



DATA CENTER PORTFOLIO · INVESTMENT THESIS

AI Needs More Than Chips. Here's How To Invest In What It's Actually Running Out Of.

Everyone owns the obvious AI names already. Gregor's got the ten stocks behind the power, memory, and fiber the buildout can't run without.



Gregor Becker

Editor-in-Chief · 1 Oct 2026 · 24 mins

CONTENTS

In this thesis

Part I: The Buildout in Numbers	6
Who is spending.	6
How far ahead does it run?	6
Who is paying.	7
What the numbers hide.	7
Part II: The Power Problem	8
How much power.	8
The grid can't keep up.	8
The queue.	9
What regulators are doing.	9
Two ways around the problem.	9
What would change the picture.	10
Part III: The Memory Squeeze	11
What HBM is and who makes it.	11
A sold-out market.	11
The price of scarcity.	11
The cycle warning.	12
Part IV: The Network Nobody Talks About	13
Fibre per rack.	13

Corning's second act.	13
A reason for humility.	13
Part V: The Compute Layer and Its Customers	15
The chips and the machines that make them.	15
The customers who are also the investors.	15
The circularity question.	16
Part VI: What's Already in the Price	17
What the market pays.	17
The bear case in numbers.	17
What a slowdown would do.	17
Part VII: How the Basket Is Built	19
Why these ten together.	19
How it is run.	19
Who it is for.	19
Final thought	20
Risks	21
The portfolio behind this thesis	23
Disclaimer	24

The Case In A Nutshell.

- 1 The spending is enormous and still rising. FactSet ([FDS](#)) puts calendar 2026 capital expenditure by the five big hyperscalers at close to \$800 billion once finance leases and customer pre-payments are counted, and Dell'Oro says global data centre capex is on track to pass \$1 trillion in 2026, with the four largest US cloud companies' data centre spending up 78% year on year in the first quarter.
- 2 Power is the binding constraint. The IEA found that global data centre electricity demand grew 17% in 2025 and that AI-focused facilities grew 50%, and it expects consumption to roughly double from 485 terawatt-hours in 2025 to 950 terawatt-hours in 2030. Wood Mackenzie estimates US supply of power transformers ran 30% short of demand in 2025 and distribution transformers 10% short.
- 3 Memory is the second shortage. Micron ([MU](#)) says it has agreed price and volume for its entire calendar 2026 high-bandwidth memory supply, and Counterpoint puts the market's supply in just three hands: SK hynix ([KRX: 000660](#)) 50% of HBM revenue in the second quarter of 2026, Samsung ([KRX: 005930](#)) 33% and Micron 18%.
- 4 The network is the third. A Corning ([GLW](#)) executive says the density of optical fibre in new AI networks is about ten times that of traditional cloud front-end networks, and Corning's optical communications sales rose 32% to \$2.07 billion in the second quarter, on top of a multiyear agreement with Meta ([META](#)) worth up to \$6 billion.
- 5 Scarcity gives pricing power, and the results already show it. Micron's gross margin reached 86.8% in its latest quarter, and Bloom Energy's ([BE](#)) revenue more than doubled year on year in the second quarter, to \$1.07 billion.
- 6 The basket pairs the compute names everyone already owns with the power, cooling, memory and fibre businesses that decide whether that compute can be built, powered and connected on schedule. It is a bet on the plumbing as much as on the product.

Ask most investors how to play AI and they'll give you the same five names: Nvidia ([NVDA](#)), Microsoft ([MSFT](#)), Alphabet ([GOOG](#)), and a couple of other hyperscalers. They're not wrong, exactly, but they're only telling you half the story. The AI buildout has quietly turned into the biggest infrastructure project in corporate history, and its real constraints have moved away from chips and toward the unglamorous stuff underneath them: electricity, memory, and the glass fibre that carries data between racks. A data centre that can't get power doesn't switch on. A GPU that can't get memory doesn't ship. A cluster that can't be wired together doesn't work as one machine. This basket holds the obvious compute names alongside the overlooked infrastructure ones, because right now the second group is where the bottleneck actually sits. It is also where the numbers are most striking: in the quarter to August, Micron ([MU](#)) booked \$54.2 billion of revenue, more than the whole of its previous fiscal year.

