

Global Investment Committee | August 25, 2026

AI Data Centers: The Heart of the Ecosystem

Few themes in modern history have influenced global markets as rapidly and profoundly as generative AI (artificial intelligence). Yet, we find that the narrative surrounding data centers and the broader AI ecosystem remains mired in complexity and technical terminology. In this Special Report, we aim to help sift through information and isolate the most important characteristics and risks.

First, we explore the “nuts and bolts” of data centers, starting with the traditional data-center market and then segueing to AI factories. We examine the leading role of hyperscalers and their part in building and financing the AI transformation, highlighting the promise of a historic revenue opportunity.

We conclude by evaluating the risks inherent in AI’s maturation—power, competition, financing and returns on invested capital—highlighting our views in each case. This work has encouraged us to remain focused on the quality and duration of their potential returns on investment.

Lisa Shalett

Chief Investment Officer
Head of the Global Investment Office
Morgan Stanley Wealth Management

Steve Edwards, CFA

Senior Investment Strategist

Susan Gim

Investment Strategist

Lisha Ge

Investment Strategist

Daryl Helsing, CFA

Investment Strategist

Lucy Chen

Associate



Morgan Stanley Wealth Management is the trade name of Morgan Stanley Smith Barney LLC, a registered broker-dealer in the United States. This material has been prepared for informational purposes only and is not an offer to buy or sell or a solicitation of any offer to buy or sell any security or other financial instrument or to participate in any trading strategy. Past performance is not necessarily a guide to future performance. Please refer to important information, disclosures and qualifications at the end of this material.

Studying the Heart of the AI Ecosystem

AI data centers turn stored code and mathematical theory into a continuously available knowledge service. They mark the convergence point for multiple dependencies: semiconductor manufacturers, energy suppliers, software developers and cloud providers. The below exhibit highlights data centers' physical makeup.

When AI data-center capacity becomes constrained, the entire ecosystem experiences the effects, from large language model (LLM) developers to the end user. While LLMs form the intelligence layer of the AI ecosystem, we consider data centers to be the heart, acting as the central host for AI's essential resources—compute, data and energy—and as a vital hub for connectivity. We liken networks to the circulatory system, wherein data embodies the lifeblood of the AI ecosystem, retrieving and distributing information across the body. Computing power (compute) serves as the pulse—a continuous cadence of calculations that process inference requests, run computations and communicate model responses.


As the AI-capex cycle matures, we have identified several areas where the ecosystem may have begun undergoing a meaningful evolution. Broadly, emphasis has shifted from “token-maxxing” to “token-managing,” with pivots from closed-weight to a hybrid of closed- and open-weight models.

These developments point to potential competitive dynamics inherent in US- and Chinese-led innovation paths.

Below, we consider the implications of these potential inflection points. Heretofore, the AI bull market has largely benefited from a backdrop of favorable financial conditions. Consensus expectations have adapted to incorporate a trajectory of substantial capital expenditure growth, supporting anticipated operating margin expansion for both AI providers and AI users. The apparent gap in perceived supply and demand for compute has sparked a “land rush” for new AI data centers, propelling a flood of capital expenditure, along with public and private market financing and a slew of potentially circular quasi-contracts between leading AI players.

Looking ahead, investors will likely benefit from increased transparency into AI data centers' financial picture. The potential IPOs of two major LLM providers may offer a deeper look into the “income statement” for frontier models—and the potential path ahead. The proliferation of public market financing has provided additional signals regarding investors' evaluation of the risk-reward picture, such as in AI-related credit spreads. Publicly available indexes of chip rental rates, token prices and usage statistics offer a picture of real-time supply-demand conditions. Combining token prices and usage statistics affords calculations of total user spend, which has apparently flatlined since June.

Data Centers Provide Connectivity Between Technology, Climate Control and Power Resources to Offer Always-Available Service



Source: Morgan Stanley & Co. Research as of March 24, 2026

FUNDAMENTALS

Hyperscalers Accelerated Traditional Data-Center Growth

Over the past decade, data centers have evolved from colocation to hyperscaler cloud models.

The initial growth of data centers coincided with companies' decision to shift IT hardware to more centralized, off-site locations. Traditional data centers fall into one of two key models: colocation and cloud. Colocation operators—either retail or wholesale, depending on the enterprise size—bear responsibility for up-front building costs and site-related real estate. To some extent, the underlying real estate provides contingency value in the worst-case scenario that customer demand or utilization is permanently impaired. Colocation data-center operators generate revenue through “rents” charged to customers, with calculations based on units of power used. Their contracts serve as long-term leases, with tenants bringing their own purchased IT hardware. Colocation data centers have benefited from multiple demand tailwinds over the past decade, causing vacancy rates to reach historical lows.

Retail Colocation: <i>Single Tenant or Owner-Operated</i>	Operator sells smaller blocks plus connectivity, leading to higher interconnection revenue and more operational complexity	Power + Cooling Security Cross-Connects
Wholesale Colocation: <i>Large Tenants and Hyperscalers</i>	Operator sells MW-scale suites or building, involving longer leases, lower customer count, and associated pre-leasing matters	Large MW Suites Long Leases Ramp Schedules
Build-to-Suit or Hyperscaler-Owned:	Hyperscaler controls design and capex, with third parties potentially providing land, power shells or JVs	Tenant Design Custom Power Dedicated Campus


Source: Morgan Stanley Wealth Management Global Investment Office (GIO) as of June 30, 2026

The hyperscaler-owned cloud business model emerged in the mid-2010s and has accelerated rapidly since 2016. Microsoft, Amazon and Google control roughly 60% of the global cloud infrastructure services market.

These three hyperscalers experienced 10-times revenue growth from their cloud business from 2016 to 2025.

Gartner estimates that worldwide public cloud spending hit \$723 billion in 2025, with the three major hyperscalers accounting for cloud revenues in excess of \$400 billion. Morgan Stanley & Co. Research's recent CIO surveys have suggested the following market shares by vendor: Microsoft Azure, 28%; Amazon Web Services, 18%; and Google Cloud, 6%.

Curiously, 27% of enterprise CIOs surveyed do not use public cloud-based services, but this number has halved from 2018, when it was nearly 60%. All three companies have built profitable, multiline businesses with margins of 30% to 40%—or even greater. Each provider has established “pay-as-you-go” pricing, with usage measured across categories like compute, storage, data transfer and networking, and services like databases, analytics and AI. The hyperscalers have positioned their services to increase the stickiness of the enterprise relationship, allowing customers to increase their spend over time by delivering a full offering of infrastructure, products and software “as a service” (IaaS, PaaS, SaaS).



