VAN LANSCHOT KEMPEN

INVESTMENT MANAGEMENT

Market Musings 07/26

For Professional Investors only

The Trillion-Dollar Machine

What happens when AI becomes a trillion-dollar capital cycle?

Since 2022, the world's largest technology companies have committed themselves to an unprecedented build-out of AI infrastructure. Combined spending could approach $1.1 trillion by 2027, reshaping not only the technology sector but also credit markets, utilities and corporate funding strategies. The question for investors is not whether AI is transformative, but whether the economics ultimately justify the scale of investment. That is the question we explore here and why our view remains constructive, though selective.

What the money buys

At its core, an AI data centre is a very large, power-hungry machine. The hyperscalers, large cloud-computing companies such as Microsoft, Amazon and Google, are investing in four main layers: buildings, power, cooling and networking, and compute. Power is increasingly the binding constraint, with grid queues running up to seven years and data centres potentially accounting for 9–17% of US electricity demand by 2030. GPUs are also being supplemented by custom chips such as Google's TPU, Amazon's Trainium and Microsoft's Maia. The final bottleneck is high-bandwidth memory, supplied by only three companies, with both Meta and Microsoft citing memory costs as a driver of higher capex guidance.

A dense web of cross-holdings

Each hyperscaler has also become a direct backer of some of its biggest AI-lab customers. Publicly reported commitments and stakes include Amazon's investments in Anthropic and OpenAI; Microsoft's stake in OpenAI, alongside a large Azure commitment from OpenAI; Alphabet's commitment to Anthropic; Meta's stake in Scale AI; and Oracle's role in Stargate. These links matter because capex, funding and earnings quality increasingly interact across the AI ecosystem.

Bottlenecks and financing

Demand has outrun supply, lifting costs at Oracle and Microsoft alike. Power is the more structural constraint; chip and memory scarcity have been more cyclical, though both currently bind. Funding differs sharply: Microsoft has funded its build-out almost entirely from cash flow, with no bond deal in five years. Alphabet and Amazon have pre-funded aggressively, including euro tranches, while Meta has pushed spending off-balance-sheet via a joint venture with Blue Owl. Oracle has funded most of its programme externally through bonds convertible preferred stock, equity, and one asset sale.

> The question for investors is not whether AI is transformative, but whether the economics ultimately justify the scale of investment.

Chart 1: Big Tech bond issuance has accelerated sharply
Senior bond issuance from the four largest hyperscaler issuers, cumulative over the past three years.

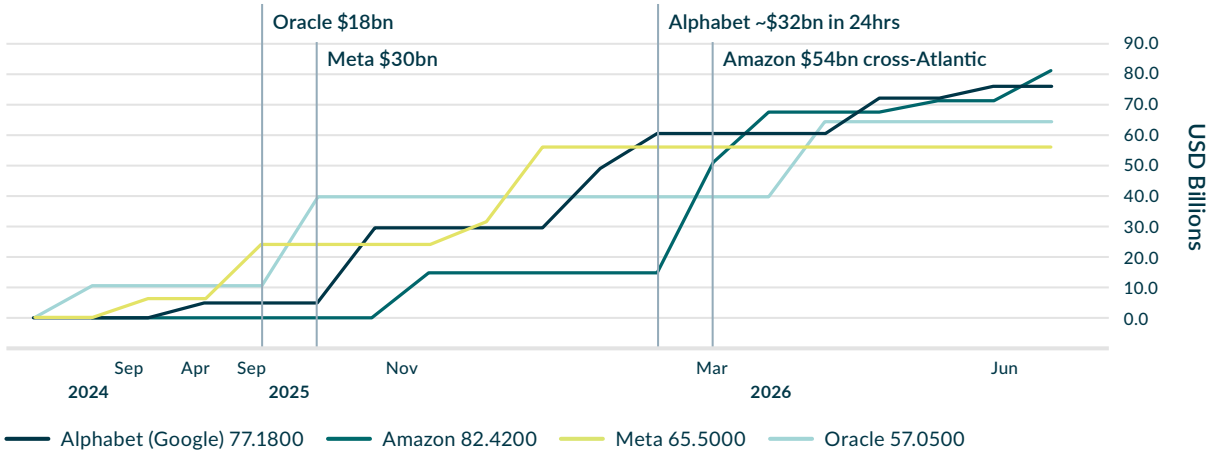

Source: Bloomberg, July 2026

Two debates worth watching

The first concerns depreciation. All five write off servers over five to six years. In late 2025, American investor and hedge fund manager Michael Burry reportedly argued, in connection with disclosed short positions in Nvidia and Palantir, that the true economic life is closer to two to three years, materially inflating reported profit; this is not independently verified here. CoreWeave, a prominent AI cloud provider, counters that older GPUs retain much value in resale once redeployed to inference work, in some cases near 95 per cent of original cost. Amazon shortened a subset of useful lives back to five years in 2025, citing faster-than-expected obsolescence.