The Buildout in Numbers

It's easy to lose the scale of what's happening in the noise of quarterly earnings, so let's start with the numbers, and with who is spending and who is paying.

Who is spending.

Yahoo Finance's roundup of the July earnings has Alphabet ([GOOG](#)) guiding 2026 capital expenditure to \$195-205 billion, raised from \$180-190 billion earlier in the year, Meta ([META](#)) lifting its 2026 floor twice to \$135-145 billion, and Amazon ([AMZN](#)) guiding to roughly \$220 billion of cash capex, up from about \$200 billion. Microsoft's ([MSFT](#)) fiscal-year cash capex was \$115.95 billion, up 80%. FactSet ([FDS](#)) adds finance leases and customer pre-payments and arrives at close to \$800 billion for calendar 2026 across the five biggest, Oracle ([ORCL](#)) included. Dell'Oro, which tracks the whole industry rather than five companies, says global data centre capex is on track to exceed \$1 trillion in 2026 and will pass \$3 trillion by 2030.

How far ahead does it run?

McKinsey's widely quoted estimate is that data centres worldwide will need \$6.7 trillion of investment by 2030, of which \$5.2 trillion is for AI-capable capacity and \$1.5 trillion for traditional workloads, built around 156 gigawatts of AI demand.

Roughly 60% of the AI total goes to technology developers, 25% to energy suppliers and 15% to builders. DCD, which reported it, headlined the estimate with a "maybe", which is the right amount of scepticism for any forecast of this length. What matters is the split: a quarter of the money goes to power, and that is the layer where the shortages are.

Who is paying.

For years the hyperscalers funded this out of operating cash flow. That has changed. FactSet ([FDS](#)) reports that debt's share of incremental capex rose from 9% in fiscal 2024 to 32% over the latest twelve months, that aggregate hyperscaler debt is around \$700 billion, and that free cash flow in fiscal 2026 is expected to be close to zero or negative for every one of the five except Alphabet ([GOOG](#)) and Microsoft ([MSFT](#)). Alphabet raised \$84.75 billion of equity in June 2026, of which \$44.75 billion was earmarked for AI capex. Oracle's ([ORCL](#)) fiscal 2026 capex was \$55.7 billion against operating cash flow of \$32.0 billion, leaving free cash flow of negative \$23.7 billion, and it raised \$43 billion of debt and \$5 billion of equity during the year.

What the numbers hide.

Dell'Oro itself is careful to say that a significant part of the capex growth is price. Its data show component revenue up 182% year on year, with DRAM and storage cost per bit more than doubling. When memory prices double, a data centre budget buys fewer servers than last year's did. That is not a reason to doubt the buildout. It is a reason to think the bottleneck is real, because rising prices are what scarcity looks like.

The Power Problem

That is not a reason to doubt the buildout.

Power has become the industry's most-discussed bottleneck, and the data now back it up from several independent directions.

How much power.

The IEA's April 2026 report found that global data centre electricity demand grew 17% in 2025 and that AI-focused data centres grew 50%, consuming 155 terawatt-hours. It expects total data centre consumption to roughly double, from 485 terawatt-hours in 2025 to 950 in 2030. In the US, Lawrence Berkeley National Laboratory estimated that data centres used 4.4% of the country's electricity in 2023, or 176 terawatt-hours, and projected that share to rise to between 6.7% and 12% by 2028.

The grid can't keep up.

The equipment to deliver that power is in short supply. Wood Mackenzie estimated that US power transformers were 30% short of demand in 2025, and distribution transformers 10% short, after demand grew 116% and 41% respectively since 2019. The US imports about 80% of its power transformers and half of its distribution transformers. A Wood Mackenzie survey of the second quarter of 2025, reported by pv magazine, put average lead times at 128 weeks for large power transformers and 144 weeks for generator step-up units, which is two and a half to nearly three years. On the generation side, GE Vernova ([GEV](#)) reported in July 2026 that its gas turbine backlog and slot reservations had reached 116 gigawatts, with 125 gigawatts expected by year end, and that it is taking reservations for delivery in 2031. About 80% of those customers are utilities and power producers and 20% data centres, which is a reminder that data centres are competing with everyone else for the same machines.

