The Semiconductor Trade

18 November 2026

by Krasni Shrivastava

Last week we added to the semiconductor holdings in our international portfolio. This note explains why, and sets out where in the semiconductor stack we think the value from AI is pooling.

The end of the monopoly. For all of history, the human mind has held a monopoly on mental work: turning information into answers. That monopoly is ending. There are now systems that can do this work, alongside us and sometimes instead of us. Our thesis is not that AI will eliminate work, but that it will change how work is done. AI is increasingly capable of synthesising information and generating answers, but remains far less capable of exercising judgement, accepting accountability and building relationships. We believe its greatest economic value will therefore come not from replacing people, but from making them substantially more productive.

Total Addressable Market. Keen investors have heard the market-sizing maths before: the world pays its knowledge workers an estimated US$35–50 trillion a year. The more useful question is what the build-out must earn to pay for itself.

The four largest hyperscalers, Microsoft, Alphabet, Amazon and Meta, spent about US$250 billion on capital investment in 2024 and about US$410 billion in 2025, and will spend about US$730 billion this year and more than US$1 trillion next. To pay for itself, that capital must earn back its cost over its useful life, about five and a half years for the chips and fifteen for the buildings, plus a return of 12% a year, comfortably above what investors require even of specialist data-centre companies. Repaying the cost and the return together comes to about 21% of the amount spent, every year. If AI businesses keep 30–50% of their revenue as cash, the AI revenue needed depends above all on how long the spending continues:

What we count	Cash returns needed a year	AI revenue needed a year
Spending in 2026 and 2027 only	about US$370 billion	US$0.75–1.25 trillion
Everything built from 2024 to 2027	about US$510 billion	US$1.0–1.7 trillion
Spending held at about US$1 trillion a year to 2030	about US$1.2 trillion	US$2.4–4.1 trillion
Spending held at about US$1 trillion a year indefinitely	about US$1.7 trillion	US$3.3–5.5 trillion

The best-known estimates point the same way. Sequoia's David Cahn, who has run this arithmetic since 2023, reckons that this year's US$1.5 trillion of AI infrastructure spending across all buyers, of which the hyperscalers above account for about half, needs roughly US$3 trillion of revenue over the life of that equipment to pay for itself; Bain estimates that AI will need about US$6 trillion of revenue a year by 2031 to fund a US$5–6.5 trillion build-out to 2030. Against that, Aswath Damodaran, the NYU valuation professor, puts today's revenue from AI products and

services at roughly US$250 billion. If the build-out continues at today's pace, AI revenue has to grow ten times or more by around 2030. That is a demanding bet. To judge it, we need to size the prize: the economic value AI could create.

McKinsey estimates that AI could create about US$2.5 trillion a year of economic value from work that already exists, equivalent to about 5–7% of what the world pays its knowledge workers. Goldman Sachs estimates that generative AI could lift US productivity by 15% and unlock revenue worth about US$8 trillion in today's money; and Bank of America's base case has AI spending earning about US$200 billion more than its cost of capital between now and 2030. The prize, in other words, is not out of proportion with the build-out. On these estimates, the value AI can create is broadly the size of what the spending needs to earn, provided enough of it can be charged for.

The size of the prize is only half the question; the other half is who captures it. Competition pushes returns down towards the cost of capital and passes the surplus to customers; only businesses that are protected from it, by scale, switching costs, intellectual property or long lead times, earn more than that for long.

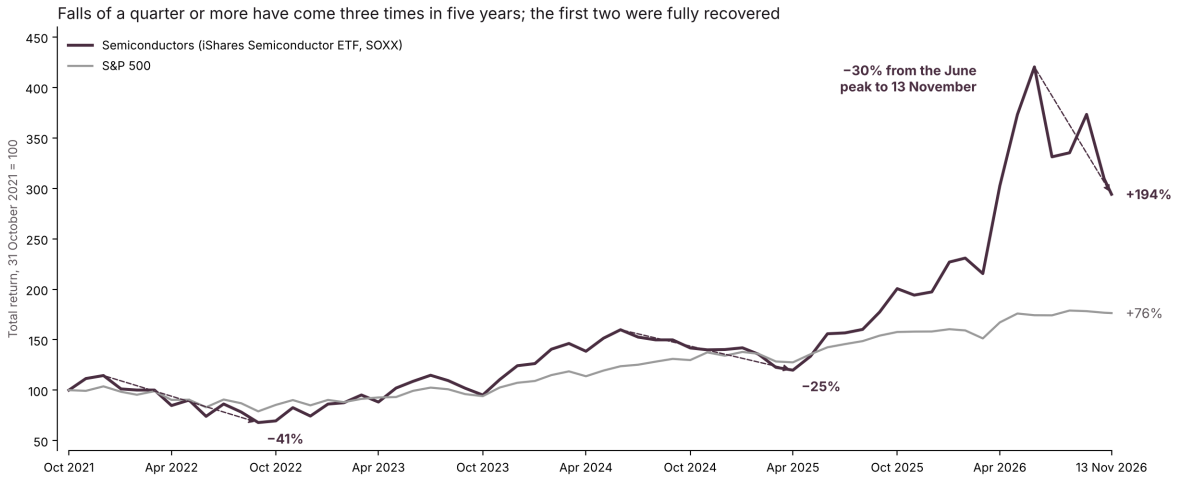
> *"Take opportunities as the price allows, and sit out when it does not.*

