SECTOR REPORT | MACROECONOMIC RESEARCH DIVISION | PUBLISHED 23 NOVEMBER 2026

AI Capital Expenditure

The Economics of the Infrastructure Supercycle

Investment Sustainability, Industry Bottlenecks, and the Path to Economic Returns

BARUCH ISLAMIC INVESTMENT RESEARCH GROUP

RESEARCH CUTOFF 23 NOVEMBER 2026

AUTHORS **Oleg Godun**, Chief Investment Officer · **Mohammad Zaman**, **Mahdeer Rahman**, **Louise Phillipe Sehi** and **Syed Muhammad Hashmi**, Analysts

Coverage: Microsoft (MSFT), Alphabet (GOOGL), Amazon (AMZN), Meta (META), NVIDIA (NVDA), Vertiv (VRT); U.S. macroeconomic data and selected international supply-chain indicators | Sector view: Constructive, with increasing selectivity

FOR RESEARCH PURPOSES ONLY

This report is independent educational and non-commercial research. It is not investment advice, a personal recommendation, an offer or a solicitation to buy or sell any security, and it should not be relied on to make investment decisions. Companies are discussed to illustrate positions in the AI infrastructure value chain. See Disclosures and Important Information, page 19.

> **Central research question.** Can the AI infrastructure investment cycle produce sufficient commercial and economic returns to justify its expanding capital intensity, financing needs and physical infrastructure demands?

EXECUTIVE SUMMARY AND INVESTMENT VIEW

Artificial intelligence has become one of the defining capital investment cycles of this decade. The largest technology companies are deploying extraordinary sums into accelerators, data centers, connectivity and power infrastructure. The commercial evidence is already visible in cloud growth, semiconductor sales and the expansion of power and thermal systems. Returns will follow a different timetable. Capital goes out today; utilization, pricing power and cash generation determine what comes back over the coming years.

We remain constructive on AI infrastructure demand and increasingly selective about who earns the returns. Semiconductor and equipment suppliers are monetizing the buildout now. Hyperscalers are taking on the capital burden and will earn their returns as installed capacity converts into profitable workloads. Microsoft continues to generate substantial cash after investment; Alphabet's accelerating purchases have tightened near-term cash conversion. Amazon's gross equipment spending now exceeds operating cash generation, while Meta's financing cushion has narrowed. The financial capacity to keep building is substantial. The investment hurdle is rising with it. [1–4]

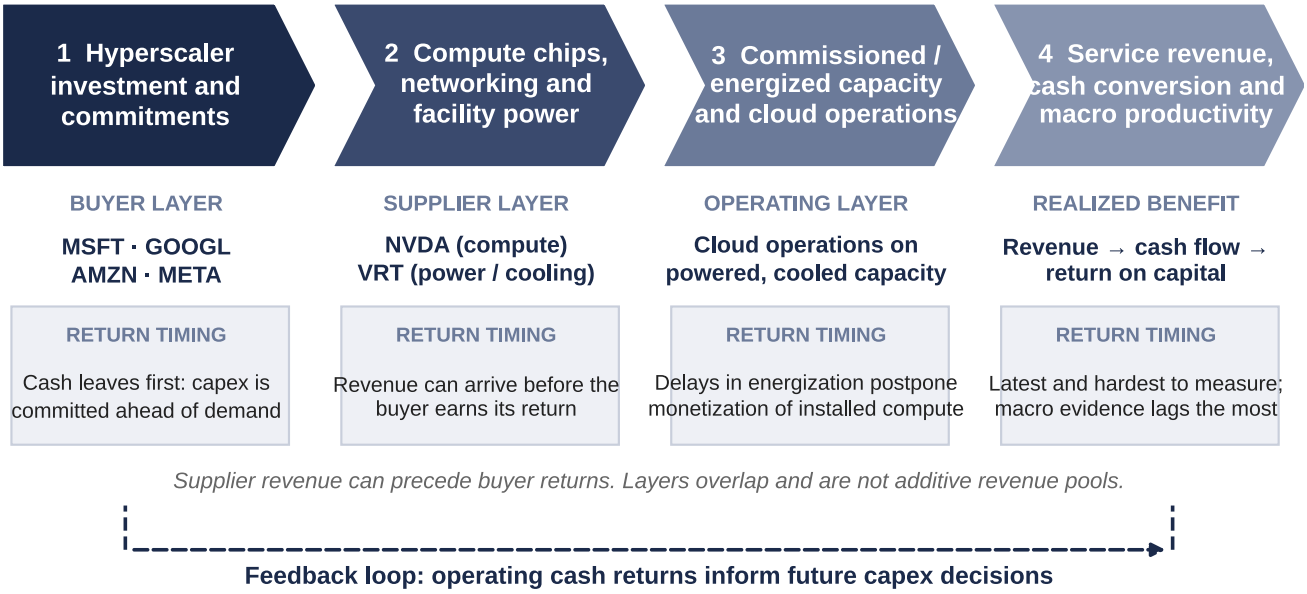
The physical buildout will help set the pace of monetization. Chips can be delivered faster than facilities can be energized. Power distribution, cooling and grid access now compete with compute availability as the practical constraints on deployment. Lawrence Berkeley National Laboratory (LBNL) estimates U.S. data centers consumed

192 TWh in 2024, roughly 4.7% of national electricity use; its 2030 reference scenario reaches 649 TWh. That scale of expansion makes **time to usable capacity** a consequential financial variable. [7–9]

The benefits are also spreading unevenly across national economies. Technology-related private investment in the United States has surged, while Taiwan and South Korea capture meaningful first-round demand through semiconductor manufacturing and components. U.S. domestic benefits are accruing through construction, integration, software and power systems, with a larger productivity payoff potentially still ahead. Official investment categories cover broader technology spending, and aggregate productivity data have yet to isolate the contribution from AI. [10–13]

Five conclusions anchor our sector view. **First**, commercial demand has moved well beyond a theoretical opportunity. **Second**, capital efficiency and cash conversion are becoming decisive measures of execution. **Third**, supplier revenue arrives earlier in the cycle than buyers' returns on deployed assets. **Fourth**, power and delivery constraints can postpone monetization despite robust order books. **Fifth**, the broader productivity upside remains substantial but has yet to emerge clearly in aggregate data. We expect the market's focus to shift from announced capacity toward utilization, incremental profit and returns on capital.

EXHIBIT 1 | The AI investment value chain and timing of economic returns



Suppliers can receive purchase orders before hyperscalers earn satisfactory asset-level returns; the layers are one value chain, not additive revenue pools.

Conceptual diagram; no numerical inputs. Illustrative.

Source: BIIRG sector synthesis.

