DSM CAPITAL PARTNERS

INVESTMENT PERSPECTIVE
November 2026

AI INFRASTRUCTURE

Demand Is Insatiable. Supply Must Catch Up.

Why the market's AI question has moved from whether demand is real to whether the supply chain, power, financing and local approvals can keep up.

+62%
GB300 contract rates since April
Source: SemiAnalysis

+103%
Google Cloud revenue growth, latest quarter
Source: Alphabet

+48%
Azure revenue growth, latest quarter
Source: Microsoft

160+ wks
Lead time for large transformers
Source: Reuters

The market has stopped asking whether AI demand is real. It is asking whether supply can keep up. Recent volatility in AI names reflects that shift: investors are weighing power, financing and local approvals, not end demand.

> **Demand is not the question. Supply is — and the shortage is raising prices for those with capacity in place.**

PART I · THE DEMAND EVIDENCE

- **Compute is getting more expensive.** GPU rental and contract prices have risen through 2026, including for older chips.
- **Cloud revenue is reaccelerating.** Azure, AWS and Google Cloud all grew faster in the latest quarter than a year earlier.
- **Suppliers are sold out.** Order visibility across the component chain extends into 2028–2030.

PART II · THE SUPPLY DEBATE

- **Power** is the leading constraint, and the industry is building around the grid.
- **Financing** needs are large, but pricing power covers higher funding costs.
- **Local and state pushback** is rising; developers are responding with local commitments.
- **Cheaper models** expand demand rather than shrink it.
- **Component equipment** to build data centers is in short supply.

PART I · THE DEMAND EVIDENCE

01 Compute is getting more expensive, not cheaper

If AI demand were fading, the price of renting compute would be falling. It is rising. Five-year GB300 contracts that cleared at $3.35–3.75 per GPU-hour in April now price at $5.50–6.00, according to SemiAnalysis, and early Vera

Rubin discussions are at $7.00–8.00. Ornn's B200 on-demand index reached a record $7.88 per GPU-hour on September 22, up 67.7% over twelve months.

Rental prices have risen across three generations of Nvidia GPUs

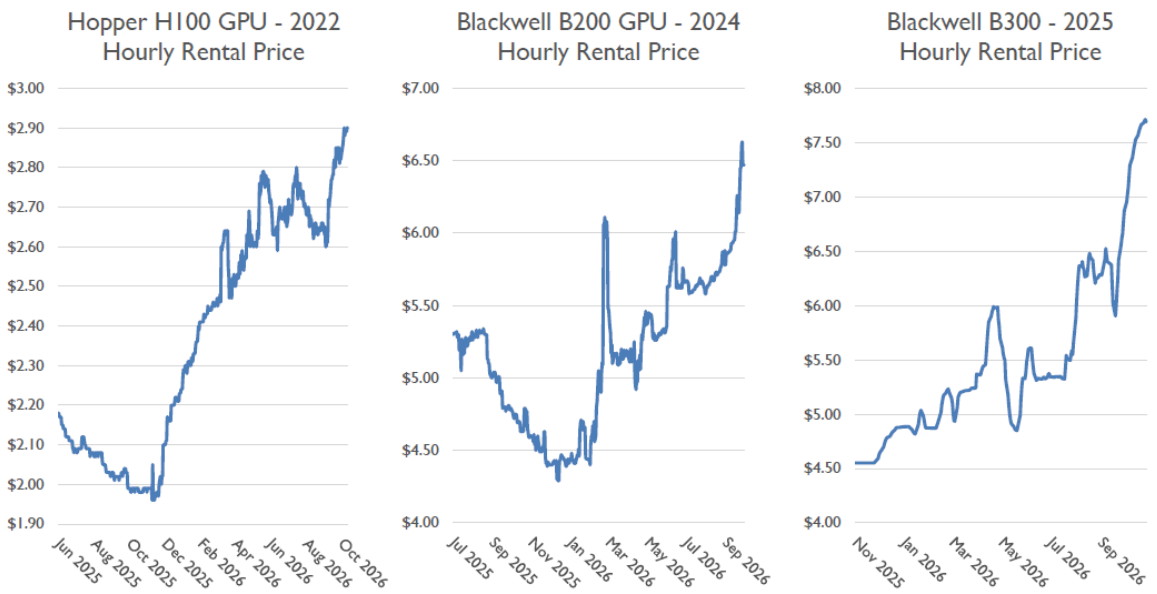

Hourly rental prices for Hopper H100 (launched 2022), Blackwell B200 (2024) and Blackwell B300 (2025) GPUs, mid-2025 to October 2026. Chart: Bloomberg, reproduced as published.

