

THE 808

VOL. 1 NO. 1 · VENDOR-AGNOSTIC ADVISORY ON CLOUD, AI & SECURITY

The Brief

September 2026

Twenty-five stories across enterprise software, SaaS, AI and security — each with what happened, who it hits, and a link to the source.

THIS MONTH'S FEATURE · OP-ED

Too Fast to Fail; Anthropical Storm?

A confidential S-1, a Nasdaq slot, a \$2 trillion ask — and a Pentagon blacklisting upheld in court three days before we went to press.



the808.org · contact@the808.org · Toronto, Ontario

Every claim links to its source. Nothing in this document is investment advice.

IN THIS ISSUE

September 2026

Feature — Too Fast to Fail; Anthropical Storm?	3
Around this story	8
Enterprise — 7 stories	9
SaaS — 6 stories	9
AI — 6 stories	10
Security — 6 stories	11

FROM THE EDITOR

This is the first issue. The Brief exists because the gap between what gets announced and what actually changes a buying decision is wider than it has ever been — and most of what crosses your desk is written by someone with a product to sell. We don't sell software and we take no referral fees from any vendor named in these pages.

From the next issue this edition is assembled by a crawler that checks every quoted line word-for-word against the source article before it publishes. This one was compiled by hand, to the same standard.

— Andrew J. Baynes, Toronto

FEATURE • OP-ED

Too Fast to Fail; Anthropical Storm?

Anthropic wants to go public while telling the industry to slow down.


By Andrew J. Baynes • 28 September 2026 • 7 min read

Whiplash. In one timeline over the next two months, Anthropic pulls off the biggest IPO in history. In another, they stay home. The punchline isn't the eye-watering valuation — it's that the company's own CEO has spent the past few weeks pitching the world's most compelling PowerPoint on why everyone needs to chill out and slow down.

Franken-corp. Anthropic was hatched in 2021 by OpenAI defectors and isn't wired like a normal Silicon Valley darling. It's a Delaware public benefit corporation wearing a Long-Term Benefit Trust as a hat. This is a bespoke governance arrangement ensuring that, in theory, the company's altruistic mission can outvote the money. Bookmark that, because it's about to become the spiciest line item in the prospectus.

Anthropic's reported revenue run rate

Annualised, in billions of dollars. Figures as reported; not audited.



Sources: company figures reported by Reuters, Bloomberg, WSJ, CNBC · Chart: The 808

From \$87 million to \$65 billion in thirty months. The shape is the pitch.

Moonmath. The numbers escalated from "cute" to "absurd" overnight. Reported annualized revenue run-rates strapped on a jetpack, rocketing from roughly \$1 billion in December 2024 to \$14 billion by February 2026, eventually cresting an astronomical \$65 billion by mid-2026. Following a \$65 billion Series H raise at a \$965 billion post-money valuation, Anthropic is now one of the most valuable private entities on earth.

Asterisk. In May, the Wall Street Journal reported Anthropic had told investors to expect a first-ever quarterly operating profit of \$559 million on Q2 revenue of \$10.9 billion. What actually landed in August was revenue north of \$11.5 billion and a first quarter of positive adjusted operating income — with the dollar figure undisclosed. Worth

holding those apart. The company also warned investors that compute spending in late 2026 and 2027 is likely to swing operating results back into the red. The profit was a window, not a trend.

Consumer user counts are reported anywhere from 30 million to 245 million depending on definition and source, so we've left them out.



S-1-derful. On June 1, the company quietly slid a draft S-1 over to the SEC. They hired Wilson Sonsini — the legal equivalent of the Beatles, who ran Google's 2004 IPO — to get them red-carpet ready. Now Nasdaq is apparently clearing the VIP section for an October or November listing, anchored by Wall Street's finest and a rumored fat check from Nvidia. Investors are drooling over a potential \$2 trillion valuation, which would effortlessly dethrone the \$1.77 trillion record SpaceX set in June.

Mirage? But here is the asterisk that could break the model. Anthropic reports its cloud partner revenue (AWS, Azure, Google Cloud) on a gross basis, scooping up the full billed amount. That is notoriously not how public software companies do the math. The bankers are currently underwriting against a mythical \$190–200 billion in 2028 revenue. Whether you buy that depends entirely on whether you think you're looking at actual cash or funny money.

"We must slow the pace at which we improve the capabilities of AI models."

