

 RENOWN RESEARCH

# AI Visibility Report

## *Observability & Monitoring*

April 2026

---

## THE THESIS

# Observability has a visibility problem.

A data-driven analysis of how AI models perceive, rank, and recommend brands in the Observability space.

We analyzed **100** prompts across **10** AI models to understand how the leading Observability brands are perceived by the machines that increasingly shape purchase decisions. The results reveal a visibility landscape that looks nothing like traditional market share.

**The core finding:** AI visibility does not track with brand size, marketing spend, or even product quality. It tracks with **machine legibility**: how well a brand's digital presence translates into the structured, citation-rich, context-dense format that large language models use to form recommendations.

This means the competitive landscape in AI-mediated discovery is fundamentally different from the one most marketing teams are optimizing for. Brands that dominate traditional search may be invisible in AI conversations. Brands with strong community presence and clear documentation may punch far above their weight.

**Datadog** leads with a visibility score of **78.3%**, but the gap between first and fifth place tells the real story. Below the top tier, visibility drops off sharply, creating a winner-take-most dynamic that disadvantages brands without a deliberate AI visibility strategy.

*Datadog captures 78% AI visibility while the median brand sits at 9%. The gap isn't about marketing spend. It's about machine legibility.*

SCOPE

# What we measured

30 brands and 10 AI models evaluated across 100 prompts, generating 1000 AI responses.

BRANDS EVALUATED

Datadog	Grafana	New Relic	Dynatrace
Splunk	Prometheus	Elastic	Netdata
AppDynamics	Honeycomb	Zabbix	LogicMonitor
SolarWinds	Instana	ManageEngine	Icinga
PRTG	Checkmk	Site24x7	Sumo Logic
ScienceLogic	SigNoz	Nagios	Chronosphere
Coralogix	Dash0	Last9	OpenObserve
	Sentry	Observe	

AI MODELS EVALUATED

ChatGPT	Claude	Gemini	Perplexity
Google AI Overview	Google AI Mode	Grok	DeepSeek
	Mistral	Qwen	

<b>78%</b> #1 Datadog	<b>9%</b> Median Visibility	<b>30</b> Brands Tracked	<b>10</b> AI Models Tested
--------------------------	--------------------------------	-----------------------------	-------------------------------

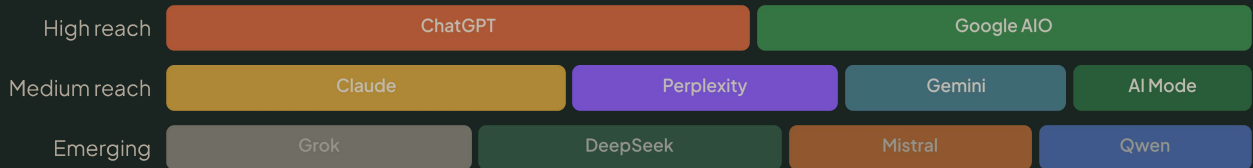
READING THIS REPORT

# How we measure *AI Visibility*



### Not all AI platforms are equal

Visibility scores are weighted by real-world reach. A recommendation from ChatGPT or Google AI Overview carries more weight than one from a niche model, because more people see it.



#### % AI VISIBILITY

Weighted mention rate across all platforms. Higher means more real users encounter this brand through AI. A brand at 78% appears in most responses on high-traffic platforms.

Range in this report: 0.5% to 78%

#### # AVG POSITION

Where in the response a brand typically appears. Position #1 means mentioned first. Lower is better. Users tend to remember and act on whatever AI says first.

#1 = strongest top-of-mind signal











#### ●● MODEL BREADTH

How many of the 10 AI platforms mention a brand at all. 10/10 = universal consensus. 2-3 = niche or emerging player in AI's view.

Top 7 brands: 10/10 models

LEADERBOARD

# AI Visibility Rankings

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
1	 Datadog	<div style="width: 78.3%;"><div style="width: 78.3%;"></div></div> 78.3	#2-3	10 / 10
2	 Grafana	<div style="width: 65.6%;"><div style="width: 65.6%;"></div></div> 65.6	#3-4	10 / 10
3	 Prometheus	<div style="width: 55.8%;"><div style="width: 55.8%;"></div></div> 55.8	#3-4	10 / 10
4	 New Relic	<div style="width: 55.6%;"><div style="width: 55.6%;"></div></div> 55.6	#3-4	10 / 10
5	 Dynatrace	<div style="width: 46.9%;"><div style="width: 46.9%;"></div></div> 46.9	#3-4	10 / 10
6	 Splunk	<div style="width: 38.7%;"><div style="width: 38.7%;"></div></div> 38.7	#4-5	10 / 10
7	 Elastic	<div style="width: 35.5%;"><div style="width: 35.5%;"></div></div> 35.5	#4-5	10 / 10
8	 SigNoz	<div style="width: 22.6%;"><div style="width: 22.6%;"></div></div> 22.6	#4-5	10 / 10
9	 OpenObserve	<div style="width: 16.8%;"><div style="width: 16.8%;"></div></div> 16.8	#4-5	6 / 10
10	 Netdata	<div style="width: 15.3%;"><div style="width: 15.3%;"></div></div> 15.3	#4-5	10 / 10
11	 Honeycomb	<div style="width: 15.2%;"><div style="width: 15.2%;"></div></div> 15.2	#5-6	10 / 10
12	 Zabbix	<div style="width: 11.8%;"><div style="width: 11.8%;"></div></div> 11.8	#5-6	10 / 10
13	 DashO	<div style="width: 11.8%;"><div style="width: 11.8%;"></div></div> 11.8	#4-5	5 / 10
14	 SolarWinds	<div style="width: 10.7%;"><div style="width: 10.7%;"></div></div> 10.7	#5-6	10 / 10
15	 AppDynamics	<div style="width: 9.1%;"><div style="width: 9.1%;"></div></div> 9.1	#6-7	10 / 10