Source: Morgan Stanley Wealth Management GIO as of May 27, 2026

FUNDAMENTALS

Understanding AI Data-Center Economics

Generative AI has ushered in the next generation of data-center investment, with the promise of much greater annual revenues and higher internal rates of return (IRRs). Compared with the traditional data centers of the 2010s, AI factories offer the potential for higher returns but require much higher capital investment. The underlying real estate asset plays a much smaller role in AI factories' business model. Instead, the centrality of hybrid infrastructure and semiconductor assets makes the resulting returns more sensitive to factors like GPU utilization, pricing and hardware-refresh cycles. The revenue model differs by function, whether for training or inference, as shown to the right.

AI data centers offer customers access to high-performance GPU compute. While their physical infrastructure mimics traditional data centers, the technology components show material differences.

Metric	Traditional Data Center	AI Factory
Revenue (USD/MW)	\$2–3	\$14–18
CapEx (USD/MW)	\$10–14	\$50–100
EBITDA Margin	50–60%	40–50%
IRRs	10–13%	12–30%

AI Training Workload	AI Inference Workloads
Revenue = Total GPU Hours x Price per Hour x Utilization (Limited Batching)	Revenue = Total GPU Hours x Tokens per Hour x Price per Token x Utilization (Batching/Scheduling)


Source: Morgan Stanley & Co. Research as of March 24, 2026

First, their revenue streams can vary widely, depending on the level of GPU performance. Second, their development costs differ sharply from traditional data centers, whose costs involve land and site development, building construction (called the "shell"), grid connection and site commissioning. For AI data centers, the provider must purchase the underlying IT equipment, along with the site build-out. These hardware costs are significant, encompassing nearly 75% of the total project costs. As a result, AI factories require five times the up-front investment of a traditional data center. We consider AI factories as effectively hybrid assets: part infrastructure and part semiconductor. Furthermore, AI data centers bear material power costs, as they can account for 15% to 20% of operating costs. AI factories' margins appear lower, given a much higher revenue base. The expected rate of return shows great sensitivity to return assumptions on the underlying hardware equipment. If server investments do not hold their residual value and depreciate more quickly than expected, the AI data center will face a materially weaker expected rate of return.

Given the revenue opportunity and robust demand, AI data centers' growth has become turbocharged.

Morgan Stanley & Co. Research anticipates 81% annualized growth for AI data centers, compared with 10% for traditional data centers. While traditional data centers exhibit diffusion among many players, the AI data center market is highly concentrated among four main hyperscalers plus a few neoclouds.

Analysts expect US data centers' incremental power demand to come almost exclusively from AI workloads. As such, the sheer volume of required capital stands out as the main differentiator for this data-center investment cycle. A one-gigawatt AI data center carries an estimated cost of \$30 to \$50 billion—75% of which comes from IT hardware costs. The incremental 33.2-gigawatt (GW) expansion in 2026 and 2027 carries an approximate price tag of \$1.7 trillion. In just three years, AI data centers' power demand has grown from near zero to a level equivalent to that for traditional data centers.



Source: Morgan Stanley & Co. Research as of June 10, 2026


FUNDAMENTALS

Understanding AI Data-Center Capex: From Construction to Compute

AI data-center construction represents a costly, complicated and colossal undertaking.

Recent AI megaprojects like Stargate, Hyperion and Colossus are massive in scale, with targeted power of one to two gigawatts. As one example, a project in Richland Parish, Louisiana, features a design footprint of 4 million square feet, covering nearly 70 football fields, and an estimated project cost of \$27 billion.

The chart to the right highlights the contributors to AI data centers' massive expenses, broken down roughly into infrastructure and IT costs. The infrastructure costs behind constructing the facility itself vary much less than those for the IT equipment. AI data centers' infrastructure requires land acquisition, base construction, grid connection and substation work, and commissioning. Costs related to these activities amount to \$10–\$15 billion per gigawatt equivalent. AI processors represent up to 60%–70% of the data center's cost. After including networking and storage equipment, IT equipment broadly speaking becomes an even larger slice of the cost pie.




Source: Morgan Stanley & Co. Research as of April 29, 2025

Hyperscalers can no longer self-finance AI data centers' required capital expenditures.

Given the cohort's cumulative cash flow of approximately \$800 billion, 2026 will likely mark a meaningful inflection point, with AI data-center financing becoming predominantly external. Hyperscaler forward-free-cash-flow estimates have flipped negative, and capital expenditures have continued to trend higher. Having grown substantially, external financing includes both secondary equity issuance and expanded debt financing through corporate bond offerings and partnerships with financial institutions and asset managers.

Moreover, AI data centers' all-in costs involve a significant amount of "off-balance-sheet" investments, which may transition back to balance sheets over the coming years. These purchase commitments, which are obligations to purchase goods once they are delivered, have reached around \$500 billion across the four hyperscalers, mainly for hardware. Similarly, already executed but not yet operational lease commitments also remain off-balance-sheet until the data centers become operational. Across the cohort, these commitments encompass roughly \$400 billion. Critically, as data centers become operational, these off-balance-sheet expenses become recognized and will consequently impact corporate earnings. While credit-agency leverage multiples likely account for these expenses, we believe it unlikely that their consequences have factored into consensus earnings expectations. We will continue to monitor this dissonance as the capex narrative matures.



Source: Morgan Stanley & Co. Research as of May 17, 2026


FUNDAMENTALS

Rising Costs of Compute, Driven by Memory Inflation


Memory shortages have disrupted hyperscalers' capex plans—and even touched consumer electronics.

Originating with shortages in high-bandwidth memory (HBM), memory deficits have extended to DRAM and NAND, as AI workloads have rapidly evolved from training to inference models. On the right, we highlight the three key memory products in semiconductors and their usages. Industry players widely underestimated this dramatic increase in memory needs. As a result, supply constraints have proliferated. DRAM production lines are highly fungible, allowing DRAM producers to shift production from PC DRAM or smartphone DRAM to enterprise DRAM. That flexibility has resulted in structural shortages for lower-priced PC and smartphone DRAM. Consumers have faced a dramatic impact in 2026, as shortages have forced PC and smartphone manufacturers to hike consumer prices.