Even the 2022-vintage H100 now rents for about $2.90 an hour, up from about $2.00 a year ago. The newest B300 has risen from about $4.55 to about $7.70 since November 2025.

GB300 contract rates are up 62% since April

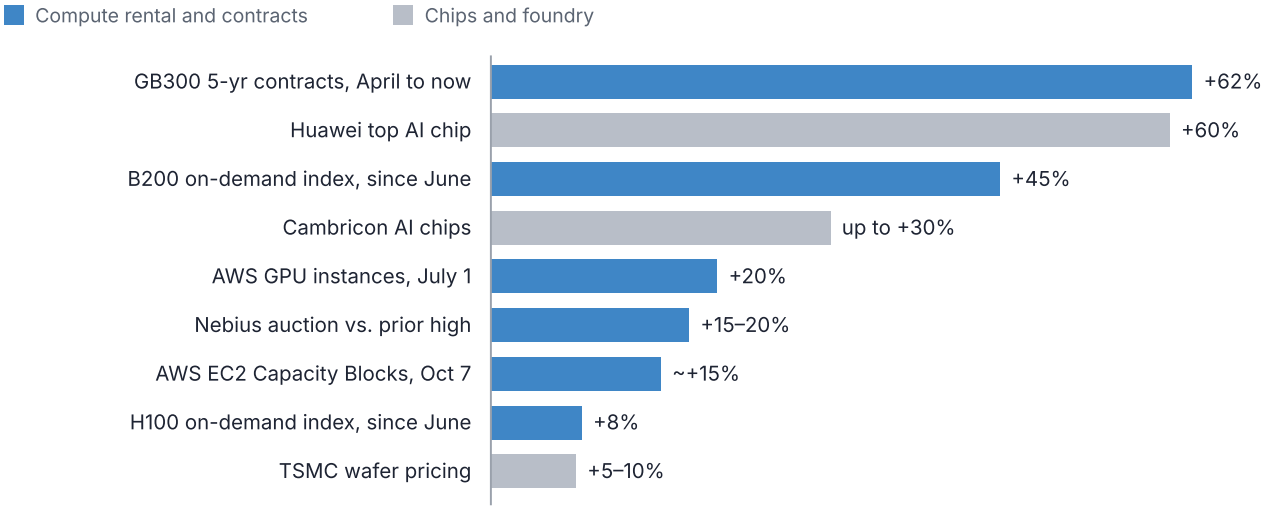

Reported 2026 price increases (%); ranges plotted at midpoint. Sources: SemiAnalysis (GB300, B200, H100); Amazon Web Services; Nebius; Bloomberg (Huawei, Cambricon); Nikkei Asia (TSMC).

Older chips are holding their value. Refurbished H100s, launched in 2022, are listed for resale at $29,000 to $58,000, at or above their reported launch price of about $25,000 to $30,000. CoreWeave recently renewed A100 capacity, a 2020 architecture, into 2029, and renewed H200 capacity at a premium to the original contract. That cuts against the bear case of rapid depreciation and stranded assets.

