Internet/e-Commerce

BofA Frontier AI Data Tracker: Model rankings, usage data, and token/GPU costs

Industry Overview

Key signals shaping Internet AI stocks sentiment

Frontier AI capabilities, AI usage & pricing, and the pace of AI investments remain key drivers of sentiment and valuation across large-cap Internet stocks. The BofA Frontier AI Tracker monitors key indicators across the AI value stack, including frontier model intelligence rankings, usage trends, token pricing, GPU rental rates & memory costs.

Intelligence: Claude Fable 5.1 & GPT-6 Astra leading

Per Artificial Analysis, Claude Fable 5.1 (max w/fallback) and GPT-6 Astra (max) both ranked highest in intelligence index with a score of 53. Meta's newest model, MuseSpark 1.3 (max) ranked 5th, while Gemini 3.8 Flash ranked 12th. Claude Fable leads on agentic, while MuseSpark is a close T-2nd. MuseSpark's rankings suggest MSL is narrowing the performance gap and we look forward to the Watermelon launch. Maintain Buy on Meta.

Usage: DeepSeek leads in token share, Anthropic on spend

Data from Vercel AI platform suggests DeepSeek gained token usage share m/m and led at 46% market share, while Anthropic ranked second for usage at 14% and first for spend at 52% (down from 64% in Aug.). OpenAI accounted for 12% of usage & 22% of spend (up from 13% in Aug.). On OpenRouter, token usage share was led by DeepSeek V4 Flash 0731 (Open) at 50.7tr tokens MTD, followed by GPT-5.6 Luna (Closed) at 45.1tr tokens and Hy4 preview (Open) at 34.6tr tokens.

Token prices stabilize per TPI following August decline

AI Token price Index (TPI) was up 5% m/m (up 104% y/y) to $2.41 as of Sep 13 while the LLM Token Expenditure Index was down 11% m/m to $0.98. Potential drivers of change in AI token price index included: 1) GPT-6 Astra joined the index matching Fable 5.1 as the highest-priced model, 2) GLM-5.3-Flash entered as the lowest-priced model, and 3) DeepSeek V4.1 Flash replaced V4 Flash at a lower price.

GPU rentals mixed; memory largely stable m/m

According to Silicon Data, B200 rental prices were up 1% m/m (+9% y/y) to $5.68 in Sept, H100 were down 4% m/m (+30% y/y) to $2.65 & A100 were down 3% m/m (+15% y/y) to $1.60. DRAM memory prices (DDR5 – 5600/6400 24GB 3Gx8) were up 1% m/m (+476% y/y) and NAND memory prices (QLC Flash 1TB) were down 1% m/m (+437% y/y). Data suggests server useful lives remain healthy (A100 shipments started in 2020).

Healthy AI demand, Meta closing gap to the frontier

Sept data suggests: 1) Token price has stabilized after August decline, 2) Anthropic & OpenAI continue to lead on intelligence, and models are closely matched, 3) Meta's latest models are narrowing the intelligence gap, and we remain constructive on upcoming AI product launches and new Watermelon LLM (expected in Oct), 4) Open models continue to gain usage share, and DeepSeek continues to lead token usage on some platforms, 5) GPU rental pricing remains firm suggesting AI infrastructure demand remains healthy, and 6) Memory costs were stable m/m but remain significantly elevated y/y. We think debate around AI safety/regulation may influence sentiment on hyperscalers in the near-term, See our note: AI safety concerns back in the headlines.

BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.
Refer to important disclosures on page 13 to 17. Analyst Certification on page 11. Price Objective Basis/Risk on page 11. 13018161

15 September 2026

Equity
United States

Justin Post
Research Analyst
BofAS
+1 415 676 3547
justin.post@bofa.com

Nitin Bansal, CFA
Research Analyst
BofAS
+1 415 676 3551
nbansal7@bofa.com

MTD: Month to Date
GPU: Graphic Processing Unit
API: Application Programming Interface
TPI: Token price Index
DRAM: Dynamic random-access memory
MSL: Meta Superintelligence Labs

BofA Frontier AI Tracker - Sept Overview

AI Index Dashboard

Exhibit 1: Anthropic models leading in frontier AI capabilities
BofA Frontier AI Tracker (September 2026 Data Overview)

Frontier AI Capabilities	
Top Model by Intelligence Score	Anthropic Fable 5.1 and OpenAI GPT-6 Astra
Top Model by Coding Agentic Score	Anthropic Fable 5.1 and OpenAI GPT-6 Astra

Model Usage	
Highest usage on OpenRouter (Closed model)	OpenAI GPT 5.6 Luna (45.1tr tokens MTD)
Highest usage on OpenRouter (Open model)	DeepSeek V4 Flash 0731 (50.7tr tokens MTD)
Highest token volume share (Vercel AI Platform)	DeepSeek
Highest spend share (Vercel AI Platform)	Anthropic

Token Pricing	
AI Token Pricing Index	Up 5% m/m (up 104% y/y)
LLM Token Expenditure Index	Down 11% m/m

GPU Rental Prices	
B200 (Blackwell)	Up 1% m/m (up 9% y/y)
H100 (Hopper)	Down 4% m/m (up 30% y/y)
A100 (Ampere)	Down 3% m/m (up 15% y/y)

Memory Prices	
DRAM	Up 1% m/m (up 476% y/y)
NAND	Down 1% m/m (up 437% y/y)

