Techstrong | Special Report

PART OF Futurum MEDIA
OCTOBER 2026

AI ECONOMICS

FINANCE · INFRASTRUCTURE · RISK

The AI Money-Go-Round

Behind the record earnings, who is really paying for the AI boom?

BY ALAN SHIMEL

The same dollar, several times
How $10 billion becomes a $27 billion headline

NVIDIA on both sides of the table
Supplier, investor and customer at once

Who is holding the risk?
Why neoclouds feel a slowdown first

$677B
Microsoft commercial remaining performance obligation

–$9.1B
Alphabet free cash flow after $51.4B in quarterly capex

~$36B
Cloud-service commitments NVIDIA has made to its own customers

INSIDE THIS REPORT

Contents

ES **Executive Summary** 03
The boom is real. Following the money is harder.

01 **The Greatest Growth Story in Technology** 05

02 **Follow the Money** 09

03 **The Same Dollar, Several Times** 15

04 **Who Is Actually Paying for AI?** 19

05 **Who Is Holding the Risk?** 23

06 **When the Music Slows** 28

07 **What the Industry Needs to Tell Us** 32

SIDEBARS

A The AI Money-Go-Round, Explained 11

B Revenue, Backlog and Capex Are Not the Same Thing 18

C The Neocloud Balance Sheet 25

D Five Questions Every AI Investor Should Ask 31

ALSO

- Methodology, notes and sources 34

A TECHSTRONG SPECIAL REPORT

AUTHOR
Alan Shimel, CEO, Techstrong Group; CSO, Futurum Media

RESEARCH CUTOFF
October 28, 2026

IMAGERY
Photo illustrations are AI-generated and conceptual. They do not depict real facilities, people or transactions.

Techstrong
a Futurum company

EXECUTIVE SUMMARY

The boom is real. *Following the money is* harder.

THE OCTOBER 2026 technology earnings season offered one of the most impressive demonstrations of growth the industry has seen. Alphabet reported Google Cloud revenue more than doubling year over year. Microsoft continued expanding Azure at an extraordinary rate while accumulating hundreds of billions of dollars in commercial commitments. NVIDIA's data center business reached revenue levels that would have seemed unimaginable only a few years ago. Oracle, Amazon, Meta and the emerging generation of specialized AI cloud providers continued investing enormous sums in the infrastructure required to support artificial intelligence.

There are legitimate reasons for the optimism surrounding these results. AI is moving into production environments, inference workloads are expanding, and businesses are discovering applications that were either impractical or impossible with earlier generations of technology. The demand for computing power is substantial, and the physical infrastructure needed to satisfy it takes years to plan, finance and construct. Companies that believe AI will become a foundational technology have little choice but to make significant investments before all the future demand is visible.

The financial relationships supporting that expansion, however, deserve considerably more scrutiny than they have received. Many of the industry's largest participants operate simultaneously as suppliers, customers, investors and financing partners. Microsoft has invested heavily in OpenAI while generating substantial commercial revenue from its relationship with the company. NVIDIA sells infrastructure to AI cloud providers while also committing to purchase services from some of those providers, investing in the ecosystem and supporting infrastructure financing. Oracle uses customer prepayments, vendor financing and other arrangements to support its capacity expansion. Specialized cloud providers are building rapidly against large contractual commitments, often supported by debt and concentrated customer relationships.

These transactions are not inherently questionable. Strategic investments, customer financing and long-term purchasing arrangements have supported major technology transitions for decades. The difficulty emerges when financial activity from different stages of the same economic chain is aggregated into an apparent measure of AI demand. Equipment sales, cloud service revenue, capital expenditures, equity investments and future contractual commitments are different financial categories. Each can be accurately reported while contributing to an inflated impression of the amount of independent purchasing power entering the AI economy.

Our examination finds that the industry has substantial, growing demand from enterprises and consumers, but publicly available disclosures do not establish how much of the enormous financial activity surrounding AI infrastructure ultimately depends on independently generated customer revenue. The problem is particularly acute where strategic investors are also major customers or suppliers, and where substantial future commitments depend on businesses that have not yet demonstrated sustainable profitability.

The risks are not distributed evenly. Diversified hyperscalers have established businesses capable of generating enormous amounts of cash, while specialized infrastructure providers often depend more directly on customer concentration, favorable borrowing conditions and the economic life of expensive computing equipment. These

differences could become increasingly important if infrastructure capacity grows faster than the market's ability to use it profitably.

The evidence does not establish that the AI boom is artificial, nor does it support a definitive estimate of recycled industry capital. It does establish a financial system whose participants are becoming increasingly dependent on one another and whose headline growth figures provide an incomplete picture of underlying demand.

AI may fulfill its enormous technological promise while some of the investments supporting its expansion produce disappointing financial returns. Understanding that possibility requires looking beyond the earnings headlines and following the money.

> *Publicly available disclosures do not establish how much of the enormous financial activity surrounding AI infrastructure ultimately depends on independently generated customer revenue.*

THE SEASON IN EIGHT NUMBERS

Growth, spending *and the ties between them*

$127.6B Alphabet Q3 revenue, up 25%	**$30.7B** Google Cloud revenue, more than doubling	**+48%** Azure revenue growth, fiscal Q1 2027	**$96.2B** NVIDIA quarterly revenue; $89.0B from Data Center
$51.4B Alphabet quarterly capital expenditures	**$34.9B** Meta quarterly capex incl. finance-lease principal	**$24.1B** Microsoft FY2026 revenue from OpenAI arrangements	**$104B** CoreWeave revenue backlog, June 30, 2026

Sources: company reports and filings; see Notes and sources. Figures describe different measures and periods and should not be added together.

01

PART ONE

The Greatest *Growth Story in* Technology

October's earnings season delivered extraordinary results — and record bills for the infrastructure behind them.

A remarkable earnings season

THERE IS NO reason to diminish what the technology industry reported during October. Alphabet, Microsoft and other major technology companies delivered results that would have been considered extraordinary at almost any point in their histories, and the accelerating demand for cloud computing and AI infrastructure has become a defining feature of the current investment cycle.

Alphabet's third-quarter revenue reached approximately $127.6 billion, an increase of 25% from the previous year. Google Cloud reported $30.7 billion in revenue, more than doubling year over year, while operating income for the cloud segment reached approximately $11.1 billion.[1,2] The growth suggests that Google's investments in infrastructure, AI models and its broader cloud platform are creating substantial commercial momentum. It also provides evidence that Google Cloud is already capable of generating meaningful operating profits rather than merely promising them at some unspecified point in the future.

Microsoft's fiscal first-quarter results were similarly impressive. Revenue increased 19% to approximately $92.3 billion, with Microsoft Cloud generating $63.9 billion and Azure revenue growing 48%. Commercial remaining performance obligations reached $677 billion, giving the company considerable contractual visibility into future business.[3] Microsoft's broad enterprise relationships, established software franchises and significant cash generation make its position particularly strong. Unlike many newer AI infrastructure companies, Microsoft can support much of its expansion through the cash produced by businesses that were highly profitable long before the current AI boom.

NVIDIA's performance illustrates the extraordinary economics available to suppliers positioned at the center of a major technology investment cycle. Its most recently reported quarterly revenue reached $96.2 billion, more than double the comparable period a year earlier, with approximately $89 billion generated by its Data Center business.[4] The company's operating profitability is remarkable, reflecting both enormous demand for its products and its ability to capture a substantial portion of the economic value associated with AI infrastructure construction. NVIDIA is benefiting from a market in which customers have strong incentives to secure computing capacity before they can be certain how much they will ultimately need.

The broader market reinforces the picture. Amazon's most recently available AWS results showed revenue growing 37% to $42.2 billion,[5] demonstrating that the largest established cloud provider continues to benefit from increasing demand. Oracle has reported exceptional infrastructure cloud growth and increasingly ambitious revenue targets, while Meta continues using AI to improve the economics of its advertising business even as it invests heavily in infrastructure and model development.

Specialized AI cloud providers are participating in the same expansion. CoreWeave and Nebius have reported rapid growth as customers seek access to accelerated computing infrastructure without constructing every data center themselves. These businesses serve an important function in an industry where

EXHIBIT 1 **The AI engines are growing at very different speeds**

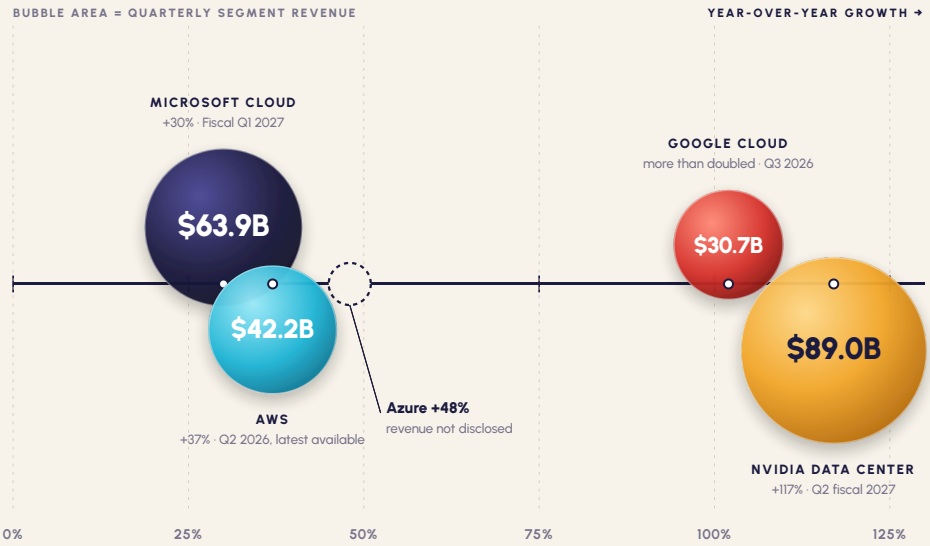

Sources: Microsoft; Amazon; Alphabet; NVIDIA. Latest reported quarter for each segment; periods differ. Notes 1,3,4,5.

demand for advanced GPUs, power and network capacity can exceed what traditional cloud providers are prepared to deliver on short notice.

Taken together, the results provide substantial evidence of an expanding market. They also demonstrate that companies with very different business models can participate in the AI investment cycle simultaneously. A chip supplier, a cloud operator, an enterprise software company and an AI model developer can all experience rapid growth as a result of the same underlying technology transition. Their financial results, however, should not be interpreted as equivalent measures of the market's economic maturity.

The cost of building ahead of demand

> *Their financial results should not be interpreted as equivalent measures of the market's economic maturity.*

The infrastructure required to support this growth is expensive, and the industry's capital expenditure figures are beginning to attract almost as much attention as its revenues. Alphabet reported approximately $51.4 billion in capital expenditures for the quarter, exceeding operating cash flow and producing negative free cash flow of roughly $9.1 billion.[1] Meta's quarterly capital expenditures, including finance lease principal payments, reached approximately $34.9 billion, while its reported free cash flow turned negative.[6]

Microsoft offers a useful contrast. It generated approximately $58 billion in operating cash flow and spent roughly $39.4 billion in cash on property and equipment, leaving substantial positive free cash flow even after its infrastructure investments.[3] The company's ability to generate cash from established businesses provides considerable flexibility, allowing it to pursue long-term investments without depending entirely on the immediate profitability of newly installed AI capacity.

