AI Capex 2.0

If You Build It, They Will Finance It - Updated Growth & Capital Markets Forecasts

The financing wave is in full swing. Since we wrote in November, total AI/Data Center–related issuance has exceeded $300 billion, with data center–related debt issuance the single largest driver of near-record debt issuance to start the year. While it's been a strong start, we expect current issuance trends to represent a baseline given the amount of financing set to hit markets in the coming years. Corporate credit markets have dominated so far, but we expect issuers to tap every single capital market to support their growth needs. Relative to our early expectations, issuers have surprised a bit by globalizing their issuance, including large non-USD IG deals for hyperscaler and debut euro HY issuance for data centers. But we expect creativity to continue to be the overall baseline for this sector for years to come.

Figure 1: YTD AI-Related Debt Issuance

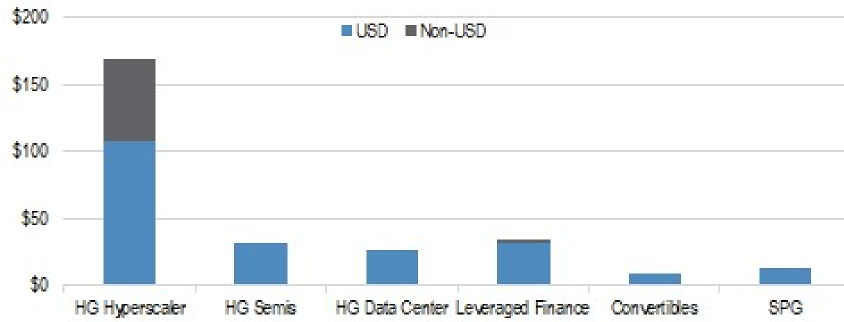

Source: J.P. Morgan.

Underlying growth forecasts nudged slightly higher. We continue to think that power remains the most critical constraint on potential data center capacity growth. Nonetheless, the amount of flexibility and creativity being shown by developers is remarkable, with bring-your-own-power and behind-the-meter power deals helping modestly debottleneck power constraints. Additionally, more power efficient cluster solutions, improved GPUs and custom ASICs are helping bring more compute per megawatt to market on a go-forward basis. One other constraint to watch is likely labor, with costs for skilled powerline electricians spiking and availability shrinking. Nonetheless, we expect 138 GW of capacity growth through 2030, up from 122 GW in our November forecast.

North American Fundamental Research

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Harlan Sur AC
(1-415) 315-6700
harlan.sur@jpmorgan.com

Nathaniel Rosenbaum, CFA AC
(1-212) 834-2370
nathaniel.rosenbaum@jpmorgan.com

Nelson Jantzen, CFA AC
(1-212) 270-1169
nelson.r.jantzen@jpmorgan.com

Chong Sin AC
(1-212) 834-2611
chong.c.sin@jpmorgan.com

Kevin L Kwan AC
(1-212) 270-9455
kevin.l.kwan@jpmorgan.com

Jeremy Tonet, CFA AC
(1-212) 622-4915
jeremy.b.tonet@jpmorgan.com

Erica R Spear AC
(1-212) 834-4143
erica.spear@jpmchase.com

Thomas Egan, CFA AC
(1-212) 270-2149
thomas.j.egan@jpmorgan.com

Kabir Caprihan, CFA AC
(1-212) 834-5613
kabir.x.caprihan@jpmorgan.com

Silvi Mantri AC
(1-212) 834-7239
silvi.mantri@jpmchase.com
J.P. Morgan Securities LLC

See page 56 for analyst certification and important disclosures.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 2: Growth to Ramp Up Sharply in 2027

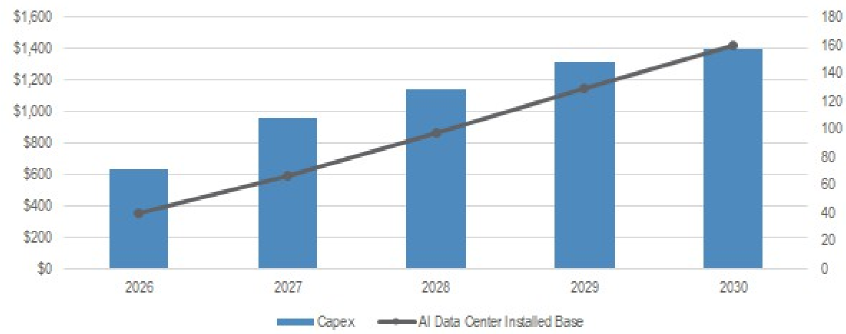

Source: J.P. Morgan.

Don't bury the lede. We expect total AI capex spending of $5.5 trillion through 2030, up from $5.1 trillion in November. Some of that increase is the GW increase noted above, partially offset by slightly lower assumed development costs per MW. However, we have also increased our assumed debt financing component of the AI capex buildout to $4.1 trillion, largely to reflect a higher overall loan-to-cost assumption. Disclosed project LTCs in our cohort have averaged over 85% with select LTCs over 90%. Some of that probably reflects the benign financing environment over the last six months. **But some of that reflects the implicit cost-to-value uplift.** To paraphrase an astute investor, if a company announces an expansion for $15mm a megawatt, the market cap tends to increase by $25mm a megawatt. **Given current market valuations, 90% Loan-to-Cost appears closer to 60% Loan-to-Value.**

Figure 3: Loan-to-Cost Ratios Have Moved Higher

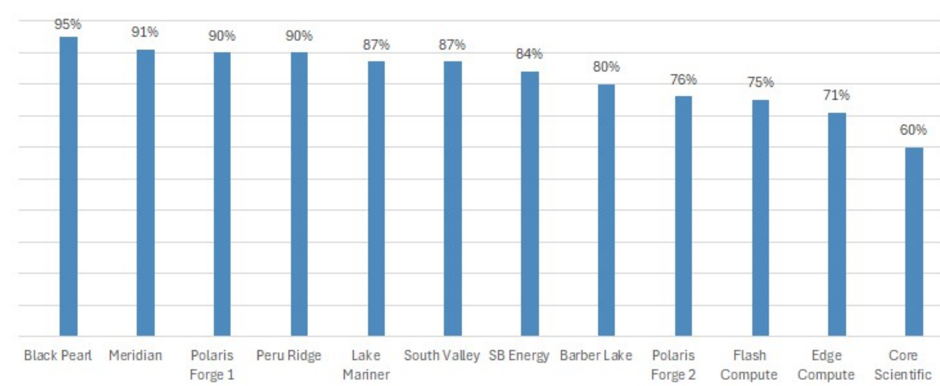

Source: Company reports and J.P. Morgan estimates.

But to be clear, the hyperscalers are still in great shape even as capex ramps. Hyperscaler capex has increased to $650 billion for 2026, and is likely headed above $1.1 trillion in 2027. However, hyperscalers remain remarkably profitable, with early returns on investment positive. We expect operating cash flow from the hyperscaler cohort to be over $900 billion in 2027. Additionally, recent equity offerings and IPOs and pending IPOs in the space should put additional equity on the sector's balance sheet. Frankly, we could argue that high recent LTC ratios across the sector also represent some husbanding of capital. Put another way, by leveraging projects now, companies reserve the right to use operating cash flow or primary equity market proceeds to reduce their debt funding needs in the future in a less attractive/ higher rate market. These management teams are smart and look to be preserving optionality for rainy days.

The rise of the GPU financing. Over the last six months, spending on data center infrastructure has been the dominant theme in markets. However, we expect capital spending on GPU and custom ASICs to start taking a disproportionate level of mind and wallet share. Over the next five years, we expect over $3 trillion of financings for AI accelerators. Interestingly, while the curve for data center financing will likely stabilize ~2028, chip financings will likely continue to grow into 2030, particularly as replacement demand becomes a theme. Finding the right silicon financing paradigm is the billion dollar question, in our view. Private credit has done a lot of work financing silicon so far, with the recent AI XPV JV between Broadcom/Apollo/Blackstone as a potential landmark deal (discussed below). However, we think public markets will eventually need to do the heavy lifting. Additionally, at some point, creativity will be needed to match the large funding needs of the AI silicon space with relatively limited capital available at the five-year duration point. Finding a way to access giant pools of capital in the Investment Grade market with duration of longer than ten years would solve a lot of funding challenges, in our view. But doing that given the underlying asset life/funding duration mismatch is not easy.

Figure 4: More Chip Financings in Our Future

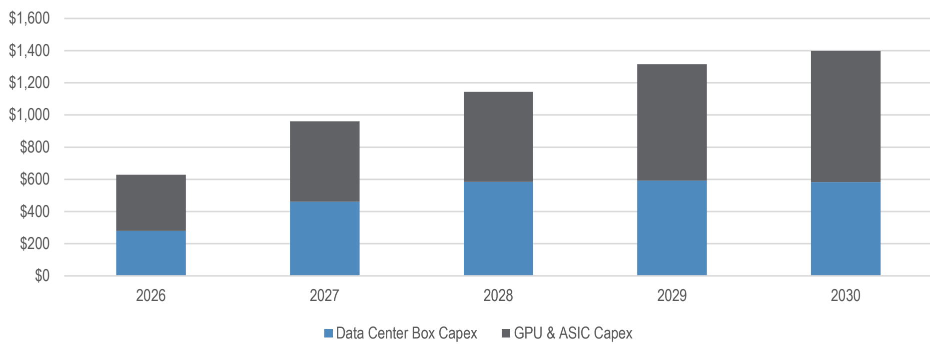

Source: J.P. Morgan.

HG Corporate markets still the likely workhorse across markets. We expect HG corporate markets to contribute over $2.1 trillion of data center financing over the next five years. That includes hyperscalers, data center owners and potential chip financings. Candidly, there will likely be at least some fungibility between public and private high grade funding, with much of the private funding CUSIP'd and potentially subject to limited syndication. Additionally, we expect approximately $350 billion of new funding in Leveraged Finance markets. One nuance to think about is the ratings upgrade cycle, particularly for data center projects. As projects are completed, ratings are likely to migrate to Investment Grade. That would push some primary High Yield paper into the High Grade markets as new Rising Stars. Our SPG numbers remain unchanged. We think the SPG markets are likely to remain opportunistic and ramp up their involvement in a wider spread/ yield environment.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 5: Anticipated AI Capex Funding Sources

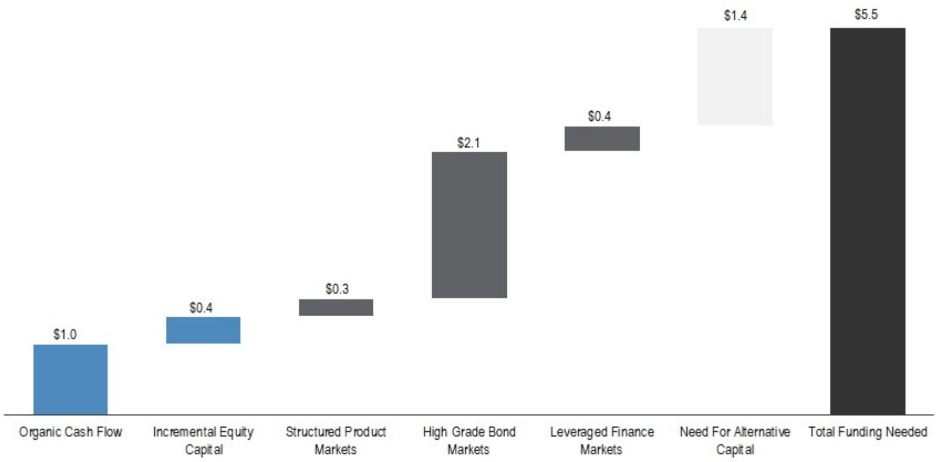

Source: J.P. Morgan.

Keep it simple stupid. The sector offers a giant universe of bonds and loans from hyperscalers to HG data center companies to HY data center companies—virtually all of which ultimately look to a hyperscaler as a tenant, customer, or partner. In a world of scarce compute, all compute is remarkably valuable and we like these instruments. In that environment, buying wider spread project deals into a likely refinancing or rise to Investment Grade presents an attractive risk/reward in our view. That may change at some point if and when compute shortages are ameliorated, but at present we remain very bullish across the stack.

Ironically, those compute shortages have really changed the monetization narrative. Six months ago, the existential question was how adoption can possibly be fast enough to allow these companies to make enough money to justify these investments. Now, the narrative is that corporate customers are huge consumers of AI but are blowing through their entire 2026 AI budgets in 1Q26. Claude Cowork and Claude Code really changed the game on monetization and adoption, with the dominant ChatGPT LLM consumer platform also performing to the upside. But we just don't have enough compute to allow for agentic adoption across the economy, and that is causing costs to the end user to increase even as per token costs fall. Our fear on growth at this point is that adoption stalls, not because the use cases are not there—they certainly are and we are using AI in our own work processes at an accelerating rate. Rather, our biggest current fear is that adoption stalls as costs are too high to justify some use cases, until such time as additional compute can further reduce token costs. Agentic workloads have triggered a logarithmic increase in token usage, while compute additions have triggered a linear reduction in token costs—we probably need a year or two for linear reductions to catch up to the agentic step change.

Tarek Hamid [AC] (1-212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC
Harlan Sur [AC] (1-415) 315-6700
harlan.sur@jpmorgan.com

Nathaniel Rosenbaum, CFA [AC] (1-212) 834-2370
nathaniel.rosenbaum@jpmorgan.com
Nelson Jantzen, CFA [AC] (1-212) 270-1169
nelson.r.jantzen@jpmorgan.com

J.P.Morgan

Table Of Contents

Data Center Growth & Inference

Tarek Hamid
(212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC

Harlan Sur
(415) 315-6700
harlan.sur@jpmorgan.com
J.P. Morgan Securities LLC

We highlight two key factors underpinning the large-scale AI infrastructure spend being planned for the next few years: (1) **AI adoption rates are soaring across consumer and enterprise end markets**, and (2) **the compute intensity of AI workloads continues to march higher, with agentic inference emerging over the past 6-9 months as a substantially larger driver of compute demand than we or the industry anticipated**. Agentic workloads—autonomous, multi-step tasks in which models plan, reason, invoke tools, and iterate—consume orders-of-magnitude more tokens per task than one-shot inference, and they have moved from the emerging use case to the marginal unit of demand growth: Google disclosed at I/O 2026 that it now processes ~3.2 quadrillion tokens per month across its surfaces (~7x Y/Y), while enterprises including Microsoft and Uber have reportedly exhausted full-year 2026 token budgets within months as agents enter production deployment. Critically, this demand is compounding even as per-token inference costs decline 60-70% per year—efficiency gains are being more than re-absorbed by rising usage, a classic Jevons dynamic. Ongoing increases in model scale and complexity (including multi-modality, i.e., audio, video, images, etc.) and the use of reasoning models that leverage test-time compute reinforce the trend. The net effect of the continuation of these trends is that significantly more compute capacity will be required to adequately serve burgeoning AI use than is currently available across the existing installed base of data center GPUs and AI ASICs ("XPUs")—with the mix of that demand shifting decisively toward inference.

Figure 6: Compute Mix Shifting from Training to Inference

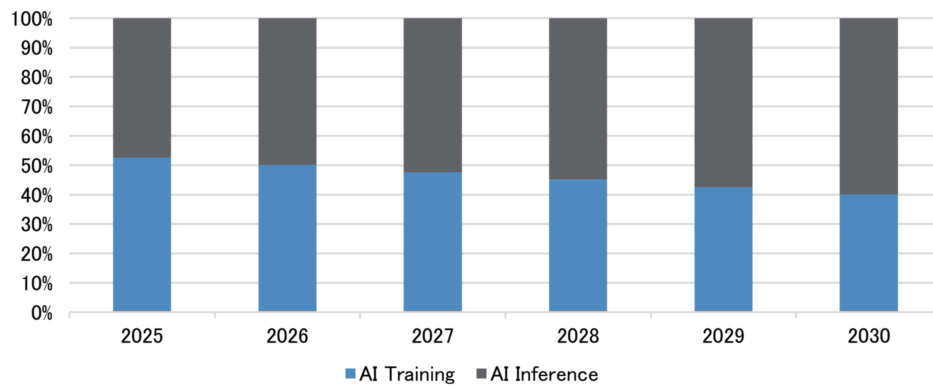

Source: McKinsey & Co. and J.P. Morgan estimates.

Even after three years of elevated spend on AI data center infrastructure (see Figure 9), which supported the deployment of an estimated ~21MM AI accelerators (GPU + XPU/ ASIC—see Figure 10) during 2023-2025, and an additional 16MM+ accelerators forecast to be deployed in 2026, providers and users of AI compute continue to face capacity shortages. Commentary we highlighted from hyperscalers/AI labs two quarters ago—Google noting it "still expects to remain in a tight demand-supply environment in Q4 and 2026", Microsoft indicating "demand remains significantly ahead of the capacity [it has] available", and OpenAI's Sam Altman describing "horrible trade-offs… [OpenAI] has better models, but just can't offer them because [it] does not have the capacity"—has been validated through the first half of 2026, with hyperscalers again raising capex outlooks over the past six months. The top four U.S. hyperscalers (GOOG,

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

AMZN, MSFT, META) have collectively guided to ~$700-725 billion of total capex in 2026, up ~75% Y/Y from ~$410 billion in 2025 and the third consecutive year of 60%+ growth, with the overwhelming majority of incremental spend directed at AI/data center infrastructure.

For OpenAI, which has seen its ChatGPT weekly active user base grow from ~100MM in October 2023 to >900MM as of March 2026 (~9x), token processing via its API grow from ~300MM per minute to >6B per minute (~20x) through late 2025, and its annualized revenue run-rate exceed $25 billion as of March 2026 (vs. ~$20B in November 2025), one pivotal solution to addressing the challenge of compute scarcity has been to secure access to vast swaths of compute capacity with key providers of AI compute (Nvidia, AMD, Broadcom) for the next several years. Those agreements remain on track, with the first ~1 GW of deployments on Nvidia's Vera Rubin platform and AMD's Instinct MI450 expected in 2H26, and the first Broadcom co-designed custom accelerators deploying through 2026-2029.

OpenAI is no longer alone in contracting multi-gigawatt, multi-vendor compute. Anthropic, whose run-rate revenue is fast-approaching $50B, up from ~$9B at YE2025, has assembled its own portfolio of capacity agreements spanning all three major silicon avenues: custom ASICs (AWS Trainium via an expanded Amazon partnership, and Google TPUs via agreements with Google and Broadcom), and merchant GPUs (Nvidia and likely AMD)—see Figure 7. Simply aggregating the estimated expenditure associated with OpenAI's capacity agreements with Nvidia, AMD and Broadcom equates to well in excess of $1T spent over the next ~4-5 years, with 50-60% of this spend expected to accrue to OpenAI's compute partners (JPMe). Layering in Anthropic's ~10+ GW of announced commitments (up to 5 GW with Amazon, >1 GW of Google TPUs by 2026 plus multiple incremental GW from 2027, and up to 1 GW initially on Azure/Nvidia) implies several hundred billion dollars of additional model-developer-driven data center capex this cycle.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 7: Recent Compute Capacity Agreements with Key Providers of AI Compute Silicon

Partner	Announced	Stated Capacity	Hardware / Platform	First Deployments	Financial / Structure Notes
OpenAI - NVIDIA	22 Sep 2025	**≥10 GW** of NVIDIA systems ("millions of GPUs")	NVIDIA systems; first wave on **Vera Rubin** platform (successors to GB200/GB300-class)	**2H 2026** (first 1 GW)	NVIDIA **intends to invest up to $100B** in OpenAI, disbursed progressively as each GW deploys
OpenAI - AMD	6 Oct 2025	**6 GW** multi-year, multi-generation	**Instinct MI450** first; subsequent Instinct generations	**2H 2026** (first **1 GW** of MI450)	Multi-year capacity agreement (terms undisclosed publicly)
OpenAI - Broadcom	13 Oct 2025	**10 GW** custom AI accelerators & rack systems	OpenAI-designed **custom accelerators**; Broadcom co-develops and deploys (Ethernet-based systems)	**2H 2026** start; **through 2029**	Multi-year collaboration; financials undisclosed publicly
OpenAI - AWS	Nov 2025	**Not specified**	NVIDIA systems; initially GB200/GB300	**Immediate**; all capacity targeted to be deployed by **YE2026**; **ability to expand further in 2027+**	**$38B** seven-year agreement
OpenAI - AWS	Feb 2026	**~2 GW**	AWS Trainium	**Starting 2027**	Multi-year strategic partnership; **~$100B+** Trainium capacity deal in addition to existing $38B agreement
OpenAI - Cerebras	Apr 2026	**750 MW initial, option for additional 1.25 GW**	Cerebras servers	**Through 2028**	**~$20B+** over three years; warrants for minority stake up to 10% if spending reaches $30B
Anthropic - Google Cloud	Oct 2025	**Up to 1MM TPUs; >1 GW**	Google TPU on GCP	**2026**	Reported value in "**tens of billions of dollars**"
Anthropic - Azure - NVIDIA	Nov 2025	**Up to 1 GW initially** + optional incremental capacity	NVIDIA Grace Blackwell and Vera Rubin systems	**2026**	**$30B** Azure compute purchase commitment
Anthropic - AWS	Apr 2026	**Up to 5 GW**	AWS Trainium + Graviton	**2026**	**$100B+** over 10 years
Anthropic - Broadcom	Apr 2026	**Multiple GW**	Next-gen TPUs	**2026**	**$35B** capacity expansion; Google backstops lease payments; capacity across five US data centers

Source: Company reports.

Figure 8: Estimated Capex Associated with Recent Nvidia, AMD and Broadcom Capacity Agreements ($ in billions)

Compute Platform Partner	Scenario	DC Capex per GW	Compute Partner Revenue per GW	Compute Capex	Networking Capex
NVIDIA	*Low case*	$45	$25	$20	$5
	High case	$55	$35	$28	$7
AMD	*Low case*	$35	$15	$14	$2
	High case	$40	$20	$18	$2
Broadcom	*Low case*	$30	$10	$7	$3
	High case	$40	$20	$14	$6

Source: Company reports, J.P. Morgan Research.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 9: Top 4 U.S. Hyperscalers' Data Center Capex

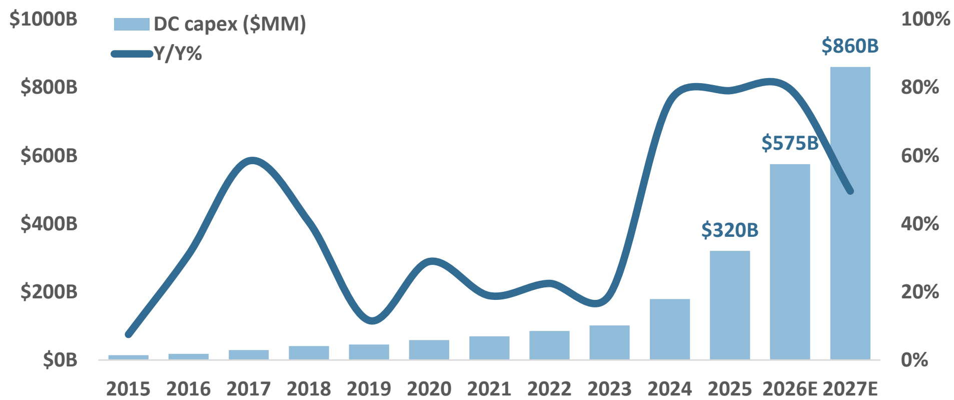

Source: Company reports, J.P. Morgan Research.

Custom AI ASICs: Adoption inflecting on power efficiency, economics and performance

A defining feature of the current phase of the AI buildout is the rising mix of hyperscaler custom silicon ("ASICs" or "XPUs") alongside merchant GPUs. Every major hyperscaler now fields a custom accelerator program—Google TPU (seven generations co-designed with Broadcom since 2014, with Ironwood/v7 ramping now), AWS Trainium (Trainium2 deployed at >1MM-chip scale, Trainium3 ramping in 2026, Trainium4 on the roadmap), Microsoft Maia (Maia 200 ramping), and Meta MTIA (multi-generation roadmap, with Broadcom planning multiple gigawatts of MTIA deployments beginning in 2027). Our data center megawatt model (Figure 10) already embeds this shift: We estimate ASICs account for ~42% of total data center AI chip unit volumes in 2026, rising to ~53% in 2027, with ASIC unit shipments growing almost 3x as fast as merchant GPU shipments this year (+109% Y/Y vs. +39% Y/Y). The trend is visible in supplier financials as well—we estimate Broadcom is positioned to generate $150B+ of AI revenue (ASIC/XPU + AI-related networking) in 2027 (management has already disclosed 2027 backlog of >$100B which we view as conservative messaging), which is more than double the ~$60B we are anticipating for 2026. A similar dynamic has also emerged in the customer base: Anthropic's capacity agreements are anchored predominantly on custom silicon (Trainium and TPU), and even OpenAI's 10 GW Broadcom program is, in effect, based on the assumption that owning the silicon roadmap improves the economics of serving its own workloads.

Figure 10: GPU + ASIC/XPU Accelerator Shipments

AI accelerators units (MM)	2023	2024	2025	2026E	2027E
NVIDIA GPUs	1.7	4.5	6.2	8.9	9.9
AMD GPUs	0.1	0.5	0.7	0.6	1.0
Google TPUs	1.2	1.7	1.8	4.5	8.0
AWS Inferenta + Trainium	0.3	0.8	1.5	1.9	3.3
Meta MTIA	0.0	0.0	0.0	0.3	0.5
Microsoft MAIA	0.0	0.0	0.0	0.2	0.7
Total	**3.3**	**7.4**	**10.1**	**16.3**	**23.3**
Growth Y/Y%		*128%*	*36%*	*62%*	*43%*
GPU % total units	55%	67%	68%	58%	47%
ASIC/XPU % total units	45%	33%	32%	42%	53%

Source: Company reports, J.P. Morgan Research.

Power efficiency is the most strategically important ASIC advantage in a power-constrained buildout. As we detail throughout this report, electric power—not capital

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

—is the binding constraint on data center growth, which makes performance-per-watt the critical currency of AI infrastructure. By dedicating die area to tensor compute and stripping out the general-purpose overhead GPUs carry, ASICs deliver structurally better efficiency on targeted workloads: Google's Ironwood TPU delivers ~2x the performance-per-watt of its prior-generation Trillium (and ~30x the efficiency of its first-generation TPU), and well-optimized TPU/ASIC deployments can achieve on the order of 2x or better performance-per-watt versus contemporary merchant GPUs on inference workloads. The arithmetic matters at scale: Every incremental point of perf/watt directly expands how much revenue-generating compute can be wrung from a fixed grid interconnect—effectively easing the single largest bottleneck in the funding equation laid out in this report.

The economics increasingly favor ASICs for at-scale, well-characterized workloads. Hyperscalers deploying in-house silicon report total cost of ownership (TCO) reductions on the order of 30-40% versus merchant GPU fleets, driven by the removal of the merchant margin stack (Nvidia gross margins remain in the ~70%+ range), cheaper Ethernet-based networking, and power savings. Recent datapoints include AWS positioning Trainium3 UltraServers at ~30% better TCO per FP8 PFLOP versus Nvidia GB300 NVL72 systems, and Microsoft claiming ~30% better performance-per-dollar from Maia 200 versus the latest hardware in its own fleet. These economics flow through our capex math: We estimate custom XPU-based data centers cost ~$30-40B per GW versus $50B+ per GW for Nvidia GPU infrastructure. Amazon CEO Andy Jassy's framing—*"our custom AI silicon offers high performance at significantly lower cost for customers, which is why it's in such hot demand"*—captures the commercial logic, and Anthropic's commitment to run Claude on Trainium for the next decade is the strongest third-party endorsement of ASIC economics to date.