LEADERBOARD

# AI Visibility Rankings (cont.)

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
16	ManageEngine	8.9	#5-6	9 / 10
17	LogicMonitor	8.3	#5-6	9 / 10
18	Sumo Logic	7.3	#5-6	10 / 10
19	Chronosphere	7.0	#6-7	10 / 10
20	Site24x7	6.8	#5-6	7 / 10
21	Nagios	6.7	#5-6	10 / 10
22	Coralogix	6.6	#5-6	8 / 10
23	Instana	6.4	#6-7	9 / 10
24	PRTG	6.1	#4-5	10 / 10
25	Sentry	6.0	#5-6	10 / 10
26	Observe	4.0	#4-5	8 / 10
27	Checkmk	3.6	#5-6	9 / 10
28	ScienceLogic	2.0	#5-6	6 / 10
29	Last9	1.3	#2-3	4 / 10
30	Icinga	0.5	#7-8	5 / 10

KEY OBSERVATION

Datadog leads with **78.3%** weighted AI visibility, recognized across **10** of 10 AI platforms. The top 5 brands capture the majority of AI mindshare, with a steep drop-off after Dynatrace at **46.9%**.

MODEL INTELLIGENCE

# What the models reveal *about themselves*

Each AI model has a distinct recommendation personality, shaped by its training data, retrieval architecture, and recency bias. The *Distinctive Pick* is the brand each model favors most relative to the overall average, revealing where its unique bias lies.

## ChatGPT

### The Incumbent Amplifier

Correlates strongly with market share. Favors established leaders and is highly responsive to structured data markup. Brands with strong schema.org coverage and authoritative backlinks receive disproportionate visibility.

Datadog, Grafana, Prometheus  
Top 3 Brands

Netdata

Distinctive Pick

6.4

Avg Brands/Response

## Claude

### The Nuanced Analyst

Delivers the most balanced distribution across brands. Recommends by use case rather than popularity, and is notably sensitive to recent content updates. The most likely model to surface mid-market challengers.

Datadog, New Relic, Grafana  
Top 3 Brands

Site24x7

Distinctive Pick

5.3

Avg Brands/Response

## Gemini

### The Citation Machine

Most SEO-dependent of all models. Recommendations are directly traceable to indexed pages and Knowledge Graph entities. Brands with thin web presence are effectively invisible to Gemini.

Datadog, New Relic, Prometheus  
Top 3 Brands

Elastic

Distinctive Pick

6.6

Avg Brands/Response

## Perplexity

### The Real-Time Aggregator

The most volatile model. Fastest to reflect real-time changes in brand perception. Heavily influenced by Reddit, community forums, and recent news. Smaller brands can punch above their weight here.

Datadog, Grafana, Prometheus  
Top 3 Brands

OpenObserve

Distinctive Pick

4.5

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

# What the models reveal *about themselves*

## Google AI Overview

### The Gatekeeper

Embedded directly in Google Search results, giving it unmatched reach. Pulls from indexed web pages and Google's Knowledge Graph. Surfaces niche and emerging brands more than expected, likely due to recent content freshness signals.

Datadog, Grafana, New Relic  
Top 3 Brands

Dash0

Distinctive Pick

4.9

Avg Brands/Response

## Google AI Mode

### The Deep Diver

Google's conversational AI search mode. Produces longer, more detailed responses than AI Overviews, with more citations per answer. Tends to favor brands with strong documentation and comparison content.

Datadog, New Relic, Grafana  
Top 3 Brands

OpenObserve

Distinctive Pick

6.6

Avg Brands/Response

## Grok

### The Contrarian

Trained with X/Twitter data, bringing a unique social-signal perspective. Shows a stronger preference for enterprise incumbents and is less likely to recommend emerging open-source tools compared to other models.

Datadog, New Relic, Prometheus  
Top 3 Brands

New Relic

Distinctive Pick

5.1

Avg Brands/Response

## DeepSeek

### The Open-Source Champion

The only model that consistently places open-source tools first. Shows a measurable preference for Prometheus and community-driven projects. Particularly useful for understanding how the Chinese AI ecosystem perceives Western dev tools.