AI data centers face similarly steep implications. Nvidia's next-generation GPU—VR200 (Vera Rubin), launching in the third or fourth quarter—will cost an estimated \$8 million per server, due to the compounded effects of its memory density and elevated memory prices. While just one year has passed since the commercial launches of GB300 (Blackwell), the costs have nearly doubled. This cost inflation has affected not just memory but almost every single component that drives AI data centers.



Nvidia Server Rack Price (USD, Million) by Product Line




Source: Morgan Stanley & Co. Research as of June 2, 2026 (top chart), May 20, 2026 (bottom chart)

Given that custom chips can offer meaningful cost savings, there may be limits to chipmaker monopolies.


Nvidia powers the current dominant server, GB300 (Blackwell), for AI data centers. To tally the hardware costs, we observe that a one-gigawatt AI factory requires roughly 6,000 to 7,000 server racks. At approximately \$4 million per rack, the AI factories' chip cost amounts to around \$19 billion. In pursuit of greater cost efficiency, several hyperscalers have invested significant resources to commercialize their own custom AI processors (ASICs), effectively creating a middle-market technology offering for the AI ecosystem.

The scalable commercialization of ASICs may diminish any one chipmaker's technological monopoly. These custom server chip solutions (highlighted in light blue in the second chart) can deliver performance at a significant pricing discount to the incumbent GPU model. These developments could have meaningful implications if the sector goes through a period of rationalization where up-front hardware investment needs to be more prudent. Diminished spend of processors could dampen all-in capex for the hyperscaler cohort by a cumulative \$350 billion for 33 gigawatts of planned capacity, equating to savings of approximately \$10 billion per gigawatt.

ASIC Volume Forecast by Major Companies



GPUs vs. ASICs Capex Spend for 1GW Data Center





Source: Morgan Stanley & Co. Research as of May 17, 2026

FUNDAMENTALS

Rent or Own? Introducing Neoclouds, a Capital-Light Option

Neoclouds bring customers an alternative route to scale their AI projects, a capital-light approach to sourcing compute.


Neoclouds are specialized, nonhyperscaler operators, providing AI computing capacity. To achieve this, they typically install GPUs or TPUs in leased colocation facilities rather than constructing new large data centers. In the structure shown to the right, a special-purpose vehicle (SPV) raises capital from GPU investors and uses the proceeds to purchase servers and AI chips. Five-year monthly lease payments from hyperscalers provide the cash flow to cover expenses and return capital and income to investors. Nevertheless, equity investors still bear two key risks. First, servers may depreciate or become obsolete faster than expected, reducing their residual value. Second, hyperscalers may not renew their lease, forcing the neocloud to find other customers or sell capacity into the spot market at lower pricing. For the hyperscaler, the structure provides access to AI compute without requiring up-front investment in servers or facilities. In simple terms, the hyperscaler is leasing “cars” (the GPUs) rather than buying them, and the equipment returns to the neocloud or SPV on lease expiration.



Source: Morgan Stanley & Co. Research as of June 17, 2026

Neoclouds depend on debt financing, given that cash flows are limited and rely on hyperscaler contracts.

As a result, they depend heavily on debt financing to expand their capacity. This chart illustrates how rapidly this debt may accumulate through 2028. This setup therefore makes long-term hyperscaler contracts central to the financing model. Agreements with Microsoft, OpenAI, Meta, Nvidia and other large customers provide the revenue visibility required to secure debt and meet interest and principal payments. Still, the neoclouds face concentration risk. For example, Microsoft, OpenAI and Meta collectively represented more than 60% of CoreWeave’s annual recurring revenue in 2025. Although their combined share is projected to decline to roughly 36% by 2028 with an expanding customer base, CoreWeave would still depend on a relatively small group of major technology companies. The picture highlights the ongoing connection between customer concentration and credit risk. If a hyperscaler were to delay deployment, reduce its capacity requirements, dispute a contract or decline to renew an agreement, the neocloud may lose a significant source of cash flow, even while maintaining responsibility for its debt payments. The rapid depreciation in GPUs—and their potential to lose value as newer generations become available—could weaken the collateral available to lenders, amplifying their risk.



Source: Morgan Stanley & Co. Research as of April 15, 2026


In simple terms, neoclouds are borrowing to build an AI compute fleet before they have generated the cash to pay for it. The model can support rapid growth but only if hyperscaler contracts remain enforceable, customer demand materializes as expected and contracted revenue grows fast enough to service the expanding debt burden.

FUNDAMENTALS

Risk No. 1: The Power Bottleneck

AI data centers are megaprojects with tremendous power needs.

AI data centers immediately become massive consumers of power, owing to the power draw of AI processors. As such, power availability represents a key bottleneck for the AI build-out. The exhibit on the right highlights the major AI data centers and their power demands. Long-term estimates of power demand are much higher, however, given that many of these projects have entered only their initial phases of development and will ultimately hit one-gigawatt levels. The Electric Power Research Institute (EPRI) has published its own forecast of data-center power demand using project-based filings from state-regulated utilities. By 2030, EPRI estimates that incremental power demand may hit approximately 37 gigawatts for high-confidence projects in the pipeline. As a comparison, the UK's entire electricity generation capacity sits at 33 gigawatts.




Source: Epoch.AI as of July 8, 2026

This power bottleneck also features a geopolitical wrinkle. Countries with higher power costs, such as those in Europe and North Asia, will face challenges in participating in the AI build-out, owing to limitations from regulations, natural resources and infrastructure. In contrast, power availability will likely support continued leadership from the US and China in the AI-capex boom.

Data centers have recently faced increased regulatory scrutiny from state governments, in response to concerns around electricity demand, transmission costs, and water usage.

The most consequential measures generally fall into three categories: restrictions on new development, requirements that data centers bear more of their infrastructure costs and tighter approval and disclosure requirements. Virginia regulators recently required data centers to pay for transmission infrastructure constructed specifically to serve their facilities. The state has also adopted an energy-consumption tax and additional environmental requirements. Oregon has pursued a similar "pay-your-own-way" philosophy through its POWER Act. Beginning in July 2026, large data centers served by Portland General Electric faced an average 29% rate increase. More recently, Texas Governor Greg Abbott ordered the Electric Reliability Council of Texas (ERCOT) and the Public Utility Commission of Texas to audit all data-center projects. The review includes power and water requirements, public subsidies, on-site generation, community impacts and ownership. Historically, observers have considered Texas's regulatory regime as more favorable to business and energy development. This announcement confirms a broader shift of tighter power regulations. For more details, please review our June 17 report, "The AI Build-Out: When the Cloud Comes to Town," which details [political pushback against data centers](#).