Performance is no longer the gating concern—frontier models are now trained and served on ASICs. Google's Gemini family has been trained on TPUs since inception; Anthropic trains and serves Claude on >1 million Trainium2 chips (Project Rainier, one of the largest compute clusters in the world) alongside TPUs; and Meta intends to train next-generation models on MTIA. Importantly, the workload mix is shifting in the ASICs' favor: Inference (which now accounts for the majority of all AI compute, and growing with reasoning models and agentic workloads) is precisely the high-volume, well-characterized workload where fixed-function silicon excels. Anthropic's explicitly multi-silicon strategy (Trainium, TPU, and Nvidia GPUs, matching workloads to the chips best suited for them) is, in our view, a preview of the steady-state architecture of AI compute: merchant GPUs for flexibility and frontier research, ASICs for scaled training and inference economics.

From Semiconductors to Gigawatts

The biggest limiting factor in growing the global data center footprint remains power. We will discuss that below in much more detail, but the difference between 'bull' and 'base' case estimates broadly for data center growth is generally the constraint around power. Without those constraints, growth above 150 GW over five years appears difficult but plausible. With those constraints, the need for creativity abounds.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 11: Data Center Growth Bear/Bull/Base Case Variance Across Consultants

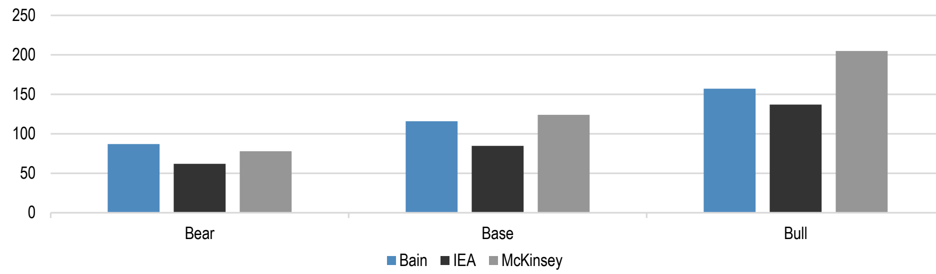

Source: Bain & Co., IEA, Mckinsey & Co and J.P. Morgan estimates.

Increasing power demand per accelerator creates a push at the data center level. The increase in wattage from the Ampere to Blackwell generations of accelerators is almost 7x (Hopper to Blackwell almost 4x). That is helping drive the need for AI factories with very large power supplies to the facility. Clustering will make this even more apparent going forward, with a Vera Rubin cluster of 72 GPUs and 36 CPUs likely to draw as much as 200,000 watts. One other consideraton is that the amount of power consumed by silicon vs. the entirety of the site is shifting, with silicon power consumption likely to move from closer to 60% to well over 70% of power consumption —creating more variability in overall power consumption. To be clear, compute efficiency per watt is increasing. Blackwell setups are potentially 100x more efficient from the standpoint of tokens per second per megawatt, but nonetheless that is a substantial tax on power consumption. Moreover, as noted above, the token economics vs. power and cost of custom silicon are potentially even more attractive going forward.

Figure 12: Wattage & Compute Efficiency per Nvidia GPU Generation

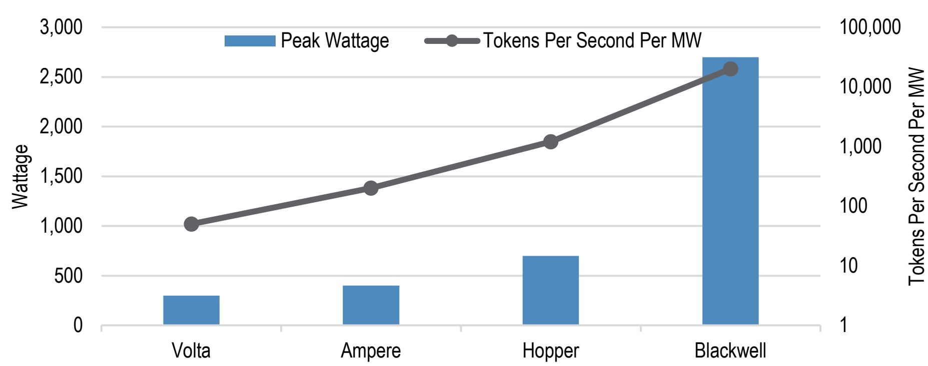

Source: Nvidia and J.P. Morgan estimates.

Overall data center growth rates are still constrained by power. As we have been writing for the last three years, power remains the key constraint for data center growth. Factoring in power constraints, our base case forecast for data center growth is now 138 GW through 2030, slightly above consensus base cases but well below unconstrained bull cases. The key limitations on power growth remain real, but creativity and bring-your-own-power and behind-the-meter solutions have slightly debottlenecked the power situation vs. six months ago.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 13: Data Center Growth and Total Capex

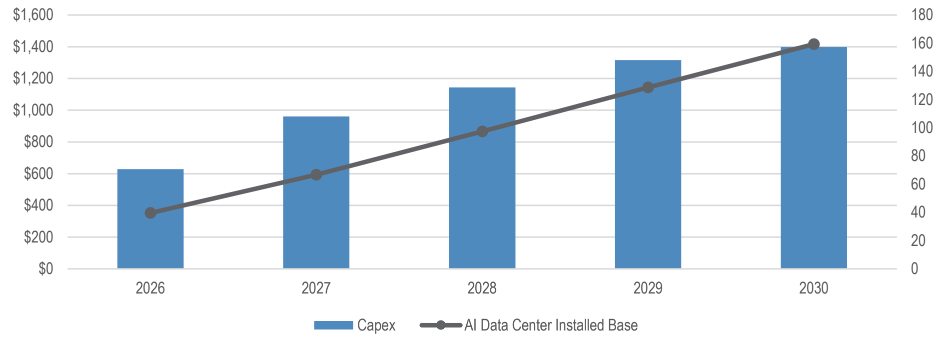

Source: J.P. Morgan estimates.

Data center box and non-silicon equipment capex likely stabilizes before all-in data center capex due to silicon needs (and eventually replacement needs). One fascinating implication of our analysis is that data center box capex of $280 billion in 2026 likely stabilizes in the $580-$590 billion range from 2028-2030. Conversely, spending on AI accelerators in the data centers likely increases from $340 billion in 2026 to over $800 billion in 2030. That has very important implications for the mix of financing going forward—the amount of deals for long-term data center boxes with long-term leases attached will potentially be dwarfed by financing for GPUs and custom ASICs.

Figure 14: Mix of AI Capex Spending Likely to Shift

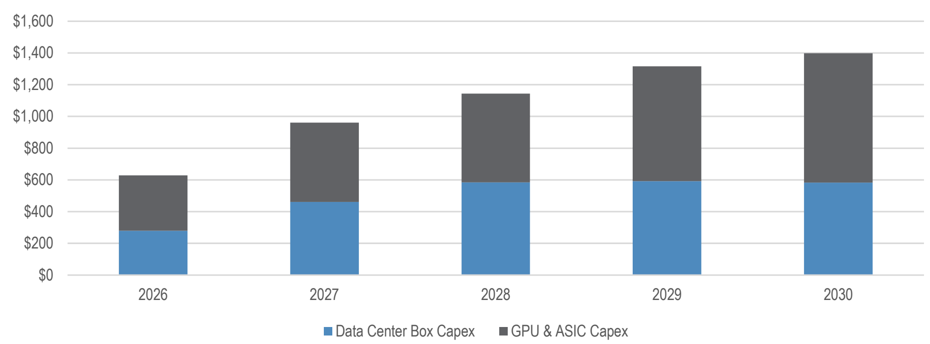

Source: J.P. Morgan estimates.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Hyperscaler Balance Sheets Strong, Capex Higher

Erica Spear
(212) 834-4143
erica.spear@jpmchase.com
J.P. Morgan Securities LLC

The funding mix behind the AI buildout has shifted decisively this year. What began as a capex story funded largely through operating cash flow has evolved into a broader mobilization of the capital markets, with hyperscalers tapping the IG bond market, project finance structures, and, in some cases, common equity to fund data center capacity at a pace that continues to exceed expectations. As a result, hyperscalers have become a structural presence on the IG new issue calendar rather than episodic issuers, with implications for index weightings, spread behavior, and relative value across the sector. **The willingness to raise equity alongside debt also suggests that management teams view the buildout as a long-duration, balance-sheet-defining event rather than a cyclical spending cycle.** In this section, we examine how the five largest spenders are funding their programs and what the capex trajectory implies for FCF generation and balance sheet health.

Year-to-date, the five hyperscalers have raised $107bn of USD senior unsecured debt, supplemented by $62bn issued across five foreign currency markets and $22bn of project finance bonds tied to individual data center campuses. The funding story has also expanded beyond debt, with Alphabet announcing an $85bn equity raise and Oracle launching a $25bn ATM offering, bringing total announced external funding to ~$240bn YTD in the USD market alone. At this point, issuing equity would be prudent for the likes of Amazon, Meta, and Microsoft, in our view. However, we note that META's and MSFT's weak YTD equity performance could give management pause. Against this, combined FY26 capex guidance across the five names exceeds $750bn, and we expect it to approach or exceed $1.1tn in FY27. Even as the combined projected OCF for the five names nears ~$750bn, OCF alone cannot fund a spending program of this size over multiple years, even for some of the most cash-generative companies in the world. As such, we expect debt and equity issuance to remain a recurring part of the funding mix over the next several years.

Foreign currency issuance has also emerged as an attractive funding tool, though only Alphabet and Amazon have used it so far. Alphabet has issued ~$31bn equivalent across five currencies this year, including ~$10bn in euros, ~$7bn in sterling, ~$6bn in CAD, ~$4bn in Swiss francs, and ~$4bn in yen. Amazon has raised a similar amount, including ~$17bn in euros, ~$10bn in CAD, and ~$4bn in Swiss francs. **The shift reflects the reality that hyperscalers are increasingly competing for the same pool of USD IG demand.** Diversifying across currencies relieves pressure on USD spreads, broadens the investor base, and in some cases lowers all-in funding costs once swapped back, particularly in yen and Swiss francs. **Beyond the near-term economics, moving first may provide a durable advantage.** Alphabet and Amazon are establishing reference curves in markets where supply from issuers of this caliber remains limited, supporting tighter concessions and deeper books. By the time Microsoft, Meta, or Oracle follow, those markets may have already absorbed meaningful hyperscaler supply relative to their depth, leaving later entrants with less favorable demand dynamics and pricing benchmarked to curves established by competitors.

Lastly, against a backdrop of capex consistently outrunning OCF for now, off-balance-sheet data center project financing has emerged as a meaningful relief valve for the hyperscaler complex. The structure is broadly consistent across the IG

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

deals we have seen so far, with a developer or SPV owning and constructing the facility while the hyperscaler signs a long-dated lease or take-or-pay arrangement as the anchor tenant, and the project-level debt is secured by those contracted cash flows and the underlying asset rather than appearing as funded debt on the hyperscaler's own balance sheet. Because the commitment surfaces as a lease or purchase obligation rather than a corporate borrowing, the approach preserves reported leverage capacity and spares the company's senior unsecured curve from absorbing the full weight of the buildout, and the paper still has cleared at IG levels because the cash flows ultimately rest on the tenant's contractual commitment. The appeal to the issuer is that it can reach a deeper and distinct pool of infrastructure and structured credit investors without crowding its own primary supply, while shifting construction risk onto the developer, and as the funding gaps widen into 2027 and 2028, we expect this channel to carry a growing share of the incremental capital needed to fund the AI buildout.

Figure 15: Hyperscaler Capex Growth

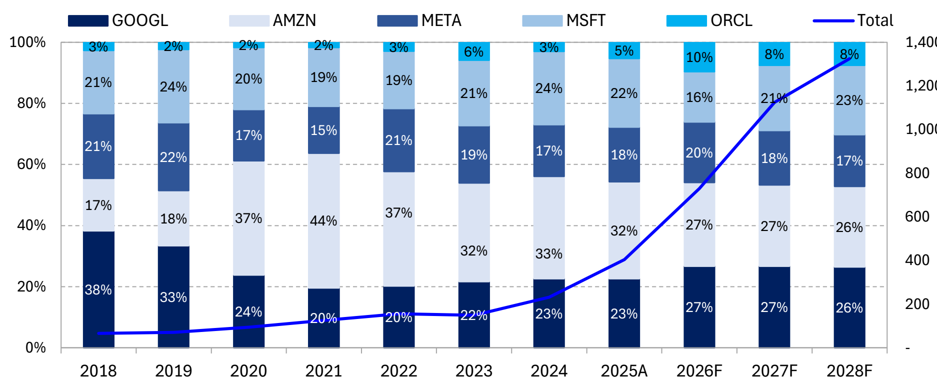

Source: Company reports and J.P. Morgan.

Alphabet

Alphabet is funding this buildout from a position of operating strength. 1Q results showed Cloud revenue accelerating 63% to $20bn with the segment operating margin expanding to 33% (up 16 points y/y), pushing back on the view that AI revenue is inherently lower margin. Cloud backlog nearly doubled q/q to $462bn, aided by the inclusion of TPU hardware sale agreements, creating a new revenue stream beyond capacity rental while positioning Alphabet as a supplier of AI infrastructure. Management continues to frame capital allocation through an ROIC lens, and demand remains strong, with CEO Sundar Pichai noting that Cloud revenue would have been higher absent compute constraints. The quarter also highlighted the expansion of Alphabet's AI strategy beyond the income statement, with the Wiz and Intersect acquisitions, a $40bn commitment to Anthropic, and a growing portfolio of credit backstops supporting third-party data centers. These backstops, which are offered in the form of financial guarantees and credit derivatives, carry maximum potential future exposure of ~$9bn and ~$28bn, respectively. Additionally, upon a default under these backstops, Alphabet retains the right to assume the underlying leases for internal use or to sublease to third parties.

That said, the credit story is increasingly defined by the scale of the spending program. Management increased prior FY26 capex guidance to $185bn at the midpoint in 1Q (up $10bn q/q), and we expect capex to reach as high as ~$300bn in FY27 and ~$350bn in FY28 as management expects FY27 capex "to significantly increase

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

compared to 2026." Based on current OCF estimates, capex is expected to consume nearly all internally-generated cash across the forecast horizon, with FY28 spending near breakeven before dividends. That dynamic helps explain the breadth of this year's funding response. Alphabet has issued ~$51bn of senior notes YTD, including $20bn in dollars and the equivalent of $31bn across five foreign currencies, while also announcing an $85bn equity raise, of which ~$30bn is earmarked for tax withholding on stock-based compensation. **Notably, this year's funding activity runs well ahead of this year's needs. With OCF expected to cover FY26 capex and dividends almost entirely, the combined debt and net equity proceeds far exceed any near-term shortfall. This indicates to us that Alphabet is aggressively pre-funding the FY27 and FY28 ramp while both debt and equity markets remain receptive, and the decision to accept equity dilution alongside new debt underscores how deliberately management is locking in capital ahead of need.** Long-term debt has increased from ~$47bn to ~$78bn in just over a quarter, and we expect that balance to continue growing over the next several years. While we remain comfortable with the credit profile given the scale of cash generation and the ~$127bn liquidity position, Alphabet has crossed the Rubicon from a company that once issued debt opportunistically into one with a structural funding requirement, and its growing off-balance-sheet commitments through leases, backstops, and contingent funding agreements warrant attention alongside senior debt.

Table 1: Alphabet Model Snapshot

Alphabet (GOOGL)	FY20A	FY21A	FY22A	FY23A	FY24A	FY25A	1Q26A (LTM)	FY26E	FY27E	FY28E
Google Cloud Revenue	13,059	19,206	26,280	33,088	43,229	58,705	66,473	102,246	168,632	252,948
Total Revenue	182,527	257,637	282,836	307,394	350,018	402,836	422,498	492,710	604,471	567,328
Y/Y % Δ	*12.8%*	*41.2%*	*9.8%*	*8.7%*	*13.9%*	*15.1%*	*17.5%*	*22.3%*	*22.7%*	*-6.1%*
Operating Margin	*22.6%*	*30.6%*	*26.5%*	*27.4%*	*32.1%*	*32.0%*	*32.7%*	*34.7%*	*35.4%*	*35.8%*
Adj. EBITDA	67,912	106,531	110,132	120,452	150,486	178,585	190,905	236,076	308,382	397,140
Operating Cash Flow	**65,124**	**91,652**	**91,495**	**101,746**	**125,299**	**164,713**	**174,353**	**206,678**	**272,947**	**351,856**
Capital Expenditures	**(22,281)**	**(24,640)**	**(31,485)**	**(32,251)**	**(52,535)**	**(91,447)**	**(109,924)**	**(195,000)**	**(300,000)**	**(350,000)**
Capex Intensity	*12.2%*	*9.6%*	*11.1%*	*10.5%*	*15.0%*	*22.7%*	*26.0%*	*39.6%*	*49.6%*	*61.7%*
Free Cash Flow	**42,843**	**67,012**	**60,010**	**69,495**	**72,764**	**73,266**	**64,429**	**11,678**	**(27,053)**	**1,856**
Free Cash Flow Conversion	*0.0%*	*-1.1%*	*0.0%*	*0.0%*	*-0.8%*	*-0.6%*	*0.0%*	*-1.0%*	*-0.4%*	*-0.7%*
Dividend Payments	-	-	-	-	(7,363)	(10,049)	(10,157)	(10,684)	(10,943)	(11,048)
Share Repurchases	(31,149)	(50,274)	(59,296)	(61,504)	(62,222)	(45,709)	(30,641)	-	-	-
Senior Maturities	-	(1,000)	-	-	(1,000)	(1,000)	-	(2,000)	(1,000)	(2,500)
Funding Gap	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**1,006**	**38,996**	**11,692**
Gross Debt	**13,932**	**14,817**	**14,701**	**13,523**	**10,883**	**46,547**	**77,501**	**97,874**	**107,874**	**122,874**
Cash & Marketable Securities	(136,694)	(139,649)	(113,762)	(110,916)	(95,657)	(126,843)	(126,840)	(246,914)	(201,918)	(188,226)
Net Debt	**(122,762)**	**(124,832)**	**(99,061)**	**(97,393)**	**(84,774)**	**(80,296)**	**(49,339)**	**(149,040)**	**(94,044)**	**(65,352)**
Term Loan Draw	-	-	-	-	-	-	-	-	-	-
USD Issuance	10,000	-	-	-	-	22,500	20,000	20,000	10,000	15,000
XCCY Issuance	-	-	-	-	-	15,330	11,210	31,327	-	-
Gross Debt Issuance	**10,000**	**-**	**-**	**-**	**-**	**37,830**	**31,210**	**51,327**	**10,000**	**15,000**
Gross Equity Issuance	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**84,750**	**-**	**-**
Total Funded	**10,000**	**-**	**-**	**-**	**-**	**37,830**	**31,210**	**136,077**	**10,000**	**15,000**
Gross Leverage	***0.2x***	***0.1x***	***0.1x***	***0.1x***	***0.1x***	***0.3x***	***0.4x***	***0.4x***	***0.3x***	***0.3x***
Net Leverage	***-1.8x***	***-1.2x***	***-0.9x***	***-0.8x***	***-0.6x***	***-0.4x***	***-0.3x***	***-0.6x***	***-0.3x***	***-0.2x***

Amounts in USD millions.
Source: J.P. Morgan Equity Research, Bloomberg Finance L.P., S&P Global.

Amazon

Amazon's recent results show the AWS reacceleration thesis playing out, with 1Q cloud revenue growth of 28%, the fastest pace in 15 quarters, on a $150bn run rate and up $2bn q/q, the largest 4Q to 1Q AWS revenue increase ever. Management highlighted a $364bn backlog that excludes the $100bn+ Anthropic expansion announced after quarter-end. Custom silicon has become central to the story, with more than $225bn of Trainium revenue commitments from Anthropic and OpenAI, while management argues the program can save tens of billions in annual capex and improve margins relative to merchant GPU reliance. Amazon has also become increasingly intertwined with the

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

model labs, combining its $15bn stake in OpenAI and additional $35bn equity commitment letter with its Anthropic investments, which could total up to ~$33bn and generated ~$17bn of mark-to-market gains in the quarter. The retail business continues to provide support, with unit growth of 15% outpacing fulfillment cost growth and consolidated operating margin reaching a record 13.1%, **giving the AI buildout a much larger earnings base than during the first AWS investment cycle.**

The funding math is where the story tightens. The firm has guided to capex of ~ $200bn this year, and we currently expect it to grow to ~$300bn in FY27 and ~$350bn in FY28. When compared to OCF estimates over the same period, this leaves Amazon outspending its cash generation by ~$30bn this year, over ~$50bn in FY27, and ~$10bn in FY28, representing a cumulative gap approaching ~$97bn, including senior maturities, though before accounting for the remaining OpenAI equity commitment and the Anthropic facility. TTM FCF has already fallen to ~$1bn from ~$26bn a year ago, and the company has responded with ~$67bn of bond issuance across four currencies in 1H. **Additionally, last week, Amazon added a $17.5bn three-year delayed draw term loan** with a September 30th commitment expiry. The facility provides low-cost liquidity that can be drawn as capex needs arise, rather than requiring another benchmark bond deal. Given Amazon's history of utilizing shorter-duration term loans to supplement bond issuance during heavy investment cycles, we expect the firm could supplement with similar funding options in the future. **Taken together, Amazon has arranged ~$85bn of new funding capacity in less than six months, without having to turn to the equity market. Against a full-year funding gap of ~$33bn, this aggression suggests to us that, like Alphabet, Amazon is pre-funding a portion of its FY27 capex ramp while markets remain receptive.** With no dividend, limited buybacks, and ~$143bn of cash on hand, the firm retains substantial financial flexibility and liquidity. However, the pace of issuance makes clear that Amazon is now a structural and recurring presence in primary markets, spreading supply across a wide range of currencies to broaden its investor base rather than concentrating issuance on its USD curve, and suggests that it will likely diversify toward other funding sources like term loans, project finance, and possibly even equity issuance in the medium-term.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Table 2: Amazon Model Snapshot

Amazon (AMZN)	FY20A	FY21A	FY22A	FY23A	FY24A	FY25A	1Q26A (LTM)	FY26E	FY27E	FY28E
AWS Revenue	43,370	62,202	80,096	90,757	107,556	128,725	137,045	170,599	221,779	283,877
Total Revenue	386,064	469,822	513,983	574,785	637,959	716,924	742,776	823,986	934,188	1,060,583
Y/Y % Δ	*37.6%*	*21.7%*	*9.4%*	*11.8%*	*11.0%*	*12.4%*	*14.2%*	*14.9%*	*13.4%*	*13.5%*
Operating Margin	*8.3%*	*8.0%*	*6.4%*	*10.7%*	*14.3%*	*14.5%*	*14.8%*	*15.3%*	*17.0%*	*18.2%*
Adj. EBITDA	57,213	72,131	75,053	110,305	144,162	169,837	180,449	222,220	289,765	362,557
Operating Cash Flow	**66,064**	**46,327**	**46,752**	**84,946**	**115,877**	**139,514**	**148,531**	**169,701**	**255,005**	**351,385**
Capital Expenditures	**(35,046)**	**(55,396)**	**(58,321)**	**(48,133)**	**(77,658)**	**(128,320)**	**(147,299)**	**(200,000)**	**(300,000)**	**(350,000)**
Capex Intensity	*9.1%*	*11.8%*	*11.3%*	*8.4%*	*12.2%*	*17.9%*	*19.8%*	*24.3%*	*32.1%*	*33.0%*
Free Cash Flow	**31,018**	**(9,069)**	**(11,569)**	**36,813**	**38,219**	**11,194**	**1,232**	**(30,299)**	**(44,995)**	**1,385**
Free Cash Flow Conversion	*47.0%*	*-19.6%*	*-24.7%*	*43.3%*	*33.0%*	*8.0%*	*0.8%*	*-17.9%*	*-17.6%*	*0.4%*
Dividend Payments	-	-	-	-	-	-	-	-	-	-
Share Repurchases	-	-	(6,000)	-	-	-	-	-	-	-
Senior Maturities	(1,000)	(1,000)	(1,250)	(3,000)	(8,500)	(4,872)	-	(2,750)	(8,750)	(11,750)
Funding Gap	**-**	**10,069**	**18,819**	**-**	**-**	**-**	**-**	**33,049**	**53,745**	**10,365**
Gross Debt	**31,816**	**48,744**	**67,150**	**58,314**	**52,623**	**65,648**	**119,074**	**153,275**	**183,275**	**211,275**
Cash & Marketable Securities	(84,396)	(96,049)	(70,026)	(86,780)	(101,202)	(123,029)	(143,089)	(174,857)	(151,112)	(165,747)
Net Debt	**(52,580)**	**(47,305)**	**(2,876)**	**(28,466)**	**(48,579)**	**(57,381)**	**(24,015)**	**(21,582)**	**32,163**	**45,528**
Term Loan Draw	-	-	-	-	-	-	-	17,500	-	-
USD Issuance	10,000	18,500	21,000	-	-	15,000	37,000	37,000	20,000	15,000
XCCY Issuance	-	-	-	-	-	-	16,777	30,377	10,000	10,000
Gross Debt Issuance	**10,000**	**18,500**	**21,000**	**-**	**-**	**15,000**	**53,777**	**84,877**	**30,000**	**25,000**
Gross Equity Issuance	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**
Total Funded	**10,000**	**18,500**	**21,000**	**-**	**-**	**15,000**	**53,777**	**84,877**	**30,000**	**25,000**
Gross Leverage	***0.6x***	***0.7x***	***0.9x***	***0.5x***	***0.4x***	***0.4x***	***0.7x***	***0.7x***	***0.6x***	***0.6x***
Net Leverage	***-0.9x***	***-0.7x***	***0.0x***	***-0.3x***	***-0.3x***	***-0.3x***	***-0.1x***	***-0.1x***	***0.1x***	***0.1x***

Amounts in USD millions.
Source: J.P. Morgan Equity Research, Bloomberg Finance L.P., S&P Global.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Meta Platforms

Meta's 1Q was operationally strong, with revenue up 33%, driven by ad impression growth of 19% and ad pricing growth of 12% as a result of increased AI engagement. Engagement improvements continue to compound, including a 10% increase in Instagram time spent and the largest quarterly gain in Facebook video time in four years. The launch of Muse Spark and the rebuilt Meta AI marked the first major outputs from Meta Superintelligence Labs while monetization initiatives continue to scale, with the value optimization suite exceeding a $20bn annual revenue run rate, partnership ads reaching a $10bn run rate, more than doubling y/y, and ads expanding across Threads and WhatsApp Status. **The challenge is that the products absorbing much of the incremental investment remain largely pre-revenue.** Business AI is free for most customers, personal agent monetization remains largely conceptual, and **CEO Zuckerberg reiterated the company's longstanding strategy of building scale before monetization.** When asked about ROIC, management emphasized model quality and product adoption rather than financial milestones, highlighting the growing gap between near-term spending and future returns. Legal risks also remain relevant, with management warning that youth-related litigation could result in material losses following adverse verdicts in several recent cases.