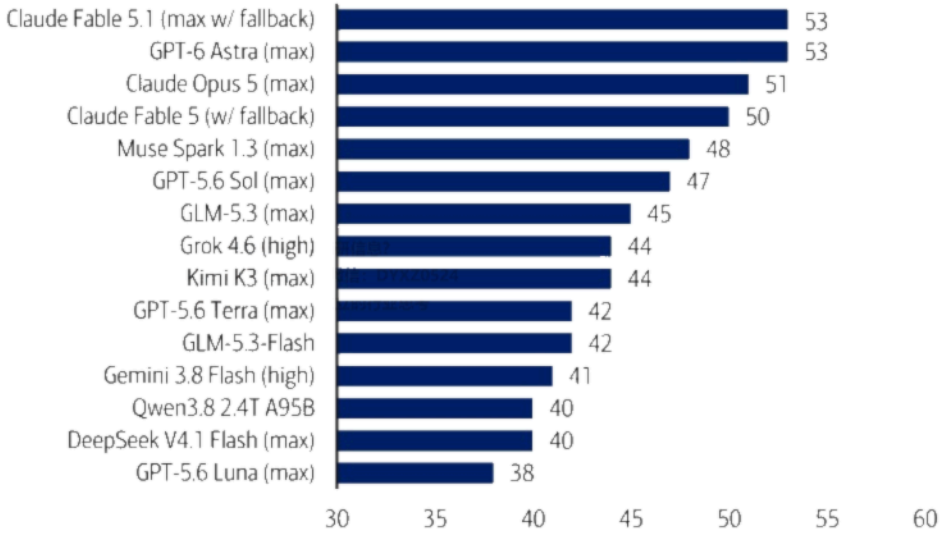
Source: Artificial Analysis, Bloomberg, Silicon Data, OpenRouter

Frontier Model Intelligence Ranking

Intelligence Index

According to Artificial Analysis, Claude Fable 5.1 (max w/fallback) and GPT-6 Astra (max) both ranked highest in intelligence index with a score of 53, followed by Claude Opus 5 (max), with a score of 52. Meta's newest model, MuseSpark 1.3 (max) ranked 5th, with intelligence index score of 48 while Gemini 3.8 Flash ranked 12th, with a score of 41.

Exhibit 2: Claude Fable 5.1 (max w/fallback) was ranked highest in intelligence index
Intelligence Index - Top 15 Models

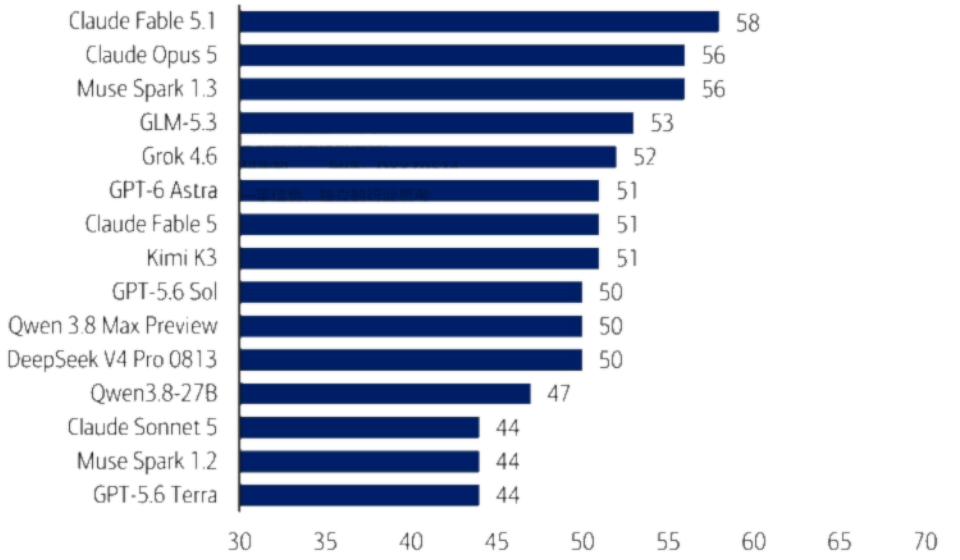


Source: Artificial Analysis
Note: Data as of Sep 14, 2026

Agentic Index

According to Artificial Analysis & BenchLM, Claude Fable 5.1 ranked highest in Agentic Index with a score of 58, followed by tie between Claude Opus 5 (max) and Meta's Muse Spark 1.3, with scores of 56.

Exhibit 3: Claude Fable 5.1 ranked highest in Agentic Index
Agentic Index – Top 15 Models

