Not for redistribution without written consent of Morgan Stanley

Morgan Stanley | RESEARCH

FOUNDATION

October 26, 2026 04:01 AM GMT

Software | North America

Token Thoughts – How Will Rising Open-Weight Model Adoption Impact Software?

Open-weight model adoption has surged over the past year creating an opportunity for software incumbents to transform business models and unlock better unit economics. However, the benefits are unlikely to be universal and value likely accrues to those with the most defensible moats.

Key Takeaways

- Open-weight model usage is surging with share jumping from 20%+ of all tokens processed a year ago to as high as 70%+ today
- While non-western models make up the majority usage today, over the next year we expect greater choice of US-based models helping fuel enterprise adoption
- Overall we see open-weight models as an opportunity for software but the biggest beneficiaries will be those with strong moats and who control complex workflows
- In horizontal and vertical SaaS we see benefits accruing to BOX, CRM, MANH, NOW, SHOP, TOST, and TEAM
- In Infra and cyber we see tailwinds for CRWD, DDOG, MSFT, NET, OKTA, PANW, PLTR and SNOW

What Does Rising Adoption of Open-weight Models Mean for Software? We see the investor debate on software turning more positive in recent months given the rise of supportive themes including: 1) consistent and material declines in the cost of intelligence, 2) a proliferation in model choice, 3) the potential for a steady versus explosive advancement in AI capabilities with use case expansion outside of coding and 4) the lack of evidence pointing to major head count displacement, particularly among developers where hiring appears to be ongoing. This edition of *Token Thoughts* expounds on the recent report by our colleague Stephen Byrd (see: **Open-Weights Models & 3 States of the World**) on the topic of open-weight models and how rising adoption of open-weight models will impact the software industry going forward. Broadly, we see the rise of open-weight adoption as constructive for the industry given the amount of AI software that can be developed and consumed given the potential lower costs. However, rising adoption by itself does not make this universally positive for incumbent software economics.

Open-weight Adoption Is Positive But the Benefits Are Not Likely to Be Universal. The promise for open-weight model adoption is not simply that models become cheaper – it is that the cost of producing useful intelligence falls while the

MORGAN STANLEY & CO. LLC

Sanjit K Singh
Equity Analyst
Sanjit.Singh@morganstanley.com

Adam Wood
Equity Analyst
Adam.Wood@morganstanley.com

Meta A Marshall
Equity Analyst
Meta.Marshall@morganstanley.com
+1 212 761-0430

Elizabeth Porter, CFA
Equity Analyst
Elizabeth.Porter@morganstanley.com

Chris Quintero
Equity Analyst
Chris.Quintero@morganstanley.com
+1 212 761-1686

Matthew Glickman
Research Associate
Matthew.Glickman@morganstanley.com

SOFTWARE
North America
Industry View **Attractive**

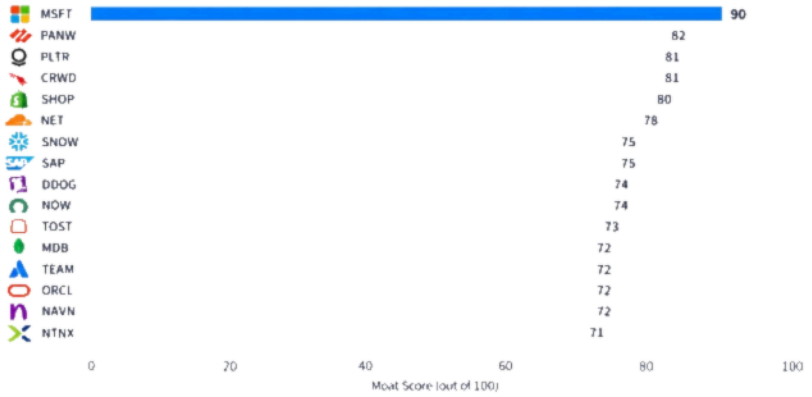
Morgan Stanley does and seeks to do business with companies covered in Morgan Stanley Research. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of Morgan Stanley Research. Investors should consider Morgan Stanley Research as only a single factor in making their investment decision.