In AI, the price of a given capability is collapsing. A model as capable as 2021's GPT-3 cost about US$60 per million tokens then and about US$0.06 now, as more efficient chips and models cut the cost of serving it, and competition passed the savings on to customers. The newest frontier models still cost far more, and so far, capability has been the price-setter. The scarcity today lies further down: in leading-edge manufacturing, lithography, high-bandwidth memory and power. Bottlenecks migrate as capacity is added, so we judge where profits are pooling by who can raise prices, and who already earn more than their cost of capital, rather than promising to. The personal computer made fortunes for both kinds of bottleneck, Intel in hardware and Microsoft in software (remember 'Wintel'?), rather than for the companies that assembled the machines. The winners in that cycle were not those who owned everything, but those who bought the bottlenecks when they were cheap and sold when the scarcity faded. Today the hardware bottlenecks are scarce; over time, value may migrate to the software 'harness' that turns models into working products. We intend to do the same: to take opportunities as the price allows, and to sit out when it does not.

Where the volatility comes from. On 5 November the US semiconductor index fell 12% in a single day, its worst since March 2020. We think that is largely the nature of a nascent technology: the range of plausible futures is very wide. AI might lift productivity across the economy within two years, or fifteen. The returns might accrue to the chipmakers, the cloud providers, the model builders or the businesses that use the tools. Demand for computing power might keep compounding, or plateau once the first wave of buildings is full.

The market cannot price all of those futures, so it prices their average, and when the range is this wide, every new piece of information can move that average a long way, taking the share price with it. The swings are not a sign that the technology is failing, but they do demand that investors understand what they own.

What has happened in the past month. Several small pieces of information have compounded. Lam Research claimed its shipments were being held back by shortages of parts and factory space, which the market read as a peak in orders. Intel, Apple and Qualcomm each blamed the rising cost of memory chips for weaker forecasts. And investors began to ask how the build-out is being paid for. Three of the four largest hyperscalers now spend more on it than their operations generate, and some chipmakers are helping to fund their own customers, which we examine below. The strain showed most clearly at CoreWeave, the largest of the neoclouds, the newer cloud companies that rent out computing for AI alone, which has borrowed heavily to build data centres. Its shares fell by more than half between May and early November, and by 6 November its bonds were yielding more than 16% and the cost of insuring its debt against default had risen to about 13% a year, levels that signal lenders doubt a borrower can keep refinancing. CoreWeave is now seeking loans secured on its contracts with Meta at roughly half that rate: lenders will fund the build-out cheaply when a hyperscaler is the ultimate payer, but not when it is a neocloud alone. Then, in the past week, NVIDIA, the largest company in the group, lost its founder and chief executive, a departure we return to later in this note.



Total return in US dollars of the iShares Semiconductor ETF (SOXX) and the S&P 500, month-end values to 31 October 2026 and the latest close on 13 November 2026, rebased to 100 at 31 October 2021. The final point includes the 5 November sell-off; from 31 October the S&P 500 is extended using the SPDR S&P 500 ETF's daily total return. Source: Morningstar; Royston Capital analysis.



Price, expected earnings and free cash flow per unit of the iShares Semiconductor ETF (SOXX), month-end to 30 September 2026 and then 13 November 2026, November 2021 = 100. Expected earnings are derived from Morningstar's forward price-to-earnings ratio for the fund's holdings, based on analysts' estimates, and free cash flow from its trailing price-to-free-cash-flow ratio. Morningstar recalculates these ratios mainly when it processes each month's holdings, so the stale 31 October ratios are left out; the 13 November point uses the ratios refreshed after the October holdings update. Source: Morningstar; Royston Capital analysis.

In many of the names, the sell-off was deserved. By June, prices across the sector seemed to be focused on today's earnings, rather than *sustainable* earnings, and the fall since has brought them closer to what we believe is a risk-weighted fair value. The worries that drove the sell-off are real, and we take each of them seriously, which is why we have been selective and taken pains to understand the issues.

Why we believe the building continues. We weigh the narrative being told all the way up the supply chain.

- **TSMC**, which makes nearly every advanced chip, plans its capacity with its customers and its customers' customers, and now sends its own teams to inspect data-centre construction and power supply in the United States; it describes the feedback as "so far so good". It does not allow customers to resell capacity they have booked, and expects to spend significantly more on new capacity from 2026 to 2028 than in the previous three years.
- **Applied Materials**, asked directly on 12 November whether any customer had sought to delay tool deliveries, said no: customers are asking it to deliver faster. Its largest customers give it rolling eight-quarter forecasts, and more than ten major fab (chip factory) projects have been added to its tracker in each of the last three quarters, many scheduled to open in 2029 and 2030.
- **ASML**, whose lithography machines, which print circuit patterns onto silicon wafers, must be ordered more than a year ahead, is taking orders for 2028, backed by down payments it describes as large enough to hurt, and customers are asking to place orders for 2029 and 2030. ASML says it checks that demand against customers' construction plans, and expects its own planned output to fall short of demand in 2027 and 2028.