1. THE STRUCTURE OF THE AI CAPITAL INVESTMENT CYCLE

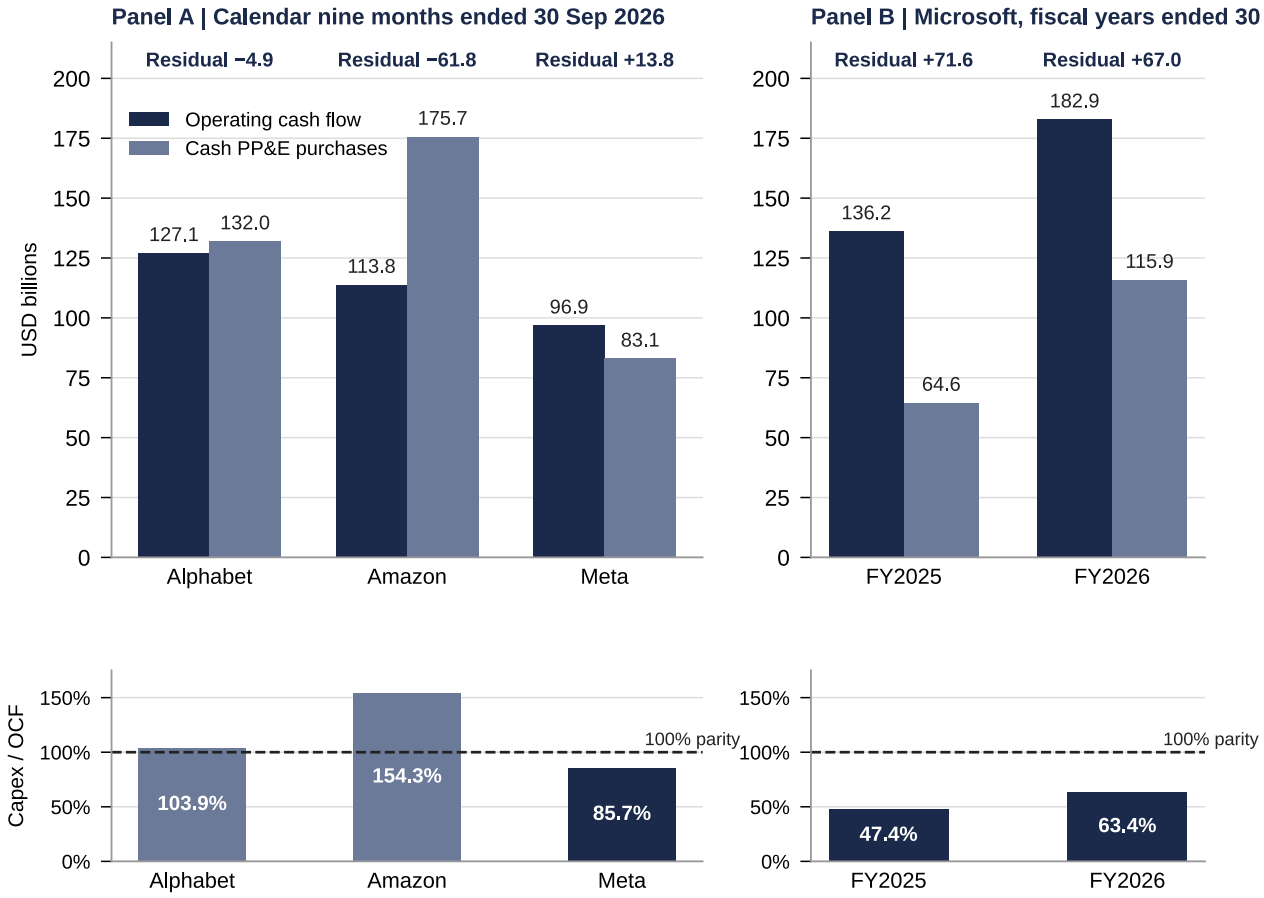
AI infrastructure is one investment cycle with sharply different economics at each stage. At the demand end, enterprises and consumers purchase applications, cloud capacity and AI-enabled services. Cloud providers and platform operators respond by commissioning data centers and purchasing accelerated computing equipment. Semiconductor companies, electrical-equipment suppliers, contractors and power providers then receive orders that ultimately support the physical deployment of capacity. Revenue reaches these businesses at different points; the spending entity bears the responsibility for earning a return on the finished asset.

For infrastructure owners, the cash burden arrives well before the earnings test. Purchases and construction consume cash immediately; depreciation flows through reported earnings over the assets' useful lives. That timing allows reported profit to rise while free cash flow contracts. Delays in commissioning extend the period between capital deployment and revenue generation, increasing the economic burden of carrying an unfinished asset.

Each link in the chain has a distinct performance test. For hyperscalers, the relevant tests are cloud demand, utilization, incremental segment profitability, maintenance and replacement capex, cash conversion and cost of capital. For semiconductor suppliers, the near-term questions are customer orders, production economics, product cycles and customer concentration. For power and thermal suppliers, order conversion, delivery schedules, margins and services attach rates matter. For the macroeconomy, the test is whether infrastructure spending translates into domestic production, higher output per hour and wider efficiency gains. Together these measures show where growth is converting into financial value.

The cycle sits inside the wider expansion of digital infrastructure. Conventional cloud workloads, enterprise storage and networking still account for part of spending and revenue. Accordingly, **company-wide capex remains company-wide capex** in our comparisons; AI-specific attribution is used only where disclosed. Cloud revenue and national-account information-processing investment provide context for demand and economic scale, respectively. [1, 2, 10–12]

EXHIBIT 2 | Reported hyperscaler capital intensity and cash generation



Similar demand backdrop, materially different cash-spending coverage.

BIIRG calculations. Residual = operating cash flow − cash PP&E purchases; not company-defined free cash flow. Amazon uses gross PP&E purchases. Microsoft's June fiscal year is shown separately from the calendar nine-month results for Alphabet, Amazon and Meta; the periods are not directly comparable. Peer annual observations for 2025 are not charted.

Source: Microsoft FY2026 Annual Report; Alphabet, Amazon and Meta Q3 2026 earnings releases and filings [1–4]; BIIRG calculations.

2. HYPERSCALER ECONOMICS: WHEN GROWTH MUST JUSTIFY INVESTMENT

Microsoft: financial coverage provides room to execute

Microsoft's fiscal 2026 results illustrate the simultaneous strength and cost of cloud expansion. Microsoft reports total revenue of $331.8 billion, up 18%, Microsoft Cloud revenue of $214.4 billion, up 27%, and Azure and other cloud services growth of 41%. Over the same period, reported cash purchases of property and equipment increased from $64.6 billion to $115.9 billion, while operating cash flow increased from $136.2 billion to $182.9 billion. On the consistent simplified definition, operating cash flow less cash capex declined from $71.6 billion to $67.0 billion. The resulting figures use cash purchases of property and equipment, rather than non-cash finance-lease additions. [1]

Microsoft has the cash generation to sustain substantial expansion. Its rising capital intensity nevertheless raises the threshold for attractive incremental returns. Capex as a share of revenue rose from 22.9% to 34.9% of annual revenue. This is a meaningful change in the financial composition of the business: a larger amount of current cash generation is committed to assets whose profitability depends on future workloads, utilization and pricing. Azure's growth confirms strong demand; the next test is how efficiently newly deployed capacity converts that demand into incremental profit. Microsoft leads this two-company comparison on cash-flow coverage, providing room to execute through a demanding investment cycle.

Alphabet: commercial momentum against higher cash absorption

Alphabet presents a somewhat different profile. Google Cloud continued to expand, with its third-quarter 2026 operating income reaching $11.1 billion versus $3.6 billion a year earlier. Segment profitability is a useful sign of commercial conversion, although the segment includes services beyond AI. [2] At the same time, its reported statements show cash PP&E purchases of $91.4 billion in calendar 2025 and $132.0 billion in the first nine months of 2026. Using operating cash flow of $127.1 billion for that nine-month period, the simplified cash-flow result is approximately negative $4.9 billion using the workbook’s rounded inputs. This simplified measure is useful for comparing cash generation and capital intensity; it should not be presented as a company-reported, all-inclusive free-cash-flow definition. [1, 2]