Three-year-old H100s are listed at or above their launch price

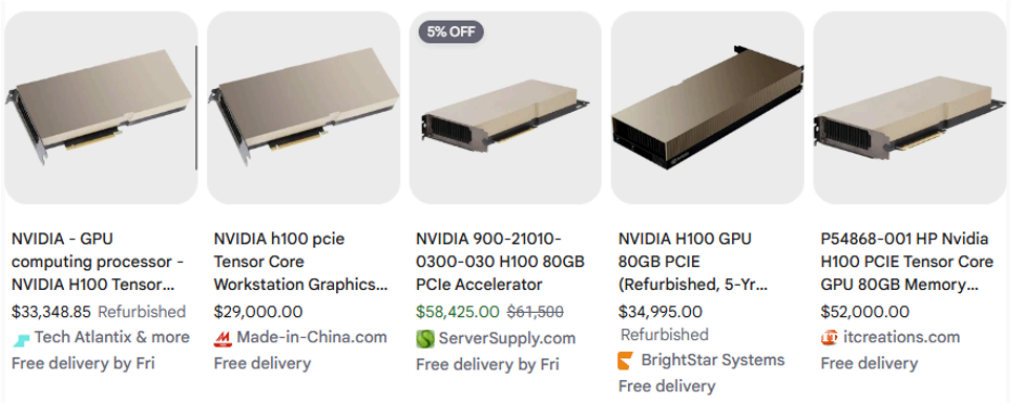

Retail listings for NVIDIA H100 80GB PCIe accelerators, new and refurbished, November 2026. Asking prices, not completed sales. Reported launch price about $25,000 to $30,000.

Better pricing is shortening paybacks. IREN reports revenue per IT megawatt has doubled, from $10 million on its Microsoft contract to $20 million on recent three-year deals, cutting GPU payback from three years to two. Alibaba puts its AI capex payback at about three years, potentially falling toward two.

02 Hyperscaler cloud revenue is reaccelerating

The clearest test of AI return on investment is whether end customers are paying for it. They are. Cloud growth at Microsoft, Amazon, Google and Oracle is higher in the latest quarter than a year earlier, at a scale where acceleration is hard to achieve.

Cloud revenue growth has accelerated at every major hyperscaler

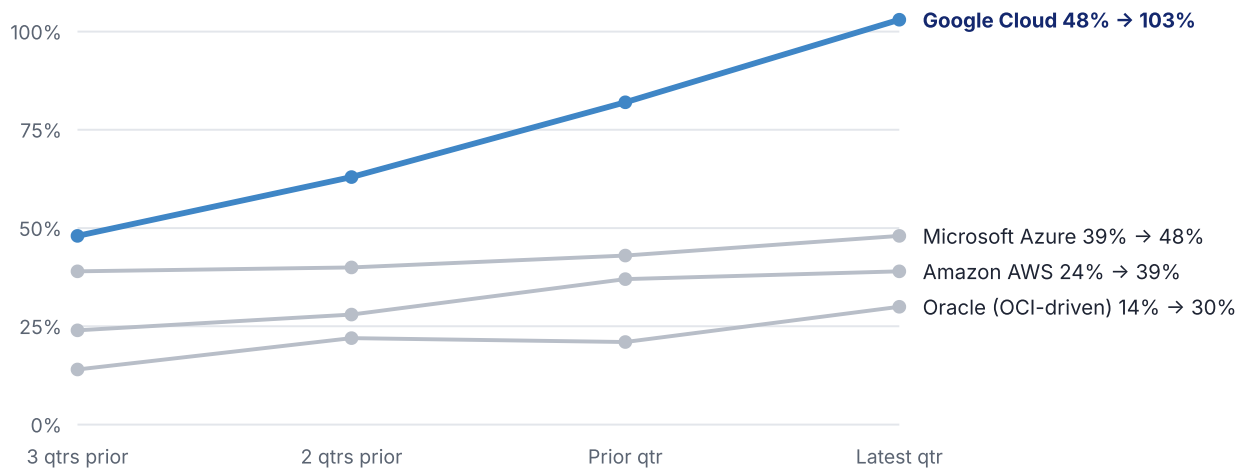

Year-over-year revenue growth, last four reported quarters. Sources: Microsoft, Amazon, Alphabet and Oracle quarterly reports; latest quarter as of October 30, 2026. Oracle figure is total company revenue.

Capex expectations for 2027 continue to rise, and free cash flow at the largest spenders may be thin or negative next year. We think out-year earnings estimates, which are starting to move up with revenue, matter more for these stocks than near-term cash flow.