These figures should be interpreted in context. Capital expenditures are not ordinary operating expenses, and negative free cash flow during a major investment cycle does not automatically indicate financial distress. Data centers, electrical systems, networking infrastructure and computing equipment are intended to support future economic activity over multiple years. Some assets will remain useful through several generations of AI hardware, while others will require more frequent replacement. The relevant financial question is whether those assets generate sufficient returns over their economic lives to compensate for their purchase, operation and financing.

There is a credible argument that the industry must build aggressively because the alternative could be worse. Data centers take years to plan and construct, grid connections can face substantial delays, and the most advanced AI systems require tightly integrated computing, networking and cooling environments. If demand continues increasing at anything approaching its recent pace, companies that fail to invest today may find themselves unable to meet customer requirements tomorrow.

The rapid improvement in model efficiency also complicates predictions of future demand. More efficient inference can reduce the cost of individual tasks,

EXHIBIT 2 **Cash in, cash out: the cost of building ahead of demand**

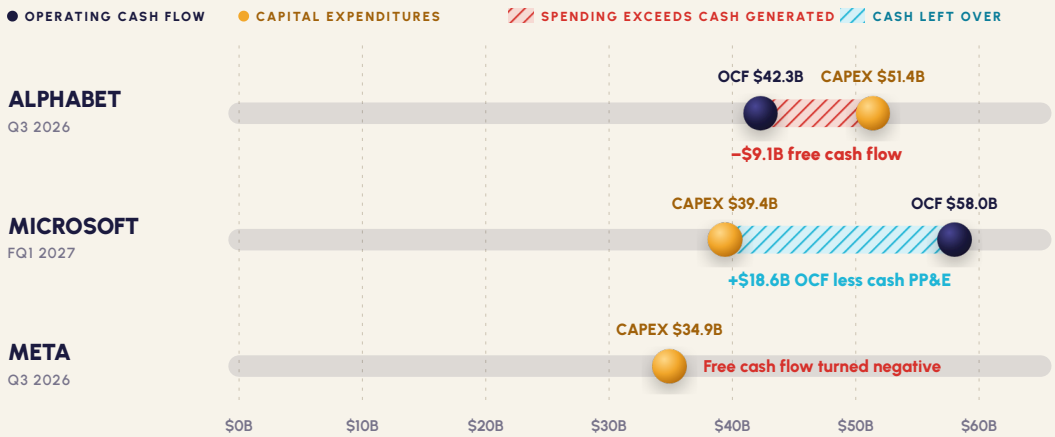

Alphabet operating cash flow implied from reported capex and free cash flow. Meta capex includes finance-lease principal. Notes 1,3,6.

but it can also make AI practical for applications that would previously have been uneconomic. Lower prices may increase total consumption as developers incorporate AI into more products, services and business processes. This is an important part of the investment thesis and one reason that judging future infrastructure requirements solely from today's usage would be shortsighted.

The case for investment ahead of demand is therefore substantial. What the earnings releases cannot establish by themselves is whether the amount being invested, the financing structures supporting those investments and the future revenues expected from them will ultimately produce attractive returns across the entire ecosystem.

Several different businesses, one investment cycle

The financial differences among AI companies become clearer when we distinguish the businesses involved. NVIDIA sells the equipment and systems needed to construct AI infrastructure. Hyperscalers operate large cloud platforms, combining traditional workloads with AI training, inference and related services. Specialized AI clouds concentrate more directly on accelerated computing infrastructure. Model developers purchase large amounts of compute to train and operate their systems, while application providers attempt to convert those capabilities into products customers will pay to use.

Each layer can recognize revenue as transactions occur, and each can report growth from its own position in the supply chain. A model developer's decision to expand its computing requirements can create business for a cloud provider, which may purchase additional NVIDIA equipment and networking infrastructure. The economic activity is legitimate, but the resulting revenues are connected to a common underlying source of demand.

The distinction becomes even more important when strategic financing is introduced. A technology company may invest in a model developer that subsequently purchases its cloud services. An infrastructure supplier may invest in a cloud operator that buys its equipment. A large customer may provide a prepayment that helps a provider finance the infrastructure required to deliver future services.

Such arrangements can make considerable commercial sense, but they also blur the lines between independently generated demand and activity supported by financing within the ecosystem. When industry forecasts aggregate these different financial measures, the risk is that the resulting numbers convey more certainty about the size of the ultimate market than the underlying evidence supports.

That is the part of the AI financial story the earnings celebrations tend to leave out.

EXHIBIT 3 **Several different businesses, one investment cycle**

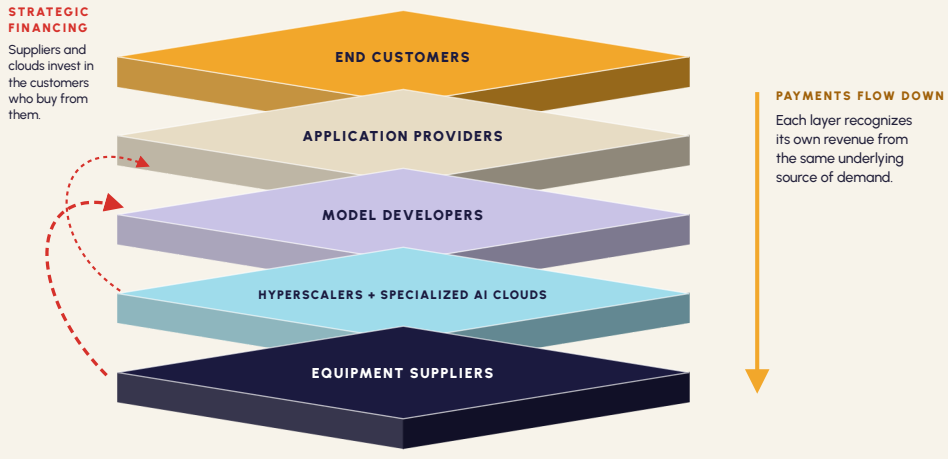

Arrows are schematic. Layers can be one company playing several roles.

Techstrong schematic of the AI value chain described in this report.

02

PART TWO

Follow *the* Money

Suppliers, customers, investors and lenders are increasingly the same companies. That changes how to read their results.

The emergence of a financially interconnected ecosystem

THE TECHNOLOGY INDUSTRY has a long history of using strategic investments and commercial partnerships to accelerate adoption. Hardware manufacturers have financed equipment purchases, software companies have invested in developers, and platform providers have offered incentives to companies building applications for their ecosystems. These arrangements can reduce the financial barriers associated with adopting new technologies and create markets that might otherwise develop more slowly.

AI is taking these practices to an entirely different scale. Building a competitive frontier model requires substantial computing resources, while constructing the infrastructure that provides those resources requires enormous investments in equipment, facilities, electrical capacity and networking. The demand for capital is so large that traditional boundaries between supplier, customer and financial sponsor are increasingly difficult to maintain.

For the companies involved, these relationships can provide important strategic benefits. Equipment suppliers secure customers and influence infrastructure architecture. Cloud providers attract major workloads. Model developers gain access to capital and computing resources. Investors participate in potential future growth, while lenders gain confidence from long-term contractual commitments.

The arrangements can also create a network in which the financial performance of one participant affects the apparent strength of several others. The difficulty for outside observers is determining how much of that strength originates with independent customers purchasing valuable products and services, and how much depends on the continued availability of financing among participants.

Microsoft's relationship with OpenAI provides an unusually clear example.

Microsoft and OpenAI: Investment meets commercial revenue

Microsoft's fiscal 2026 annual report disclosed approximately $11.9 billion in funding provided against $13 billion in investment commitments related to OpenAI as of June 30, 2026. It also disclosed approximately $24.1 billion in revenue from commercial arrangements with OpenAI during the fiscal year, including revenue-sharing payments, along with $6 billion in receivables from OpenAI at year-end. Microsoft reported holding an approximately 25% interest in OpenAI on an as-converted basis under its relevant investment arrangement.[7]

The amounts are significant, but their meaning requires careful consideration. Microsoft's investment represents capital provided in exchange for an economic interest. Its commercial revenue arises from contractual arrangements that include different services and payment obligations. Receivables represent amounts owed under applicable accounting rules. These figures are related to the same strategic relationship, but they are not interchangeable, and the financial statements do not establish which specific investment dollars funded particular commercial payments.

It would therefore be inappropriate to subtract the investment from Microsoft's revenue and describe the remainder as independent demand. That calculation would impose a relationship on the figures that the disclosures do not support. Equally, there is no basis for dismissing the commercial revenue as fictitious simply because Microsoft has invested in OpenAI. The companies have a substantial business relationship involving real products, services and financial obligations.

The importance of the disclosure lies elsewhere. Microsoft is simultaneously a major strategic investor and a recipient of significant commercial revenue from a company whose operations require extraordinary computing resources. OpenAI's ability to sustain and expand its purchasing depends ultimately on its own commercial success, access to financing and strategic relationships. Microsoft's investment position therefore creates an economic relationship with the customer that extends beyond an ordinary arms-length cloud contract.

The same relationship can be described through several legitimate financial measures: equity investment, cloud-related commercial revenue, revenue-sharing payments, receivables and future business commitments. Each provides useful information, but none independently establishes the amount of revenue ultimately generated by customers outside the relationship.

> *Each measure provides useful information, but none independently establishes the revenue ultimately generated by customers outside the relationship.*

Microsoft has substantial financial resources and a diversified business, making it well positioned to manage such exposure. Its commercial relationship with OpenAI may become one of the most successful strategic partnerships in technology history. Nevertheless, the

SIDEBAR A

The AI Money-Go-Round, Explained

When suppliers, customers and investors become the same companies

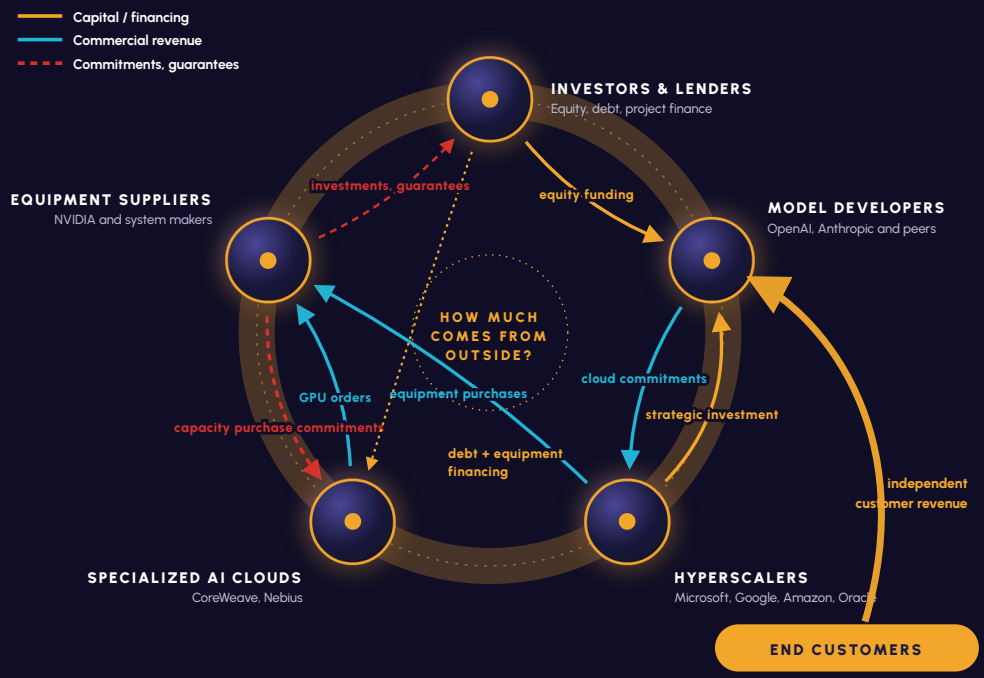

The financial relationships supporting today's AI infrastructure boom resemble arrangements that have existed in technology for decades, but the scale and complexity have become exceptional. Major technology companies are investing in model developers that purchase their cloud services, financing infrastructure providers that buy their equipment and contracting to purchase capacity from companies that also depend on them as suppliers.