The orders for 2029 and 2030 are requests that ASML has not yet accepted, because it wants to gauge total demand first, and of course, it is in the interest of the companies quoted here to sound confident. But commitments backed by cash and construction are harder evidence than forecasts.

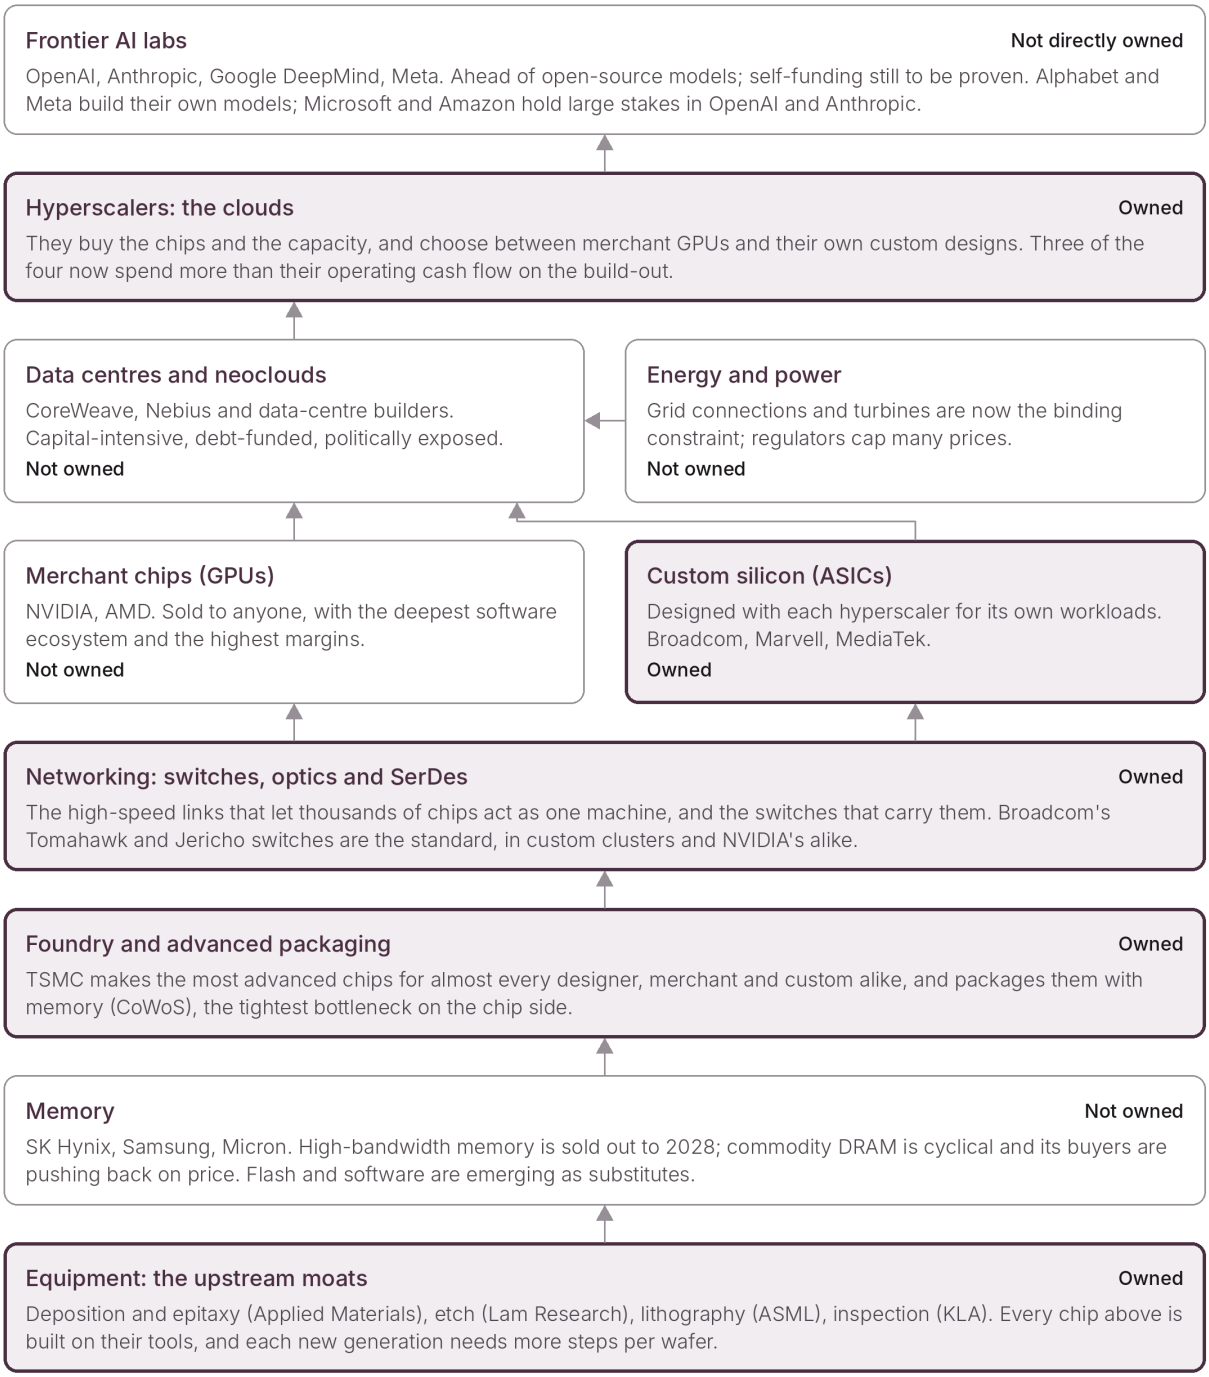
The coordination test. The supply chain is not a set of independent suppliers responding to demand. TSMC, the memory makers and the equipment makers must all commit capital at the same time, years ahead, on the strength of their confidence in future demand. NVIDIA has played a central role in coordinating that investment, aligning suppliers around its roadmap and giving them confidence in the scale and continuity of its orders. Its founder ran an unusually flat company, with some 60 people reporting to him directly and no public succession plan, so the question after his departure is how much of that coordination was built into the company and how much rested on his personal relationships and authority. We will look for the answer in suppliers' willingness to commit capacity, take down payments and invest ahead of demand.