Management commentary matches the numbers. Microsoft's Scott Guthrie describes “insatiable demand” for AI. AWS CEO Matt Garman says Amazon could sell every GPU it has to frontier labs and fills only about 60% of GPU requests. Google CEO Sundar Pichai calls securing compute for customers in 2027 “extremely difficult.”

03 The supply chain is sold out

Suppliers across the chain report demand ahead of capacity. Taiwan's AI server and component makers posted September sales growth of roughly 40% to 200% year over year. More telling is how far out they can see:

Through 2030	TSMC's order visibility extends to 2030, according to Digitimes; SK Hynix sees demand above supply through 2030.
Through 2028–30	Broadcom's Hock Tan expects compute demand to outrun supply at least through 2028–29.
Through 2027	Lumentum can meet only about 30% of the laser demand it sees for 2027.

TSMC raised its 2026 capex guidance in July to $60–64 billion, from $52–56 billion, and held it there in October. It expects to spend more on capex from 2026 to 2028 than in the previous three years. Suppliers do not commit capital at that scale against demand they doubt.

The shortage reaches the equipment that goes inside and around a data center. Large power transformers carry lead times of more than 160 weeks, and Micron says its customers will buy everything it makes. Infineon says demand for its AI power chips exceeds supply, and Vertiv's orders doubled last quarter.

PART II · THE SUPPLY DEBATE

04 Power is the leading constraint

The pace of AI build-outs is now set by power and physical infrastructure more than by chips. Large data center operators report more demand than they can serve for the next three to five years, and they consistently name power as the top constraint.

Power, memory and foundry capacity is booked into 2029 and beyond

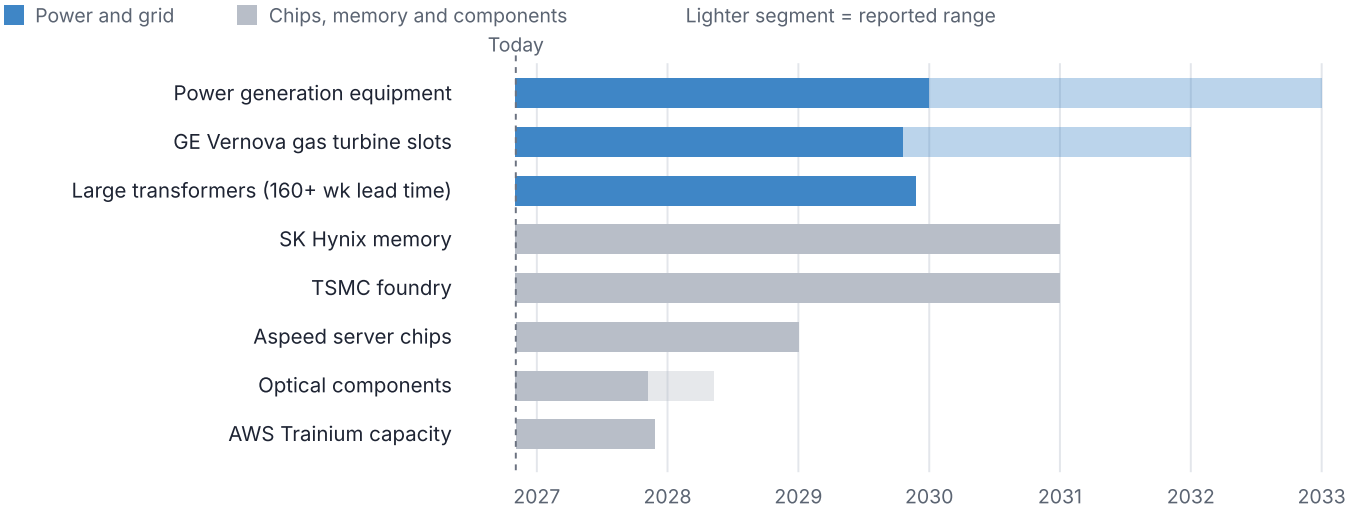

How far ahead capacity is booked or visible. Horizons placed at the end of the year cited. Sources: GE Vernova, SK Hynix, Lumentum, Amazon and industry commentary; Reuters (transformers); Digitimes (TSMC); Liberty Times (Aspeed).

