

How to Read and Digest The Mind Domain Gets an Engine

A reader's guide to the 167-page report

DATE	12 October 2026
AUTHOR	Treadstone 71
STATUS	Final

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

Document control

Date	12 October 2026
Author	Treadstone 71
Status	Final

Copyright 2026 Treadstone 71 - www.treadstone71.com

Contents

1 What you hold 5
2 The answer in a minute 5
3 Choose a route 5
4 The map 6
5 Reading the grammar 7
5.1 Two ratings on every judgment 7
5.2 Grades on the evidence 8
5.3 Codes you will meet 8
6 Habits that keep you from over-reading 8
7 A full-read order 9
8 How to trace a claim to its source 9
9 After you read: decisions and dates 10
10 What the report does not claim 10

Copyright 2026 Treadstone 71 - www.treadstone71.com

List of tables

Table 1. The five paths to 31 December 2027 5
Table 2. Routes by time available 6
Table 3. Routes by role 6
Table 4. What each part does 6
Table 5. The Sherman Kent bands, with half-open edges 7
Table 6. Identifiers and their home sections 8
Table 7. Eight reading habits 9
Table 8. A two-pass plan for the long sections 9
Table 9. Dates that follow from the report 10

Copyright 2026 Treadstone 71 - www.treadstone71.com

1 What you hold

The report runs 167 pages across 17 numbered sections, a glossary, 75 tables and 25 figures. Few readers take it front to back in one sitting, and the report does not ask them to.

Three layers sit inside. The top one answers the question in four pages. Judgments make the middle: 20 of them, each with a likelihood and a confidence rating. Proof sits at the bottom, in scoring engines, nine analytic tests, a forecast model and a 63-row evidence ledger. Each layer explains the one above it. Read down until you have what you need. Then stop. Most readers need the answer and one section.

Page numbers in the guide follow the PDF of the report. A Word copy shifts by a page or two. Section numbers hold either way.

2 The answer in a minute

Iran, Russia and China use AI as standing staff. Models draft, translate, run personas, profile targets and write reports. North Korea uses them for fake identities and lure documents. No source measures the savings in cost or the gain in output, and belief change goes unmeasured too. Documented reach stays small in most cases.

The sharper change sits at the edges. Chatbots now serve as an audience. AI projects serve as targets. Stolen model capability is a prize, and the AI supply chain is a way in. The report calls the enterprise exposure sharper than the influence story, and it is the part a reader is able to act on in October.

Labs and vendors that sell the models supply almost all the public evidence. The report reads that as a map of what they see, and the map misses whatever runs on an adversary's own hardware. Local models appear in 2 of 16 non-provider ledger rows (3 to 36 percent). With the extended set the count reads 4 of 34 (12 percent).

The outlook through 2027 splits five ways. Confidence in every share is Low, and the order of A and B does not hold as a finding.

Table 1. *The five paths to 31 December 2027.*

Path	What it means	Share (percent)	Band
A	AI becomes ordinary staff equipment, with no measured effect	39	Unlikely (15 to <45)
B	Local-model rows reach 30 percent of the counted non-provider rows in 2027	30	Unlikely (15 to <45)
C	A poisoned-machine era, with measured distortion of chatbot answers	18	Unlikely (15 to <45)
D	A shock event with comparison evidence	5	Highly unlikely (5 to <15)
E	The rules cannot resolve the year	8	Highly unlikely (5 to <15)

Four clocks run at once. The war with Iran began on 28 February 2026. Midterms fall on 3 November 2026. Anthropic published on 10 September, two days after the NSA, CISA and FBI advisory on distillation. Russia's Law 243-FZ starts its sovereign and national model classes on 1 March 2027.

3 Choose a route

Start with the time you have. Pick one row. The pages in each row are the minimum, and every row builds on the one above it.

Copyright 2026 Treadstone 71 - www.treadstone71.com

Table 2. Routes by time available.

Time	Read	Pages	You leave with
15 minutes	The front summary, then the wrap-up in Section 17	9 to 12 and 164 to 165	The answer, the headline numbers and the first actions
1 hour	Add Section 2, the five paths in 12.3 and the opening of 12.18	16 to 20, 94 to 95 and 119 to 120	The 20 judgments, the path shares and the election window
Half a day	Add the section for your actor or concern, then Section 8, 11.1, 12.2 and 12.16	See the role table. 60 to 63, 78 to 81, 90 to 93 and 115 to 116	The evidence behind the judgments you rely on, and how far it bends
Full read	Front to back, in the order of Section 7 of this guide. Treat Section 16 as a lookup table	9 to 165	The whole case, with the proof

Then match the route to your job. The table gives the first stop and the second.