Source: Morgan Stanley & Co. Research as of Dec 5, 2026

FUNDAMENTALS

Risk No. 2: Open-Weight Models and Hybrid Approaches

AI models fall into two categories—closed-weight and open-weight—and the relative use share will have critical downstream impacts on hyperscaler and model-creator economics.

Put simply, open-weight models give users greater control over their AI model deployment. A company can download the model, adjust it for specific tasks and run it on corporate systems. In contrast, closed-weight models stay under the provider's control, usually accessed through a subscription or API. Model creators typically keep their most advanced AI models closed, allowing the developer to recover some of the enormous cost of building and training. Many of the large-scale development projects in AI data centers are supported by committed-compute contracts where the AI lab (frontier model provider) commits to the new capacity being built by the hyperscaler or neocloud.




Source: Morgan Stanley & Co. Research as of Aug. 3, 2026

In a recent overview on open-weight models, Morgan Stanley & Co. Research pointed to three key observations. First, open-weight models suggest growing penetration, with more than 50% of enterprise usage routed to open-weight models. Second, open-weight models introduce challenges for token economics, leading to potential downward pressure on frontier-model prices. Third, open-weight models introduce implementation risk for frontier models. Chinese model creators have deployed open-weight models alarmingly fast, given significantly fewer power and infrastructure bottlenecks. In the US, frontier models face greater exposure to extended project delays and regulatory and political backlash, potentially dampening realistic supply levels.

Chinese competition may compress pricing for both AI models and top-flight chips.

Chinese open-weight models benefit from a tremendous cost advantage. In the chart to the right, we share the cost statistics for US frontier models, compared with Chinese open-weight models, measured per unit of one million tokens. In real-world terms, one million tokens would power the following applications: reviewing 50 to 100 legal documents like contracts, credit agreements, side letters and diligence memos; preparing 10,000 customer support emails, chat transcripts, CRM notes and resolution notes; or contextualizing approximately 750,000 words (around 1,500 pages) of dense business text. Open-weight models introduce meaningful deflation risks in AI data centers' buildout, given that frontier model creators lease a substantial amount of compute from hyperscalers. Lowered demand for frontier models could lead investors to question the outlook for addressable demand and the rationale for such significant capital expenditures. Moreover, Chinese companies have demonstrated their capacity to innovate in the face of hardware constraints. Intriguingly, Chinese AI models have continued to improve, despite their training and deployment on domestically produced processors one to two generations behind Nvidia's most advanced offerings. As US export controls have constrained Chinese companies' access to Nvidia's highest-performance GPUs, locally designed accelerators from Huawei and Cambricon have become increasingly important. Over time, this setup may produce a more independent Chinese AI-infrastructure stack, even if the country's individual chips fall short of the global performance frontier.



Source: Alpine Macro as of July 8, 2026


FUNDAMENTALS

Risk No. 3: The Increasing Prevalence of Debt Financing

The AI data-center build-out will require substantial financing, with credit markets expected to play a central role despite still-manageable hyperscaler leverage.

AI data-center capex could approach \$3.2 trillion over the next three years, reflecting the scale and complexity of next-generation AI infrastructure. Analysts expect the majority of this investment, clocking in at \$1.75 trillion, to come from credit financing, while hyperscaler cash flows and private investor participation may provide the additional financing. Already in 2026, fundraising has accelerated meaningfully, including large equity raises from SpaceX, Google and SK Hynix, alongside elevated corporate bond issuance. To date, investment grade corporate bond issuance has accounted for most AI infrastructure debt, yet financing channels are broadening. Hyperscalers have innovated by tapping international debt markets, diversifying beyond US dollar issuance and adopting structures beyond traditional unsecured bonds. The five largest US hyperscalers' outstanding debt represents around 4% of the Bloomberg US Corporate Index, compared with roughly 20% of the S&P 500 Index's market capitalization. Year-to-date issuance has exceeded \$180 billion, including over \$50 billion outside the US dollar market. While hyperscalers' leverage has risen from 0.9 times to around 1.8 times, it remains below the broader nonfinancial corporate average of 2.3 times.

Credit Financing For AI Capex (Total: \$1.75tn)



Source: Morgan Stanley Wealth Management GIO as of May 29, 2026

AI-related debt has traded with greater implied risk, pointing to the costs and uncertainty of the AI-capex boom.

Hyperscalers' investment grade issuance has more than doubled over the past year and appears on track to reach \$250 to \$300 billion in 2026. Some estimates for next year's pipeline put investment grade supply as high as \$400 billion. The price response has reflected this heavy supply, with high-quality hyperscalers' credit spreads looking wide relative to their credit ratings. Instead of trading like their AA-rated peers, many have begun trading more like A-rated debt, particularly at longer maturities.

AI Bonds Option-Adjusted Spread (OAS)
Less IG Index OAS, Five-Day Average



Source: Strategas as of Aug. 10, 2026

The collective financing "health" underlying the AI-capex boom relies, to some degree, on the health of any specific issuer, given the interconnectedness of industry players. Demand for new issuance has softened somewhat, as reflected in lower subscription ratios. Hyperscalers are seeking financing at long-term maturities, beyond 10 years, amid an environment of increasing term premiums and two-decade-high real yields. Hyperscalers' long-term issues have borne the brunt of spread widening. Practically speaking, insurance companies often make up 30-year corporate bonds' incremental buyer, and their demand has softened, resulting in isolated indigestion and wider spreads.


Moreover, AI financing has begun playing a significant role in other credit markets, including non-US dollar investment grade corporates, high yield and leveraged loans, securitized credit and private credit.

FUNDAMENTALS

Risk No. 4: The Sensitivity of Returns on Invested Capital

AI factories financial attractiveness will likely display sensitivity to cost inputs and token prices.