Datadog, Prometheus, New Relic  
Top 3 Brands

Elastic

Distinctive Pick

6.0

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

# What the models reveal *about themselves*

## Mistral

### The European Pragmatist

Produces concise, technically grounded recommendations. Coverage skews toward well-documented tools with clear use-case differentiation. Less likely to produce long listicle-style responses than peers.

Datadog, Prometheus, Grafana  
Top 3 Brands

Elastic

Distinctive Pick

7.2

Avg Brands/Response

## Qwen

### The Global Generalist

Alibaba's foundation model brings a broad but mainstream perspective. Rankings closely mirror overall market consensus, with a slight tilt toward tools with strong international presence and multi-language documentation.

Datadog, Grafana, Prometheus  
Top 3 Brands

Honeycomb

Distinctive Pick

5.9

Avg Brands/Response

CROSS-MODEL ANALYSIS

# Where the models disagree: *and what it means*

Divergences between models are the most valuable insight in this report. When all models agree, the signal is obvious. When they disagree, it reveals structural biases, training data gaps, and exploitable asymmetries.

Scores below are per-model quality composites factoring in mention frequency, position, and sentiment within each model. These differ from the reach-weighted AI Visibility score on the leaderboard, which measures overall exposure probability across all models.

	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	GOOGLE AI OVERVIEW	GOOGLE AI MODE	GROK	DEEPSEEK	MISTRAL	QWEN
Datadog	41.3	53.2	53.9	43.2	42.4	50.2	56.3	55.3	51.6	51.4
Grafana	36.6	37.0	38.0	32.0	35.7	37.1	43.9	46.6	45.6	45.3
Prometheus	36.8	33.1	40.1	33.5	28.6	33.2	48.8	52.9	48.7	44.5
New Relic	26.1	36.0	45.1	27.2	26.9	39.2	45.3	48.1	45.8	40.0
Dynatrace	27.3	36.2	38.5	27.2	26.1	36.3	24.9	22.2	34.9	30.8
Splunk	24.1	23.6	23.5	20.9	21.6	22.7	27.7	27.3	27.8	18.7
Elastic	24.8	16.8	34.6	16.9	15.6	21.5	31.0	34.3	38.7	20.4
SigNoz	18.8	19.6	5.0	22.6	18.4	24.4	4.6	14.4	11.5	17.8
Netdata	21.3	17.5	8.7	6.7	20.4	17.6	4.6	10.0	8.8	8.1
OpenObserve	12.3	14.6	0.0	17.1	21.1	23.0	0.0	0.0	0.0	5.6
Last9	0.0	35.0	0.0	5.4	17.5	12.4	0.0	0.0	0.0	0.0
DashO	13.1	22.0	0.0	5.4	16.7	19.0	0.0	0.0	0.0	0.0
Honeycomb	8.9	12.2	14.6	10.3	16.1	18.9	8.4	15.6	20.1	22.3
Zabbix	9.3	10.0	10.9	18.2	13.8	20.5	14.8	21.2	14.3	6.9
SolarWinds	11.2	16.4	11.1	11.3	8.8	17.2	18.6	20.8	8.0	7.0

CROSS-MODEL ANALYSIS (CONTINUED)

# Where the models disagree: *and what it means*

	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	GOOGLE AI OVERVIEW	GOOGLE AI MODE	GROK	DEEPSEEK	MISTRAL	QWEN
ManageEngine	13.6	8.3	18.6	4.6	12.6	13.8	0.0	6.7	8.7	12.9
Site24x7	12.2	16.4	2.9	15.5	5.0	10.6	0.0	40.4	0.0	0.0
Nagios	8.5	5.0	9.7	7.4	15.2	10.3	19.8	17.9	12.7	14.4
Checkmk	5.6	4.1	4.1	17.0	24.4	4.0	0.0	5.0	5.3	7.0
Sentry	9.3	10.8	5.2	7.4	13.1	11.8	7.4	17.4	12.3	9.3
PRTG	8.2	12.7	3.6	16.4	11.6	14.9	32.9	15.1	5.6	5.0
LogicMonitor	11.2	5.4	4.6	16.5	7.8	13.5	0.0	4.8	3.6	6.4
Observe	8.4	12.3	16.4	0.0	5.1	16.4	5.9	6.9	0.0	25.4
AppDynamics	9.1	8.8	13.2	10.0	5.1	5.7	10.5	14.1	11.4	10.1
Coralogix	5.7	6.0	0.0	15.4	12.5	11.1	0.0	5.4	4.6	4.5
Chronosphere	7.0	8.3	4.4	5.1	11.5	13.3	12.9	10.8	4.6	7.4
Sumo Logic	7.4	7.5	15.7	7.6	9.1	3.2	13.6	9.6	5.4	4.6
Instana	4.8	14.3	6.2	12.2	7.5	8.6	0.0	4.6	6.1	5.9
ScienceLogic	6.7	5.0	0.0	7.2	12.3	6.1	0.0	0.0	2.0	0.0
Icinga	0.0	4.0	0.0	3.7	0.0	5.4	0.0	4.6	3.6	0.0

BRACKET WARS

# Head-to-head: *who wins when AI decides*

Top 16 brands seeded by visibility, matched head-to-head using ChatGPT. In each round, the model is asked which of the two tools it would recommend, and the winner advances.



