Copyright 2026 Treadstone 71. Reproduction without written permission is prohibited.

The Mind Domain Gets an Engine

Plain edition. How Iran, Russia, China and North Korea put AI to work, and why the AI a company runs now sits on the battlefield. Treadstone 71, 12 October 2026.

DATE	19 October 2026
AUTHOR	Treadstone 71
STATUS	Final
EDITION	Plain executive edition. Companion to the academic version of the same title.

On 10 September 2026, Anthropic named the state bodies and labs it caught using its model. The report shows Iranian offices drafting soft-war doctrine with Claude, Russian state media filing Claude-written copy, and Chinese police running sentiment work on a model. AI crossed a line that month. That shift sits at the center of this edition. The pages read the four states one by one, then turn to the harder exposure, the AI that sits inside a company.

The month mattered.

TL;DR

AI is now ordinary office gear for state influence and monitoring work. The catch is plain. Making volume is cheap. Effect almost never gets measured. The public record counts posts, videos, accounts and stolen tokens, and close to none of it shows a changed mind or a changed vote. A sharper risk sits inside the enterprise, where any firm that builds, buys or hosts AI holds something an adversary treats as an audience, a target and a doorway at the same moment.

One owner has to hold all three.

BLUF

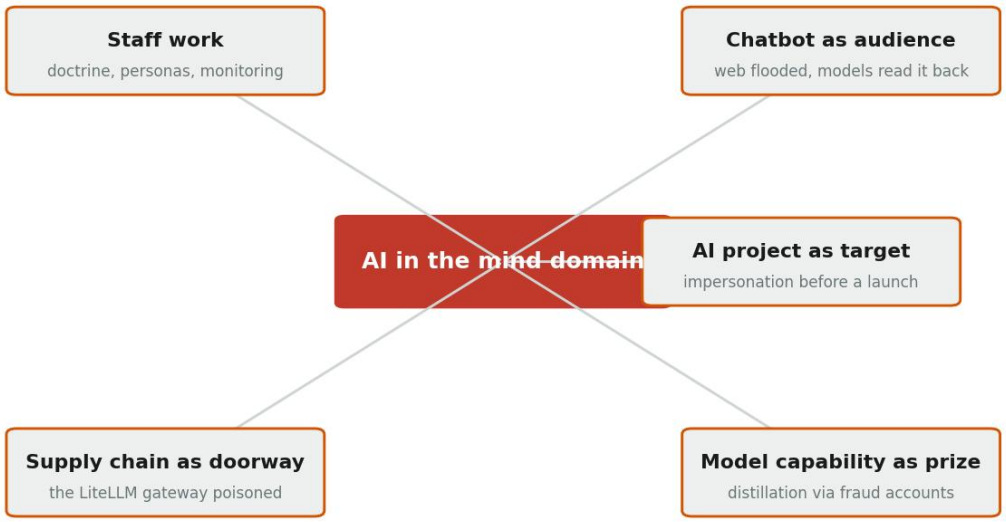


Figure 1. The five lines of adversary AI use, around one center. Staff work sits beside four lines that touch the enterprise directly.

Who. Four states and the hands they hire. Iran works through the Islamic Culture and Communications Organization, the Islamic Propaganda Office of Khorasan Razavi, and the Bina Cultural Observatory. Russia runs state-media pipelines at Sputnik and RIA Novosti, plus two networks, CopyCop, tracked as Storm-1516 and tied by researchers to the GRU, and Matryoshka. China fields municipal public security and state security bureaus, a police-linked cyber unit, and private firms that sell influence to provincial clients. North Korea runs Kimsuky and its IT-worker networks. The evidence shows no joint plan. Shared use of Western commercial AI makes the pattern look common, and no single command stands behind it.

What. AI moved from experiment to staff work across all four states. Anthropic documents Iranian bodies using Claude for doctrine, persona systems and posting schedules, Russian state-media output traced to Claude-written articles and broadcast tickers, and Chinese police units running sentiment monitoring on a model. Four more lines run beside that staff work. Chatbots became an audience, as a Russian network floods the open web with material that AI systems later read back. AI projects became a target, when CopyCop ran impersonation video against an Armenian AI facility before it opened. Model capability became a prize, as Chinese labs drain frontier systems through fraudulent accounts. Supply chains became a doorway, after the FBI found that a criminal group had poisoned the LiteLLM gateway that holds provider keys.

So what. Volume costs almost nothing to make or repeat. A defender who budgets against raw output chases a cost curve the adversary sets. Whether any of it lands is a different story, and close to no one runs that measure. The enterprise exposure cuts sharper than the influence story, because one owner must cover the audience, the target and the channel together.

Volume is cheap, and belief is not.