Utilities face long lead times on electrical equipment, permitting and construction, and new nuclear capacity is unlikely before the 2030s. Necessity is driving invention. GE Vernova is booking gas turbine slots into 2029 and 2030, some into 2031, and sees no clear ceiling on pricing. Our research finds buyers refurbishing old equipment, relocating power plants from overseas and repurposing aircraft engines, because time to market now outweighs cost.

The most important response is going behind the meter: operators build generation next to the data center rather than wait for the grid. Others are adding fuel cells and engines adapted from heavy equipment. We expect the power constraint to persist through the end of the decade, which favors suppliers of generation and grid equipment.

Orders confirm the trend. ABB, Schneider Electric and Vertiv each reported strong data center orders, with ABB and Schneider citing triple-digit growth. JPMorgan expects the top four U.S. cloud providers to add more than $400 billion of data center capex in 2027, after more than $270 billion of growth in 2026.

05 Financing: pricing power covers the cost

The concern is that the buildout outgrows the money available to fund it. We expect Meta, Microsoft, Amazon, Alphabet and Oracle, taken together, to be free-cash-flow negative for the next year or two, so more of their data center spending will be funded in the debt and equity markets. Neoclouds, utilities and contractors are raising capital at the same time. Some investors worry the total cannot be raised, or only at much higher cost and on tighter terms.

The offset is pricing. Because demand exceeds supply, companies with compute in service are raising prices, often as contracts renew. CoreWeave, Nebius and Amazon have all done so (Section 1). Our estimate is that like-for-like renewals of contracts signed in 2023–2025 could reprice 30% to 100% higher. That would more than cover higher borrowing and supply-chain costs. We are watching funding conditions for the supply chain through 2028.

06 Local and state pushback: developers are responding

Local opposition to data centers is now a regular news story. Residents and their elected officials worry about higher electric bills, noise, and limited local economic benefit.

Pushback is moving from town halls to state capitals. In the November 3 elections, governors elected or re-elected in six states had called for pausing or tightening rules on data centers, according to Bloomberg. Votes have not been finalized.

Governors elected in six states have called for data center pauses or tighter rules

AI Buildout to Confront Blue Wave
Democratic gains in local races threaten more regulation, moratoriums

State	House	Senate	Governor	New governor's data-center stance
Texas	14D, 24R	D/R	D*	Called for an immediate moratorium on data centers, end tax giveaways
Pennsylvania	11D, 6R	D/R	D	Signed executive order to reinforce local decision-making while requiring stricter standards for energy use
Georgia	5D, 9R	D	D*	Called for a statewide temporary "pause" on new data center construction
New Mexico	3D	D	D	Called for an executive order upon taking office to pause all new large-scale data center proposals
Ohio	8D, 7R	D/R	D*	Supports conditional moratorium on data centers to ensure they cover utility costs, live by Ohio's environmental standards
Wisconsin	4D, 4R	D/R	D	Called for data-center regulation, supports repealing existing sales tax incentives and a pause of data centers during environmental review

Note: * indicates a flipped seat from Republican incumbents. Votes have yet to be finalized
Source: Bloomberg, local media
Bloomberg

Bloomberg, local media. * indicates a seat that changed party; votes not yet finalized. Chart: Bloomberg, reproduced as published.

If enacted, statewide pauses would delay projects in these states and could shift new capacity elsewhere.

The Texas pause looks conditional, not permanent. Texas holds about a quarter of planned US AI data center capacity. Its governor-elect, Gina Hinojosa, said she is "not ruling data centers out" and wants new construction paused until laws protect residents and give communities a say. The current pause began under Governor Abbott, and polls show about three-quarters of Texans oppose unregulated data center growth, so new rules look likely

whoever governs. We read this as a rulebook, not a ban: projects that pay their own way and bring their own power should be able to clear it.

Developers are responding directly rather than assuming acceptance. We see developers being explicit about the construction and permanent jobs a project creates and the local tax revenue it generates. Some offer more: affordable housing, paying the full cost of utility upgrades, or supplying their own power so local ratepayers are not affected. We think these commitments will go a long way toward easing local concerns. State-level pauses are a newer risk that we are watching.