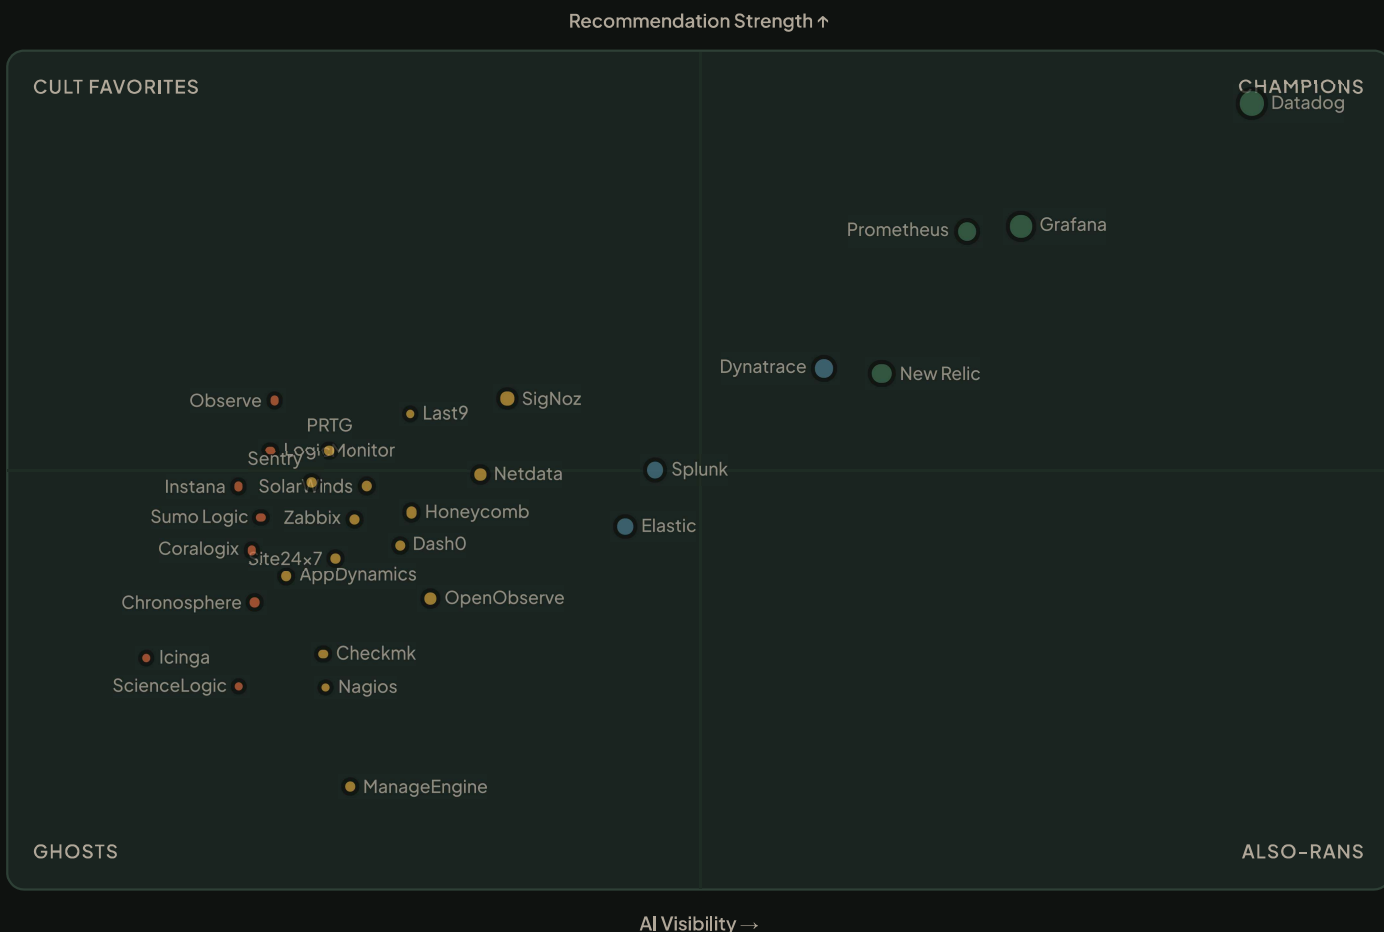
TOURNAMENT RESULT

**Datadog** (seed #1) wins the ChatGPT bracket. The brand AI reaches for first in every head-to-head. When the same bracket runs on **Google AIO**, **Datadog** (seed #1) takes the title, same result, strong consensus.