Set against the operating strength, the balance sheet and cash flow trajectory remains very much in focus. Management raised the midpoint of FY26 capex guidance by $10bn to $135bn, **primarily reflecting higher component and memory costs,** and we forecast spending of ~$185bn in FY27 and ~$225bn in FY28. Against expected OCF over the same period, Meta's FCF will turn negative this year, and that trend is expected to deepen materially in the near-term, particularly in FY27. 1Q results already reflected the pressure, with $12.4bn of FCF against $19.8bn of capex. **As a result, Meta is facing a cumulative estimated funding gap of ~$97bn over the next three years, and the firm's funding strategy has yet to fully reveal itself.** Additionally, the obligations building quietly beyond the income statement are arguably more significant than the capex itself, as multi-year cloud agreements and infrastructure purchase commitments drove a $107bn increase in contractual commitments during the quarter to $238bn, with another $24bn added in April, alongside $183bn of leases yet to commence. Meta has also been the most aggressive hyperscaler in shifting spend off balance sheet, including a 20% stake in the Hyperion data center venture backed by residual value guarantees of ~$28bn. The funding response is already underway, though the direction remains uncertain. **Meta returned to the bond market with a $25bn deal following 1Q earnings,** increasing gross debt to ~$84bn, though management indicated it does not intend to issue debt in consecutive quarters. Additionally, the *Financial Times* recently reported that Meta is considering an equity raise worth "tens of billions of dollars," though management dismissed the report and stated no banks had been hired. With buybacks paused in 1Q and ~$81bn of liquid cash on hand, we expect future funding to increase with periodic benchmark debt issuance, continued SPV and project finance structures, and potentially equity. The growing share of obligations, both on- and off-balance sheet, complicates the analysis, but with a franchise generating north of $130bn in annual operating cash flow at a 41% operating margin, Meta has ample flexibility to navigate the spending ramp from here in our view.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Table 3: Meta Platforms Model Snapshot

Meta Platforms (META)	FY20A	FY21A	FY22A	FY23A	FY24A	FY25A	1Q26A (LTM)	FY26E	FY27E	FY28E
Advertising Revenue	84,169	114,934	113,642	131,948	160,633	196,174	209,806	247,891	297,193	344,466
Total Revenue	85,965	117,929	116,609	134,902	164,501	200,966	214,962	254,057	304,893	353,878
Y/Y % Δ	*20.8%*	*37.2%*	*-1.1%*	*15.7%*	*21.9%*	*22.2%*	*26.2%*	*26.4%*	*20.0%*	*16.1%*
Operating Margin	*38.0%*	*39.6%*	*24.8%*	*34.7%*	*42.2%*	*41.4%*	*41.3%*	*34.7%*	*32.2%*	*32.0%*
Adj. EBITDA	46,246	64,089	54,470	77,553	102,276	122,684	132,002	147,415	181,280	218,376
Operating Cash Flow	**38,747**	**57,683**	**50,475**	**71,113**	**91,328**	**115,800**	**123,993**	**139,668**	**174,216**	**213,547**
Capital Expenditures	**(15,115)**	**(18,567)**	**(31,431)**	**(27,266)**	**(37,256)**	**(69,691)**	**(75,747)**	**(145,000)**	**(200,000)**	**(225,000)**
Capex Intensity	*17.6%*	*15.7%*	*27.0%*	*20.2%*	*22.6%*	*34.7%*	*35.2%*	*57.1%*	*65.6%*	*63.6%*
Free Cash Flow	**23,632**	**39,116**	**19,044**	**43,847**	**54,072**	**46,109**	**48,246**	**(5,332)**	**(25,784)**	**(11,453)**
Free Cash Flow Conversion	*61.0%*	*67.8%*	*37.7%*	*61.7%*	*59.2%*	*39.8%*	*38.9%*	*-3.8%*	*-14.8%*	*-5.4%*
Dividend Payments	-	-	-	-			(5,282)	(5,331)	(5,395)	(5,401)
Share Repurchases	(6,272)	(44,537)	(27,956)	(19,774)	(30,125)	(26,248)	-	-	-	-
Senior Maturities	-	-	-	-	-	-	-	-	(2,750)	(1,500)
Funding Gap	**-**	**5,421**	**8,912**	**-**	**-**	**-**	**-**	**10,663**	**33,929**	**18,354**
Gross Debt	**-**	**-**	**11,000**	**19,500**	**30,000**	**60,000**	**60,000**	**85,000**	**97,250**	**95,750**
Cash & Marketable Securities	(61,954)	(47,998)	(40,738)	(65,403)	(77,815)	(81,592)	(81,180)	(95,929)	(107,000)	(88,646)
Net Debt	**(61,954)**	**(47,998)**	**(29,738)**	**(45,903)**	**(47,815)**	**(21,592)**	**(21,180)**	**(10,929)**	**(9,750)**	**7,104**
Term Loan Draw	-	-	-	-	-	-	-	-	-	-
USD Issuance	-	-	11,000	8,500	10,500	30,000	-	25,000	15,000	-
XCCY Issuance	-	-	-	-	-	-	-	-	-	-
Gross Debt Issuance	**-**	**-**	**11,000**	**8,500**	**10,500**	**30,000**	**-**	**25,000**	**15,000**	**-**
Gross Equity Issuance	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**30,000**	**-**
Total Funded	**-**	**-**	**11,000**	**8,500**	**10,500**	**30,000**	**-**	**25,000**	**45,000**	**-**
Gross Leverage	***0.0x***	***0.0x***	***0.2x***	***0.3x***	***0.3x***	***0.5x***	***0.5x***	***0.6x***	***0.5x***	***0.4x***
Net Leverage	***-1.3x***	***-0.7x***	***-0.5x***	***-0.6x***	***-0.5x***	***-0.2x***	***-0.2x***	***-0.1x***	***-0.1x***	***0.0x***

Amounts in USD millions.
Source: J.P. Morgan Equity Research, Bloomberg Finance L.P., S&P Global.

Microsoft

Among the five hyperscaler issuers discussed, Microsoft is the only name that has so far funded its AI buildout entirely with internal cash flow, though, given the magnitude of the numbers discussed at F3Q, we expect that will soon change. In F3Q, Microsoft Cloud revenue reached $55bn, up 29% y/y; the AI business surpassed a $37bn revenue run rate; and Azure grew 40%. However, management noted that Azure growth benefited from capacity coming online earlier in the quarter and cautioned that growth rates will continue to fluctuate based on capacity timing and contract mix, making it difficult to assess the durability of the current pace once supply constraints ease. The headline commercial RPO of $627bn, up 99% y/y, also warrants context, as much of the increase reflects Azure commitments from OpenAI, while growth excluding OpenAI was 26%, closer to historical levels. M365 Copilot has surpassed 20mn paid seats, with seat additions up 250% y/y, but adoption remains modest relative to the broader M365 installed base. At the same time, the shift toward consumption-based pricing across Copilot, GitHub, and Dynamics introduces greater revenue variability into businesses that have historically been valued for their predictability. Management also expects capacity constraints to persist through at least CY26 and has repeatedly increased spending expectations after initially forecasting slower capex growth, most recently adding ~$25bn of component cost inflation to the CY26 outlook.

On the numbers, Microsoft appears to be ~1 year behind Alphabet and Amazon on the funding curve and approaching a similar inflection point. Capex, which includes cash payments for finance leases per Microsoft's definition, has increased from ~$88bn in FY25 to an expected ~$140bn in FY26; a $190bn CY26 outlook implies a FY27 run-rate of at least ~$240bn in FY27; and we model ~$300bn capex in FY28. OCF still covers current spending levels, but estimates for FY27 and FY28 imply that the funding gap expands to ~$90bn and ~$130bn once dividends and buybacks are accounted for, though we would expect the firm to pause buybacks as spending continues to accelerate. **However, unlike Alphabet and Amazon, Microsoft has not pre-funded this ramp. The company has not tapped the IG bond market in over five years, issuing no**

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

debt or equity during the buildout; remains a net debt repayer with outstanding balances reduced to ~$46bn; and has instead relied on finance leases, which increased to ~$63bn in F3Q, alongside another ~$197bn of data center leases not yet commenced through FY31. That approach is already increasing agency-adjusted leverage ahead of any return to the bond market, and it likely will be unable to absorb a funding need of this magnitude on its own. **In our view, the math points to sizable new supply from Microsoft in C2H26 and beyond.** While its rare AAA rating and modest debt load provide substantial capacity, the ultimate size and mix of the funding program remains the key variable. We often find it useful to note that Microsoft acquired Activision for ~$70bn in cash in 2023 and funded the transaction entirely with its commercial paper program and cash on hand, **reinforcing that the firm has historically been very selective about its use of permanent debt financing.** A strategy that pre-funds the FY28 step-up towards $300bn of capex, similar to Alphabet's approach, would push issuance well beyond the immediate funding gap and into territory that has not been tested for a credit the market has long treated as effectively risk-free.

Table 4: Microsoft Model Snapshot

Microsoft (MSFT)	FY20A	FY21A	FY22A	FY23A	FY24A	FY25A	F3Q26A (LTM)	FY26E	FY27E	FY28E
Intelligent Cloud Revenue	48,251	60,080	74,965	72,944	87,464	106,265	114,195	136,585	174,760	212,227
Total Revenue	143,015	168,088	198,270	211,915	245,122	281,724	294,544	329,112	381,373	441,303
Y/Y % Δ	*13.6%*	*17.5%*	*18.0%*	*6.9%*	*15.7%*	*14.9%*	*4.6%*	*16.8%*	*15.9%*	*15.7%*
Operating Margin	*37.0%*	*41.6%*	*42.1%*	*42.3%*	*44.6%*	*45.6%*	*46.8%*	*46.6%*	*46.7%*	*46.6%*
Adj. EBITDA	72,591	89,061	107,946	114,509	143,722	161,155	148,957	174,023	205,831	249,769
Operating Cash Flow	**60,675**	**76,740**	**89,035**	**87,582**	**118,548**	**136,162**	**170,141**	**182,940**	**217,527**	**259,409**
Capital Expenditures	**(15,441)**	**(20,622)**	**(23,886)**	**(28,107)**	**(44,477)**	**(64,551)**	**(97,225)**	**(120,000)**	**(240,000)**	**(300,000)**
Capex Intensity	*10.8%*	*12.3%*	*12.0%*	*13.3%*	*18.1%*	*22.9%*	*33.0%*	*36.5%*	*62.9%*	*68.0%*
Free Cash Flow	**45,234**	**56,118**	**65,149**	**59,475**	**74,071**	**71,611**	**72,916**	**62,940**	**(22,473)**	**(40,591)**
Free Cash Flow Conversion	*74.6%*	*73.1%*	*73.2%*	*67.9%*	*62.5%*	*52.6%*	*42.9%*	*34.4%*	*-10.3%*	*-15.6%*
Dividend Payments	(15,137)	(16,521)	(18,135)	(19,800)	(21,771)	(24,082)	(25,856)	(26,009)	(28,089)	(30,336)
Share Repurchases	(22,968)	(27,385)	(32,696)	(22,245)	(17,254)	(18,420)	(22,238)	(21,000)	-	-
Senior Maturities	(5,500)	(3,750)	(7,750)	(2,750)	(5,250)	(2,250)	-	(3,000)	(9,080)	-
Funding Gap	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**59,642**	**70,927**
Gross Debt	**63,327**	**58,146**	**49,781**	**47,237**	**49,381**	**40,152**	**40,262**	**37,152**	**88,072**	**143,072**
Cash & Marketable Securities	(135,627)	(130,334)	(104,754)	(111,262)	(75,543)	(94,565)	(78,272)	(107,496)	(182,854)	(166,927)
Net Debt	**(72,300)**	**(72,188)**	**(54,973)**	**(64,025)**	**(26,162)**	**(54,413)**	**(38,010)**	**(70,344)**	**(94,782)**	**(23,855)**
Term Loan Draw	-	-	-	-	-	-	-	-	-	-
USD Issuance	10,270	8,185	-	-	3,344	-	-	-	40,000	40,000
XCCY Issuance	-	-	-	-	-	-	-	-	20,000	15,000
Gross Debt Issuance	**10,270**	**8,185**	**-**	**-**	**3,344**	**-**	**-**	**-**	**60,000**	**55,000**
Gross Equity Issuance	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**-**	**75,000**	**-**
Total Funded	**10,270**	**8,185**	**-**	**-**	**3,344**	**-**	**-**	**-**	**135,000**	**55,000**
Gross Leverage	***0.9x***	***0.7x***	***0.5x***	***0.4x***	***0.3x***	***0.2x***	***0.3x***	***0.2x***	***0.4x***	***0.6x***
Net Leverage	***-1.0x***	***-0.8x***	***-0.5x***	***-0.6x***	***-0.2x***	***-0.3x***	***-0.3x***	***-0.4x***	***-0.5x***	***-0.1x***

Amounts in USD millions. Note: Microsoft's fiscal year ends in June.
Source: J.P. Morgan Equity Research, Bloomberg Finance L.P., S&P Global.

Oracle

Oracle closed out FY26 with a strong quarter that went some way toward steadying a story that has drawn more scrutiny than any other in the group. F4Q revenue of ~$19bn grew 21% y/y, driven by ~$6bn OCI revenue, up 93% and accelerating from 84% in F3Q, while cloud applications maintained double-digit growth at 10% and SaaS deferred revenue increased 16%. RPO grew by $85bn q/q to $638bn, supported by $67bn of AI infrastructure contracts signed during the quarter and four customers each committing more than $8bn, an important diversification point for a company that has faced concentration concerns since the OpenAI announcements last year. Operating metrics also continue to improve, with 1.2GW delivered in FY26, nearly another GW expected in F1Q27, GPU utilization of 97.5%, and multicloud database revenue up 404%. Management reaffirmed FY27 revenue guidance of $90bn, raised adjusted EPS guidance, and reiterated its FY30 framework of 31% revenue CAGR. New

CFO Hilary Maxson emphasized disciplined capital allocation, balance sheet strength, and preservation of the firm's IG ratings, while framing steady-state project-level ROIC in the high 20% range. Execution remains the key variable, but, one year into the buildout, the delivery record has held up better than many expected.

The funding picture is improving in composition even as Oracle remains the most reliant on external funding among the hyperscaler cohort. FY26 capex of ~$56bn against OCF of ~$32bn resulted in negative FCF of ~$24bn, though that figure increasingly overstates the true funding requirement. **A growing share of large AI contracts are structured as customer prepay or bring-your-own-hardware arrangements,** now totaling $75bn and representing the majority of the $67bn signed in the quarter. Management indicated these contracts carry margins at or above the rest of the business, and the distinction led Oracle to introduce a new net capex metric, which totaled ~$48bn in FY26 and is expected to reach ~$70bn in FY27, with reported capex running ~$20-25bn higher. Against the consensus OCF estimates for FY27, that still implies a funding gap of ~$37bn before dividends. Based on our estimates of $85bn and $100bn of net capex in FY28 and FY29 against consensus OCF, the gap narrows but does not close. Oracle raised $43bn of debt and $5bn of mandatory convertible preferreds in FY26 and **plans ~$40bn of combined debt and equity funding in FY27,** including the ~$20bn remaining on the previously-announced ATM program, with no incremental debt issuance planned in CY26. **Notably, Oracle was the first hyperscaler to incorporate equity funding into its capital plan, more out of necessity than desire, though execution has been limited to date, with just $5bn of the $25bn plan executed through May.** We expect that equity will need to remain a recurring funding source in the near-term in order to keep ratings and spreads in check against an already massive $130bn gross debt balance and lease obligations that have continued to grow rapidly alongside debt, with $261bn of lease commitments not yet reflected on the balance sheet. The credit remains supported by $638bn of contracted revenue visibility and improving capital efficiency from prepay structures, but we view Oracle as the clearest example where the trajectory depends on the revenue ramp arriving on schedule.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Table 5: Oracle Model Snapshot

Oracle (ORCL)	FY20A	FY21A	FY22A	FY23A	FY24A	FY25A	FY26A	FY27E	FY28E	FY29E
Cloud & Software Revenue	34,098	34,098	36,052	41,087	44,464	49,231	58,530	80,214	123,785	174,968
Total Revenue	40,479	40,479	42,441	49,955	52,961	57,399	67,357	90,000	130,430	183,332
Y/Y % Δ	*3.6%*	*3.6%*	*4.8%*	*17.7%*	*6.0%*	*8.4%*	*17.3%*	*56.8%*	*44.9%*	*40.6%*
Operating Margin	*46.9%*	*46.9%*	*46.2%*	*41.8%*	*43.5%*	*43.6%*	*42.9%*	*40.1%*	*36.6%*	*36.8%*
Adj. EBITDA	20,537	20,537	21,565	23,428	26,184	28,900	36,548	50,370	72,709	99,068
Operating Cash Flow	**15,887**	**15,887**	**9,539**	**17,165**	**18,673**	**20,821**	**31,977**	**43,759**	**64,165**	**86,463**
Capital Expenditures	**(2,135)**	**(2,135)**	**(4,511)**	**(8,695)**	**(6,866)**	**(21,215)**	**(55,663)**	**(70,000)**	**(85,000)**	**(100,000)**
Capex Intensity	*5.3%*	*5.3%*	*10.6%*	*17.4%*	*13.0%*	*37.0%*	*82.6%*	*77.8%*	*65.2%*	*54.5%*
Free Cash Flow	**13,752**	**13,752**	**5,028**	**8,470**	**11,807**	**(394)**	**(23,686)**	**(26,241)**	**(20,835)**	**(13,537)**
Free Cash Flow Conversion	*86.6%*	*86.6%*	*52.7%*	*49.3%*	*63.2%*	*-1.9%*	*-74.1%*	*-60.0%*	*-32.5%*	*-15.7%*
Dividend Payments	(3,063)	(3,063)	(3,457)	(3,668)	(4,391)	(4,743)	(5,787)	(5,868)	(5,989)	(6,114)
Share Repurchases	(20,934)	(20,934)	(16,248)	(1,300)	(1,202)	(600)	(95)	-	-	-
Senior Maturities	(2,250)	(2,250)	(8,250)	(3,750)	(3,500)	(8,000)	(4,500)	(5,250)	(5,500)	(5,500)
Funding Gap	**12,495**	**12,495**	**22,927**	**248**	**-**	**13,737**	**34,068**	**37,359**	**32,324**	**25,151**
Gross Debt	**84,245**	**84,245**	**75,859**	**90,481**	**86,869**	**92,568**	**129,541**	**144,291**	**163,791**	**178,291**
Cash & Marketable Securities	(46,554)	(46,554)	(21,902)	(10,187)	(10,661)	(11,203)	(31,894)	(29,253)	(26,577)	(21,728)
Net Debt	**37,691**	**37,691**	**53,957**	**80,294**	**76,208**	**81,365**	**97,647**	**115,038**	**137,214**	**156,563**
Term Loan Draw	-	-	-	-	-	-	-	-	10,000	-
USD Issuance	20,000	15,000	-	12,250	-	14,000	43,000	20,000	20,000	10,000
XCCY Issuance	-	-	-	-	-	-	-	-	-	-
Gross Debt Issuance	**20,000**	**15,000**	**-**	**12,250**	**-**	**14,000**	**43,000**	**20,000**	**25,000**	**20,000**
Gross Equity Issuance	**-**	**-**	**-**	**-**	**-**	**-**	**5,000**	**20,000**	**10,000**	**10,000**
Total Funded	**20,000**	**15,000**	**-**	**12,250**	**-**	**14,000**	**48,000**	**40,000**	**35,000**	**30,000**
Gross Leverage	***4.1x***	***4.1x***	***3.5x***	***3.9x***	***3.3x***	***3.2x***	***3.5x***	***2.9x***	***2.3x***	***1.8x***
Net Leverage	***1.8x***	***1.8x***	***2.5x***	***3.4x***	***2.9x***	***2.8x***	***2.7x***	***2.3x***	***1.9x***	***1.6x***

Amounts in USD millions. Note: Oracle's fiscal year ends in May.
Source: J.P. Morgan Equity Research, Bloomberg Finance L.P., S&P Global.

Power Bottlenecks Remain Critical

Kevin Kwan
(212) 270-9455
kevin.l.kwan@jpmorgan.com
J.P. Morgan Securities LLC

Jeremy Tonet
(212) 622-4915
jeremy.b.tonet@jpmorgan.com
J.P. Morgan Securities LLC

Power constraints continue to evolve as hyperscalers consider workarounds in their push for "speed-to-power." One of the early themes of the data center boom was the shift of the binding constraint from compute towards power, driven primarily by a grid interconnection queue of up to 7 years, a major mismatch compared to the 2- to 3-year data center planning process. However, some of the constraints have been moving downstream, including distribution, regulatory and cost-sharing pushback from states, as well as public concern around affordability. New Jersey, Pennsylvania, Virginia, and Maryland are included among a handful of states where legislators have been particularly outspoken on the need for hyperscalers to pay their fair share for generation. In turn, **hyperscalers have quickly pivoted to a variety of grid workarounds, including behind-the-meter (BTM) gas, co-located generation, flexible interconnects, as well as bring-your-own generation (BYOG)**. While the number of BTM arrangements has increased for co-located projects, this type of arrangement does not completely alleviate regulatory scrutiny if existing capacity is taken off the grid. We saw this happen when FERC rejected an arrangement between Talen's Susquehanna nuclear plant and AWS in a first-of-its kind BTM co-location agreement. Following the rejection by FERC in 2024, Talen pivoted to a front-of-meter arrangement to avoid further FERC scrutiny. In response, hyperscalers **are increasingly more willing to consider building their own generation for BTM projects as speed-to-power remains a priority**. Since most hyperscalers are not looking to own grid infrastructure long-term, build-to-transfer models seem to be gaining traction. Build-to-transfer models allow utilities to offload generation construction risk to hyperscalers, which then transfer the completed generation asset back to the utility, which commits to serving the large-load entity.

Utilities and hyperscalers continue to move towards the middle. McKinsey's commentary from our equity team's recent fireside chat highlights an expected ~260-300 GW of total compute demand by 2030, which is clearly far above what is plausible from a supply standpoint. The generation supply deficit should gradually decline as additional gas generation comes online in 2029 and an increase in inferencing compute should allow smaller, distributed generation options over time. Still, utilities and hyperscalers are gradually closing the gap as both parties' expectations continue to evolve: Hyperscalers have become more comfortable with cost-sharing, demand response, and alternative generation procurement, while utilities have been moving forward with tariffs that satisfy regulatory concerns around residential bills.

Interconnection queues continue to grow in ERCOT. One of the most visible constraints for new data center projects is the continually growing interconnection queue. A significant amount of capacity in the interconnect queue consists of storage and solar assets that have been trying to get ahead of the tax credit sunset, but incremental gas additions in ERCOT, SPP and MISO coincide with the recent uptick in data center demand. We expect the cadence of gas additions to increase as more turbines enter the market in the 2029 timeframe.

Figure 16: Interconnection Requests by Region and Generator Type, 2014-2025

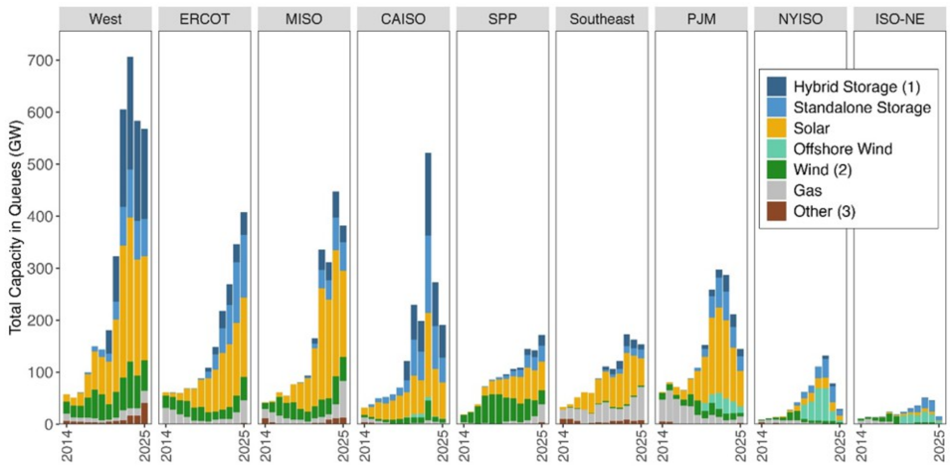

Source: Lawrence Berkeley National Laboratory.

Gas generation continues to be the largest addition to planned generation through 2031. The figures below show planned generation capacity that will come online over the next several years. Notably, much of the planned generation is increasingly coming from natural gas, as it remains one of the more cost competitive dispatchable technologies. However, the lag in building gas turbines in the US means most of this supply will not be available until closer to 2029. Separately, the amount of planned generation grew by over 80 GW from April 2024 to April 2025, underscoring the amount of new generation that is slotted to be put in play through the end of the decade. Notably, there continues to be a large uptick in planned generation capacity beyond 2029, and 2030 stands out with gas generation representing 71% of the 25.5GW of total planned capacity for the year. We expect this trend to continue as GE Vernova and other turbine manufacturers point to a large backlog of generation capacity that won't begin coming online until the end of the decade.

Figure 17: Natural Gas % of Total Planned Generation Capacity in the U.S.

MW

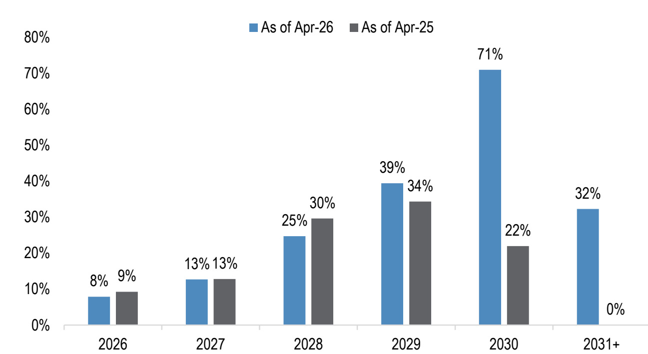

Source: EIA Preliminary Monthly Electric Generator Inventory (based on Form EIA-860M as a supplement to Form EIA-860).

Figure 18: Total Planned Generation Capacity in the U.S. – All Technologies

MW

	As of Apr-26	As of Apr-25
2026	70,404	71,053
2027	86,302	65,753
2028	64,150	47,977
2029	31,615	15,388
2030	25,562	6,037
2031+	10,352	675
Total	**288,386**	**206,883**

Source: EIA Preliminary Monthly Electric Generator Inventory (based on Form EIA-860M as a supplement to Form EIA-860).