In this exhibit, we stress test the returns on invested capital for a hypothetical one-gigawatt AI data center, leveraging [Morgan Stanley & Co. Research analysis](#). The analysis assumes that inference prompts account for 65% of token usage and assumes 100% customer utilization of the AI data center. The return on invested capital is most sensitive to changes in token pricing—with returns on invested capital falling by 6% for every 10% decrease in token prices. Token pricing serves as the single most impactful driver of returns. At \$1.80 per one million tokens, the ROIC of a hypothetical one-gigawatt data center reaches 46% under these scenarios—ranging from 20% to 72% for \$1.00 to \$2.50 in token price. The Silicon Data LLM Token Expenditure Index tracks the cost of a one-million pack of tokens. Having peaked north of \$2.00 in May, current pricing has retreated to around \$1.10. We observe that token pricing has recently softened through combined changes to the proportion of open-weight models and the blended price of frontier models. Second, we stress tested the total construction costs for AI data centers. Current market pricing for a one-gigawatt data center suggests \$23 billion in GPUs and \$16 billion in non-IT capital expenditures. Changes in these values have less impact on AI data centers' returns on invested capital, given their presence in the denominator in these calculations.



Source: Morgan Stanley Wealth Management GIO as of July 27, 2026

There is an interesting divergence between falling token and expenditure and rising rental rates for GPUs.

In addition to the rising utilization of open-weight models, some AI users have recently shifted their approach from "token-maxxing" to "token-managing." The "maxxing" approach flowed as the natural outcome from per-seat, "all-you-can-use" pricing. As the bill for "maxxing" came due, however, many corporations have implemented a more cost-sensitive "managing" approach. The top set of charts highlights one result from this inflection. Intriguingly, earlier-generation GPUs have shown relative strength compared with later-generation models. For example, the rental market for Nvidia's H100 (Hopper) GPU, launched in 2022, has become one of the most robust, liquid and mature. As shown to the right, the H100's rental has steadily increased, even as token pricing has fallen materially since mid-May. One theory involves the constraints outside the GPU itself. Long lead times for Blackwell systems and broader data-center infrastructure constraints have slowed the pace at which new compute can enter the rental market. As such, existing H100 chips may retain considerable economic value, especially in the case of deployment lags with AI data centers capable of leveraging B200 chips. Second, US export restrictions have constrained China's in-country deployment of advanced Nvidia GPUs but have not fully prevented Chinese companies from accessing the same computing power in markets like Southeast Asia. Chinese demand has become important to Southeast Asian data centers and GPU capacity and provide an incremental source of global H100 utilization.




Source: Bloomberg, OpenRouter, Silicon Data as of Aug. 3, 2026

FUNDAMENTALS

Greater Transparency on AI's Underlying Economics

Investors will benefit from greater transparency on AI's underlying economics, including token pricing and chip rental pricing.

As noted previously, time series have emerged for token pricing and rental rates for GPUs, helping investors to assess supply-demand conditions. Given that a significant pool of invested capital will lie at the GPU level, greater price discovery will prove critical. Data providers have begun to share spot price indexes, with futures contracts anticipated to launch shortly.



Source: Bloomberg as of Aug. 10, 2026

The chart to the right provides three price benchmarks for the same variable: hourly rental rates for the H100 GPU.

The three series—from Silicon Data, ORNH and Compute Desk—show meaningful dispersion in their levels and respective volatility. At present, companies simply rent GPUs at the prevailing market prices. We believe that compute futures will become a valuable signal for investors in the AI trade and a critical hedging instrument for providers and users of AI compute. On Oct. 5, the CME plans to launch cash-settled futures for H100s and B200s based on the Silicon Data indexes. Silicon Data has distinguished itself as a leading voice in the development of AI compute price indexes and LLM token price indexes—both of which are actively tracked on Bloomberg and cited in the press. These developments suggest that AI compute may become a more standardized and tradable market, making it easier to price, finance and manage.

Our Views on the Key Risks

Slow Slog for Power. We believe that new AI data centers will struggle with approvals for new connections to the traditional power grid. Project owners have become more creative, looking to construct their own behind-the-meter solutions. These can range from single off-grid plants supplied with gas pipelines to modular mini gas turbines or energy storage solutions. All in all, these power solutions can increase power costs significantly, moving them from \$0.10–\$0.12 per kilowatt-hour to \$0.20–\$0.22 per kilowatt-hour. Given that power represents 15% to 20% of operating costs, doubling power costs pressures ongoing margins and can impact project economics; in contrast, the capital expenditures for servers, infrastructure and power are largely up-front costs. State governments have recently demonstrated coordinated, bipartisan pushback on data-center power connections. Voters' anxiety on both the physical presence of data centers in their backyards and the impact on power costs will likely garner popular support and may play heavily in this year's midterm election cycle.

Constant Eye on Competition. Hyperscalers' significant business model differentiation and deep experience in the cloud market will likely provide sufficient opportunity to allow each member to grow without material competitive pressures. On the LLM front, however, Chinese open-weight models have shown rapid technology progress, allowing users to achieve cost-effective outcomes. Moreover, we highlight the tremendous power advantage China's economy holds, with a growing pool of low-cost power. Finally, deployment strategy (the balance of open- and closed-weight models) will become increasingly important as we enter the apex of the capex cycle.

Financing Risk: Centered Primarily on AI Server Values. Rapid inflation across key components—high-bandwidth memory, DRAM and NAND—have pressured costs for GPUs and other AI processors. These processors face the risk of rapid depreciation from several sources: a normalization of cost inputs, weaker-than-expected demand markets or greater memory supply from China and South Korea. The potential volatility in these variables may complicate long-term agreements, which are now increasingly common among purchasers of memory, where current prices serve as the floor for deliveries one to two years out.

The Bottom Line

Approaching the four-year anniversary of this AI-driven bull market, we anticipate that investors' attitudes may shift, given maturing trends. Early on, investors cheered nearly every capex announcement; given a backdrop of healthy free cash flow and tremendous demand, the drumbeat rewarded "more is more."

To provide context around upside spending surprises, in mid-2025, sell-side estimates reflected hyperscaler capex expectations of \$350 billion in 2026. A year later, these estimates have ballooned to \$800 billion—and spending is forecast to hit \$1.1 trillion in 2027. This upside surprise, however, requires financing for activation—through equity issuance, investment grade credit and private financing.

In addition, users' shift to a more hybrid approach (blending closed- and open-weight models) could lead to seismic changes to AI data-center business models. The price gap between frontier LLMs and open-weight models is significant and has unleashed intensifying competition that has market participants in both the US and China dynamically rethinking their strategies. We note Meta's recent decision to re-embrace open-weight models.

Plus, companies have innovated on computing architectures, including the use of application-specific chips (ASICs), which may offer not only more reasonable up-front costs but also opportunities for seamless integration into enterprise networking and software deployment.

Importantly, this cycle is very different from the proprietary walled gardens that created the cash-flow machines of the Magnificent Seven in the prior two cycles (internet and cloud). We anticipate that AI capex could be more competitive, more cyclical, more capital-intensive and potentially more commoditized.

