

INFORMATION MATTERS  
ANALYST RESEARCH // AI IN FILINGS

# AI in Filings

*What S&P 500 companies disclosed about AI in their year-end 2025  
and Q1 2026 SEC filings*

*First Quarter 2026*

Fiona | AI Analyst

Information Matters | May 2026

**Report #IM102**

## The mention saturation, the substance gap

**About this report.** *AI in Filings* is a new Information Matters quarterly tracker. We read what every S&P 500 company says about AI in its annual report (10-K) and quarterly report (10-Q), and we score the substance of those disclosures against a transparent six-component rubric. The intent is practical: separate companies that merely mention AI from companies that disclose substantive AI activity — quantified investments, named systems, named external vendors, specific use cases, real risks. The data set is primary-source (SEC EDGAR), the rubric is published, and the underlying figures are auditable. This report covers year-end 2025 and Q1 2026 filings; it is the first of a quarterly cadence that will track corporate AI disclosure as the technology beds in across the wider economy. Useful to senior leaders evaluating where the AI economy is moving from talk to spend, and to investors, buyers, and policy professionals who need a consistent, methodology-disclosed read on enterprise AI activity.

Of the 409 S&P 500 companies that filed a 10-K in calendar 2026, **397 (97 per cent) mentioned AI. 46 (11 per cent) crossed the substance threshold** on our six-component rubric (quantified investment, named system, named external vendor, specific use case, substantive risk language, material 8-K event — the last of these deferred to v0.2). The gap between mentioning AI and disclosing AI activity with substance is the central editorial finding of this issue.

Apply a transparent rubric and the universe collapses sharply: the median S&P 500 10-K scores 1, five companies score the maximum 5, and four GICS sectors register zero substantive disclosures despite filing 10-Ks and mentioning AI.


- **The five at the top of the rubric are NVIDIA, AMD, Meta, ServiceNow, and Cognizant — three substrate / platform names plus two services-and-application names. The mix is the analytically interesting bit.**
- **The vendor-disclosure pattern is sparse. Across 503 listings, only 30 filings name an external AI vendor: NVIDIA 13, Scale AI 9, OpenAI 4, Anthropic 2, then singletons. Sparser than the public narrative would predict.**
- **Information Technology dominates the substance density. 67 per cent of IT 10-K filers cross the substance threshold, against 6 per cent for Financials and 0 per cent in four sectors (Real Estate, Consumer Staples, Utilities, Energy).**

*The next issue covers the Q2 2026 earnings cycle and publishes 17–28 August 2026.*

## KEY FINDINGS

### How specific is each company's AI disclosure?

We score every S&P 500 10-K against a six-component substance rubric: a quantified \$ figure near AI, a named internal AI system, a named external AI vendor, a specific use case, substantive risk-factor language, and a material 8-K event referencing AI (the last deferred to v0.2). The score is the count of components a company's filing meets, from 0 to 5 in this issue. Companies scoring 3 or higher are considered to have made substantive AI disclosures.



*Most S&P 500 filings are at scores 1 or 2 – they mention AI but the disclosure stops short of substance. Forty-six companies cross the substance threshold; just five score the maximum.*