Regulated utility capex on generation is increasing, but is unlikely to keep up with data center needs alone. Regulated utility investment in new generation is increasing and made up 30% of the overall $208 billion utility spend in 2025, as shown below, but

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

much of the increase in utility capex over the last several years is driven by state-directed electrification and grid hardening mandates, as opposed to adapting to large loads. Data center–related capex and generation build has only accelerated in the last few years, and planned capacity additions from regulated utilities are unlikely to keep up with the total compute demand expected through 2030. As a result, hyperscalers are increasingly looking at “all-of-the-above” solutions for additional power procurement, including bridge power solutions that we’ve seen from companies such as VoltaGrid and PowerBridge, as well as traditional oilfield services companies that have since pivoted to power generation.

Figure 19: EEI Projected Functional Capital Expenditures, 2024-2025 (Regulated Utilities)

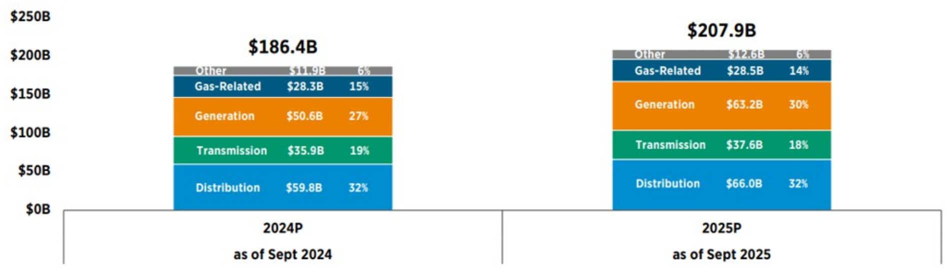

Source: EEI Financial Analysis Department, EEI member company reports.

Data center announcements remain dominated by Google, Amazon, Microsoft, and Meta. The cadence and scale of hyperscaler-utility partnerships announced over the past few years underscore a structural inflection in US power demand. In estimates published by the Electric Power Research Institute in February, data centers could account for 9% to 17% of total US electricity demand by 2030, up from about 3% to 4% today. Google has dominated the data center and grid connection battle and became the only tech giant to own a power company with the $4.75 billion acquisition of Intersect, a wind and solar developer. In addition to the contracts Google has announced for roughly 15 GW of capacity since January 2024, the company recently announced an agreement with OG&E to build out two solar facilities with a total capacity of 1.2 GW. Many of the top AI and cloud firms are committed to renewable energy targets, but due to supply constraints in the market and an administration less partial to renewable energy, energy sources including nuclear, gas-fired and coal are in equally high demand. To fill the speed-to-power gap, some hyperscalers are building out their own behind-the-meter gas and fuel cells with companies such as VoltaGrid, Bloom Energy, and PowerBridge, which continue to serve as bridge generation solutions until more permanent baseload generation is put into service. While substantial agreements have been made and many GWs of capacity have been contracted, factors such as regulatory treatment of large-load tariffs, cost allocation, interconnection queue throughput, and supply-chain bottlenecks are creating significant headwinds for these agreements to be put into service, with projects being pushed back to the late 2020s/early 2030s.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Table 6: Select Data Center / Utility Agreement Announcements

Announcement Date	Details	Utility / Power Provider	Hyperscaler	Generation Plant	Capacity	Term
5/7/26	Generation Agreement	Black Hills	Not Announced	Wyoming	1.8 GW	
4/30/26	ESAs and Capacity Purchase Agreements	OG&E	Google	Two solar facilities under construction	1.2 GW	
4/13/26	Expanded PPA	Bloom Energy	Oracle	Bloom's fuel cell systems	2.8 GW	
4/2/26	non-binding LOI to form JV	NewEra Energy	Stream Data Centers	Texas Critical Data Center campus	1 GW	
4/2/26	Development Agreement	PowerBridge	LandBridge	Co-located power generation under development	2 GW	
3/31/26	Exclusivity Agreement	Chevron & Engine No.1	Microsoft	No terms	2.5 GW	
3/20/26	New Transmission Investment	AEP Ohio	SB Energy		10 GW	
3/17/26	Development Agreement	DTE Energy	Google		2.7 GW	
3/16/26	Acquisition of Monarch Compute Campus	Nscale	Microsoft		1.35 GW	
2/24/26	ESA for new data center	Xcel Energy	Google		1.9 GW	
2/24/26	PPA for new data center	AES Corp	Google	Building infrastructure in Wilbarger County, TX		20-yr
2/11/26	2 PPAs	TotalEnergies	Google	Sites under development in Wichita and Mustang Creek, TX	1 GW	15-yr
1/15/26	3 PPAs	Clearway Energy	Google	Projects in Missouri, Texas, West Virginia	1.24 GW	20-yr
1/9/26	Development Agreement	Oklo	Meta	New power campus in Pike County, OH	1.2 GW	
1/9/26	Development Agreement	TerraPower	Meta	8 Natrium reactor and storage system plants	2.8 GW to 4 GW capacity	
1/9/26	PPA	Vistra	Meta	Beaver Valley, Davis-Besse, Perry nuclear power plants	2.6 GW	20-yr
12/8/25	Expanded PPA	NextEra Energy	Google		multi GW	25yr
12/8/25	11 PPAs and 2 ESAs	NextEra Energy	Meta		2.5 GW	
10/30/25	Power Supply Agreement	DTE Energy	Not Announced		1.4 GW	19-yr
10/27/25	PPA for data center	NextEra Energy	Google	Duane Arnold Energy Center	615 MW	25yr beginning 1Q29
10/15/25	Modular gas power supply	VoltaGrid	Oracle	VoltaGrid gas platform in Texas	2.3 GW	
10/2/25	PPA for new data center	Entergy	Google	New Solar Plant	600 MW solar + 350 MW battery storage	
7/24/25	PPA	Bloom Energy	Oracle	Bloom's fuel cell systems	1.2 GW	
7/15/25	Hydro Framework Agreement + PPAs	Brookfield Renewable	Google	Holtwood and Safe Harbor hydroelectric faciilities initially	3 GW	20-yr PPAs
6/11/25	Expanded PPA	Talen Energy	Amazon	Susquehanna	1.92 GW	20-yr ending 2042
6/3/25	PPA	Constellation Energy	Meta	Clinton Nuclear Facility	1.1 GW	20-yr beginning June 2027
2/11/25	On-site Generation Partnership	VoltaGrid	Vantage Data Centers		>1 GW	
12/5/24	ESA for data center	Entergy	Meta	Adding power plants	1,5 GW	
3/4/24	Data Center PPA	Talen Energy	Amazon	Susquehanna	960 MW	

Source: Company IR sites, J.P. Morgan.

Creativity in Behind-the-Meter Solutions

Longer regulatory approval timelines, combined with extended lead times for utility-scale power and equipment, are pushing project owners to develop more creative solutions. The single 'easiest' solution for data center developers is to bring your own power via behind-the-meter (BTM) power solutions. Announced BTM solutions so far exceed 90 GW, although that should be heavily discounted given that it does include low utilization renewables that themselves require grid power or substantial diesel backup to support. Texas remains a key market for BTM arrangements given the lack of FERC oversight, which has caused more regulatory lag in other states such as PA.

Figure 20: Announced BTM Power Solutions by State (GWs)

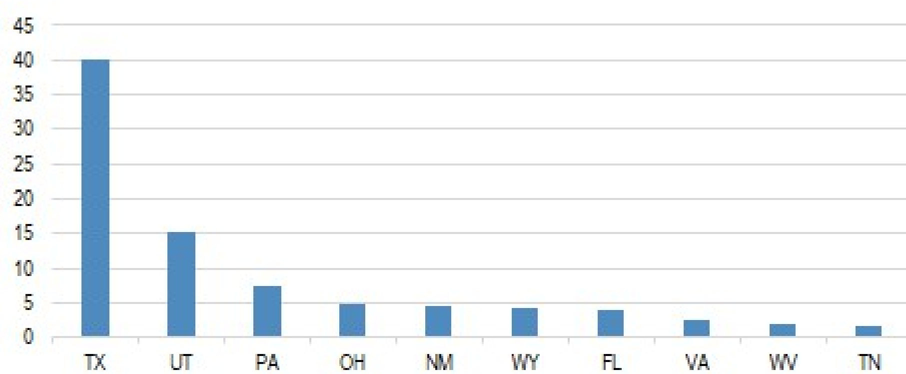

Source: Cleanview.

Effective lead times for utility-scale turbines are approaching 5 years, and even smaller turbines have lead times approaching 3 years. As a result, data centers are looking at reciprocating engines, fuel cells and aeroderivatives—basically anything that can be used to move copper and generate electricity. These "bridge" power solutions are expected to satisfy much of hyperscalers' speed-to-power priorities for the next several years as more baseload CCGTs come online in the 2029 timeframe.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 21: Select Generation Equipment Lead Times and Capacity

Technology	Size per Unit	Lead Time	Land Use	Electrical Efficiency
	MW	*Months*	*MW/Acre*	*%*
Aeroderivative GT	30-60	18-36	30-50	35-40%
Industrial GT	5-50	12-36	20-40	35-40%
Small CCGT	40-100	18-36	20-30	40-55%
Medium Speed RICE	7-20	15-24	8-15	40-50%
High-Speed RICE	3-5	15-24	5-10	40-50%
Fuel Cells	0.325	3-4	30-100	50-55%
H-Class CCGT	600-1000	36-60	20-30	50-60%

Source: SemiAnalysis and J.P. Morgan estimates.

All of the above (or whatever is available) is turning into the go-to BTM solution. We are seeing data center power solutions being deployed using a mix of different size and producer gas turbines and reciprocating engines all on the same site. Producers are essentially jury-rigging solutions using whatever generation capacity they can find, and for now appear willing to stomach the maintenance challenges, operational complexity, and inefficiency of multiple power solutions from multiple vendors. **Basically, the best power is the power you can find.**

Figure 22: Meta Data Center Power Solutions

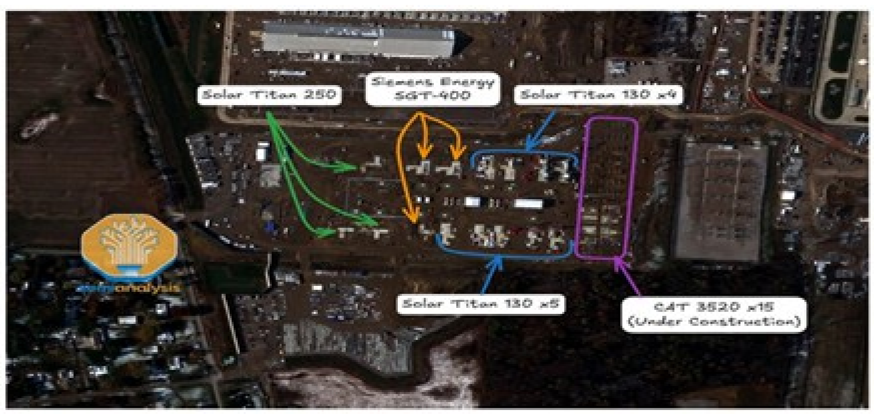

Source: SemiAnalysis.

GPU & ASIC Financing Needs and Solutions

Harlan Sur
(415) 315-6700
harlan.sur@jpmorgan.com
J.P. Morgan Securities LLC

Tarek Hamid
(212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC

The capital markets have proven remarkably adept at financing hyperscalers and data centers directly. However, the next leg of financing needs are arguably chunkier and potentially more complex in the form of GPUs and custom ASICs. Roughly 60% of the fully burdened cost of a new leading-edge data center using Blackwell equivalent GPUs is silicon and servers. That ratio should come down as the share of custom ASICs increases in the next couple of years. Nonetheless, we have effectively only been financing the data center boxes in public markets so far, with the exception of term loans and bonds for neocloud CoreWeave.

Figure 23: New Datacenter Capital Spending Mix

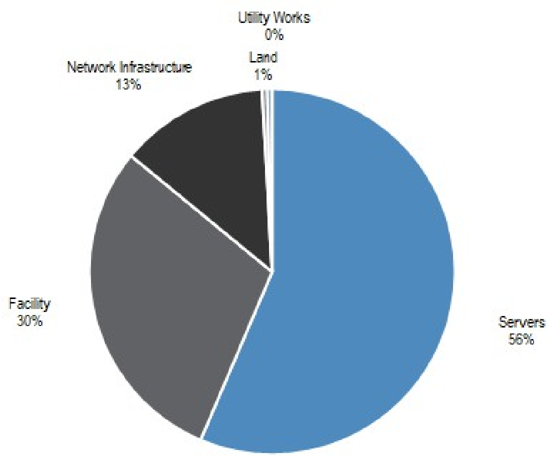

Source: Epoch AI.

Figure 24: New 1GW Datacenter Using Blackwell Illustrative Costs ($bns)

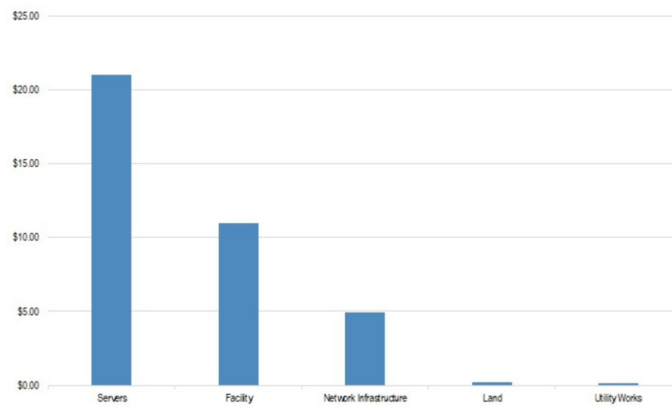

Source: Epoch AI.

On a go-forward basis, we expect to see some very large funding needs for GPUs and ASICs. Some of that is going to be captured directly by the hyperscalers themselves, while some of it will likely be financed on neocloud balance sheets or SPVs. But the aggregate funding need for AI silicon directly in the next five years is estimated at well over $2 trillion.

Financing Silicon Has Its Challenges

- **We don't know how long silicon will physically last, and not every chip will be the same.** The difference between a 3-year and 7-year economic life is profound in terms of debt capacity. Utilization and environmental factors should drive substantial variation in mechanical degradation of silicon. Put another way, chips used in intensive model training applications are likely to suffer quicker mechanical failure rates than chips being used in inference or batch processing applications. The Meta Llama 3 study noted that over 54 days of training, the run encountered 419 unexpected component failures with 262 of those failures related to GPUs or GPU onboard memory or sensors. While that sounds like a very large number, the training

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

cluster included 16,384 GPUs—so failures related to less than 2% of GPUs.

- **We don't know how well current generations of technology will hold their value as new generations of silicon come to market.** Intellectually, it is hard to see the world go from short compute to long compute in the next five years, which arguably makes the obsolescence argument less relevant in the very near-term. Moreover, the oft-cited GPU cascade from training for 1-2 years to inference for 2-3 years to batch processing is likely to extend economic life as chips move to less computationally demanding tasks. Finally, we think recent rental prices for Hopper (H100/H200) chips are a very positive indicator that some of the draconian obsolescence assumptions people have been using are flawed:

Figure 25: Neocloud H100 Rental Price Index

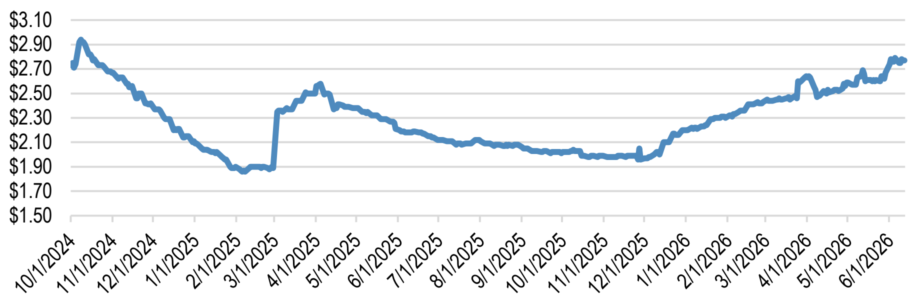

Source: Bloomberg Finance L.P., Silicon Data.

- **There is a fundamental duration mismatch between where markets want to lend (i.e. longer duration) and the economic lives of GPUs.** Regardless of whether silicon lasts 3 years, 5 years or 7 years, it probably doesn't last 10 years or 30 years. The depth of the High Yield bond market and Leveraged Loan markets at the 5-year duration point is very large—but those markets are collectively small as compared to the High Grade bond market. For context, the High Grade bond market is 3x the size of the combined Leveraged Finance market. The combined 3-year and 5-year duration buckets of the High Grade market constitute just 43% of the market, and more importantly investors consistently express a preference for more duration in the primary market. Creating market depth for 5-year assets represents a real challenge for the data center buildout, in our view.

Leveraged Loans and High Yield Are Natural GPU Homes, Especially for the Neoclouds

We have already seen $15 billion of largely delayed draw term loan GPU-related financing in the Leveraged Loan market. We expect that to continue as the Leveraged Loan market is a very natural home for amortizing 5-year paper, in our view. Unlike the bond markets, convexity and prepayability are materially less of a concern for the loan market. Additionally, repricing optionality allows borrowers to be somewhat less sensitive to timing the market. The downside of the Leveraged Loan market is likely scale. The single largest tranche currently in the Leveraged Loan market is just $13 billion and the largest borrower in the Leveraged Loan market is that same $13 billion. Relative to the potential demand for GPU financings, that is a proverbial drop in the bucket. On the other hand, for neoclouds like Coreweave and Iren, market capacity is likely the right number for a relatively straightforward and easy financing structure.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

The Broadcom/Apollo/Blackstone JV Is a Paradigm to Bring Private IG Capital to Bear

The recently announced Broadcom/Apollo/Blackstone JV is a landmark transaction in the space. Under the broad contours of the deal, Broadcom is manufacturing Google TPUs for use by Anthropic. Those chips are being financed on the newly formed AI XPV platform. An entity controlled by Apollo, Atlas SP, is contributing $800 million of equity to the SPV. The SPV then issued $6 billion of A1 Notes, reportedly at T+100bp, largely to banks, and $24 billion of A2 Notes at a reported 5.75% interest rate to longer-term investors, including including insurance companies. Both the A1 and A2 Notes benefit from a residual value guarantee from Broadcom. Per Broadcom's 10-Q, *"On June 8, 2026, we arranged for an investor partner to take on certain agreements to purchase AI racks based on custom AI accelerators designed by us and the related lease agreements with a customer that enable access to compute capacity. In connection with the arrangement, we entered into a backstop agreement with the investor partner for the customer's lease obligations over 5-year terms. The backstop will increase over time as the AI racks are deployed and decrease as the customer makes payments on its lease obligations, with a maximum exposure of $29 billion. In the event of default by the customer, we have various remedies, including the assumption of the lease or effecting a sale of the AI racks, which would reduce our maximum exposure."*

The residual value guarantee from Broadcom allowed both the A1 and A2 Notes to reportedly benefit from mid-BBB equivalent private investment grade rating. The Joint Venture also issued $4.5 billion of Series B Notes that do not benefit from the residual value guarantee. Those Notes reportedly cleared at 8.5%. Intuitively, that means the cost of capital for the 85-98% LTC tranche of the GPU joint venture capital structure was 8.5%—which appears quite attractive. That implicitly puts the all-in pre-tax cost of capital in the low 6% area, and the post-tax cost of capital likely at or below 6%.

Figure 26: AI XPV Schematic

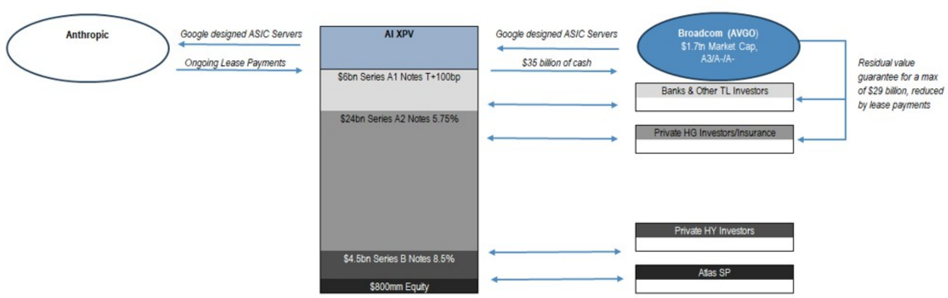

Source: Company reports and J.P. Morgan Research estimates.

But the Billion Dollar Solution Is Solving the Duration Mismatch

The AI XPV joint venture is a landmark structure and we think could easily be ported over to public markets. However, it doesn't necessarily solve the duration mismatch. Part of the challenge is that at the end of day the pool of assets (chips) is static. As lease payments are received, notes are amortized and the residual value guarantee and overall funding pool shrink. That functionally limits the amount of capital

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

that can play in the deal to short-duration capital.

An idealized solution in the future, in our view, would allow a joint venture to engage in active management of a chip portfolio by reinvesting lease payments into new chips to replace damaged or obsolete chips, allowing the joint venture to tap longer-duration capital. We believe that argues for a much larger equity check as a proportion of the capital stack, at a minimum. Additionally, proving a residual value guarantee to achieve high grade ratings for an active pool of assets is likely more cumbersome. That would likely require a minimum LTC guarantee on the asset pool that matches the agreed-to capital stack; that is, if 60% of the capital stack at cost is High Grade, does a chip provider or hyperscaler provide a residual value guarantee up to 60% of the acquisition cost of the pool of chips under management (which really means a guarantee likely of a particular depreciation curve with LTC covenants in the debt stack to potentially force amortization or reinvestment)?

Conversely, a better solution might very well be a diversified pool of chips (GPUs, ASICs) from different manufacturers with a portfolio of different hyperscalers and end users. Diversification benefits and low enough leverage might allow for investment grade ratings in that structure. But keeping the cost of capital of such a structure competitive could be a challenge. Moreover, that business model frankly looks to us a lot like a ‘super’ neocloud at some point.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

High Grade Bond Markets The AI Workhorse

Nathaniel Rosenbaum
(212) 834-2370
nathaniel.rosenbaum @jpmorgan.com
J.P. Morgan Securities LLC

Silvi Mantri
(212) 834-7239
silvi.mantri@jpmchase.com
J.P. Morgan Securities LLC

It is hard to spell High Grade in 2026 without 'AI'. Our key AI capex forecasts for the High Grade market are as follows:

- From a top-down perspective, **we expect a total of $2.1tr of HG borrowing over the next 5 years** (inclusive of what has already come to market YTD) across all hyperscalers, datacenters and associated GPU financing activity globally. This equates to $420bn/year over 5 years, though we don't expect the pace to be linear but rather accelerating as the buildout progresses.
- From a bottom-up perspective, in 2026 we are forecasting $150bn of hyperscaler supply in the USD market in 2026, with a further $100bn USD-equivalent in the increasing number of non-USD markets participating in this trend. This leaves another ~$170bn of data center and semiconductor financing structures (e.g. from issuers that may not yet currently be a part of the public HG market; approximately $60bn of such structures have already priced YTD). These issuers have also made inroads into HG-adjacent markets such as the Municipal prepay market, but we believe this issuance will be relatively small in 2026 (on the order of $5-$15bn).

The Heavily Anticipated Hyperscaler Issuance Surge Is Upon Us

The amount of AI-related capex that has been financed in the HG bond market YTD is, by some measures, quite astounding. We are not yet at the midpoint of the year and the hyperscalers, data centers and semiconductor financings total $165bn or already $27bn more than what was done in the *full-year* 2025. This universe now makes up 6.2% of the HG bond market versus just 3.9% at the end of 2024. If we expand this scope to a slightly wider grouping of issuers that are directly involved in the AI capex ecosystem (Utilities contracting directly with hyperscalers, Industrials making generators going into datacenters, etc., as explained in this research note), these related issuers now account for 15.5% of the HG market ($1.4tr), effectively making it the largest 'sector' and one that continues to expand.

As capex expectations continue to move higher, the associated financing need is scaling up rapidly, particularly within hyperscalers, driving a step-change in issuance across both USD and non-USD markets. In this context, we expect about $2.1tr of HG issuance across both USD and non-USD markets tied to AI/Data Center financing over the next 5 years in total, which implies a $420bn/year cadence across the TMT ecosystem, though we don't expect the pace to be linear but rather accelerating as the buildout progresses. Importantly, the recent step-up in issuance has been driven almost entirely by hyperscalers, with USD supply volumes already surpassing last year's total of $107bn and tracking toward our FY26 forecast of $150bn, while issuance from chipmakers has been broadly in-line with prior-year trends so far.

Figure 27: 15.5% of IG debt is tied to AI, making it in a sense the biggest 'sector' in the market

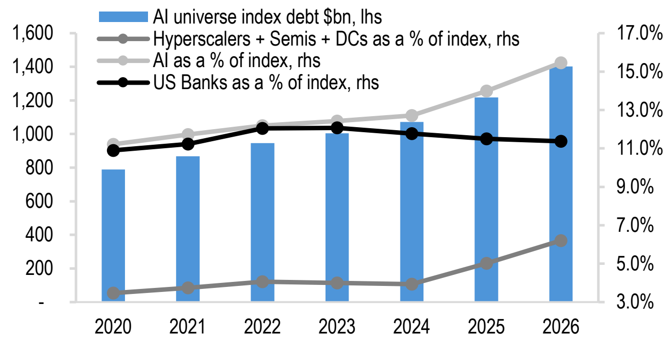

Source: J.P. Morgan, Dealogic.

Figure 28: USD hyperscaler issuance YTD 2026 has already exceeded full year 2025

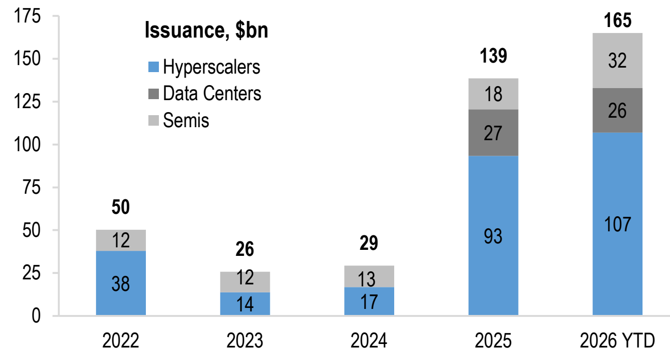

Source: J.P. Morgan, Dealogic.

The market impact on Tech sector spreads of this issuance ramp-up has been quite pronounced, with Tech spreads going from 34bp *tight* to the JULI index in late 2023 to 10bp *wide* to the JULI currently, or close to the cheapest Tech has traded to the index since 2004. The most recent widening was not a function of the existing cohort of hyperscaler debt going wider, but rather our decision to include the recent crop of Datacenter issuance into the index at April month-end (see discussion here). These data centers have continued to trade relatively wide of associated hyperscaler debt (see Relative Value section of this note further below). For the broader cohort of hyperscaler senior unsecured debt, credit spreads have largely stabilized at these wider levels and, thus, we recently upgraded the sector from Underweight to Neutral (see here).

Figure 29: Tech spreads have widened significantly versus the market but now finding a renewed bid as investors move up in quality

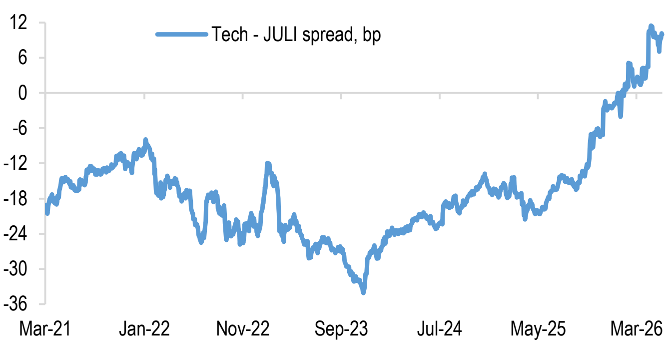

Source: J.P. Morgan.