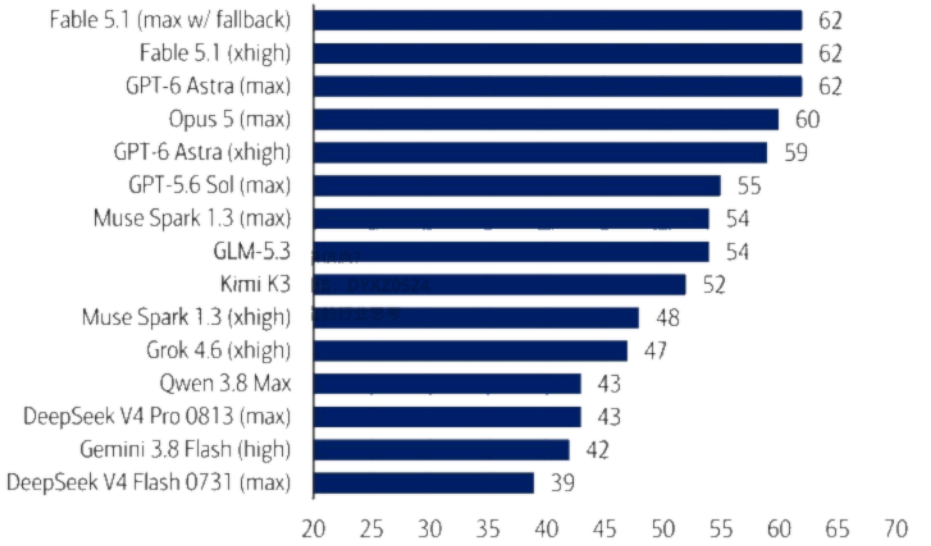


Source: Artificial Analysis, BenchLM
Note: Data as of Sep 14, 2026

Coding Agent Index

According to Artificial Analysis, Claude Fable 5.1 (max w/fallback) and GPT-6 Astra (max), both ranked highest in Coding Agent index with a score of 62, followed by Claude Opus 5 (max) with a score of 60. Meta's MuseSpark 1.3 (max) ranked 7th, with a score of 54 and Gemini 3.8 Flash (high) ranked 14th, with a score of 42.

Exhibit 4: Claude Fable 5.1 and GPT-6 Astra, both ranked highest in Coding Agent index
Coding Agent Benchmark - Top 15 Models



Source: Artificial Analysis
Note: Data as of Sep 14, 2026

Top Models by Intelligence Score by BenchLM

According to BenchLM, Anthropic's Claude Fable 5.1 and OpenAI's GPT-6 Astra lead the overall intelligence index, with a score of 84. Google's Gemini 3.8 Flash ranked 6th with a score of 78 and Meta's Muse Spark 1.2 ranked 13th with a score of ~71. Note: Meta's newest model MuseSpark 1.3 was not included in BenchLM analysis.

Exhibit 5: Claude Fable 5.1 and GPT-6 Astra models lead in the intelligence score
Top 20 Models by Intelligence by BenchLM

	Model			Pricing		Intelligence Score					
Rank	Model	Creator	Type	Input Price	Output Price	Overall Score	Coding	Agentic	Reasoning	Knowledge	Math
1	Claude Fable 5.1	Anthropic	Proprietary	$10	$50	84	84	80	79	87	-
2	GPT-6 Astra	OpenAI	Proprietary	$10	$50	84	75	70	90	82	85
3	Claude Opus 5	Anthropic	Proprietary	$5	$25	82	76	78	76	82	-
4	Claude Fable 5	Anthropic	Proprietary	$10	$50	81	77	75	78	83	-
5	GPT-5.6 Sol	OpenAI	Proprietary	$4	$20	81	75	70	71	81	97
6	Gemini 3.8 Flash	Google	Proprietary	$1	$4	78	68	67	77	76	-
7	Kimi K3	Moonshot AI	Proprietary	$3	$15	75	68	72	79	72	-
8	Claude Opus 4.8	Anthropic	Proprietary	$5	$25	72	67	62	56	72	65
9	GPT-5.5	OpenAI	Proprietary	$5	$30	72	68	62	64	73	70
10	Qwen3.8 Max	Alibaba	Open Weight	-	-	72	61	67	87	69	-
11	GPT-5.6 Terra	OpenAI	Proprietary	$3	$12	71	67	60	64	69	97
12	GPT-5.4	OpenAI	Proprietary	$3	$15	71	54	54	57	69	65
13	Muse Spark 1.2	Meta	Proprietary	$1	$4	71	60	61	75	71	-
14	Muse Spark 1.1	Meta	Proprietary	-	-	70	60	60	74	71	-
15	Claude Opus 4.7	Anthropic	Proprietary	$5	$25	70	63	56	-	65	61
16	Grok 4.6	xAI	Proprietary	$2	$6	70	67	69	76	70	-
17	Gemini 3.1 Pro	Google	Proprietary	$2	$12	70	46	40	51	66	54
18	Claude Sonnet 5	Anthropic	Proprietary	$2	$10	70	64	66	77	67	-
19	Claude Opus 4.7 (Adaptive)	Anthropic	Proprietary	$5	$25	69	58	60	50	61	-
20	Claude Opus 4.6	Anthropic	Proprietary	$5	$25	69	57	50	67	61	59

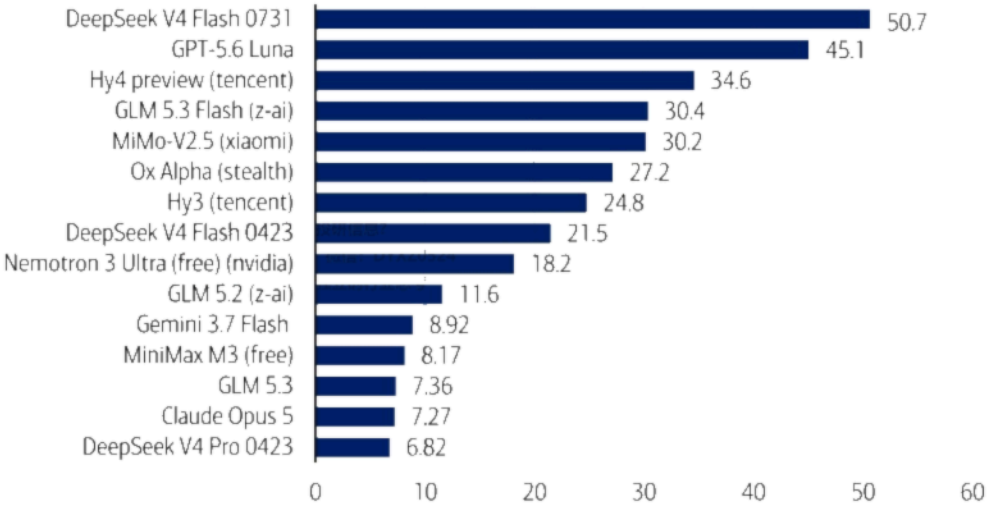
Source: benchlm.ai
Note: Data as of Sep 10, 2026 and BenchLM analysis could be excluding some models.