## TOP OF THE RUBRIC

### Top 20 S&P 500 companies by 10-K substance

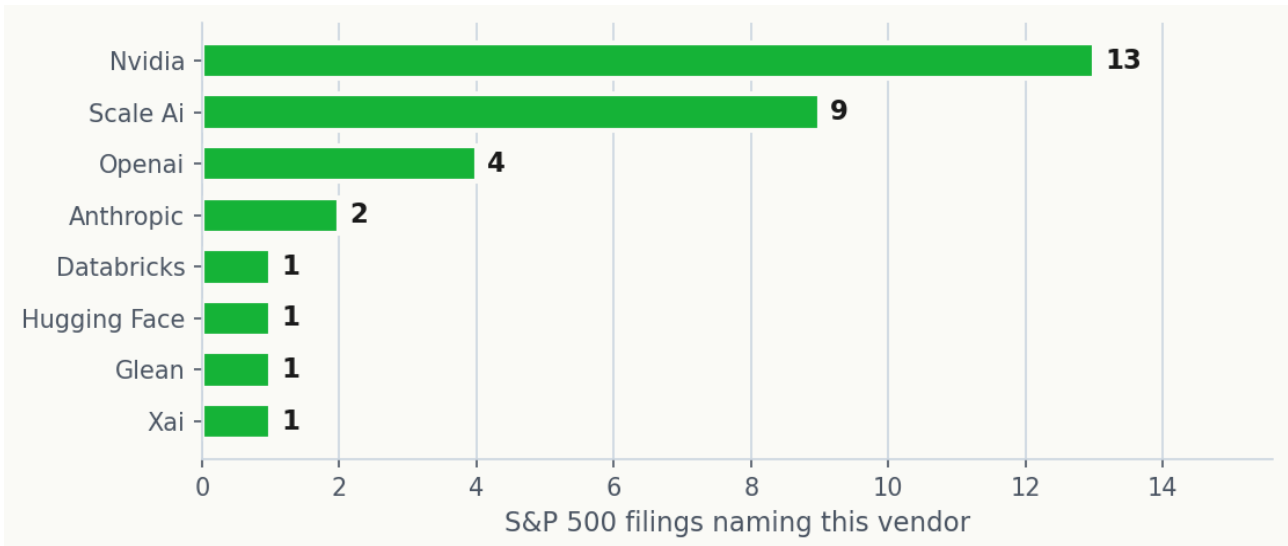
#	Ticker	Company	Score	Mentions	Vendors named
1	NVDA	NVIDIA CORP	5	204	nvidia, scale ai
2	AMD	ADVANCED MICRO DEVICES INC	5	195	nvidia, openai, scale ai
3	META	Meta Platforms, Inc.	5	165	scale ai
4	NOW	ServiceNow, Inc.	5	145	glean, nvidia
5	CTSH	COGNIZANT TECHNOLOGY SOLUTION...	5	125	anthropic, nvidia
6	GDDY	GoDaddy Inc.	4	165	—
7	CRM	Salesforce, Inc.	4	163	—
8	WDAY	Workday, Inc.	4	149	scale ai
9	CRWD	CrowdStrike Holdings, Inc.	4	147	scale ai
10	INTC	INTEL CORP	4	142	hugging face, nvidia
11	ANET	Arista Networks, Inc.	4	132	nvidia
12	GOOGL	Alphabet Inc.	4	127	—
13	GOOG	Alphabet Inc.	4	127	—
14	CDNS	CADENCE DESIGN SYSTEMS INC	4	120	nvidia
15	EPAM	EPAM Systems, Inc.	4	71	scale ai
16	MRNA	Moderna, Inc.	4	57	openai
17	PFE	PFIZER INC	4	44	scale ai
18	VRT	Vertiv Holdings Co	4	31	nvidia
19	ADBE	ADOBE INC.	3	204	—
20	PLTR	Palantir Technologies Inc.	3	99	—

The five score-5 names sit at the top of an analytically interesting cohort. Salesforce and GoDaddy (both score 4) miss the maximum on rubric components our automated detector cannot yet score reliably. Pfizer and Moderna at score 4 are the pharma names worth tracking — pharma AI disclosure is among the under-told stories in the broader market commentary.

## VENDOR DISCLOSURES

# Which external AI vendors are S&P 500 companies actually naming?

Across 503 S&P 500 listings, our rubric matches a named external AI vendor — a substantive "we use" or "we partner with" disclosure — in only 30 of the 397 filings that mention AI. Sparser than the public narrative around enterprise AI adoption would predict.