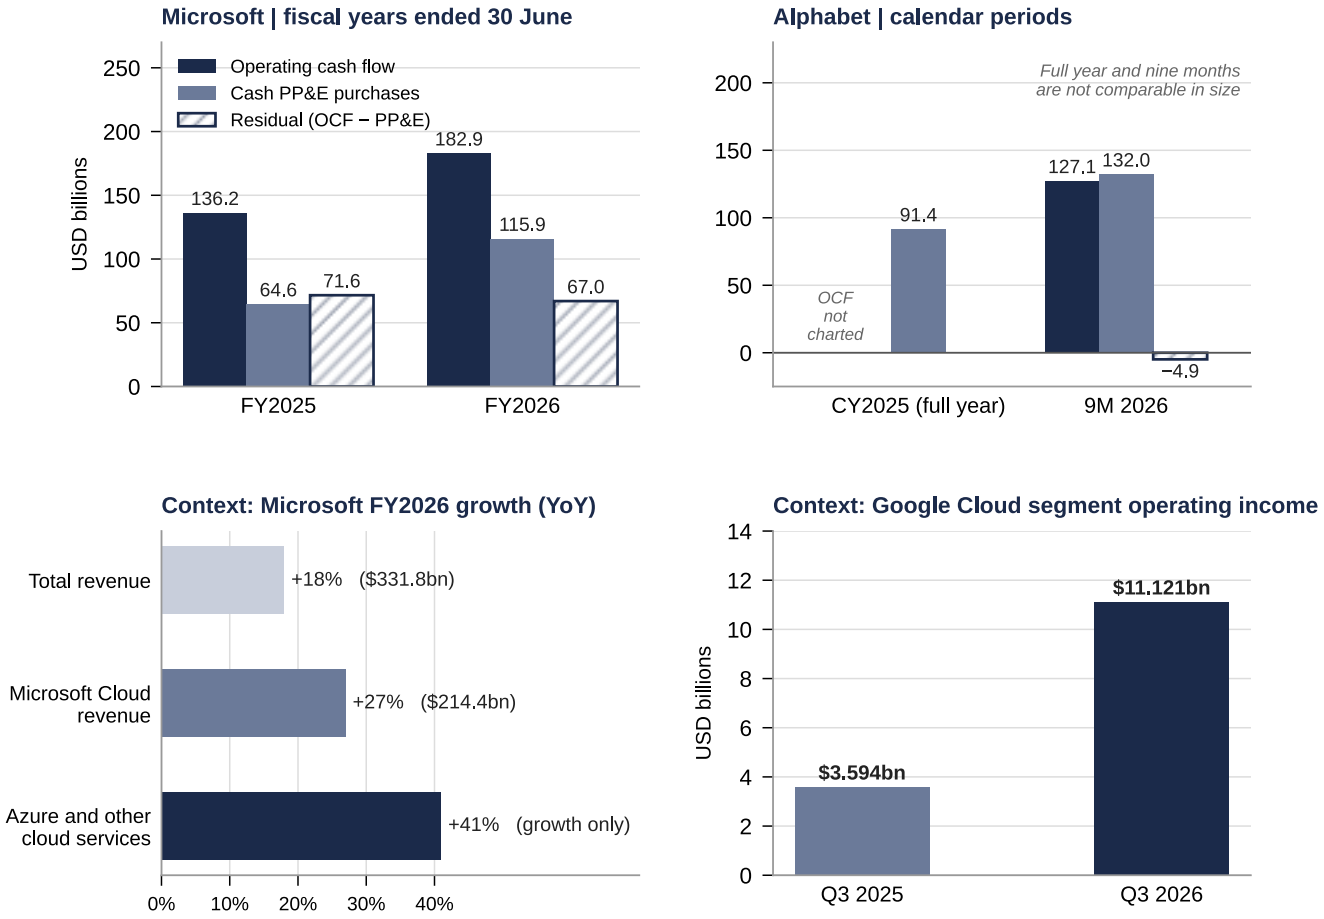
Alphabet's cloud business is delivering stronger commercial results, supported by a diversified earnings base. The pressure point is the pace of infrastructure spending: new capacity is consuming cash faster than operating profit is currently translating into cash generation. As equipment becomes operational, higher depreciation and operating costs may also affect margins. If utilization eventually increases and spending growth slows, free cash flow could recover quickly; if demand disappoints or the assets prove expensive to operate, capital intensity could remain structurally elevated.

The equity comparison: financial flexibility and operating leverage

Microsoft currently offers stronger cash-flow coverage, while Alphabet has greater near-term pressure on cash conversion alongside rapid cloud-profit growth. Both benefit from installed customer relationships, scale and diversified revenue. The most revealing test in coming quarters will be whether cloud operating profit begins to outgrow the cash required to maintain and expand infrastructure. Precise AI-only returns remain unavailable in public disclosures, making that broader measure the practical benchmark.

From an operating standpoint, Microsoft enters this phase with greater financial flexibility. Alphabet offers meaningful upside to cash conversion if rapid cloud-profit growth persists as investment normalizes. Their relative stock performance will ultimately reflect valuation and market expectations as well as these fundamentals. [1, 2]

EXHIBIT 3 | Microsoft and Alphabet: cloud growth, capex and simplified FCF



Strong cloud growth must be read alongside rising cash intensity.

BIIRG calculations. Microsoft (June fiscal year) and Alphabet (calendar periods) are separate panels; no continuous timeline is implied. Capex is company-wide and cloud is not AI-only. Alphabet 2025 operating cash flow is not charted. The Alphabet 9M 2026 residual of −$4.9bn uses rounded inputs (127.1 − 132.0). Google Cloud segment operating income is a context callout, not a capex or return calculation.

Source: Microsoft FY2026 Annual Report [1]; Alphabet Q3 2026 earnings release and FY2025 annual filing [2]; BIIRG calculations.

3. FINANCING THE BUILDOUT: LIQUIDITY, LEVERAGE AND CAPITAL ALLOCATION

The financing cycle is entering a more demanding phase. Amazon and Meta retain powerful cash-generating businesses and access to capital, yet infrastructure investment is absorbing a growing share of their financial flexibility. The consequence is visible in borrowing requirements, contractual obligations and the scope for discretionary distributions. [3, 4]

Amazon: investment is outrunning operating cash flow

For the first nine months of 2026, Amazon reported operating cash flow of $113.8 billion against gross purchases of property and equipment of $175.7 billion, equivalent to roughly 154% of operating cash flow. On that basis, the cash spending gap is approximately $61.8 billion. Amazon's disclosures also identify substantial proceeds from long-

term debt and increasing non-current borrowings, together with meaningful lease liabilities. The buildout is increasingly reliant on Amazon's balance sheet and access to financing markets. Company financing flows fund the enterprise as a whole, but their scale makes the change in funding posture unmistakable. [3]

Amazon's funding posture has important strengths. It retains substantial liquidity and a large operating cash flow base, and AWS provides a visible monetization mechanism through recurring cloud revenues. However, continued investment at a rate exceeding operating cash generation makes the duration of this gap consequential. The balance sheet buys time for infrastructure projects to reach profitable scale. It also makes their eventual returns more consequential. Amazon also discloses a substantial multi-purpose investing cash-flow line that cannot be allocated to individual AI projects from the available detail. This limits attempts to attribute the entirety of its financing activity to data centers. [3]

Meta: less cash-flow slack and more complex financing arrangements

Meta's nine-month 2026 cash property purchases amounted to approximately 86% of operating cash flow, compared with roughly 38% in 2023. The firm still covered those purchases internally over the nine months, but the third-quarter comparison was tighter, and free cash flow declined materially. Debt issuance, leases and a data-center development arrangement are also part of the funding picture. These contractual structures matter because headline debt does not capture every obligation attached to the physical infrastructure footprint. [3, 4]

The capital-allocation trade-off is especially visible in discretionary buybacks. Company disclosures show a sharp reduction in repurchases while infrastructure investment accelerated, with the dividend continuing. Meta is increasingly directing financial capacity toward infrastructure rather than repurchases. That choice preserves room to fund strategic assets, but it carries a substantial opportunity cost: every dollar committed to new facilities competes with alternative investment, shareholder distributions and balance-sheet strength. Whether that reallocation adds shareholder value depends on future project economics.