Figure 30: Data centers continue to trade relatively wider to their associated hyperscalers

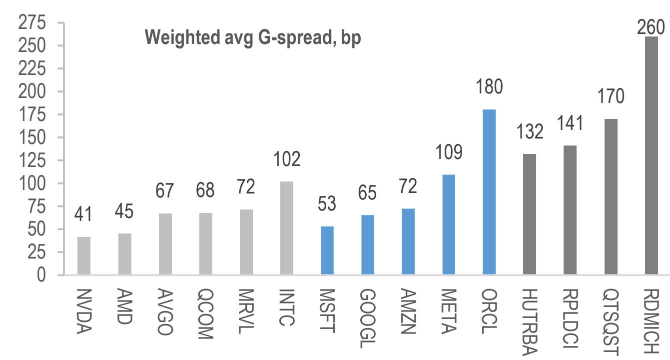

Source: J.P. Morgan.

Hyperscalers Are Doing a Remarkable Job of Diversifying Their Funding Globally

One of the key reasons we believe hyperscaler spreads have stabilized is because of the surging share of non-USD issuance. This is broadening the potential investor base and therefore improving overall absorption of their USD debt too. USD hyperscaler issuance is just 63% of their global issuance YTD versus 86% in 2025: Hyperscalers have priced $62bn USD-equivalent in non-USD markets, which is more than 4x what they did there in full-year 2025. Crucially, this supply is entering markets that up until now had structurally low Tech exposure. In other words, rather than saturating a relatively Tech-heavy USD market, incremental issuance is being distributed across markets that had

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

very little Tech exposure, improving overall absorption capacity.

Figure 31: Hyperscaler issuance: Non-USD portion rapidly expanding

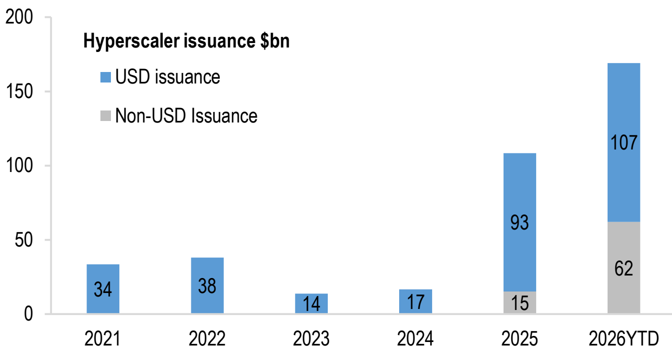

Source: J.P. Morgan, Dealogic, Bloomberg Finance L.P.

Figure 32: Most non-USD markets had very little exposure to the Tech sector

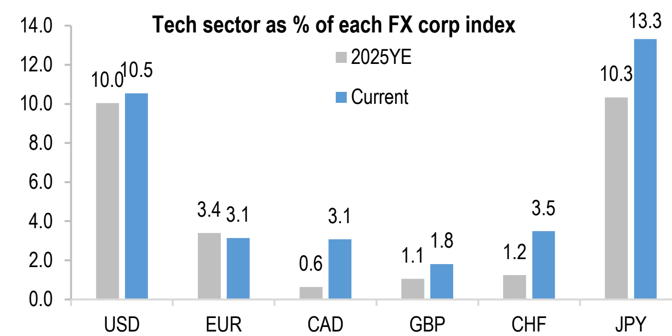

Source: J.P. Morgan, Bloomberg Finance L.P.

Issuers are actively leaning into this dynamic: Alphabet has been the most diversified to date, issuing $20bn in USD alongside $32bn (61%) across five other currencies, while Amazon has issued $37bn in USD versus more than $30bn across EUR, CAD and CHF. In contrast, Meta and Oracle have funded exclusively in the USD market YTD, printing $25bn each. Overall, the evolution of funding channels suggests the market is better equipped to digest elevated Tech supply than in prior cycles. The one major non-USD market that has been notably absent from this surge has been AUD, despite its recent rapid growth (which we discussed in more detail here); we expect it will join the funding fray soon enough. We believe hyperscalers will be able to continue leveraging the non-USD markets significantly, though to varying degrees depending on the depth of these various markets: In the chart below, we show how much has been done in each currency YTD and the associated performance since issuance. These are not all over the exact same time frame, but they do point towards some markets having a lot more depth (EUR and GBP) than others (CHF or JPY), which seems sensible in light of the relative market sizes.

Figure 33: Hyperscaler issuance by FX and associated new issue performance: Most markets have room for more

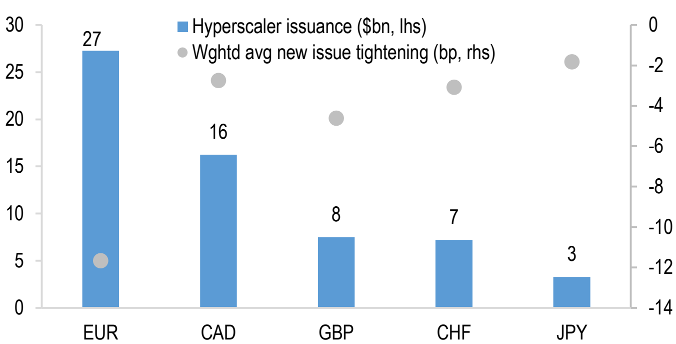

Source: J.P. Morgan, Dealogic, Bloomberg Finance L.P.

Figure 34: DM IG corporate markets by FX: AUD only major market not yet tapped YTD

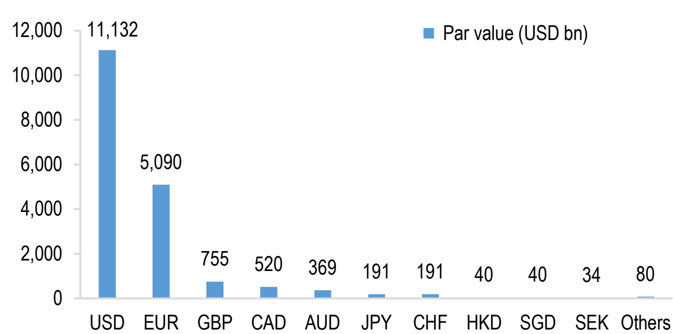

Source: J.P. Morgan, Bloomberg Finance L.P.

Some Capacity Constraints in USD Market Emerging in the Form of Steeper Credit Curves

The average maturity of non-USD hyperscaler issuance YTD is 14 years (and 12 years excluding the GOOGL 100yr GBP bond) versus 17 years for USD issuance; in other

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

words, hyperscalers have been steering more of the long-dated supply to the market that is best equipped to absorb it. This long-end USD issuance isn't really competing against too many other issuers as the 30yr supply dynamic has become very bifurcated lately as other issuers are shying away from the long-end. As such, 30yr non-hyperscaler supply is only 9% of the total YTD versus 12% in 2025 whereas if we include hyperscalers this figure jumps to 12% for 2026 YTD. That has driven hyperscaler spread curves 7bp steeper y/y versus the rest of the market 2bp flatter to a very flat 12bp already (curves got as flat as 10bp last summer, prior to the hyperscaler issuance surge).

Figure 35: 30yr supply dropping rapidly ex-hyperscalers

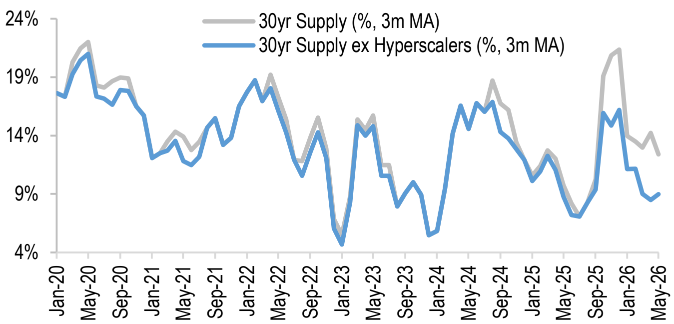

Source: J.P. Morgan.

Figure 36: 10s30s spread curve flatter ex-hyperscalers which are steeper on account of their heavy 30yr issuance

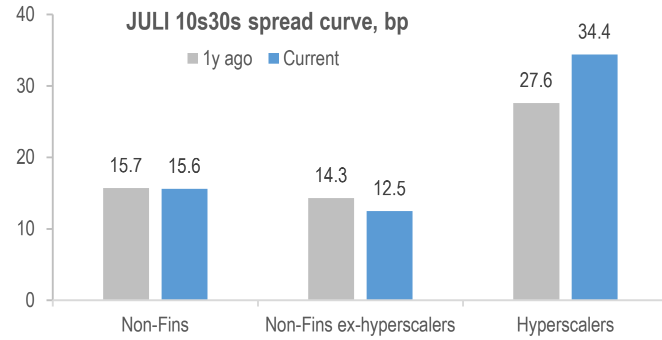

Source: J.P. Morgan.

This suggests there is little room for 10s30s spread curves to flatten further unless core long-end buyers raise their exposure to hyperscalers. Our recent analysis of US life insurance holdings for 2025 shows insurance companies are underweight versus the index on all five hyperscalers, but the smallest UW is in MSFT and the largest is in META: META +$1.8bn YoY (-0.3% UW versus the index), GOOGL +$1.1bn YoY (-0.18% UW) and ORCL +$703mn YoY (-0.28% UW) while staying about unchanged in AMZN (+$165mn YoY, -0.30% UW) and decreasing holdings in MSFT (-$616mn YoY, -0.11% UW). This implies US life insurance companies came into 2026 well set up to absorb the significant primary market activity from the hyperscalers, though we suspect most are now at more neutral levels given heavy 2026 YTD issuance and with the prospect of more issuance from wider-trading data center and GPU financing structures, we suspect there is little incentive for these insurers to overweight the hyperscalers and flatten their still steep 10s30s spread curve.

Figure 37: Insurance companies came into 2025 underweight the hyperscalers but may not be anymore given heavy issuance

Issuer	Sector	Insurance Holdings ($mn)	Insurance Holdings YoY Change ($mn)	2025 issuer index weight	Insurance holdings weight	OW/UW versus index debt	June 2026 issuer index weight
Alphabet Inc	Technology	2,135	1,079	0.34%	0.16%	-0.18%	0.55%
Amazon.com Inc	Technology	5,290	165	0.70%	0.40%	-0.30%	1.02%
Meta Platforms Inc	Media / Entertainment	3,873	1,763	0.61%	0.29%	-0.31%	0.92%
Microsoft Corp	Technology	4,496	-616	0.45%	0.34%	-0.11%	0.37%
Oracle Corp	Technology	9,959	703	1.04%	0.76%	-0.28%	1.29%

Source: J.P. Morgan, NAIC.

High Grade markets globally are excelling at financing the lion's share of the AI capex surge, leaving scope for these issuers to continue raising a staggering amount of debt over the coming years.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Leveraged Finance Markets Upside

Nelson Jantzen
(212) 270-1169
nelson.r.jantzen@jpmorgan.com
J.P. Morgan Securities LLC

The High-Performance Computing (AI) Subsector in the High-Yield Bond Index has grown remarkably YTD. That is causing investors and issuers to wonder how much additional capacity there is in Leveraged Finance to fund the AI capex buildout.

Figure 38: Across leveraged finance, we forecast $80bn of HPC issuance in 2026 and $350bn cumulatively by the end of 2030

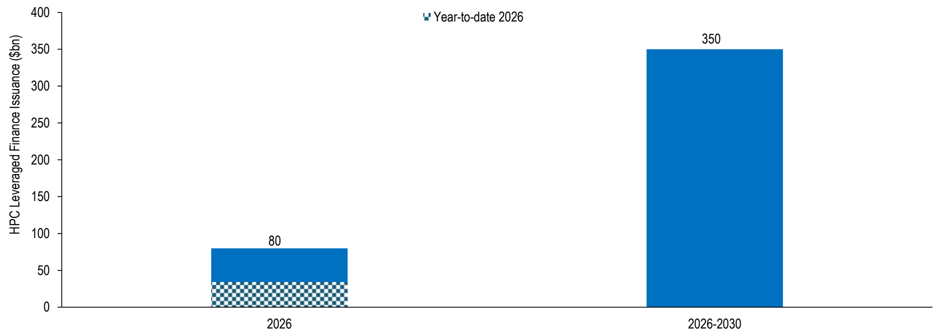

Source: J.P. Morgan.

A large contributor to the growth in net high-yield bond issuance in 2026 can be attributed to the High-Performance Computing (AI) subsector. **For context, HPC (AI) bond issuance now totals $31.1bn year-to-date.** Accounting for 41% of the $76.9bn non-refi-related HY issuance year-to-date, HY net issuance is tracking well ahead of last year's pace (+125% yoy). Coupled with a rise in fallen angel volume, the high-yield bond market is expanding again in 2026 (+$48bn) after expanding at its fastest pace in 2025 since 2020's Fallen Angel wave. **How big has the HPC subsector become in the J.P. Morgan High-Yield Bond index? The HPC subsector has expanded from 1.07% at YE25 to 2.75%.** This includes $39.8bn of bonds outstanding via 12 issuers with an average duration inside 3.5yrs. For reference, HPC now accounts for 33% of the HY Technology Sector, which is now 8.31% of the overall High-Yield Bond Index. For context, Technology already has a far greater presence in the Leveraged Loan index (predominately Software) with a weighting of 14.07%. With spreads in HPC 256bp tighter year-to-date to 281bp, this rapidly growing sector is contributing to the overall tight level of HY index spreads. **As well, the HPC sector has been a significant source of alpha for credit managers in 2026, with total returns of +10.64% easily surpassing the HY index (+1.63% YTD).** Of note, CRWV's 5.5yr $3.1bn leveraged loan priced in early May at SOFR+450bp was the first for the High-Performance Computing (AI) subsector in the Leveraged Loan index (X.AI paid down). We believe the leveraged loan market is going to fund an increasing part of the AI capex buildout in Leveraged Finance over the next year. **For Leveraged Finance, our forecast for HPC issuance in 2026 is $80bn ($65bn HY, $15bn LL). And over the next five years (2026-2030), we estimate $350bn of HPC issuance syndicated across bonds and loans with a potential peak in 2027.**

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 39: HPC (AI) bond issuance now totals $31.1bn year-to-date following $8.0bn of issuance in 4Q

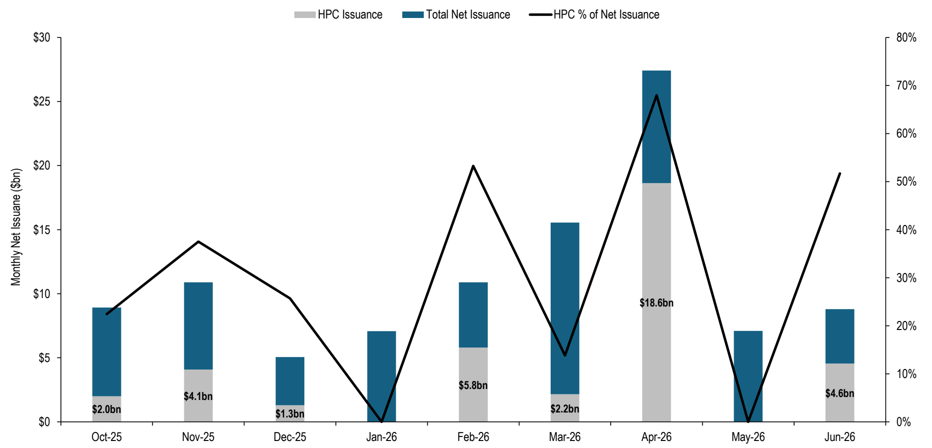

Source: J.P. Morgan.

Figure 40: The HPC subsector has expanded from 1.07% at YE25 to 2.75%

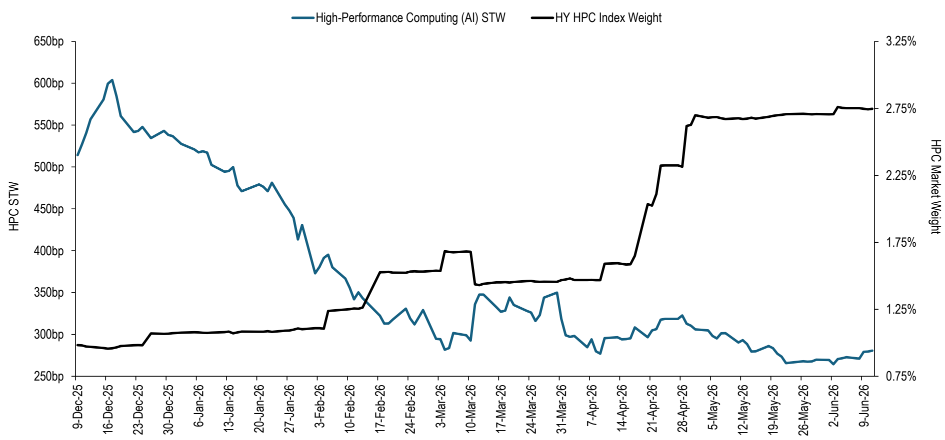

Source: J.P. Morgan.

Figure 41: High Performance Computing Stats vs Index

	HPC SubSector	HY Index	HPC - Broader HY
STW	281bp	313bp	-32bp
YTW	6.94%	7.24%	-0.30%
Price	$101.37	$96.95	+$4.42
YTD Return	10.64%	1.63%	+9.01%
Duration	3.30yrs	3.16yrs	+0.14yrs

Source: J.P. Morgan.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

How Much Capacity Is There in Leveraged Finance to Fund the AI Capex Buildout?

Size of the leveraged credit universe: The syndicated leveraged credit universe is expanding again. **The US$ developed high-yield bond and institutional loan universes currently stand at $1.60trn and $1.65trn, respectively.** Expanding $48bn year-to-date (or 3%), the Developed US$ HY bond market has grown by $173bn or +12% since YE23 after contracting by $232bn (or 14%) in 2022-23. For context, the high-yield bond universe expanded by $355bn (or 27%) in 2020 and 2021 amid a wave of fallen angels and heavy new-issue activity. Meanwhile, institutional loans outstanding have grown a modest $9bn year-to-date amid a rise in M&A issuance and a relatively balanced mix of private and public market takeouts. For context, the loan market contracted by $84bn (or 5%) from June 2022 to YE23 after expanding by 11% per annum over the preceding decade. Constrained capital markets, the imbalance of rising stars/fallen angels (HY), and cannibalization by the private credit space (LL) were the largest contributors to the contraction of the syndicated bond and loan base in 2022-23. Lastly, global private credit drawdown fund AUM stands at $1.81trn, which consists of $543bn dry powder and $1.27trn invested. This compares to $1.1trn at YE20. For context, the private credit universe has grown at a CAGR of 14% since 2020 compared to syndicated loans at a CAGR of 6%. **This puts the syndicated leveraged finance DM universe at $3.25trn, double where it stood coming out of the GFC, or close to $5trn after including private debt.** What was the largest 5-year stretch for growth of high-yield bonds and loans outstanding post GFC? In the 5yr period ending 2015, the syndicated Leveraged Finance universe expanded by as much as 50% or $825bn to $1.9trn (or $650bn ex-Fallen Angels).

Figure 42: The Developed US$ high-yield bond and leveraged loan universe has grown to $3.2trn, with US$ Leveraged Finance close to $5trn after including private debt

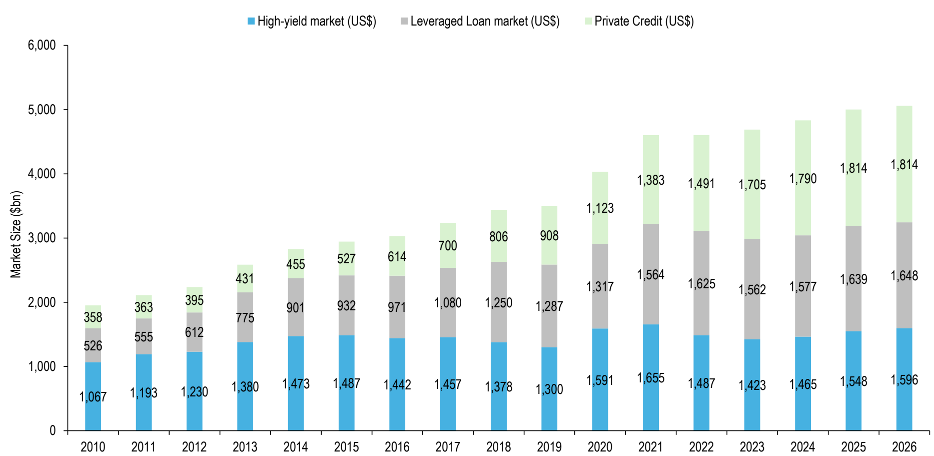

Source: J.P. Morgan, Preqin.

Historical industry growth and current composition: Historically, rapid industry growth in leveraged finance has been accompanied by an overextension in borrowing and an eventual default cycle. **Between 1995 and 1999, the Telecom industry grew from 5% to as much as 23% of the HY index.** And in 2001 and 2002, the Telecom industry produced sector default rates of 26% and 45%, respectively. **Between 2009 and 2014, the Energy industry grew from 10% to as much as 17% of the HY index.** This equated to growth of $135bn of par outstanding over 5 years. And when Oil prices collapsed in 2015-2016 and again in 2020, the Energy industry produced double-digit annualized sector default rates. Also post GFC, the Healthcare sector reached a double-digit market-weight in both the HY and LL indices and contributed as much as 40% of default activity in 2022-23. **And more recently, the Financial and Technology sectors have experienced the largest growth in Leveraged Finance.** For example, since 2020 the Financials sector has grown from 6% to 10% of the HY index and from 6% to 10% of the LL index. This equates to growth over 5yrs of $56bn and $73bn in par outstanding bonds and loans, respectively ($129bn combined). And the Technology sector expanded steadily in the leveraged loan universe to as much as 16% of the Leveraged Loan index by 2021 (has held steady since). The max 5yr stretch in growth in par outstanding for Tech loans was $130bn, matching the growth in Energy bonds outstanding post-GFC.

Figure 43: Largest Industry Weightings in the High-Yield Bond and Loan Indices

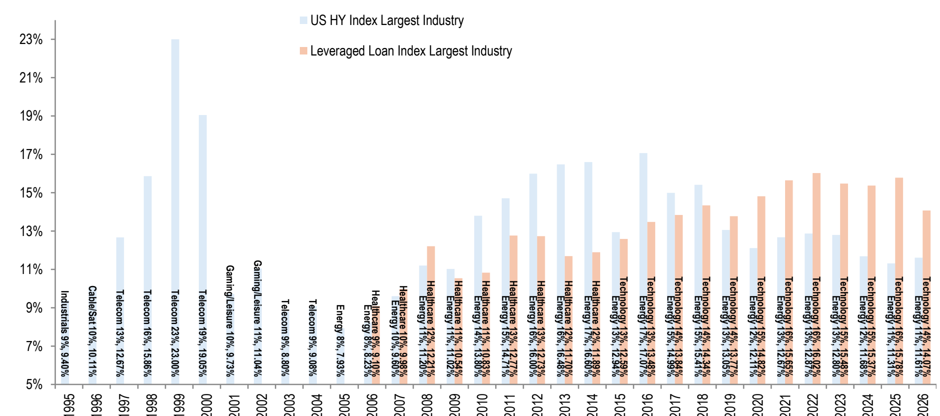

Source: J.P. Morgan.

Figure 44: Largest Industry Weightings in the High-Yield Bond and Loan Indices

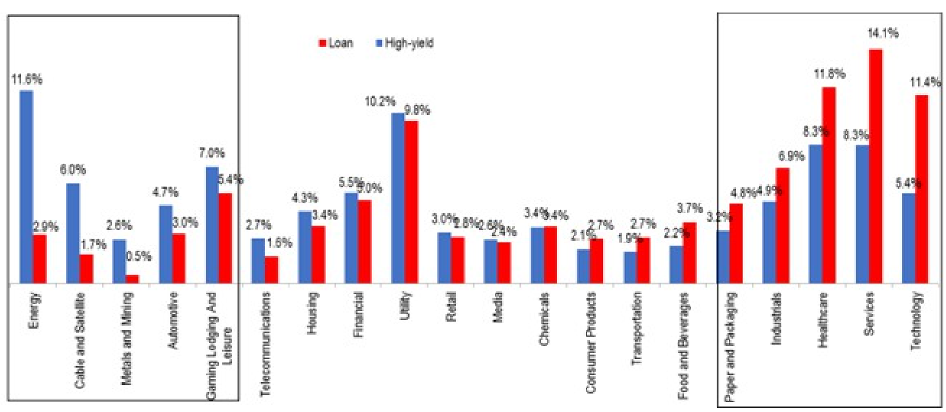

Source: J.P. Morgan.

Figure 45: From 2001 to 2003 75% of the Telecom sector defaulted

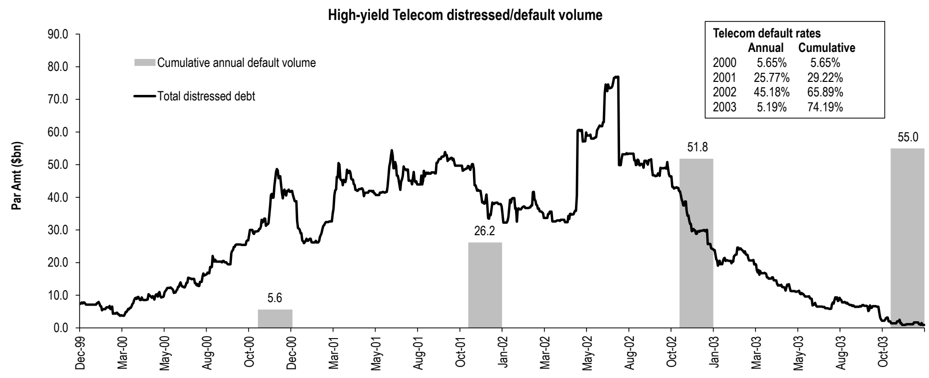

Source: J.P. Morgan.