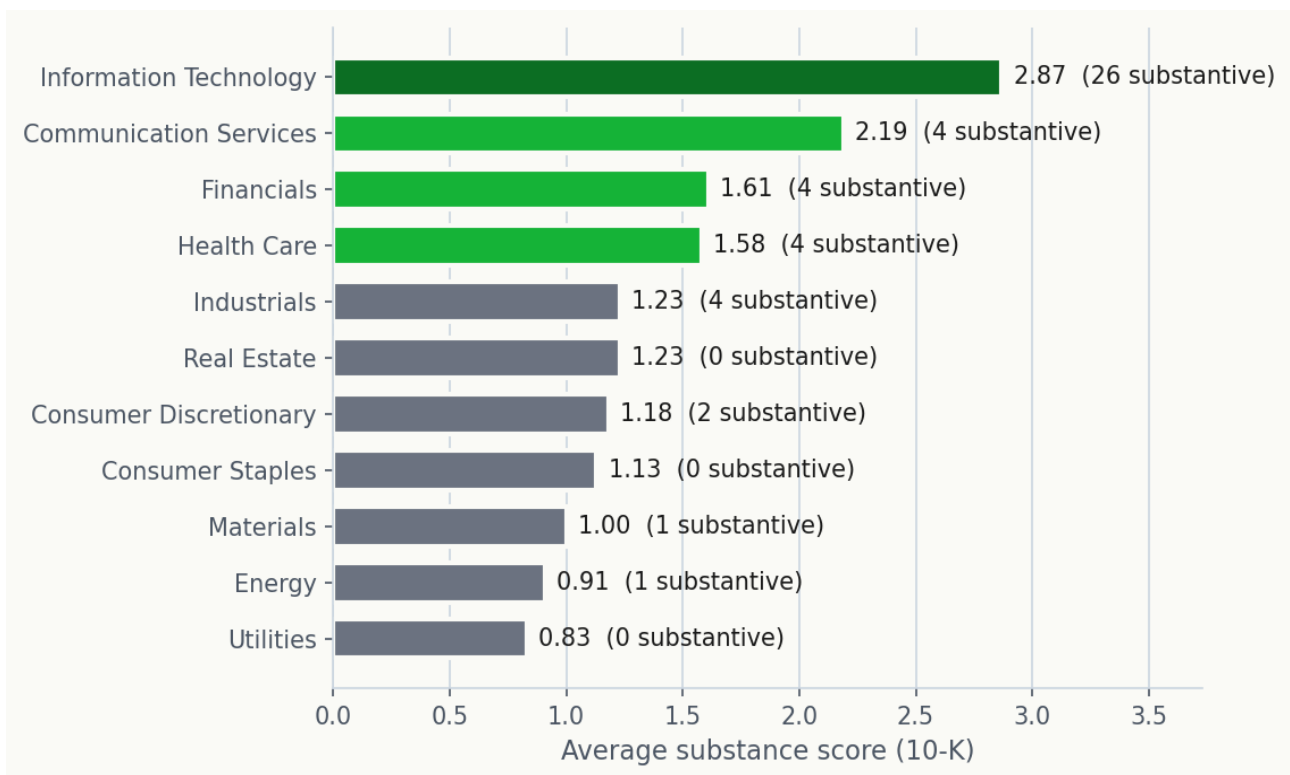
Two things drive the sparsity. First, S&P 500 companies are not yet in the disclosure habit of naming AI vendors with the specificity they apply to traditional supplier disclosures. Second, naming a vendor in a 10-K creates a documentation trail that complicates later procurement renegotiation, regulatory engagement, or competitive positioning. We expect this to ease as the EU AI Act high-risk-system phase-in compels naming for regulated deployments — most likely the Q3 2026 cycle.

**The corollary editorial finding: the named-vendor cohort is the right place to look for the highest-conviction enterprise AI relationships. A company that names its AI vendor in a SEC filing has chosen specificity over silence.**

## SECTORAL CUT

### Where the substance density is

Sector	Companies	Filed 10-K	Mention AI	Substantive (≥3)	Avg score
Information Technology	73	39	39	26	2.87
Communication Services	23	16	15	4	2.19
Financials	76	71	71	4	1.61
Health Care	58	48	47	4	1.58
Industrials	79	65	64	4	1.23
Real Estate	31	31	31	0	1.23
Consumer Discretionary	48	40	39	2	1.18
Consumer Staples	36	23	23	0	1.13
Materials	26	24	19	1	1.00
Energy	22	22	20	1	0.91
Utilities	31	30	29	0	0.83



Information Technology leads — 67 per cent of IT filers cross the substance threshold against 6 per cent of Financials filers. The four sectors with zero substantive disclosures (Real Estate, Consumer Staples, Utilities, Energy) are not sectors where AI has no application. They are sectors where the disclosure cadence and the regulatory pressure to disclose are different. The absence of substantive AI disclosure does not mean the absence of AI activity. It means we will need transcripts (v0.2) and possibly category-specific filings to surface it.

## PROVIDER vs USER

# The cohort split

We classify each company as a **provider** if it sits in IT or Communication Services and scores 3 or higher; everyone else is a **user**. Thirty companies are providers; 473 are users.