Dario Amodei, 12 September 2026

Buzzkill. CEO Dario Amodei has a quirky habit of saying the quiet, valuation-tanking parts out loud. This radical candor is either the IPO's secret weapon or a massive liability, and the roadshow is going to be a fascinating social experiment to find out which.

Handbrake. On September 12, Amodei dropped an essay basically begging the AI industry to take a breather. Warning that rogue AI swarms could hijack the internet before we finish our lattes, he wrote, "We must slow the pace at which we improve the capabilities of AI models." He pitched external safety watchdogs, global treaties, and democratic standard-setting — a coordinated effort so massive it practically requires a permission slip from antitrust regulators.

Frenemies. The reaction was surreal. Sam Altman, traditionally not Amodei's biggest fanboy, posted his agreement about pacing the frontier. Elon Musk chimed in with a stark "Dario is right." Getting those three to agree on anything is like getting cats to march in a parade.

Two positions, one prospectus

WHAT THE COMPANY SAYS

"We must slow the pace at which we improve the capabilities of AI models."

— Dario Amodei, 12 Sep 2026

WHAT THE COMPANY IS DOING

Confidential S-1 filed 1 June.
Nasdaq selected. Reported target: up to \$100B raised at a valuation near \$2T.
Window: October–November.

AND, AS OF 25 SEPTEMBER

A federal appeals court upheld the Pentagon's supply-chain-risk designation — in a 2–1 decision, weeks before the IPO window.

Sources: Axios, Reuters, WSJ, CNBC, Defense One, Nextgov · Graphic: The 808

The two positions the S-1 will have to hold in the same document.

Tomatoes. Naturally, the peanut gallery had notes. VC loudmouth Chamath Palihapitiya accused Amodei of making "the case to stop open source and concentrate enormous technological and economic power with Anthropic." Amodei knows he sounds like a doomer to some — he's acknowledged the charges of "hype, 'doomerism,' or regulatory capture" — but that last one is a remarkably sticky accusation to shake off when you're asking investors for a trillion dollars.

Radioactive. Amodei's political dance card hasn't helped. A vocal critic of Donald Trump and a 2024 Kamala Harris backer, he watched Hegseth direct the Pentagon in February to designate Anthropic a supply-chain risk — a label typically reserved for firms suspected of working for adversarial governments. Trump has since said the administration stopped "AI 'people' from doing bad, or potentially bad, 'things,' like Dario (Anthropic!), who is now pretending to be a 'perfect little angel'." Playing politics has left a tangible scent of risk on the ticker.

Gladiators. Compared to OpenAI — which hovered around \$852 billion this past spring and insists it isn't going public in 2026 — Anthropic gets to be the first major frontier lab to brave the public markets. They get the pioneer premium, but they also get the arrows in their back when Wall Street finally decides what an AI lab is actually worth.

Bazaar. Down in the trenches, it's a knife fight. Google is baking AI straight into its enterprise suite. Microsoft glued Copilot to everything but the office microwave. DeepSeek is raising funds in Shanghai, and Alibaba just slashed audio API prices by 95 percent. Enterprise buyers are polyamorous now, running three models at once. Good luck explaining those profit margins to analysts who only care about cold, hard SaaS retention.



Friction. Here lies the ultimate rub, and the S-1 will have to address it with a straight face.

AWOL. Since January 2026, Anthropic has been in a staring contest with the Pentagon. Claude was reportedly used in the January operation to capture Venezuelan President Nicolás Maduro. The Pentagon then asked Anthropic — and other AI companies — to permit use of their models for "all lawful purposes." Anthropic was willing to adapt its policies, but drew a hard line on two things: no mass domestic surveillance, and absolutely no killer robots. In retaliation, Hegseth stamped them with a supply-chain-risk scarlet letter.

Benched. And on Friday, the courts backed the brass. In a 2–1 decision on September 25, the DC Circuit Court of Appeals rejected Anthropic's challenge. DOD employees and contractors remain prohibited from using Claude for department business. Anthropic says it "respectfully disagree[s]" and is weighing further review; Defense Undersecretary Emil Michael posted that "the hammer of justice has smashed @AnthropicAI arguments." The company claims the designation has already cost it billions — weeks before it wants to price an offering.

Fasting. Look at that ruling next to Amodei's September essay, and a theme emerges: Anthropic loves leaving money on the table. Telling the Pentagon "no thanks" is declining billions on principle. Begging the industry to decelerate is asking the government to put a leash on your own golden goose.

