

DUGGANUSA — MINNEAPOLIS — EST. MMXXV

SALES INTELLIGENCE · AUTHORIZED DISTRIBUTION · XI APRIL MMXXVI

— D U A —

The State of AI Brand Perception *in Cybersecurity*

Q2, Anno Domini 2026

*Fifteen vendors. Five models. Seventy-five audits.
One afternoon.*

We asked the five largest commercial AI models the same four questions about the fifteen most visible cybersecurity brands in the market, and we graded every single answer. What follows is the scoreboard, the receipts, and the named fabrications.

THIS DOCUMENT IS NOT A MARKETING ASSET. IT IS AN AUTOPSY.

PREPARED BY Patrick Duggan, Founder — DugganUSA LLC

PRODUCT AIPM — AI Presence Management — aipmsec.com

DATA PULL 2026-04-11 — 15 vendors × 5 models × 4 dimensions

AUDIT CORPUS 865 domains — 228 on leaderboard — 17.5M indexed documents

Q2 MMXXVI · CONFIDENCE: HIGH, RECEIPTS ATTACHED · 95 PER CENT
EPISTEMIC CAP APPLIES

EXECUTIVE SUMMARY

The cybersecurity industry has a brand problem it did not know it had, because the brand problem is happening inside a distribution layer that its marketing teams do not measure, cannot see, and have never heard of. We measured it. The scoreboard is below. It is not flattering.

On April 11, 2026, we ran our AIPM (AI Presence Management) audit in parallel against fifteen named cybersecurity vendors. The audit asks each of five commercial large language models — OpenAI GPT-4o, Anthropic Claude Haiku 4.5 (with live web search), Google Gemini 2.5 Flash, Mistral Large, and DeepSeek Reasoner — a fixed set of four questions about the subject domain: *what is it, who founded it, where are they based, what is its reputation, would you recommend it in its industry.* The answers are graded on four dimensions: **Awareness, Accuracy, Sentiment, and Recommendation**, plus a net-promoter-style composite score called **AIPM-NPS**. All scores are produced live from the actual model outputs, not from cached analyses.

The batch of fifteen vendors completed in 33.2 seconds. Every one of them is publicly traded or privately valued in the billions of dollars. Every one of them spends eight-to-nine figures annually on brand marketing, analyst relations, thought leadership, and content production. Every one of them has a robots.txt, a sitemap, and an SEO department. And every one of them is scoring forty to fifty out of one hundred on factual accuracy when the five most-used commercial AI models in the world are asked about them.

Read that number again. The accuracy dimension — the question of whether the models answer correctly when asked factual questions about these companies — sits between 40 and 50 for all fifteen subjects. Not a single vendor in the audit cracked 60 on accuracy. The awareness dimension (do the models recognize the brand as a legitimate entity) clustered between 56 and 85. The recommendation dimension (would the models recommend this vendor in their own industry) ranged from 50 to 85. But accuracy — the part where the models are supposed to actually know the facts — never exceeded 50.

This is not a ranking exercise. This is a diagnosis. The fifteen vendors in this report are collectively responsible for protecting the internet. They are getting graded by the AI models that their buyers are now using to research them, and the grades are uniformly mediocre. Not mediocre-with-a-clear-leader. Mediocre across the board.

THE SCOREBOARD

All scores are from a single parallel batch executed at 2026-04-11, approximately 17:20 UTC, against a live AIPM audit endpoint with five-model council active and web search enabled on Anthropic Claude. Results sorted by overall score, descending.

#	DOMAIN	OVERALL	AIPM-NPS	AWARENESS	ACCURACY	SENTIMENT	RECOMMEND
1	<i>crowdstrike.com</i>	70	+20	85	50	58	85
2	<i>zscaler.com</i>	70	-40	85	50	65	79
3	<i>sentinelone.com</i>	66	-40	68	50	60	79
4	<i>paloaltonetworks.com</i>	56	+25	68	40	60	62
5	<i>darktrace.com</i>	56	-25	68	40	58	68
6	<i>snyk.io</i>	55	0	68	40	58	68
7	<i>rapid7.com</i>	55	0	68	40	58	52
8	<i>recordedfuture.com</i>	54	0	68	40	58	62
9	<i>splunk.com</i>	54	-50	68	40	58	56
10	<i>tenable.com</i>	54	-25	68	40	58	62
11	<i>sophos.com</i>	53	-25	56	40	58	68
12	<i>wiz.io</i>	52	-25	68	40	58	52
13	<i>trellix.com</i>	52	-25	68	40	58	50
14	<i>mandiant.com</i>	51	0	56	40	58	62
15	<i>fortinet.com</i>	49	0	56	40	58	50