The second concerns earnings quality. Each hyperscaler holds an AI-lab stake marked to fair value, so a new funding round can inject a large, non-cash gain into the income statement. Alphabet's 'other income' rose to $37.7bn in Q1 2026, from $11.2bn, almost entirely a markup on its Anthropic stake; Amazon's rose to $16.8bn similarly; Microsoft swung to a near-$10bn gain when OpenAI restructured. Oracle's gain was real cash: $2.7bn from selling its Ampere stake. Operating income remains the more reliable gauge.

A structurally larger capital base

Combined capex/revenue has risen from roughly 12% in 2023 to a forecast 38% this year, easing only to about 34 per cent by 2029, a materially higher base rather than a temporary spike. What matters is the split between building and land capex (long-lived, largely one-off) and silicon and memory capex (five-to-six-year life, recurring regardless of demand).

Chart 2: AI investment has lifted capital intensity to a new level
Capital expenditure as a percentage of revenue, 2012-2028, 2026-2028 are estimates.

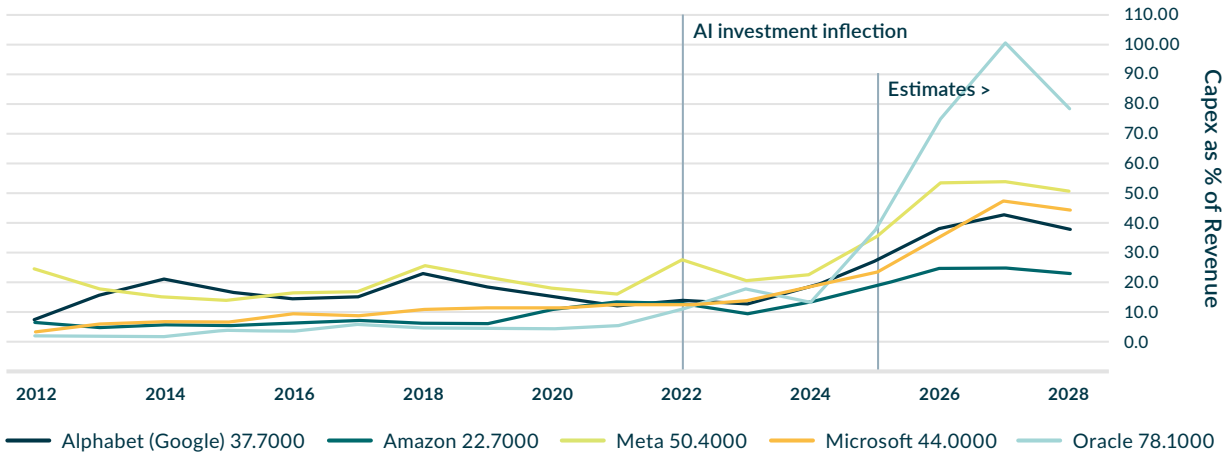

Source: Bloomberg, July 2026

The credit-rating dimension

S&P Global Ratings[1] centres its framework on free cash flow, expecting it "for every company except Microsoft, will weaken, and in some cases turn negative" but says ratings on four of the five "are not impacted... as we forecast cash flows to improve beginning in 2028," conditionally: "ratings would be at risk if we don't believe these improvements... are likely to materialise." Microsoft and Alphabet enter this period with "stronger balance-sheet headroom"; Oracle's FCF is "more constrained," its leverage forecast already breaching its 2027 downgrade threshold.

The chart below shows free cash flow as a percentage of revenue, including projections: relatively resilient at Microsoft, comfortable at Alphabet and Meta, thinner at Amazon, and deeply negative at Oracle before an assumed recovery.

Chart 3: Free cash flow pressure differs sharply across hyperscalers

Free cash flow as a percentage of revenue, historical and projected, by company.