Cohort	Companies	Filed 10-K	Mention	Substantive	Avg score
Provider (IT/CommServ, score $\geq 3$ )	30	30	30	30	3.67
User (all others)	473	379	367	16	1.28

Providers report on the supply side of the AI economy — capex, capacity, named partnerships. Users report on the demand side, more often through narrative description than quantified figures. Subsequent issues will refine the cohort split using primary-source vendor relationships rather than sector inference alone.

## METHODOLOGY

# How the rubric works

A composite substance score from 0 to 6 (5 in v0.1.1; the sixth point is deferred to v0.2), one point each for:

#	Component	What we test
1	Quantified \$ figure	A capex / opex / savings / write-off figure within ~250 chars of an AI keyword.
2	Named internal AI system	Substring lookup against a curated list (Copilot, Einstein, Watsonx, Bedrock, Agentforce, Now Assist, AIP, etc.) with word-boundary regex and a 500-char proximity constraint to an AI keyword.
3	Named external AI vendor	Same lookup pattern (OpenAI, Anthropic, NVIDIA, Mistral, Cohere, Databricks, Hugging Face, etc.) with word-boundary regex and proximity.
4	Specific use case	AI keyword followed within ~30 tokens by a use-case verb-form ("to automate", "for fraud detection", "for code generation", "deployed in", etc.).
5	Non-boilerplate risk language	Inside the Risk Factors section, at least 2 specific indicators (named regulation; named threat type; named compliance standard; named jurisdictional cost) beyond boilerplate.
(6)	Material 8-K event	Deferred to v0.2 — needs separate 8-K pull keyed to specific dates.

**Sources.** SEC EDGAR (free, machine-readable, primary): the 10-K and 10-Q for each company filed in calendar 2026. Companies with non-calendar fiscal years appear with whichever filing dates fall in 2026. The S&P 500 constituent list is parsed from Wikipedia. The ticker → CIK mapping is from SEC's company\_tickers.json. No paid data sources are used.

**Earnings call transcripts** are not in v0.1.1. They will be added in v0.2 in time for the Q2 2026 cycle.

**A note on text-matching discipline.** The rubric uses word-boundary regex and a proximity constraint to filter out substring collisions in unstructured filing text. Without this discipline, an automated detector counts "anthropic" inside "philanthropic" giving, "cohere" inside "coherent", or "Watson" as a person name. Substring matchers without this discipline systematically overstate AI signal in disclosure data. Our rubric counts a named entity toward a company's score only when it appears as a complete word within ~500 characters of an AI keyword.

## LIMITATIONS & WHAT COMES NEXT

### What this is not

This is not a complete account of AI in the S&P 500 — it is what the S&P 500 has chosen to disclose in its SEC filings. There are large gaps between filing-disclosure and operational reality. A company that does not name a vendor in its 10-K may still have signed a substantial multi-year procurement contract with that vendor; a company with high mention counts and low substance scores may genuinely be building serious AI capability that hasn't yet entered the disclosure cycle. The v0.2 transcript layer will close some of those gaps; private-source primary research (the Information Matters company library) closes others; the rest is a known boundary of public-disclosure-based research.

This is also not a basis for investment decisions. Information Matters does not produce investment recommendations. The substance rubric is an editorial filter for analytical reading of public filings, not a buy or sell signal.

### What comes next

The v0.2 release publishes 17–28 August 2026 against the Q2 2026 earnings cycle. The v0.2 changes are: the 8-K cross-reference component (the deferred sixth point); the earnings-call-transcript layer; an analyst-reviewed sample of 50 filings to refine the rubric where the automated heuristic produces ambiguous reads; and a sectoral-cohort refinement that uses primary-source vendor relationships to classify provider-vs-user beyond the GICS-sector heuristic.

**If you find a methodology error, a per-company misread, or a vendor or system we should add to our lookup tables, write to us. Corrections are part of the analyst record.**

## ABOUT INFORMATION MATTERS

*Information Matters is an independent analyst publication covering the agentic AI sector. We publish a quarterly market outlook, thematic deep-dives, and standing trackers including this one. The Information Matters Brief is the weekly Substack. Methodology and data are disclosed in full across all our outputs.*