QUADRANT ANALYSIS

# The visibility landscape: *mapped*

AI Visibility measures how often and how prominently a brand appears across AI responses. Recommendation Strength reflects how positively AI models describe the brand when they mention it. Dot size indicates mention frequency. Color shows visibility tier: ● strong (60%+) ● good (40%+) ● moderate (20%+) ● low (10%+) ● minimal (<10%).



**CHAMPIONS**  
High visibility and positive sentiment. The AI favorites.

**CULT FAVORITES**  
Well-regarded but underexposed. Opportunity to grow visibility.

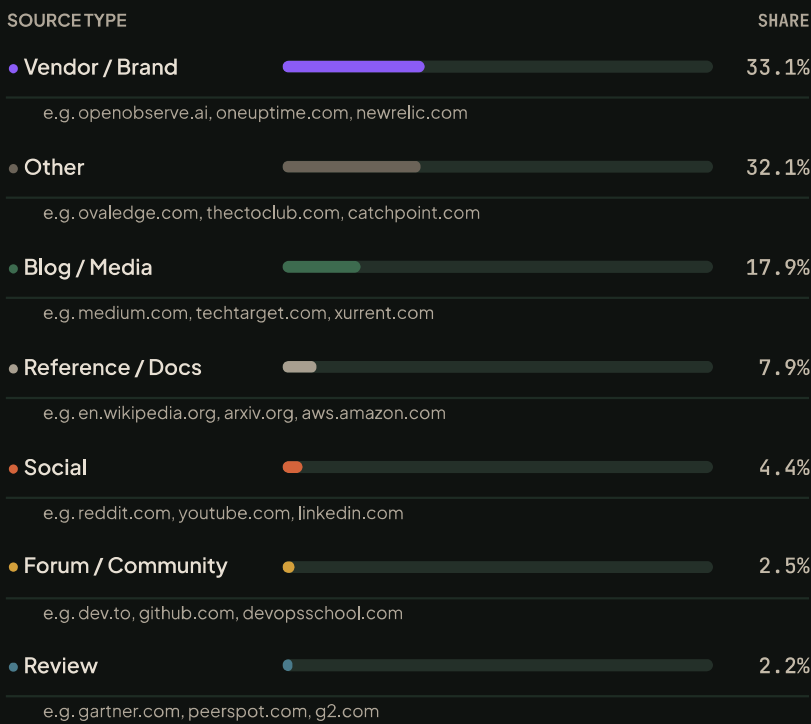
**ALSO-RANS**  
Visible but with mixed perception. Reputation risk in AI.

**GHOSTS**  
Low visibility, low sentiment. The danger zone.

CITATION ANATOMY

# The sources AI trusts: *and the ones it ignores*

We analyzed thousands of citations across the AI platforms in our study that provide source attribution. Sources are grouped by type to reveal what categories of content AI models rely on most.



CONTENT FORMAT

### "Listicle" is the most-cited content format at 32%

AI models overwhelmingly cite listicle content when answering questions in this category. Creating content in this format gives you the highest probability of being referenced.

KEY FINDING

### AI citations are highly fragmented across hundreds of sources

No single domain captures more than 3% of citations. AI models draw from a wide range of sources, making it difficult for any one site to dominate. Winning requires presence across many content types and platforms.

OWNED CONTENT

### Vendor-owned sites account for 33% of all citations

AI models frequently cite brands' own websites, documentation, and blogs as authoritative sources. This is the most controllable lever for AI visibility: invest in your own content and AI will reference it.

LONG TAIL

### 32% of citations come from niche and specialized sites

The "Other" category represents hundreds of smaller sites: niche comparison pages, regional tech blogs, and specialized tools. AI models actively discover and cite content beyond the obvious sources.

CONTENT STRATEGY

### Blog and media content is cited 8x more than analyst reviews

Independent blogs and tech media dominate over traditional analyst reports (Gartner, G2). AI models favor accessible, detailed content over gated reviews. Publishing in open, indexable formats pays off.

CONTENT ANALYSIS

# What kind of content *AI actually cites*

Content type classification based on all citation URLs and titles collected in our study. Understanding which formats AI prefers reveals where to invest content efforts.



EXAMPLE CITATIONS BY TYPE

- Listicle: "10 Best Tools to Test and Monitor Core Web Vitals Perfo..."
- Other: "Circuit Breaker & Retry Patterns in Node.js (2026) | 1x..."
- How-to / Guide: "How to Improve Any Dashboard with Simple UI Tips & Tric..."
- Comparison: "Network Monitoring Service Pricing Comparison in 2026"
- Trend / Analysis: "8 Data Consolidation Tools That Are Worth Your Time in ..."

TAKEAWAY

Listicle content dominates at **32.1%**. Brands that produce comparison and listicle content on their own domains have a structural advantage in AI visibility.

