Techstrong
a Futurum company

SEPTEMBER 2026

A TECHSTRONG SPECIAL REPORT

AI Is Now a Pillar of the American Economy.

The Business Model Is Still Unproven.

America's growth, infrastructure, energy system, financial markets and industrial policy increasingly depend on AI investment. The applications have not yet proved they can carry the cost.

BY ALAN SHIMEL

READER'S GUIDE

America's growing dependence on AI spending.

America is receiving the benefits of building AI before proving that its use can generate enough value to pay for the buildout.

BY ALAN SHIMEL · SEPTEMBER 2026

AI investment, enterprise adoption and the economics of the infrastructure buildout. Chart definitions, survey comparisons and methodological limitations appear alongside the figures and in the closing notes.

THE OPENING MARKER

America's investment priorities are changing.

America's investment priorities are changing.

In the second quarter of 2026, inflation-adjusted investment in information-processing equipment reached an annual rate of approximately $752.1 billion. Real private residential investment came in at approximately $748 billion. By that measure, the United States is now investing more in computers and related equipment than it is in housing.

The difference is only about $4 billion, small enough that a future revision could temporarily reverse the order. Information-processing equipment includes computers, communications equipment and other hardware. It is not a pure measure of artificial intelligence, and it does not include the buildings we call data centers. In current dollars, housing remains larger, approximately $1.186 trillion compared with $794.9 billion for information-processing equipment.

Housing also has its own problems. High mortgage rates, affordability, zoning restrictions and construction costs have weighed on residential investment. AI did not cause America's housing shortage. Every dollar invested in a server was not taken from a new home.

Real investment in information-processing equipment increased roughly 22% from the second quarter of 2025. Real residential investment fell about 4%. Computing investment is surging while housing struggles.

The shift extends beyond housing to business investment, nonresidential construction, manufacturing, electricity demand, corporate borrowing, stock-market earnings and federal policy.

AI is now a pillar of the American economy.

The applications and services running on this infrastructure have yet to demonstrate returns commensurate with the cost of building it.

AI's economic contribution today comes substantially from construction and capital formation. Its future justification depends on productivity, applications and services. America is receiving the benefits of building AI before proving that its use can generate enough value to pay for the buildout.

The wager may pay off spectacularly. America is already depending on it.

THE HOUSING CROSSOVER

Computing equipment edges past housing.

Quarterly investment, billions of chained 2017 dollars, seasonally adjusted annual rate.[1 2]

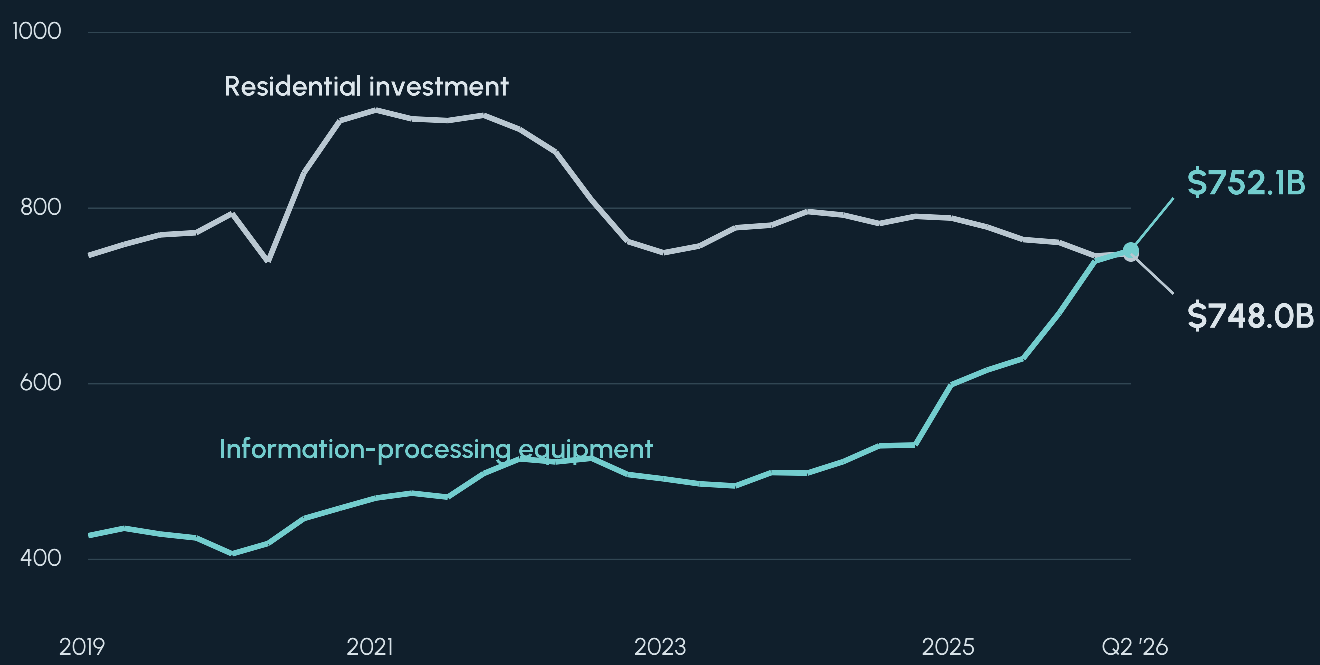

What the Housing Crossover Does and Does Not Mean

The crossover compares two inflation-adjusted annualized series in the national accounts. It shows real information-processing equipment investment at approximately $752.1 billion and real residential investment at approximately $748 billion in Q2 2026.