There are practical reasons for these arrangements. AI model developers need enormous amounts of computing power before their businesses necessarily generate enough revenue to fund it. Infrastructure providers need substantial upfront capital before facilities become operational, while equipment suppliers benefit from customers having the financing needed to deploy their systems. Strategic investments and long-term purchasing commitments can help bridge those gaps, creating infrastructure that might otherwise take much longer to develop.

The financial complication arises when the same underlying capital supports several rounds of commercial activity. An investor provides funding to a model developer, which signs a cloud computing agreement, enabling the cloud provider to purchase equipment and finance additional capacity. The original investment may be reflected in a financing announcement, while the cloud agreement appears in reported backlog and the equipment purchase becomes revenue for the supplier. Each transaction represents a different economic event, yet public commentary may treat the amounts as separate evidence of independent AI demand.

Reciprocal arrangements can make the relationships even harder to interpret. A supplier may invest in a cloud provider, sell the provider equipment and separately agree to purchase capacity from it. Such agreements may create commercial efficiencies and support real customer needs, but they also make the financial prospects of the companies more interdependent.

The AI money-go-round is not a claim that these transactions are fake or improperly recorded. It describes a financial network in which capital, contractual commitments and revenue circulate among companies that increasingly serve several roles at once. The key issue is whether enough independent demand eventually develops to support the investments and obligations created along the way.

A successful outcome could justify much of the infrastructure being constructed and reward the companies that financed it. A disappointing outcome could leave some participants with underutilized assets and obligations that are harder to support than anticipated. Understanding the difference requires tracing the sources of money rather than relying on the combined size of financial announcements.

financial disclosures illustrate why analysts should distinguish revenue earned from strategically supported counterparties from independently funded end-customer demand when evaluating the broader AI market.

That distinction is especially important when other companies are establishing similarly complex relationships.

NVIDIA expands beyond selling the machinery

NVIDIA's position in the AI economy is extraordinary. Its products provide a foundation for much of the industry's most demanding computing infrastructure, giving it considerable influence over system architecture, networking, software and the economics of AI factories. Its financial success is also among the strongest evidence that the infrastructure investment cycle is already producing enormous returns for at least some participants.

Yet NVIDIA's role is expanding beyond that of an equipment supplier. Its July 2026 regulatory filing described approximately $36 billion in future cloud-service commitments involving AI infrastructure partners that purchase NVIDIA equipment.[8] Under these arrangements, NVIDIA can contract to use computing capacity supplied by companies whose infrastructure depends on NVIDIA systems, with provisions that may allow third-party customers to use capacity under specified circumstances.

There are practical reasons for these agreements. NVIDIA requires substantial computational capacity for its own research, software development and model-related activities. Its infrastructure customers benefit from contractual demand that can make new projects easier to finance. Some arrangements may provide mechanisms for directing unused capacity toward outside customers, potentially improving overall utilization.

From a financial perspective, however, the arrangements put NVIDIA on both sides of related commercial activity. It may recognize revenue from equipment sales to an AI cloud provider while committing to purchase services from that provider in the future. The provider can use the expected revenue from those services as part of its business and financing plans, creating a relationship in which NVIDIA's purchasing commitments support infrastructure built with NVIDIA equipment.

NVIDIA's filing also identified substantial planned equity investment commitments across AI-related businesses and contingent guarantees associated with infrastructure development.[9] These arrangements do not constitute realized losses or immediate cash expenditures, and their contractual details determine how much exposure ultimately materializes. They do establish that NVIDIA is assuming a wider financial role across the AI infrastructure ecosystem than would be apparent from its hardware sales alone.

For NVIDIA, the arrangements may improve the resilience and scale of the market for its products. Supporting infrastructure development can help customers bring systems online and make additional computing capacity available to the broader industry. There is also a potential advantage in maintaining

EXHIBIT 4 **Microsoft and OpenAI: one relationship, four measures**

CIRCLE AREA PROPORTIONAL TO DOLLAR VALUE · DO NOT NET OR ADD

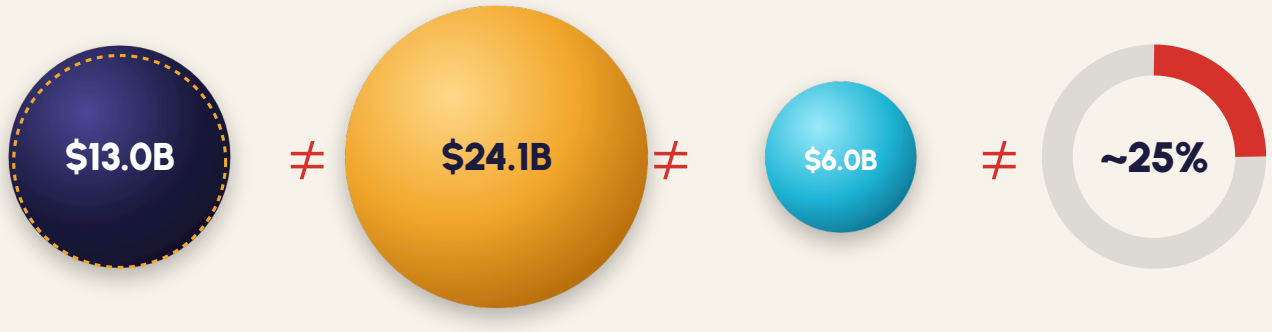

FUNDING COMMITTED
$11.9B funded by June 30, 2026

COMMERCIAL REVENUE
FY2026, incl. revenue sharing

RECEIVABLES
Owed by OpenAI at year-end

ECONOMIC INTEREST
As-converted, equity method

Source: Microsoft fiscal 2026 Form 10-K. Values as of or for the year ended June 30, 2026. Note 7.

access to cloud resources and participating in the growth of businesses that operate NVIDIA-powered infrastructure.

The financial risk is that the relationships can become mutually dependent. A customer whose expansion depends partly on financing commitments or capacity purchases from NVIDIA may be more exposed if NVIDIA changes its strategy, if compute prices decline or if independent customer demand develops more slowly than expected. NVIDIA could face risks through several channels, including equipment demand, investment value, purchasing obligations and contingent guarantees.

> *NVIDIA is becoming not only a supplier to the AI infrastructure economy, but also one of the companies helping shape and finance that economy.*

These are not the same exposures, and they should never be combined into a single unsupported liability figure. But examining them individually reveals an important transformation: NVIDIA is becoming not only a supplier to the AI infrastructure economy, but also one of the companies helping shape and finance that economy.

Oracle: Contracted demand and the cost of delivery

Oracle's expanding role in AI infrastructure offers another example of how commercial demand and financial engineering are becoming increasingly connected. The company has announced

EXHIBIT 5 **NVIDIA on both sides of the table**

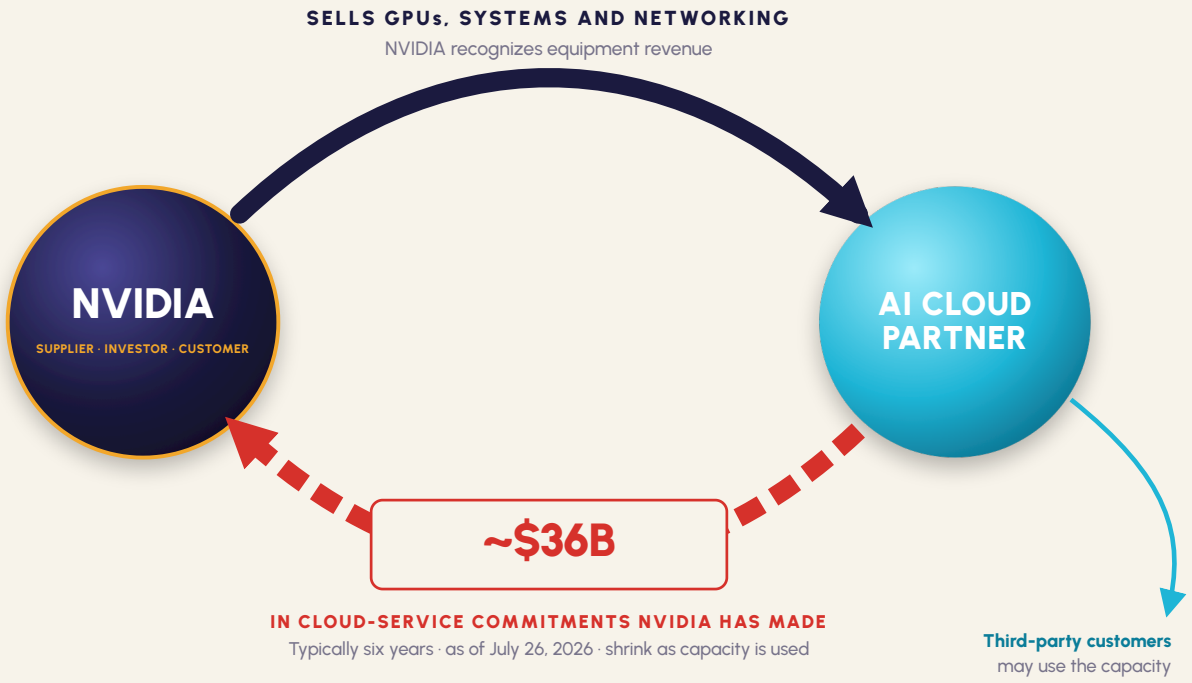

FOUR EXPOSURE CHANNELS, KEPT SEPARATE

01 EQUIPMENT DEMAND	02 INVESTMENT VALUE	03 PURCHASE OBLIGATIONS	04 CONTINGENT GUARANTEES
Orders from customers who rely on NVIDIA support	Equity stakes across AI businesses	Cloud capacity NVIDIA has agreed to buy	Payable only if specified conditions occur

Sources: NVIDIA Form 10-Q for the quarter ended July 26, 2026; CNBC. The channels are different exposures and should not be summed. Notes 8,9.

substantial customer commitments and ambitious long-term revenue expectations, placing cloud infrastructure at the center of its growth strategy. Its October 28 Investor Day provided additional detail on the financial arrangements intended to support that expansion.[10]

Oracle described the use of vendor financing, customer-provided hardware and customer prepayments to reduce the amount of upfront cash required for certain infrastructure projects. These structures can be commercially sensible because they align the funding of capacity with customers that have committed to using it. A customer seeking long-term access to computing resources may prefer to help finance dedicated infrastructure rather than wait for a cloud provider to fund and construct that capacity independently.

The accounting and financial implications depend on the details of each arrangement. Equipment purchased directly by a cloud provider can create different balance-sheet and cash-flow effects from equipment supplied by a customer. Vendor financing can reduce immediate cash expenditures while creating future obligations. Prepayments can provide liquidity before services are delivered, but they also create commitments that the provider must satisfy.