Model Usage / Adoption

Token Usage (OpenRouter)

Model usage on OpenRouter was led by DeepSeek V4 Flash 0731 at 50.7tr tokens MTD. GPT-5.6 Luna ranked 2nd at 45.1tr tokens, followed by Hy4 preview at 34.6tr tokens.

Exhibit 6: DeepSeek V4 Flash 0731 and GPT-5.6 Luna lead in token usage
Top 15 Models by Token Usage (September MTD, trillion token)

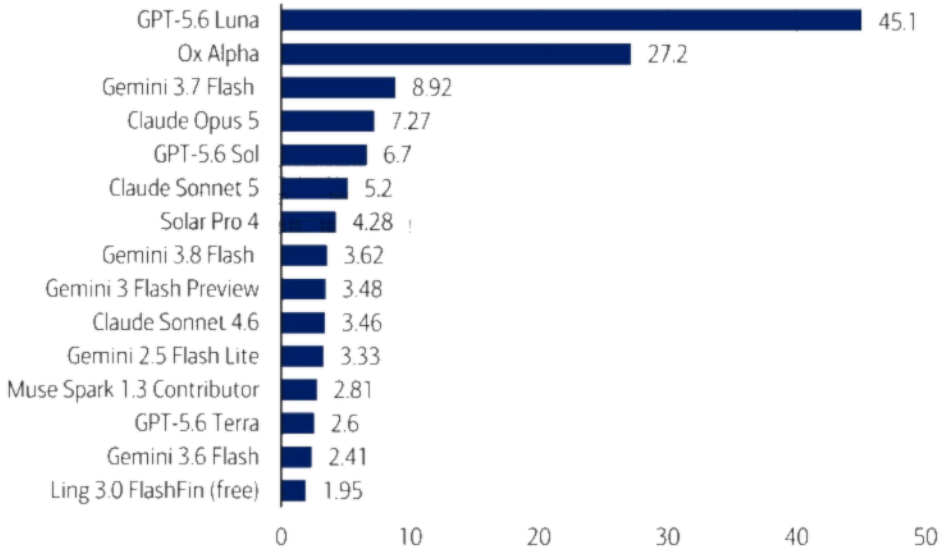


Source: Openrouter.ai
Note: OpenRouter usage data of Sep 13, 2026

Closed Models Token Usage (OpenRouter)

Among closed models GPT 5.6 Luna led in usage with 45.1tr token MTD, followed by Ox Alpha with 27.2tr token and Gemini 3.7 Flash with 8.92tr tokens.

Exhibit 7: GPT 5.6 Luna and Ox Alpha lead in token usage
Top 15 Closed Models by Token Usage (September MTD, trillion token)



Source: Openrouter.ai
Note: OpenRouter usage data of Sep 13, 2026

How to interpret OpenRouter model usage data?

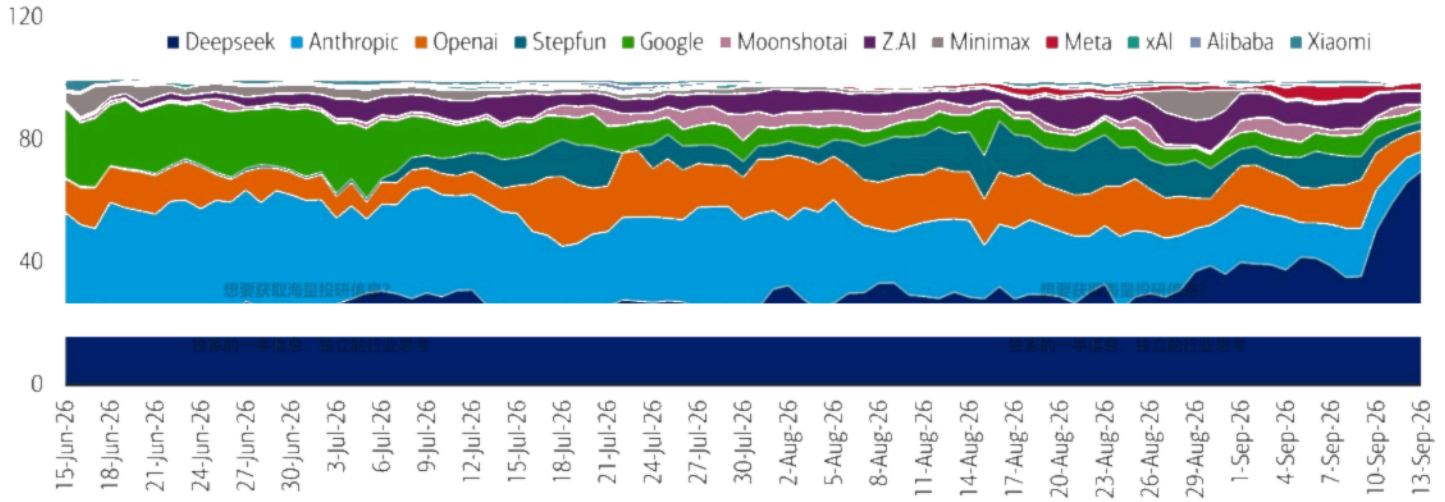
OpenRouter is an AI model aggregation and routing platform that provides developers with access to multiple LLMs through a single API. OpenRouter LLM usage rankings offer a real-time view of model adoption based on token volume routed through the OpenRouter platform, making the data a useful indicator of developer mindshare, cost-performance trade-offs and commercial traction within the frontier AI ecosystem.