It does not mean that America spends more current dollars on data centers than homes. Nominal residential investment remains larger. It does not mean all information-processing equipment is AI. It also does not include the buildings, power plants and transmission systems supporting data centers.

Investment in computing equipment has accelerated while housing investment has weakened. The two categories can also compete at the margin for electricians, engineers, construction workers, transformers, steel, power, land and capital. The crossover shows a changing investment mix. It does not establish that AI caused the housing shortage.

1. https://fred.stlouisfed.org/data/Y034RX1Q020SBEA.txt
2. https://fred.stlouisfed.org/data/PRFIC1.txt

THE SEVEN PILLARS

What Makes Something an Economic Pillar?

Consumer spending still accounts for roughly two-thirds of the American economy. AI does not approach that scale. Nor is it larger than healthcare, housing or the services economy. Calling AI a pillar is not the same as calling it the entire building.

An economic pillar affects several systems when it expands or contracts. AI now does that.

It now influences at least seven important areas: business fixed investment, construction, semiconductor and equipment demand, electricity generation, corporate finance, equity-market performance and federal industrial policy.

The Treasury Department's review of the second quarter of 2026 shows how important capital spending has become. Real GDP grew at a 1.5% annual rate. Business fixed investment added 1.2 percentage points to that growth after expanding 8.4%. Equipment investment grew 15.2%, with strong spending on information-processing equipment. Data-center construction continued growing at a double-digit rate even as structures investment declined overall.

Consumer spending was still the largest contributor, adding 2.1 percentage points. Net exports, inventories and government spending pulled the headline number back down. AI did not produce all the quarter's growth. But without strong business investment, the result would have looked materially weaker.

Data-center construction, excluding the machines installed inside, reached an annual rate of approximately $75 billion in July 2026, up 57% from the previous year. An industry analysis found that data centers accounted for all the month's increase in nonresidential construction.

The Federal Reserve's August 2026 Beige Book described an economy in which residential construction declined while nonresidential activity was heavily concentrated in data-center projects. Manufacturing demand was supported by data centers and defense. Skilled construction and technical workers were difficult to find. In Cleveland, data centers and defense drove robust manufacturing demand while consumer spending declined for a fourth consecutive period.

AI has also become energy policy. The International Energy Agency expects global data-center electricity consumption to more than double by 2030. The United States accounts for the largest share of the increase, with data centers responsible for nearly half of projected U.S. electricity-demand growth through the end of the decade.

After years of relatively flat electricity use, AI is helping restart investment in generation, transmission, natural gas, nuclear power, turbines, transformers and grid equipment. The digital buildout is becoming a physical-energy buildout.

AI investment also supports market earnings. Goldman Sachs has estimated that it could generate approximately 40% of S&P 500 earnings growth in 2026, with the largest cloud providers planning hundreds of billions of dollars in capital spending.

The AI investment cycle supports growth across enough industries that a slowdown would hurt well beyond the technology sector.

AN INTERCONNECTED ECONOMY

One investment cycle. Seven systems exposed.

01 **Business investment**
Capital formation and equipment spending

02 **Construction**
Data centers and their physical footprint

03 **Semiconductors**
GPUs, custom silicon, CPUs and memory

04 **Electricity**
Generation, transmission and grid equipment

05 **Credit**
Bonds, leases, project finance and private lenders

06 **Equity markets**
Earnings expectations and market valuations

07 **Federal policy**
Permitting, exports and national-security strategy

Calling AI a pillar is not the same as calling it the entire building.

Conceptual map of the economic channels affected by AI investment and adoption, not a quantitative allocation of GDP.

THE INGREDIENTS

The Buildout Has Become the Business Model

Companies supplying AI infrastructure already have a working business model.

Chipmakers sell GPUs, CPUs, custom accelerators and memory. Systems companies sell servers, racks and integrated computing platforms. Networking vendors sell switches, optics and interconnects. Data-center developers sell space, power and cooling. Utilities sell electricity. Cloud companies rent compute. Model developers buy training capacity. Banks, bondholders and private-credit funds finance the machinery around them.