Net Issuance Trends: To fund an AI capex buildout, leveraged finance is going to need to produce more net issuance. **For context, leveraged finance has produced below-average net issuance in each of the past four calendar years.** Specifically, high-yield originated new-issue volume ex-refi totaled $98bn in 2025, which was up from $71bn in 2024, $59bn in 2023, and $56bn in 2022. These figures equate to 3-6% of beginning outstanding. **For context, HY capital markets produced net volume totaling $153bn in 2020 and $193bn in 2021, or 12% each of beginning outstanding. Notably, high-yield net new issuance year-to-date totals $77bn, or well above pace of the last four years.** And the institutional loan market originated new-issue volume ex-refi/repricing totaling $199bn in 2025, which was up from $170bn in 2024, $82bn in 2023, and $163bn in 2022. These figures equate to 5-10% of beginning outstanding. **For context, loan capital markets produced net volume totaling $408bn in 2021, or 31% of beginning outstanding.** Year-to-date, loan new issuance ex-refi/repricing totals $100bn, or only slightly above 2025's pace, in contrast to bonds. Note between 2010 and 2021, high-yield and loan capital markets produced net issuance equating to an average 10% and 23% of beginning outstanding, respectively. **Certainly, there is scope for net issuance to double from 2025 levels (from $300bn in 2025 to $600bn) and by no means be unprecedented (2021 net volume $602bn).** Of course, a rise in M&A activity should also account for a sizeable piece of that growth (alongside AI funding). **Note $403bn of 2021's $602bn of net volumes were used for M&A/LBO with $199bn earmarked for general corporate purposes or to fund a dividend.**

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 46: Non-refi-related proceeds across bonds and loans could double over the next 12 months (+$300bn to $600bn) and still not be unprecedented

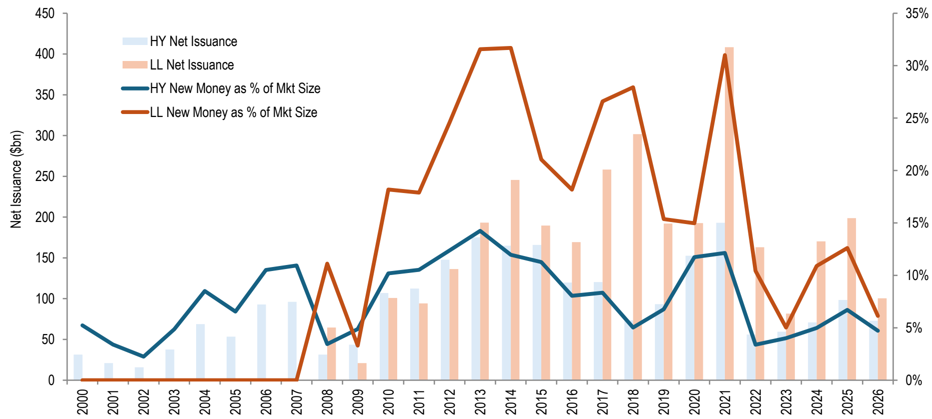

Source: J.P. Morgan.

Is there capacity to fund an AI capex buildout in leveraged finance? Yes, a lot, but is there enough? As mentioned previously, one could double the past 12 months' net-new issuance (non-refi) in syndicated markets over the next year and only reach 2021's M&A/LBO-induced net volumes of $602bn. The equivalent net new volumes equating to 2021's 21% of syndicated leveraged finance would translate into $672bn of net volume in today's terms. **There is also a record-high $543bn of dry powder in private credit drawdown funds in the latest reading from Preqin.** In other words, capital markets in Leveraged Finance appear well equipped to manage a rise in AI capex funding. That said, a lot will also depend on how much of that rise in net issuance is earmarked for new M&A/LBO funding (2021 67%). Where one could see resistance, however, is that a lot of these funding needs will be in the Technology sector. **Currently, Technology accounts for a sizeable 22% of the private credit universe and 14% of the leveraged loan universe.** As an aside, Technology represents 30% of the S&P 500 (market cap) and 15% of the Russell 2000. At 14% of the leveraged loan universe, Technology could expand another ~9% to reach the current weighting within private credit and the historical peak industry weighting within Leveraged Finance (Telecom 23% in 1999). **Despite Software accounting for the bulk of Technology loans and facing challenges, we expect Technology could easily grow to 17-18% of the leveraged loan universe over the coming years.** Where is there capacity for Technology to grow? **Technology is presently the fourth-largest industry at 8.3% of the HY bond index.** Assuming no industry growth elsewhere, the HY Technology sector could easily double in size over the next 5 years (+$100bn) and not reach the peak stretch for growth in the Energy sector post-GFC. For Leveraged Finance, our forecast for HPC issuance in 2026 is $80bn ($65bn HY, $15bn LL). **This would place the HPC subsector at about 4.5% of the High-Yield Bond index by year-end 2026, versus a modest 1% of the Leveraged Loan index.** And over the next five years (2026-2030), we estimate $350bn of HPC issuance syndicated across bonds and loans with a potential peak in 2027.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 47: Dry powder in private credit has risen to a record $543bn amid a 4yr slump in M&A activity

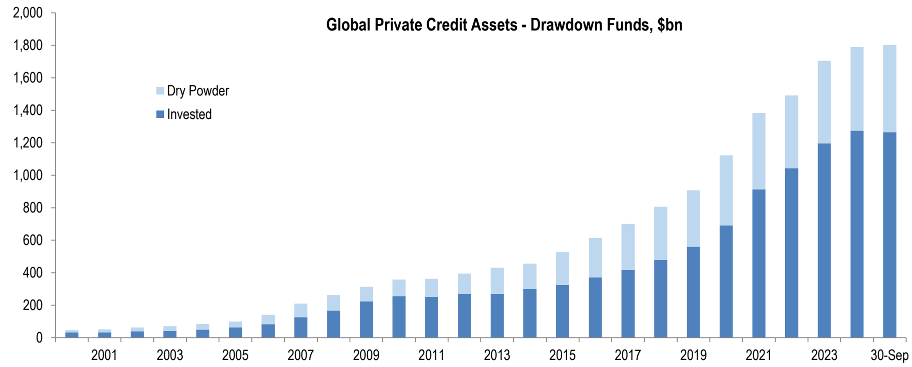

Source: Preqin.

Figure 48: Software accounts for 21% of the KBRA DLD Direct Lending Index

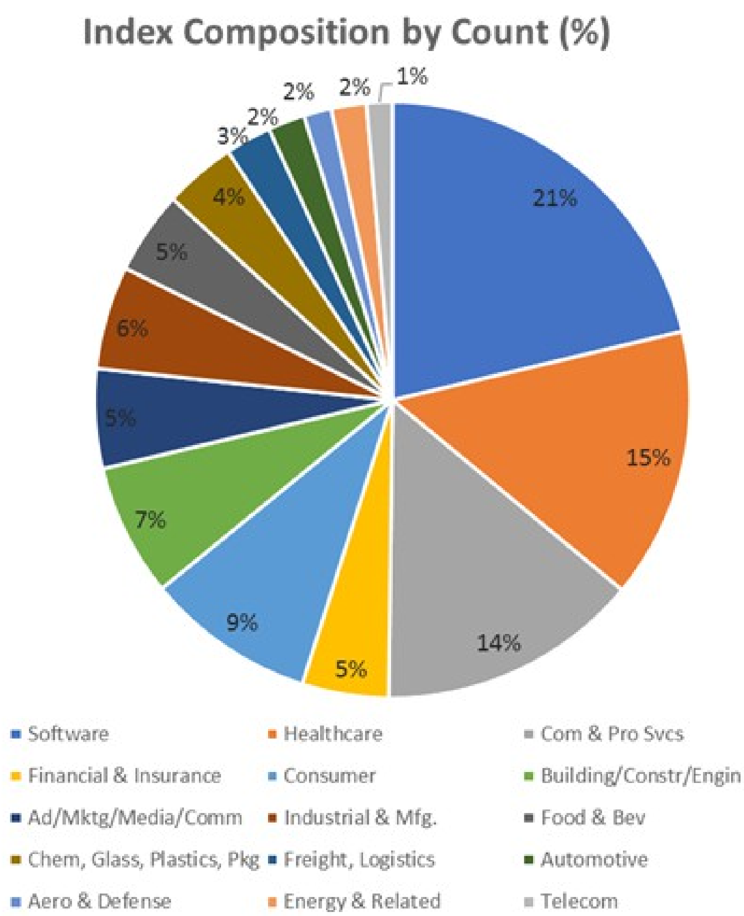

Source: KBRA DLD.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

SPG Growth More Deliberate

Chong Sin
(212) 834-2611
chong.c.sin@jpmorgan.com
J.P. Morgan Securities LLC

The US securitization market has issued $11.2bn of data center bonds across CMBS and ABS year-to-date ($13.5bn including ABS 4(a)(2) bonds) and remains on pace to hit our full-year target of $30-40bn (Figure 49). Securitizations continue to see issuance of bonds backed by fully occupied data centers leased to hyperscalers, enterprises, and retail, with CMBS more skewed to hyperscaler tenants and ABS more skewed to enterprise and retail colocation tenants (though ABS is seeing more hyperscalers recently). Investor skepticism remains significant as handicapping residual value risk remains a tough task. As such, only about 30% of regular single asset/ borrower (SASB) CMBS buyers have purchased data center SASB in the last 2 years. While money managers and insurance investors constitute most of data center securitization demand in CMBS and ABS, CMBS has garnered more alternative manager demand (Figure 50).

Figure 49: We continue to look for $30-40bn of US data center securitization issuance this year

US data center securitization issuance ($bn)

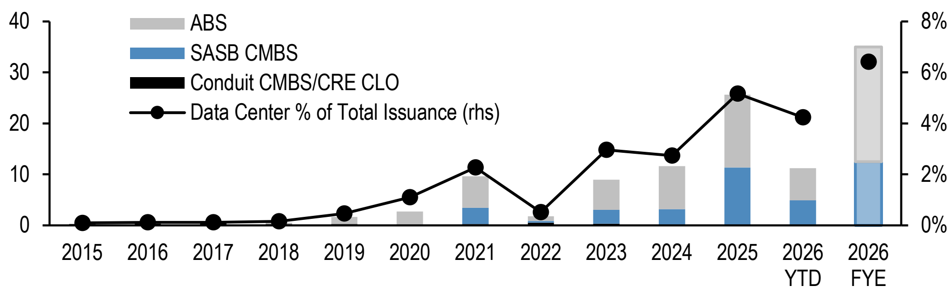

Source: J.P. Morgan, Bloomberg Finance L.P.

Figure 50: Real money accounts continue to dominate data center securitization demand

US data center securitization new issue participation by investor type

	Data Centers				All Types			
Investor Type	2025 CMBS	2026 YTD CMBS	2025 ABS	2026 YTD ABS	2025 SASB CMBS	2026 YTD SASB CMBS	2025 ABS	2026 YTD ABS
Money Manager	56%	37%	69%	81%	54%	52%	64%	66%
Insurance/Pension	9%	21%	22%	18%	14%	14%	21%	20%
Bank	5%	7%	1%	0%	4%	3%	8%	7%
Alternative Manager	30%	34%	6%	1%	28%	31%	4%	4%
Other	0%	0%	2%	0%	0%	0%	4%	3%
Top 5%	38%	48%	34%	64%	28%	28%	20%	21%
Top 10%	54%	72%	50%	84%	41%	43%	34%	33%

Note: Data is limited to deals in which J.P. Morgan played a syndication role.
Source: J.P. Morgan.

For securitization investors who are active in data centers, the underwriting focus falls squarely on tenant credit, lease terms including termination options, SLA termination clauses, location, and construction year. For hyperscale data center securitizations, investors are relying on the creditworthiness and durability of the long leases that extend far beyond the financing maturity or anticipated repayment date (ARD) for comfort. As such, the residual value question is largely being side-stepped.

In CMBS, however, two recent hyperscale-tenanted data center deals are challenging investors to think more about residual value risks. Both deals have

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

ARDs and a long tail to final maturity (9-10yrs) with final maturities that are coterminous with the IG hyperscale tenants' lease terms. The typical post-ARD cashflow sweep is as follows: 1) original interest payment, 2) principal payments until principal is fully paid down, 3) accrued step-up interest (typically a +2-3% step-up in CMBS). If the ARD is tripped, neither deal will have enough net cashflow from the contracted rents to fully amortize the AAA bonds, leaving outstanding principal + accrued interest as the refinancing basis. Should these bonds even be considered AAAs? Some argue "no" even if the risk of an ARD trip is small. Mitigating factors to the risk of tripping ARD include the starting leverage, termination covenants, and location. That said, these two deals have drastically different starting leverage, but it's in these details that the debate is occurring among informed investors.

But what about the ABS data center ARD? This has been a prominent feature of ABS data centers for a while and the ABS post-ARD tail is even longer (~25yrs). The difference is that for at least some ABS master trust issuers, losing ABS market access could be an existential problem given their heavy ABS financing dependency. In CMBS, the ARD option is layered on top of non-recourse property-level financing, which means borrowers can simply walk if the property economics fail. Taking a step back, however, it's questionable whether ARD is worth much at all as a structural feature as issuers will naturally honor the ARD when the market/issuer is functioning well, while they may not when the market/issuer is in crisis. It's not clear that CMBS or ABS investors are considering the long post-ARD extension scenarios when bidding on these bonds.

What seems clear is that leverage and structural optionality is being pushed in data center securitizations (at least in CMBS). Fortunately, this is against a market backdrop that is showing extremely strong fundamentals. As we recently noted, the for-lease North American data center market has grown by at least 48% year-over-year as of Q1 2026, and vacancy rates have fallen by 49bp to 1.37% (Figure 51)! In no other core commercial real estate market do we see this kind of strength. In CMBS, 83% of deals by outstanding balance are backed by properties that sit in markets with sub-1% vacancy rates (Figure 52). In the same publication, we proposed a checklist of factors investors should consider when weighing relative value across CMBS data centers.

Figure 51: North American leased commissioned data center capacity grew by at least 48% year-over-year, but, remarkably, vacancy rates fell

North American data center market size (GW) and vacancy rate (%), as of Q1 2026

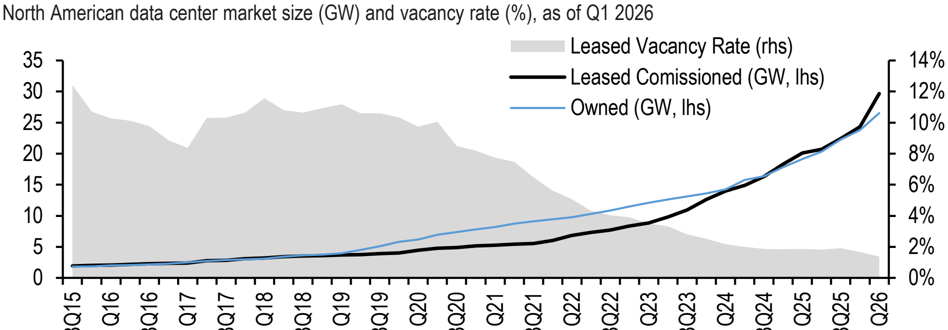

Note: datacenterHawk's North American leased commissioned total includes only established "insight" markets. Tertiary markets like those in West Texas have yet to be included.
Source: datacenterHawk, J.P. Morgan.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 52: The vast majority of outstanding SASB CMBS data centers sit in sub-1% vacancy data center markets

Outstanding SASB data center deals mapped to datacenterHawk markets by market vacancy rates as of Q1 2026

		datacenterHawk Market Vacancies					Deal Tenant Type Split	
Deal	**MW**	**<1%**	**1-2%**	**3-5%**	**5-10%**	**N/A**	**Hyperscale %**	**Hybrid Colocation %**
BX 2025-COPT	376	100%					100%	0%
BX 2025-VOLT	202	83%		6%		11%	34%	66%
BX 2025-VLT6	183	100%					71%	29%
CONE 2026-DFW3	153	100%					67%	33%
BX 2025-VLT7	138	100%					100%	0%
BX 2024-VLT4	122	66%			34%		35%	60%
BX 2026-VLT9	108	61%	39%				100%	0%
SWCH 2025-DATA	104	100%					0%	100%
BX 2024-VLT5	100	100%					100%	0%
DC 2023-DC	77	100%					100%	0%
DATA 2023-CNTR	67		100%				91%	9%
VDCM 2025-AZ	64		100%				100%	0%
JPMCC 2022-DATA	46	100%					100%	0%
CONE 2024-DFW1	45	100%					0%	100%
GSMS 2025-800D	30		100%				100%	0%
DATA 2024-CTR2	27		100%				100%	0%
All deals	**1,841**	**83%**	**13%**	**1%**	**2%**	**1%**	**74%**	**25%**

Source: datacenterHawk, deal documents, J.P. Morgan.

As a general view, we continue to like AAA data center SASB CMBS floaters that offer 164DM over 1m Term SOFR to full extension, a 24bp pickup to traditional industrial warehouse AAAs and a 7bp pickup to hotel AAAs (Figure 53). Our ABS strategists remain neutral on ABS data centers as other more established ABS sectors offer comparable spreads.

Figure 53: We continue to like AAA data center SASB floaters, which offer a 24bp pickup to industrial warehouse AAA floaters and a 7bp pickup to hotel AAAs

J.P. Morgan floating-rate SASB index discount margin to 1m Term SOFR to full extension (bp)

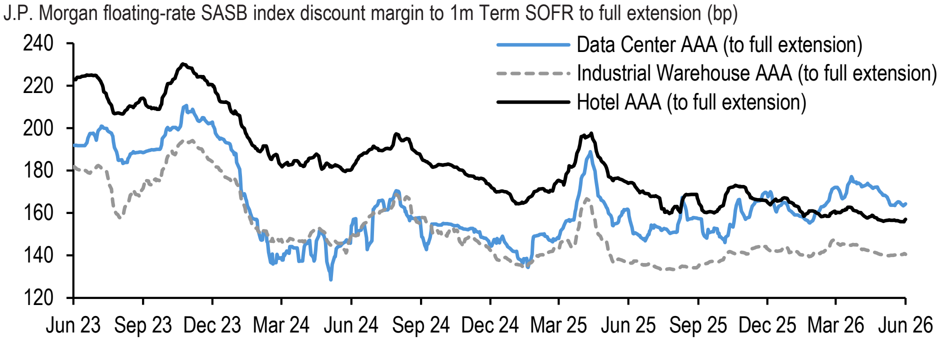

Source: J.P. Morgan.

The path to $250bn of data center securitization issuance by 2030 remains possible in our view, but the up- and down-side risks to this view are contingent on how successfully the market can absorb the HPC campus construction loans that are looking for a securitization exit. The mammoth 1GW+ campuses are difficult to

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

squeeze into the securitization market currently because of their sheer size. A 1GW campus would likely require a $10-15bn takeout financing today, and the largest deal the CMBS market has done to date was only $3.2bn and that was a portfolio of 10 buildings. Instead, tackling a mega campus construction loan take-out is likely to involve other sources, including corporates that have proven to be able to absorb large data center financings. Interestingly, we may see multiple sources of financing tackle a single mega campus construction loan takeout. Smaller data halls may end up in CMBS or ABS while the building housing the super computer may end up in corporates.

Alternative Asset Managers & Private Credit

Nelson Jantzen
(212) 270-1169
nelson.r.jantzen@jpmorgan.com
J.P. Morgan Securities LLC

Tarek Hamid
(212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC

Kabir Caprihan
(212) 834-5613
kabir.x.caprihan@jpmorgan.com
J.P. Morgan Securities LLC

The alternative asset industry continues to view digital infrastructure and the build-out of data centers as a core focus and has correctly surmised that the build-out cannot be funded by public funds or the leveraged loan market alone. **Fund raising for infrastructure assets has been very strong with AUMs of nearly $1.8tn. This asset class is one of the few bright spots for the alternative asset managers given the trials and tribulations within private equity and direct lending fund raising.** According to Preqin, there is also about $400bn of dry powder within the Infrastructure asset class which should support data center build-out. **Notably, most managers estimate that digital infrastructure accounts for nearly 20% of infrastructure assets and that is expected to grow given the funding needs of the asset class.** Every major manager views this as a multi-decade, structural opportunity with private capital providing a solution to fill the gap between hyperscaler demand and the available financing capacity. We continue to estimate a call on private credit/alternative capital of as much as $1.4tn and continue to believe that to be the case.

Figure 54: Infrastructure AUMs have increased to nearly $1.8tn

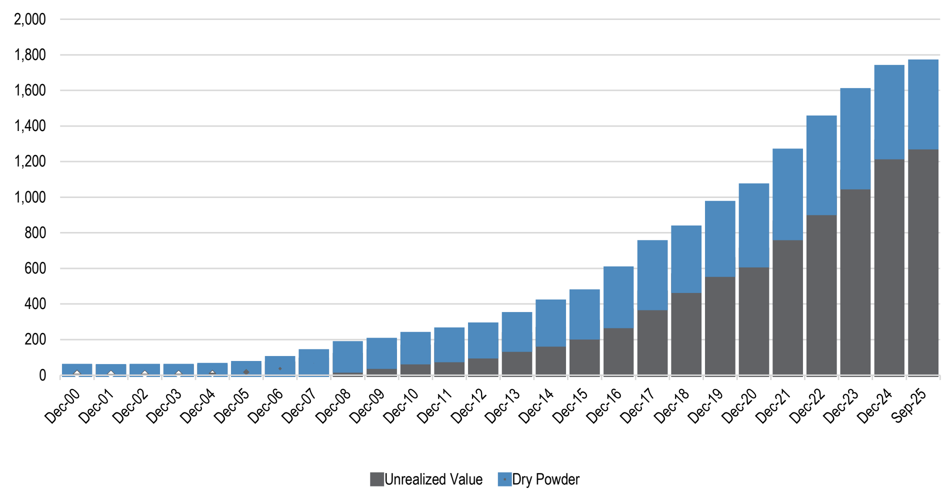

Source: Preqin and J.P. Morgan.

The approach from large asset managers varies across the board. While Blackstone and Blue Owl primarily participate in the build-out through equity investments, Apollo has taken the approach of providing credit and financing solutions. **One of the largest data center build-outs by private credit is the Hyperion project, where Blue Owl provided about $7bn of equity for a ~$27bn project.** Other large projects funded by private credit include the $10bn project for Firmus Technologies by a consortium of Blackstone managed funds. More recently, as discussed above, the $35bn financing led by Apollo and Blackstone for the Anthropic AI expansion is a potential milestone. Apollo and Blackstone are the anchor investors for the AI XPV platform, which was established by Broadcom to enable more than 20GW of compute capacity for Anthropic and OpenAI through 2028. Every manager is very focused on this asset class and whichever approach is taken, we do believe that the fund-raising momentum implies that $1.4tn target is highly achievable. Also note that some of the solutions are being funded on balance sheet, i.e., Apollo and KKR through their insurance affiliates, while

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

others rely on external funds; however, given the asset nature of these financings, they have access to ABF third-party money as well as the credit sleeve.

Blackstone is the largest private investor in data centers globally and operates with the most vertically integrated platform. The manager acquired QTS in 2021, and it is now the largest independent operator of data centers globally. Away from its development through QTS, recent transactions include a 49% stake in Rowan Digital Infrastructure at a $3.8bn valuation in the first quarter of 2026, the AirTrunk acquisition, a $25bn-plus Pennsylvania digital and energy hub through a joint venture with PPL, and a $5bn Google neocloud partnership tied to Anthropic. BREIT now carries 23% data center exposure. More recently, Blackstone filed for the Blackstone Digital Infrastructure Trust (BXDC), a public data center REIT targeting a $2bn raise. Infrastructure AUM, which includes data centers as well as energy assets, increased 41% y/y to nearly $84bn.

Apollo's approach is to provide large-scale, investment-grade private credits that are potentially too complex or bespoke for syndicated markets. In February and April 2026, Apollo led two AI-related financings totaling more than $8bn to support a client's acquisition and lease of data center infrastructure to a large investment-grade counterparty. The firm also structured a $3.5bn capital solution for Valor Compute Infrastructure, which supported a $5.4bn acquisition and lease of data center compute to an xAI subsidiary using a triple-net lease structure with NVIDIA as the anchor limited partner. While Apollo has not launched a dedicated data center fund, capital is being deployed through its credit strategies, which have about $50bn in dry powder.

Ares's acquisition of GCP added digital development capabilities, and the company is now raising a dedicated global data center equity fund with an expected close in summer 2026. The fund is seeded with GCP's existing portfolio and builds on Ares's track record of over $10bn historically deployed across digital infrastructure including towers, networks, and data centers. Available capital stands at over $158bn including undrawn lines of capital, with more than $100bn in credit dry powder.