However, we note OpenRouter data should not be interpreted as total model usage across the broader AI market. The dataset captures usage from developers and applications that route requests through OpenRouter but does not include workloads that go directly through provider-specific APIs or cloud platforms such as OpenAI API, Anthropic API, Google Vertex AI/Gemini API, AWS Bedrock or Azure OpenAI. It also excludes usage from enterprise private deployments, self-hosted inference stacks, hyperscaler internal workloads and company-owned data centers unless those workloads are routed through OpenRouter.

Token Volume Share on Vercel AI Platform

DeepSeek led the average MTD token usage share on the Vercel AI platform at 45.7% (vs. 30.1% in August), followed by Anthropic with average MTD token market share at 13.6% (vs. 22.1% in August), OpenAI's 12% (vs. 14.9% in August) and Z.AI 7.2% (vs. 6.8% in August). The most notable share gains were DeepSeek (+15.5pt) and Meta (+1.9pt), followed by Z.AI (+0.4pt), while the largest declines were Anthropic (down 8.5pt), Stepfun (down 4.5pt), and OpenAI (down 2.9pt).

Exhibit 8: DeepSeek lead AI Token Usage on Vercel, capturing approximately 45.7% share MTD in September (vs. 30.1% in August)
Daily Token Volume Share on Vercel AI Platform (%)



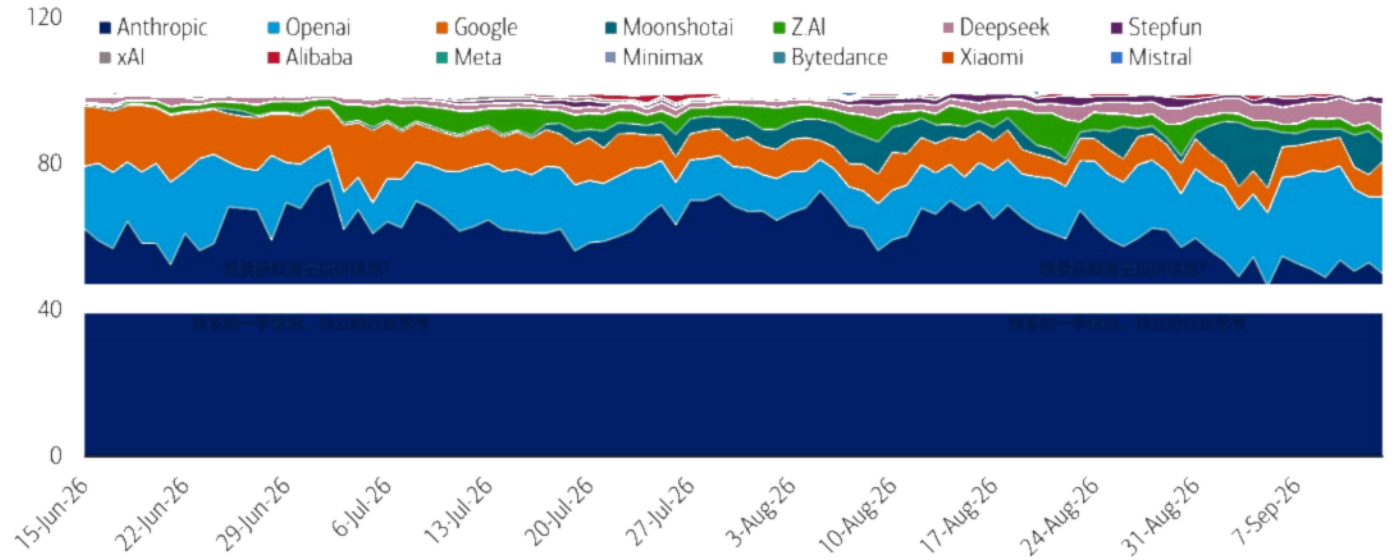
Source: Vercel.com
Note: Vercel AI data as of Sep 13, 2026

Vercel's daily token share data captures each AI lab's share of the total tokens routed through the Vercel AI Gateway, regardless of the price paid per token. Similar to OpenRouter, the data reflects only activity routed through the Vercel gateway and excludes AI inference activity handled directly through the respective labs' APIs or through third-party cloud providers.

Spend Share on Vercel AI Platform

Anthropic led the average MTD spend share on the Vercel AI platform at 52.0% (vs. 63.9% in August), followed by OpenAI at 21.8% (vs. 13.3% in August), and Moonshot.ai at 8.2% (vs. 4.5% in August). The most notable spend share gains were OpenAI (+8.5pt), Moonshot.ai (+3.7pt), and DeepSeek (+2.3pt), while the largest declines were Anthropic (-11.9pt), Z.AI (-2.3pt), and xAI (-0.3pt).

Exhibit 9: Anthropic lead AI model spend on Vercel, capturing approximately 52% share
Daily Spend Share on Vercel AI Platform (%)



Source: Vercel.com

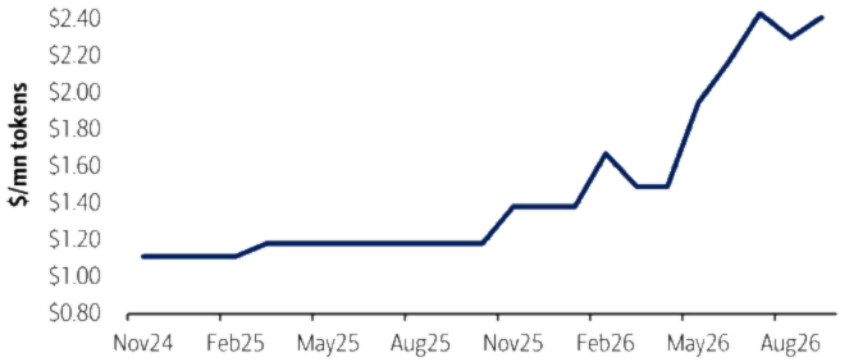
Vercel's daily spend share represents each AI lab's share of total AI Gateway spend. Vercel calculates spend using market-rate pricing based on providers' published list prices, providing a normalized measure across teams, including those using their own API keys. The metric therefore reflects both the volume of tokens routed to each lab and the market-rate price of those tokens. The data covers only activity routed through Vercel AI Gateway and excludes AI inference handled directly through the respective labs' APIs or through other cloud providers.