Credit capacity and the investment hurdle

The credit test centers on leverage relative to cash flow, available liquidity, the marginal cost of debt and obligations embedded in leases and partnerships. Higher market rates are an additional potential hurdle for long-duration assets. Higher yields can raise financing costs and investors' required returns even when incremental debt-interest expense remains manageable relative to current cash flow. Companies have raised spending plans despite elevated borrowing costs. For now, rates raise the required return on new projects more visibly than they restrict access to funding. [3, 4]

Amazon retains greater aggregate financing flexibility through its larger cash generation, liquidity and absence of recurring shareholder distribution commitments. Meta has stronger current capex coverage but a thinner financial cushion after debt, leases and shareholder distributions. Both can sustain investment; their vulnerabilities sit in different places.

EXHIBIT 4 | Amazon and Meta: funding coverage, liquidity and non-debt obligations

Measure (nine months ended 30 Sep 2026, USD billions)	Amazon	Meta
Operating cash flow	113.841	96.873
Gross cash PP&E purchases	175.670	83.055
Capex / operating cash flow	**154.3%**	**85.7%**
Simplified cash residual (OCF − PP&E)	**−61.829**	**+13.818**
Cash and marketable securities	114.226	82.449
Non-current long-term debt	173.125	80.922
Reading	Greater absolute cash generation and liquidity, but weaker current capex coverage.	Higher current cash coverage, but less unused financial slack and lease / JV commitments.
Obligations beyond debt	Substantial long-term debt proceeds and lease liabilities; a large multi-purpose investing cash-flow line cannot be allocated to projects.	Debt issuance, leases and a data-center development arrangement; buybacks sharply reduced, dividend continuing.

Funding coverage and financial flexibility are different dimensions; ratios alone do not indicate credit distress.

BIIRG calculations. Cash and debt columns are not all-inclusive net debt; finance leases, maturities and off-balance-sheet-like commitments complicate comparison. Financing flows cannot be assigned to specific projects. Residual is not company-defined free cash flow.

Source: Amazon Q3 2026 Results [3]; Meta Q3 2026 Earnings Release and Form 10-Q [4].

4. FROM SILICON TO ENERGIZED CAPACITY: THE PHYSICAL INFRASTRUCTURE CONSTRAINT

The limiting factor in AI deployment is increasingly the synchronization of the physical buildout. Accelerators and networking equipment operate inside facilities that require construction, electrical distribution, cooling, grid access and reliable power. The value chain is grouped into compute, facilities, power and thermal management, and grid interconnection. The practical risk is a mismatch between the speed at which computing equipment can be procured and the time required to commission usable, powered capacity. [7–9]

The spending trend is already material. U.S. construction statistics show substantial growth in data-center construction relative to the earlier decade, though data-center construction also supports workloads unrelated to AI. Electricity is a more direct physical constraint, albeit one with wide forecasting uncertainty. LBNL estimates U.S. data centers consumed 192 TWh in 2024, or approximately 4.7% of national electricity use. The LBNL model gives a 2030 reference projection of 649 TWh, or 11.8% of projected U.S. electricity use, and a broader 521–843 TWh scenario range. The 2030 figures are modeled demand scenarios and show the scale of infrastructure expansion under consideration. [7–9]

Interconnection and permitting impose a timing constraint. LBNL interconnection-queue evidence indicates that large generation projects commonly take several years to progress from an interconnection request to commercial operation. The queue data measure generation-project timelines; they serve as evidence of system-wide delivery

friction rather than a direct estimate of data-center connection waits. LBNL's modeled 148 GW of data-center interconnection capacity also reflects utilization and redundancy assumptions, with the mix of grid supply and on-site generation determining the actual build requirement. [7–9]

The investment consequences are asymmetric. For facility owners, a delay in energization postpones revenue from the installed compute fleet and may raise carrying costs. For grid-equipment and cooling suppliers, infrastructure scarcity can strengthen demand, yet production and commissioning delays can also defer revenue recognition. Some projects may be relocated to power-rich jurisdictions, supplemented with on-site generation or redesigned for greater efficiency. As a result, investment is likely to migrate toward locations and designs that can bring usable capacity online fastest.

NVIDIA: supplier exposure to computing procurement

NVIDIA monetizes the industry's demand for accelerated computing, networking and integrated systems. Its position allows it to capture revenue as customers purchase infrastructure, before hyperscalers have necessarily earned an adequate return on their completed data centers. This generates leverage to the pace of procurement, but also exposure to a relatively concentrated set of large buyers and rapid product transitions. If computing purchases run ahead of powered data-center capacity, customers could defer orders or stretch deployment schedules. [7–9]

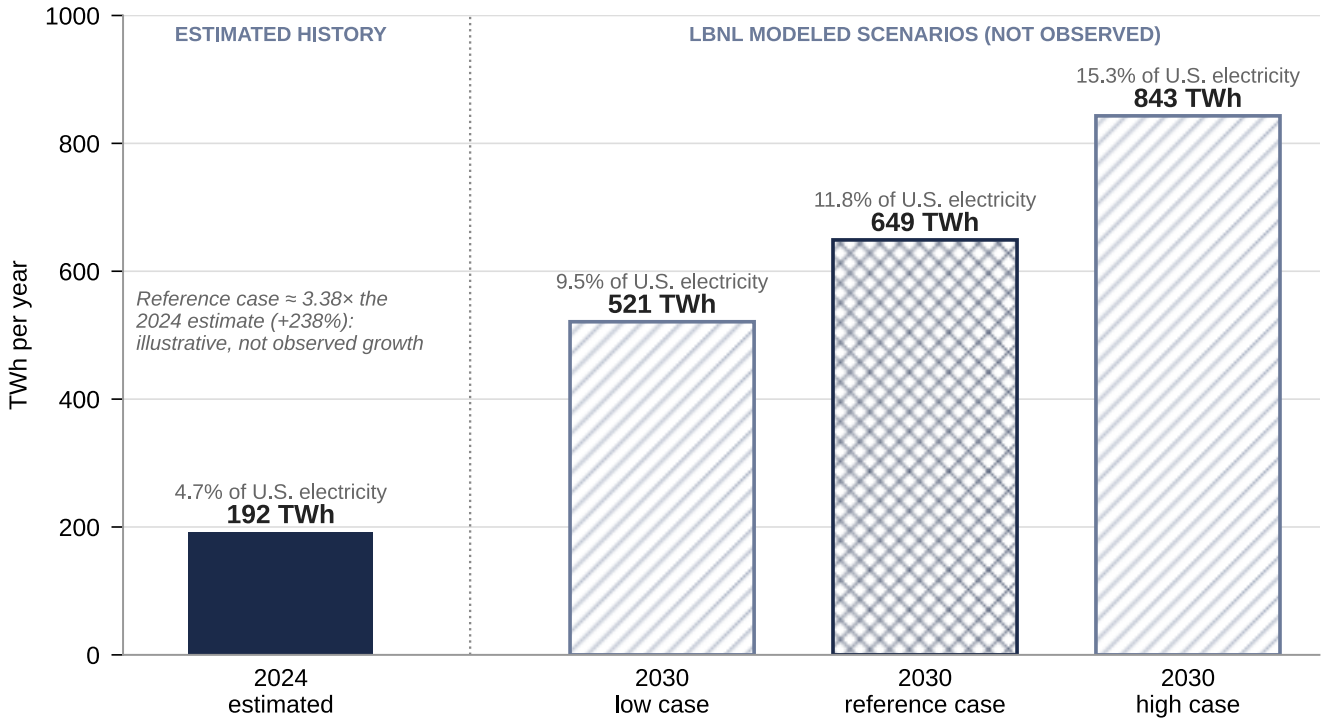
In its third quarter of fiscal 2027, reported on November 17, 2026, NVIDIA generated $115.8 billion of total revenue, including $107.6 billion from Data Center, compared with $96.2 billion and $89.0 billion respectively in the previous quarter. Those results show how quickly supplier revenues are scaling as hyperscalers purchase equipment, well before the buyers' long-term returns on that equipment become evident. [6] The economic case remains that demand for AI compute supports a major supplier opportunity, while investment cycles can be volatile and customer order timing is material.