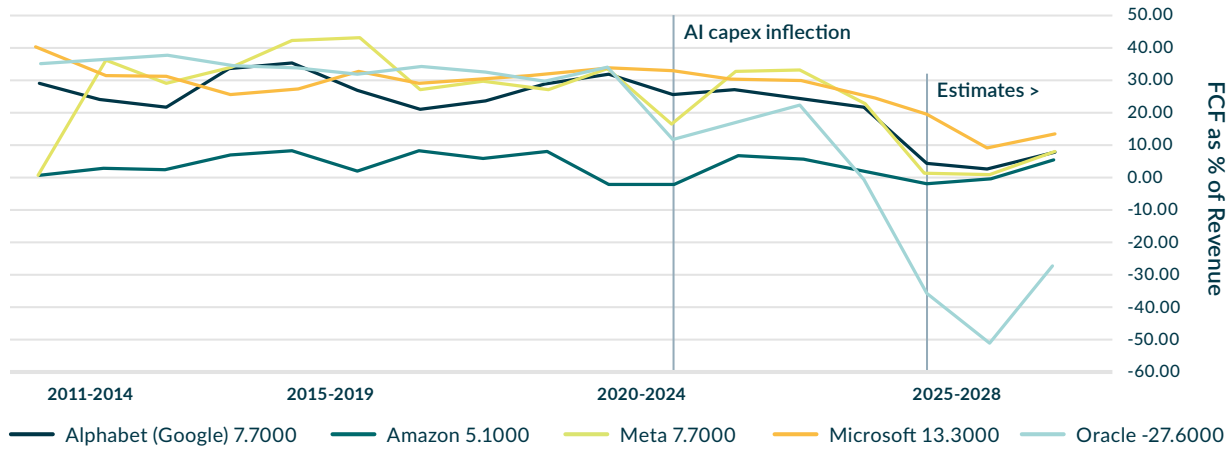

Source: Bloomberg, July 2026

Spending spreads to utilities

Before the verdict, let's switch focus to another sector doing much of the heavy lifting needed to solve the energy bottleneck: utilities. These companies have also had to increase capex to fund grid and power-generation expansion, and to add capacity to produce and deliver the energy needed by power-hungry data centres. European names already had heavier capital intensity in recent years because of green-energy-related spending; in the US, the change of pace has been more pronounced.

Chart 4: European utilities remain structurally capital-intensive

European utilities capital expenditure as a percentage of revenue, historical and projected, by company.

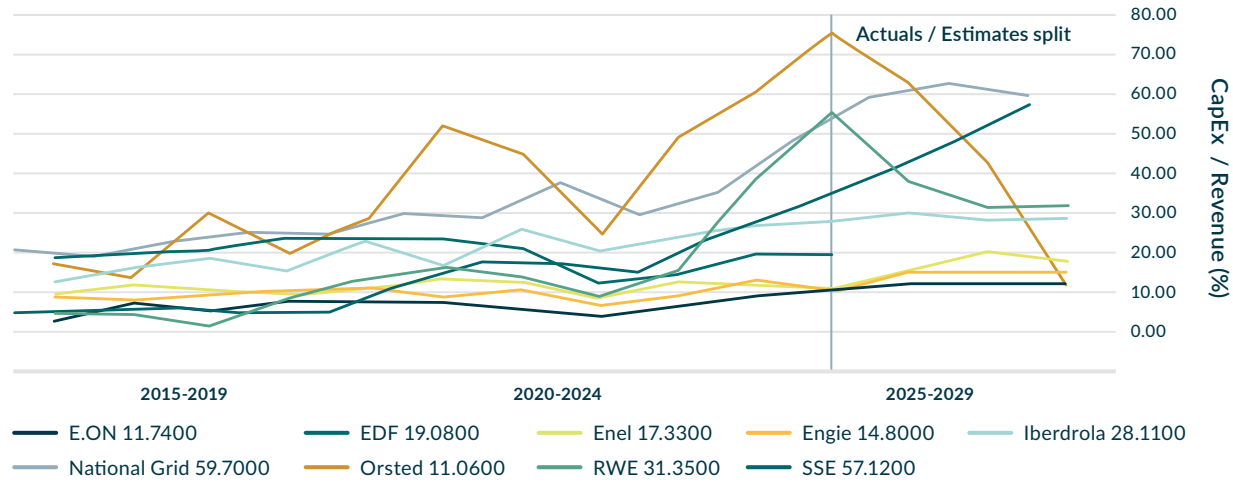

Source: Bloomberg, July 2026

[1] S&P Global Ratings Credit FAQ: Will Weakened Cash Flow Hurt Ratings On The Hyperscalers?, April 2026

US utilities capital intensity has in some cases reached levels that are even more extreme than for hyperscalers (>90% capex/revenue). New bond supply is expected not only in the USD market but also in other currencies, with some issuers already tapping the EUR market.