07 Cheaper models expand demand

The concern is that open-weight models and Chinese competitors will drive down the cost of AI. Frontier labs would then miss revenue targets and need less compute.

We think the opposite is more likely, because of Jevons paradox: when the cost of a resource falls, total use rises by more than the price fell. OpenRouter's CEO has said a 10x price cut on an OpenAI model was followed by a 13x rise in usage. A large pharmaceutical company described the same dynamic at an industry conference: of 50 AI projects under consideration, a 50% fall in AI costs would make about 40 economically viable.

Token prices have fallen by more than half since June

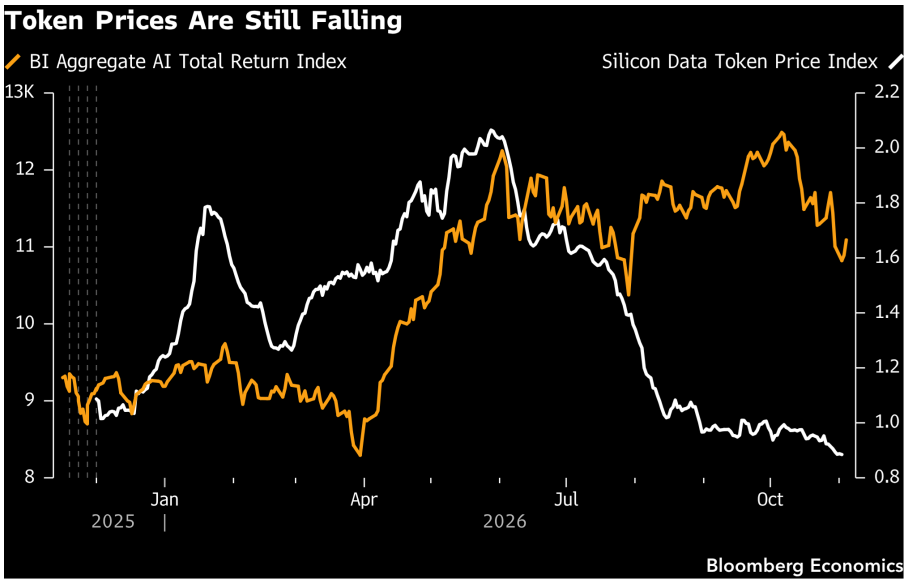

Silicon Data Token Price Index (right axis, white) and BI Aggregate AI Total Return Index (left axis, orange). Chart: Bloomberg Economics, reproduced as published.

Token prices are still falling. The Silicon Data Token Price Index has dropped by more than half since its June peak, while the cost of renting the GPUs that produce tokens has risen (Section 1). Cheaper tokens alongside dearer compute is consistent with Jevons paradox: usage is growing fast enough to absorb more compute even as the price per token falls.

Competition among model providers will stay intense. The most demanding workloads will keep using the newest models, and routine work will go to the cheapest option. Both need compute. Enterprise adoption is still early, which leaves room for demand to grow.

OUR VIEW

Demand is unending. Supply constraints are real. But each constraint is being addressed, and the shortage is raising prices and profits for those with capacity in place. We believe data centers will continue to be built in Texas, elsewhere in the United States and around the world. The debate is mainly about eliminating/minimizing “excessive” tax breaks, as well as creating and implementing zoning rules, not stopping construction.

We believe these companies are extremely undervalued. Their fundamentals have not changed: order books, pricing and earnings outlooks remain intact, and the recent selloff has been driven by policy headlines and positioning. We believe it is an opportune time to be a buyer of these names.

We manage the portfolio position by position with intensive fundamental analysis. Our approach is to own businesses with durable, predictable earnings power across the data center buildout: compute, cloud platforms, software, as well as the power and grid companies. Our valuation discipline governs buying, selling and sizing.