The queue.

The wait to connect is longer still. Lawrence Berkeley's most recent "Queued Up" analysis counts around 2,290 gigawatts of generation and storage waiting for interconnection at the end of 2024, found that the typical project that reached operation in 2024 had spent 55 months in the queue, and found that only 13% of the capacity that entered queues between 2000 and 2019 had been built by the end of 2024. The grid operator PJM, which serves the largest data centre market in the US, now forecasts its summer peak demand growing 3.6% a year to around 222 gigawatts by 2036, and its capacity auction for 2026-27 cleared at \$329.17 per megawatt-day, against \$28.92 for 2024-25.

What regulators are doing.

Policy is moving, slowly. In June 2026 the Federal Energy Regulatory Commission issued show-cause orders to the six regional grid operators on how they connect large loads and co-located generation, and the Department of Energy asked it to open a rulemaking in October 2025. States are acting too: the Ohio tariff approved for AEP ([AEP](#)) in July 2025 requires data centres above 25 megawatts to pay for at least 85% of contracted capacity under agreements of up to twelve years. Within months, AEP Ohio's large-load forecast fell from 30 gigawatts to 13, which shows how much of the queue was speculative.

Two ways around the problem.

There are really only two answers to a power shortage: make your own, or use what you have far more efficiently. Both are businesses in this basket.

Bloom Energy ([BE](#)) builds fuel cells that generate electricity on site. Its chief executive says all the major US hyperscalers and more than a dozen neoclouds, AI labs and colocation operators have approved its power solutions. Second-quarter revenue was \$1.065 billion, up 165.5%, with GAAP net income of \$196.3 million against a loss of \$42.6 million a year earlier, and full-year guidance is \$3.9-4.2 billion. Oracle ([ORCL](#)) has contracted 1.2 gigawatts of Bloom systems with the option of up to 2.8 gigawatts, and AEP ([AEP](#)) has agreed a 20-year, 1-gigawatt arrangement. This is also a business with a long record of

losses: Bloom made net losses of \$302 million in 2023, \$29 million in 2024 and \$88 million in 2025, and a short-seller report published in July 2026 by Hunterbrook challenged the quality of its reported backlog, which Bloom called false and misleading. The turn to profit is real but new, and we treat it that way.

Vertiv ([VRT](#)) sits at the other end of the same problem. It describes its offer as end-to-end infrastructure in which "power, cooling, IT, and services work in unison", and sells uninterruptible power supplies, energy storage, thermal management and liquid cooling for AI environments. The need is obvious in the rack: Nvidia's ([NVDA](#)) GB200 and GB300 racks draw roughly 120 to 130 kilowatts, the coming Vera Rubin generation roughly 190 to 230, and Nvidia has shown a 600-kilowatt rack for 2027. Uptime Institute's 2025 survey found only 22% of operators using direct liquid cooling, with 61% of the rest saying they would consider it. Vertiv's backlog was \$15.0 billion at the end of 2025, up 109%, with fourth-quarter orders up about 252% and a book-to-bill of around 2.9. It guides 2026 sales to \$13.8-14.2 billion, up 30 to 32% organically. Execution is not guaranteed, though: second-quarter revenue of \$3.27 billion missed consensus, management cited "multiphase execution and supply chain congestion", and the shares fell about 17%.

What would change the picture.

Power is the constraint most likely to ease first if the buildout slows, and the one most exposed to political and regulatory change. New generation and transmission take years but they do arrive. We should be honest that part of what we are paying for in the power names is scarcity that will not last forever.

The Memory Squeeze

Memory is the second constraint that almost nobody outside the industry was watching a few years ago. It is also the one that shows how a shortage in one corner of the supply chain can ripple through all of electronics.

What HBM is and who makes it.