Blue Owl's focus within digital infrastructure has been through joint-ventures with hyper scalers. Blue Owl is the partner for the Hyperion data center project. The $27bn project was funded with about $7bn of equity from Blue Owl managed funds, while the remaining capital was raised as project financing. The Blue Owl digital infrastructure platform (BODI) has invested $17.1 bn in digital infrastructure and has about 10.6 GW of power leased and owned.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Cross Product Relative Value Key

Silvi Mantri
(212) 834-7239
silvi.mantri@jpmchase.com
J.P. Morgan Securities LLC

The birth of a new universe of investment alternatives has created many different flavors of risk for investors to choose from. Nonetheless, almost every investment alternative links back in some way, shape or form to the core High Grade Hyperscalers and other giant High Grade participants. The gap between an obligation, a direct guarantee, a residual guarantee, a direct lease and an intermediated lease are profound from a credit (and accounting) standpoint. Nonetheless, particularly in our currently gross undersupplied compute market, spread to ultimate hyperscaler is a fairly compelling relative value metric, in our view. Below we show the relative value for each pair of High Performance Computing bond versus the corresponding hyperscaler. For IG, we compare bonds based on WAL and for HY, we show comparable bonds both on a YTM and YTC basis.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 55: IG and HY High Performance Computing versus Hyperscaler Bond Relative Value

Note: The blue rows are IG HPC bonds. The gray rows are HY HPCs which are compared on both YTM and *YTC* basis

Security	Issuer	Par amount outstanding $bn	Rating (Mdy / SP / Fitch)	Current G-Spread (bp)	Rating notch difference / Basis, bp
RDMICH 7 1/2 03/30/45	RD Michigan Property Owner I LLC	14.0	Baa3 / - / BBB	338	1 notches
ORCL 5 3/8 07/15/40	Oracle Corp	2.2	Baa2 / BBB / BBB	186	152
RPLDCI 6.581 05/30/49	Beignet Investor LLC	27.3	- / A+ / -	161	1 notches
META 6.2 05/15/46	Meta Platforms Inc	4.0	Aa3 / AA- / -	113	48
QTSQST 5.7 04/15/36	QTS Fayetteville	4.6	Baa2 / - / -	170	8 notches
MSFT 3.45 08/08/36	Microsoft Corp	1.7	Aaa / AAA / -	27	143
HUTRBA 6.192 11/15/42	Hut 8 DC LLC	3.3	- / BBB- / BBB-	152	7 notches
GOOGL 4.8 02/15/36	Alphabet Inc	4.3	Aa2 / AA+ / -	54	98
HUTBPA 6.129 11/30/42	Beacon Point DC LLC	4.3	Baa2 / - / -	149	5 notches
NVDA 3 1/2 04/01/40	Nvidia Corp	1.0	Aa1 / AA- / -	55	94
BLKPRL 6 1/8 02/15/31	Black Pearl Compute	2.0	Ba2 / - / BB-	158	9 notches
AMZN 4 1/4 03/13/31	Amazon.com Inc	5.0	A1 / AA / AA-	38	120
AMZN 3.85 03/13/28	*Amazon.com Inc*	*2.3*	*A1 / AA / AA-*	*12*	*146*
PFORGE 6 3/4 03/15/31	APLD ComputeCo 2 LLC	2.2	Ba3 / - / BB-	243	4 notches
ORCL 4.95 02/04/31	Oracle Corp	3.5	Baa2 / BBB / BBB	120	123
ORCL 2.3 03/25/28	*Oracle Corp*	*2.0*	*Baa2 / BBB / BBB*	*76*	*168*
FLASHC 7 1/4 12/31/30	Flash Compute LLC	1.3	Ba3 / - / -	210	10 notches
GOOGL 4.1 11/15/30	Alphabet Inc	2.5	Aa2 / AA+ / -	24	186
GOOGL 0.8 08/15/27	*Alphabet Inc*	*1.0*	*Aa2 / AA+ / -*	*-3*	*213*
WULF 7 3/4 10/15/30	Wulf Compute LLC	3.2	Ba2 / BB / BB	184	9 notches
GOOGL 4.1 11/15/30	Alphabet Inc	2.5	Aa2 / AA+ / -	24	160
GOOGL 0.8 08/15/27	*Alphabet Inc*	*1.0*	*Aa2 / AA+ / -*	*-3*	*187*
MERIDI 6 1/4 04/30/31	Meridian Arc Holdco LLC	5.7	Ba2 / BB- / BB	202	9 notches
GOOGL 4.1 02/15/31	Alphabet Inc	3.0	Aa2 / AA+ / -	28	174
GOOGL 3 7/8 11/15/28	*Alphabet Inc*	*1.0*	*Aa2 / AA+ / -*	*7*	*196*
CIFR 7 1/8 11/15/30	Cipher Compute LLC	1.7	Ba3 / - / BB-	190	10 notches
GOOGL 4.1 11/15/30	Alphabet Inc	2.5	Aa2 / AA+ / -	24	166
GOOGL 0.8 08/15/27	*Alphabet Inc*	*1.0*	*Aa2 / AA+ / -*	*-3*	*193*
TRACTC 5 7/8 03/01/31	SV RNO Property Owner 1	3.8	Ba1 / BB+ / BB	199	7 notches
NVDA 2 06/15/31	Nvidia Corp	1.3	Aa1 / AA- / -	18	181
NVDA 1.55 06/15/28	*Nvidia Corp*	*1.3*	*Aa1 / AA- / -*	*2*	*197*
TRACTD 6 1/2 05/01/31	PR RNO Property Owner 1	4.6	Ba1 / BB / BB	229	7 notches
NVDA 2 06/15/31	Nvidia Corp	1.3	Aa1 / AA- / -	18	211
NVDA 1.55 06/15/28	*Nvidia Corp*	*1.3*	*Aa1 / AA- / -*	*2*	*227*
APLD 9 1/4 12/15/30	APLD ComputeCo LLC	2.4	- / BB- / BB-	249	9 notches
META 4.2 11/15/30	Meta Platforms Inc	4.0	Aa3 / AA- / -	38	210
META 4.6 05/15/28	*Meta Platforms Inc*	*1.5*	*Aa3 / AA- / -*	*13*	*236*
CORZ 7 3/4 05/15/31	Core Scientific Finance	3.3	- / BB / BB-	304	-2 notches
CRWV 9 02/01/31	Coreweave Inc	1.8	B1 / B / BB-	470	-166
SECMOS 8 7/8 05/01/31	SE Cosmos LLC	1.0	- / BB+ / BB-	358	0 notches
SOFTBK 8 1/4 10/22/31	Softbank Group Corp	0.6	- / BB+ / -	323	35
SOFTBK 5 04/15/28	*Softbank Group Corp*	*1.1*	*Ba2 / BB+ / -*	*164*	*194*
EDGCOM 7 1/2 04/30/31	Edged Compute LLC	1.3	- / BB+ / BB-	361	-3 notches
CRWV 9 02/01/31	Coreweave Inc	1.8	B1 / B / BB-	470	-109
ELKGVP 7 1/2 06/15/31	Elk Grove Village Proper	0.9	Ba3 / BB- / BB-	298	-1 notches
CRWV 9 3/4 10/01/31	Coreweave Inc	2.8	B1 / B / BB-	502	-204
STNGRY 6 06/15/31	Stingray Compute	0.8	Ba2 / - / BB-	178	9 notches
AMZN 3.7 06/12/31	Amazon.Com Inc	2.5	A1 / AA / AA-	49	129
AMZN 1.65 05/12/28	*Amazon.Com Inc*	*2.3*	*A1 / AA / AA-*	*10*	*168*
ELNFOR 7 06/15/31	APLD ComputeCo 3 LLC	1.6	- / BB- / BB-	275	9 notches
META 4.55 05/15/31	Meta Platforms Inc	3.0	Aa3 / AA- / -	42	233
META 4.6 05/15/28	*Meta Platforms Inc*	*1.5*	*Aa3 / AA- / -*	*13*	*263*

Source: J.P. Morgan, Bloomberg Finance L.P.

The average relationship between a hyperscaler and an IG-rated project for which it is the lessor is 107bp. The average relationship between a hyperscaler and a HY-rated

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

project for which it is lessor is 172bp. The relationship between HY lessees and HY project bonds is complex. We suspect those relationships will always be nuanced, as investors weigh NeoCloud and GPU-heavy business models vs. relatively simple data center lessor business models.

Figure 56: Hyperscalers and the HPC Universe Credit Linkages Simplified

Note: Solid arrows: single-tenant projects. Dotted arrows: multi-tenant or intermediary structures; % reflects estimated revenue share

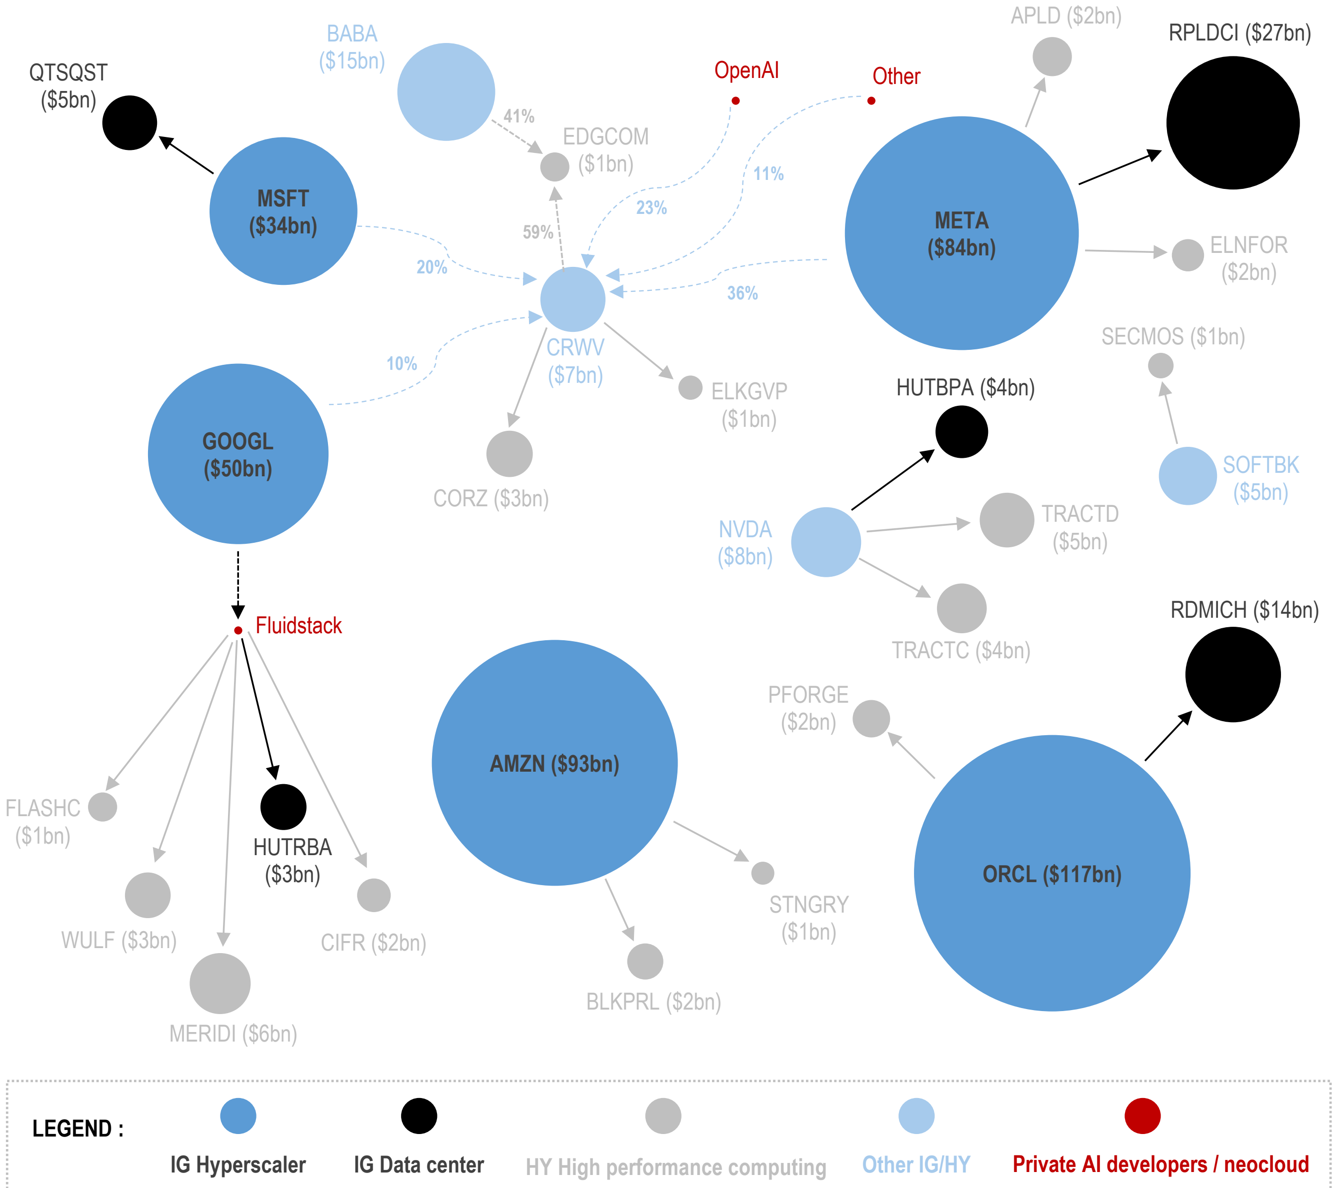

Source: J.P. Morgan, Bloomberg Finance L.P.

Monetization From Belief To Cost

Thomas Egan
(212) 270-2143
thomas.j.egan@jpmorgan.com
J.P. Morgan Securities LLC

Tarek Hamid
(212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC

In our November 2025 report, we discussed the daunting task of generating enough revenue to achieve a reasonable return on the trillions of dollars we expected would be spent on AI infrastructure. Back then we estimated that it would require annual economic benefit of $650bn to achieve a 10% return on those investments. To illustrate the difficulty in accomplishing that, we offered a somewhat imaginative example of the $35 per month charge that would be required of every iPhone user to hit that level of revenue. The point being that while it was clear that hyperscalers and AI labs could spend the money to build AI models, it was less clear exactly how well they could monetize those offerings.

Today, the appetite for AI is growing dramatically. Since the release of ChatGPT in late 2022, the world has seen an extraordinary acceleration in the consumption of AI "tokens" (the metered units of text and data processing used to price large language model inference). While many of the AI model providers do not regularly disclose customer token usage, there is plenty of evidence to suggest that usage is booming. For example, Google's CEO, Sundar Pichai, disclosed at Google I/O that monthly token usage across Google's AI platforms surged **sevenfold over the past year, reaching 3.2 quadrillion tokens.** And, according to OpenAI CEO Sam Altman, OpenAI's top internal user once consumed roughly 100,000 tokens per month several years ago—**that same top user now burns through approximately 100 billion tokens per month.** For an international context, China's daily token usage rose from100 billion at the start of 2024 to **140 trillion tokens per day** by March 2026, according to China Daily.

The growth in volumes has been accompanied by equally remarkable declines in the per-token price. Middle-of-the-road models used to cost roughly **$30 per million tokens in 2022-2023.** But competition from Anthropic, Google, and Chinese providers has now driven entry-level pricing down to a few cents per million tokens, while Enterprise blended cost per million tokens dropped from roughly $18.40 to $2.31 year-over-year—**an 87.4% reduction,** according to MENAFN. And those rates may go even lower, as the WSJ reported that both OpenAI and Anthropic are considering making significant cuts to token pricing. Figure 57 shows a representative set of current price-per-token rates for the top US models.

Tarek Hamid [AC]
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Figure 57: Representative AI Token Pricing

Provider	Model tier (2026)	Approx. input price (USD / 1M tokens)	Approx. output price (USD / 1M tokens)
OpenAI	GPT-5 (standard)	~1.25	~10
OpenAI	GPT-5.4 Nano / 5.x mini variants	~0.20–0.55	~1.25–4.40
OpenAI	GPT-5.5 / 5.x Pro reasoning	~5.00	~30.00
Anthropic	Claude Haiku 4.5	~1.00	~5.00
Anthropic	Claude Sonnet 4.6	~3.00	~15.00
Anthropic	Claude Opus 4.7/4.8	~5.00	~25.00
Anthropic	Claude Fable 5	~10.00	~50.00
Google	Gemini 3.1 Pro	~2.00	~12.00
Google	Gemini 3 Flash	~0.50	~3.00
xAI	Grok 4.1 Fast	~0.20	~0.50

Based on recent public pricing.
Source: Company reports.

Despite falling per-token prices, enterprise spending continues to rise sharply because volume growth is far outpacing price declines. In 2026, **71% of surveyed companies reported AI cost overruns**, and many large organizations are seeing monthly AI bills in the tens of millions of dollars, according to Reuters. Specific firm-level disclosures include Uber's CTO admitting that the company **exceeded its full-year 2026 AI budget by April**, Microsoft restricting LLM access for engineers and canceling Claude Code licenses to control costs, and Amazon and Meta each reportedly incurring hundreds of millions of dollars in token costs over short windows.

Why the dramatic increase in spending? The answer likely comes down to two things: "tokenmaxxing" and increased complexity. Tokenmaxxing is the practice of using as many tokens as possible to boost productivity regardless of whether that generates a good return on investment. This can come about when an organization encourages the increased use of AI by its workers without measuring what the specific gains are from that use. The increase in the number of complex tasks such as coding or creating agentic AI workflows—whereby AI agents carry out multi-step requests with minimal human intervention—are another reason for token volume expansion. Multiple independent sources confirm that agentic AI workflows consume dramatically more tokens than simple chatbot interactions with Time Online citing estimates that agentic workflows use **5 to 30 times more tokens.** The most precise figure we've seen,

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

according to MENAFN, comes from platform telemetry tracking of billions of API calls which indicate that the **median token consumption per completed agentic task is 23.4x higher** than for conversational workloads.

Figure 58: Agentic vs. Simple LLM Token Consumption

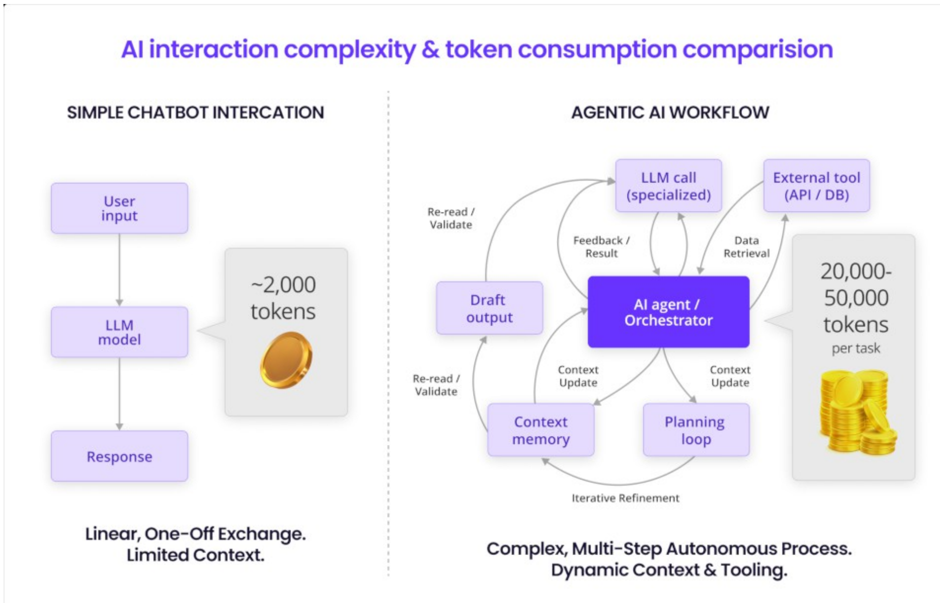

Source: Cockroach Labs.

The bottom line is that tokenmaxxing and complex AI requests appear to be driving usage volumes that more than offset any declines in token prices. That dynamic creates both an opportunity and a risk for hyperscalers and AI labs trying to earn a return on massive investments: The opportunity is higher revenue as usage grows, while the risk is that usage-based token pricing can squeeze customer budgets and, in turn, slow adoption. How these firms manage that tradeoff may be a key determinant of how effectively they hit their monetization goals.

Figure 59: Illustrative Token Economics

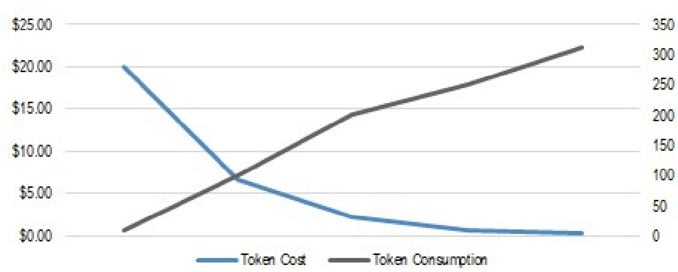

Source: J.P. Morgan.

Figure 60: Illustrative Cost to End Consumer Curve

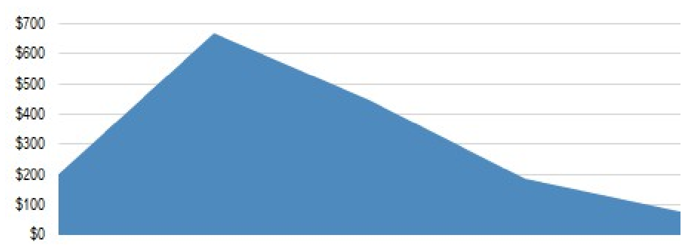

Source: J.P. Morgan.

Longer term, this likely points to a shift in pricing—potentially toward flat-rate tiers that let companies plan AI budgets more predictably. We saw a similar arc in earlier technology waves like the internet and wireless phone services: As prices dropped, adoption surged, lifting revenues, but not at the pace many had expected. Over time, flat-rate tiered plans became the standard model.

Tarek Hamid AC
(1-212) 834-5468
tarek.x.hamid@jpmorgan.com

Sensitivity analysis around monetization is extremely hard and more of an effort in 'what it would take'. Obviously, a key variable is how much is ultimately invested in the space, as well as how much profit and earnings are realized during the investment period prior to stabilization. Given the numbers that are being disclosed in the public markets, those current profit numbers are already highly relevant. Additionally, we have to assume a margin structure into perpetuity (we are assuming 50% stabilized for now, which obviously would not include additional reinvestment like growth capex or R&D, but would include maintenance). Perhaps more important is the 'pure margin' cost savings to hyperscalers and model developers, as some of the largest and most successful consumers already of AI. Additionally, it is probably relevant to think about the returns on equity given the amount of leverage being used to fund the buildout. That requires assumptions around tax rates (we are using 30%) and our existing assumptions around loan-to-cost. Regardless, given all of those highly compounding assumptions, the sector needs roughly $700 billion of revenue under our revised estimates to earn a 10% return on capital on the AI/data center buildout. That would imply an over 30% return on equity on the buildout with peak stabilized leverage of 7..4x. Given the size of the stabilized cash flow, those leverage numbers could fully amortize over 5-6 years.

Figure 61: Hypothetical AI Monetization Sensitivity Analysis

Hyperscaler / Model Developer Cost Savings ($bn)	Steady State AI Revenue ($bn)	Implied Return On Capital	Implied Return On Equity	Implied Peak AI Leverage
$200	$600	8%	29%	8.1x
$200	$700	10%	33%	7.4x
$200	$800	11%	37%	6.8x
$200	$900	12%	41%	6.3x
$200	$1,000	13%	44%	5.8x

Source: J.P. Morgan.

A final thought on monetization. It is entirely possible that the companies behind all the spending on AI will never achieve the hoped-for return on their AI infrastructure. We believe that was the case when, about 10 years ago, US wireless companies promised that spending billions to upgrade their networks to 5G would generate all sorts of new revenue from growth in the "Internet of Things" business. That incremental business never amounted to a material change in revenues, but the wireless companies likely had no choice but to make the investment for competitive reasons, given the loss of market share had they not done so. There are differences here, of course, as the level of spending and the revenue it is expected to generate are both significantly higher. But the idea that hyperscalers will need to invest just to keep pace still seems to make sense. The most obvious illustration of that dynamic is AI search—traditional Google search has declined, but AI search has surged, and Google's adoption of AI Mode and Gemini has increased overall searches on the platform.

Tarek Hamid AC (1-212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC
Harlan Sur AC (1-415) 315-6700
harlan.sur@jpmorgan.com

Nathaniel Rosenbaum, CFA AC (1-212) 834-2370
nathaniel.rosenbaum@jpmorgan.com

Nelson Jantzen, CFA AC (1-212) 270-1169
nelson.r.jantzen@jpmorgan.com

Companies Discussed in This Report (all prices in this report as of market close on 15 June 2026, unless otherwise indicated)
Advanced Micro Devices(AMD/$547.26/N), Alphabet(GOOG/$367.11/OW), Alphabet Inc.(GOOGL/$369.35/OW), Amazon.com(AMZN/$246.02/OW), Apollo Global Management Inc(APO/$136.11/NR), Ares Management(ARES/$135.36/OW), Blackstone(BX/$124.63/N), Blue Owl(OWL/$9.61/N), Broadcom Inc(AVGO/$393.94/OW), CoreWeave(CRWV/$106.71/N), IREN(IREN/$60.85/UW), Meta Platforms Inc(META/$593.48/N), Microsoft(MSFT/$399.76/OW), NVIDIA Corporation(NVDA/$212.45/OW), Oracle(ORCL/$192.64/OW)

Analyst Certification: The Research Analyst(s) denoted by an "AC" on the cover of this report certifies (or, where multiple Research Analysts are primarily responsible for this report, the Research Analyst denoted by an "AC" on the cover or within the document individually certifies, with respect to each security or issuer that the Research Analyst covers in this research) that: (1) all of the views expressed in this report accurately reflect the Research Analyst's personal views about any and all of the subject securities or issuers; and (2) no part of any of the Research Analyst's compensation was, is, or will be directly or indirectly related to the specific recommendations or views expressed by the Research Analyst(s) in this report. For all Korea-based Research Analysts listed on the front cover, if applicable, they also certify, as per KOFIA requirements, that the Research Analyst's analysis was made in good faith and that the views reflect the Research Analyst's own opinion, without undue influence or intervention.

All authors named within this report are Research Analysts who produce independent research unless otherwise specified. In Europe, Sector Specialists (Sales and Trading) may be shown on this report as contacts but are not authors of the report or part of the Research Department.

Important Disclosures

Company-Specific Disclosures: J.P. Morgan does and seeks to do business with companies 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. Important disclosures, including price charts and credit opinion history tables (if applicable), are available for compendium reports and all J.P. Morgan–covered companies, and certain non-covered companies, by visiting https://www.jpmm.com/research/disclosures, calling 1-800-477-0406, or e-mailing research.disclosure.inquiries@jpmorgan.com with your request.

Company-Specific Disclosures: Important disclosures, including price charts and credit opinion history tables (if applicable), are available for compendium reports and all J.P. Morgan–covered companies, and certain non-covered companies, by visiting https://www.jpmm.com/research/disclosures, calling 1-800-477-0406, or e-mailing research.disclosure.inquiries@jpmorgan.com with your request.

Explanation of Equity Research Ratings, Designations and Analyst(s) Coverage Universe:
J.P. Morgan uses the following rating system: Overweight (over the duration of the price target indicated in this report, we expect this stock will outperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe); Neutral (over the duration of the price target indicated in this report, we expect this stock will perform in line with the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe); and Underweight (over the duration of the price target indicated in this report, we expect this stock will underperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe. NR is Not Rated. In this case, J.P. Morgan has removed the rating and, if applicable, the price target, for this stock because of either a lack of a sufficient fundamental basis or for legal, regulatory or policy reasons. The previous rating and, if applicable, the price target, no longer should be relied upon. An NR designation is not a recommendation or a rating. Some stocks under coverage have a rating but no price target; in these cases, we expect the stock will outperform/perform in line/underperform the average total return of the stocks in the Research Analyst's, or the Research Analyst's team's, coverage universe of the relevant duration of the region. In our Asia (ex-Australia and ex-India) and U.K. small- and mid-cap Equity Research, each stock's expected total return is compared to the expected total return of a benchmark country market index, not to those Research Analysts' coverage universe. If it does not appear in the Important Disclosures section of this report, the certifying Research Analyst's coverage universe can be found on J.P. Morgan's Research website, https://www.jpmorganmarkets.com.

Coverage Universe: Sur, Harlan : Advanced Micro Devices (AMD), Analog Devices (ADI), Applied Materials (AMAT), Arm Holdings Plc (ARM), Astera Labs Inc (ALAB), Broadcom Inc (AVGO), Cadence Design Systems (CDNS), GlobalFoundries (GFS), Intel (INTC), KLA Corporation (KLAC), Lam Research (LRCX), Marvell Technology Inc (MRVL), Microchip Technology (MCHP), Micron Technology (MU), NVIDIA Corporation (NVDA), NXP Semiconductors (NXPI), ON Semiconductor Corporation (ON), Sandisk (SNDK), Synopsys Inc (SNPS), Texas Instruments (TXN)**Tonet, Jeremy** : AltaGas Ltd (ALA.TO), Ameren Corporation (AEE), American Electric Power (AEP), American Water Works Company, Inc. (AWK), Antero Midstream Corp (AM), Atmos Energy (ATO), CMS Energy Corporation (CMS), CenterPoint Energy, Inc. (CNP), Cheniere Energy Partners, L.P. (CQP), Cheniere Energy, Inc. (LNG), Consolidated Edison Inc. (ED), Constellation Energy (CEG), DT Midstream Inc (DTM), DTE Energy Company (DTE), Dominion Energy Inc (D), Duke Energy Corp. (DUK), Enbridge Inc (ENB.TO), Energy Transfer LP (ET), Entergy Corp. (ETR), Enterprise Products Partners L.P. (EPD), Eversource Energy (ES), Exelon Corp. (EXC), FirstEnergy (FE), Gibson Energy Inc. (GEI.TO), H2O America (HTO), Hess Midstream LP (HESM), Kinder Morgan, Inc. (KMI), Kinetik Holdings Inc (KNTK), MPLX LP (MPLX), NextEra Energy Inc. (NEE), ONEOK, Inc. (OKE), Oklo Inc. (OKLO), PPL Corporation (PPL), Pembina Pipeline Corp (PPL.TO), Plains All American Pipeline, L.P. (PAA), Plains GP Holdings, L.P. (PAGP), Public Service Enterprise Group (PEG), Rockpoint Gas Storage Inc. (RGSI.TO), South Bow Corporation (SOBO), Sunoco L.P. (SUN), TC Energy Corp (TRP.TO), Talen Energy (TLN), Targa Resources Corp. (TRGP), The Southern Company (SO), The Williams Companies, Inc. (WMB), Venture Global, Inc.