Chart 5: US utilities are entering a heavier capex cycle

US utilities capital expenditure as a percentage of revenue, historical and projected, by company.

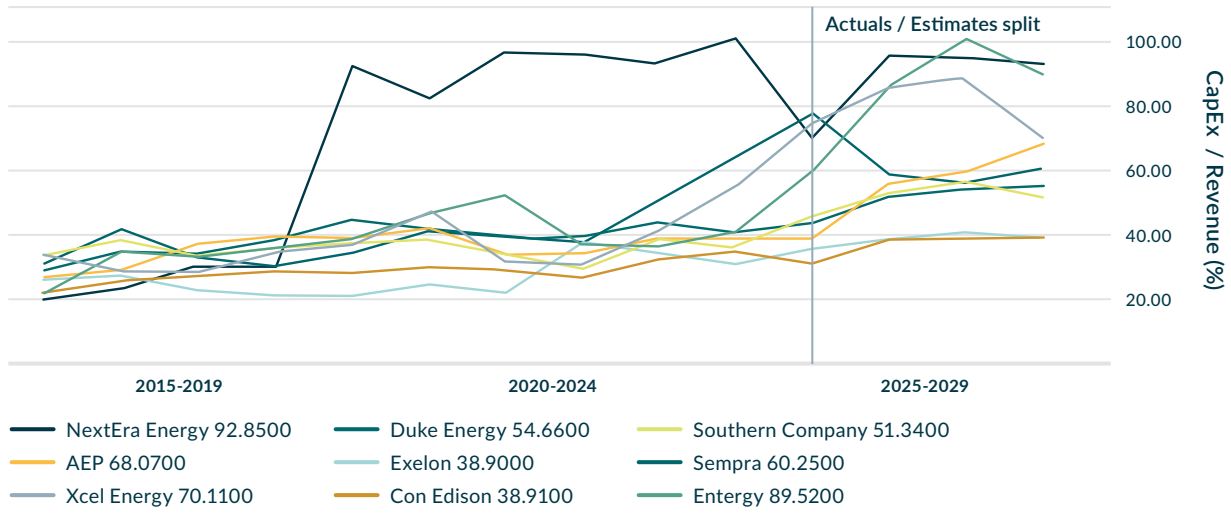

Source: Bloomberg, July 2026

Chart 6: US utilities are increasingly tapping the euro bond market

Senior bond issuance in euros by US utilities over the past five years.

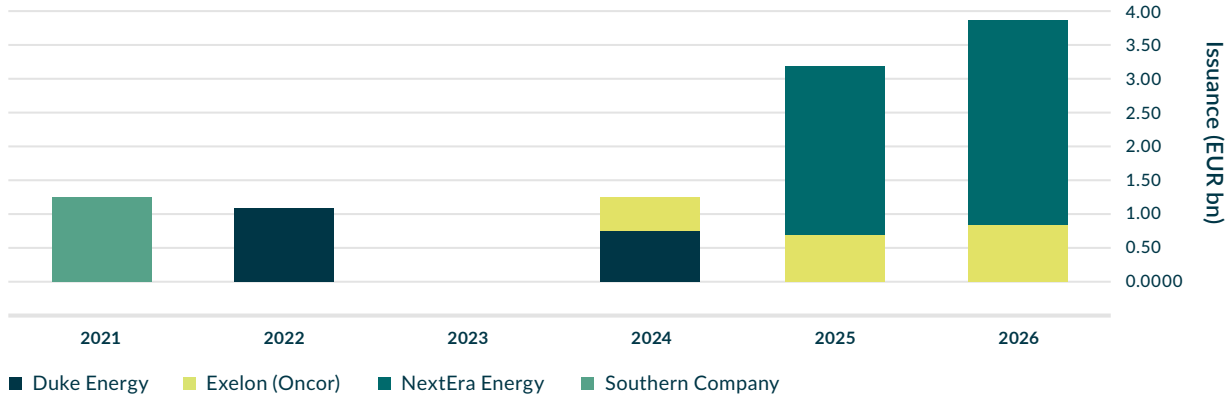

Source: Bloomberg, July 2026

Constructive, but selective

To be clear, this is not a bubble call, nor a prediction of widespread credit deterioration. The leading hyperscalers remain among the most profitable and highly rated corporate issuers globally. Yet AI has fundamentally changed the capital intensity of the sector. Balance sheets are growing, free cash flow is coming under pressure and an increasing share of reported earnings reflects valuation gains rather than cash generation.