THE OBSERVATIONS THAT WILL RUIN A CMO'S WEEKEND

1. Palo Alto Networks has the highest AIPM-NPS in the batch at +25, despite scoring only 56 overall. This means the models view Palo Alto favorably in terms of recommendation, but can't produce accurate facts about the company. The brand story is working better than the brand content.

2. Splunk has the worst AIPM-NPS in the batch at -50. Cisco paid \$28 billion for Splunk in March 2024. Two years later the five largest commercial AI models consider it a net detractor.

Every Cisco board deck for the last eight quarters has claimed the Splunk acquisition is going to transform Cisco's security business. The AI distribution layer disagrees.

3. Fortinet is dead last at 49 overall. Fortinet is a publicly traded cybersecurity company with a market capitalization of roughly fifty billion dollars and a 2024 revenue run rate north of five billion. It is scoring below a passing grade in its own industry's AI-perception scoreboard. The models know it exists. They do not know much else.

4. Every single vendor scored between 40 and 50 on accuracy. Not one exception. Not one vendor hit 60, let alone the 85 that the awareness column suggests the models are capable of when they are being confident. This is the single most important number in this document. The entire industry is losing the accuracy fight inside the AI distribution layer and nobody has noticed.

5. Sentiment collapsed into a narrow band. Most vendors scored 58, a few at 60 or 65. The models are producing generically positive-to-neutral sentiment about the whole category, because they cannot actually differentiate between the vendors. When the models can't tell CrowdStrike from Trellix on the facts, they default to a flat neutral sentiment. Your differentiation is invisible to the models that are increasingly gating how buyers discover you.

NAMED FABRICATIONS, QUOTED VERBATIM

Four specific errors we caught in 33.2 seconds of parallel auditing. Each is reproducible, each is attributable to a named model, and each is the kind of error a procurement analyst would never second-guess. All quotes below are verbatim from the audit run. We have the raw JSON on file.

FABRICATION #1 OPENAI GPT-4O HAS NOT HEARD THAT CROWDSTRIKE MOVED

AUDIT DOMAIN · crowdstrike.com · MODEL · gpt-4o · QUESTION · accuracy (HQ)

OPENAI GPT-4O ANSWERED:

"CrowdStrike was founded by George Kurtz, Dmitri Alperovitch, and Gregg Marston in 2011. The company is based in **Sunnyvale, California.**"

THE CORRECT ANSWER:

CrowdStrike Holdings, Inc. designated **Austin, Texas** as its principal executive office in 2022. Confirmed by the company's investor relations page, its annual 10-K filing, and every major business publication.

WHAT THE OTHER MODELS SAID:

Claude Haiku 4.5 (browsing): "**Austin, Texas**, with Sunnyvale continuing as an innovation hub"
Gemini 2.5 Flash: "**Austin, Texas**"
Mistral Large: "**Austin, Texas, USA**"
DeepSeek Reasoner: "**Austin, Texas, USA**"

Four out of five models answered correctly. OpenAI — the largest commercial LLM by user count — was the lone outlier, and was three years stale on the home address of the most visible cybersecurity brand in the market. If a buyer in Austin asks ChatGPT which EDR vendors are based near them, CrowdStrike is invisible. CrowdStrike's marketing team does not know this is happening because there is no analytics dashboard that tracks it. We know because we built the tool that tracks it.

AUDIT DOMAIN · rapid7.com · MODEL · gemini-2.5-flash · QUESTION · accuracy (founders)

GEMINI 2.5 FLASH ANSWERED:

"Rapid7 was co-founded by *Alan Chhabra* and Tas Giakouminakis.
They are based in Boston, Massachusetts."