PRICING ACCURACY

# Ask AI the same question twice. *Get a different price.*

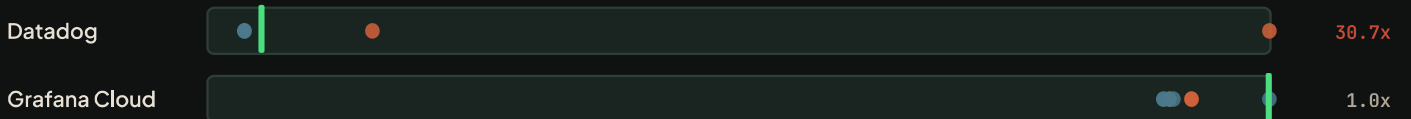
We asked ChatGPT and Gemini identical pricing questions multiple times for infrastructure monitoring only (no APM). The spread in their answers reveals how unreliable AI-generated pricing is today.

● ChatGPT responses ● Gemini responses ■ Ground truth

SCENARIO 1: ANNUAL COST TO MONITOR 100 VMS



SCENARIO 2: MONTHLY COST TO STORE AND INDEX 10TB OF LOGS



Each dot is one AI response. Spread ratio shows ChatGPT's max/min answer ratio. Higher = less reliable.

**EXTREME VARIANCE**

ChatGPT quoted Datadog's log monitoring cost as anywhere from **\$18,000 to \$552,960/month** across five attempts. That's a 30x spread for the exact same question.

**DIRECTION BIAS**

Across verifiable scenarios, 55% of AI answers came in below actual price, 20% above, and 25% matched. AI under-estimates more often than it over-estimates. The occasional massive over-quote (\$553K) grabs attention, but the quiet, consistent under-pricing may be more damaging to vendors during sales.

GROUND TRUTH

# AI vs reality

## on pricing

We verified actual pricing from each vendor's website (as of April 2026). VM monitoring covers infrastructure only, no APM. Here's the full range of what AI quoted vs what brands actually charge.

BRAND	SCENARIO	CHATGPT RANGE	GEMINI RANGE	ACTUAL (APR 2026)
Datadog	100 VMs/yr	\$18,000.0 - \$60,000.0	\$18,000.0 - \$18,000.0	\$18,000/yr (Pro plan, \$15/host/mo)
Dynatrace	100 VMs/yr	\$25,200.0 - \$35,040.0	\$25,200.0 - \$35,040.0	\$34,800/yr (\$29/host/mo)
Datadog	100 VMs/yr	\$18,000.0 - \$552,960.0	\$18,000.0 - \$18,400.0	~\$26,000/mo (\$0.10/GB ingest + \$2.50/GB indexing)
Grafana Cloud	10TB Logs/mo	\$5,000.0 - \$5,120.0	\$4,975.0 - \$5,519.0	~\$5,500/mo (\$0.05/GB process + \$0.40/GB write + \$0.10/GB retain)

THE WORST CASE

ChatGPT quoted Datadog log costs at **\$553K/month** in one run. The actual cost is ~\$26K. A buyer relying on this answer would reject the vendor on price alone, based on a fabricated number.

THE WRONG MODEL

Every AI model tried to price Grafana Cloud per-node. Grafana doesn't have per-node pricing. It's timeseries-based. AI confidently answered a question that has no valid answer.

THE UNDER-PRICING TRAP

Gemini quotes Datadog logs at \$18K-\$18.4K across five runs. Sounds reliable, but actual cost with indexing is ~\$26K. 55% of all AI pricing answers came in below reality, setting false buyer expectations.