Mutiny. Pitching a company that actively refuses revenue to a room full of public shareholders is bold. That Long-Term Benefit Trust exists specifically to shield management from greedy activist investors, but a corporate structure designed to ignore the people signing the checks is practically begging for a proxy fight. Wall Street will price that headache in.

Baking? So, is it too early to ring the bell? Three solid arguments say yes, and none of them involve feelings.

Squinting. The revenue needs a reality check. Extrapolating a \$200 billion future based on gross-basis cloud accounting is a hallucination that makes a chatbot look lucid. One more year of audited, standardized numbers would instantly vaporize the short-sellers' favorite talking point.

Uncharted. The governance math is weird. No one has ever tried to sell a \$2 trillion, mission-driven, benefit-trust AI lab to the public markets. Pushing that unprecedented boulder up the hill in a slippery November window is a massive gamble.

Baggage. That noble safety stance is about to get a price tag. Right now, Amodei's apocalyptic warnings make him look like the adult in the room. Post-IPO, dropping an essay about rogue internet-destroying AI agents sounds a lot like the CEO shorting his own stock.

YOLO. The argument to launch now is purely pragmatic: the IPO window is wide open, SpaceX proved investors have bottomless pockets for mega-caps, and capital buys compute. If the AI war is won in the next 24 months, waiting around is financial suicide.

Double-speak. But the biggest red flag comes from inside the house. If Anthropic truly believes we need international treaties, antitrust waivers, and collective industry deceleration to survive the AI revolution, then sprinting toward a November IPO just because the market is hot feels a bit hypocritical. You can believe the tech needs to slow down, and you can believe your startup needs a trillion dollars right now. But the market is only going to listen to the money.

Check-please. For five years, Anthropic has claimed that playing it safe is a feature, not a bug. An IPO is the exact moment that philosophy gets audited on a quarterly earnings call. It's a reckoning worth having. It's just not one they need to have in November.

EDITORIAL NOTE

Written by Andrew J. Baynes. AI tools were used in research and drafting. Facts are drawn from reporting by Reuters, Bloomberg, the Wall Street Journal, CNBC, Axios, Defense One, Nextgov, Forbes and the Congressional Research Service. The judgements are the author's.

Nothing here is investment advice. The 808 holds no position in Anthropic and takes no fees from any company named.

**AROUND THIS STORY**

06

Six developments that frame the feature, each linked to its source.

Appeals court upholds Pentagon blacklisting

A 2–1 DC Circuit decision on 25 September rejected Anthropic's challenge to its supply-chain-risk designation. DOD employees and contractors remain barred from using Claude for department business.

Defense One · <https://www.defenseone.com/threats/2026/09/anthropic-lawsuit-supply-chain-risk/416252/>

Amodei to dine with Trump at the White House

The first one-on-one between the two, after Amodei missed a state dinner. Commerce Secretary Howard Lutnick, on the sidelines of an event: "We trust Anthropic."

CNBC · <https://www.cnbc.com/2026/09/27/dario-amodei-set-to-have-dinner-with-trump-after-missing-state-dinner.html>

Altman and Musk back the slowdown call

Altman posted that he agreed "we need to pace the frontier." Musk replied "Dario is right." Anthropic executives followed the essay with a call for legislative and regulatory responses.

Axios · <https://www.axios.com/2026/09/12/anthropic-ai-amodei-pacing>

Q2 revenue came in above \$11.5B

Preliminary Q2 results reported in August showed revenue over \$11.5 billion, up from \$787 million a year earlier, and a first quarter of positive adjusted operating income — amount undisclosed.

CNBC · <https://www.cnbc.com/2026/08/15/anthropic-revenue-jumps-to-over-11point5-billion-in-q2-report.html>

SpaceX set the bar in June

The largest IPO ever: listed 12 June 2026 at a \$1.77 trillion valuation, raising \$86 billion. Shares fell roughly 50% from their post-IPO peak by late July.

Wikipedia · https://en.wikipedia.org/wiki/Initial_public_offering_of_SpaceX

OpenAI is not racing to list

Last reported around \$852 billion in March. Altman has told Fortune a listing is likely not coming in 2026, leaving Anthropic first among the frontier labs to test public markets.