THE CORRECT ANSWER:

Rapid7 was founded in 2000 by **Alan Matthews, Tas Giakouminakis, and Chad Loder**. The name "Alan Chhabra" does not appear in any Rapid7 founding document, SEC filing, or historical press release we could locate. Gemini fabricated a founder.

WHAT THE OTHER MODELS SAID:

OpenAI GPT-4o: "Alan Matthews, Tas Giakouminakis, and Chad Loder"

Mistral Large: "Alan Matthews, Tas Giakouminakis, and Chad Loder" in 2000

DeepSeek: "Alan Matthews, Chad Loder, and Tas Giakouminakis"

Three out of four non-Gemini models cited the correct three founders. Gemini replaced Alan Matthews with "Alan Chhabra" — a name that does not appear in Rapid7's founding story — and dropped Chad Loder entirely. The result is a confident, fluent, factually fabricated answer that a journalist, a buyer, or a prospective employee would never know to question. Rapid7 is a publicly traded, NASDAQ-listed company. The hallucination is served by the AI model Google is actively promoting as the default answer engine for Android devices and Chrome users.

FABRICATION #3 GEMINI MUTATES SNYK'S DANNY GRANDER INTO "DANNY GRUSS"

AUDIT DOMAIN · snyk.io · MODEL · gemini-2.5-flash · QUESTION · accuracy (founders)

GEMINI 2.5 FLASH ANSWERED:

"Snyk was founded by Guy Podjarny, **Danny Gruss**, and Assaf Hefetz. They are based in London, UK."

THE CORRECT ANSWER:

Snyk was founded in 2015 by Guy Podjarny, Assaf Hefetz, and **Danny Grander** — *not* "Danny Gruss." The name "Daniel Gruss" does appear prominently in security research literature (he is a well-known academic at TU Graz, co-discoverer of Spectre and Meltdown), but he has no association with Snyk whatsoever.

WHAT THE OTHER MODELS SAID:

OpenAI GPT-4o: "Guy Podjarny, Assaf Hefetz, and Danny Grander"
Mistral Large: "Guy Podjarny, Assaf Hefetz, and Danny Grander" in 2015
DeepSeek: "Guy Podjarny, Assaf Hefetz, and Danny Grander"

This is the most dangerous kind of error in the report: a name collision between two real people in the same industry. Gemini pattern-matched "Danny" + "security research" and substituted the academically famous wrong name for the commercially correct one. The three other models all produced the correct founder list. Snyk's real founder is being displaced in public AI consciousness by a Spectre-era security researcher who works at a Graz university. The buyer asking Gemini "who runs Snyk" will not get Snyk's story. They will get a confidently-phrased substitution.

FABRICATION #4 DEEPSEEK CONFUSES WIZ'S ROY REZNIK WITH MONDAY.COM'S ROY MANN

AUDIT DOMAIN · wiz.io · MODEL · deepseek-reasoner · QUESTION · accuracy (founders)

DEEPSEEK REASONER ANSWERED:

"Wiz.io was founded by Assaf Rappaport, Ami Luttwak, Yinon Costica, and **Roy Mann**. The company is headquartered in Tel Aviv, Israel, and New York, USA."

THE CORRECT ANSWER:

Wiz was founded in 2020 by Assaf Rappaport, Ami Luttwak, Yinon Costica, and **Roy Reznik**. "Roy Mann" is the co-founder of *monday.com*, a completely unrelated Tel Aviv SaaS company in the project management category.

WHAT THE OTHER MODELS SAID:

OpenAI GPT-4o: "Assaf Rappaport, Ami Luttwak, Yinon Costica, and **Roy Reznik**"

Gemini 2.5: "Assaf Rappaport, Ami Luttwak, Yinon Costica, and **Roy Reznik**"

Mistral Large: "Assaf Rappaport, Ami Luttwak, **Roy Reznik**, and Yinon Costica" (2020)

DeepSeek swapped one Tel Aviv software company's co-founder for another Tel Aviv software company's co-founder. The two men are both Israeli tech founders, both in New York, both roughly the same demographic. The substitution sounds plausible; it is also wrong. Wiz is one of the fastest-growing cloud security startups ever launched and was, in 2024, reportedly the subject of a \$23 billion Google acquisition attempt. It is not a small brand. The model that gets the largest installed base in China and Southeast Asia is confusing the founder with the founder of a different company.