High-bandwidth memory is the specialised memory that sits right next to a GPU and feeds it data fast enough to keep it working. Only three companies make it at volume. Counterpoint's second-quarter 2026 data put SK hynix ([KRX: 000660](#)) at 50% of HBM revenue, Samsung ([KRX: 005930](#)) at 33% and Micron ([MU](#)) at 18%. IDC measured SK hynix at 63.2% of 2025 sales, so the order of the three can shift quickly, but the number of players cannot.

A sold-out market.

Micron ([MU](#)) says it has completed agreements on price and volume for its entire calendar 2026 HBM supply, including HBM4, and that its HBM3E and HBM4 volume is booked into 2027. Its own estimate is that the HBM market grows from about \$35 billion in 2025 to about \$100 billion in 2028, a milestone it has pulled forward by two years. Samsung ([KRX: 005930](#)), for its part, expects HBM to take close to 30% of the industry's DRAM wafer capacity in 2027, up from about 20% now, and says HBM needs about three times the wafer area of conventional DRAM. That ratio is the reason the squeeze spills over.

The price of scarcity.

Every wafer that goes into HBM is a wafer that doesn't go into ordinary memory. TrendForce put conventional DRAM contract price increases at roughly 90 to 95% quarter on quarter in the first quarter of 2026 and 58 to 63% in the second, and forecasts a further

13 to 18% in the third, with server memory supply growing 15 to 20% a year. The result is visible in Micron's (MU) accounts. In the quarter ended August 2026 it booked revenue of \$54.23 billion and a gross margin of 86.8%. Fiscal 2026 revenue was \$133.19 billion and GAAP net income \$84.97 billion. It guides the next quarter to \$61.5 billion. It is spending to respond, with net capital expenditure of \$27.37 billion last year, and it is collaborating with Nvidia (NVDA) on a custom HBM4E.

The cycle warning.

Every wafer that goes into HBM is a wafer that doesn't go into ordinary memory.

Micron (MU) is also a cyclical business, and the market knows it. On Yahoo Finance's September figures it trades at around six times forward earnings, a multiple that says investors expect today's profits to be temporary. Memory shortages have a history of ending abruptly once new supply arrives. That is why we care less about this quarter's price spike than about how long the allocations are locked in and how fast new capacity can come on stream.

The Network Nobody Talks About

The last overlooked layer is the physical network that moves data between all those GPUs, and it is a much bigger business than it used to be.

Fibre per rack.

A Corning (GLW) executive told investors that the density of fibre in these new AI networks is "about 10x what you might have seen in enterprise, traditional cloud, kind of front-end networks". The reason is that training and running large models means thousands of chips acting as one machine, and that only works if they can talk at enormous bandwidth. Dell'Oro reports that sales of AI back-end network switches overtook front-end switches for the first time in the second quarter of 2026, with 800-gigabit ports the dominant speed.

Corning's second act.

Corning's (GLW) second-quarter core sales were \$4.74 billion, up 17%. Optical communications was the largest segment at \$2.07 billion, up 32%, with segment net income up 77% to \$438 million, while enterprise networks grew 65%. It has signed a multiyear agreement with Meta (META) worth up to \$6 billion through 2030, a multiyear multibillion-dollar supply agreement with Amazon (AMZN), and a partnership with Nvidia (NVDA) under which US optical connectivity manufacturing capacity grows tenfold and US fibre capacity by more than half. Its Springboard plan targets \$40 billion of annual sales by the end of 2030.

A reason for humility.

Corning (GLW) also lived through the last great fibre buildout. Its revenue rose from \$3.5 billion in 1998 to \$7.1 billion in 2000, about 70% of it from telecoms. Its shares peaked at \$113 in September 2000 and fell to \$1.54 by August 2002, optical revenue

dropped from \$5 billion to \$1.6 billion, 16,400 jobs went and it booked a \$5.5 billion loss, according to the Baltimore Sun and Forbes at the time. It is easier to believe this time is different when the customers are the best-funded companies in the world, but we hold Corning for its contracts and order book, not its story, and we remember the shares can fall by almost half in a month: they did in July 2026, when a revenue miss of under 3% knocked 46% off the price.

The Compute Layer and Its Customers

None of this makes the obvious AI names the wrong place to be. It makes them half of the answer.