Token Pricing

AI Token Price Index- Up 5% m/m aided by GPT-6 addition

AI Token Price Index increased 5% m/m (up 104% y/y) to reach $2.41 as of Sept 10. Potential drivers of the m/m increase included: 1) GPT-6 Astra joined the index at $10 per mn input tokens and $50 per mn output tokens, matching Fable 5.1 as the highest-priced model, 2) GLM-5.3-Flash entered at $0.15 per mn input tokens and $0.50 per mn output tokens, establishing a new price floor, and 3) DeepSeek V4.1 Flash replaced V4 Flash at a lower price.

Exhibit 10: In September, Token price index went up 5% m/m (+104% y/y).
Token Price Index – Avg cost per 1mn tokens, aggregated across leading providers



Source: Token Price Index
Note: Data as of Sep 10, 2026

The Token Price Index (TPI) tracks the average cost of accessing leading frontier AI models across major providers, including Anthropic, OpenAI, Google, xAI, DeepSeek, Mistral, and Alibaba's Qwen. The index is designed to measure how frontier AI pricing evolves over time as new models are introduced and providers adjust pricing. TPI serves as a barometer for the economics of frontier AI. Index levels can rise when higher-performing premium models gain share or enter the market and fall when lower-cost alternatives or pricing reductions improve the overall cost efficiency of accessing advanced AI capabilities.

LLM Token Expenditure Index (Silicon Data)

LLM Token Expenditure Index declined 11% m/m to $0.98 as of Sep 13 (vs $1.10 in August).

Exhibit 11: In Sep, LLM Token Expenditure Index declined 11% m/m to $0.98
Silicon Data LLM Token Expenditure Index



Source: Bloomberg, Silicon Data LLM Token Expenditure Index

The Silicon Data LLM Token Expenditure Index (SDLLMTK) is a daily benchmark that tracks the effective cost of AI inference across the LLM ecosystem, expressed as a normalized blended rate in USD per million tokens. The index is calculated using pricing observations across more than 20 leading models and is weighted towards where inference usage is concentrated, providing a market-based view of what users are actually paying for AI inference. The index incorporates data from frontier API providers, open-weight inference platforms, dedicated-instance markets, and self-hosted deployments, offering a broader measure of AI inference economics.

Cost per Task (Artificial Analysis)

According to Artificial Analysis, Claude Fable 5 with fallback has the highest cost per Intelligence Index task among leading frontier AI models at $8.75, followed by Claude Fable 5.1 Max with fallback at $7.63 and Claude Opus 5 Max at $5.86. GPT-6 Astra Max costs $3.26 per task, and Open AI suggests the model can deliver compelling price performance vs leading competition.

Exhibit 12: Cost per task ranking among leading AI models
Weighted average Dollar Cost per Intelligence Index task, by model



Source: Artificial Analysis
Note: Data as of Sep 14, 2026

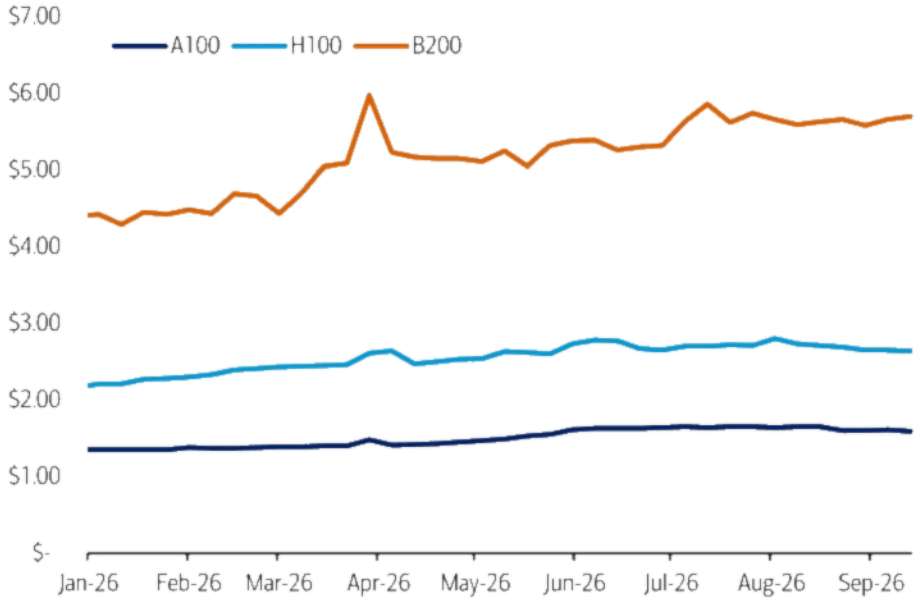
Artificial Analysis' Cost per Intelligence Index Task measures the effective cost of delivering a unit of AI capability, adjusting for model quality as captured by the Artificial Analysis Intelligence Index. The metric estimates the weighted average cost (USD) of completing a standardized AI task by incorporating all relevant inference costs, including input, cache, reasoning, and output tokens, and weighting each evaluation according to its contribution to the Intelligence Index. Lower values indicate greater cost efficiency, meaning a model can achieve a given level of intelligence at a lower inference cost.