Air pockets and Zitron. Not every chip that has been bought is switched on. Data centres take longer to build and connect to power than chips take to deliver, so the largest buyers have been taking delivery early and storing equipment until their buildings are ready. Business Insider has reported, from internal emails in August, that Amazon was staging racks of NVIDIA chips in warehouses because some sites lacked the power to run them, and that at least one of its teams had no clear forecasted path for getting some of them installed. Amazon says staging equipment near its sites is standard practice and disputes that it lacked a plan. The technology writer Ed Zitron, one of the most widely read sceptics of the AI build-out, estimates that US$200–300 billion of NVIDIA's chips are sitting in warehouses or unpowered buildings. As their buildings come online, those buyers can fill them with chips they already own rather than order new ones, so orders for chips, and for the tools that make them, could pause for several months even while the use of AI keeps growing. We would expect any pause to be temporary, but we are not complacent about it.

How quickly the stockpile can be absorbed. The answer depends on how fast new data centres are connected to power, and that is now something chipmakers track as closely as investors. TSMC has begun sending its own teams to check on US data-centre construction and power supply before committing capacity. On that measure, the build-out is accelerating. Goldman Sachs Research, drawing on project schedules, estimates that 6.4 gigawatts of new US data-centre capacity came online in 2024 and 8.5 gigawatts in 2025, with 13.6 gigawatts scheduled for 2026 and 36.3 gigawatts for 2027. Even if, as Goldman expects, only around half of next year's schedule is met on time, more capacity would come online in 2027 than this year, and more than twice as much as in 2025. This year, chips have arrived faster than power, which is why some sit in storage; next year, the buildings to house them arrive at a much faster pace. We therefore see the stockpile as a timing gap rather than a glut. The risk is that schedules slip further: Texas, the largest growth market, has paused approvals to connect new large data centres pending an audit due to report on 10 December, and the pace at which power comes online is the signal we watch most closely.

The AI stack: each layer supplies the one above, but risks are not evenly spread.

Shaded: where our international portfolio is invested.



The AI hardware stack, from the frontier labs at the top to the equipment at the bottom; shaded layers are where our international portfolio is invested. Source: company filings; Royston Capital analysis.

What we think of each layer.

- **Frontier labs:** primacy over open-source models is in question; self-funding is still to be proven.
- **Hyperscalers:** carry the "rails" on which AI models will run: corporate data, cloud infrastructure and security.
- **Data centres and neoclouds:** capital-intensive and politically exposed; lenders charge CoreWeave double-digit yields on its own credit, but about half that when a hyperscaler's contract underwrites the loan.
- **Energy and power:** grid connections are the binding constraint, and Texas has paused new data-centre connections pending an audit; utilities are signing take-or-pay contracts years ahead, but regulated returns cap who profits.
- **Merchant chips:** circular financing obscures true demand; custom silicon threatens share and margins; NVIDIA's new leadership is reportedly reining in customer financing; governance keeps us out. We see general-purpose GPUs as most exposed to step-changes in efficiency gains.
- **Custom silicon:** a tailwind of well-capitalised hyperscalers designing their own chips to reduce reliance on NVIDIA, with Google now selling its TPUs to outside customers. Custom chips built to run models efficiently are how the largest buyers capture efficiency gains for themselves.
- **Networking:** paid whichever accelerator wins, as bandwidth per chip rises with every generation.
- **Foundry and packaging:** TSMC makes nearly everyone's advanced chips, checks demand down to its customers' data centres and bars the resale of booked capacity; packaging remains the tightest bottleneck. Efficiency gains could flow through as more advanced, more profitable fabrication.
- **Memory:** high-bandwidth memory is sold out, but commodity DRAM is cyclical and being substituted.
- **Equipment:** efficiency adds steps per wafer, a tailwind; but this is the most cyclical layer, so the question is how long the building of new chip factories lasts. Customers' eight-quarter forecasts, ASML's orders for 2028 and fab projects scheduled into 2029 and 2030 are how we track it.