For analyst certification and other important disclosures, refer to the Disclosure Section, located at the end of this report.

volume of intelligence consumed can increase dramatically. For software incumbents, lower-cost intelligence creates two distinct opportunities: 1) expanding the range of economically viable workflows and 2) increasing the value of assets the vendor uniquely controls. The first can broaden adoption and usage, but is available to competitors as well. The second is more differentiated: cheaper intelligence can make enterprise context, trusted data, telemetry, secure execution, transaction processing and governance more useful to customers. Where these capabilities remain difficult to replicate, vendors should be better positioned to maintain pricing —or raise it as customer value increases—even as delivery costs fall. The biggest beneficiaries should therefore combine a broader monetizeable opportunity with the pricing power to retain a meaningful share of the savings, rather than see those benefits competed away through lower prices.

This perspective aligns with our Moat & Journey ranking that we highlighted in our Insight report – **The Moat & The Journey – Assessing Software Durability + Growth in the Age of AI**. Among the software companies having strongest combination of defensible moat while also best-positioned to benefit from secular AI tailwinds, we see Atlassian, Cloudflare, Crowdstrike Datadog, Microsoft, MongoDB, Nutanix, Palantir, Palo Alto Networks, SAP and ServiceNow as the strongest beneficiaries from rising open-weight model adoption.

Exhibit 1: Software Companies with Strongest Combination of Moat & Journey Scores (Top Quintile)

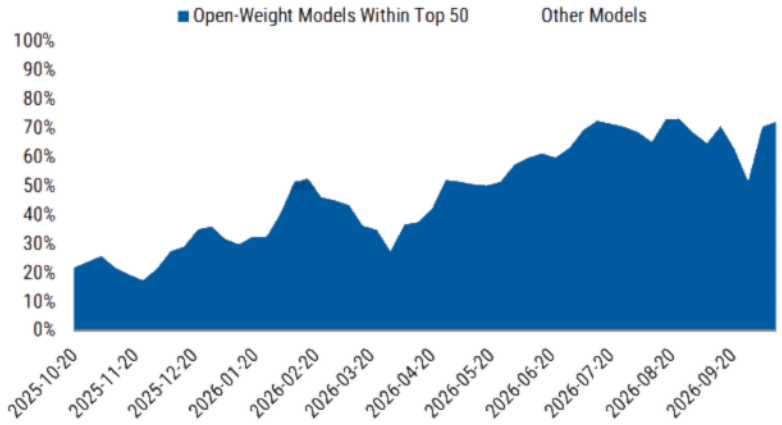


Source: Morgan Stanley Research

Evidence Points to a Surge in Open-weight Adoption Over the Past Year. OpenRouter is a third-party API marketplace and routing layer that gives developers access to a broad range of open and closed-weight models. Across the 50 most popular models on the platform (~90% of all tokens tracked), open-weight models have steadily gained share of tokens processed over the past year with share jumping from ~22% of all tokens processed on OpenRouter a year ago to ~72% today (see Exhibit 1). Data from application/agent platform provider Vercel, shows a similar trend with open-weight model usage accounting for 63% of token usage – up from just 7% in December 2025. This suggests that developers are increasingly attracted to lower-cost per token alternatives for a broader set of routine workflows and tasks. We note that these data should not be viewed as a comprehensive measure of industry-wide model share. OpenRouter's customer base

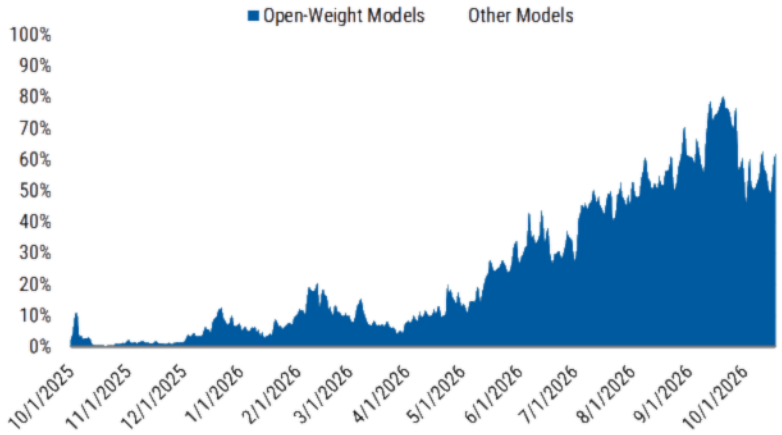
likely skews towards cost-conscious, sophisticated developers and smaller workloads, which may bias usage towards open-weight models.