Table 3. Routes by role.

Reader	First stop	Second stop
Security owner or chief of an AI program	Section 7 on the enterprise exposure (pp. 52 to 59) and 13.1 on nine actions (p. 126)	7.6 on the decision test (p. 58) and 13.6 on cost (pp. 130 to 131)
Buyer of counter-influence work	13.2 (p. 126)	11.5 on the effect ceiling (p. 84)
Analyst or editor	1.4 to 1.6 on stance, bias and assumptions (pp. 14 to 16) and 13.3 (p. 126)	Section 10 on patterns (pp. 71 to 77), Section 11 on tests (pp. 78 to 88) and 16.2 on rows near a band edge (p. 146)
Election and policy watcher	9.1 on the American record (p. 64) and 12.18 on the election window (pp. 119 to 125)	12.12 on ten dated forecasts (pp. 111 to 112)
Operator in Europe, Australia, Singapore or the Gulf	13.5 (pp. 127 to 129)	9.3 on the European services (p. 67)
Reader tracking one state	Iran in Section 3 (pp. 21 to 26), Russia in 4 (pp. 27 to 36), China in 5 (pp. 36 to 45) or North Korea in 6 (pp. 45 to 52)	The matching row in 12.9 and 12.10 on tilts and scored futures (pp. 105 to 109)

4 The map

Evidence by actor and by problem sits in Sections 3 to 7. Scoring, testing and projection fill 8 to 12. Action comes in 13. Section 16 proves the whole case, row by row. A reader who only wants to check one claim needs three stops: the Basis column in Section 2, the section cited there, and the ledger row behind the sentence, which prints the source, the grade and the limits on a single line.

Table 4. What each part does.

Part	Pages	What it does	How to use it
Front matter	1 to 8	Document control, contents and the lists of 75 tables and 25 figures	Skip, except to find a table by number
Summary	9 to 12	A lead paragraph, then seven headed blocks: Who, What, So What, Why Now, Impacts So Far, Outlook, Recommendations	Read first. It stands alone
1 Scope, method and assumptions	12 to 16	The question, six priority requirements, the method stack, stance, bias and eight assumptions	Read 1.4 to 1.6 and skim the rest

Copyright 2026 Treadstone 71 - www.treadstone71.com

Part	Pages	What it does	How to use it
2 Key Judgments	16 to 20	20 judgments, each with a band, a confidence and a basis. A second table removes Anthropic's rows and shows what survives	Read all of it. It is the core
3 Iran, 4 Russia, 5 China, 6 North Korea and hired hands	21 to 52	The evidence for each state, including what its media say and how it is likely to react	Read the state you care about. Skim the rest
7 Enterprise AI	52 to 59	Gateways, accounts, assistants and agents as target and channel, with the counterweights	Read in full if you run AI
8 Scoring	60 to 63	Ten threats ranked, eight campaigns scored, four audiences mapped, one decision tested	Read 8.1 and 8.5
9 Governments	64 to 70	What allies and agencies say, and what they leave out	Read 9.1, 9.3 and 9.5
10 Patterns	71 to 77	How the evidence itself behaves: who reports, when, and how claims lose support as they climb	Read if you judge sources for a living
11 Analytic tests	78 to 88	Nine checks on the judgments, from competing hypotheses (11.1) to the hindcast audit (11.9)	Read 11.1, 11.5 and 11.9
12 Outlook	88 to 125	Counted baselines, five paths, indicators, ten dated forecasts and the election window	Read 12.2, 12.3, 12.12, 12.16 to 12.18 first
13 So what	126 to 131	Actions for three readers, a dated calendar, a regional read and the cost of each decision	Read the part with your name on it
16 Evidence ledger	136 to 163	63 graded rows, then adversary statements and sources read after the freeze	Look up by row ID
17 Keywords and wrap-up	164 to 165	A closing read of the whole	Read second, after the summary
Glossary	166 to 167	Every term, method and abbreviation	Keep open while you read

5 Reading the grammar

The report speaks in a fixed grammar. Learn it in ten minutes and every page gets faster.