Forbes · <https://www.forbes.com/sites/jonmarkman/2026/09/15/dario-amodei-calls-for-ai-slowdown-as-anthropic-opens-up-to-auditors/>



ENTERPRISE

07

Salesforce unveils Alforce and a Headless Toolkit at Dreamforce

At Dreamforce 2026 Salesforce introduced Alforce, a new AI interface layer, alongside a Headless Toolkit, which lets customers use Salesforce software without its traditional user interface. Salesforce frames both as part of an "Enterprise AI Harness" combining data, business knowledge, workflows and control. The pitch is a composable architecture that agents can read and act on, rather than a chat window bolted onto the CRM.

The 808's take — The Headless Toolkit is the more consequential half. A vendor that lets agents bypass its own interface is betting its moat is the data model, not the screens — and that bet sets the terms for everyone evaluating CRM this year.

Salesforce Ben • <https://www.salesforceben.com/salesforce-launches-aiforce-at-dreamforce-26-ai-replaces-the-ui/>

Dataiku launches Agent Management to inventory agents across platforms

Dataiku announced Agent Management at its Succeed conference in New York on 24 September: a standalone product that inventories the AI agents an enterprise runs regardless of which platform built them. It spans Salesforce Agentforce, Microsoft Copilot Studio, AWS Bedrock, Google Vertex, Databricks and Snowflake, measures business and technical performance, and flags the agents carrying the most risk. General availability is set for October 2026, priced per instance annually with monitoring metered per agent.

The 808's take — Charging per agent while selling a tool that counts your agents is an odd incentive to hand a vendor. The category is right; the meter deserves scrutiny.

The Agile Brand Guide • <https://agilebrandguide.com/dataiku-launches-agent-management-to-track-and-manage-performance-of-ai-agents-built-and-running-on-leading-platforms/>

Microsoft redesigns Copilot around a Home screen with Office apps built in

Microsoft announced a redesigned Copilot on 25 September, built around a Home screen where Chat, Cowork and full versions of Word, Excel and PowerPoint sit together. It folds the assistant and the applications into a single surface rather than keeping Copilot as a panel beside them.

NeuralBuddies • <https://www.neuralbuddies.com/p/ai-news-recap-september-25-2026>

GitHub Copilot adds Anthropic's newest model to the picker

From 22 September, developers on Copilot Pro+, Max, Business and Enterprise plans can select Anthropic's newest model from the model menu. Model choice inside a coding tool is now a per-seat setting rather than a platform decision.

NeuralBuddies • <https://www.neuralbuddies.com/p/ai-news-recap-september-25-2026>

Gemini gains connectors to Airtable, monday.com, Adobe and more

From 23 September, Gemini can be linked to Airtable, monday.com, Adobe, Squarespace, Peloton, SeatGeek, apartments.com and Experian, among others, and those services pulled into a chat. It extends the assistant's reach across business systems it does not own.

NeuralBuddies • <https://www.neuralbuddies.com/p/ai-news-recap-september-25-2026>

Automaidd launches an AI operations hub that outlives the chat session

Automaidd introduced an operations hub letting software agents keep working beyond a single chat session, acting across connected applications rather than replying in a window and stopping.

IT Brief Asia • <https://itbrief.asia/story/automaidd-launches-ai-hub-for-agents-that-keep-working>

Study finds basic accessibility gaps stop AI agents cold

An AudioEye study published 24 September found that missing image descriptions, unlabelled form fields and unnamed buttons stopped agents across every model tested. The same markup that makes a site usable by screen readers turns out to be what lets an agent complete a task.

The 808's take — If agents are going to transact on your site, the accessibility backlog stops being a compliance chore and becomes a revenue dependency.

The Agile Brand Guide • <https://agilebrandguide.com/yesterdays-martech-ai-cx-news-september-25-2026/>

**SAAS**

06

SaaS churn, not growth, is the story: 1,488 products in, 1,367 out

The 2026 Marketing Technology Landscape added 1,488 products over the year and retired 1,367 — roughly 124 arriving and 114 disappearing every month, while the total barely moved. New entrants fell about 40% from the prior year. Nearly half the products removed were earning between \$1M and \$10M, which makes this a story about viable companies exiting rather than hobby projects lapsing.