GPU Pricing

Silicon Data GPU Rental Index

According to Silicon Data GPU Rental index, as of Sep 13, B200 (Blackwell) rental prices were up 1% m/m (+9% y/y) to $5.68, H100 (Hopper) rental prices were down 4% m/m (+30% y/y) to $2.65, and A100 (Ampere) rental prices were down 3% m/m (+15% y/y) to $1.60.

Exhibit 13: In Sep, B200 rental prices were up 1% m/m, H100 -4% m/m and A100 -3% m/m
Silicon Data GPU Rental Index ($ rental price/hour)



Source: Bloomberg, Silicon Data Indices
Note: Data as of Sep 13, 2026

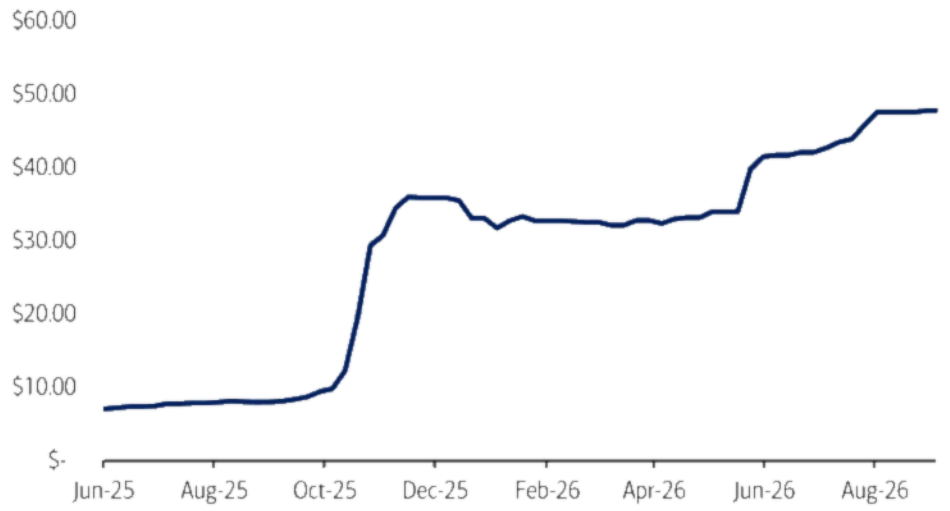
The Silicon Data GPU Rental Index tracks the market rental price of AI GPUs. Prices are expressed as a standardized $USD per GPU-hour and are derived from data across neo-cloud providers, hyperscalers, colocation markets, brokered cluster sales, and private rental platforms.

Memory Pricing

DRAM spot prices

According to Bloomberg, DRAM memory prices (DDR5 – 5600/6400 24GB 3Gx8) were up 1% m/m (+476% y/y) as of Sep 13.

Exhibit 14: DRAM memory prices were up 1% m/m in September
Spot Price of DRAM DDR5 – 5600/6400 24GB 3Gx8

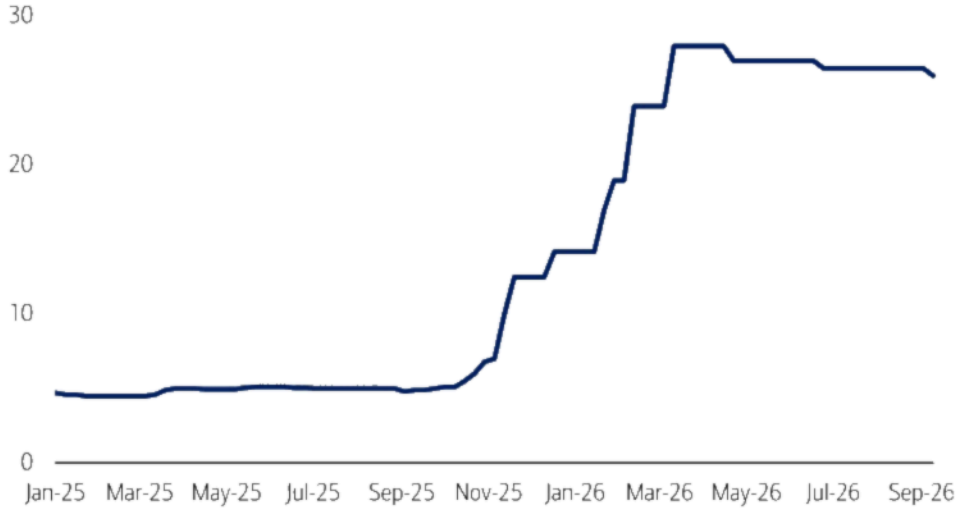


Source: Bloomberg, inSpectrum Tech
Note: Data as of Sep 13, 2026

NAND spot prices

According to Bloomberg, NAND memory prices (QLC Flash 1TB) were down 1% m/m (+437% y/y) as of Sep 13.

Exhibit 15: NAND memory prices were down 1% m/m in September
Spot Price of NAND DDR5 - QLC Flash 1TB



Source: Bloomberg, inSpectrum Tech
Note: Data as of Sep 13, 2026

Exhibit 16: Companies mentioned
Companies mentioned in this report

BofA Ticker	Bloomberg ticker	Company name	Price	Rating
GOOGL	GOOGL US	Alphabet A	US$ 349.39	B-1-7
GOOG	GOOG US	Alphabet C	US$ 345.71	B-1-7
AMZN	AMZN US	Amazon.com	US$ 253.54	B-1-9
META	META US	Meta Platforms	US$ 665.6	B-1-7

Source: BofA Global Research

Price objective basis & risk

Alphabet (GOOGL / GOOG)
Our price objective of $430/$430 is based on 27x 2027E core Google GAAP EPS plus cash per share. Alphabet has traded at an average multiple of 22x GAAP P/E over the last ten years, and we think our target multiple is reasonable vs history given expectations for double-digit revenue growth, Cloud margin expansion, and opportunity to capitalize on strong AI assets.