5.1 Two ratings on every judgment

Each judgment carries a likelihood and, apart from it, a confidence. Treat them as two gauges on one instrument. The likelihood says how probable the statement is. Confidence says how good the evidence behind that estimate is. A judgment rated Likely with Low confidence is a reasonable reading on thin ground, and the page treats it as one. Read the two together.

Table 5. *The Sherman Kent bands, with half-open edges.*

Band	Range (percent)
Almost certain	95 to 99
Highly likely	85 to <95
Likely	55 to <85
Roughly even chance	45 to <55
Unlikely	15 to <45
Highly unlikely	5 to <15
Almost no chance	1 to <5

Copyright 2026 Treadstone 71 - www.treadstone71.com

Confidence follows the ASA T-Score of the weakest row that supports the judgment. Strong evidence (80 and above) earns High. Adequate (60 to 79) earns Moderate. Weak (40 to 59), or a single belligerent source, earns Low. Three rules then adjust the print. Adequate evidence caps the likelihood at Likely. A gap in the record, or the report's own reasoning, steps confidence down once. Absence of a measure takes Roughly even chance when a working path exists and nothing shows use at scale.

The result is a cautious set. Of the 20 judgments, none reaches Highly likely. Thirteen stand at Likely, five at Roughly even chance and two at Unlikely. Seven carry Moderate confidence and thirteen carry Low. None carries High.

5.2 Grades on the evidence

Every ledger row has an Admiralty anchor such as B2. The letter grades the source, A to F. Numbers 1 to 6 grade the information. Letter and number print apart, and the report never averages them. Rows at Level 2 and Level 3 add a T-Score for quality and a D-Score for the likelihood of deception, while Level 1 rows carry the anchor alone.

Two cautions come from the report itself. All T, D and F scores are expert ratings, and no validation numbers exist yet. Twenty-seven of the 48 scored rows sit within five points of a band edge, so a one-point difference is noise. Two judgments, KJ1 and KJ8a, sit on exactly that kind of cliff at T 60. A second analyst who scores the weakest row one point lower moves each from Moderate to Low.

5.3 Codes you will meet

Table 6. Identifiers and their home sections.

Code	Meaning	Where
KJ1 to KJ14	Key judgments. A lettered judgment such as KJ3a is a companion claim, rated on its own	Section 2
IR, RU, CN, KP	Ledger rows for Iran, Russia, China and North Korea	Section 16
PX, EN, GV	Ledger rows for proxies and hired hands, for enterprise AI and for governments	Section 16
RC	A row from the extended set, coded once by one reader	16.1
H1, H2, H3	The rival hypotheses: Efficiency Amplifier, Capability Shift and Signal Artifact	11.1
Paths A to E	The five ways 2027 closes, each with a resolution rule	12.3
T, L, M	Tripwires (T1 to T8), leading indicators (L1 to L10) and one simulated-count indicator (M1)	12.6
Forecasts 1 to 10	Dated forecasts with point probabilities	12.12
W1 to W20, WE1 to WE6, WS1 to WS4	Calls, estimates and scenarios for the window to 9 November	12.18
PIR 1 to 6	The six priority intelligence requirements	1.2
A1 to A8	The eight assumptions, each with what follows if it fails	1.6

The method names (ATCRI, CWIA, CWC, HTIM, CARM, ACS, DBNA, HET, MRA and ATVA) each appear in the glossary on pages 166 to 167. Keep that page open on the first pass.

6 Habits that keep you from over-reading

The report works hard to avoid overclaiming. A reader who skims is able to undo that work in one step. Eight habits prevent it. Each takes seconds.

Table 7. Eight reading habits.

Habit	Why it matters	Where the report sets it
Read presence, not persuasion	The 2026 record proves that operations exist. It almost never shows that anyone changed their mind	1.4 and 11.5
Read the labs as a floor	A lab sees its own users and publishes what it chooses. Local use stays unseen	1.5, A1 in 1.6 and 13.3
Count the Anthropic rows	One lab with a commercial interest supplies the heaviest evidence. Read the with-and-without table at the end of Section 2	1.5 and Section 2
Keep fact and claim apart	A ministry statement or a speech enters as a claim at grade C, and never shares a sentence with a verified fact	1.4
Sort every number into three kinds	Counted, modeled or judged. Percentages in Section 12 are judgments, and the model adds arithmetic, not evidence	12.1
Treat A and B as a tie	A leads B by 8.6 points, and modest changes in coding or in the rule threshold reverse the order. Read the headline as "A and B are close"	12.3 and 12.16
Read silence as capacity	A thin American record fits a changed posture. No source shows adversary activity falling	9.1 and 13.3
Weigh the house interest	Treadstone 71 sells analysis and training, and findings that raise urgency favor those sales	1.5