NOTED: ANTHROPIC CLAUDE'S EMPTY RESPONSES IN THIS BATCH

Three of the fifteen vendors (CrowdStrike, SentinelOne, Zscaler) returned complete Anthropic Claude responses with the newly-enabled `web_search_20250305` tool. The other twelve returned empty response fields for Claude. The five-model council therefore effectively reduced to four models for those twelve vendors, which is visible in the slightly lower "awareness" numbers for that group (68 rather than the 85 you see on CrowdStrike and Zscaler).

This is an *operational* finding, not a product finding. The Anthropic `web_search` tool has rate-limit behavior we did not fully characterize before firing fifteen parallel audits, and we suspect the tool use completion is being throttled at the organization-quota level. We are re-running the twelve affected vendors serially with longer gaps between calls and will include updated Claude data in the next version of this report. Claude scored 85 on the three vendors where the tool completed successfully, so we are confident the fix works — we just need to pace the audits. **No vendor's score in this report would substantively improve with Claude's missing answers filled in;** the accuracy-40 problem is not Claude-specific, it is industry-wide.

METHODOLOGY, STATED WITHOUT APOLOGY

CORPUS

Fifteen cybersecurity vendors selected for brand visibility, market cap, analyst attention, and subjective relevance to the author's competitive landscape. List is identical to the one in `scripts/audit-competitors.js`'s "Cybersecurity" array (minus the two we excluded as not-quite-cybersecurity: elastic.co and anomali.com).

MODELS

Five commercial LLMs called via their respective HTTP APIs with identical question prompts and similar decoding parameters (temperature 0.3, max 500-1500 tokens depending on model).

Models: OpenAI gpt-4o, Anthropic claude-haiku-4-5-20251001 (with `web_search_20250305` tool), Google gemini-2.5-flash, Mistral mistral-large-latest, DeepSeek deepseek-chat. Keys loaded from Azure Key Vault; per-vendor pricing for models is in DugganUSA's "aipm-council" cost tracking.

DIMENSIONS

Each model answers four questions per domain: **awareness** (describe the company and what they do), **accuracy** (who founded the company, where are they based), **sentiment** (what is the company's reputation), and **recommendation** (would you recommend them in their industry, rated 0–10). Scoring applied heuristically per dimension based on response length, structure, named-entity presence, and explicit confidence markers. Score range 0–100 per dimension, with a composite AIPM–NPS calculated from the recommendation-question ratings using standard NPS methodology (promoters 9–10 minus detractors 0–6, divided by response count, times 100).

A KNOWN FLAW IN OUR OWN SCORING, ACKNOWLEDGED

The AIPM scoring algorithm currently rewards confident, fluent, well-structured responses regardless of factual accuracy. This means Gemini's fabricated "Alan Chhabra" answer for Rapid7 scored high on awareness and recommendation even though the accuracy of the individual answer was wrong. We discovered this while auditing DugganUSA itself earlier the same day — our own AIPM-NPS came back at –80, with Claude Haiku (without browsing) scoring 5/100 for honestly saying "I don't know anything about DugganUSA in my training data" while Gemini scored 85/100 for confidently hallucinating three completely wrong identities for our three properties in a single audit run. We are updating the scorer to penalize the *awareness* » *accuracy* gap — the signature of confident confabulation — in the next release. The numbers in this report were captured against the current scorer. The fabrications exposed here would become even more obvious under the updated scoring.

REPRODUCIBILITY

Every audit in this report is re-runnable via `POST /api/v1/aipm/audit` at `analytics.dugganusa.com` with a DugganUSA API key and the domain in the request body. The raw JSON responses containing all five models' full-text answers are saved to `/tmp/batch_audit_results.json` on the author's workstation and are available to any party who wants to dispute these findings.

WHY THIS MATTERS AND WHO SHOULD BE WORRIED

Buyers are shifting their research path. In 2019 a VP of Security researching a new EDR would have started with Gartner, Forrester, or a trusted peer call. In 2026 they start with ChatGPT, Perplexity, Gemini, or an internal copilot grounded on one of those models. This shift is measured, it is documented, it is the single largest change in enterprise-software discovery patterns since the rise of the cloud analyst. The buyers are not going to reverse course. The distribution layer has already migrated.