This makes headline comparisons of capital expenditures potentially misleading. Two providers can develop similar amounts of capacity while assuming different asset ownership, debt, lease and payment obligations. The provider with the lower immediate cash outlay is not necessarily making a smaller economic commitment. Nor does a large contracted backlog establish the profitability of the infrastructure required to fulfill it.

Oracle has an established enterprise business and substantial customer relationships that differentiate it from newer infrastructure companies. Its opportunity is significant, particularly as large AI customers seek providers capable of delivering extensive computing capacity. The company's long-term projections nevertheless depend on converting commitments into operating revenue and generating sufficient cash after the costs of financing, operating and maintaining the required infrastructure.

The relevant measure of success will be the returns produced by those commitments, not merely the total value announced when agreements are signed.

The neoclouds and the changing financing model

Specialized AI cloud providers occupy an important position between hardware suppliers and the model developers or enterprises that require computing capacity. CoreWeave and Nebius have expanded rapidly because they offer access to advanced infrastructure without requiring every customer to build and operate its own facilities. Their growth demonstrates that there is a market for specialized computing services beyond the traditional hyperscaler model.

Unlike the largest cloud companies, however, specialized providers generally lack equally large portfolios of established businesses producing cash independently of AI infrastructure. Their expansion can therefore depend more directly on borrowing, equipment financing, equity issuance, customer commitments and the continuing commercial attractiveness of the hardware they operate.

CoreWeave illustrates both the opportunity and the risk. Its substantial revenue growth and long-term customer commitments provide evidence of strong demand, while its financing costs, customer concentration and infrastructure spending demonstrate the capital intensity of the business. A large backlog can make it easier to finance capacity, but the quality of the backlog depends on the financial condition of the customers, the terms of the agreements and the provider's ability to deliver services profitably.

When infrastructure providers depend on a small number of exceptionally large customers, the financial condition of those customers becomes central to understanding the provider's risk. A contract with a well-funded AI developer can support significant investment, but the developer's funding may itself depend on strategic investments and future revenue expectations. The resulting chain creates several businesses whose commercial prospects are connected to the same ultimate economic assumptions.

This does not make the neocloud model inherently unstable. Specialized operators may develop important advantages in infrastructure deployment, performance optimization, cost management and customer responsiveness. Some may become exceptionally valuable businesses. The distinction is that their capital structures and customer concentration can leave them more sensitive to changes in demand, financing conditions and equipment economics than the largest diversified technology companies.

The industry increasingly depends on these specialized providers to expand capacity. Understanding how they are financed has become an essential part of understanding the economics of AI itself.

PART THREE

The Same *Dollar,* Several Times

> Financial announcements move through the AI supply chain. Adding them all together can make one dollar look like three.

What the headlines can conceal

THE POSSIBILITY OF multiple counting becomes clearer when we consider how financial announcements move through the AI supply chain. Imagine a hypothetical model developer that raises $10 billion from investors and commits that funding to future computing services. The cloud provider receiving the commitment purchases $7 billion in GPUs, networking systems and related equipment to construct the necessary infrastructure. Equipment suppliers recognize revenue as they deliver their products, while the cloud provider records revenue as it fulfills its contract.

A market observer could describe these developments as $10 billion of investment, $10 billion of cloud commitments and $7 billion of equipment sales, producing $27 billion in announced AI-related financial activity. That description would be accurate only if the categories were clearly identified and kept separate. It would be misleading to characterize the combined figure as $27 billion in new, independently funded AI demand because the transactions are connected to the original financing and represent different stages of the same economic process.

Nor would it be correct to claim that the entire $27 billion is somehow fictitious. The equipment supplier has produced valuable systems, the cloud operator has invested in infrastructure, and the model developer has secured the

EXHIBIT 6 How $10 billion becomes a $27 billion headline

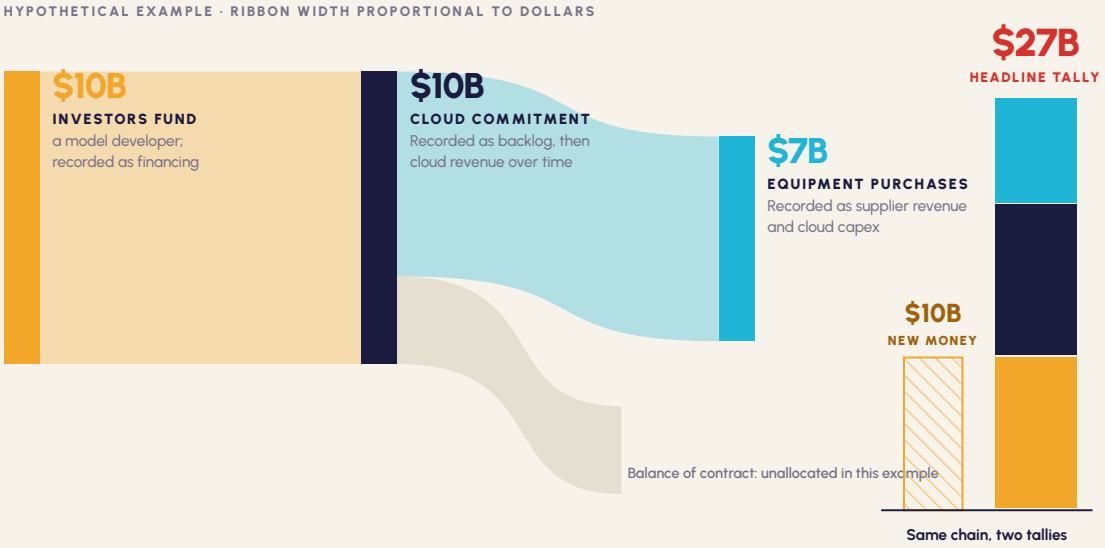

Hypothetical example from this report. Each figure is accurate for its own category; adding them overstates new, independently funded demand.

computing services it expects to need. Different participants contribute products, capital and services that may create substantial additional economic value over time.

The difficulty is that the total value of the underlying economic activity cannot be determined by adding every financial announcement together. A complete analysis must distinguish financing transactions from purchases, purchases from future commitments, and the revenues of intermediate suppliers from the spending of ultimate customers.

This distinction is familiar in economics. National income accounting differentiates gross economic output from value added so that transactions occurring at multiple stages of production do not create an inflated estimate of total economic output. Capital investment also receives specific treatment because the purchase of a productive asset differs from the consumption of an intermediate input.

The comparison is useful, but it should not be taken too far. Public companies report revenue under established accounting standards, and legitimate revenue at different stages of a supply chain does not automatically constitute improper double counting. The problem arises when industry observers combine unlike numbers and present the result as a measure of the market's size or maturity.

AI makes that temptation particularly attractive because the figures are enormous and the boundaries between infrastructure investment, cloud services, model development and commercial applications are still evolving.

When financing and demand become intertwined

The more consequential issue is not simply the possibility of counting the same economic activity more than once. It is the degree to which financing relationships may be influencing purchasing decisions and supporting revenue growth within the ecosystem.

A model developer that raises substantial capital can purchase cloud services long before its own customer revenues are sufficient to fund those purchases. The cloud provider can use the resulting demand to justify infrastructure investment. The equipment supplier recognizes revenue from the expansion, while investors and lenders provide additional capital based on expectations of future growth.

This is a normal feature of venture-backed technological development, and it can be enormously productive. Companies frequently operate at a loss during periods when they are establishing products, building capacity and acquiring customers. Some of the most valuable technology businesses followed precisely that path.

> *Both are real commercial activity. Their sustainability depends on different assumptions.*

The financial challenge becomes more significant as the commitments grow and the period required to achieve sustainable profitability lengthens. If the original customer cannot finance its future purchases through operating revenue or continued investment, the consequences can extend far beyond the customer itself. Infrastructure providers may face underutilized capacity or contract renegotiations, lenders may reassess collateral values, and suppliers may experience weaker orders.

This is the economic relationship that deserves investigation. We should not assume that every transaction involving strategically invested capital is artificial demand, but neither should we treat revenue supported by continued financing as economically equivalent to revenue generated by profitable end customers.

Both are real commercial activity. Their sustainability depends on different assumptions.

The limits of backlog as evidence

Backlog has become one of the most frequently cited indicators of the AI industry's future success. Microsoft, Oracle, CoreWeave and other providers have reported or announced extraordinary amounts of contracted business, with some commitments extending years into the future. These figures can provide valuable information about customer intentions and expected revenue, particularly in a capital-intensive market where providers need confidence before making major infrastructure investments.

Backlog is not cash already collected, however, and the amount ultimately realized depends on contract performance, timing, customer creditworthiness and applicable terms. Different companies use different definitions, and some reported figures include commitments that are not directly comparable with accounting measures such as remaining performance obligations.

A diversified enterprise software company's contracted backlog may reflect thousands of customers purchasing multiple products under established arrangements. A specialized infrastructure provider's backlog may depend on a relatively small number of customers agreeing to purchase substantial quantities of computing capacity over many years. The headline amounts may be similar while the underlying risk profiles differ considerably.

Backlog should be evaluated according to how reliably it becomes revenue, how profitably that revenue can be delivered and how much cash the provider ultimately receives. The size of the commitment is one part of that analysis, but it cannot substitute for the others.

There is also the possibility that commitments announced by different companies are economically related. A model provider may contract with a cloud operator, which finances equipment purchases and maintains its own agreements with hardware suppliers. The resulting contractual announcements describe different obligations within the same chain rather than separate pools of final customer demand.

Any attempt to estimate the total market from these commitments must account for those relationships or risk overstating the scale of independently supported demand.

EXHIBIT 7

Backlog is a promise, not a payment

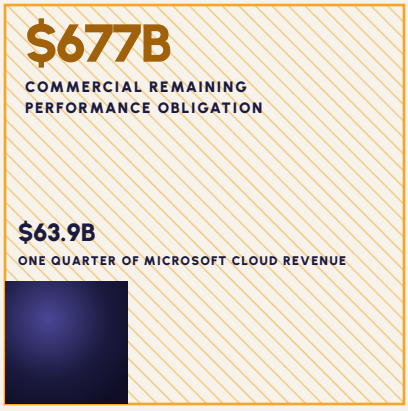

Backlog quality depends on who signed, on what terms, and whether the provider can deliver the capacity profitably — not on the headline size.

Source: Microsoft fiscal Q1 2027 results (left). Right panel is illustrative. Note 3.