7 A full-read order

Read in the printed order, with four exceptions. Take 1.4 and 1.5 before any actor section, since they set the rules of the grammar. Skip nothing in Section 2. Read it before Sections 3 to 7, so each section arrives with a job: it has to support a judgment you already know. Section 10 comes before Section 11. Finish 12.1 to 12.3 before the rest of Section 12.

Table 8. A two-pass plan for the long sections.

Section	First pass	Second pass
3 to 6	Opening paragraph and the closing reaction read of each state	The full section for the states that matter to you
8	8.1 and 8.5	8.2 to 8.4
11	11.1, 11.5 and 11.9	11.2 to 11.4, 11.6 to 11.8
12	12.1 to 12.3, 12.12, 12.16 and 12.17, then 12.18	12.4 to 12.11, 12.13 to 12.15
13	The subsection for your role	13.4 to 13.6
16	The opening paragraph and 16.2	Row lookups only

Sections 12.4 to 12.11 run 15 pages of drivers, cones, indicator tables and scored futures. They record how the path shares came together. A reader who accepts the paths as stated loses little by leaving them for a second pass. Someone who reaches Section 12 with an hour left should stop at 12.3, return to 12.12 and 12.18 for the dated calls, and keep the indicator tables, the cones and the scored futures for the day the path shares need defending.

8 How to trace a claim to its source

Every claim in the report leads back to a row. The route has four steps. None is hard.

First, find the judgment. Section 2 lists 20, and its Basis column names the sections and the ledger rows behind each one. Second, open the section it names and read the paragraph that carries the claim. Third, go to Section

Copyright 2026 Treadstone 71 - www.treadstone71.com

16 and find the row by its ID. Each row prints the source, the date, the anchor, the T-Score, the D-Score and a note on limits. Fourth, follow the link in the row to the source itself.

Two examples show the route. KJ4 says criminal groups compromise AI gateways now. The basis points to EN-1, the TeamPCP and LiteLLM compromise, at T 80, and to EN-6, the Nx and Salesloft precedent, at T 81. Each sits within five points of the 80 edge, so the edge rule prints the judgment one band lower, at Likely with Moderate confidence.

KJ5a says the Chinese government knew of the distillation. The advisory states the awareness judgment and gives no basis, no source and no confidence level. Roughly even chance, Low confidence: that is the print.

To test a number rather than trace it, go to 12.16, which shows what moves the headline, and to 12.17, which lists four observations that break the picture and the review dates that follow.

9 After you read: decisions and dates

Section 13.1 sets nine actions for the enterprise that runs AI. Person-days of effort for each sit in 13.6, and 13.4 holds the full calendar. Five dates matter in the next three months.

Table 9. *Dates that follow from the report.*

Date	What falls due	Source
31 Oct 2026	Name one owner for AI exposure. Inventory every gateway, key store and agent runtime	13.1, actions 1 and 2
3 Nov 2026	Midterms. Forecast 1 resolves	12.12
9 Nov 2026	The election window closes. Calls W1 to W20 and estimates WE1 to WE6 are scored	12.17 and 12.18
30 Nov 2026	Strip agent permissions and trace what feeds the retrieval layer	13.1, actions 3 and 4
31 Dec 2026	Grade AI analytic output, add live-video checks to hiring, require a callback on video-call payments	13.1, actions 5, 7 and 9

Two actions carry no date. Rehearse an impersonation attack before the next AI launch, and watch account markets and infostealer logs as a standing task.

10 What the report does not claim

Five absences define the edge of the work. The record holds no government measure of belief or behavior change from a live operation. Western government sources report no observed adversary use of local or self-hosted models. Nothing in the ledger shows a joint plan among the four states. Neither does any row show an election outcome that AI changed. The public pages of seven executive bodies hold no item from 1 September onward on AI-enabled influence aimed at the midterms, as of 7 October.

Read the five absences before you quote the report, since a sentence lifted without them claims more than the pages do.

Copyright 2026 Treadstone 71 - www.treadstone71.com