Who keeps the money. On reported figures, gross margins fall as you move down the stack. NVIDIA earned a gross margin of 75% in its July quarter and Broadcom about 69% in its August quarter; Broadcom's custom chips earn less than NVIDIA's GPUs, but its networking and software businesses lift the average. TSMC earned 66% in its September quarter, remarkable for a manufacturer, and Applied Materials about 50% in its October quarter. But high margins are not the same as safe ones. NVIDIA's is the very target custom silicon was built to attack, while there is currently little threat to a leading-edge fab or a decades-deep equipment portfolio. We accept lower margins in exchange for that durability.

Why we do not own NVDA (for now). Few companies have built anything as remarkable as NVIDIA. Its general-purpose GPUs can run almost any AI workload, and CUDA, the software platform it has nurtured for nearly two decades, is why most AI developers write for its chips first. That flexibility is what its customers pay for, and why it earns the margins it does. But governance is a line we do not cross, whatever the quality of the business. NVIDIA's board is now investigating its founder and chief executive after reports that he was targeted by a foreign intelligence effort; the FBI is looking into the matter, shareholders have filed a proposed class action, and Congress has asked for a briefing. On 12 November he resigned, and the board appointed the chief financial officer, Colette Kress, as interim chief executive while it searches for a successor. Since then, the US Treasury Secretary has called for an outside chief executive and a "dramatic shake-up", and wants answers from the board within ten days on whether NVIDIA complied with export controls, while Huang is reported to be weighing a return. It will probably be years before anyone outside knows what passed between NVIDIA, the US administration, China and the US intelligence agencies.

> *"Governance is a line we do not cross, whatever the quality of the business.*

The episode also changes how we read NVIDIA's numbers. The market has valued NVIDIA as a company selling out its supply while Chinese demand waited outside, an option that would pay if China reopened. NVIDIA reports that China is about 8% of its revenue and that its H200 shipments there were under 1% of data centre revenue last quarter. Yet IDC estimates that NVIDIA still held about a third of China's data centre AI chip market in the first half of

2026, and the now famous short seller Culper Research alleges that more than 20% of last year's computing revenue, about US$32 billion, reached Chinese customers through diverted chips and overseas data centres. NVIDIA disputes this and we cannot verify it. But if even part of it is true, Chinese demand is already in the numbers. To be fair to NVIDIA, there was a strategic logic to selling into China. Keeping Chinese developers on its chips and CUDA software was the surest way to stop a domestic rival growing up behind the export ban, and Beijing has since steered its own companies towards local chips for exactly that reason. The strategy was defensible; what concerns us is how it was carried out: allegations that chips reached China through intermediaries such as Megaspeed in Singapore, now under investigation by US authorities, and questions from the US Treasury over how senior in the company knowledge of them went. Export controls are also tightening: on 16 November the US restricted overseas access to advanced chips further. The Remote Access Security Act, passed by the US House of Representatives 369–22 in January and now before the Senate, would extend export controls to cloud access for the first time, and NVIDIA's own filings warn that it could have a material impact on the business.

Then there is financing. NVIDIA increasingly funds the companies that buy its chips: it now holds about US$110 billion of investments, largely in AI companies and customers, and has guaranteed up to US$105 billion of a developer's obligations on a data-centre campus being built for OpenAI. In August, Colette Kress, now interim chief executive, said AI labs it supports would be roughly a quarter of its business next year. The danger sits at the margin. NVIDIA has promised to buy any capacity CoreWeave cannot sell, up to US$6.3 billion. A buyer of last resort makes capacity no customer has signed for easier to finance, so a promise like this can help unlock borrowing well beyond its own size, spent largely on NVIDIA's own chips, in effect, a multiplier on NVIDIA's revenue. If new leadership pulls back, the neoclouds are hit first, and a slowdown would cut NVIDIA's sales, squeeze its cyclically high gross margins and cause losses on the stakes behind them.

We hold ourselves to the same test. Broadcom, which we do own, also finances purchases of its own chips, but, we think, in a sufficiently different way. Its support goes to the end customer rather than to leveraged middlemen, and it is secured on the hardware. It is, however, larger relative to Broadcom's size and concentrated in one customer rather than spread across many. Its filings show two backstops on AI-rack leases, of up to US$29 billion and about US$43 billion, and Anthropic's prospectus shows a further US$42 billion of convertible notes, undrawn by August; all are contingent rather than money already lent. Only two of its six large AI customers, Anthropic and OpenAI, use its financing, but Mizuho estimates Anthropic alone will bring Broadcom about US$42 billion of AI revenue in 2027, more than a third of its expected AI total. Take that financing away and the sales are unlikely to disappear, but would come later and on tougher terms.