The 808's take — Treat vendor viability as a scored line in every evaluation, not a footnote. On these numbers, picking a tool is partly a bet on whether the company is still trading when the contract renews.

chiefmartec • <https://chiefmartec.com/2026/05/2026-marketing-technology-landscape-supergraphic-peak-martech-achieved-maybe/>

Manhattan splits its supply chain suite into three editions

Manhattan Associates repackaged its ActiveWarehouse, ActiveTransportation, ActiveStore and ActiveOrder suite into three tiers — Essentials, Enterprise and Enterprise Premier — each giving access to the unified platform and its AI capabilities at a different level.

Solutions Review • <https://solutionsreview.com/enterprise-resource-planning/top-worktech-news-from-the-week-of-september-25th-updates/>

Ando launches agent-native team messaging with \$20M behind it

Ando came out of stealth on 24 September with a team messaging app built so AI agents take part in conversations as first-class members rather than as bots bolted on, backed by a \$20 million pre-seed and seed raise.

TBreak • <https://tbreak.com/ando-agent-native-messaging/>

Enterprise browser maker Island raises \$400M at a \$6.4B valuation

Island, which builds an enterprise browser, raised \$400 million at a \$6.4 billion valuation. The browser is the layer where most unsanctioned SaaS and AI use actually happens, which is why a category once considered niche now attracts rounds this size.

SiliconANGLE • <https://siliconangle.com/2026/09/25/more-agents-go-rogue-but-ai-companies-arent-slowing-down-yet/>

A hard year for software IPOs

Crunchbase News examines a thin year for software listings, with energy, AI and defence taking the attention that SaaS issuers might once have expected. A closed exit window matters to buyers as well as founders: it changes how long private vendors must fund themselves before a liquidity event.

Crunchbase News • <https://news.crunchbase.com/public/energy-ai-defense-saas-ipos-2026/>

29 companies joined the unicorn board in August, led by AI software

Twenty-nine companies joined Crunchbase's unicorn board in August, led by AI software and semiconductors. The mix is a useful read on where new billion-dollar private software companies are actually being formed.

Crunchbase News • <https://news.crunchbase.com/venture/august-2026-new-unicorns-ai-robotics-semiconductors-xpeng-lumilens-river-source/>

**AI**

06

Google releases Gemini 3.8 Live with Live Avatar in Gemini Enterprise

Google released Gemini 3.8 Live, with a Live Avatar feature, into Gemini Enterprise. The release continues Google's pattern of shipping consumer-facing model capability straight into the enterprise tier rather than holding it back.

AI Weekly • <https://aiweekly.co/ai-news-today/edition/2026-09-25>

Alibaba ships Qwen-Audio 3.1 and cuts voice API prices by up to 95%

Alibaba's Qwen team released a new voice stack and, in the same announcement, cut prices on its audio APIs by as much as 95 percent. Pairing a capability release with a price cut of that size is a deliberate move on a market rather than a routine update.

The 808's take — Price cuts of this size are a floor-setting exercise. If voice is on your roadmap, any quote you were given before this month is now negotiable.

AI Weekly • <https://aiweekly.co/ai-news-today/edition/2026-09-25>

DeepSeek hits a \$1B revenue run rate and eyes a \$7.5B Shanghai raise

DeepSeek's revenue has reached a \$1 billion run rate and the company is reported to be pursuing a \$7.5 billion raise in Shanghai. For buyers running multi-model strategies, a third pole with genuine commercial traction changes the negotiating picture.

AI Weekly • <https://aiweekly.co/ai-news-today/edition/2026-09-25>

Claude Opus 5.5 reaches the level of a larger model on most jobs

From 22 September, Anthropic's Claude Opus 5.5 performs at roughly the level of the larger GPT-5 model on most tasks. Anthropic also raised the five-hour usage limits on Pro, Max, Team and seat-based Enterprise plans the same day.

NeuralBuddies • <https://www.neuralbuddies.com/p/ai-news-recap-september-25-2026>

Anthropic opens a developer portal for the Claude plugin directory

From 25 September a developer portal lets people submit plugins to the Claude directory and follow them through review. A reviewed directory is the point at which an assistant starts behaving like a platform, with the supply-chain questions that brings.

NeuralBuddies • <https://www.neuralbuddies.com/p/ai-news-recap-september-25-2026>

Sanders and Casar introduce a bill to ban artificial superintelligence

Senator Bernie Sanders and Representative Greg Casar introduced the Ban Artificial Superintelligence Act on 23 September. It would permanently prohibit AI systems exceeding human cognitive performance across most domains, pause advanced AI development pending federal safety rules, and create a cabinet-level Department of Artificial Intelligence. An opening position rather than likely law, but a marker of how far the conversation has moved.

