Hardware & Networking

Data Center Watch: Token Spend Growth Accelerates in October Amid Mixed GPU Rental and Memory Pricing

IT Hardware/ Telecom & Networking Equipment

Joseph Cardoso AC
(1-212) 622-9036
joseph.cardoso@jpmchase.com

Manmohanpreet Singh
(1-212) 622-4527
manmohanpreet.singh@jpmchase.com

Marc Vitenzon
(1-212) 622-3342
marc.vitenzon@jpmchase.com
J.P. Morgan Securities LLC

In the Data Center Watch report, we outline key data points relevant to AI infrastructure demand, including LLM token volumes and prices, GPU rental prices, and memory costs. Taken together, the latest data on LLM token spend growth suggest AI infrastructure demand remains strong, but pricing signals from GPU rentals and memory spot prices were more mixed. LLM token spend growth accelerated to +40% m/m in October from +28% m/m in September, driven by strong volume growth and a rebound in volume-weighted pricing, primarily reflecting a mix shift toward higher-priced models, particularly Claude Opus 5.5. Meanwhile, sequentially, GPU rental prices declined further for A100s, while H100 pricing rebounded and B200 price growth accelerated. With regards to memory, DRAM spot prices rebounded following September's modest decline, while NAND prices declined for a second consecutive month. Growth in contract prices for both DRAM and NAND is expected to moderate further to +13% and +18% q/q, respectively, in 4Q26E, versus +16% and +21% q/q in 3Q26. *Key highlights from Data Center Watch report:*

- **Token spend growth accelerated sequentially on OpenRouter in October.** LLM growth trends on OpenRouter in October were strong: token volume tracked to +25% m/m and 33x y/y, while average pricing tracked to -5% m/m and +9% y/y. Net-net, overall spending was robust, with growth tracking to +40% m/m and 20x y/y. Volume-weighted pricing increased +12% m/m and declined -38% y/y, with the m/m increase driven by: (1) mix shift toward higher-priced models, particularly Opus 5.5, which added ~7.9 tn tokens at a premium price; and (2) same-model repricing, with GPT-5.6 Sol pricing doubling as its 50% discount lapsed while Kimi K3 and DeepSeek v4 Pro saw increased output prices.
- **A100 GPU rental prices declined further in October, while H100 prices rebounded and B200 price growth accelerated.** Sequential growth trends for GPU rental prices remained mixed in October for Nvidia GPUs with regards to non-hyperscaler capacity. A100 prices declined -3.5% m/m, while H100 prices rebounded +6.5% m/m and B200 prices increased +4.7% m/m. Pricing for B200s was 2.1x that of H100s, down sequentially, while H100 pricing was 1.9x that of A100s, up sequentially.
- **Growth in contract prices for both DRAM and NAND is expected to moderate further in 4Q26.** Per TrendForce, conventional DRAM and NAND contract prices are expected to grow +10–15% q/q and +15–20% q/q in 4Q26, respectively vs. +16% q/q for DRAM and +21% q/q for NAND in 3Q26. TrendForce attributed the moderation to prior price increases, pressure on consumer demand, sufficient buyer inventories in client SSDs, and long-term pricing agreements for some server DRAM suppliers. However, AI server demand and capacity allocation toward server products continue to support pricing, with enterprise SSDs the only category expected to see price growth accelerate in 4Q26.

See page 6 for analyst certification and important disclosures.

Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

J.P.Morgan

LLM Token Tracker

Our LLM Token Tracker captures a single week of usage and pricing data each month, providing a monthly snapshot or sampling of trends across ~300 LLMs listed on OpenRouter, an aggregation platform that routes API requests to models through a single endpoint. Keep in mind, OpenRouter skews towards developer, startup, and agentic-coding traffic and excludes first-party API volume (direct OpenAI/Anthropic, hyperscaler-hosted endpoints).

Key takeaways:

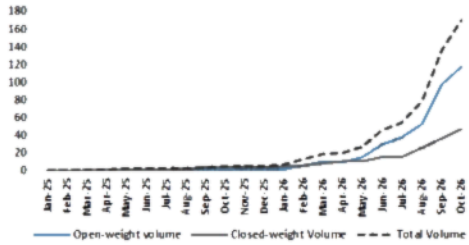
- **Token volume growth moderated in the month of October, following strength in September.** Token volume growth moderated on a sequential basis in October, expanding +25% m/m (vs. +71% in September and +47% in August), and 33x y/y (vs. 30x in September and 28x in August). Volumes for closed-weight models (OpenAI, Anthropic, etc.), which represented 28% of the total (vs. 27% in September and 33% in August), saw growth moderate in October, expanding +30% m/m (vs. +39% in September and +66% in August) and 12x y/y (vs. 10x in September and 14x in August). Volumes for open-weight models (DeepSeek, Moonshot AI, etc.) saw growth moderate more sharply in October, expanding +21% m/m (vs. +83% in September and +39% in August) and 97x y/y (vs. 101x in September and 56x in August).
- **Token VWAP rebounded in October.** Volume-weighted pricing increased +12% m/m and declined -38% y/y, with the m/m increase driven by: (1) mix shift toward higher-priced models, particularly Opus 5.5, which added ~7.9 tn tokens at a premium price; and (2) same-model repricing. Volume-weighted pricing for open-weight models declined -5% m/m (vs. -11% in September and -17% in August) and rose +24% y/y (vs. +30% in September and +60% in August) vs. +14% m/m in October for closed-weight models (vs. -15% in September and -37% in August) and +37% y/y (vs. ~flat in September and -29% in August). At the same time, average token pricing decreased -5% m/m in October (vs. +4% in September and -5% in August), while year-over-year growth moderated, with pricing expanding +9% y/y (vs. +20% in September and +5% in August).
- **Token spending growth accelerated in October.** Net-net, token spending growth strengthened in October, expanding +40% m/m (vs. +28% in September and +7% in August) and 20x y/y (vs. 14x in September and 12x in August). Spending for closed-weight models, which represented 77% of the total (vs. 72% in September and 78% in August), saw growth accelerate in October, increasing +49% m/m (vs. +18% in September and +4% in August) and 16x y/y (vs. 10x in September and 10x in August). By comparison, spending for open-weight models saw growth moderate sharply in October, increasing +15% m/m (vs. +64% in September and +16% in August) and 121x y/y (vs. 132x in September and 89x in August).
- **One model overlapped in the top five by spend versus usage.** The top five models in October based on token usage were: (1) DeepSeek V4.1 Flash, (2) DeepSeek V4.1 Flash 1017, (3) GLM 5.3 Flash, (4) Tencent's Hy4 preview, and (5) MiMo-V2.6-Flash, accounting for 47%+ of total token volume, down from 53%+ and from 50%+ in August. However, on a spending basis, the top five models in October were: (1) Claude Opus 5.5, (2) Tencent's Hy4 preview, (3) Claude Sonnet 5.5, (4) GPT-6 Astra, and (5) GPT-6.1 Sol, accounting for 54% of total token spend, up from 50% in September and 48% in August.



Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

Figure 1: Token Volumes

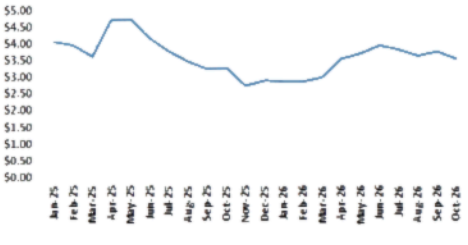
Weekly Tokens in Trillion



Source: OpenRouter, J.P. Morgan Estimates

Figure 2: Average Token Pricing

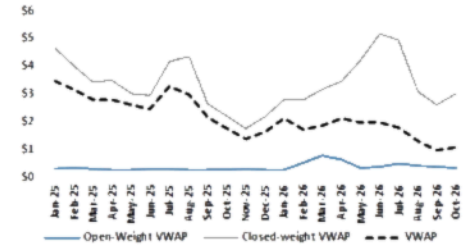
$ per Million Tokens



Source: OpenRouter, J.P. Morgan Estimates
* Includes models without observable volumes on OpenRouter

Figure 3: Volume-Weighted Average Token Pricing

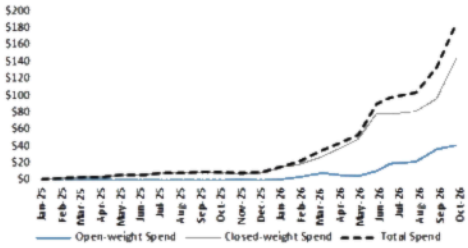
$ per Million Tokens



Source: OpenRouter, J.P. Morgan Estimates

Figure 4: Snapshot of Weekly Token Spend

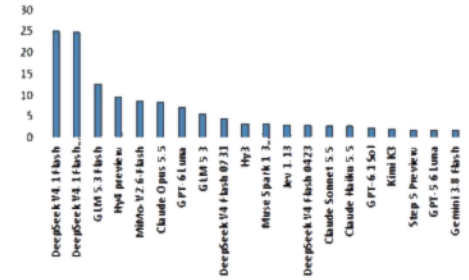
$ in Million



Source: OpenRouter, J.P. Morgan Estimates

Figure 5: Top 20 Models Based on Volume

Weekly Tokens in Trillion



Source: OpenRouter, J.P. Morgan Estimates

Figure 6: Token Pricing vs. Volume



Source: OpenRouter, J.P. Morgan Estimates

GPU Rental Price Tracker

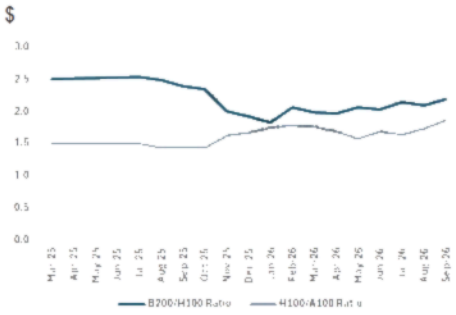
Key takeaways based on Bloomberg indices:

- **A100 rental prices declined further m/m in October.** A100 rental prices tracked to $1.53/GPU-hour on average in October, down -3.5% m/m (vs. -3.0% m/m in September and -0.7% m/m in August).
- **H100 rental prices rebounded m/m in October; H100-to-A100 ratio moved higher.** H100 rental prices tracked to $2.83/GPU-hour on average in October, up +6.5% m/m (vs. -2.0% m/m in September and +0.4% m/m in August). The H100-to-A100 price ratio tracked to 1.85x on average in October (vs. 1.67x in September and 1.65x in August).
- **B200 rental price growth accelerated m/m in October; B200-to-H100 ratio moved lower.** B200 rental prices tracked to $6.00/GPU-hour on average in October, up +4.7% m/m (vs. +1.8% m/m in September and -1.5% m/m in August). The B200-to-H100 price ratio tracked to 2.12x on average in October (vs. 2.16x in September, 2.08x in August,

Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

2.12x in July, and 2.58x around index launch in Sep-25).

Figure 7: Avg. Monthly Rental Prices (Non-Hyperscaler)



Source: Bloomberg Finance L.P.

Figure 8: Avg. Monthly Rental Prices (M/M % Change)



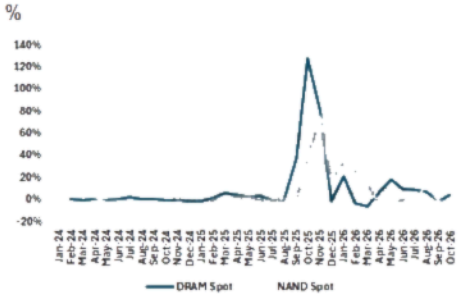
Source: Bloomberg Finance L.P.

Memory Price Tracker

Key takeaways based on Bloomberg data:

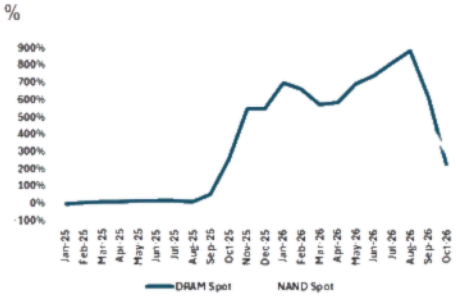
- **DRAM spot prices rebounded in October, with DDR5 16Gb up more than 3x y/y.** DDR5 16Gb spot prices reached $51.80 in October, increasing +4.4% m/m (vs. $49.60 in September and $50.20 in August) and +227% y/y (vs. $15.85 in October 2025), marking a return to monthly increases following September's modest decline.
- **NAND spot prices declined further in October; NAND 1Tb remains up nearly 4x y/y.** NAND 1Tb spot prices tracked to $29.40 in October, declining -2.6% m/m (vs. $30.20 in September and $30.50 in August) and increasing +298% y/y (vs. $7.38 in October 2025), marking the second consecutive monthly decline following August's increase.
- **DRAM and NAND contract prices expanded +16% and +21% q/q in 3Q26; contract prices to expand into 4Q26E as well, albeit at a moderated pace.** As per TrendForce, DRAM contract prices expanded +16% q/q in 3Q26 (vs. +56% q/q in 2Q26), while NAND contract prices increased +21% q/q (vs. +58% q/q in 2Q26). For 4Q26E, DRAM and NAND contract prices are expected to expand +13% q/q and +18% q/q, respectively.

Figure 9: Memory Prices (M/M % Change)



Source: Bloomberg Finance L.P.

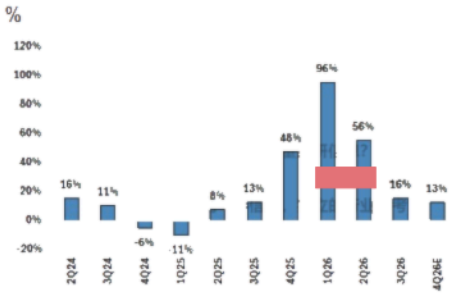
Figure 10: Memory Prices (Y/Y % Change)



Source: Bloomberg Finance L.P.

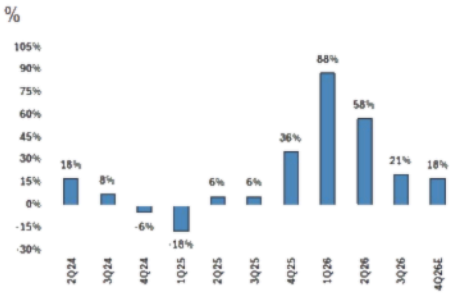
Joseph Cardoso [AC]
(1-212) 622-9036
joseph.cardoso@jpmchase.com

Figure 11: DRAM Contract Price Q/Q growth (%)



Source: TrendForce.

Figure 12: NAND Contract Price Q/Q growth (%)



Source: TrendForce.