For credit investors, the key questions from here are whether revenue growth can keep pace with investment, whether backlogs translate into cash flow and whether free cash flow recovers as expected over the coming years. Those factors will matter more than headline AI spending or market enthusiasm.

Our view therefore remains constructive on high-quality issuers across the AI ecosystem, but selective on valuation, funding structures and exposure. For credit investors, the key question is not whether AI will succeed, but how this investment cycle reshapes balance sheets, funding needs and future cash flows. The AI buildout is no longer just a technology story. It is a balance-sheet story, and increasingly a credit story.

At a glance: Nine questions, nine answers

A reader's guide to the arguments above, for anyone arriving here first.

1 What does AI capex actually buy?

Five things, in descending order of durability: land and buildings (25-40 year life, effectively permanent once built); power infrastructure (increasingly financed off-balance-sheet through utilities and power purchase agreements); cooling and networking (a rising share of cost as chips run hotter and clusters grow); and compute silicon plus the memory that feeds it, including graphics processing units, custom application-specific integrated circuits and highbandwidth memory, which is the largest and shortest-lived component of the bill.

2 Where are the main bottlenecks?

Chips, power and memory. Power is generally viewed as the more structural constraint, since grid interconnection queues cannot be shortened by capital alone, while chip and memory scarcity have historically been somewhat more cyclical, though both have already produced measurable, disclosed cost inflation in this cycle.

3 How have hyperscalers funded their AI buildout?

Mostly still from operating cash flow, though the mix is diverging. Microsoft remains the least reliant on external funding. Alphabet and Amazon are pre-funding through bonds and, at Alphabet, equity, including sizeable euro-denominated tranches. Meta is using off-balance-sheet joint ventures. Oracle is funding the large majority of its programme externally.

4 Have hyperscalers become a structurally more capital-intensive business?

Yes, in our assessment. Combined capex/revenue has roughly tripled since 2023 and is expected, on current forecasts, to settle at a materially higher base rather than revert quickly to pre-AI norms.

5 What part of capex is recurring?

The silicon and memory layer, given useful lives of five to six years that require continuous replacement regardless of whether demand keeps growing. Inference capacity and custom-chip programmes add a second, usage-linked recurring layer on top.

6 What part is more likely to normalise?

Land acquisition, initial shell construction, and the current wave of scarcity-driven forward-purchasing and interim power solutions, spending that should moderate as supply chains and grid connections mature.

7 Is 'recurring' capex understated: the depreciation debate?

An open question, and one we are monitoring rather than taking a firm view on. Some investors argue useful lives of five to six years overstate GPU longevity; industry counter-evidence on resale values suggests otherwise. Amazon's own 2025 adjustment to a subset of useful lives is the most concrete data point either way.

8 Are current earnings sustainable, or is the 'one-off' component too large?

Operating income is growing steadily and genuinely at all five companies. Reported net income, however, has been significantly affected in some recent quarters by unrealised markups on private AI-lab stakes. Most visibly at Alphabet and Amazon, a distinction worth tracking separately rather than a sign of weaker underlying performance.

9 Will there be pressure on credit ratings?

We don't expect this currently for Microsoft, Alphabet, Amazon or Meta, provided free cash flow recovers broadly on the timeline S&P has assumed from around 2028. Oracle already carries a Negative outlook from Moody's and less headroom than the rest of the group. Our credit stance across the cohort remains constructive but selective. Recent high-grade euro issuance from names such as Alphabet and Amazon may offer relative value versus ratings, while monetisation progress should be watched closely before extending that view further down the credit spectrum.

Market Musings authors

Kim Lubbers
Portfolio Manager Credits

Follow Kim on LinkedIn

Giacomo Mattioli
Portfolio Manager Credits

Follow Giacomo on LinkedIn

VAN LANSCHOT KEMPEN

Beethovenstraat 300
1088 WZ
Amsterdam
The Netherlands

vanlanschotkempen.com/investment-management

Follow us on LinkedIn