IMPORTANT LEGAL INFORMATION

This material is provided for informational purposes only. It is not intended to reflect a current or past specific recommendation, investment, legal, tax or accounting advice of any kind, or an offer or solicitation of an offer to buy or sell any securities or investment services in any jurisdiction where or to any person to whom it would be unauthorized or unlawful to do so. Except as otherwise specified, any companies, sectors, securities and/or markets discussed are solely for illustrative purposes regarding economic trends and conditions or investment process and may or may not be held by DSM Capital Partners LLC ("DSM") or other investment vehicles or accounts managed by DSM.

Price, EPS, and price-to-earnings ratio information discussed herein is provided only to illustrate market trends, and is not intended to represent the performance of any investment vehicle or account managed by DSM, nor any portion of any such investment vehicle or account. The information provided herein is current as of the date of issuance (or such other date as referenced herein) and is subject to change without notice.

This information should not be construed as research or investment advice. Returns are historical and past performance is no guarantee of future results and individual accounts and results will vary. The value of investments and the income derived from investments can go down as well as up. Investing entails risks, including possible loss of principal. Future returns are not guaranteed. There is also special risk considerations associated with international and global investing (especially emerging markets), small and mid-capitalization companies, or other growth and/or concentrated investment strategies. Investors are urged to consult with their financial advisors before buying or selling any securities.

This information may not be current, and DSM has no obligation to provide any updates or changes. Although the information has been obtained from sources believed to be reliable, there are no guarantees of accuracy, completeness or fairness. DSM has relied upon and assumed without independent verification the accuracy and completeness of some of the information. Opinions expressed are current opinions as of the date written. No part of this material may, without DSM's prior written consent, be (i) copied, photocopied or duplicated in any form, by any means, or (ii) distributed to any person that is not an employee, officer, director, or authorized agent of the recipient.

Certain statements herein are based on current expectations, estimates, projections, opinions and/or beliefs constituting "forward-looking statements," which can be identified by the use of forward-looking terminology such as "may," "expect," "anticipate," "project," "estimate," "intend," "target," or "believe" or similar terminology. No representation or warranty is made with respect to such statements and future events may differ materially from those contemplated herein. Projected earnings growth is shown for informational purposes only and is based on various assumptions, including historical performance for similar investments and/or current market conditions. Risks and uncertainties mean the actual growth could differ materially from the projected earnings growth. There is no guarantee that the projected earnings growth will occur. Projected earnings growth is theoretical and does not represent actual trading or the impact of economic or market factors and is not a projection of the return on any investment or of any client account.

Earnings per share, price to earnings ratios and other valuation metrics/ratios do not guarantee investment success. DSM may not be successful in predicting EPS growth or P/E ratio and other ratios/metrics, and as a result, client portfolios may fail to meet investment objectives and investors may experience losses. Investment outcomes may be materially worse than presented herein. The price-to-earnings ratio is an equity valuation metric defined as market price per share divided by annual earnings per share. It can be simplified as the ratio of total market capital value over earnings. Earnings per share is a company's total earnings or net income divided by its shares outstanding. Earnings per share can be found by taking the net income and dividing it by the basic or diluted number of shares outstanding.

Certain information contained herein is based on DSM proprietary models. The use of models and/or tools does not guarantee investment success. Models/tools apply statistical methods and a series of fixed assumptions to derive estimates of asset class performance. Reasonable people may disagree about the appropriate assumptions. Models/tools also have limitations. For instance, assumptions may not be consensus views, or the models/tools may not be updated to reflect current economic, market or political conditions. Models/tools should not be relied upon to make predictions of actual future performance. DSM has no obligation to provide updates or changes to such data. DSM projections are not guarantees of future results and there is no representation that these securities were, would have been or will be profitable. Performance-based criteria was not used to compile the contribution to return lists. The holdings were selected over the measurement period shown in a mechanical and objective manner by using a calculation to show their relative impact on overall performance; they were not included or excluded for any other reason. The holdings identified my not represent all of the securities purchased, sold, or recommended. The calculation computes the contribution of each holding by calculating the weight (i.e., percentage of the total account) invested in each holding multiplied by the rate of return for that holding during the measurement period. The result shows each holding's contribution to the overall return during the measurement period. A complete list of all recommendations made by DSM is available upon request.