Exhibit 2: Open-Weight Models are Gaining Share of OpenRouter Tokens



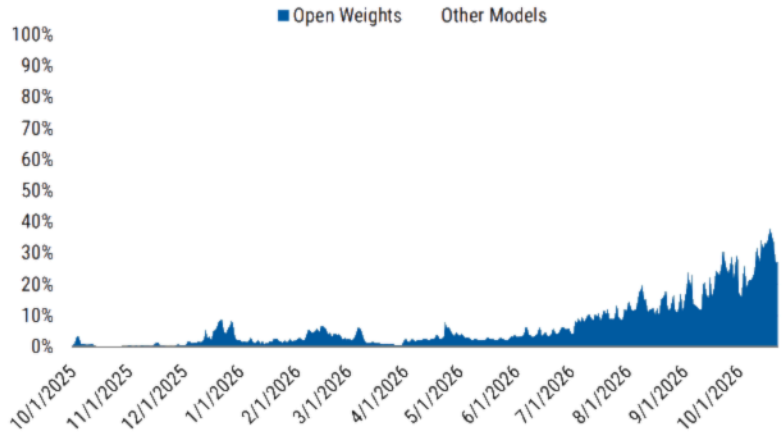
Source: OpenRouter, Morgan Stanley Research

Exhibit 3: Open-Weight Models are Gaining Share of AI Gateway Tokens on Vercel's Platform



Source: Vercel AI Gateway, Morgan Stanley Research

Exhibit 4: Open-Weight Models are Gaining Share of AI Gateway Spend But Overall Share is Somewhat Modest



Source: Vercel AI Gateway, Morgan Stanley Research

Key Debates

Debate #1: Will Enterprises Adopt Open-Weight Models Despite Data Sovereignty Risks?

- *Bull Case:* Western open-weight models (Nemotron 4, Muse Spark, Mistral Large 4, Replication) could come close to closing the capability gap with leading international alternatives over the coming months. This should help unlock adoption among enterprises who have been reluctant to incorporate non-western models into their environments. As competitive options emerge, enterprises stand to benefit from greater deployment flexibility, post-training opportunities and lower inference costs that open-weight models provide.
- *Bear Case:* Enterprises, particularly in Europe and North America, may remain unwilling to adopt non-western open-weight models due to geopolitical, reputation or security concerns
- *Our View:* We expect enterprise adoption to broaden as new Western models expand capabilities in the coming months/years and hyperscalers make secure deployment more accessible. The key distinction is between where a model is developed and how it is hosted, governed and monitored. Even for non-western models, there is the potential for Hyperscaler platforms to provide regional hosting, private connectivity, encryption, granular access controls, sensitive-information filtering and ongoing monitoring to help address concerns of enterprise customers.

Debate #2: Can Open-Weight Models Unlock Better AI Unit Economics?

- *Bull Case:* Open-weight models can provide software incumbents several advantages including: 1) less dependence on expensive proprietary model APIs, giving vendors greater negotiating leverage, 2) the ability to fine-tune models for particular workflows and route straightforward tasks to cheaper models and 3) greater control over deployment environments, data

residency, latency, and inference optimization.

- *Bear Case:* The key tenet of the bear case is that these same advantage described above are also available to AI-native challengers, hyperscalers, and enterprise customers building internally. This could translate to lower prices, falling customer switching costs, and a shift from per-seat pricing to per-task or per-outcome monetization. A related issue is who captures the value of lower inference costs – the software provider or the end customer. Historically, most software has not followed a cost-plus pricing model meaning that any efficiencies gained in delivering the service was retained by the software provider. Should customers demand that lower costs of inference be passed on (under the threat of switching to competitive alternatives or building their own solutions), software gross margins could be pressured.
- *Our View:* In aggregate, we think broader open-weight model adoption should provide stability for gross margins. That said, we note that some categories such as analytical data infrastructure has shown a tendency to pass on at least some of the portion of efficiencies enabled by technology advancement to end customers. While this can represent a temporary headwind to revenue/gross profit dollars, this is typically outweighed by the unlock in usage and workflows over time. Furthermore, companies more exposed to competition due to lack of proprietary data and enterprise context and a reliance on simple workflows or being viewed as a point solution are unlikely to see gross margin tailwinds as a result of open-weight model adoption.