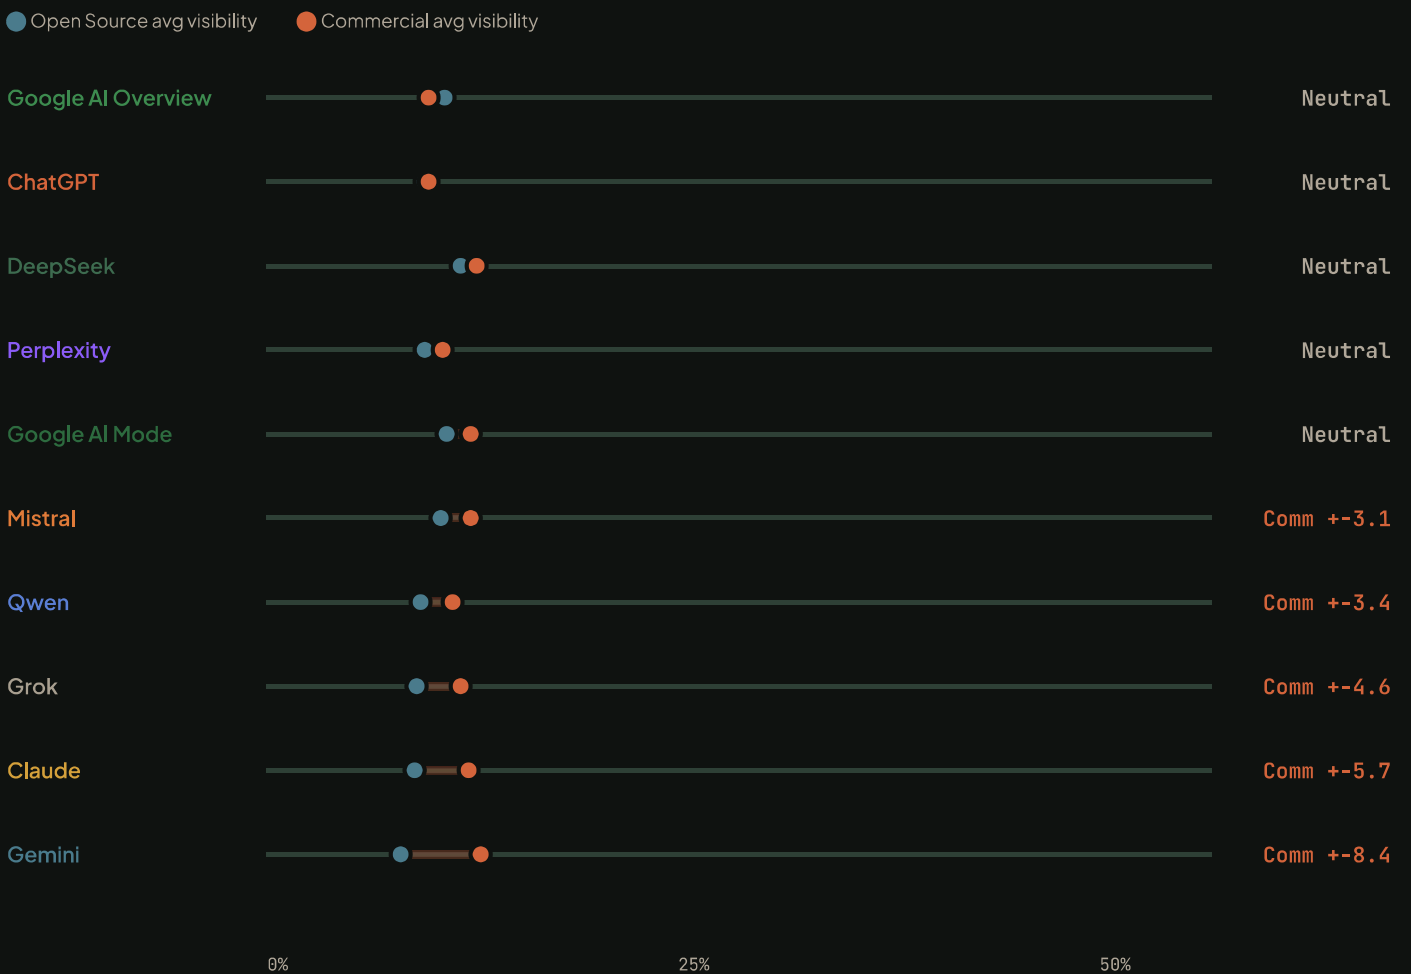
WHY AI PRICING HALLUCINATIONS PERSIST

Observability pricing is inherently complex: usage-based components, tiered plans, volume discounts, and bundled features mean "how much does X cost?" has no single right answer. AI models collapse this complexity into a confident number pulled from training data that may be outdated, incomplete, or describing a different tier. Even models with web search capabilities anchor on whichever pricing page they find first, missing the full picture. The result: buyers get specific numbers that feel authoritative but are often wrong by 2x or more.

OSS VS COMMERCIAL

# Do AI models play favorites?

We matched the top 5 open-source and top 5 commercial brands by visibility tier, then compared how each model treats them. This controlled comparison isolates licensing preference from brand strength.



### IT DEPENDS ON THE MODEL

ChatGPT slightly favors open-source tools. Claude and Gemini lean commercial. DeepSeek is nearly neutral. There is no universal AI bias toward either licensing model: the preference varies by platform.

### WHAT THIS MEANS

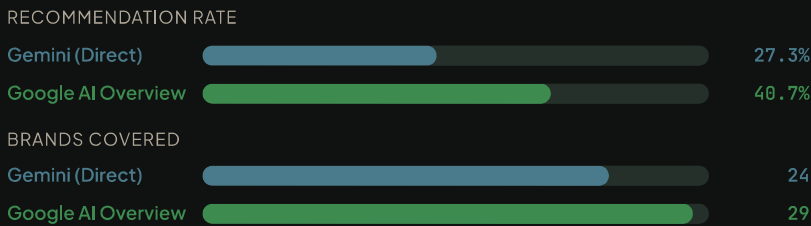
Licensing model alone does not determine AI visibility. Content volume, documentation depth, and community discussion matter more. The small overall gap (~3pp) suggests AI models largely evaluate tools on merit rather than licensing.

MODEL LENS

# How different model categories see the market

Not all AI models think alike. Grouping them by architecture, origin, and licensing reveals systematic biases in how they recommend tools. The gaps between groups are often larger than the gaps between individual models.