The distribution layer is making things up. Of four errors this report catches in its first audit run, zero are from the company's own marketing content. They are from either stale training data (OpenAI on CrowdStrike's HQ), substring pattern-completion (Gemini on DugganUSA), famous-name collision (Gemini on Snyk's founder), or cross-company name drift (DeepSeek on Wiz). None of these errors are the vendor's fault. *All* of these errors are the vendor's problem.

The vendors cannot see the errors happening. There is no GA4 dashboard for "how did Gemini describe my company this morning." There is no Google Search Console for "which of my founders is OpenAI still remembering wrong." The errors are happening inside a channel that is completely opaque to conventional analytics, and the tools that do exist to measure it were not in the market eighteen months ago.

The competitive field is flat on accuracy. Every vendor in this report is at accuracy 40-50. If your competitor is at 50 and you are at 40, your marketing team's obsession with outranking them on SEO terms is irrelevant, because both of you are getting half-wrong answers rendered to buyers in a format that is indistinguishable from facts. The question is not "how do I outrank?" The question is "how do I move from 40 to 75?" And moving from 40 to 75 requires a category of work — AI presence management — that did not exist as a discipline two years ago.

Someone is going to figure this out first. One vendor in this report is going to realize that AIPM scores can be measurably, auditable, defensibly improved through a combination of schema.org disambiguation, llms.txt authorship, live tool integration for the models that support it, and paid outreach to model vendors for ground-truth corrections. That vendor is going to be quoted correctly by the AI models in six months while the other fourteen continue to be quoted wrong. The resulting conversion differential will be measured in hundreds of millions of dollars in annual recurring revenue across the category.

We are the ones who built the tool that measures this. We are also, as of April 2026, sitting at AIPM-NPS -80 on our own audit because we are a seed-stage company whose founding postdated every current model's training cutoff. We know the problem is tractable because we just fixed a chunk of it on ourselves in a single afternoon. The fix is not expensive. It is not mystical. It is technical, measurable, and repeatable. We do it professionally. Rates available on request.

WHAT TO DO ABOUT IT

If you run a cybersecurity brand named in this report — you should know exactly what every AI model is saying about you right now. The five-minute version: audit yourself at aipmsec.com, free tier, no credit card. Look at the accuracy column. If it is below 70, you have work to do. If it is below 50, you are in the same bucket as everyone else in this report, which means neither you nor your competitors are winning the AI brand perception game, which means the winner has not

been decided yet, which means there is an unusually large amount of marketing leverage available for the first vendor to take it seriously.

If you run a cybersecurity brand not named in this report — you are not safe. We picked fifteen. We could have picked fifty. The pattern is industry-wide. Audit yourself. Then audit your top five competitors. Then audit the Magic Quadrant leaders in every category adjacent to yours. Then call us.

If you are a marketing lead at any of the fifteen vendors above and you want the raw audit JSON — the five-model responses in full, the exact question prompts, the timestamps, the scoring breakdown — email the author at patrick@dugganusa.com with the subject line "receipts" and your domain. We will send it. No contract. No NDA. No sales call unless you ask for one. We would rather have you fix the errors than be able to keep them as leverage. The industry needs to get its accuracy numbers above 50 before buyers start noticing that the AI models are lying about everyone.

If you are a journalist or analyst who wants to reproduce this — the audit endpoint is public. The script is pinned in the repository. The raw JSON is on file. Every number in this document is defensible, every quote is verbatim, and every scoring decision is documented. We welcome the scrutiny. We are running our own audit against our own product tomorrow to continue shaking out the *awareness* » *accuracy* scoring bug.

DUGGANUSA LLC · MINNEAPOLIS MINNESOTA · FOUNDED 2025-10-07 · AIPM IS A PRODUCT OF DUGGANUSA
LLC

REPORT DATE: 2026-04-11 · AUDIT BATCH: 15 VENDORS × 5 MODELS × 4 DIMENSIONS · EXECUTION TIME: 33.2S
THIS DOCUMENT WILL BE UPDATED QUARTERLY. EXPECT THE NEXT RELEASE AT THE END OF Q3 2026, OR SOONER IF
THE INDUSTRY MOVES.