Downside risks are 1) loss of search traffic to AI tools from competitors, 2) LLM integration in search may take longer than expected or negatively impact search revenues, 3) revenue pressure from compliance with the EU Digital Markets Act (DMA), and 4) potential for increasing Capex and lower FCF-given AI investments.

Amazon.com (AMZN)
Our $320 PO is based on our SOTP analysis that values AWS at 9x 2027E sales, 1P Retail at 1.0x, 3P Retail at 2.5x, and Advertising at 5.0x. Our 2027 sales multiples compare to 5.7x for SaaS comps, 1.3x for Retail comps, 2.4x for Marketplace comps, and 3.9x for Online Media comps. Our PO implies 3.6x blended 2027 P/Sales, 13x '27 EBITDA, and 31x '27 EPS.

Downside risks to our price objective are increasing competition from Cloud providers and big box retailers including Walmart, increasing AWS infrastructure costs that will pressure FCF, and Agentic AI usage that could pressure direct traffic and high-margin ad revenues. The stock has been subject to heavy volatility in the past, based on margin trends, and this volatility could increase due to economic uncertainty.

Meta Platforms Inc (META)
Our $810 price objective is based on 24x 2027E GAAP EPS. On a total company basis, including Metaverse investments, our valuation is at a slight premium to the S&P 500 given Meta's relatively high growth rate and AI opportunity. Historically, Meta has traded at an average premium of 3pts to the S&P 500.

Downside risks are (1) a significant dependence on digital advertising revenue, which remains highly sensitive to macroeconomic conditions, (2) elevated spend on AI initiatives that we expect to hurt margins, (3) a growing base of fixed assets that could reduce cost flexibility in a macro downturn, (4) intensifying competition for user engagement and advertising budgets from AI-native platforms, and (5) elevated regulatory overhang from new teen safety cases or unfavorable outcomes in ongoing litigation.

Analyst Certification

I, Justin Post, hereby certify that the views expressed in this research report accurately reflect my personal views about the subject securities and issuers. I also certify that no

part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or view expressed in this research report.



US - Internet Coverage Cluster

Investment rating	Company	BofA Ticker	Bloomberg symbol	Analyst
BUY				
	Alphabet	GOOGL	GOOGL US	Justin Post
	Alphabet	GOOG	GOOG US	Justin Post
	Amazon.com	AMZN	AMZN US	Justin Post
	Booking Holdings Inc	BKNG	BKNG US	Justin Post
	Chewy Inc	CHWY	CHWY US	Michael McGovern
	DoorDash	DASH	DASH US	Justin Post
	Ethos	LIFE	LIFE US	Michael McGovern
	Etsy, Inc.	ETSY	ETSY US	Michael McGovern
	Expedia	EXPE	EXPE US	Justin Post
	LegalZoom	LZ	LZ US	Michael McGovern
	Life360	LIF	LIF US	Nitin Bansal, CFA
	Magnite, Inc.	MGNI	MGNI US	Omar Dessouky, CFA
	Meta Platforms Inc	META	META US	Justin Post
	Navan, Inc.	NAVN	NAVN US	Nafeesa Gupta
	Peloton	PTON	PTON US	Nafeesa Gupta
	Take-Two Interactive	TTWO	TTWO US	Omar Dessouky, CFA
	Tripadvisor	TRIP	TRIP US	Nafeesa Gupta
	Uber	UBER	UBER US	Justin Post
	Unity	U	U US	Omar Dessouky, CFA
	Wayfair	W	W US	Michael McGovern
	Wix.com	WIX	WIX US	Michael McGovern
NEUTRAL				
	Airbnb	ABNB	ABNB US	Justin Post
	AppLovin	APP	APP US	Omar Dessouky, CFA
	Carvana	CVNA	CVNA US	Michael McGovern
	Coursera Inc.	COUR	COUR US	Nafeesa Gupta
	eBay	EBAY	EBAY US	Michael McGovern
	Instacart	CART	CART US	Justin Post
	Match Group	MTCH	MTCH US	Omar Dessouky, CFA
	Pinterest	PINS	PINS US	Justin Post
	Reddit	RDDT	RDDT US	Justin Post
	Roblox Corp. Class A	RBLX	RBLX US	Omar Dessouky, CFA
	Snap	SNAP	SNAP US	Justin Post
	Zillow	Z	Z US	Omar Dessouky, CFA
	Zillow Group	ZG	ZG US	Omar Dessouky, CFA
UNDERPERFORM				
	Bending Spoons SpA	BSP	BSP US	Omar Dessouky, CFA
	Bumble	BMBL	BMBL US	Omar Dessouky, CFA
	Cardlytics, Inc	CDLX	CDLX US	Omar Dessouky, CFA
	Duolingo	DUOL	DUOL US	Omar Dessouky, CFA
	Ibotta	IBTA	IBTA US	Nitin Bansal, CFA
	Lyft, Inc.	LYFT	LYFT US	Justin Post
	StubHub Holdings	STUB	STUB US	Justin Post
	Yelp	YELP	YELP US	Nitin Bansal, CFA
RSTR				
	Global Business Travel Group, Inc.	GBTG	GBTG US	Nafeesa Gupta

Disclosures

Important Disclosures