Vertiv: power and thermal infrastructure as a scarce input

Vertiv occupies a different part of the buildout. It supplies electrical power management, thermal infrastructure and services that help make computing capacity operational. In third-quarter 2026 results, Vertiv reported $4.313 billion of revenue, up 61% from the prior year, while organic orders grew by more than 100%. Revenue timing remains sensitive to construction schedules and delivery execution. [7] The larger point is that power density and thermal requirements increase the importance of supporting infrastructure per installed computing system. [7–9]

Vertiv's growth profile may benefit from sustained facility construction, equipment replacement and the needs of higher-density racks. Order conversion and execution will determine how much of Vertiv's demand pipeline becomes reported revenue. Facility readiness, delivery schedules and customer financing remain the variables to watch. A broad hyperscaler spending slowdown would affect Vertiv as well as NVIDIA. Both remain exposed to hyperscaler spending decisions, with different sensitivities to procurement and commissioning schedules.

EXHIBIT 5 | Historical and projected U.S. data-center electricity use

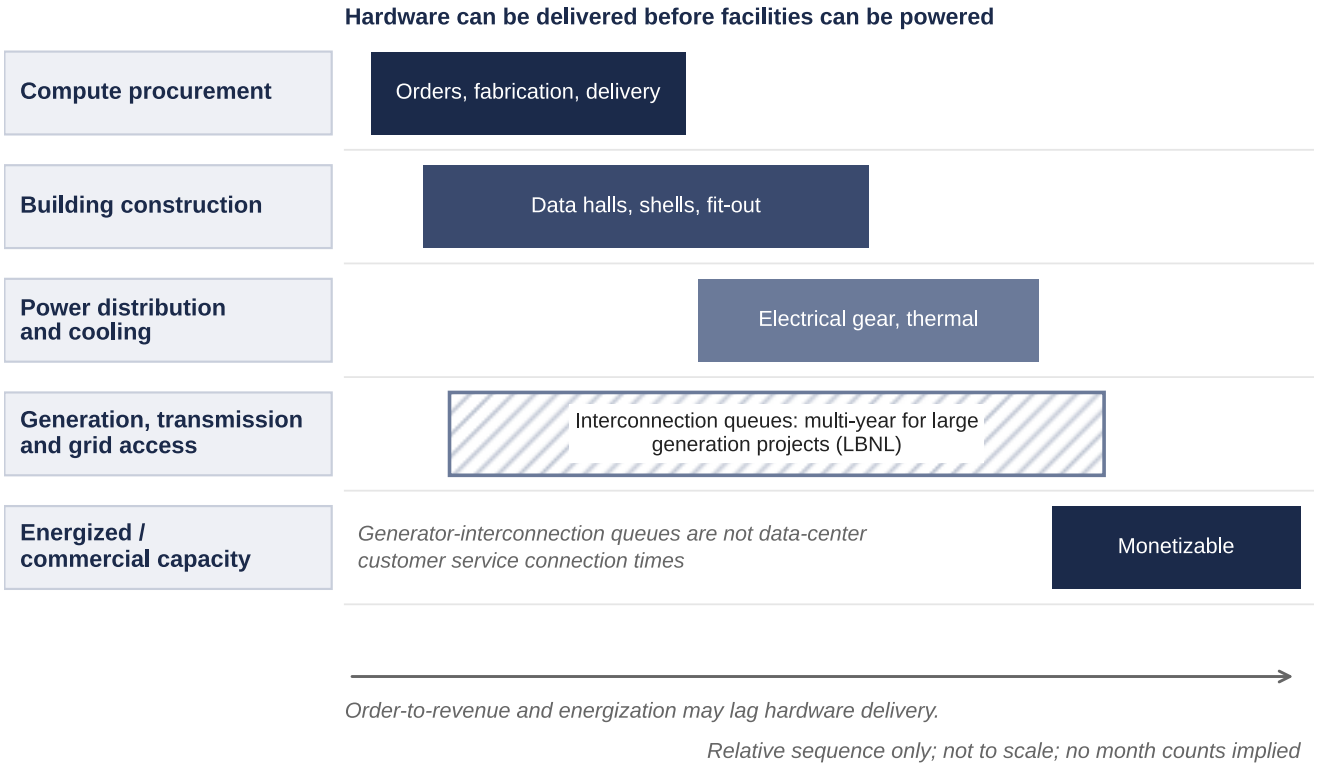


Electrical infrastructure requirements could rise rapidly by 2030, subject to substantial scenario uncertainty.

Legend: solid = estimated history (2024); hatched = LBNL modeled 2030 scenarios, not observed consumption. Shares are of total U.S. electricity consumption (2030: projected consumption). Data-center load includes non-AI workloads. The 3.38× ratio (+238%) compares the reference scenario with the 2024 estimate and is not observed growth.

Source: Lawrence Berkeley National Laboratory, United States Data Center Energy Usage Report: 2025 Update [8].

EXHIBIT 6 | Compute and power/cooling: the timing mismatch



Installed compute does not automatically equal monetizable capacity.

Conceptual and illustrative; relative timing only, no month counts. Generator-interconnection queue studies describe generation projects and should not be read as data-center utility connection times. Alternative responses include relocation to power-rich jurisdictions, on-site generation and efficiency gains.

Source: LBNL, Queued Up [9]; company disclosures [6–7]; BIIRG synthesis.

5. GLOBAL SUPPLY CHAINS AND THE DISTRIBUTION OF GROWTH

The macroeconomic impact of the AI buildout is becoming substantial, although national accounts reveal it through broader technology categories. U.S. national-account figures show that U.S. private fixed investment in information-processing equipment and software increased to $1,645.7 billion at a seasonally adjusted annual rate in the second quarter of 2026, approximately 5.05% of nominal GDP on a same-period basis. The computers-and-peripherals component reached $420.1 billion at an annual rate. These categories cover information technology beyond AI; they serve as measures of the broader investment environment. [10–13]