Why now. Four clocks run at once. The war with Iran opened on 28 February 2026, and Iranian outlets frame the fight as cognitive war. Midterms fall on 3 November 2026. A joint advisory of 8

September from NSA, CISA and the FBI warned of industrial-scale model distillation by China-based companies, named six firms, and judged the activity likely proceeded with Chinese government awareness while stating no basis. Russia signed Federal Law 243-FZ on 26 July 2026, a framework with no penalties, in general effect from 1 September, with its sovereign model classes set to start on 1 March 2027. Part of the clock is a publication calendar, since 25 of 61 dated ledger rows fall in September and October 2026.

Impacts so far. Documented reach is mostly small. OpenAI found no meaningful breakout for a PRC-linked operation aimed at American AI debates and rated it Category One. ISD found that Matryoshka made basic errors and kept limited reach in Germany. Two cases pull the other way. OpenAI put the Russian-run Burke Institute operation at the lower end of Category Three, with some breakout to real audiences, and tied its weight to infrastructure over reach. Anthropic reported a Category Four case, tracked as GTG-04001, in which a Russian-speaking operator ran daily content on Radio Lengo Songo, an FM station in the Central African Republic. The biggest audience numbers belong to networks no source ties to AI output. Two pro-Iran X networks pulled more than a billion views in the first month of the war, and ISD found no sign of state backing. The real losses sit on the enterprise side. Anthropic counted over 16 million exchanges and about 24,000 fraudulent accounts across three Chinese labs in February, and more in September, led by over 151 million exchanges for Alibaba. Those are provider counts, and no government verified them. No reviewed source shows an election outcome that AI changed.

Outlook. The baseline leads through 2027, and it leads thinly. AI becomes ordinary staff equipment, effect stays unproven, and the loudest numbers stay with networks no source links to AI. That path holds 39 percent, Unlikely (15 to <45). A shift toward local models that labs cannot watch holds 30 percent, Unlikely (15 to <45). The poisoned-machine era, with confirmed distortion of chatbot or retrieval answers, takes 18 percent, Unlikely (15 to <45). Shock events with comparison evidence take 5 percent, Highly unlikely (5 to <15). The rest will not resolve. Confidence in every share is Low. China leans toward the local shift, Russia toward the poisoned machine, Iran toward the shock.

The headline judgments

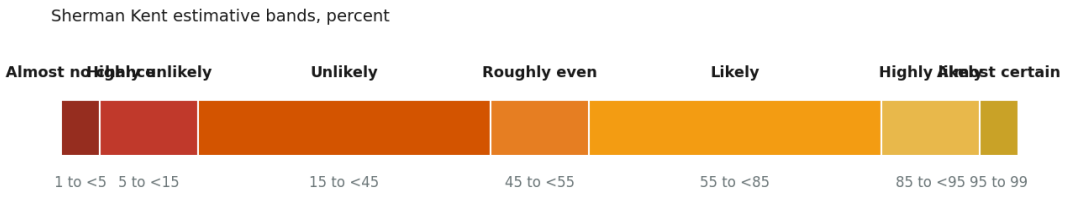


Figure 2. Sherman Kent estimative bands used across the judgments below.

Every judgment below carries a likelihood band and a separate confidence rating. The weakest supporting source sets the confidence, and a judgment resting on a gap or on our own reasoning steps down.

Iran

Iran runs the most open case on the page. Its operators named their own offices inside the prompts. Anthropic shows those offices running Claude for doctrine, personas and posting work, which lands at Likely (55 to <85), Moderate beside Russia and China. State direction of the two pro-Iran networks that drew over a billion views is a separate matter. Open sources show alignment and no order, so that claim prints at Unlikely (15 to <45), Low. Strip Anthropic from the file and the Iranian staff-work claim falls to Roughly even chance, Low.

Iran is the clearest case, and the loudest claim about it is the weakest.

Russia

Russian state-linked content sits in the Common Crawl data that trains models, at Likely (55 to <85), Moderate. Whether that content reaches chatbot answers through retrieval and bends what users believe, at scale, stays unproven. ISD found state-attributed sources in nearly 18 percent of 300 test answers. Citation is not belief. The reach-to-belief claim holds at Roughly even chance (45 to <55), Low. Researchers tie CopyCop to the GRU, and that network aimed fake clips at an Armenian AI site before it opened.

China

A US advisory of 8 September says Chinese labs distilled American frontier models at industrial scale through fraudulent accounts, which comes in at Likely (55 to <85), Moderate. Six firms appear. The list runs DeepSeek, Moonshot, Alibaba, MiniMax, StepFun and Z.AI. Whether the Chinese government knew sits at Roughly even chance (45 to <55), Low, since the advisory gives no basis for its awareness line. Hired firms that serve thin-capacity states surface through GoLaxy, whose own documents name PLA Unit 31040 in a 2020 procurement letter. That hired-hand claim prints at Roughly even chance (45 to <55), Low.