SIDEBAR B

Revenue, Backlog *and* Capex Are Not the Same Thing

Why AI financial headlines require a closer reading

WHERE EACH NUMBER SITS IN TIME

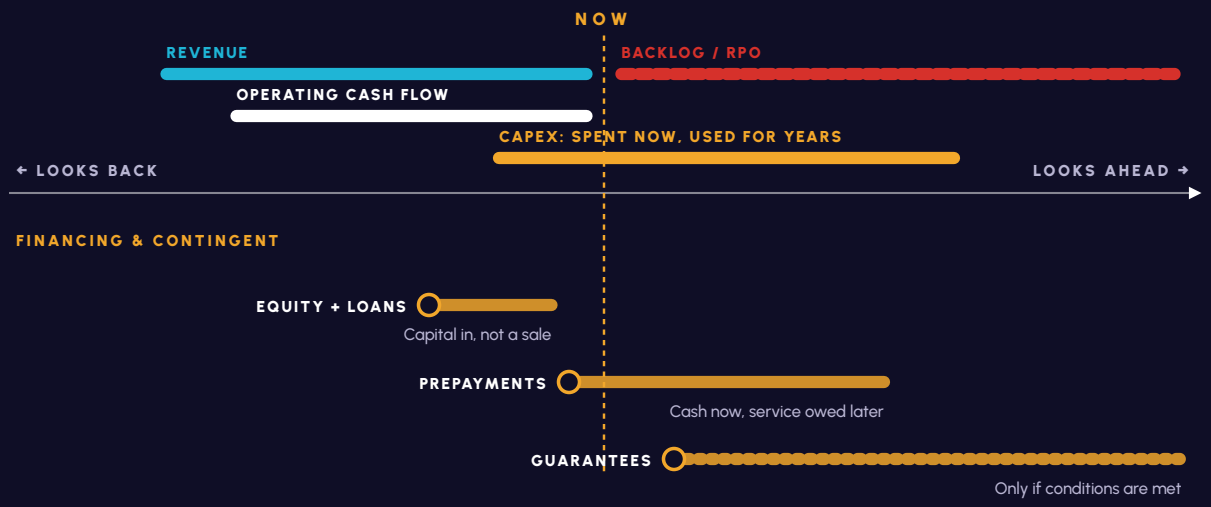

The numbers associated with AI infra-structure have become so large that nearly every major announcement attracts attention. A company reports tens of billions in revenue, another announces hundreds of billions in future commitments, and a third outlines an enormous capital expenditure program. All three developments may be significant, but they describe different economic activities and cannot be combined into a meaningful total without understanding what each measures.

Revenue is recognized when a company satisfies the applicable requirements for reporting sales of products or services. It does not necessarily mean that the customer paid cash during the same period, and it tells us little by itself about the profitability of the business. Operating income provides another perspective by reflecting the company's reported operating profitability, although it does not directly measure cash available after infrastructure spending.

Capital expenditures represent investments in productive assets such as computing equipment, data centers, networking and electrical infrastructure. These assets are generally expected to contribute to the business over multiple years. Operating cash flow measures the cash generated by business operations, while free cash flow is commonly calculated after deducting cash capital expenditures. Different definitions, particularly involving leases and financing structures, can complicate comparisons across companies.

Backlog and remaining performance obligations look forward rather than backward. They reflect contracted or otherwise defined future business, not revenue that has already been earned or cash that has necessarily been collected. The amount and timing of eventual revenue depend on contractual provisions, service delivery and the financial ability of customers to fulfill their obligations.

Equity investments and loans add another category. They provide financing rather than payment for products or services. Customer prepayments can supply cash before the related services are delivered, while guarantees create potential financial obligations if specified conditions are met. Each arrangement can be economically significant without representing current revenue.

The analytical danger is treating all these numbers as equivalent evidence of AI demand. A $10 billion investment, a $10 billion cloud contract and a $7 billion equipment purchase may be different stages in the deployment of infrastructure supported by the same initial financing. Adding them produces a large figure, but not a reliable measure of independently funded customer spending.

For investors and technology leaders, the first question should be what each figure actually represents. Only then can the numbers help establish the strength and sustainability of the business behind them.

04

PART FOUR

Who Is *Actually* Paying for AI?

Real customers are buying. The hard part is measuring how much of the money starts with them.

The end-customer market is developing

THE FINANCIAL RELATIONSHIPS among major AI companies would be considerably more concerning if there were little evidence of real customers purchasing AI products and services. Fortunately, that is not the situation. Enterprises are spending meaningful amounts on AI applications, consumers are paying for subscriptions, and software developers are incorporating model APIs into products that solve actual business problems.

AI-assisted software development has become one of the more visible examples of commercially valuable adoption. Businesses are also deploying AI for customer service, document analysis, security operations, data processing, knowledge retrieval and increasingly sophisticated automation. Some applications remain experimental, but others have become part of routine production workflows.

Research firms tracking enterprise adoption have identified substantial growth in generative AI spending, including purchases of applications, model access and related services. Consumer subscriptions and mobile AI applications provide further evidence that demand extends beyond infrastructure suppliers and model developers purchasing services from one another.

The difficulty is obtaining a consistent picture of the entire market. Enterprise spending surveys often include categories that overlap with cloud revenues or software subscriptions. Consumer payment estimates may omit web purchases and enterprise contracts. Some AI capabilities are bundled into existing products, making it difficult to distinguish new AI revenue from a change in the features offered under established subscriptions.

Internal corporate AI deployments create another complication. A manufacturer using AI to improve production efficiency or a financial services company reducing processing costs may generate genuine economic value without producing any identifiable external AI revenue. The benefits appear through productivity, expense reductions, improved output or other business measures rather than a separate payment to an AI application provider.

This means that independent demand cannot be measured solely by adding subscription revenues and API fees. Nor can it be inferred from the total capital expenditure of the companies providing the infrastructure.

We need to understand the commercial value being created at the point where AI is actually used.

Consumption, monetization and economic value

The AI industry has developed an impressive collection of operational measures. Companies routinely report token volumes, active users, API calls, applications built, model performance, deployed accelerators and available computing capacity. Each measure describes some part of the industry's progress, and many provide legitimate evidence that adoption is expanding.

These figures are less useful when they are treated as substitutes for revenue or profitability. Token consumption may grow rapidly while prices decline, and active users do not necessarily become paying subscribers. Enterprise deployments can generate substantial usage without delivering acceptable financial returns, particularly when the costs of integration,

EXHIBIT 8 **Three different ways to size the AI economy**

Techstrong schematic. Circle sizes are conceptual and not to scale.

governance, security and ongoing inference are included.

The relationship can also work in the other direction. A relatively small AI application may deliver exceptional economic value by eliminating expensive manual processes, improving the accuracy of important decisions or allowing employees to accomplish work that previously required significantly more time. Its contribution to enterprise productivity may be much greater than the amount paid to the AI provider.

Consequently, the size of the AI economy cannot be evaluated through a single operating metric. We need to distinguish the amount of computing being consumed, the revenue being generated from that consumption and the value customers receive after accounting for their own costs.

These measures will become increasingly important as inference grows. Training a frontier model can require enormous but episodic computing expenditures. Inference potentially creates recurring demand as applications serve customers continuously. The value of that recurring demand depends on pricing, efficiency, competitive alternatives and the willingness of users to pay for the services involved.

Continued progress in model efficiency may improve profitability, but it can also create price competition. Lower inference costs may encourage much greater usage, potentially producing significant aggregate demand even as the price of individual computations falls. Whether that expansion produces sufficient revenue to justify the infrastructure already being installed remains an open financial question.

Model developers connect infrastructure to customers

OpenAI, Anthropic and other leading model developers occupy a particularly important position in the economic chain. They are among the largest consumers of advanced AI computing, while also being responsible for commercializing much of the capability that creates demand for that computing.

Their operating economics therefore offer a potential bridge between the money being spent upstream and the value being generated downstream. Revenue from subscriptions, APIs and enterprise products provides evidence that customers are purchasing AI capabilities. Spending on training, inference and infrastructure indicates the cost of producing those capabilities.

> *A market can be growing rapidly while remaining difficult to measure accurately.*

What investors and industry observers need to understand is how quickly the first set of figures is catching up with the second.

The October 2026 Anthropic filing provides a potentially valuable opportunity to examine a frontier model developer's financial position more closely, including its recognized revenue, operating expenses, cash flows and long-term compute commitments.[11] Such information can improve our understanding of the relationship between rapid commercial growth and the capital required to maintain a competitive position.

Any comparison must account for the timing and terms of the underlying agreements. Long-term infrastructure commitments may extend over many years, while reported revenue describes a much shorter period. Comparing the face value of one with the other without accounting for contract duration would produce a distorted conclusion.

At the same time, a company cannot indefinitely support infrastructure commitments through expectations of future growth alone. Ultimately, commercial revenue, available financing or some combination of the two must fund the obligations as they become due.

The model developers' ability to build sustainable businesses is therefore central to the financial outlook for the providers supplying their computing capacity. If model revenues expand rapidly and unit economics improve, today's

IN PLAIN ENGLISH

Three ways to measure the AI economy

CONSUMPTION
How much computing is being used: tokens processed, API calls, active users, deployed accelerators.

MONETIZATION
The revenue that consumption actually produces. Token volumes can soar while prices fall, and active users are not necessarily paying subscribers.

ECONOMIC VALUE
What customers gain after their own costs of integration, governance, security and inference. A small application can create value far larger than what it pays its AI provider.

WHY IT MATTERS
No single operating metric sizes the AI economy. Usage, revenue and value have to be read separately.

infrastructure commitments may prove well justified. If growth slows or the cost of providing services remains too high, financial pressure could move through the ecosystem.

The demand bridge we still cannot see

A useful measure of the AI economy would trace independent customer payments through applications and model providers to the underlying infrastructure that supports them. It would account for the value created at each stage while distinguishing operating revenue from capital investment and financing transactions.

Public disclosures do not currently provide enough consistent information to construct that measure with confidence. Hyperscalers do not always identify AI revenue separately from traditional cloud workloads. Model developers vary considerably in the information they disclose. Infrastructure providers may report capacity and commitments without enough detail to determine how much is independently utilized at commercially attractive prices.

The absence of a complete measurement doesn't establish that the demand is insufficient. Infrastructure can be built productively ahead of customer consumption, and some of the most important benefits of AI may not appear in provider revenue until new applications mature.

It does mean that confident claims about the aggregate financial size of the AI economy should be treated cautiously, particularly when those claims combine hardware sales, financing commitments, cloud revenues and future contracts.

A market can be growing rapidly while remaining difficult to measure accurately. The AI industry appears to be demonstrating both characteristics.

> *Whether that expansion produces sufficient revenue to justify the infrastructure already being installed remains an open financial question.*

EXHIBIT 9 **The demand bridge we still cannot see**

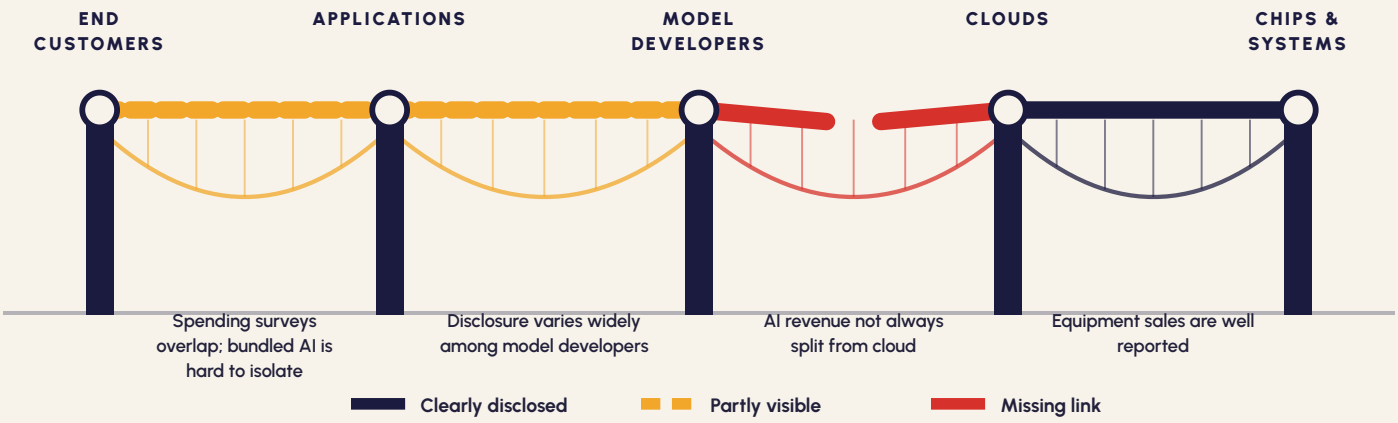

Techstrong editorial assessment of public disclosure across the AI value chain, based on the analysis in this report.