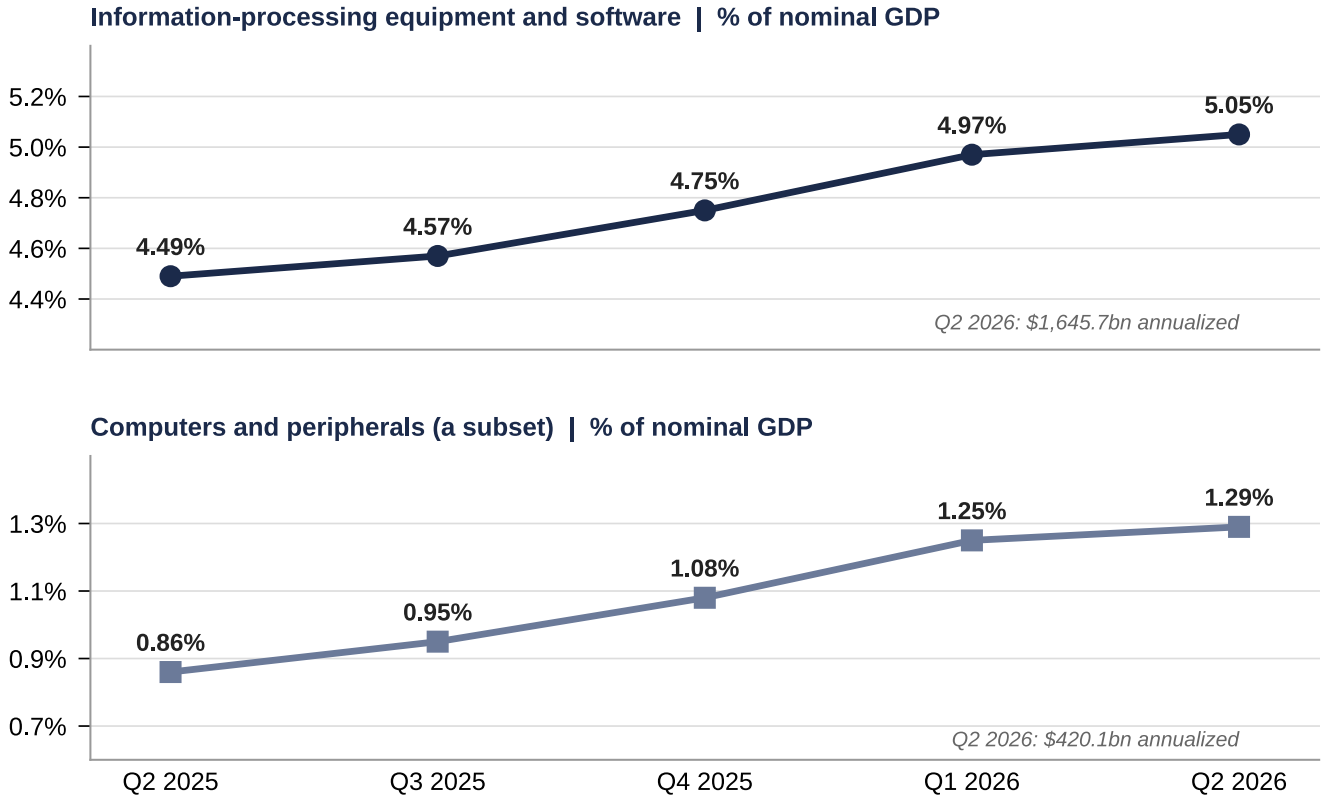
The geography of production matters. When a U.S. company purchases a server manufactured abroad, the expenditure enters the investment component of U.S. GDP accounts, but the corresponding import is subtracted. In expenditure-side accounting, the import offsets the investment entry, leaving U.S. value added dependent on domestically produced inputs and subsequent activity. Domestic value added can arise from locally produced components, construction, installation, software, electricity and subsequent operation. The profits earned by U.S.-headquartered firms are economically important but should not be confused with domestic production: profits generated abroad are not automatically additions to U.S. GDP. [10–13]

Taiwan and South Korea are important nodes in this cross-border division of production. The relevant supply-chain structure includes Taiwanese advanced chip fabrication and packaging, Korean memory and components, and the multi-country assembly and equipment supply chain as channels through which U.S. capital investment supports international manufacturing demand. Export and order data from Taiwan and South Korea point to unusually strong electronics demand, but published gross export values include imported intermediate inputs, non-AI products and demand from customers outside the United States. The export figures indicate strong electronics demand, while domestic value added and AI-specific demand require a narrower accounting treatment. [10–13]

The international structure produces both efficiency and vulnerability. Specialization gives hyperscalers access to advanced suppliers and scale, while geographic concentration exposes deployment to geopolitical disruption, trade restrictions, logistics problems and component-price shocks. Diversifying manufacturing capacity may lower concentration over time but can create higher initial costs and slower execution. The distribution of long-run returns will depend on pricing power, production location and which firms control the scarce capabilities in the chain.

EXHIBIT 7 | Information-processing investment as a share of U.S. GDP

BROAD TECHNOLOGY INVESTMENT, NOT AI-ONLY SPENDING

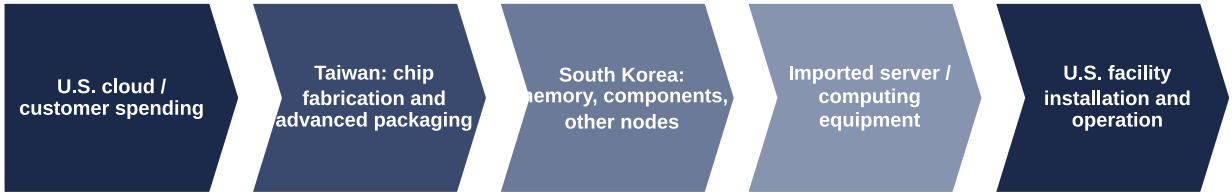


Technology capital formation is a larger share of GDP, but its AI component is not directly observed.

BIIRG calculations: investment ÷ nominal GDP, seasonally adjusted annual rates. Computers and peripherals are a subset of the broader series; the two shares are not additive.

Source: U.S. Bureau of Economic Analysis via FRED (A679RC1Q027SBEA, B935RC1Q027SBEA, GDP) [10–12]; BIIRG calculations.

EXHIBIT 8 | International AI hardware supply chain and GDP-accounting bridge



Arrows show value-chain relationships, not necessarily chronological shipping paths

ACCOUNTING BRIDGE | ILLUSTRATIVE ONLY



Domestic construction, software, installation, electricity and future services are separate additions to U.S. value added, not part of this illustration.

Spending by U.S. companies does not imply equal domestic manufacturing output.

Conceptual. The accounting bridge is illustrative only. No AI-only Taiwan or South Korea export figures are shown: gross exports include imported inputs, non-AI products and demand from outside the United States.

Source: BEA national-accounts identity; BIIRG sector supply-chain analysis [10–12].

6. PRODUCTIVITY: THE NEXT LEG OF THE RETURN CYCLE

The largest macroeconomic prize from AI is higher productivity across the wider economy. AI technologies may reduce the cost of information-intensive tasks, improve decision processes, automate components of service work and expand output from existing labor and capital. Such benefits would make the investment cycle more economically significant than a one-time construction or hardware procurement boom. Their timing and scale will depend on adoption, complementary investment and the actual tasks improved by AI.

Bureau of Labor Statistics data indicate U.S. nonfarm-business labor productivity growth of approximately 2.2% year over year in the second quarter of 2026, close to the post-2019 annual pace in the underlying research. The aggregate data have yet to identify a distinct AI-driven acceleration. Productivity statistics aggregate across industries and technologies, and gains may follow a lag as companies install systems, retrain workers and redesign processes. The concept of a productivity **J-curve** is relevant: technology investment may initially require complementary organizational effort before its benefits register in measured output per hour. [10–13]

A useful framework distinguishes three stages. **Capital formation** measures investment in physical and software assets. **Commercial monetization** measures revenue and profit generated from the services that use those assets. **Productivity realization** measures improvements in the economy's ability to produce output with labor and capital. Progress from one stage to the next is the economic conversion we expect markets to scrutinize. Current evidence is strongest on the first and selective parts of the second, and weaker on the third.

A delayed productivity response would be consistent with earlier general-purpose technologies. The useful evidence will come from adoption at scale, revenue per unit of installed capacity, segment profitability and measurable improvements in output per hour. Forecasts of the eventual addressable market will carry less weight as operating data accumulate.