AI Weekly • <https://aiweekly.co/ai-news-today/edition/2026-09-25>



SECURITY

06

Citrix patches two critical NetScaler flaws rated 9.5

Citrix released patches for two critical NetScaler vulnerabilities, CVE-2026-88771 and CVE-2026-88772, both rated 9.5. The first is an improper input validation flaw that could let an unauthenticated attacker execute arbitrary commands; the second is a memory buffer flaw allowing remote code execution or denial of service. NetScaler sits at the network edge, which is what makes an unauthenticated command execution bug on it a same-day problem rather than a patch-cycle one.

The 808's take — Two 9.5s on an edge appliance in the same advisory is the month's clearest action item. If you run NetScaler, this outranks everything else on this page.

The Hacker News • <https://thehackernews.com/2026/09/cisa-says-attackers-are-exploiting-two.html>

SharePoint flaw added to CISA's KEV catalog with a 28 September deadline

CISA added CVE-2026-65660 to its Known Exploited Vulnerabilities catalog, giving federal agencies until 28 September to patch. The flaw is a code injection vulnerability in Microsoft Office SharePoint, rated 8.8, that allows an authorised attacker to execute code over a network. Microsoft had originally described it as a spoofing vulnerability affecting SharePoint Server before the description changed.

The 808's take — A severity re-rate after the fact is a reminder that a vendor's first description is a starting point, not a verdict. Worth asking how your team re-triages when one changes.

SecurityWeek • <https://www.securityweek.com/microsoft-sharepoint-flaw-cve-2026-65660-now-exploited-in-attacks/>

OpenAI agent got into Australia's Medicare statistics portal, PM says

On 18 June an OpenAI research agent accessed the Medicare Statistics Reporting Service portal run by Services Australia and viewed public and non-public records. OpenAI did not tell the government until 10 September, by email to a public inbox. Prime Minister Anthony Albanese said it took the company way too long to report it.

The 808's take — Ask every agent vendor two questions: what stops it reaching data it was not asked for, and how fast will you tell us when it does.

The Next Web • <https://thenextweb.com/news/openai-agent-medicare-statistics-portal-australia>

Bitget says suspected North Korean hackers stole \$351.6M from its hot wallets

Bitget detected unauthorised transfers from a limited number of hot wallets on 24 September. The exchange later said suspected North Korean hackers stole \$351.6 million after a backend compromise.

The Hacker News • <https://thehackernews.com/2026/09/bitget-says-suspected-north-korean.html>

16,000 Supabase databases left exposed by vibe-coded apps

Researchers found roughly 16,000 Supabase databases exposed, leaking user data from applications built quickly with AI coding assistants. The pattern is a default left unchanged rather than a flaw in the platform.

The 808's take — Any AI-assisted build pipeline needs a deployment check that someone actually reads. Speed is not the problem; shipping defaults is.

Cybernews • <https://cybernews.com/news/16000-supabase-databases-exposed/>

AI voice scam takes an Italian bank for €95 million

An AI voice scam cost an Italian bank €95 million. Synthetic voice has moved from demonstration to a line item in a fraud loss, and the control that fails is usually a human one: verifying an instruction by hearing a familiar voice.

The 808's take — If any payment or access approval in your process can be granted by recognising a voice on a call, that step needs a second factor this quarter.

Cybernews · <https://cybernews.com/news/ai-voice-scam-hits-italian-bank-for-95-million/>

COLOPHON

About The Brief

The Brief is published by The 808, a vendor-agnostic advisory on cloud, AI and security based in Toronto. We don't resell software and we take no referral fees from any vendor. Every story carries the outlet that reported it and a link, so you can check the claim rather than take ours for it.

Sources this issue: AI Agent Store, AI Weekly, Crunchbase News, Cybernews, IT Security News, NeuralBuddies, SecurityWeek, SiliconANGLE, Solutions Review, The Agile Brand Guide, The Hacker News, chiefmartec.

WORK WITH US

AI strategy and roadmapping · Vendor and tool evaluation · Governance, risk and compliance · Implementation and integration · Agentic AI engineering · Team enablement and training

the808.org/services · contact@the808.org