Tarek Hamid AC (1-212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC
Harlan Sur AC (1-415) 315-6700
harlan.sur@jpmorgan.com
Nathaniel Rosenbaum, CFA AC (1-212) 834-2370
nathaniel.rosenbaum@jpmorgan.com
Nelson Jantzen, CFA AC (1-212) 270-1169
nelson.r.jantzen@jpmorgan.com

(VG), Vistra (VST), WEC Energy Group (WEC), Western Midstream Partners LP (WES), X-energy, Inc. (XE), Xcel Energy (XEL)

J.P. Morgan Equity Research Ratings Distribution, as of April 04, 2026

	Overweight (buy)	Neutral (hold)	Underweight (sell)
J.P. Morgan Global Equity Research Coverage*	51%	37%	12%
IB clients**	83%	79%	74%
JPMS Equity Research Coverage*	49%	39%	13%
IB clients**	94%	93%	85%

***Please note that the percentages may not add to 100% because of rounding.**
****Percentage of subject companies within each of the "buy," "hold" and "sell" categories for which J.P. Morgan has provided investment banking services within the previous 12 months.**
For purposes of FINRA ratings distribution rules only, our Overweight rating falls into a buy rating category; our Neutral rating falls into a hold rating category; and our Underweight rating falls into a sell rating category. Please note that stocks with an NR designation are not included in the table above. This information is current as of the end of the most recent calendar quarter.

Equity Valuation and Risks: For valuation methodology and risks associated with covered companies or price targets for covered companies, please see the most recent company-specific research report at http://www.jpmorganmarkets.com, contact the primary analyst or your J.P. Morgan representative, or email research.disclosure.inquiries@jpmorgan.com. For material information about the proprietary models used, please see the Summary of Financials in company-specific research reports and the Company Tearsheets, which are available to download on the company pages of our client website, http://www.jpmorganmarkets.com. This report also sets out within it the material underlying assumptions used.

History of Investment Recommendations:
A history of J.P. Morgan investment recommendations disseminated during the preceding 12 months can be accessed on the Research & Commentary page of http://www.jpmorganmarkets.com where you can also search by analyst name, sector or financial instrument.

Explanation of Credit Research Valuation Methodology, Ratings and Risk to Ratings:
J.P. Morgan uses a bond-level rating system that incorporates valuations (relative value) and our fundamental view on the security. Our fundamental credit view of an issuer is based on the company's underlying credit trends, overall creditworthiness and our opinion on whether the issuer will be able to service its debt obligations when they become due and payable. We analyze, among other things, the company's cash flow capacity and trends and standard credit ratios, such as gross and net leverage, interest coverage and liquidity ratios. We also analyze profitability, capitalization and asset quality, among other variables, when assessing financials. Analysts also rate the issuer, based on the rating of the benchmark or representative security. Unless we specify a different recommendation for the company’s individual securities, an issuer recommendation applies to all of the bonds at the same level of the issuer’s capital structure. We may also rate certain loans and preferred securities, as applicable. This report also sets out within it the material underlying assumptions used. We use the following ratings for bonds (issues), issuers, loans, and preferred securities: Overweight (over the next three months, the recommended risk position is expected to outperform the relevant index, sector, or benchmark); Neutral (over the next three months, the recommended risk position is expected to perform in line with the relevant index, sector, or benchmark); and Underweight (over the next three months, the recommended risk position is expected to underperform the relevant index, sector, or benchmark). For CDS, we use the following rating system: Long Risk (over the next three months, the credit return on the recommended position is expected to exceed the relevant index, sector or benchmark); Neutral (over the next three months, the credit return on the recommended position is expected to match the relevant index, sector or benchmark); and Short Risk (over the next three months, the credit return on the recommended position is expected to underperform the relevant index, sector or benchmark).

J.P. Morgan Credit Research Ratings Distribution, as of April 04, 2026

	Overweight (buy)	Neutral (hold)	Underweight (sell)
Global Credit Research Universe*	27%	57%	16%
IB clients**	84%	84%	89%

***Please note that the percentages may not add to 100% because of rounding.**
****Percentage of subject companies within each of the "Overweight," "Neutral" and "Underweight" categories for which J.P. Morgan has provided investment banking services within the previous 12 months.**
For purposes of FINRA ratings distribution rules only, our Overweight rating falls into a buy rating category; our Neutral rating falls into a hold rating category; and our Underweight rating falls into a sell rating category. The Credit Research Rating Distribution is at the issuer level. Issuers with an NR or an NC designation are not included in the table above. This information is current as of the end of the most recent calendar quarter.

Analysts' Compensation: The research analysts responsible for the preparation of this report receive compensation based upon various factors, including the quality and accuracy of research, client feedback, competitive factors, and overall firm revenues.

Registration of non-US Analysts: Unless otherwise noted, the non-US analysts listed on the front of this report are employees of non-US affiliates of J.P. Morgan Securities LLC, may not be registered as research analysts under FINRA rules, may not be associated persons of J.P. Morgan Securities LLC, and may not be subject to FINRA Rule 2241 or 2242 restrictions on communications with covered companies, public appearances, and trading securities held by a research analyst account.

Other Disclosures

J.P. Morgan is a marketing name for investment banking businesses of JPMorgan Chase & Co. and its subsidiaries and affiliates worldwide.

UK MIFID FICC research unbundling exemption: UK clients should refer to UK MIFID Research Unbundling exemption for details of J.P. Morgan's implementation of the FICC research exemption and guidance on relevant FICC research categorisation.

All research material made available to clients are simultaneously available on our client website, J.P. Morgan Markets, unless specifically permitted by relevant laws. Not all research content is redistributed, e-mailed or made available to third-party aggregators. For all research material available on a particular stock, please contact your sales representative.

Any long form nomenclature for references to China; Hong Kong; Taiwan; and Macau within this research material are Mainland China; Hong Kong SAR (China); Taiwan (China); and Macau SAR (China).

J.P. Morgan Research may, from time to time, write on issuers or securities targeted by economic or financial sanctions imposed or administered by the governmental authorities of the U.S., EU, UK or other relevant jurisdictions (Sanctioned Securities). Nothing in this report is intended to be read or construed as encouraging, facilitating, promoting or otherwise approving investment or dealing in such Sanctioned Securities. Clients should be aware of their own legal and compliance obligations when making investment decisions.

Any digital or crypto assets discussed in this research report are subject to a rapidly changing regulatory landscape. For relevant regulatory advisories on crypto assets, including bitcoin and ether, please see https://www.jpmorgan.com/disclosures/cryptoasset-disclosure.

The author(s) of this research report may not be licensed to carry on regulated activities in your jurisdiction and, if not licensed, do not hold themselves out as being able to do so.

Exchange-Traded Funds (ETFs): J.P. Morgan Securities LLC ("JPMS") acts as authorized participant for substantially all U.S.-listed ETFs. To the extent that any ETFs are mentioned in this report, JPMS may earn commissions and transaction-based compensation in connection with the distribution of those ETF shares and may earn fees for performing other trade-related services, such as securities lending to short sellers of the ETF shares. JPMS may also perform services for the ETFs themselves, including acting as a broker or dealer to the ETFs. In addition, affiliates of JPMS may perform services for the ETFs, including trust, custodial, administration, lending, index calculation and/or maintenance and other services.

Options and Futures related research: If the information contained herein regards options- or futures-related research, such information is available only to persons who have received the proper options or futures risk disclosure documents. Please contact your J.P. Morgan Representative or visit https://www.theocc.com/components/docs/riskstoc.pdf for a copy of the Option Clearing Corporation's Characteristics and Risks of Standardized Options or https://www.finra.org/sites/default/files/2020-08/Security_Futures_Risk_Disclosure_Statement_2020.pdf for a copy of the Security Futures Risk Disclosure Statement.

Changes to Interbank Offered Rates (IBORs) and other benchmark rates: Certain interest rate benchmarks are, or may in the future become, subject to ongoing international, national and other regulatory guidance, reform and proposals for reform. For more information, please consult: https://www.jpmorgan.com/global/disclosures/interbank_offered_rates

Private Bank Clients: Where you are receiving research as a client of the private banking businesses offered by JPMorgan Chase & Co. and its subsidiaries ("J.P. Morgan Private Bank"), research is provided to you by J.P. Morgan Private Bank and not by any other division of J.P. Morgan, including, but not limited to, the J.P. Morgan Corporate and Investment Bank and its Global Research division.

Legal entity responsible for the production and distribution of research: The legal entity identified below the name of the Reg AC Research Analyst who authored this material is the legal entity responsible for the production of this research. Where multiple Reg AC Research Analysts authored this material with different legal entities identified below their names, these legal entities are jointly responsible for the production of this research. Where more than one legal entity is listed under an analyst's name, the first legal entity is responsible for the production unless stated otherwise. Research Analysts from various J.P. Morgan affiliates may have contributed to the production of this material but may not be licensed to carry out regulated activities in your jurisdiction (and do not hold themselves out as being able to do so). Unless otherwise stated below in the legal entity disclosures, this material has been distributed by the legal entity responsible for production, or where more than one legal entity is listed under the analyst's name, the first legal entity will be responsible for distribution. If you have any queries, please contact the relevant Research Analyst in your jurisdiction or the entity in your jurisdiction that has distributed this research material.

Legal Entities Disclosures and Country-/Region-Specific Disclosures:
Argentina: JPMorgan Chase Bank N.A Sucursal Buenos Aires is regulated by Banco Central de la República Argentina ("BCRA"- Central Bank of Argentina) and Comisión Nacional de Valores ("CNV"- Argentinian Securities Commission - ALYC y AN Integral N°51).

Australia: J.P. Morgan Securities Australia Limited ("JPMSAL") (ABN 61 003 245 234/AFS Licence No: 238066) is regulated by the Australian Securities and Investments Commission and is a Market Participant of ASX Limited, a Clearing and Settlement Participant of ASX Clear Pty Limited and a Clearing Participant of ASX Clear (Futures) Pty Limited. This material is issued and distributed in Australia by or on behalf of JPMSAL only to "wholesale clients" (as defined in section 761G of the Corporations Act 2001). A list of all financial products covered can be found by visiting https://www.jpmm.com/research/disclosures. J.P. Morgan seeks to cover companies of relevance to the domestic and international investor base across all Global Industry Classification Standard (GICS) sectors, as well as across a range of market capitalisation sizes. If applicable, in the course of conducting public side due diligence on the subject company(ies), the Research Analyst team may at times perform such diligence through corporate engagements such as site visits, discussions with company representatives, management presentations, etc. Research issued by JPMSAL has been prepared in accordance with J.P. Morgan Australia's Research Independence Policy which can be found at the following link: J.P. Morgan Australia - Research Independence Policy.

Brazil: Banco J.P. Morgan S.A. is regulated by the Comissao de Valores Mobiliarios (CVM) and by the Central Bank of Brazil. Ombudsman J.P. Morgan: 0800-7700847 / 0800-7700810 (For Hearing Impaired) / ouvidoria.jp.morgan@jpmchase.com.

Canada: J.P. Morgan Securities Canada Inc. is a registered investment dealer, regulated by the Canadian Investment Regulatory Organization and the Ontario Securities Commission and is the participating member on Canadian exchanges. This material is distributed in Canada by or on behalf of J.P.Morgan Securities Canada Inc.

Chile: Inversiones J.P. Morgan Limitada is an unregulated entity incorporated in Chile.

China: J.P. Morgan Securities (China) Company Limited has been approved by CSRC to conduct the securities investment consultancy business.

Colombia: Banco J.P. Morgan Colombia S.A. is supervised by the Superintendencia Financiera de Colombia (SFC). Any reference in this material to products or services offered abroad by entities other than the Bank in Colombia is included exclusively for descriptive purposes. Such references do not constitute, and should not be construed as, promotional activity or the provision of financial products or services within Colombian territory, as defined under applicable Colombian regulation.

Dubai International Financial Centre (DIFC): JPMorgan Chase Bank, N.A., Dubai Branch is regulated by the Dubai Financial Services Authority (DFSA) and its registered address is Dubai International Financial Centre - The Gate, West Wing, Level 3 and 9 PO Box 506551, Dubai, UAE. This material has been distributed by JP Morgan Chase Bank, N.A., Dubai Branch to persons regarded as professional clients or market counterparties as defined under the DFSA rules.

European Economic Area (EEA): Unless specified to the contrary, research is distributed in the EEA by J.P. Morgan SE ("JPM SE"), which is authorised as a credit institution by the Federal Financial Supervisory Authority (Bundesanstalt für Finanzdienstleistungsaufsicht, BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). JPM SE is a company headquartered in Frankfurt with registered address at TaunusTurm, Taunustor 1, Frankfurt am Main, 60310, Germany. The material has been distributed in the EEA to persons regarded as professional investors (or equivalent) pursuant to Art. 4 para. 1 no. 10 and Annex II of MiFID II and its respective implementation in their home jurisdictions ("EEA professional investors"). This material must not be acted on or relied on by persons who are not EEA professional investors. Any investment or investment activity to which this material relates is only available to EEA relevant persons and will be engaged in only with EEA relevant persons.

Hong Kong: J.P. Morgan Securities (Asia Pacific) Limited (CE number AAJ321) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission in Hong Kong, and J.P. Morgan Broking (Hong Kong) Limited (CE number AAB027) is regulated by the Securities and Futures Commission in Hong Kong. JP Morgan Chase Bank, N.A., Hong Kong Branch (CE Number AAL996) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission, is organized under the laws of the United States with limited liability. Where the distribution of this material is a regulated activity in Hong Kong, the material is distributed in Hong Kong by or through J.P. Morgan Securities (Asia Pacific) Limited and/or J.P. Morgan Broking (Hong Kong) Limited.

India: J.P. Morgan India Private Limited (Corporate Identity Number - U67120MH1992FTC068724), having its registered office at J.P. Morgan Tower, Off. C.S.T. Road, Kalina, Santacruz - East, Mumbai – 400098, is registered with the Securities and Exchange Board of India (SEBI) as a 'Research Analyst' having registration number INH000001873. J.P. Morgan India Private Limited is also registered with SEBI as a member of the National Stock Exchange of India Limited and the Bombay Stock Exchange Limited (SEBI Registration Number – INZ000239730) and as a Merchant Banker (SEBI Registration Number - MB/INM000002970). Telephone: 91-22-6157 3000, Facsimile: 91-22-6157 3990 and Website: http://www.jpmipl.com. JPMorgan Chase Bank, N.A. - Mumbai Branch is licensed by the Reserve Bank of India (RBI) (Licence No. 53/ Licence No. BY.4/94; SEBI - IN/CUS/014/ CDSL : IN-DP-CDSL-444-2008/ IN-DP-NSDL-285-2008/ INBI00000984/ INE231311239) as a Scheduled Commercial Bank in India, which is its primary license allowing it to carry on Banking business in India and other activities, which a Bank branch in India are permitted to undertake. For non-local research material, this material is not distributed in India by J.P. Morgan India Private Limited. Compliance Officer: Ashutosh Sharma; ashutosh.j.sharma@jpmchase.com; +912261575002. Grievance Officer: Ramprasadh K, jpmipl.research.feedback@jpmorgan.com; +912261573000. Registration granted by SEBI and certification from NISM in no way guarantee performance of the intermediary or provide any assurance of returns to investors. Please visit Terms and Conditions and Most Important Terms and Conditions (MITC). The annual Compliance audit report is available at http://www.jpmipl.com/#research.

Indonesia: PT J.P. Morgan Sekuritas Indonesia is a member of the Indonesia Stock Exchange and is registered and supervised by the Otoritas Jasa Keuangan (OJK).

Korea: J.P. Morgan Securities (Far East) Limited, Seoul Branch, is a member of the Korea Exchange (KRX). JPMorgan Chase Bank, N.A., Seoul Branch, is licensed as a branch office of foreign bank (JPMorgan Chase Bank, N.A.) in Korea. Both entities are regulated by the Financial Services Commission (FSC) and the Financial Supervisory Service (FSS). For non-macro research material, the material is distributed in Korea by or through J.P. Morgan Securities (Far East) Limited, Seoul Branch.

Japan: JPMorgan Securities Japan Co., Ltd. and JPMorgan Chase Bank, N.A., Tokyo Branch are regulated by the Financial Services Agency in Japan.

Malaysia: This material is issued and distributed in Malaysia by JPMorgan Securities (Malaysia) Sdn Bhd (18146-X), which is a Participating Organization of Bursa Malaysia Berhad and holds a Capital Markets Services License issued by the Securities Commission in Malaysia.

Mexico: J.P. Morgan Casa de Bolsa, S.A. de C.V., J.P. Morgan Grupo Financiero is member of the Mexican Stock Exchange (“Bolsa Mexicana de Valores”) and the Institutional Stock Exchange (“Bolsa Institucional de Valores”), and it is authorized to act as a broker dealer by the National Banking and Securities Exchange Commission (“Comisión Nacional Bancaria y de Valores”).

New Zealand: This material is issued and distributed by JPMSAL in New Zealand only to "wholesale clients" (as defined in the Financial Markets Conduct Act 2013). JPMSAL is registered as a Financial Service Provider under the Financial Service providers (Registration and Dispute Resolution) Act of 2008.

Philippines: J.P. Morgan Securities Philippines Inc. is a Trading Participant of the Philippine Stock Exchange and a member of the Securities Clearing Corporation of the Philippines and the Securities Investor Protection Fund. It is regulated by the Securities and Exchange Commission.

Singapore: This material is issued and distributed in Singapore by or through J.P. Morgan Securities Singapore Private Limited (JPMSS) [MDDI (P) 057/08/2025 and Co. Reg. No.: 199405335R], which is a member of the Singapore Exchange Securities Trading Limited, and/or JPMorgan Chase Bank, N.A., Singapore branch (JPMCB Singapore), both of which are regulated by the Monetary Authority of Singapore. This material is issued and distributed in Singapore only to accredited investors, expert investors and institutional investors, as defined in Section 4A of the Securities and Futures Act, Cap. 289 (SFA). This material is not intended to be issued or distributed to any retail investors or any other investors that do not fall into the classes of “accredited investors,” “expert investors” or “institutional investors,” as defined under Section 4A of the SFA. Recipients of this material in Singapore are to contact JPMSS or JPMCB Singapore in respect of any matters arising from, or in connection with, the material.

South Africa: J.P. Morgan Equities South Africa Proprietary Limited and JPMorgan Chase Bank, N.A., Johannesburg Branch are members of the Johannesburg Securities Exchange and are regulated by the Financial Services Conduct Authority (FSCA).

Taiwan: J.P. Morgan Securities (Taiwan) Limited is a participant of the Taiwan Stock Exchange (company-type) and regulated by the Taiwan Securities and Futures Bureau. Material relating to equity securities is issued and distributed in Taiwan by J.P. Morgan Securities (Taiwan) Limited, subject to the license scope and the applicable laws and the regulations in Taiwan. **To the extent that J.P. Morgan Securities (Taiwan) Limited produces research materials on securities not listed on the Taiwan Stock Exchange or Taipei Exchange (“Non-Taiwan Listed Securities”), these materials shall not constitute securities recommendations for the purpose of applicable Taiwan regulations, and, for the avoidance of doubt, J.P. Morgan Securities (Taiwan) Limited does not act as broker for Non-Taiwan Listed Securities.** According to Paragraph 2, Article 7-1 of Operational Regulations Governing Securities Firms Recommending Trades in Securities to Customers (as amended or supplemented) and/or other applicable laws or regulations, please note that the recipient of this material is not permitted to engage in any activities in connection with the material that may give rise to conflicts of interests, unless otherwise disclosed in the “Important Disclosures” in this material.

Thailand: This material is issued and distributed in Thailand by JPMorgan Securities (Thailand) Ltd., which is a member of the Stock Exchange of Thailand and is regulated by the Ministry of Finance and the Securities and Exchange Commission. The registered address is 548 One City Center Building, 50th Floor, Ploenchit Road, Lumphini, Pathum Wan, Bangkok 10330.

UK: Research is produced in the UK by J.P. Morgan Securities plc (“JPMS plc”) which is a member of the London Stock Exchange and is authorised by the Prudential Regulation Authority and regulated by the Financial Conduct Authority and the Prudential Regulation Authority or J.P. Morgan Markets Limited (“JPMML Ltd”) which is authorised and regulated by the Financial Conduct Authority. Unless specified to the contrary, this material is distributed in the UK by JPMS plc and is directed in the UK only to: (a) persons having professional experience in matters relating to investments falling within article 19(5) of the Financial Services and Markets Act 2000 (Financial Promotion) (Order) 2005 (“the FPO”); (b) persons outlined in article 49 of the FPO (high net worth companies, unincorporated associations or partnerships, the trustees of high value trusts, etc.); or (c) any persons to whom this communication may otherwise lawfully be made; all such persons being referred to as "UK relevant persons". This material must not be acted on or relied on by persons who are not UK relevant persons. Any investment or investment activity to which this material relates is only available to UK relevant persons and will be engaged in only with UK relevant persons. A description of J.P. Morgan EMEA’s policy for prevention and avoidance of conflicts of interest related to the production of Research can be found at the following link: J.P. Morgan EMEA - Research Independence Policy.

U.S.: J.P. Morgan Securities LLC (“JPMS”) is a member of the NYSE, FINRA, SIPC, and the NFA. JPMorgan Chase Bank, N.A. is a member of the FDIC. Material published by non-U.S. affiliates is distributed in the U.S. by JPMS who accepts responsibility for its content.

General: Additional information is available upon request. The information in this material has been obtained from sources believed to be reliable. While all reasonable care has been taken to ensure that the facts stated in this material are accurate and that the forecasts, opinions and expectations contained herein are fair and reasonable, JPMorgan Chase & Co. or its affiliates and/or subsidiaries (collectively J.P. Morgan) make

Tarek Hamid AC (1-212) 834-5468
tarek.x.hamid@jpmorgan.com
J.P. Morgan Securities LLC
Harlan Sur AC (1-415) 315-6700
harlan.sur@jpmorgan.com

Nathaniel Rosenbaum, CFA AC (1-212) 834-2370
nathaniel.rosenbaum@jpmorgan.com
Nelson Jantzen, CFA AC (1-212) 270-1169
nelson.r.jantzen@jpmorgan.com

no representations or warranties whatsoever to the completeness or accuracy of the material provided, except with respect to any disclosures relative to J.P. Morgan and the Research Analyst's involvement with the issuer that is the subject of the material. Accordingly, no reliance should be placed on the accuracy, fairness or completeness of the information contained in this material. There may be certain discrepancies with data and/or limited content in this material as a result of calculations, adjustments, translations to different languages, and/or local regulatory restrictions, as applicable. These discrepancies should not impact the overall investment analysis, views and/or recommendations of the subject company(ies) that may be discussed in the material. Artificial intelligence tools may have been used in the preparation of this material, including assisting in data analysis, pattern recognition, and content drafting for research material. J.P. Morgan accepts no liability whatsoever for any loss arising from any use of this material or its contents, and neither J.P. Morgan nor any of its respective directors, officers or employees, shall be in any way responsible for the contents hereof, apart from the liabilities and responsibilities that may be imposed on them by the relevant regulatory authority in the jurisdiction in question, or the regulatory regime thereunder. Opinions, forecasts or projections contained in this material represent J.P. Morgan's current opinions or judgment as of the date of the material only and are therefore subject to change without notice. Periodic updates may be provided on companies/industries based on company-specific developments or announcements, market conditions or any other publicly available information. There can be no assurance that future results or events will be consistent with any such opinions, forecasts or projections, which represent only one possible outcome. Furthermore, such opinions, forecasts or projections are subject to certain risks, uncertainties and assumptions that have not been verified, and future actual results or events could differ materially. The value of, or income from, any investments referred to in this material may fluctuate and/or be affected by changes in exchange rates. All pricing is indicative as of the close of market for the securities discussed, unless otherwise stated. Past performance is not indicative of future results. Accordingly, investors may receive back less than originally invested. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client circumstances, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies to particular clients. This material may include views on structured securities, options, futures and other derivatives. These are complex instruments, may involve a high degree of risk and may be appropriate investments only for sophisticated investors who are capable of understanding and assuming the risks involved. The recipients of this material must make their own independent decisions regarding any securities or financial instruments mentioned herein and should seek advice from such independent financial, legal, tax or other adviser as they deem necessary. J.P. Morgan may trade as a principal on the basis of the Research Analysts' views and research, and it may also engage in transactions for its own account or for its clients' accounts in a manner inconsistent with the views taken in this material, and J.P. Morgan is under no obligation to ensure that such other communication is brought to the attention of any recipient of this material. Others within J.P. Morgan, including Strategists, Sales staff and other Research Analysts, may take views that are inconsistent with those taken in this material. Employees of J.P. Morgan not involved in the preparation of this material may have investments in the securities (or derivatives of such securities) mentioned in this material and may trade them in ways different from those discussed in this material. This material is not an advertisement for or marketing of any issuer, its products or services, or its securities in any jurisdiction.

Confidentiality and Security Notice: This transmission may contain information that is privileged, confidential, legally privileged, and/or exempt from disclosure under applicable law. If you are not the intended recipient, you are hereby notified that any disclosure, copying, distribution, or use of the information contained herein (including any reliance thereon) is STRICTLY PROHIBITED. Although this transmission and any attachments are believed to be free of any virus or other defect that might affect any computer system into which it is received and opened, it is the responsibility of the recipient to ensure that it is virus free and no responsibility is accepted by JPMorgan Chase & Co., its subsidiaries and affiliates, as applicable, for any loss or damage arising in any way from its use. If you received this transmission in error, please immediately contact the sender and destroy the material in its entirety, whether in electronic or hard copy format. This message is subject to electronic monitoring: https://www.jpmorgan.com/disclosures/email

MSCI: Certain information herein ("Information") is reproduced by permission of MSCI Inc., its affiliates and information providers ("MSCI") ©2026. No reproduction or dissemination of the Information is permitted without an appropriate license. MSCI MAKES NO EXPRESS OR IMPLIED WARRANTIES (INCLUDING MERCHANTABILITY OR FITNESS) AS TO THE INFORMATION AND DISCLAIMS ALL LIABILITY TO THE EXTENT PERMITTED BY LAW. No Information constitutes investment advice, except for any applicable Information from MSCI ESG Research. Subject also to msci.com/disclaimer

Sustainalytics: Certain information, data, analyses and opinions contained herein are reproduced by permission of Sustainalytics and: (1) includes the proprietary information of Sustainalytics; (2) may not be copied or redistributed except as specifically authorized; (3) do not constitute investment advice nor an endorsement of any product or project; (4) are provided solely for informational purposes; and (5) are not warranted to be complete, accurate or timely. Sustainalytics is not responsible for any trading decisions, damages or other losses related to it or its use. The use of the data is subject to conditions available at https://www.sustainalytics.com/legal-disclaimers. ©2026 Sustainalytics. All Rights Reserved.

"Other Disclosures" last revised May 16, 2026.

Completed 16 Jun 2026 06:39 AM EDT Disseminated 16 Jun 2026 06:39 AM EDT