7. RISKS AND THESIS BREAKERS

We identify five groups of risks. **Monetization risk** arises if AI-enabled products and cloud services do not generate sufficient incremental revenue or pricing power. **Capital efficiency risk** emerges if the useful lives, maintenance requirements or replacement rates of compute assets prove less favorable than expected. **Financing risk** increases if capex persistently outruns operating cash flow, yields rise or less visible lease and joint-venture obligations become more burdensome. **Execution risk** includes slow grid connections, construction bottlenecks and delayed commissioning. **Supply-chain risk** reflects concentrated semiconductor production, changing export rules, higher component costs and customer concentration. [1–10]

These risks interact. Physical delays can reduce utilization and postpone monetization while debt service and contractual commitments continue. Higher financing costs can increase the hurdle rate for capital-intensive projects. An inventory cycle in accelerators can amplify swings in supplier revenue. A shift toward local production may protect supply resilience but require more capital per unit of capacity. Consequently, analysis that treats company capex, semiconductor supplier growth, grid capacity and productivity as independent variables is incomplete.

There is a serious constructive counterargument. The most important infrastructure buyers have diversified cash-generating franchises, demand for cloud services is expanding in company disclosures, and AI applications may create new revenue pools rather than simply replace existing workloads. Infrastructure constraints could lead to relocation or greater investment in power equipment rather than permanent demand destruction. Moreover, general-purpose technologies have historically required complementary investment before their full productivity effects became visible. These considerations support persistence in the investment cycle even as cash generation becomes less comfortable.

The more difficult scenario is overbuilding. Investment decisions might be driven by competitive fear rather than demonstrated economics, leading to duplicated infrastructure, aggressive financing and a future period of excess capacity. Falling hardware prices can improve the economics of adoption while simultaneously compressing returns on assets and suppliers whose valuations assume persistent scarcity. Both outcomes could coexist: cheaper compute could accelerate adoption while compressing supplier margins and the value of older equipment. We see durable demand, with increasing differentiation between companies that convert it into cash returns and those that depend on another round of spending.

EXHIBIT 9 | Risk-to-economic-failure-mode matrix

Risk	Economic channel	Indicator
1 Monetization	Weak incremental revenue and pricing	Cloud operating profit vs cash investment
2 Capital efficiency	Depreciation and replacement pressure	Capex plus depreciation and utilization
3 Financing	Cash coverage and credit risk	Operating cash flow / capex, borrowing and leases
4 Physical execution	Delayed energization and construction	Project timing and supplier order conversion
5 Supply chain	Concentration, trade and component inflation	Lead times and customer dependence

Counterargument. Existing cash flows from diversified franchises, expanding cloud demand, relocation and efficiency responses, and adoption lags consistent with a normal productivity J-curve all support persistence in the investment cycle even as cash generation becomes less comfortable.

Different failure modes affect different layers of the value chain.

Conceptual taxonomy; no probabilities, risk scores or heat-map rankings. Illustrative.

Source: BIIRG synthesis [1–10].

8. SECTOR OUTLOOK AND MONITORING FRAMEWORK

Base case: elevated investment with increasingly selective returns

In our base case, cloud and AI-related workloads continue expanding, infrastructure investment remains high, and some constraints in power availability and delivery persist. Hyperscaler financial results remain mixed because the timing of cash expenditure, depreciation and monetization differs. Equipment suppliers continue to participate in the buildout, but company results diverge based on order quality, product positioning and execution. Productivity benefits emerge gradually and remain difficult to isolate at the aggregate level. This is our central qualitative scenario.

Upside: utilization and monetization accelerate

A stronger outcome would involve sustained demand growth, faster time-to-power, rising infrastructure utilization and a more favorable relationship between incremental cloud operating profits and capex. Free cash flow could improve as newly installed assets reach productive scale, while AI-enabled business processes begin to yield measurable operating efficiencies. The broadest market beneficiaries would likely include both infrastructure suppliers and the platforms converting capacity into recurring profitable services. The key signal would be rising returns on deployed capital alongside sustained demand.

Downside: slower monetization meets fixed investment commitments

A weaker outcome would combine softer AI workload demand, falling utilization or pricing, higher asset replacement requirements and continued cash spending against precommitted projects. Hyperscaler free cash flow would remain under pressure, with greater reliance on credit markets and reduced flexibility for alternative capital

allocation. Supplier growth could slow sharply as customers defer new orders, while equipment and grid projects already underway might remain costly to complete. Aggregate productivity gains would also be delayed. Even useful technology can disappoint investors when capacity is built ahead of economically realizable demand.

The indicators that will decide the thesis

We recommend a deliberately limited dashboard. For hyperscalers: (1) cash capex relative to operating cash flow; (2) cloud growth and operating profit; (3) depreciation, margin and free-cash-flow trends; and (4) financing commitments including debt and leases. For suppliers: (5) order growth, backlog conversion and customer concentration; and (6) equipment lead times and deployment schedules. For the physical economy: (7) power availability, grid connection times and data-center utilization. For the macroeconomy: (8) technology investment relative to GDP, import content, and productivity indicators. Read together, these indicators capture the progression from capital deployment to commercial and economic returns. [1–10]

EXHIBIT 10 | What would validate the cycle? Sector monitoring dashboard

1 HYPERSCALER	2 HYPERSCALER	3 HYPERSCALER	4 HYPERSCALER
Cash PP&E / OCF	**Cloud segment operating profit**	**Simplified residual (OCF – PP&E)**	**Debt / lease obligations relative to cash generation**
▼ **Lower, or improving coverage**	▲ **Sustained growth**	▲ **Improvement**	► **Stable**
Source: Company SEC filings	Source: Company SEC filings	Source: Company SEC filings	Source: Company SEC filings

5 SUPPLIER	6 PHYSICAL	7 MACRO	8 MACRO
Supplier order-to-sales conversion	**Power availability / interconnection delays**	**Technology fixed investment and import content**	**Broad productivity**
▲ **Improving**	▼ **Shortening**	● **Contextual; no mechanically positive direction**	▲ **Sustained improvement**
Source: NVIDIA and Vertiv filings	Source: LBNL; EIA	Source: BEA; FRED	Source: BLS; FRED

Validation requires evidence across cash returns, deployment readiness and macro benefits.

Directions only: no absolute thresholds, probabilities or weighted composite scores. Illustrative.

Source: Company SEC filings [1–7]; LBNL and EIA [8–9]; BEA and FRED [10–12]; BLS [13].

CONCLUSION

AI infrastructure has entered a period of sustained industrial-scale investment. Demand is visible in cloud revenue, semiconductor orders, data-center construction and the power systems needed to commission new facilities. The market's next test will be how quickly this installed capacity generates profits and productivity sufficient to justify its full economic cost.

We remain constructive on additional compute and supporting infrastructure, with greater preference for durable monetization, disciplined capital allocation and reliable project execution. We expect competition for economic returns to sharpen as the buildout matures. Firms with strong financial coverage, durable monetization, delivery

capacity and disciplined investment decisions should be better positioned than businesses relying principally on forecasts of future demand.

For investors, the next phase should be approached through **economic conversion**: the rate at which commitments become operating assets, operating assets become revenue, revenue becomes cash flow, and cash flow becomes an attractive return on capital. The infrastructure supercycle is underway. Its winners will be determined by the return cycle that follows.

RESEARCH METHODOLOGY AND MEASUREMENT CONVENTIONS

The report combines company financial disclosures, public macroeconomic statistics and industry research available through November 23, 2026. Company examples are selected to demonstrate distinct positions in the AI infrastructure value chain rather than to constitute a comprehensive equity coverage universe. Reported financial statements, corporate guidance, modeled industry scenarios and BIIRG calculations are distinct evidence categories. Whenever these measures appear together, the unit and reporting period should be shown explicitly.