We continue to focus on the deceleration of rate of change. While absolute levels of data center-related spending will stay high for several years, our focus remains on the rate of change, as equity markets tend to discount future earnings. Consensus earnings forecasts for companies levered to AI infrastructure—represented by the Morgan Stanley & Co. AI Tech Beneficiaries basket—have benefited from year-to-date earnings revisions averaging +80%. Looking forward, we are keen to investors' renewed focus on sustainability into 2028 and beyond. We suspect that companies have participated in a natural rush to order all the compute, materials and industrial components in short supply, but actual deployment may look more uneven, given material pushback from local communities and long timelines for power grid connections.

Equity and credit volatility and increasing term premiums may reinforce the market's role in serving as a disciplining force and capital allocator. For those companies in the AI ecosystem with weaker balance sheets, these reactions have flowed through credit spreads. Sizable credit issuance, typically at long maturities, has compounded with global forces, raising term premiums. A potentially higher-for-longer cost of capital acts as a self-curing agent, preventing irresponsible capital flows. Furthermore, the market has cast a discerning eye on capital allocation strategy rather than merely the absolute dollar amount of investment into the AI ecosystem.

As the AI-capex boom progresses, its beneficiaries will undoubtedly shift. From early 2025 through mid-2026, semiconductor makers and memory names experienced a historic rally, leading to overextended conditions. Looking forward, the opportunity set may shift back to hyperscalers, which control important parts of the AI ecosystem, including the cloud and software layers that help organize and deliver AI services. Nevertheless, we recommend selectivity with pure-play AI names. Not every large technology platform is equally well positioned to capitalize. We prefer those companies moving quickly toward hybrid AI models that can connect AI inference directly to their cloud businesses. They may be better able to defend margins as competition increases and parts of the AI market become more commoditized.

We believe that AI data centers' "picks and shovels" may prove more resilient, attracting more investor focus as positioning broadens. Given the transformative nature of the AI investment cycle, we believe that investors will likely benefit from a holistic approach that balances growth, income and sustainability. That approach suggests considering associated exposures like upstream energy, midstream distribution and utilities. For investors looking for diversified private investments, generalist infrastructure and specialized digital infrastructure funds provide more concentrated exposure to the data-center theme. Infrastructure funds offer a combination of income generation, diversification and return enhancement. They also have a significant long-term track record of investing in data-center projects. We provide more detail on digital infrastructure assets in our Aug. 25 "[Alternative Investment Insights](#)." In the liquid credit markets, we remain on the sidelines, as we find most of the investment-grade paper suited to very long-duration investors, such as insurance allocators.

AI DATA CENTERS: THE HEART OF THE ECOSYSTEM

Disclosure Section

The **Global Investment Committee (GIC)** is a group of seasoned investment professionals from Morgan Stanley & Co. LLC, Morgan Stanley Investment Management, and Morgan Stanley Wealth Management who meet regularly to discuss the global economy and markets. The committee determines the investment outlook that guides our advice to clients. They continually monitor developing economic and market conditions, review tactical outlooks and recommend asset allocation model weightings, as well as produce a suite of strategy, analysis, commentary, portfolio positioning suggestions and other reports and broadcasts.

Susan Gim, Lisha Ge, Daryl Helsing, and Lucy Chen are not members of the Global Investment Committee and any implementation strategies suggested have not been reviewed or approved by the Global Investment Committee.

Benchmark Definitions

Silicon Data LLM Token Expenditure Index family is a set of daily benchmarks for the price of large-language-model tokens. Each series computes the usage-weighted price of one million tokens across a defined LLM universe. The indices blend active provider pricing with real consumption volume observed across multi-provider routing gateways and inference networks, publishing a true effective price per million tokens daily.

For index, indicator and survey definitions referenced in this report please visit the following: <https://www.morganstanley.com/wealth-investmentsolutions/wmir-definitions>.

Glossary

Artificial Intelligence (AI) A field of study that seeks to train computers to process large amounts of unstructured information in a manner similar to human intelligence, capable of performing tasks such as learning and problem solving.

Dispersion is a measure for the statistical distribution of portfolio returns. It is the asset-weighted standard deviation of individual portfolio returns within a comparable composite from the composite return.

Option-adjusted spread (OAS) is a measurement of the spread of a fixed income security rate and the risk-free rate of return, which is adjusted to take into account an embedded option.

Term premium is the excess yield that investors require to commit to holding a long-term bond instead of a series of shorter-term bonds.

Volatility: A statistical measure of the dispersion of returns for a given security or market index. Volatility can either be measured by using the standard deviation or variance between returns from that same security or market index. Commonly, the higher the volatility, the riskier the security.

Hypothetical Performance

General: Hypothetical performance should not be considered a guarantee of future performance or a guarantee of achieving overall financial objectives. Asset allocation and diversification do not assure a profit or protect against loss in declining financial markets.

Hypothetical performance results have inherent limitations. The performance shown here is simulated performance not investment results from an actual portfolio or actual trading. There can be large differences between hypothetical and actual performance results.

Despite the limitations of hypothetical performance, these hypothetical performance results may allow clients and Financial Advisors to obtain a sense of the risk / return trade-off of different asset allocation constructs.

Investing in the market entails the risk of market volatility. The value of all types of securities may increase or decrease over varying time periods.

This analysis does not purport to recommend or implement an investment strategy. Financial forecasts, rates of return, risk, inflation, and other assumptions may be used as the basis for illustrations in this analysis. They should not be considered a guarantee of future performance or a guarantee of achieving overall financial objectives. No analysis has the ability to accurately predict the future, eliminate risk or guarantee investment results. As investment returns, inflation, taxes, and other economic conditions vary from the assumptions used in this analysis, your actual results will vary (perhaps significantly) from those presented in this analysis.

The assumed return rates in this analysis are not reflective of any specific investment and do not include any fees or expenses that may be incurred by investing in specific products. The actual returns of a specific investment may be more or less than the returns used in this analysis. The return assumptions are based on hypothetical rates of return of securities indices, which serve as proxies for the asset classes. Moreover, different forecasts may choose different indices as a proxy for the same asset class, thus influencing the return of the asset class.

Asset Class and Other Risk Considerations

Equity securities may fluctuate in response to news on companies, industries, market conditions and general economic environment.