PART FIVE

Who Is *Holding* the Risk?

> The same slowdown that a hyperscaler can absorb could destabilize a specialized provider built on debt and a few large customers.

The hyperscalers have financial advantages

ONE REASON THE AI investment cycle has expanded so quickly is that several of its largest participants are among the strongest cash-generating businesses in the world. Microsoft, Alphabet, Amazon and Meta possess established franchises capable of producing substantial operating cash flow independently of their newest AI investments.

That strength provides important protection against uncertainty. Microsoft can fund new computing infrastructure through cash generated across enterprise software and cloud services. Alphabet benefits from its advertising businesses while growing Google Cloud. Amazon has established commerce and cloud operations, while Meta continues generating substantial advertising revenue.

Their financial resources allow them to take risks that would be much more difficult for companies with limited operating histories or concentrated revenue sources. A hyperscaler can tolerate an extended period of infrastructure investment while waiting for demand to mature, and it can absorb underperforming projects without necessarily jeopardizing its broader business.

The largest companies are not immune to economic mistakes. Spending tens or hundreds of billions on infrastructure that produces disappointing returns could eventually affect margins, shareholder returns and capital allocation. However, the consequences would differ substantially from those

experienced by a highly leveraged provider whose principal business depends on the infrastructure itself.

This is why the distribution of financial risk matters as much as the total spending. The same technological transition can be manageable for a diversified hyperscaler and financially destabilizing for a more specialized participant.

The financing sensitivity of specialized clouds

AI infrastructure requires substantial investment before revenue can be generated. Computing equipment must be purchased, deployed and connected to power and networking resources, often long before customer demand reaches its expected level. Facilities may require extensive construction, specialized cooling and electrical systems, while operating expenses begin accumulating as capacity becomes available.

For specialized AI cloud providers, these requirements can create significant financing exposure. Borrowing allows infrastructure to be built more quickly, but interest payments, collateral requirements and refinancing obligations introduce costs that must be supported by future operating cash flow.

Customer contracts help reduce uncertainty, particularly when they include meaningful commitments from financially strong counterparties. Concentrated customers can also create vulnerabilities. If one or two companies represent a substantial portion of revenue, changes in their demand or financial condition can affect the provider's entire business.

CoreWeave's reported customer concentration and financing expenses make it an important case study.[12] Its commercial growth demonstrates the market opportunity, while its balance sheet shows how much capital can be required to pursue that opportunity.

The key issue is whether infrastructure revenues remain sufficient to cover operating expenses, financing

KEY TERMS

How specialized AI clouds are financed

SPECIALIZED AI CLOUD
Also called a neocloud: a provider that rents accelerated computing capacity to model developers and enterprises, often with a small number of very large customers.

CUSTOMER CONCENTRATION
When one or two customers represent a large share of revenue, a change in their demand or finances can affect the whole business.

COLLATERAL AND REFINANCING
Debt-funded equipment is often pledged to lenders. Interest, collateral requirements and refinancing must be covered by future operating cash flow.

EXHIBIT 10 **Where the financial risk concentrates**

	CUSTOMER CONCENTRATION	DEPENDENCE ON OUTSIDE FINANCING	SENSITIVITY TO CHEAPER COMPUTE	TIES TO PARTNERS' FINANCES
Diversified hyperscalers	● LOWER	● LOWER	●● MODERATE	●● MODERATE
Specialized AI clouds	●●● HIGHER	●●● HIGHER	●●● HIGHER	●●● HIGHER
Model developers	●● MODERATE	●●● HIGHER	●● MODERATE	●●● HIGHER
Equipment suppliers	●● MODERATE	● LOWER	● LOWER	●● MODERATE
Lenders & strategic investors	●● MODERATE	● LOWER	●● MODERATE	●●● HIGHER

Qualitative editorial assessment drawn from the analysis in this report; not a credit rating.

Qualitative Techstrong editorial assessment based on the analysis in this report.

The Neocloud Balance Sheet

SIDEBAR C

Why extraordinary growth doesn't guarantee attractive returns

COREWEAVE, SECOND QUARTER 2026 · NOTE 12

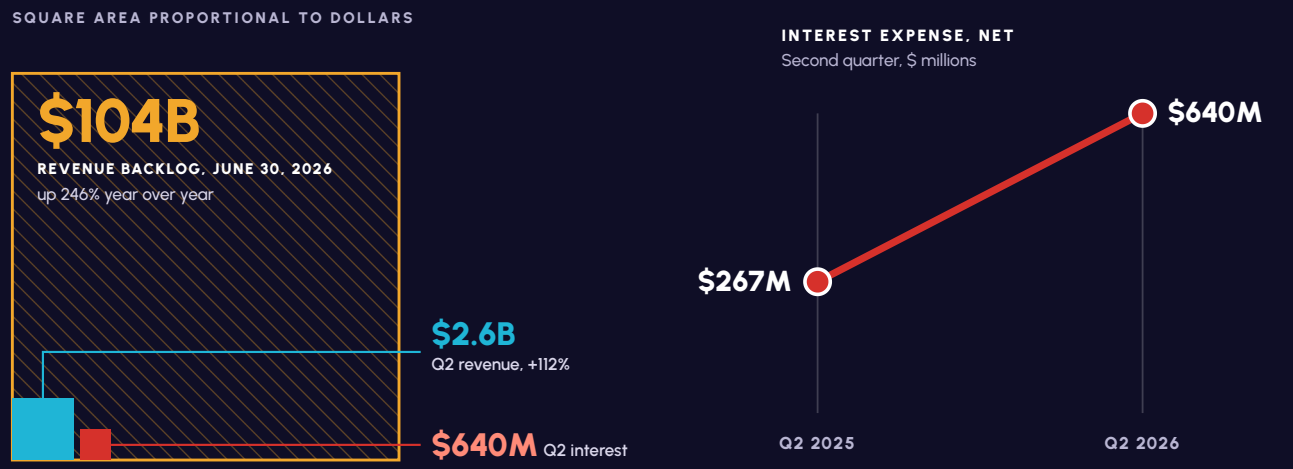

Specialized AI cloud providers have emerged as important participants in the infrastructure buildout, offering advanced computing capacity to model developers, enterprises and other organizations that need access to large numbers of accelerators. Their rapid expansion reflects genuine demand for infrastructure that is expensive, difficult to deploy and increasingly important to businesses developing AI products.

The financial model is considerably more demanding than that of a conventional software company. A specialized cloud provider must secure facilities, electrical capacity, cooling infrastructure and high-performance networking before it can deliver computing services. It must purchase or finance expensive hardware, install and maintain the equipment, and attract sufficient customer demand to generate revenue throughout the assets' economic lives.

Long-term customer contracts can support financing by providing lenders and investors with some visibility into expected future payments. Those arrangements can accelerate growth, allowing providers to develop capacity more quickly than would be possible using internally generated cash alone. They also create dependence on the customers signing the contracts and on the infrastructure remaining commercially competitive throughout the contract period.

Customer concentration is particularly important. A provider whose revenue depends on a handful of large AI companies can experience substantial financial pressure if one customer reduces spending, renegotiates obligations or struggles to obtain additional financing. Even a binding contract may require considerable time and expense to enforce, and contractual remedies do not necessarily compensate for every economic consequence of an interruption.

Financing costs introduce further uncertainty. Equipment purchases funded through debt create interest and repayment obligations that continue regardless of whether the infrastructure achieves expected utilization. Refinancing may become more expensive if credit conditions tighten, equipment values decline or investors begin questioning the returns available from older generations of accelerators.

The pace of technological improvement creates both opportunity and risk. More efficient hardware can reduce computing costs and expand demand, but it may also place downward pressure on the commercial value of previously installed systems. A provider must generate sufficient revenue over the useful life of its assets to cover operating expenses, financing costs and the economic depreciation of its equipment.

None of these challenges makes the specialized AI cloud business inherently unsustainable. Efficient operators with strong customer relationships and disciplined financing may build valuable long-term businesses. The lesson is that rapid revenue growth, impressive backlog and expanding installed capacity do not by themselves establish attractive returns. For neoclouds, the financial performance of the infrastructure matters as much as the volume of demand it attracts.

obligations and the economic consumption of equipment over time. Growth in recognized revenue is encouraging, but it cannot establish that the returns on invested capital will ultimately be attractive.

The strongest specialized providers may succeed by delivering superior operational efficiency, better utilization and close alignment with customer requirements. Others may discover that rapid growth financed on optimistic assumptions creates difficult choices when demand, pricing or capital markets change.

The useful life of AI infrastructure

The rapid pace of hardware development introduces another layer of financial uncertainty. AI factories contain assets with very different economic lives. Buildings, electrical systems and some networking infrastructure may support multiple technology generations, while accelerators and other components can become less competitive as newer systems improve performance and energy efficiency.

The relevant question is not simply how long equipment can operate. It is how long that equipment can provide commercially competitive computing services at prices sufficient to cover its costs.

A provider may be able to run an older GPU cluster for many years, but customers may prefer newer systems if they deliver substantially more computing capacity per watt or per dollar. Falling compute prices can make AI applications more affordable and encourage adoption, while also reducing the revenues available from equipment acquired at earlier cost levels.

These developments create potential tension between the industry's technological success and the financial performance of specific infrastructure investments. A new generation of efficient systems can be good for AI adoption while making some older assets less profitable.

Networking and system architecture add another dimension. Expensive accelerators deliver limited economic value when they cannot be kept productive because of communication delays, data movement bottlenecks or insufficient connectivity. The efficiency of

IN PLAIN ENGLISH

Physical life vs. economic life

PHYSICAL LIFE

How long equipment can keep running. Buildings, electrical systems and some networking may serve several hardware generations.

ECONOMIC LIFE

How long equipment can sell competitive computing at prices that cover its costs. Accelerators can lose ground quickly as newer systems improve performance and energy efficiency.

WHY IT MATTERS

If economic life is shorter than the financing behind the asset, the revenue may stop covering the debt before the debt is repaid.