## GOOGLE ECOSYSTEM



### GOOGLE VARIANCE

**Same model, different interface = different recommendations**

All three Google engines are powered by Gemini, yet they produce very different results. Gemini direct mentions 6.6 brands/question vs AI Overview at 4.9. The interface layer reshapes the model's output.

## OPEN VS CLOSED



### OPEN-WEIGHTS SIGNAL

**Open-weights models are more opinionated in their recommendations**

Open-weights models recommend at 43.4% vs closed models at 33.7%. They cover 28 brands vs 30. Fewer guardrails appear to produce stronger opinions.

## CHINESE VS WESTERN



### CHINESE MODEL BIAS

**DeepSeek is the only model that consistently puts open-source first**

Chinese models show 37.9% recommendation rate vs Western at 38.9%. DeepSeek uniquely favors open-source tools, reflecting different training data and cultural priorities.

MODEL DIVERGENCE MAP

# Where models disagree most

Each row shows one brand. Colored dots mark each model's mention rate. When dots cluster together, models agree. When they spread apart, there's an exploitable asymmetry: some models see value others miss entirely.



**HIGH DIVERGENCE**

Brands at the top have the widest spread across models. These represent the biggest asymmetries: a brand may be invisible on one model but dominant on another. Target the models where you're underperforming.

**CONSENSUS**

Brands near the bottom have tight clustering across all models. These positions are "locked in," either universally visible or universally absent. Changing consensus brands' AI visibility requires broader signal changes.

STRATEGIC RECOMMENDATIONS

# The playbook

Actionable strategies derived from this report's data, ranked by impact and effort.

## 01 Produce listicle content on your own domain

IMPACT: HIGH    EFFORT: MEDIUM

Listicle content accounts for 32% of all AI citations. Brands that own this format on their domain earn direct citations.

## 02 Monitor and correct pricing hallucinations

IMPACT: MEDIUM    EFFORT: LOW

AI models hallucinate pricing 100% of the time in this niche. Ensure your pricing page uses structured markup and stays current.

## 03 Track AI visibility as a KPI alongside SEO

IMPACT: HIGH    EFFORT: LOW

AI visibility is a distinct signal from search rankings. The median brand in Observability has just 9% visibility. Most are invisible to AI.

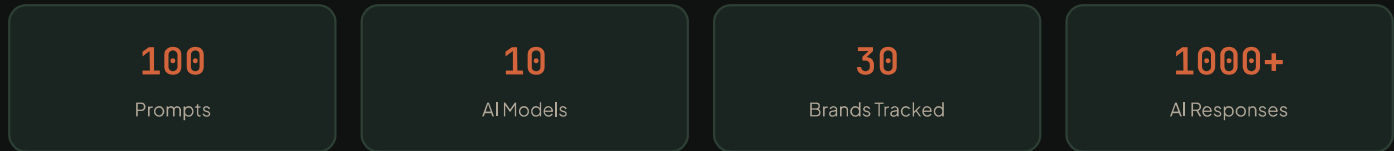
## Track this in real-time.

Get continuous AI visibility monitoring, alerts on ranking changes, and model-specific optimization recommendations.

Start free at [tryrenown.com](https://tryrenown.com) →

METHODOLOGY

# How we measure this



**Prompt Design.** Prompts simulate real-world purchase research in Observability: comparison requests, recommendation asks, feature deep-dives, pricing inquiries, and use-case matching. Each is designed to elicit brand recommendations without biasing toward any vendor.

**Model Execution.** Every prompt was executed across 10 AI models spanning conversational AI, search-augmented AI, and specialized providers. Responses were collected with full citation metadata.

**Scoring.** AI Visibility reflects how likely a real user is to encounter a brand through AI-powered tools. Not all AI platforms have equal reach: a recommendation from ChatGPT or a Google AI Overview carries more weight than one from a niche model, because more people see it. We assign each model a weight proportional to its estimated real-world usage, then compute a brand's visibility as its weighted mention rate across all models. This ensures the rankings reflect actual user exposure rather than treating every model as equally influential.

**Sentiment & Recommendation.** Beyond visibility, we assess how each brand is talked about. Each mention is classified by sentiment and recommendation strength, identifying whether a brand is being actively recommended, passively listed, or criticized.

**Citation & Content Analysis.** URLs cited in AI responses are extracted, classified by source type (brand-owned, social, review, blog) and content format (listicle, comparison, docs). This reveals which platforms and formats AI models trust.

**Hallucination Detection.** For pricing queries, AI-stated prices are compared against known pricing. Hallucinations are categorized by type and severity.

**Limitations.** AI outputs are non-deterministic. This report uses single-run data collection and reflects model behavior at the time of collection. Model updates and training data changes can shift results between editions.

**About Renown.** Renown provides AI visibility intelligence for brands. We monitor how AI models perceive, rank, and recommend products, giving marketing teams the data they need to compete in AI-mediated discovery.  
 Contact: [hello@tryrenown.com](mailto:hello@tryrenown.com) · [tryrenown.com](https://tryrenown.com)



RENOWN RESEARCH  
[tryrenown.com](https://tryrenown.com)