We hold Broadcom because its advantages run deeper than the AI cycle. It leads the cycle-agnostic shift to custom silicon on the strength of decades of networking technology, chip-design expertise its customers cannot easily replicate and priority access to TSMC, and it earns on networking whichever chip wins. Beneath that, its software business, mostly VMware, grew 29% last quarter, with recurring revenue up 15%, at an operating margin of about 84%.

The Nortel-ephant in the room. In the late 1990s Lucent and Nortel, then the two largest makers of telecommunications equipment, lent billions of dollars to their own customers so they could keep buying. The borrowers were mostly young network operators with little revenue. At the same time, multiplexers made the technology far more efficient, so capacity grew far faster than traffic, much of the fibre laid was never lit, and the price of bandwidth collapsed. When capital markets closed in 2000 and 2001, the customers failed together: the orders disappeared, the loans were written off, and both companies lost more than 90% of their value, and Nortel eventually collapsed.

How we avoid the same trap. The lesson of 2000 is not that suppliers should never lend. It is that a boom becomes fragile once the marginal buyer depends on the seller's credit rather than its own cash flow. From that point, reported revenue measures the supplier's willingness to lend more than the customer's need for the product, and the two diverge until funding stops. We avoid the leveraged builders who occupy that position today, quantify how much of each holding's revenue its own balance sheet supports, and size accordingly. We judge the risk by four questions:

Question	What we watch	What would concern us
Is demand real?	AI revenue at end customers; how much of what has been built is in use	Commitments growing faster than use
Can the buyers pay?	Customers' cash flow, borrowing costs, supplier guarantees	Growing reliance on supplier support or costly debt
Is scarcity holding?	Prices, margins and lead times for chips and computing capacity	Prices falling while capacity keeps expanding
Are the advantages durable?	Alternative chips, portable software, suppliers' commitments	Customers or suppliers needing our holdings less

Conclusion. Until recently, turning information into understanding had been a human faculty alone. We are now looking at a very different future. Yet the AI economy can keep growing while those who invest in it still earn poor returns. Value pools where supply stays scarce and returns are already earned rather than promised: today in the chips, the factories and the tools that make them, and perhaps tomorrow in the software harnesses that turn models into products. So we own the bottlenecks, avoid the buyers who depend on borrowed money, and watch for the moment the scarcity moves. The worries of the past month are real, but none of them is beyond analysis. We cannot resolve them in advance, but we can watch them closely, and take advantage when the price overstates the uncertainty.

We close with something a dear friend once said to us: "We have a tendency to believe we're living in the most probable iteration of the future, but just because something is probable doesn't mean it comes to pass." Investors price the most probable future, but they are rewarded, or punished, by the one that actually arrives. Semiconductors will remain among the most volatile holdings in any portfolio, falling further than the market when the outlook darkens and rising faster when it clears, so how much of them to own should reflect each investor's appetite for that ride as much as our conviction in the companies.

Sources

Capital spending: Microsoft, Alphabet, Amazon and Meta; all figures include equipment acquired under finance leases, consistent with company guidance; 2024 and 2025 figures are approximate totals from the companies' annual reports (on a cash basis, excluding leases, about US$228 billion and US$376 billion), and 2026 and 2027 figures reflect company guidance from the September-quarter 2026 results, October 2026, and Applied Materials' outlook for cloud spending in 2027.

Semiconductor index: Philadelphia Semiconductor Index, close of 5 November 2026.

NVIDIA financing, all from the same Form 10-Q: page 8 (cash flow statement: purchases of equity securities US$42.4 billion, operating cash flow US$74.4 billion); page 18 (Note 10, Commitments: equity investment commitments US$25 billion; AI cloud agreements US$36 billion); page 19 (Note 10, Guarantees: SB Energy guarantees for OpenAI's leases capped at US$105 billion, total maximum exposure US$108.5 billion); page 26 (Management's Discussion and Analysis: equity investments of US$99 billion).

NVIDIA October quarter: third-quarter fiscal 2027 results, 17 November 2026 (equity securities of about US$110 billion: US$65.5 billion non-marketable and US$44.2 billion marketable; purchases of equity securities of US$58.1 billion in the nine months to 25 October 2026; accounts receivable of US$81.0 billion, up from US$38.5 billion at 25 January 2026).

CoreWeave capacity backstop: CoreWeave filing, as quoted by I/O Fund, 2026 (NVIDIA obligated to purchase CoreWeave's residual unsold capacity through 13 April 2032, about US$6.3 billion).

NVIDIA China: same Form 10-Q, page 23 (Note 13: revenue from customers headquartered in China, US$7.9 billion of US$96.2 billion in the quarter) and page 27 (H200 shipments under 1% of Data Center revenue); IDC data as reported by Reuters, 2026; Culper Research, a short seller, as reported by CNBC and Yahoo Finance, November 2026.

NVIDIA leadership: Form 8-K dated 12 November 2026, Item 5.02 and Exhibit 99.1 (resignation of Jensen Huang; appointment of Colette Kress as interim chief executive).