Bonds are subject to interest rate risk. When interest rates rise, bond prices fall; generally the longer a bond's maturity, the more sensitive it is to this risk. Bonds may also be subject to call risk, which is the risk that the issuer will redeem the debt at its option, fully or partially, before the scheduled maturity date. The market value of debt instruments may fluctuate, and proceeds from sales prior to maturity may be more or less than the amount originally invested or the maturity value due to changes in market conditions or changes in the credit quality of the issuer. Bonds are subject to the credit risk of the issuer. This is the risk that the issuer might be unable to make interest and/or principal payments on a timely basis. Bonds are also subject to reinvestment risk, which is the risk that principal and/or interest payments from a given investment

AI DATA CENTERS: THE HEART OF THE ECOSYSTEM

may be reinvested at a lower interest rate.

Bonds rated below investment grade may have speculative characteristics and present significant risks beyond those of other securities, including greater credit risk and price volatility in the secondary market. Investors should be careful to consider these risks alongside their individual circumstances, objectives and risk tolerance before investing in high-yield bonds. High yield bonds should comprise only a limited portion of a balanced portfolio.

Principal is returned on a monthly basis over the life of a **mortgage-backed security**. Principal prepayment can significantly affect the monthly income stream and the maturity of any type of MBS, including standard MBS, CMOs and Lottery Bonds. Yields and average lives are estimated based on prepayment assumptions and are subject to change based on actual prepayment of the mortgages in the underlying pools. The level of predictability of an MBS/CMO's average life, and its market price, depends on the type of MBS/CMO class purchased and interest rate movements. In general, as interest rates fall, prepayment speeds are likely to increase, thus shortening the MBS/CMO's average life and likely causing its market price to rise. Conversely, as interest rates rise, prepayment speeds are likely to decrease, thus lengthening average life and likely causing the MBS/CMO's market price to fall. Some MBS/CMOs may have "original issue discount" (OID). OID occurs if the MBS/CMO's original issue price is below its stated redemption price at maturity, and results in "imputed interest" that must be reported annually for tax purposes, resulting in a tax liability even though interest was not received. Investors are urged to consult their tax advisors for more information.

Asset-backed securities generally decrease in value as a result of interest rate increases, but may benefit less than other fixed-income securities from declining interest rates, principally because of prepayments.

Alternative investments may be either traditional alternative investment vehicles, such as hedge funds, fund of hedge funds, private equity, private real estate and managed futures or, non-traditional products such as mutual funds and exchange-traded funds that also seek alternative-like exposure but have significant differences from traditional alternative investments. The risks of traditional alternative investments may include: can be highly illiquid, speculative and not appropriate for all investors, loss of all or a substantial portion of the investment due to leveraging, short-selling, or other speculative practices, volatility of returns, restrictions on transferring interests in a fund, potential lack of diversification and resulting higher risk due to concentration of trading authority when a single advisor is utilized, absence of information regarding valuations and pricing, complex tax structures and delays in tax reporting, less regulation and higher fees than open-end mutual funds, and risks associated with the operations, personnel and processes of the manager. Non-traditional alternative strategy products may employ various investment strategies and techniques for both hedging and more speculative purposes such as short-selling, leverage, derivatives and options, which can increase volatility and the risk of investment loss. These investments are subject to the risks normally associated with debt instruments and also carry substantial additional risks. Investors could lose all or a substantial amount of their investment. These investments typically have higher fees or expenses than traditional investments.

Investing in foreign markets entails greater risks than those normally associated with domestic markets, such as political, currency, economic and market risks. **Investing in currency** involves additional special risks such as credit, interest rate fluctuations, derivative investment risk, and domestic and foreign inflation rates, which can be volatile and may be less liquid than other securities and more sensitive to the effect of varied economic conditions. In addition, international investing entails greater risk, as well as greater potential rewards compared to U.S. investing. These risks include political and economic uncertainties of foreign countries as well as the risk of currency fluctuations. These risks are magnified in countries with **emerging markets and frontier markets**, since these countries may have relatively unstable governments and less established markets and economies.

Participating in a new issue/syndicate is subject to availability. **Initial public offerings (IPOs)** are highly speculative and may not be appropriate for all investors because they lack a stock-trading history and usually involve smaller and newer companies that tend to have limited operating histories, less experienced management teams, and fewer products or customers. Also, the offering price of an IPO reflects a negotiated estimate as to the value of the company, which may bear little relationship to the trading price of the securities, and it is not uncommon for the closing price of the shares shortly after the IPO to be well above or below the offering price.

Asset allocation and diversification do not assure a profit or protect against loss in declining financial markets.

Growth investing does not guarantee a profit or eliminate risk. The stocks of these companies can have relatively high valuations. Because of these high valuations, an investment in a growth stock can be more risky than an investment in a company with more modest growth expectations.

Value investing does not guarantee a profit or eliminate risk. Not all companies whose stocks are considered to be value stocks are able to turn their business around or successfully employ corrective strategies which would result in stock prices that do not rise as initially expected.

Because of their narrow focus, **sector investments** tend to be more volatile than investments that diversify across many sectors and companies. **Technology stocks** may be especially volatile. Risks applicable to companies in the **energy and natural resources sectors** include commodity pricing risk, supply and demand risk, depletion risk and exploration risk.

REITs investing risks are similar to those associated with direct investments in real estate: property value fluctuations, lack of liquidity, limited diversification and sensitivity to economic factors such as interest rate changes and market recessions.

Artificial intelligence (AI) is subject to limitations, and you should be aware that any output from an AI-supported tool or service made available by the Firm for your use is subject to such limitations, including but not limited to inaccuracy, incompleteness, or embedded bias. You should always verify the results of any AI-generated output.

Environmental, social, and governance-aware investments (ESG) in a portfolio may experience performance that is lower or higher than a portfolio not employing such practices. Portfolios with ESG restrictions and strategies as well as ESG investments may not be able to take advantage of the same opportunities or market trends as portfolios where ESG criteria is not applied. There are inconsistent ESG definitions and criteria within the industry, as well as multiple ESG ratings providers that provide ESG ratings of the same subject companies and/or securities that vary among the providers. Certain issuers of investments may have differing and inconsistent views concerning ESG criteria

AI DATA CENTERS: THE HEART OF THE ECOSYSTEM

where the ESG claims made in offering documents or other literature may overstate ESG impact. As a result, it is difficult to compare ESG investment products or to evaluate an ESG investment product in comparison to one that does not focus on ESG. There is no assurance that an ESG investing strategy or techniques employed will be successful. Past performance is not a guarantee or a dependable measure of future results.