North Korea

North Korean units run AI for fake identities and lure documents, inside the same staff-work judgment at Likely (55 to <85), Moderate. Kimsuky built decoys with an agent and left local tooling behind. The case points to a model Western labs cannot watch. That is the blind spot.

The enterprise

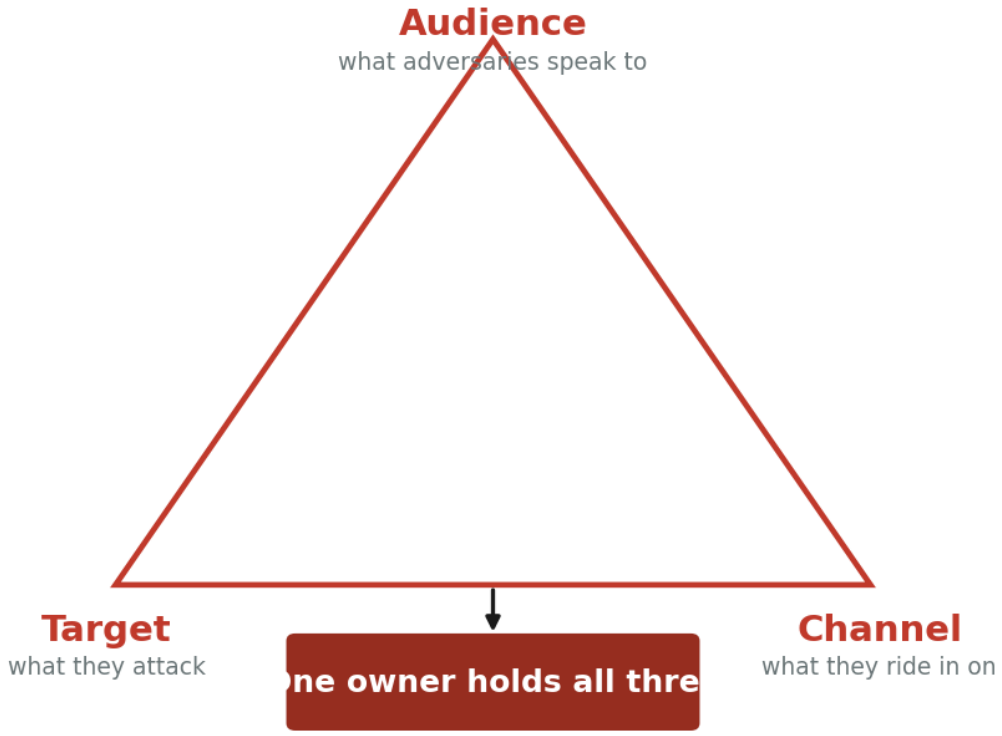


Figure 3. The enterprise exposure. One firm holds the audience, the target and the channel together, so one owner must cover all three.

The strongest judgment on the page sits here. Criminal groups compromise AI gateways, provider keys and agent frameworks now, and one break reaches many keys, at Likely (55 to <85), Moderate. The FBI, Palo Alto Unit 42 and Wiz each documented the LiteLLM case from a separate seat, and the FBI has nothing to sell. Reports put the exposure window near three hours. The gap stayed open that long. Wiz finds LiteLLM in 36 percent of the cloud setups in its data. Agentic intrusion runs at a pace that shrinks a defender's warning time, at Likely (55 to <85), Low. Deepfake payment fraud against finance teams stays live, at Roughly even chance (45 to <55), Low, on a single 2024 case in which Arup booked a 25.4 million pound loss.

Here the report pushes hardest.