Cash-spending coverage. For cross-company discussion, cash capital expenditure means cash purchases or additions of property and equipment. The simplified cash residual is operating cash flow less cash PP&E purchases. This is not necessarily the same as company-reported free cash flow: definitions may include proceeds, incentives, payments on finance leases or other adjustments. In the figures, Microsoft's fiscal year ends June 30 and is intentionally presented separately from calendar nine-month 2026 results for Alphabet, Amazon and Meta. The three companies' financial periods are not equivalent to Microsoft's annual period. Cash coverage also measures *funding capacity*, not return on invested capital.

AI attribution. Hyperscaler capex encompasses non-AI cloud facilities and other investments; cloud segment revenue includes conventional enterprise services; data center electricity supports non-AI computing; national-account information-processing investment encompasses non-AI software, communications equipment and hardware. None is presented as a pure measure of AI expenditure or economic output. We do not add supplier revenues to customer capex when describing sector-wide investment, which would double-count transactions along the value chain.

Scenario interpretation. The baseline, upside and downside frameworks are qualitative illustrations, not probability-weighted forecasts. The LBNL 2030 electricity values are modeled scenarios and must not be labeled observed electricity consumption. The grid-connection sequence in Exhibit 6 is conceptual and is not a measurement of project-level completion times.

REFERENCES AND DATA SOURCES

1. **Microsoft Corporation.** *FY2026 Annual Report*, fiscal year ended June 30, 2026. Revenue, Microsoft Cloud growth, operating cash flow and additions to property and equipment. https://www.microsoft.com/investor/reports/ar26/index.html
2. **Alphabet Inc.** *Third-Quarter 2026 Earnings Release* and FY2025 annual filing. Google Cloud segment revenue and operating profit, consolidated cash flows, and equipment investment. https://www.sec.gov/Archives/edgar/data/1652044/000165204426000080/googexhibit991q32026.htm ; https://www.sec.gov/Archives/edgar/data/1652044/000165204426000018/goog-20251231.htm
3. **Amazon.com, Inc.** *Third-Quarter 2026 Results*, October 29, 2026. Operating cash flow, gross PP&E purchases, debt proceeds and company-defined free cash flow. https://ir.aboutamazon.com/news-release/news-release-details/2026/Amazon-com-Announces-Third-Quarter-

Results/default.aspx

4. **Meta Platforms, Inc.** *Third-Quarter 2026 Earnings Release* and *Form 10-Q*. Operating cash flow, capital purchases, finance-lease treatment, guidance, and financing arrangements. https://investor.atmeta.com/investor-news/press-release-details/2026/Meta-Reports-Third-Quarter-2026-Results/default.aspx ; https://www.sec.gov/Archives/edgar/data/1326801/000162828026066773/meta-20260930.htm
5. **NVIDIA Corporation.** *Second-Quarter Fiscal 2027 Results*, August 26, 2026. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Second-Quarter-Fiscal-2027/
6. **NVIDIA Corporation.** *Third-Quarter Fiscal 2027 Results*, November 17, 2026. https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Third-Quarter-Fiscal-2027/default.aspx
7. **Vertiv Holdings Co.** *Third-Quarter 2026 Results*, October 28, 2026. https://www.sec.gov/Archives/edgar/data/1674101/000162828026067707/q32026exhibit991vrt10282026.htm
8. **Lawrence Berkeley National Laboratory.** *United States Data Center Energy Usage Report: 2025 Update*, 2026. Observed/estimated historical electricity consumption and 2030 reference, low and high cases. https://www.energy.gov/documents/united-states-data-center-energy-usage-report-2025-update
9. **Lawrence Berkeley National Laboratory.** *Queued Up* interconnection studies. These describe generator-project queues and should not be confused with direct data-center utility service waiting times. https://emp.lbl.gov/queues
10. **U.S. Bureau of Economic Analysis**, via **FRED**. Private fixed investment in information-processing equipment and software, series A679RC1Q027SBEA. https://fred.stlouisfed.org/series/A679RC1Q027SBEA
11. **U.S. Bureau of Economic Analysis**, via **FRED**. Investment in computers and peripheral equipment, series B935RC1Q027SBEA. https://fred.stlouisfed.org/series/B935RC1Q027SBEA
12. **U.S. Bureau of Economic Analysis**, via **FRED**. Gross domestic product, series GDP. https://fred.stlouisfed.org/series/GDP
13. **U.S. Bureau of Labor Statistics.** *Productivity and Costs*, second quarter 2026, and subsequent productivity releases. https://www.bls.gov/productivity/
14. **Brynjolfsson, Erik, Daniel Rock, and Chad Syverson.** "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *American Economic Journal: Macroeconomics* 13(1), 2021. https://www.aeaweb.org/articles?id=10.1257/mac.20180386
15. **David, Paul A.** "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80(2), 1990. https://www.jstor.org/stable/2006600

Research attribution. This report was authored jointly by Oleg Godun, Chief Investment Officer, and analysts Mohammad Zaman, Mahdeer Rahman, Louise Phillipe Sehi and Syed Muhammad Hashmi of the BIIRG Macroeconomic Research Division, using its November 2026 collaborative research project: the equity markets and earnings workstream (Microsoft and Alphabet), the financing and credit workstream (Amazon and Meta), the infrastructure workstream (NVIDIA and Vertiv), and the macroeconomic growth and international spillovers workstream. Analytical interpretations are BIIRG's; company disclosures and official statistics are identified separately above.

DISCLOSURES AND IMPORTANT INFORMATION

FOR RESEARCH PURPOSES ONLY

This publication is independent educational and non-commercial research prepared by the BIIRG Macroeconomic Research Division. It is provided for informational and educational purposes only. It does not constitute investment, legal, tax or accounting advice; an individualized or personal recommendation; an offer; or a solicitation to buy, sell or hold any security or financial instrument. It contains no price targets and no buy, sell or hold ratings.

No advisory relationship. BIIRG is a student research organization. It does not provide investment advisory, brokerage or other regulated financial services and does not manage client assets. Reading this report does not create an advisory or other professional relationship. Readers should consult a qualified professional before making any financial decision.

Companies discussed. Microsoft, Alphabet, Amazon, Meta, NVIDIA and Vertiv are discussed solely to illustrate distinct positions in the AI infrastructure value chain. They do not form a comprehensive coverage universe, and their mention is not a recommendation to transact in their securities.

Sources and data. The report combines company financial disclosures, public macroeconomic statistics and industry research available through 23 November 2026. Third-party information has been drawn from sources believed to be reliable but has not been independently audited, and BIIRG does not guarantee its accuracy, completeness or timeliness. Figures may be revised after publication. Certain figures are calculated or normalized by BIIRG and may differ from company-reported results, guidance or third-party estimates. Rounding may produce minor inconsistencies.

Evidence categories. *Observed:* company cash-flow and revenue disclosures where the company reports the period. *Calculated:* ratios, cash residuals and GDP-share measures derived from source numbers ("BIIRG calculations"). *Estimated:* the 2024 LBNL electricity figure, an analytical estimate rather than a metered census. *Scenario:* LBNL 2030 electricity cases and the report's qualitative outlooks, never to be read as observed history. *Illustrative:* conceptual flows, the accounting example and scenario mechanisms, with no implied statistical calibration.

Forward-looking statements. Base, upside and downside outlooks are qualitative scenarios, not calibrated or probability-weighted forecasts. Forward-looking assessments are subject to substantial uncertainty, and actual outcomes may differ materially.

AI attribution. Hyperscaler capex, cloud segment revenue, data-center electricity use and national-account information-processing investment include non-AI components. None is presented as a pure measure of AI expenditure or economic output.

Independence of views. The views expressed are those of the authors and do not represent the views of Baruch College or the City University of New York. Any relevant author holdings or conflicts are to be disclosed by the authors.