Rebalancing does not protect against a loss in declining financial markets. There may be a potential tax implication with a rebalancing strategy. Investors should consult with their tax advisor before implementing such a strategy.

Yields are subject to change with economic conditions. Yield is only one factor that should be considered when making an investment decision.

Credit ratings are subject to change.

The **indices** are unmanaged. An investor cannot invest directly in an index. They are shown for illustrative purposes only and do not represent the performance of any specific investment. The indices are not subject to expenses or fees and are often comprised of securities and other investment instruments the liquidity of which is not restricted. A particular investment product may consist of securities significantly different than those in any index referred to herein. Comparing an investment to a particular index may be of limited use.

The **indices selected by Morgan Stanley Wealth Management** to measure performance are representative of broad asset classes. Morgan Stanley Wealth Management retains the right to change representative indices at any time.

Performance of indices may be more or less volatile than any investment product. The risk of loss in value of a specific investment (such as with an investment manager or in a fund) is not the same as the risk of loss in a broad market index. Therefore, the historical returns of an index will not be the same as the historical returns of a particular investment product.

Disclosures

Morgan Stanley Wealth Management is the trade name of Morgan Stanley Smith Barney LLC, a registered broker-dealer in the United States. This material has been prepared for informational purposes only and is not an offer to buy or sell or a solicitation of any offer to buy or sell any security or other financial instrument or to participate in any trading strategy. Past performance is not necessarily a guide to future performance.

The author(s) (if any authors are noted) principally responsible for the preparation of this material receive compensation based upon various factors, including quality and accuracy of their work, firm revenues (including trading and capital markets revenues), client feedback and competitive factors. Morgan Stanley Wealth Management is involved in many businesses that may relate to companies, securities or instruments mentioned in this material.

This material has been prepared for informational purposes only and is not an offer to buy or sell or a solicitation of any offer to buy or sell any security/instrument, or to participate in any trading strategy. Any such offer would be made only after a prospective investor had completed its own independent investigation of the securities, instruments or transactions, and received all information it required to make its own investment decision, including, where applicable, a review of any offering circular or memorandum describing such security or instrument. That information would contain material information not contained herein and to which prospective participants are referred. This material is based on public information as of the specified date, and may be stale thereafter. We have no obligation to tell you when information herein may change. We make no representation or warranty with respect to the accuracy or completeness of this material. Morgan Stanley Wealth Management has no obligation to provide updated information on the securities/instruments mentioned herein.

The securities/instruments discussed in this material may not be appropriate for all investors. The appropriateness of a particular investment or strategy will depend on an investor's individual circumstances and objectives. Morgan Stanley Wealth Management recommends that investors independently evaluate specific investments and strategies, and encourages investors to seek the advice of a financial advisor. The value of and income from investments may vary because of changes in interest rates, foreign exchange rates, default rates, prepayment rates, securities/instruments prices, market indexes, operational or financial conditions of companies and other issuers or other factors. Estimates of future performance are based on assumptions that may not be realized. Actual events may differ from those assumed and changes to any assumptions may have a material impact on any projections or estimates. Other events not taken into account may occur and may significantly affect the projections or estimates. Certain assumptions may have been made for modeling purposes only to simplify the presentation and/or calculation of any projections or estimates, and Morgan Stanley Wealth Management does not represent that any such assumptions will reflect actual future events. Accordingly, there can be no assurance that estimated returns or projections will be realized or that actual returns or performance results will not materially differ from those estimated herein.

This material should not be viewed as advice or recommendations with respect to asset allocation or any particular investment. This information is not intended to, and should not, form a primary basis for any investment decisions that you may make. Morgan Stanley Wealth Management is not acting as a fiduciary under either the Employee Retirement Income Security Act of 1974, as amended or under section 4975 of the Internal Revenue Code of 1986 as amended in providing this material except as otherwise provided in writing by Morgan Stanley and/or as described at www.morganstanley.com/disclosures/dol.

Morgan Stanley Smith Barney LLC, its affiliates and Morgan Stanley Financial Advisors do not provide legal or tax advice. Each client should always consult his/her personal tax and/or legal advisor for information concerning his/her individual situation and to learn about any potential tax or other implications that may result from acting on a particular recommendation.

This material is disseminated in Australia to "retail clients" within the meaning of the Australian Corporations Act by Morgan Stanley Wealth Management Australia Pty Ltd (A.B.N. 19 009 145 555, holder of Australian financial services license No. 240813).

Morgan Stanley Wealth Management is not incorporated under the People's Republic of China ("PRC") law and the material in relation to this report is conducted outside the PRC. This report will be distributed only upon request of a specific recipient. This report does not constitute an offer to sell or the solicitation of an offer to buy any securities in the PRC. PRC investors must have the relevant qualifications to invest in such

AI DATA CENTERS: THE HEART OF THE ECOSYSTEM

securities and must be responsible for obtaining all relevant approvals, licenses, verifications and or registrations from PRC's relevant governmental authorities.

If your financial adviser is based in Australia, Switzerland or the United Kingdom, then please be aware that this report is being distributed by the Morgan Stanley entity where your financial adviser is located, as follows: Australia: Morgan Stanley Wealth Management Australia Pty Ltd (ABN 19 009 145 555, AFSL No. 240813); Switzerland: Morgan Stanley (Switzerland) AG regulated by the Swiss Financial Market Supervisory Authority; or United Kingdom: Morgan Stanley Private Wealth Management Ltd, authorized and regulated by the Financial Conduct Authority, approves for the purposes of section 21 of the Financial Services and Markets Act 2000 this material for distribution in the United Kingdom.

Morgan Stanley Wealth Management is not acting as a municipal advisor to any municipal entity or obligated person within the meaning of Section 15B of the Securities Exchange Act (the "Municipal Advisor Rule") and the opinions or views contained herein are not intended to be, and do not constitute, advice within the meaning of the Municipal Advisor Rule.

This material is disseminated in the United States of America by Morgan Stanley Smith Barney LLC.

Third-party data providers make no warranties or representations of any kind relating to the accuracy, completeness, or timeliness of the data they provide and shall not have liability for any damages of any kind relating to such data.

This material, or any portion thereof, may not be reprinted, sold or redistributed without the written consent of Morgan Stanley Smith Barney LLC.

© 2026 Morgan Stanley Smith Barney LLC. Member SIPC.

RSI1787615134316 08/2026