EXHIBIT 11 **AI factories contain assets with very different lives**

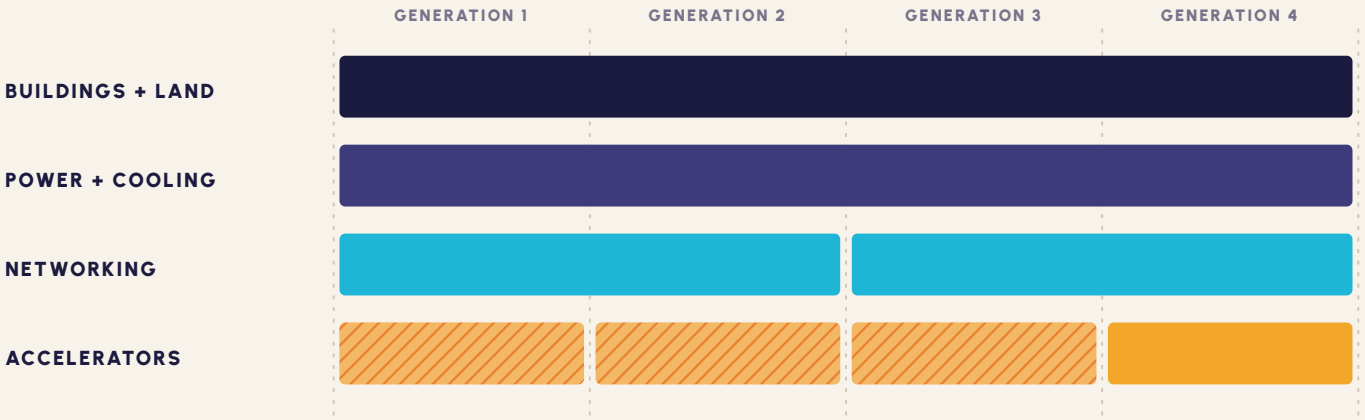

Illustrative. Older accelerators can keep running long after they stop being priced competitively.

Illustrative. Hardware generations are not drawn to a time scale.

scale-up, scale-out and other network layers can materially affect the output produced by an installed cluster.

This supports the argument that AI infrastructure economics should be evaluated at the system level rather than through GPU specifications alone. A well-designed AI factory may remain commercially valuable longer than a poorly integrated facility containing comparable accelerators.

> *Technological usefulness does not guarantee an attractive return for every investor, lender or infrastructure operator.*

The amount invested in infrastructure is only one part of the calculation. Its productivity over time is what determines the return.

How financial stress can spread

The AI ecosystem's interconnected financing arrangements can reduce risks during periods of strong growth. Strategic investments provide capital, long-term purchase commitments improve revenue visibility, and financing guarantees can make infrastructure projects possible that might otherwise struggle to attract funding.

Those same relationships can create additional exposure if demand disappoints.

Consider a model developer that faces slower-than-expected customer revenue growth and difficulty raising additional financing. The company may seek to limit future compute purchases or renegotiate contractual obligations. An infrastructure provider depending on those purchases could then experience lower utilization, reduced revenue or increased financing pressure.

Lenders might respond by reassessing collateral values or requiring additional protection. Equipment suppliers could see weaker orders, while strategic investors and counterparties could face losses associated with their own commitments.

The result would not necessarily be a collapse of the AI industry. Demand from other customers could continue growing, and financially stronger providers might absorb capacity or acquire distressed assets. Nevertheless, the commercial relationships that helped accelerate expansion could also transmit financial pressure across participants.

We have seen technology investment cycles in which infrastructure remained valuable even after some of the companies that financed it failed. Telecommunications networks, internet infrastructure and other capital-intensive sectors provide useful historical comparisons, although the specific technology and market conditions differ.

AI may follow its own path, but the financial lesson is familiar. Technological usefulness does not guarantee an attractive return for every investor, lender or infrastructure operator involved in building the market.

EXHIBIT 12 **How financial stress can spread**

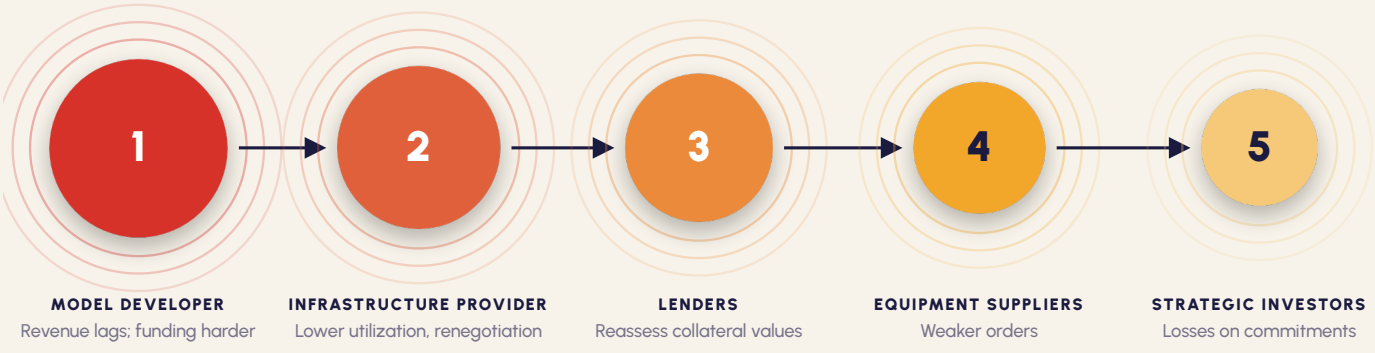

Techstrong schematic of the transmission path described in this report. Not a forecast.

06

PART SIX

When *the* Music Slows

Three ways the cycle could play out — and why even the best case won't reward every company that financed it.

The successful expansion

THE MOST FAVORABLE outcome is that AI adoption expands sufficiently to absorb the infrastructure being constructed. Inference becomes an increasingly common component of enterprise applications, consumer services and industrial systems, while improvements in efficiency make new use cases economically practical. Model developers convert usage into commercial revenue, infrastructure providers achieve strong utilization, and the cost of computing declines without undermining the profitability of the most efficient operators.

Under those conditions, much of today's aggressive investment could prove justified. Strategic financing relationships would be remembered as mechanisms that accelerated the development of productive infrastructure, and companies that established strong positions early could enjoy meaningful advantages.

This possibility is consistent with the history of major computing transitions. Early investment often appears excessive when measured against existing demand, particularly when it supports technologies that enable entirely new categories of applications.

The difficulty is that successful long-term adoption does not automatically validate every investment decision made during the expansion. Some projects may still produce poor returns because of excessive construction costs, weak utilization, unfavorable financing or equipment that becomes uneconomic sooner than expected.

Even the optimistic scenario should therefore distinguish the success of AI as a technology from the financial success of every company participating in the buildout.

AT A GLANCE

Three ways the cycle could play out

A · SUCCESSFUL EXPANSION
Adoption grows fast enough to absorb the capacity being built. Even then, some projects may earn poor returns because of cost, timing or competition.

B · DIGESTION CYCLE
Demand keeps rising, but new capacity arrives faster. Utilization dips, prices compete, and projects and orders slow while customers benefit.

C · FINANCING DISRUPTION
A key participant struggles to raise money or meet commitments, and the problem travels through contracts, investments, guarantees and customer concentration.

The infrastructure digestion cycle

A second scenario involves continued growth accompanied by an adjustment in the pace of infrastructure investment. Demand for AI services could rise substantially while new capacity becomes available even faster, creating periods of lower utilization and increased pricing competition.

Under these conditions, some planned projects might be delayed, equipment orders could moderate, and infrastructure providers might become more selective about expansion. Customers would benefit from more competitive pricing, potentially increasing adoption while reducing the profitability of capacity installed at higher costs.

This would resemble the adjustment periods that have followed earlier investment booms. The industry's installed infrastructure could remain valuable, while the financial returns to some of its owners fall short of expectations.

Such a correction would not necessarily mean that companies had been wrong about AI's long-term importance. It might mean they misjudged the timing of demand, the amount of capacity required or the price customers would ultimately pay for computing services.

A period of consolidation could also improve the industry's longer-term economics if stronger operators acquired facilities and equipment from companies unable to finance their obligations.

> *Even the optimistic scenario should distinguish the success of AI as a technology from the financial success of every company in the buildout.*

A disruption in the financing network

A third scenario involves more significant financial stress originating with one or more major participants. The trigger could be slower revenue growth at a model developer, a change in credit conditions, difficulties refinancing infrastructure debt or a sharp decline

in the market value of computing assets.

In a system with extensive reciprocal financial relationships, the consequences could extend beyond the company where the difficulty began. Contractual obligations, investments, guarantees and customer concentration create channels through which changes in one company's financial condition can affect others.

The severity would depend on the amount of exposure, the enforceability of commitments, available liquidity and the financial strength of the counterparties. A guarantee does not necessarily result in a payment, and a commercial contract does not necessarily become worthless when a customer experiences difficulty. Each arrangement has its own terms, protections and potential remedies.

Nevertheless, an ecosystem in which major participants finance and purchase from one another can create correlated exposures that are difficult to appreciate by reading individual financial reports.

This scenario should not be presented as inevitable. The companies involved have substantial commercial opportunities, and many possess significant financial resources. The relevant concern is that the connections supporting rapid expansion can also affect how financial stress is distributed if market conditions deteriorate.

The industry would benefit from greater transparency about those relationships before they become the subject of an unexpected financial adjustment.

> *Successful long-term adoption does not automatically validate every investment decision made during the expansion.*

EXHIBIT 13 **Three scenarios for the AI infrastructure cycle**

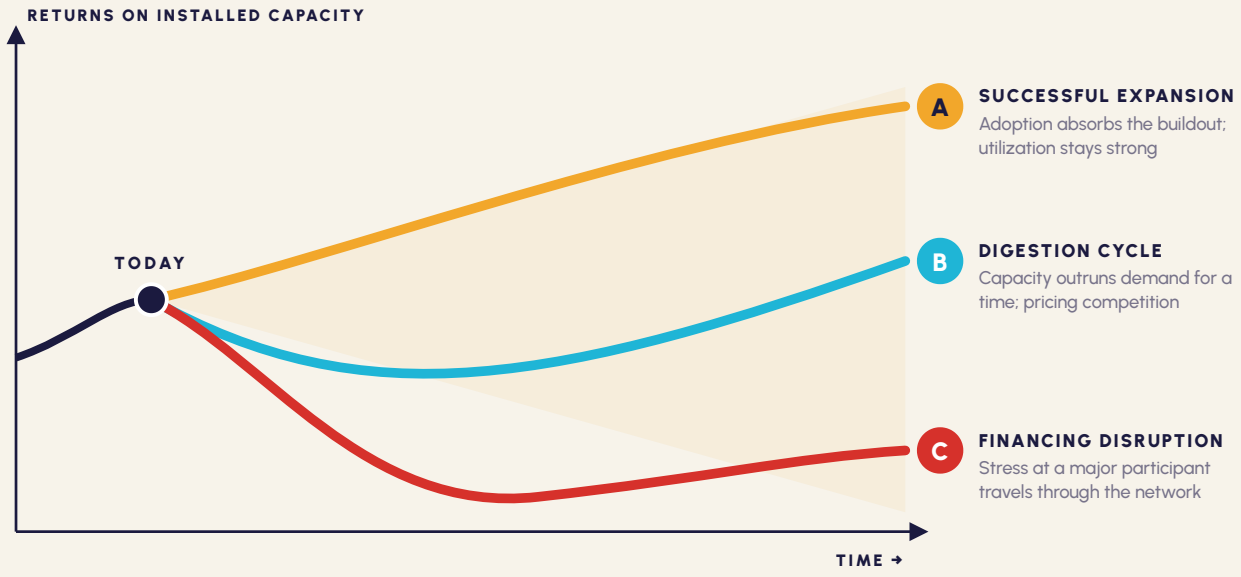

Illustrative scenario paths, not forecasts.

Illustrative scenario paths drawn from this report. Not forecasts.