Futurum's Data Center Semiconductors Market Model, 1H 2026 estimates the global market for data-center GPUs, custom accelerators, CPUs and off-chip memory at $241 billion in 2025. The base forecast reaches $1.213 trillion by 2030. The bull case reaches $2.2 trillion. Even the bear case reaches $747 billion.

The United States accounted for an estimated $109.5 billion in 2025, or 45.4% of the global market. Futurum's base case puts the U.S. market at $477.9 billion by 2030.

Hyperscalers represented $136.1 billion, or 56.5%, of 2025 data-center semiconductor spending. Nvidia held approximately 95.5% of the merchant data-center GPU market by revenue in the fourth quarter of 2025. That figure excludes custom silicon such as Google's TPU, Amazon's Trainium, Broadcom-designed accelerators and Huawei's Ascend products. Futurum sizes the 2025 custom-accelerator market at another $37.4 billion.

These figures describe semiconductor vendor revenue. They do not include complete server systems, networking, cooling, software, power generation or construction. They also do not tell us whether the buyers of those chips will earn an adequate return.

The $1.213 trillion semiconductor forecast measures expected vendor revenue, not the incremental value enterprises will get from the equipment. It tells us nothing about utilization, pricing power, replacement cycles or return on invested capital.

Futurum's $747 billion bear case and $2.2 trillion bull case differ by almost threefold. All three scenarios assume a large, growing AI market but imply very different infrastructure requirements.

Demand for chips and data centers can remain strong without every investment earning its cost of capital.

FUTURUM SEMICONDUCTOR SCENARIOS

2030 forecasts range from $747B to $2.2T.

Global data-center semiconductor vendor revenue. 2025 estimate and 2030 scenario endpoints; no annual path is implied.

$241B COMMON 2025 STARTING POINT

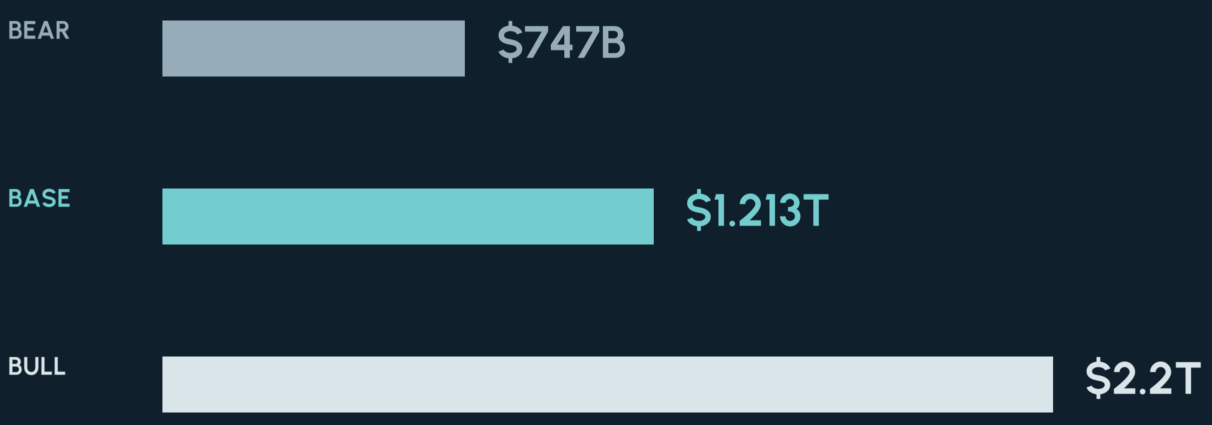

2030 SCENARIOS

Nvidia's 95% Share, and What It Excludes

Futurum estimates that Nvidia held 95.5% of the merchant data-center GPU market by revenue in CY25Q4. The denominator is Nvidia plus AMD. It does not include custom accelerators such as Google TPU, Amazon Trainium, Microsoft Maia, Huawei Ascend or ASICs designed by Broadcom and others.

Those custom accelerators represented an additional $37.4 billion market in 2025. Nvidia's merchant-GPU share should not be confused with its share of every accelerator deployed for AI. Hyperscalers use custom silicon to lower costs, reduce dependence and optimize inference.

Source: Futurum Data Center Semiconductors Market Model, 1H 2026. The chart compares the 2025 estimate with 2030 scenario endpoints; intervening annual paths are not shown.

THE SPENDING TEST

The Investment Is Ahead of the Return

Enterprise technology budgets show how quickly AI commitments are growing.

Futurum's July 2026 CIO and Technology Buyers Decision Maker Panel surveyed 1,636 respondents. It asked how actual AI spending during the previous six months compared with plan.

A total of 35.6% said spending was moderately above plan. Another 11.3% said it was substantially above plan. Combined, 46.9% of organizations were spending more than expected. Only 5.6% were below plan.

Among Fortune 500 respondents, 51.3% were over budget.

Among the 767 organizations exceeding plan, 47.6% sought supplemental approval. Another 43.3% accepted the overrun and planned to reconcile it during the next budget cycle. