest acquirers with the broadest portfolios. A big company appearing in two adjacent markets no more proves the markets are one war than Amazon being in both retail and cloud proves retail and cloud are one market.*

That is the null hypothesis at full strength, and the mere fact of overlap does not refute it. Coincidental scale really is the default explanation for why the three largest, best-capitalized acquirers turn up everywhere. Shared incumbency on its own proves nothing about whether the two arcs are one contest.

What breaks the analogy is not the *fact* of the deals. It is the **disclosed rationale across two architectural layers**. Amazon never announces that retail and cloud are one problem. Palo Alto and CrowdStrike announce the equivalent, in their own dated press, across two layers.

Palo Alto justified the *identity*-layer move — CyberArk — on the rationale that it can now "secure every identity across the enterprise — human, machine, and agentic," citing the 80-to-1 ratio in the same closing release. That identity logic is the acquirer's own dated words: the 80-to-1 agent explosion. Its data-layer move (Dig, a DSPM acquisition) is tied to that same agent force not by Palo Alto's Dig press but by this report's Part-1 §4.6 Pattern Claim 2. That claim reads the data-source-enforcement pivot as the agent-driven response. The identity side comes in the acquirer's words; the data side comes in the report's reading of the *same fact seen on two layers* (as the data chapter shows, and the Part-2 identity convergence).[5]

CrowdStrike corroborates the identical pattern. Flow sits on the data side, SGNL on the identity side. SGNL grants and revokes access for "human, non-human (NHI), and AI identities" in real time: the same triad, the same window, the same buyer (see the Part-2 identity convergence).[6]

The 80-to-1 ratio is the literal hinge. It is footnoted in *both* reports — the one number that ties a data-source-governance problem and an agent-identity-governance problem to one fact: the enterprise's actor population is now overwhelmingly non-human.

I am not asserting the acquirers are *correct* that the two markets are one; markets are not defined by press releases, and a stated rationale can be post-hoc narrative. I am asserting the weaker, sturdier claim that defeats the null. The same acquirers converge on the same agent rationale across two layers, in one ~13-month window (from Palo Alto's Protect-AI intent in April 2025 to its CyberArk close in February 2026), hinged on a single number both reports independently found central. That is a dated public signal that the cross-report overlap is *vectored* — pointed at one force — not coincidental scale.

Opportunistic roll-ups share a balance sheet; vectored ones share a thesis. The Amazon analog never tells you retail and cloud are one problem. These acquirers tell you, on two layers, that data governance and identity governance are two faces of the agent.

The pricing-power thesis is the harder competitor to dismiss, because it predicts the *same* observable outcome — shared incumbency — and differs only on the driver. And it has real, dated evidence: the **E5 estate that ships Purview in Part 1 is the same E5 estate that ships Defender, Sentinel, and Entra in Part 2.** One enterprise agreement, one distribution surface, the marginal cost of each adjacent module trending toward zero across both layers. I do not deny that. I absorb it as a **co-driver.**

Bundle economics is real, and it explains *how fast and how cheaply* the shared incumbency runs across both arcs. The bundle is the distribution mechanism. What it does not explain is *why the same agent rationale appears in the acquirers' M&A press across two different architectural layers.* A pure pricing-power story predicts the incumbents would justify each acquisition on suite-attach economics. Instead they justify the identity acquisitions on agent identity and the data acquisitions on agent-driven data sprawl. The bundle accounts for the speed of attach; it cannot account for the rationale convergence.

So this is a both/and, not a dichotomy — there is no "not because of the bundle" here. The bundle is the distribution mechanism; the agent is *why the categories are worth assembling into one platform at all.* The shared incumbents win because they can bundle cheaply *and* because the agent makes the categories worth bundling. The second half is the half the pricing-power thesis cannot supply on its own.

And the win is **bounded** — this is the second qualification of the pre-commit, and the evidence forces it. The shared incumbents win the **mid-market and greenfield decisively** across both reports, where the E5 estate is already paid for and a small security team has every reason to retire consoles. But the **sophisticated enterprise tier is genuinely contested in both arcs**, and a capstone that claims it is settled is falsified by both convergence chapters' own buyer's-decision sections.

On the Part-2 side, best-of-breed scales independently in public material. CrowdStrike, Zscaler, and Netskope command premium public multiples as focused platforms; SailPoint re-IPO'd into a $1.16B-ARR governance business; Saviynt took a $700M KKR round at roughly $3B (see the Part-2 convergence).[7] On the Part-1 side, the mirror holds: **Cyera reached $9B post-money in January 2026 and stayed standalone** — not absorbed — the data war's enterprise best-of-breed pole, alongside Cyberhaven and BigID (see the Part-1 convergence).[8]

The honest read of the deal field underneath sharpens the bound further, because the Part-1 data layer is consolidating from four different starting positions, and several of its biggest absorbers — Veeam (Securiti, $1.725B), Google (Wiz, $32B), IBM (Polar), Rubrik (Laminar), Proofpoint (Normalyze) — are *not* the Part-2 platform incumbents (see the Part-1 convergence).[8] That is exactly why the shared-incumbency claim is bounded and not "every big name wins both." Only Palo Alto and CrowdStrike, plus Microsoft on the build side, demonstrably span both reports' M&A and Gravity record *and* carry the agent rationale across both layers. The enterprise tier is contested in both wars; the standalone leaders persisting on both sides are the evidence.

Cross-Report Pattern Claims

Two claims close the Grand Unification, and they are the only two claims in the whole report that span Part 1 (the IRM / DLP / DSPM data war) *and* Part 2 (the SOC / SSE / AI Security / Identity platform war). Neither could be made inside either report alone. Part 1's convergence chapter resolved at the data layer and explicitly deferred the cross-report question. Part 2's convergence chapter named the shared incumbency in one bridge sentence and handed the claim forward to this capstone — as that chapter put it, whether the cross-report shared incumbency is itself a thesis is the Grand Unification's claim to argue, not Part 2's.

The two claims below are the discharge of that deferral. Each follows the report's proven Observation → My read → Conditional prediction → Sources structure, with a co-located diagram and a falsifiable-test footer grounded on public signals only.

Both claims are bounded by the capstone discipline the CV-1 gate fixed: the defensible inference is that **the same force drives both arcs**, not that the two wars are the same problem. The two arcs resolve on two *different* enforcement primitives — data-source enforcement (Part 1's chosen answer) and agent-identity governance (Part 2's chosen answer). The force that drives both is one force: the autonomous AI agent and the 80-to-1 machine-vs-human identity inversion.

The win is bounded too. The shared incumbents win the mid-market and greenfield decisively across both reports; the sophisticated enterprise tier is genuinely contested in both. And the through-line is asymmetric across the two legs — agentic AI is Part 2's *chosen* vector but Part 1's *secondary* pattern claim (§4.6 PC2), not Part 1's chosen thesis. Both claims state that asymmetry rather than hide it.

Observation. - The same three names span both reports. They span them not by coincidental scale but by *vectored* buying, justified in the acquirers' own dated press on one rationale. - **Microsoft** holds Gravity-tier placement across all three Part-1 data fronts: Purview ships IRM, DLP, and DSPM as modules of the same M365 E5 / Purview Suite SKU (see the Part-1 convergence). It *also* anchors the Part-2 platform fronts — Defender/Sentinel in the SOC, Entra Agent ID in Identity, Defender-for-Cloud in AI (see the Part-2 convergence).[9] - **Palo Alto Networks** bought into the Part-1 data layer with **Dig Security (DSPM)**, named in Part-1's six DSPM-relevant M&A events. It then moved across the Part-2 theater, closing **CyberArk (~$25B, 2026-02-11)** on the stated rationale that it can now "secure every identity across the enterprise — human, machine, and agentic," and folding **Protect AI** into its Prisma AIRS runtime line (see the Part-1 and Part-2 convergence chapters).[10] - **CrowdStrike** bought into the Part-1 data layer with **Flow Security (DSPM)**, then the Part-2 platform, acquiring **SGNL ($740M, 2026-01-08)** to grant and revoke access for "human, non-human (NHI), and AI identities" in real time (see the Part-1 and Part-2 convergence chapters).[11] - The hinge that makes this one pattern and not three coincidences is a single figure that carries weight in *both* reports: the **80-to-1 machine-vs-human identity ratio**, cited in the same PANW/CyberArk closing release that appears in Part-1's §4.2 Force-1 footnote and in Part-2's §9.6 Pattern Claim 1.[12]

My read. - I read the cross-report overlap as vectored, not coincidental — and that is the whole answer to the disciplined skeptic, who deserves to be stated at full strength first. The skeptic's best line is the report's own "battery-firm" argument lifted one level up: *Microsoft, Palo Alto, and CrowdStrike show up in both reports for the same reason they show up in every cyber sub-market — they are the biggest, best-capitalized acquirers, and Amazon being in retail and cloud does not make retail and cloud one market.* - That null hypothesis is not refuted by the mere fact of overlap; on its own, shared incumbency proves nothing. What breaks it is the **disclosed rationale across two architectural layers**. Coincidental-scale acquirers do not justify a *data-layer* move (Dig, Flow — data sprawl driven by AI agents) *and* an *identity-layer* move (CyberArk, SGNL — the 80-to-1 agent-identity explosion) with the *same* agent rationale in the *same* dated press. - Amazon never announces that retail and cloud are one problem; PANW and CrowdStrike announce — in their own words — that the agent is the common object spanning both layers. Opportunistic roll-ups share a balance sheet; vectored ones share a thesis. That is the inference the skeptic cannot touch. - But I bound it twice. The win is a mid-market-and-greenfield win, not an enterprise rout: standalone leaders persist on both sides — Cyera at $9B, Cyberhaven, BigID in Part 1; CrowdStrike, Zscaler, Netskope, SailPoint, Saviynt in Part 2. And the through-line is asymmetric — it is Part 2's chosen thesis but only Part 1's secondary pattern claim, so the report is one report because the *force* is one force, not because the two wars are one problem.

Conditional prediction. *This claim is falsifiable against two competing reads on public signals only.* - ***Force-test (vs the null — two unrelated arcs):*** *if the next cross-layer move by a shared incumbent — or by a fourth platform giant beyond Microsoft / PANW / CrowdStrike — carries ONE agent-governance rationale spanning data AND identity in dated acquirer press, the shared force is confirmed and the coincidental-scale null falsifies. If instead the incumbents' data moves are justified on data-sprawl/compliance and their identity moves on classic SSO/PAM with no common agent rationale, the arcs read as independently driven and the null wins.* - ***Driver-test (vs pricing-power):*** *if cross-report consolidation continues to be justified on the human-*

machine-agentic governance problem, the architectural through-line holds. If it reverts to suite-attach economics and E5-bundle pricing with no agent-architecture rationale, the bundle is the better explanation and the unification resolves at the commercial layer. - ***Cross-report overlap-persistence:*** *if Microsoft stays Gravity-across-both and PANW/CrowdStrike keep extending across both layers, the pattern confirms; if data-security leaders and platform leaders cleanly separate and scale independently, the null wins. Primary source: acquirer press releases, SEC Form 8-K filings, earnings disclosure, and named-outlet M&A coverage — never regulatory dates or private signals.*

Sources. [9] [10] [11] [12] (Part-1 §4.1 + §4.2 + §4.3; Part-2 §9.3 + §9.6 PC1)

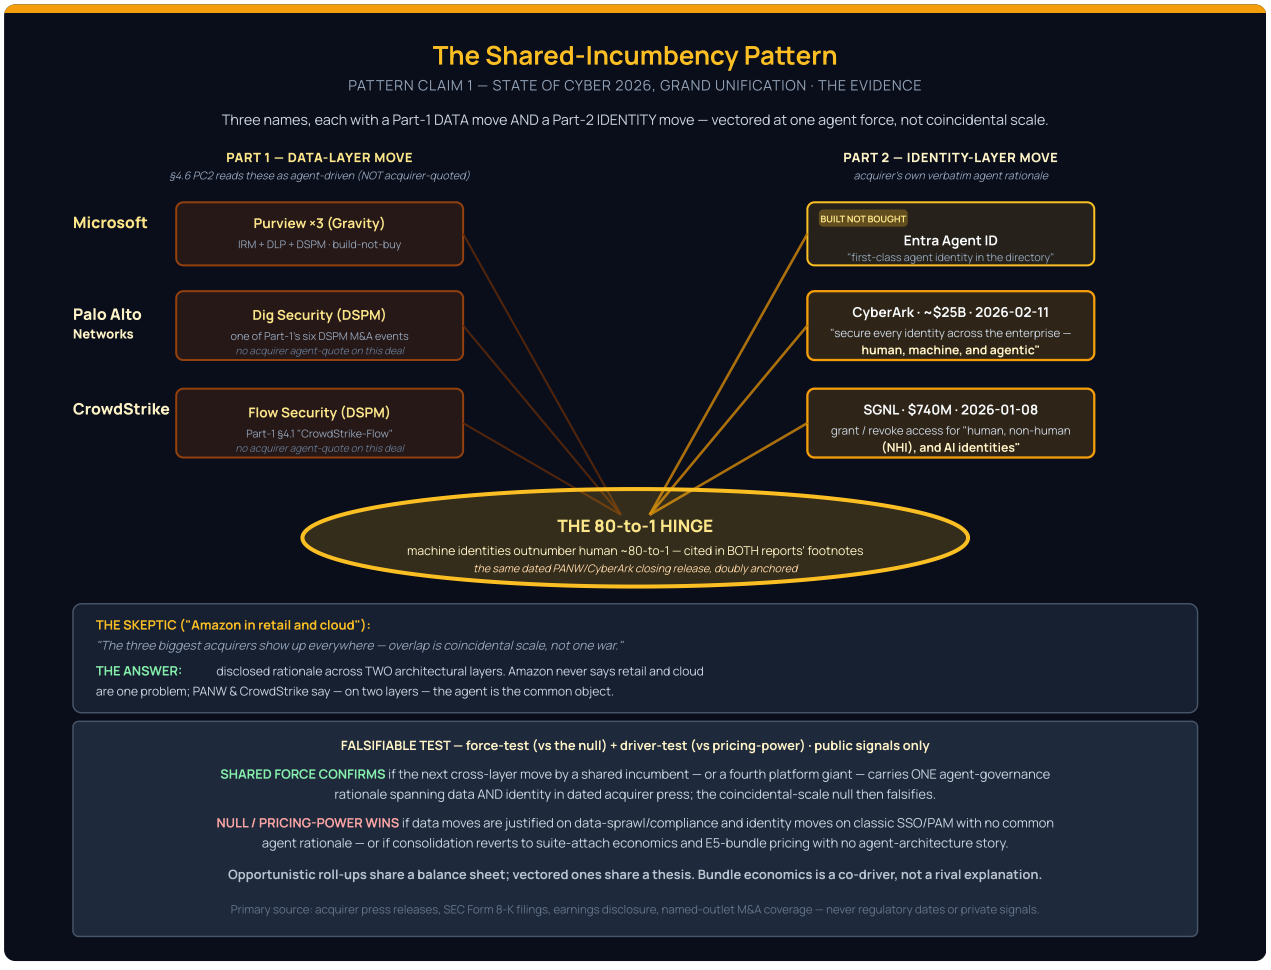

Observation. - The two reports resolved on two genuinely *different* enforcement primitives, and the public record shows a third sitting between them. - **Part 1 resolved on enforcement pulled to the DATA SOURCE.** Its §4.6 Pattern Claim 2 states it verbatim: "when the actor is a machine identity at machine speed, the only enforcement point that scales is the data store itself." Veeam's DataAI thesis is the same: control must shift "to the data source, not at the agent, so known and unknown agents cannot access sensitive data if that data is governed at the source" (see the Part-1 second pattern claim).[13] - **Part 2 resolved on consolidation through AGENT IDENTITY** — who the actor is, what it may authorize, the alert it generates. That is the vector the platform incumbents named in their own dated press (see the Part-2 identity convergence).[12] - And **Front 7 (AI Security) owns the AI-RUNTIME frontier between them** — securing the agent where it *executes*, distinct from where its data sits and from who it is. The runtime primitive is real and dated: Protect AI folded into PANW's Prisma AIRS line "to deploy autonomous AI agents safely," and Lakera into Check Point (reported ~$300M, closed 2025-10-22) as the foundation of its AI-security center of excellence (see the Part-2 convergence and Front 7).[14] - Three distinct primitives — **data** (Part 1), **runtime** (F7), **identity** (Part 2) — each owned by its own front, none reducible to the others.

My read. - I read these as three governance primitives of ONE force — the autonomous agent and the 80-to-1 inversion. The structural consequence is the spine of the whole report: an autonomous agent is *simultaneously* a thing-that-touches-data (Part 1's question), a thing-that-executes-in-a-runtime (F7's question), and a thing-that-must-be-identified-and-authorized (Part 2's question). - So only a platform spanning all three primitives can govern it end to end. That is precisely why the shared incumbents of Pattern Claim 1 are the structural favorites: they are the vendors assembling all three. - But the discipline that keeps this claim honest is to draw the three primitives *cleanly* and never collapse them into one undifferentiated "agent problem" — the over-claim the capstone exists to avoid. Data-source enforcement and identity governance are adjacent architectural layers, not one layer; the runtime frontier is a third. - And the leg-asymmetry must be stated, not buried. Agent-force is Part 2's *chosen* thesis, argued directly. But in Part 1 it is a *cross-front pattern claim* (§4.6 PC2) elevated to capstone level alongside Part 2's chosen vector — Part 1's own chosen thesis was DSPM as the data layer, a data-economics claim, not an agent-force claim. - I am elevating a Part-1 pattern claim, not rewriting Part 1 as "about agent identity." Stated that way — one force, three cleanly-drawn primitives, one set of incumbents spanning all of them — the architecture holds.

Conditional prediction. *Falsifiable on public buyer and vendor signals only. - If vendors and buyers increasingly demand governance spanning all three primitives — data AND runtime AND identity — from one platform, and the shared incumbents win the enterprise by being the only ones able to supply all three, the three-layer convergence holds and the spine is confirmed. - If instead the three primitives stay in separate best-of-breed lanes — a data-security leader that never crosses into runtime or identity, a runtime/AI-security specialist that scales independently, an identity pure-play that holds its own — with no single platform spanning all three winning the sophisticated enterprise, then the layers remain distinct markets and the "one force, three primitives" architecture is bounded to the mid-market where the bundle already spans them cheaply. - Primary source: vendor product-page heroes and analyst category coverage naming cross-primitive (data + runtime + identity) platform coverage; acquirer disclosed rationale across*

> *all three layers; earnings-disclosed enterprise logo mixes and cross-sell/attach metrics at the standalone leaders — public signals an outsider can watch, not regulatory timing or private buyer documents.*
>
> **Sources.** [12] [13] [14] (Part-1 §4.6 PC2; Part-2 §9.2 + §9.3 + §9.6 PC1 + Front 7 §7.6)

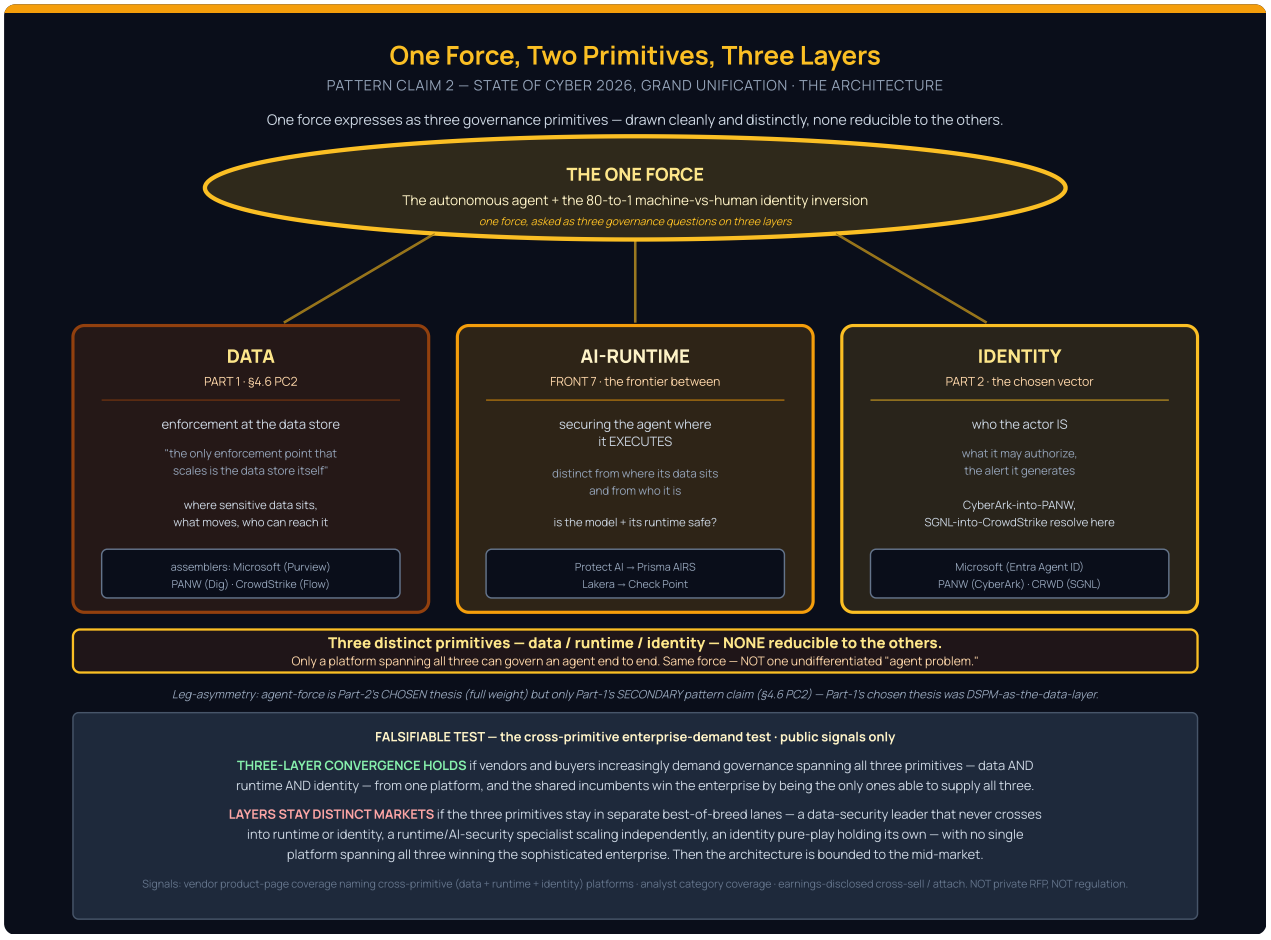

The two cross-report claims bracket the capstone from its two decisive sides. Pattern Claim 1 is the *evidence* — the same incumbents winning both wars, on a vectored agent rationale disclosed across two architectural layers, which is what the coincidental-scale skeptic cannot explain. Pattern Claim 2 is the *architecture* — one force expressing as three governance primitives (data / runtime / identity), only jointly governable by a platform spanning all of them.

Together they make the report one report. Not because the data war and the platform war are the same problem — they resolve on different primitives, and Part 1 chose the data primitive while Part 2 chose the identity one — but because the same force drives both, and the same handful of incumbents are assembling across all three layers to govern it. The win is bounded to the mid-market and greenfield in both reports; the sophisticated enterprise tier is genuinely contested in both. That is the single thesis of the whole report, stated at exactly the altitude the evidence will carry and no higher.

10.5 The Campaign Ahead — What Would Falsify the Capstone

A capstone earns its place only if it can be watched losing. The Grand Unification rests on a single inference — that the cross-report incumbent overlap is vectored at one force rather than coincidental scale or bundle convenience — and that inference is falsifiable on public signals alone, no privileged access required. Three discriminator signals decide it, and each separates the chosen thesis from a specific competitor.

The first is the **force-test against the null.** The watchable signal: does the next cross-layer move by a shared incumbent — or, more decisively, by a *fourth* platform giant beyond Microsoft, Palo Alto, and CrowdStrike — carry one agent-governance rationale spanning data *and* identity? If a fourth giant makes both a data-layer acquisition (DSPM or DLP) and an identity-layer acquisition and justifies both on the agent and non-human-identity problem in dated press, the shared force is confirmed and the coincidental-scale null falsifies. If instead the incumbents' data moves are justified on data-sprawl and compliance while their identity moves cite classic single-sign-on and privileged access with no common agent rationale, the arcs read as independently driven and the null wins. The signal class is the dated public record both convergence chapters already monitor: acquirer press releases, SEC Form 8-K filings, and named-outlet M&A coverage.

The second is the **driver-test against pricing-power.** The watchable signal: do the incumbents and buyers cite agent-governance *architecture*, or *price and bundle*? If cross-report consolidation continues to be justified on the human-machine-agentic governance problem — the PANW/CyberArk and CrowdStrike/SGNL pattern repeating — the architectural through-line holds and the bundle stays a co-driver. If the rationale increasingly reverts to suite-attach economics, procurement convenience, and E5-bundle pricing with no agent-architecture story, the bundle is the better explanation and the unification resolves at the commercial layer rather than the architectural one. The signal class is the acquirers' disclosed rationale and the buyers' stated reasons in dated, named-outlet coverage.

The third is the **cross-report overlap-persistence test** — the capstone's own falsifier, the one neither Part-1 nor Part-2 convergence could state alone because it is inherently cross-report. The watchable signal has two parts. Does Microsoft remain, on this report's own Gravity grading (Part-1 §4.3), the only vendor it places at Gravity across *both* all-three-Part-1-data-fronts and the Part-2-platform-fronts? And do Palo Alto and CrowdStrike keep extending across both layers, or does the overlap *narrow*? Persistent, deepening cross-layer overlap on one rationale hardens the shared-incumbency spine. The weakening signal is a clean separation: a Part-1 data-security leader such as Cyera scaling independently and never crossing into the Part-2 platform, or a Part-2 platform leader never touching the data layer. Data-security leaders and platform leaders separating cleanly confirms the null — two unrelated arcs that happened to share a few big names. The signal class is the next-refresh Gravity-placement audit across both reports, plus the same dated M&A record the first two tests watch.

These three signals are what make the Grand Unification a capstone and not a bridging epilogue. The claim is bounded — one force, two primitives, mid-market decided and the enterprise tier contested in both wars — and it is watchable. The report is one report because the force is one force, not because the two wars are one problem; and the next several quarters of public signal will either harden that spine or separate the two arcs back into two reports that happen to share three names.

References & Footnotes

1. The 80-to-1 machine-vs-human identity ratio is cited in the Palo Alto Networks closing release on its ~$25B CyberArk acquisition, 2026-02-11, and is footnoted as central in BOTH reports: Part-1 §4.2 Force 1 and Part-2 §9.6 Pattern Claim 1. Palo Alto Networks, "Palo Alto Networks Completes Acquisition of CyberArk to Secure the AI Era," 2026-02-11, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2026/palo-alto-networks-completes-acquisition-of-cyberark-to-secure-the-ai-era See Part-1 §4.2 Force 1 and Part-2 §9.6 PC1. ↩

2. Part-1 §4.6 Pattern Claim 2 ("Agentic AI Pulls Enforcement Back to the Data Source"), verbatim: "when the actor is a machine identity at machine speed, the only enforcement point that scales is the data store itself." This is the Part-1 bridge — a secondary cross-front pattern claim, NOT Part 1's chosen thesis (which is DSPM-Absorption-Substrate, a data-economics claim). See Part-1 §4.6 Pattern Claim 2. ↩

3. The AI-runtime primitive (Front 7): Palo Alto Networks folded Protect AI (acquisition intent 2025-04-28, completed 2025-07-22; reported >$500M, officially undisclosed) into its Prisma AIRS runtime line; Check Point acquired Lakera (~$300M, announced 2025-09-16, closed 2025-10-22) as the foundation of its Global Center of Excellence for AI Security. Carried from Part-2 §9.6 Pattern Claim 1 / Front 7 §7.6. See Part-2 §9.6 PC1. ↩

4. Microsoft spans both reports by building, not buying: Purview at Gravity across all three Part-1 data fronts (IRM + DLP + DSPM as M365 E5 / Purview Suite modules, with Agent 365 AI observability), per Part-1 §4.3; Defender/Sentinel in the SOC, Entra Agent ID (GA April 2026, agent identities a first-class directory construct inside Agent 365), and Defender-for-Cloud in AI on the Part-2 side (see Part-2 §9.3). Microsoft Learn, "Microsoft Purview Data Security Posture Management (DSPM)" and "Microsoft Entra Agent ID," accessed 2026-06-17. https://learn.microsoft.com/en-us/purview/data-security-posture-management-learn-about ; https://learn.microsoft.com/en-us/entra/identity/ ↩

5. Palo Alto Networks spans both reports: data layer — Dig Security (DSPM), named in Part-1 §4.1 as one of six DSPM-relevant M&A events ("PANW-Dig"); identity layer — CyberArk (~$25B, closed 2026-02-11), verbatim rationale "secure every identity across the enterprise — human, machine, and agentic," citing the 80-to-1 ratio in the same closing release; AI-runtime — Protect AI into Prisma AIRS. Palo Alto Networks, "Palo Alto Networks Completes Acquisition of CyberArk to Secure the AI Era," 2026-02-11, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2026/palo-alto-networks-completes-acquisition-of-cyberark-to-secure-the-ai-era ; GovCon Wire, "Palo Alto Networks Closes $25B Acquisition of Identity Security Company CyberArk," accessed 2026-06-17. https://www.govconwire.com/articles/palo-alto-networks-cyberark-25b-acquisition See Part-1 §4.1 and §4.2 Force 1, and Part-2 §9.3 and §9.6 PC1. ↩↩

6. CrowdStrike spans both reports: data layer — Flow Security (DSPM), named in Part-1 §4.1 as "CrowdStrike-Flow"; identity layer — SGNL (~$740M, announced 2026-01-08), verbatim rationale to grant and revoke access for "human, non-human (NHI), and AI identities" in real time, with "Continuous Identity for AI Agents" shipped at Identiverse 2026-06-15. CrowdStrike, "CrowdStrike to Acquire SGNL to Transform Identity Security for the AI Era," 2026-01-08; CNBC, "CrowdStrike buys identity security startup SGNL for $740 million," 2026-01-08, accessed 2026-06-17. https://www.crowdstrike.com/en-us/press-releases/crowdstrike-to-acquire-sgnl-to-transform-identity-security-for-ai-era/ ; https://www.cnbc.com/2026/01/08/crowdstrike-ai-cybersecurity-sgnl-acquisition.html See Part-1 §4.1 and Part-2 §9.6 PC1. ↩↩

7. Part-2 enterprise-tier best-of-breed persistence: CrowdStrike, Zscaler, and Netskope scale as premium-multiple public platforms; SailPoint re-IPO'd into a $1.16B-ARR (reported $1.163B, +26%) governance business; Saviynt took a $700M KKR-led round at ~$3B valuation (December 2025). Carried from Part-2 §9.5 / Front 8 §8.6. See Part-2 §9.5. ↩

8. Part-1 enterprise-tier best-of-breed persistence and the data-war M&A field: Cyera reached $9B post-money (Series F, 2026-01-08) and stayed standalone — not absorbed — alongside Cyberhaven and BigID. The Part-1 data layer consolidates from four starting positions; several of its largest absorbers are NOT Part-2 platform

incumbents — Veeam-Securiti ($1.725B, closed Dec 2025), Google-Wiz ($32B, early 2026), IBM-Polar, Rubrik-Laminar, Proofpoint-Normalyze. Carried from Part-1 §4.1, §4.3. Fortune / BusinessWire / TechCrunch / Calcalist on the Cyera round, accessed 2026-06-17. https://techcrunch.com/2026/01/08/data-security-startup-cyera-hits-9b-valuation-six-months-after-being-valued-at-6b/ See Part-1 §4.1 and §4.3. ↩↩

9. Microsoft's cross-report span. Part 1: Microsoft Purview holds Gravity-tier placement across all three Part-1 data fronts (IRM + DLP + DSPM as modules of the M365 E5 / Purview Suite SKU), per Part-1 §4.1, §4.2 Force 2, and §4.3 (the only single vendor at Gravity in all three Part-1 fronts). Microsoft Learn, "Learn about Microsoft Purview Data Security Posture Management (DSPM)," accessed 2026-05-14. https://learn.microsoft.com/en-us/purview/data-security-posture-management-learn-about . Part 2: Defender/Sentinel anchor the SOC, Entra Agent ID the Identity front, Defender-for-Cloud the AI front, per Part-2 §9.3. See the Part-1 convergence (§4.1/§4.2/§4.3) and the Part-2 convergence (§9.3). ↩↩

10. Palo Alto Networks' cross-report span. Part 1 (data layer): PANW acquired Dig Security (DSPM), named as one of the six DSPM-relevant M&A events in Part-1 §4.1 ("PANW-Dig"). Part 2 (identity + runtime): PANW completed its ~$25B acquisition of CyberArk on 2026-02-11, stating the addition lets it "secure every identity across the enterprise — human, machine, and agentic"; and folded Protect AI into its Prisma AIRS runtime line. Palo Alto Networks, "Palo Alto Networks Completes Acquisition of CyberArk to Secure the AI Era," 2026-02-11, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2026/palo-alto-networks-completes-acquisition-of-cyberark-to-secure-the-ai-era ; Palo Alto Networks, "Palo Alto Networks Announces Intent to Acquire Protect AI," 2025-04-28, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-announces-intent-to-acquire-protect-ai-a-game-changing-security-for-ai-company . See the Part-1 convergence (§4.1) and the Part-2 convergence (§9.3 + §9.6 PC1). ↩↩

11. CrowdStrike's cross-report span. Part 1 (data layer): CrowdStrike acquired Flow Security (DSPM), named in Part-1 §4.1 ("CrowdStrike-Flow"). Part 2 (identity): CrowdStrike announced its ~$740M acquisition of SGNL on 2026-01-08, to grant and revoke access for "human, non-human (NHI), and AI identities" in real time, and shipped "Continuous Identity for AI Agents" at Identiverse on 2026-06-15. CrowdStrike, "CrowdStrike to Acquire SGNL to Transform Identity Security for the AI Era," 2026-01-08, accessed 2026-06-17. https://www.crowdstrike.com/en-us/press-releases/crowdstrike-to-acquire-sgnl-to-transform-identity-security-for-ai-era/ ; CNBC, "CrowdStrike buys identity security startup SGNL for $740 million," 2026-01-08, accessed 2026-06-17. https://www.cnbc.com/2026/01/08/crowdstrike-ai-cybersecurity-sgnl-acquisition.html . See the Part-1 convergence (§4.1) and the Part-2 convergence (§9.6 PC1). ↩↩

12. The 80-to-1 hinge, central in BOTH reports. The PANW/CyberArk closing release cites an 80-to-1 machine-vs-human identity ratio. In Part 1 this same release anchors §4.2 Force 1: "Machine identities now outnumber human identities 80-to-1, per the Palo Alto Networks closing release on its $25B CyberArk acquisition." In Part 2 it anchors §9.6 Pattern Claim 1 as the figure cited in the same dated breath as the human-machine-agentic rationale. Palo Alto Networks, "Palo Alto Networks Completes Acquisition of CyberArk to Secure the AI Era," 2026-02-11, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2026/palo-alto-networks-completes-acquisition-of-cyberark-to-secure-the-ai-era . See Part-1 §4.2 Force 1 and Part-2 §9.6 PC1. ↩↩↩↩

13. Part-1 DATA-SOURCE primitive — the decisive Part-1 bridge. Part-1's §4.6 Pattern Claim 2 ("Agentic AI Pulls Enforcement Back to the Data Source") states verbatim: "when the actor is a machine identity at machine speed, the only enforcement point that scales is the data store itself." Veeam's DataAI Command Platform thesis (VeeamON 2026, 2026-05-13): enforcement must shift "to the data source, not at the agent, so known and unknown agents cannot access sensitive data if that data is governed at the source." This is a Part-1 cross-front *pattern claim*, NOT Part-1's chosen thesis (DSPM-Absorption-Substrate, a data-economics claim) — the capstone elevates the pattern claim. Veeam Press Release via Morningstar, 2026-05-13. https://www.morningstar.com/news/business-wire/20260512305482/veeam-launches-dataai-command-platform-the-industrys-first-unified-data-and-ai-trust-infrastructure-for-the-agentic-era . See the Part-1 convergence (§4.6 PC2). ↩↩

14. F7 AI-RUNTIME primitive — the third, distinct governance layer. Securing the agent where it executes, separate from data-source (Part 1) and identity (Part 2). Palo Alto Networks announced intent to acquire Protect AI on 2025-04-28 (completed 2025-07-22), folding it into the Prisma AIRS runtime line whose tagline is to "deploy autonomous AI agents safely." Check Point acquired Lakera (reported ~$300M, announced 2025-09-16, closed 2025-10-22) as the foundation of its Global Center of Excellence for AI Security. Part-2 §9.3 frames F7 absorption as "capability acquisition" — distinct from F8's category-leader absorption — and the two must not be blurred. Palo Alto Networks, "Palo Alto Networks Announces Intent to Acquire Protect AI," 2025-04-28, accessed 2026-06-17. https://www.paloaltonetworks.com/company/press/2025/palo-alto-networks-announces-intent-to-acquire-protect-ai-a-game-changing-security-for-ai-company ; Check Point / GlobeNewswire, "Check Point Acquires Lakera to Deliver End-to-End AI Security for Enterprises," 2025-09-16, accessed 2026-06-17. https://www.globenewswire.com/news-release/2025/09/16/3150869/0/en/Check-Point-Acquires-Lakera-to-Deliver-End-to-End-AI-Security-for-Enterprises.html . See the Part-2 convergence (§9.3 + §9.6 PC1) and Part-2 Front 7 (§7.6). ↩↩

Disclosures

DISCLOSURE.

This is ProductBeacon's independent, impartial market research, based solely on publicly verifiable signals, published on the open web. No vendor sponsors, no paywalled data, no analyst-firm reuse. It reflects ProductBeacon's independent view and is not investment advice. Full independence and sourcing methodology: Methodology →.

AFFILIATION.

By Yohay Etsion, Head of Product (Fractional), AXIA. AXIA is an AI-powered data-security (Insider Risk Management / DLP) company, which competes in the DLP segment covered in this report. This analysis is based only on publicly verifiable material.

NOT ADVICE.

This report does not constitute investment, legal, tax, or accounting advice. No claim should be relied upon as the basis for any investment decision. The author holds no trading position in any named public security and is not compensated by any named vendor. Readers who use this report in investment contexts bear sole responsibility for their decisions.