Remote Access Security Act: US House of Representatives roll call, 12 January 2026 (369–22); NVIDIA Form 10-Q for the quarter ended 26 July 2026, Part II, Item 1A (Risk Factors).

Broadcom: Form 10-Q for the quarter ended 2 August 2026, Condensed Consolidated Statements of Operations (revenue of US$71.1 billion for the three quarters to August, US$45.9 billion a year earlier) and Note 10, Commitments and Contingencies (lease backstop of up to US$29 billion; convertible note facility of up to US$42 billion); fiscal 2025 revenue of US$63.9 billion from the company's September 2026 overview; second rack backstop of about US$43 billion from Broadcom's Form 8-K filed 27 October 2026, Item 2.03 (maximum potential liability once the racks are deployed, undiscounted, equal to any shortfall between 72.5% of the amount owed on the lease and the value of the racks). The Broadcom filings are SEC web documents without printed page numbers, so notes are cited by number.

Semiconductor ETF: Morningstar, iShares Semiconductor ETF (SOXX) monthly total return, November 2021 to October 2026; month-end portfolio forward price-to-earnings and trailing price-to-free-cash-flow ratios, November 2021 to September 2026.

Applied Materials: fourth-quarter fiscal 2026 earnings call prepared remarks, 12 November 2026 (2027 outlook and fab projects into 2029 and 2030).

CoreWeave: bond yields and credit default swap levels as reported by 24/7 Wall St. and others, November 2026; CoreWeave third-quarter 2026 results, 10 November 2026.

Lucent and Nortel: company histories and contemporary reporting, 1999–2002.

AI market: Aswath Damodaran, "AI's Bar Mitzvah Moment: From Hype & Hope to Business Questions!", Musings on Markets, 20 August 2026 (estimate of the current AI products and services market at roughly US$250 billion).

Build-out hurdle: Royston Capital arithmetic. Assumes silicon at about 60% of AI data-centre cost (Applied Materials, fourth-quarter fiscal 2026 call), a useful life of about five and a half years for chips and fifteen years for buildings and power, a 12% required return, comfortably above the cost of capital of even specialist data-centre companies, calculated as an annual payment over each asset's life (about 26% a year of the cost of chips and 15% of buildings, 21% blended). The table applies this to about US$1.73 trillion of spending in 2026–27; about US$2.39 trillion in 2024–27; spending held at about US$1 trillion a year to 2030, counting each year's chips and buildings still in use; and the same spending continuing indefinitely. Cash margins on AI revenue of between about 30% and 50% give the upper and lower ends of each range. All figures are illustrative.

Published estimates of the revenue hurdle: David Cahn, Sequoia Capital, "AI's $600B Question", June 2024, updated in July 2026 to about US$3 trillion of revenue needed against about US$1.5 trillion of 2026 AI infrastructure spending, as reported by TechCrunch, 9 July 2026; Bain & Company, Technology Report 2026, 29 September 2026 (about US$6 trillion of annual AI revenue needed by 2031, against US$5–6.5 trillion of data-centre spending by 2030).

The bull case: McKinsey, "The AI value curve", October 2026 (about US$2.52 trillion of annual economic value from existing work; about US$7 trillion of AI infrastructure investment by 2030 needing about US$760 billion of annual profits); Goldman Sachs, as cited by J.P. Morgan Asset Management, Eye on the Market, 2026 Outlook (15% boost to US productivity; about US$8 trillion net present value of revenue unlocked by generative AI); Bank of America, report of 16 October 2026, as reported (base-case excess returns of about US$200 billion on AI spending in 2026–30).

Build-out evidence: TSMC third-quarter 2026 earnings call, 15 October 2026 (capacity planning with customers' customers, data-centre checks, no resale of capacity, 2026–28 capital spending); Applied Materials fourth-quarter fiscal 2026 earnings call, 12 November 2026, question-and-answer session (no customer delays, eight-quarter forecasts, fab projects for 2029 and 2030); ASML third-quarter 2026 earnings call, 14 October 2026 (orders for 2028, requests for 2029 and 2030, down payments, validation against construction plans, 2027–28 capacity).

Stored chips: Ed Zitron, newsletter and public posts, September–October 2026 (estimate of US$200–300 billion of NVIDIA GPUs in warehouses or unpowered data centres, derived by him from hyperscalers' construction-in-progress disclosures); Business Insider reporting on Amazon's GPU inventory, October 2026.

Data-centre capacity coming online: Goldman Sachs Research, "US Data Center Power Demand Projected to Double by 2027", 2026, drawing on Aterio development schedules (realised additions of 6.4 gigawatts in 2024 and 8.5 gigawatts in 2025; 13.6 gigawatts scheduled for 2026 and 36.3 gigawatts for 2027; Goldman expects about 60% of capacity scheduled for the next year, and about 50% of capacity scheduled two years out, to arrive on time).

Texas: ERCOT, pause on approvals to energise new data centres of 75 megawatts or more pending the verification and audit ordered by Governor Abbott on 3 August 2026; verification and community-impact reports due to the Public Utility Commission of Texas on 10 December 2026.