Debate #3: Can an Open-Weight Strategy Help Win New Customers and Expand User Pools?

- *Bull Case:* A key element of the growth strategy for many software companies is to accelerate customer acquisition and widen the audience by reaching new types of users. However, with customers often demanding that vendors prove out value first before signing a deal, software vendors often have to engage in lengthy proof-of-concepts or offer free trials to spur interest in their AI offerings which can often be expensive. Greater use of open-weight models is likely to make usage-based and outcome-based AI products more economically viable expanding software's addressable market into budgets historically allocated to human labor. In conjunction these forces could materially improve gross margins for AI products, particularly vendors that embed substantial inference within fixed-price subscriptions. Furthermore, open-weight models can also make product-led growth more economical allowing companies to grow their customer bases and pursue new users.

- *Bear Case:* Similar to the debate on unit economics, the same open-weight strategies can be pursued by competitors particularly well-funded AI startups who can execute aggressive customer acquisition plans without having the need to generate profitability in the near-term.
- *Our View:* Lower inference costs enabled by open-weight models should

make product-led-growth (PLG) strategies more viable compared to those using proprietary models alone. With the same open-weight models available to the competition (particularly AI startups), the biggest returns to software will be with those companies having strong moats that arise from attributes such as proprietary data, deep knowledge and context of their customer workflows and processes, regulatory benefits and deep vertical/ domain expertise.

Exhibit 5: Assessment of Open-weight Model Adoption Across Software Markets

Software Category	Summarized View	Key Beneficiaries
Cybersecurity	More AI agents expand the attack surface and drive security demand. Cheaper open-weight AI intensifies automated SOC competition but unlocks new monetization opportunities for data rich cyber vendors	CRWD, PANW, OKTA
Infrastructure software / SaaS	Cheaper AI drives more applications and higher demand for data infrastructure, observability, orchestration and governance, with more agents generating data, queries and monitoring needs	MSFT, NET, DDOG, DT, ESTC, MDB, NTNX, PLTR, SNOW
Horizontal application SaaS	Lower inference costs can boost margins, accelerate agent adoption, enable outcome/usage-based pricing, and help SaaS incumbents become the system of action, especially those with proprietary data and deep workflow expertise	TEAM, CRM, NOW, WIX, ZM, FIG, BOX
Vertical SaaS	Lower inference costs make it economical to automate manual, industry-specific workflows, with incumbents best positioned when they have proprietary data, domain expertise and strong integrations	SHOP, TOST, MANH, NAVN, TTAN

Source: Morgan Stanley Research

Implications for Infrastructure Software

We see open-weight model adoption as generally positive, particularly for mission-critical data infrastructure and observability. Cheaper intelligence makes more AI applications economically feasible, and those applications still require persistent data, retrieval, monitoring, security, execution infrastructure and operational governance. For infrastructure incumbents, this creates several opportunities.

- Core data platforms (Microsoft, Databricks, MongoDB, Oracle, Palantir, Snowflake): provide AI agents real-time, permissioned enterprise context, database and data-platform vendors should see higher query volume, expanded use cases and growing demand for governed data access.
- Observability (Datadog, Dynatrace, Elastic): Agents create new traces, tool invocations, execution paths, errors, latency issues, token costs and quality failures that expand demand for monitoring, remediation and policy enforcement.
- Model gateways and orchestration (Cloudflare, Palantir, Vercel): Enterprises using growing proprietary and open models will need routing, evaluation, governance, budgeting and performance monitoring.
- Search and retrieval (Elastic, MongoDB, Palantir) : Proprietary company knowledge becomes more valuable when lower-cost models can query and act upon it.

Key Beneficiaries: Microsoft, Cloudflare, Datadog, Dynatrace, Elastic, MongoDB, Nutanix, Palantir, Snowflake

- Key risks: Include the potential for open-weight models to be coupled with open-source infrastructure, allowing customers to build cheaper alternatives

to proprietary monitoring, data processing and orchestration products. In addition, agents have the potential for making existing infrastructure more efficient. For example, if automated optimization reduces wasted queries, duplicate data processing, unnecessary telemetry or manual operational activity, usage-based vendors might see lower consumption for a given workload.