The chips and the machines that make them.

Nvidia ([NVDA](#)) reported second-quarter fiscal 2027 revenue of \$96.2 billion, up 106%, of which data centre was \$89.0 billion, up 117%, at a 75.0% gross margin, and guided the next quarter to \$108 billion. It has \$279 billion of supply commitments, mostly memory for its Vera Rubin generation, which shows how dependent the chip leader is on the memory makers. AMD's ([AMD](#)) second-quarter revenue was \$11.5 billion, up 50%, with data centre at \$6.7 billion, up 107%.

ASML ([ASML](#)) is the quieter chokepoint behind them. It is the world's only manufacturer of extreme-ultraviolet lithography systems, according to its annual report, and shipped 48 of them among 327 lithography systems in 2025 on net sales of €32.7 billion. It guides 2026 sales to €43-45 billion and plans around 30% more low-numerical-aperture EUV capacity in 2027 than the roughly 65 units of 2026. Advanced packaging is the other chokepoint: TSMC's ([TSM](#)) chief executive said in June that CoWoS capacity was sold out through 2026, and TrendForce expects the gap between supply and demand to narrow from about 20% to about 10% by the end of the year.

The customers who are also the investors.

Microsoft ([MSFT](#)), Alphabet ([GOOG](#)) and Oracle ([ORCL](#)) are the hyperscalers spending the hundreds of billions described above. Oracle in particular has made the biggest bet in proportion to its size. It reports \$638 billion of remaining performance obligations, up 363%, of which \$75 billion is prepaid or customer-supplied hardware, against the negative \$23.7 billion of free cash flow and the \$48 billion raised in the year. It said it does not expect to issue additional debt in calendar 2026. That is the most levered balance sheet in the group to a continued build-out, in both directions.

The circularity question.

Here is where we should be most candid. Nvidia ([NVDA](#)) announced in September 2025 a letter of intent to invest up to \$100 billion in OpenAI as capacity was deployed; reports in early 2026 said that had become a \$30 billion equity stake. AMD ([AMD](#)) agreed to supply six gigawatts of GPUs to OpenAI and granted it a warrant to buy up to 160 million AMD shares at a cent each, vesting as milestones are met. Oracle ([ORCL](#)) signed a \$300 billion five-year contract with OpenAI. OpenAI's chief executive has talked about \$1.4 trillion of commitments over eight years, and the company later told investors it was targeting around \$600 billion of compute spending by 2030. HSBC ([HSBC](#)) calculated that OpenAI would need to find \$207 billion to meet its commitments. None of this is a sign of trouble in itself, but it does mean the system is more tightly coupled than the headline demand numbers suggest. Owning the infrastructure layer as well as the compute layer spreads that exposure across more points in the chain.

What's Already in the Price

It would be dishonest to write about this basket without saying that a lot of the story is no secret.

What the market pays.

FactSet's (FDS) data show the S&P 500 on a forward price-to-earnings ratio of 19.2 in late September, in line with its 5-year average of 19.8, with the Magnificent Seven on 23.0 against 17.3 for everything else. Yahoo Finance's screens put Nvidia (NVDA) on roughly 24 times forward earnings, AMD (AMD) on 72 and Micron (MU) on around 6. The market is not paying peak multiples for peak growth. It is paying a low multiple for the memory cycle and a high one for AMD's growth, which is a sign that investors are discriminating between what looks durable and what looks cyclical.

The bear case in numbers.

The sceptics' argument is specific. Bain has estimated that around \$2 trillion of new annual revenue is needed by 2030 to justify the spending, leaving a gap of around \$800 billion a year. Goldman Sachs (GS) has said the hyperscalers need around \$300 billion of annual AI revenue to break even. And the depreciation debate matters: if the five biggest buyers are depreciating AI hardware over five to six years when its economic life is two to three, their reported earnings, and their ability to keep spending, are overstated.

What a slowdown would do.

If hyperscaler capex guidance were cut, the high-multiple names would fall first and hardest. That is the central risk to the basket, and no amount of analysis of any single company makes it go away. What the infrastructure names offer is a different mix of risks:

more tied to physical constraints and contracted order books, less to sentiment about the next model release. That is an argument for owning both layers, not for owning either one alone.