SIDEBAR D

Five Questions Every AI Investor Should Ask

Looking beyond record revenue and enormous backlogs

The extraordinary growth of the AI industry has encouraged investors and business leaders to focus on revenue acceleration, announced contracts and the scale of infrastructure expansion. Those measures provide useful information, but they do not fully explain the quality of demand or the financial obligations associated with serving it. Five questions can help establish a more complete picture.

1 Where does the customer's money originate?

Revenue generated from purchases by established enterprises and consumers may have a different risk profile from revenue generated by customers depending heavily on strategic investments or repeated fundraising. Both can represent legitimate sales, but the durability of the underlying demand depends on the customers' ability to continue paying. Understanding the sources of funding is particularly important when buyers are committing to infrastructure services over several years.

2 Does the company invest in, finance or purchase services from its major customers?

Strategic investments, reciprocal purchasing commitments and supplier financing can support valuable commercial relationships. They can also create economic dependencies that make reported revenue difficult to evaluate in isolation. Investors should understand whether significant customers are also investees, financing counterparties or suppliers, and how those relationships could affect the business if market conditions change.

3 How much of the reported backlog will become cash, and when?

Large contractual commitments provide visibility into expected business, but contract terms determine how much certainty that visibility actually provides. Investors should examine the timing of revenue recognition, payment schedules, cancellation rights, customer creditworthiness and the infrastructure spending required to fulfill agreements. A large backlog supported by concentrated customers may carry different risks from a smaller backlog spread across thousands of enterprises.

4 What happens if computing capacity becomes cheaper faster than expected?

Lower inference costs and improving hardware performance can accelerate AI adoption while reducing the revenues available from older equipment. Infrastructure operators need utilization levels and pricing sufficient to cover operating expenses, financing obligations and the economic consumption of their assets. Investors should understand the assumptions being made about depreciation, useful equipment life, future pricing and replacement costs.

5 Can the company sustain its operations if financing becomes more expensive or demand takes longer to develop?

The strongest hyperscalers generate substantial cash from diversified businesses, while some specialized infrastructure providers depend more directly on debt markets, equity financing and concentrated customer commitments. Evaluating liquidity, interest costs, debt maturities and financial guarantees provides a clearer picture of how a company might perform if infrastructure growth slows or capital markets tighten.

These questions are not intended to undermine the investment case for AI. The technology is creating substantial commercial opportunities, and many companies are already benefiting. They are intended to ensure that the size of reported financial activity does not become a substitute for understanding its underlying economics.

07

PART SEVEN

What the Industry Needs to Tell Us

THE AI INDUSTRY has become extraordinarily effective at communicating scale. Companies announce enormous capital expenditures, multiyear infrastructure agreements, growing backlogs, increasing token consumption and the construction of data centers measured in hundreds of megawatts or even gigawatts. These figures demonstrate the ambition of the buildout and the extraordinary resources being committed to it.

What remains less transparent is the relationship between that spending and the independent customer demand expected to support it. Public financial statements contain considerable relevant information, but the disclosures are distributed across different categories, reporting periods and contractual arrangements. The resulting picture is difficult to assemble even for sophisticated investors and industry analysts.

A more useful framework would identify revenue earned from strategic investees and other financially supported customers, explain material reciprocal purchasing commitments and clarify the composition and expected timing of contractual backlogs. Infrastructure operators could provide greater visibility into utilization, equipment life and the economics of new capacity, while model developers could help investors understand the relationship between commercial revenue, compute expense and long-term infrastructure obligations.

Some information is already required under existing accounting and

THE FRAMEWORK

Four kinds of AI money

INDEPENDENT DEMAND
Payments from customers outside the strategic financing relationships.

ECOSYSTEM ACTIVITY
Legitimate sales among suppliers, clouds and model developers.

FINANCING
Equity, loans and vendor financing: capital, not payment.

FUTURE COMMITMENTS
Backlog and performance obligations not yet revenue or cash.

securities disclosure rules. Other details may be commercially sensitive or difficult to separate from broader operations. Improving transparency does not necessarily require an entirely new reporting regime, but it does require recognizing that conventional financial comparisons can conceal important relationships in an ecosystem where the same companies frequently occupy several economic roles.

The broader strategic implications extend beyond accounting. As companies become indispensable to one another as suppliers, customers and financing partners, they can establish powerful relationships that accelerate innovation and make the industry difficult for competitors to challenge. Those relationships can also concentrate financial risk and make the success of one participant increasingly dependent on the assumptions supporting others.

This is an important aspect of the transformation now underway. AI is not simply producing a new generation of software or an unprecedented demand for computing. It is changing the financial relationships through which technology infrastructure is developed, owned and consumed.

> *The earnings releases tell us how quickly the money is moving.*

The October earnings season provided convincing evidence that the industry is growing rapidly. The revenue is real, the infrastructure is being built and the demand for AI capabilities is substantial. We should recognize those achievements without accepting that the size of the reported financial activity establishes the sustainability of the entire investment cycle.

To make that judgment, we need a better understanding of the money entering the ecosystem from independent customers, the economic value created as it moves through the supply chain and the obligations being accumulated along the way.

AI may ultimately justify an infrastructure investment cycle larger than anything the technology industry has previously experienced. But an outcome that is good for the technology does not guarantee an equally favorable result for every company financing it.

The earnings releases tell us how quickly the money is moving. Understanding where it originates, how often it changes hands and who ultimately bears the risk will tell us much more about the financial future of AI.

EXHIBIT 14 Five disclosures that would make AI finances legible

		HYPERSCALERS	SUPPLIERS	INFRA OPERATORS	MODEL DEVELOPERS
01	Revenue from strategic investees and financially supported customers	✓	✓	✓	·
02	Material reciprocal purchasing commitments	✓	✓	✓	✓
03	Composition and expected timing of contractual backlog	✓	✓	✓	✓
04	Utilization, equipment life and the economics of new capacity	✓	·	✓	·
05	Commercial revenue vs. compute expense vs. long-term obligations	·	·	·	✓

Techstrong recommendations drawn from Part Seven, mapped to the participants best placed to provide each disclosure.

METHODOLOGY

How we *followed the money*

This report examines four broad categories of AI-related financial activity: independent customer spending, commercial revenue generated within the AI ecosystem, financing transactions and future contractual commitments. These categories are used to distinguish different economic functions rather than create mutually exclusive accounting classifications.

Independent customer demand includes identifiable payments for AI products and services from enterprises, consumers and other customers outside the immediate strategic financing relationships under examination. Ecosystem commercial activity includes legitimate transactions among infrastructure suppliers, cloud providers, model developers and related businesses. Financing transactions include equity investments, borrowing, vendor financing and other arrangements that provide capital. Future commitments include backlog, remaining performance obligations and contractual purchasing arrangements that have not necessarily resulted in current revenue or cash payments.

Financial measures are not added across categories to calculate total AI demand. Different reporting periods, accounting definitions and contractual terms limit direct comparisons between companies. Hypothetical examples are identified as such, and potential financial outcomes are discussed as scenarios rather than forecasts.

The report does not attempt to calculate an industry-wide total for recycled AI capital because publicly available disclosures are insufficient to establish the origins and subsequent uses of all relevant cash flows. Documented overlapping relationships should not be interpreted as evidence of improper accounting or artificial transactions without additional supporting facts.

Research cutoff: October 28, 2026. Figures and contractual characterizations require final source reconciliation before publication.

NOTES AND SOURCES

1. Alphabet third-quarter 2026 results: revenue $127.58 billion (+25%), Google Cloud revenue $30.7 billion, capital expenditures $51.36 billion, free cash flow –$9.07 billion. **CNBC, Oct. 28, 2026**.
   https://www.cnbc.com/2026/10/28/alphabet-googl-earnings-q3-2026.html
2. Google Cloud third-quarter operating income of approximately $11.1 billion; Alphabet Q3 2026 earnings remarks. **Google, Q3 2026 earnings call remarks**.
   https://blog.google/company-news/inside-google/message-ceo/alphabet-earnings-q3-2026/
3. Microsoft fiscal 2027 first-quarter results: revenue $92.3 billion (+19%), Microsoft Cloud $63.9 billion, Azure +48%, commercial RPO $677 billion; net cash from operations $58.0 billion and additions to property and equipment $39.4 billion. **Microsoft, Oct. 28, 2026**.
   https://news.microsoft.com/source/2026/10/28/microsoft-announces-first-quarter-results-2/
4. NVIDIA second-quarter fiscal 2027 results: revenue $96.2 billion; Data Center revenue $89.0 billion, up 117% year over year. **NVIDIA**.
   https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2027/default.aspx
5. Amazon second-quarter 2026 results: AWS net sales up 37% to $42.2 billion. **Amazon**.
   https://ir.aboutamazon.com/news-release/news-release-details/2026/Amazon-com-Announces-Second-Quarter-Results/
6. Meta third-quarter 2026 results: capital expenditures, including principal payments on finance leases, of $34.91 billion. **Meta, Exhibit 99.1, SEC**.
   https://www.sec.gov/Archives/edgar/data/1326801/000162828026066700/meta-09302026xexhibit991.htm
7. Microsoft fiscal 2026 Form 10-K: total OpenAI funding commitments of $13.0 billion, $11.9 billion funded as of June 30, 2026; FY2026 revenue from commercial arrangements with OpenAI of $24.1 billion; receivables of $6.0 billion; approximate 25% interest on an as-converted basis. **Microsoft Form 10-K, fiscal 2026**.
   https://microsoft.gcs-web.com/static-files/8a9beeed-0e9e-48e6-b4fb-593d558e2693
8. NVIDIA Form 10-Q for the quarter ended July 26, 2026: AI cloud-service commitments totaling $36 billion, typically six years in duration, which decrease as capacity is used by third-party customers or by NVIDIA. **NVIDIA Form 10-Q via MarketScreener**.
   https://www.marketscreener.com/news/nvidia-quarterly-report-for-quarter-ending-july-26-2026-form-10-q-ce7858d9d18df624
9. Reporting on NVIDIA's role as an AI financier, including off-balance-sheet cloud agreements and guarantees. **CNBC, Oct. 24, 2026**.
   https://www.cnbc.com/2026/10/24/nvidias-role-as-ai-financier-faces-first-big-test-as-neoclouds-take-market-hit.html
10. Oracle Investor Day 2026: vendor financing, customer-provided hardware and customer prepayments expected to reduce upfront cash requirements. **Oracle, Oct. 28, 2026**.
    https://investor.oracle.com/investor-news/news-details/2026/Oracle-Investor-Day-2026-Accelerating-Value-Creation/default.aspx
11. Coverage of Anthropic's October 2026 S-1 registration statement. **CNBC, Oct. 26, 2026**.
    https://www.cnbc.com/2026/10/26/anthropic-ipo-live-updates-ai-s-1-filing.html
12. CoreWeave second-quarter 2026 results: revenue of $2.6 billion (+112%), revenue backlog of approximately $104 billion (+246%), net interest expense of $640 million versus $267 million a year earlier. **CoreWeave**.
    https://investors.coreweave.com/news/news-details/2026/CoreWeave-Reports-Strong-Second-Quarter-2026-Results/default.aspx

All photographs in this report are AI-generated photo illustrations created for editorial purposes. Exhibits are original Techstrong graphics.

An outcome that is good for the technology does not guarantee an equally favorable result for every company financing it.

Techstrong

Techstrong is the technology media, research and events network for builders and buyers of cloud-native, DevOps, security and AI infrastructure. Techstrong is part of Futurum Media.