Implications For Cybersecurity

Open-weight model adoption is structurally positive for cybersecurity as greater model availability accelerates the number of AI agents operating within enterprises. With each agent requiring access to enterprise applications, private data, cloud infrastructure, identity credentials, source code repositories and external services, the attack surface area likely expands resulting in newer threats including:

- Model supply-chain integrity and untrusted model artifacts.
- Prompt injection and malicious tool outputs.
- Excessive agent permissions and machine-identity proliferation.
- Autonomous execution of unauthorized actions.
- Data exfiltration across models, tools and workflows.
- Continuous auditing, behavioral monitoring and policy enforcement.

The need to develop guardrails with the open-weight models likely creates a much larger opportunity for cyber vendors that instead have to secure a much more complex hybrid world.

Beyond the notion that agents create a more risky environment, open-weight model adoption is likely to play a role in reducing the cost of processing telemetry, investigating alerts, classifying threats and automating responses. We see this being most challenging to automated SOC solutions, as the cost to provide ongoing insights gets reduced with open weight model adoption (with Hugging Face incident notably using open-weight models for resolution). While this creates more competition for Automated SOC solutions, as we noted coming out of Black Hat this year, it also gives cyber vendors another way to monetize their data solutions (like CRWD's SafeMind).

Key Beneficiaries: Crowdstrike, Palo Alto Networks, Okta

- Key Risks: Include the potential for customers or competitors to use open-weight models to automate capabilities that replace some functionality currently sold as premium capabilities, such as correlating alerts, writing detection rules, analyzing logs, and prioritizing potential threats.

Implications For Horizontal SaaS

Overall, we see the rise of open-weight models as an opportunity for hortizontal SaaS vendors. The bullish scenario is that open-weight models provide a path to lower inference costs which can drive gross margins and stimulate incremental revenue. Specifically, growing utilization of open-weight models has the potential to

lower AI delivery costs. foster adoption of agentic solutions, transition pricing models from seat based to outcome or usage-based pricing and help position SaaS incumbents to emerge as the system of action. We see the biggest beneficiaries in SaaS are those with proprietary data, deep knowledge of complex workflows and core strengths in integrations, identity, permissions, and enterprise distribution. Near-term, companies may reinvest profitability upside driven by lower model costs into Sales and Marketing and Freemium models to capture share in a rapidly shifting competitive landscape.

Key beneficiaries (large cap): Atlassian, Salesforce, ServiceNow

Key beneficiaries (smid cap): Wix, Zoom Figma, and Box

- Key Risks: Low-cost models have the potential to reduce the cost of building new workflow applications, encouraging greater competition which threaten gross margin gains. In addition, customers may become less willing to pay premium prices for functionality available through generic agents.

Implications for Vertical Software

We see rising open-weight model adoption as a meaningful tailwind for Vertical Software, given the combination of 1) Under-automated end markets, 2) Outcome-aligned monetization, and 3) Deep ownership of industry workflows. Lower inference costs should make it increasingly economical for vertical software vendors to automate historically manual, labor-intensive processes, while the ability to customize and post-train open-weight models on proprietary industry data and context creates opportunities to improve performance and further differentiate their offerings. Importantly, we believe the biggest beneficiaries will be vendors with deep domain expertise and the ability to develop domain-specific evaluations and feedback loops that measure successful task completion and continuously improve agent performance. These capabilities should allow vendors to extract greater value from increasingly commoditized intelligence, potentially achieving superior outcomes at lower costs without developing proprietary foundation models from scratch. While open-weight models also lower barriers to entry for AI-native challengers, we believe incumbents with proprietary workflow data, deep industry integrations, established distribution, and superior domain-specific evaluation capabilities are well positioned to defend their competitive moats and capture incremental AI-driven growth.

Key Beneficiaries: SHOP, TOST, MANH, NAVN, TTAN

- Key Risks: Vertical software incumbents that move too slowly to develop and monetize agentic capabilities risk ceding ownership of the System of Action to AI-native competitors, potentially weakening their competitive moats and relegating them to lower-value systems.