How the Basket Is Built

This is a concentrated portfolio on purpose. It is built around a single theme, the AI build-out, and it should be understood and sized as such.

Why these ten together.

The basket pairs the obvious AI compute names with the infrastructure layer underneath them. Nvidia ([NVDA](#)) and AMD ([AMD](#)) are the chips the whole build-out runs on. Microsoft ([MSFT](#)), Alphabet ([GOOG](#)) and Oracle ([ORCL](#)) are the customers spending the money. ASML ([ASML](#)) is the lithography chokepoint. Micron ([MU](#)) covers the memory squeeze directly. And Bloom Energy ([BE](#)), Vertiv ([VRT](#)) and Corning ([GLW](#)) are the overlooked third of the basket: the power, cooling and networking businesses that determine whether all that compute can actually be built, powered and connected on schedule, rather than whether the demand for it exists.

How it is run.

Every position is bought at an equal share of the portfolio at the time it is added, and winners are allowed to run, so a holding that does well becomes a larger share of the book on purpose. Names leave when the thesis breaks, not when they get expensive or fall in the rankings. The result is a portfolio that behaves like a conviction book rather than an index, which is what you should expect from something with this few names.

Who it is for.

It is for investors who already believe in the AI build-out and want exposure to the parts of it that the crowd is not already holding. It is not a substitute for a diversified core, and it will fall harder than the market if sentiment on AI spending turns.

Final thought

The AI story that's already priced into the obvious chip and cloud names is real, but it isn't the whole story. Right now the industry's own constraint has moved to power, memory and fibre, and the companies solving those problems are only starting to be valued as AI infrastructure plays rather than as an industrial-gas company, a legacy IT-hardware supplier, a memory maker or a glassmaker. The evidence is in the results: record margins at the memory maker, a tripling of revenue at the fuel-cell company, a \$15 billion backlog at the cooling company. The risks are in the same results. If you are going to be exposed to AI at all, be exposed to the layer that decides how fast it can grow, and size it knowing the cycle can turn.

RISKS

What could go wrong.

- 1 Capex cycles reverse. Hyperscaler data centre capex grew just 9% in 2022 after a 20% average over the prior two years, and Meta ([META](#)) cut roughly \$4 billion of data centre spending that year, according to Bisnow's reporting. In February 2025 analysts at TD Cowen reported that Microsoft ([MSFT](#)) had cancelled leases for a couple of hundred megawatts.
- 2 The buildout is increasingly funded with debt, not cash flow. FactSet ([FDS](#)) found that debt's share of incremental hyperscaler capex rose from 9% in fiscal 2024 to 32% over the latest twelve months, with aggregate hyperscaler debt around \$700 billion, and free cash flow close to zero or negative for all of the five except Alphabet ([GOOG](#)) and Microsoft.
- 3 Some of the money is circular. Nvidia ([NVDA](#)) agreed to invest in OpenAI, which buys Nvidia systems. AMD ([AMD](#)) granted OpenAI a warrant for up to 160 million shares as part of a six-gigawatt chip deal. Oracle's ([ORCL](#)) contract with OpenAI is worth \$300 billion over five years from 2027. If end demand disappoints, that loop unwinds both ways.
- 4 A lot of the capex growth is price, not volume. Dell'Oro notes that a significant portion of reported capex growth reflects higher equipment prices, with memory and storage costs per bit more than doubling, so spending can rise while physical capacity grows far less.
- 5 Scarcity does not last. Memory prices that rose by roughly 90% in a single quarter, transformer shortages and turbine backlogs all invite new capacity, and each of those new factories eventually brings pricing power back down. Yahoo Finance's September screens put Micron ([MU](#)) at around six times forward earnings, which tells you how the market treats peak-cycle profits.
- 6 Accounting and valuation. Meta, Amazon ([AMZN](#)), Microsoft, Alphabet and Oracle depreciate AI hardware over five to six years, while critics such as Michael Burry argue two to three. Bain has estimated that around \$2 trillion of new annual revenue is needed by 2030 to justify the spending. If either argument proves right, the earnings of the whole supply chain are worth less than they look.