NVIDIA's self-financed share: NVIDIA, second-quarter fiscal 2027 earnings call, 26 August 2026, corrected transcript (Colette Kress: "we expect demand from the AI labs for which we expect to leverage our balance sheet to contribute toward roughly a quarter of our business next year").

Broadcom's financed customer: Broadcom third-quarter fiscal 2026 earnings call, 2 September 2026 (financing structures mainly for Anthropic and OpenAI; Anthropic expected to be its largest custom-chip customer in fiscal 2027; about US$115 billion of AI revenue guided for fiscal 2027); Mizuho, Vijay Rakesh, as reported by CNBC, 2026 (Broadcom AI revenue from Anthropic of about US$21 billion in 2026 and US$42 billion in 2027).

Broadcom's US$42 billion facility: Anthropic's IPO prospectus, as reported by Reuters, 1 October 2026 (agreement to lend up to US$42 billion through convertible notes, proceeds restricted to TPU lease obligations; no notes issued as of 2 August 2026).

AI prices: HPCwire, "The enterprise AI cost reckoning", 19 August 2026 (a GPT-3-level model at about US$60 per million tokens in 2021 against about US$0.06 now).

NVIDIA after Huang: Scott Bessent, Fox News, 15 November 2026, as reported by CNBC and Bloomberg (call for an outside chief executive; answers within ten days); Financial Times, 15 November 2026, as reported by Reuters (Huang weighing a return); Huang's roughly 60 direct reports, as widely reported.

Margins: reported gross margins for the latest quarter. NVIDIA, second-quarter fiscal 2027 results, 26 August 2026 (75.0%, GAAP and non-GAAP alike); Broadcom, third-quarter fiscal 2026 results, 2 September 2026 (69.1% GAAP; about 75% excluding the amortisation of acquired intangibles and stock-based pay; infrastructure software operating margin of about 84%); TSMC, third-quarter 2026 results, 15 October 2026 (66.4%, IFRS); Applied Materials, fourth-quarter fiscal 2026 results, 12 November 2026 (50.4% GAAP).

Counterparties: where a counterparty is not named in a company filing, the identification comes from press reports, including Bloomberg, the Financial Times and Reuters, and has not been confirmed by the companies concerned. Broadcom's US$42 billion note facility is identified in Anthropic's IPO prospectus; its two lease backstops are not identified by customer. NVIDIA's self-financed share is the company's own figure; Mizuho's estimate of the financed customer's revenue is as reported, compared with Broadcom's own fiscal 2027 AI revenue outlook.

Disclaimer. This information was prepared by Royston Capital, ABN 98 158 028 392 AFSL 438262. The information in this presentation is current as of 18/11/2026 unless otherwise stated. Material contained in this publication is an overview or summary only and it should not be considered a comprehensive statement on any matter or relied upon as such. This publication is confidential and intended only for current clients and prospective clients of Royston Capital and their professional institutional advisors only and is not to be provided to other persons. This publication does not constitute an offer to sell, or the solicitation of an offer to buy, any securities or any interests in investment funds advised by, or the advisory services of, Royston Capital. This publication contains general information only and does not take into account your personal objectives, financial situation or needs and so you should consider its appropriateness having regard to these factors before acting on it. No legal or tax advice is provided. Any taxation position described in this publication should be used as a guide only and is not tax advice. You should consult a registered tax agent for specific tax advice on your circumstances. Recipients should independently evaluate specific investments and trading strategies. By accepting receipt of this publication, the recipient agrees not to distribute, offer or sell this presentation or copies of it and agrees not to make use of the publication other than for its own general information purposes. The views expressed in this publication represent the opinions of the persons responsible for it as at its date, and should not be construed as guarantees of performance with respect to any investment. Royston Capital has taken reasonable care to ensure that the information contained in this publication has been obtained from reliable sources, but no representation or warranty, express or implied, is provided in relation to the accuracy, completeness or reliability of such information. Royston Capital does not undertake and is under no obligation to update or keep current the information or content contained in this publication for future events. Past performance is not necessarily indicative of future results. In addition the price and/or value of and income derived from any particular investment may vary because of changes in interest rates, foreign exchange rates, operational or financial conditions. Investors may therefore get back less than originally invested. Furthermore, these investments may not be eligible for sale in all jurisdictions or to certain categories of investors. Portfolio performance is based on a theoretical model portfolio, returns for individual investors may differ. Royston Capital, nor its associates, does not accept any liability in negligence or otherwise for any loss or damage suffered by any party resulting from reliance on this publication.

© 2026 Royston Capital Pty Ltd. All rights reserved. No part of this publication may be reproduced by any means, whether graphically, electronically, mechanically or otherwise howsoever, including without limitation photocopying and recording on magnetic tape, or included in any information store and/or retrieval system without prior written permission of Royston Capital Pty Ltd.

Royston Capital Pty Ltd · ABN 98 158 028 392 · Level 22, 459 Collins St, Melbourne VIC 3000.