What the labs cannot see

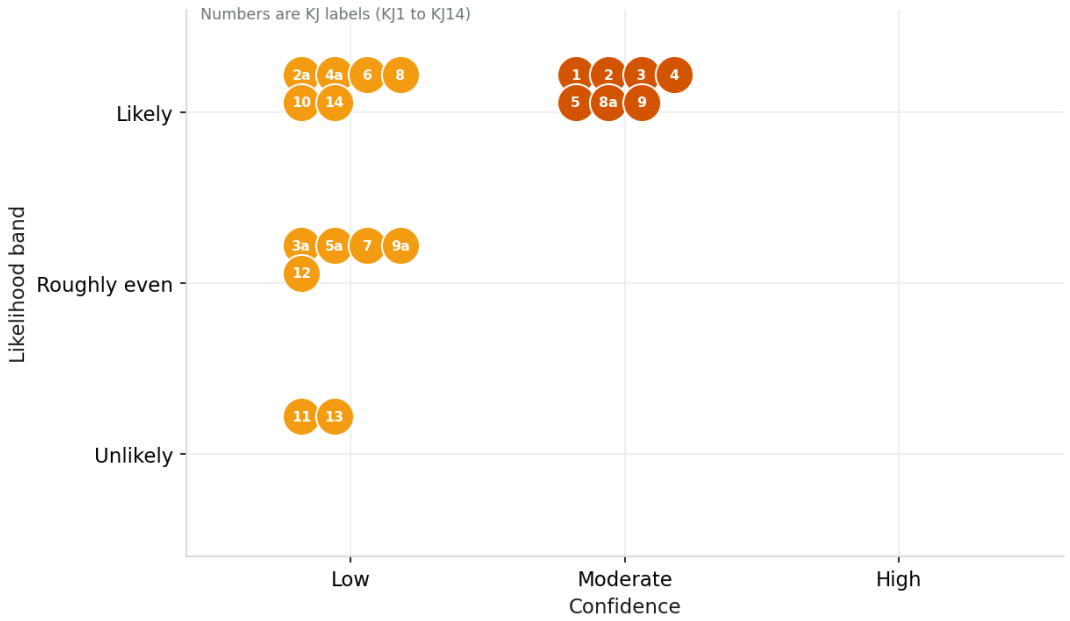


Figure 4. The key judgments by likelihood and confidence. No judgment reaches High confidence, and the weakest sit at Roughly even or Unlikely.

Operators in Russia, China and North Korea run work on models the labs cannot watch, so lab reports undercount, at Likely (55 to <85), Low. The measured local-model share reads 2 of 16 non-provider rows, a 3 to 36 percent interval, or 4 of 34 once the extended set joins. Local use stands below one third of state-linked AI operations today, at Likely (55 to <85), Low. That ceiling leans on a middle setting, since a lower one drops the chance from 73 percent to 42.

The outlook in paths

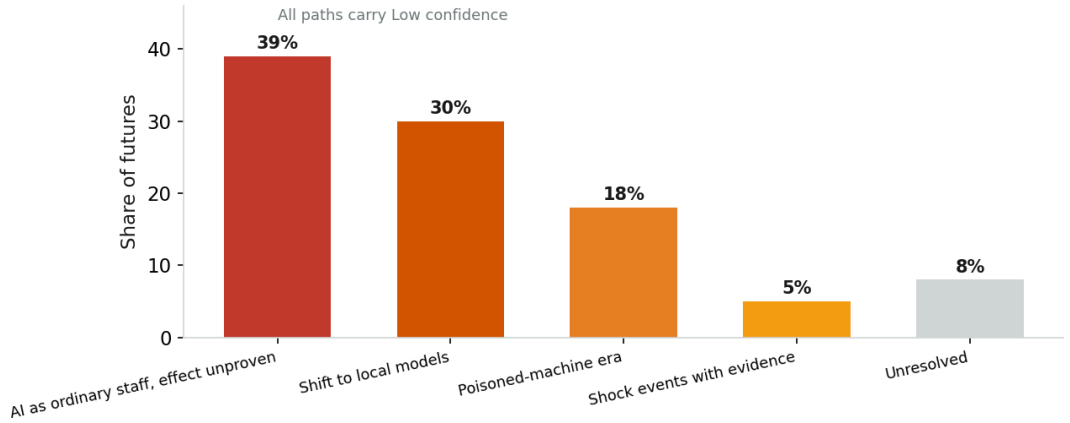


Figure 5. The four paths through 2027, plus the share that will not resolve. The top two are close enough to read as a tie.

The order of the top two paths does not hold as a finding. Three coding choices decide how far apart they sit, and a single recoded OpenAI row closes or flips the lead. Quiet from the labs would not mean calm, since a silent provider route pushes the local-model path from 30 up to 50.

Read the top two as a tie.

What to do

Name one owner for AI exposure across influence, supply chain and agent permissions, and inventory every gateway and credential store, by 31 October. Strip tool permissions from agents until each earns them back, and trace what feeds the retrieval layer, by 30 November. Grade every AI-made analytic output on the ASA risk clusters, add live-video checks to remote hiring, and require a callback on any payment instruction that arrives by video call, by 31 December. Rehearse an impersonation attack against the next AI launch.

The gaps

Four holes mark the government record. No government source measures belief or behavior change from a live operation. None reports observed adversary use of local or self-hosted models. Seven executive bodies, Justice and State and the White House among them, carry no midterm AI-influence item from 1 September forward. A War Department directive of 22 September names no AI, and CISA's 2026 election plan carries no AI text and no foreign-influence task.

How to read this

Count the output if you must, but price the effect. The firm that runs AI owns the exposure that matters, and the owner has to be one person holding the whole picture. Watch three signals: the share running on local models, the first government measure of effect, and the first clean audit of a poisoned chatbot.

Those three move the forecast.