7

Execution and volatility. Vertiv's ([VRT](#)) shares fell about 17% on the day it reported second-quarter results that missed revenue expectations, and Corning's ([GLW](#)) shares fell roughly 46% across July 2026 on a revenue miss of under 3%. The overlooked names can fall as sharply as the obvious ones.

THE PORTFOLIO

The portfolio behind this thesis.



Gregor Becker

Editor-in-Chief

Gregor runs the Data Center Portfolio: from the screen that picks the holdings to the weekly call on each one.

WHAT A SUBSCRIBER GETS



Weekly Signals

A full deep-dive on every stock we add: financials, valuation, and entry logic.



The Oxford Stance

A clear Add, Hold, Trim or Exit call on every conviction holding, every week.



Catalyst Alerts

The earnings, data releases and catalysts that could move each position, mapped to your holdings.



Every holding, unlocked

Every ticker, every weight, and the full trade history from day one.

WHY IT IS WORTH HAVING



A view, not a hedge. Every holding carries a clear call, so you always know where we stand and why.



A portfolio, not a watchlist. A concentrated set of conviction holdings, screened from a much larger universe and reviewed every Monday.



Research, not management. We never touch your money. Your assets stay in your own brokerage account, in your name, and you decide what to act on.

[Read more and subscribe](#)

oxford-intelligence.com/datacenter

Disclaimer

Disclaimer: These articles are provided for information purposes only. Occasionally, an opinion about whether to buy or sell a specific investment may be provided. The content is not intended to be a personal recommendation to buy or sell any financial instrument or product, or to adopt any investment strategy as it is not provided based on an assessment of your investing knowledge and experience, your financial situation or your investment objectives. Past performance is not indicative of future results. The value of your investments, and the income derived from them, may go down as well as up. You may not get back all the money that you invest. The investments referred to in this article may not be suitable for all investors, and if in doubt, an investor should seek advice from a qualified investment advisor.

Information only. This document is published by Oxford Market Intelligence for general information. It is not an offer, or a solicitation of an offer, to buy or sell any security or financial instrument, and it is not investment, legal, tax or accounting advice. It does not take into account the objectives, financial situation or needs of any particular reader.

Risk of loss. All investing involves risk, including the possible loss of the money you invest. Equity prices can fall as well as rise, sometimes sharply and without warning. A concentrated portfolio built around a single theme will usually be more volatile than a diversified one, and a setback affecting the theme can hit every holding at the same time. Nothing in this document is a guarantee, and no particular result is promised or implied.

Past performance and modelled results. Any figures describing the past performance of a portfolio, a stock or a sector are historical and are not a reliable indicator of future results. Where a portfolio's track record is calculated from a rules-based model rather than from trades made with real money, it does not reflect the impact of dealing costs, spreads, taxes or the timing of an individual reader's trades, and actual results would have differed.

Forward-looking statements. Statements about the future, including market forecasts, growth projections and views on how a company or sector may perform, reflect opinion at the date of writing and are subject to risks and uncertainties. Actual outcomes may differ materially. Opinions may change without notice, and we are under no obligation to update this document.

Sources and accuracy. Information is taken from sources we believe to be reliable, including company filings, market data providers and third-party research, but we do not guarantee its accuracy or completeness. Data is not audited. Figures quoted from third parties remain the responsibility of those sources.

Companies named. References to listed companies, and the tickers shown beside them, are included to identify the securities discussed and do not amount to a recommendation to buy, hold or sell them. Companies may be added to or removed from any portfolio at any time.

Your decisions. You are solely responsible for your own investment decisions. Before acting on anything in this document, consider whether it is suitable for you in light of your own circumstances, and seek independent advice from a qualified financial adviser where appropriate. Oxford Market Intelligence does not manage your money and does not hold or have access to your brokerage account.

Copyright. This document is the property of Oxford Market Intelligence and its authors. It is provided to subscribers and to those they share it with for personal, non-commercial reading. It may not be reproduced, sold or republished without permission.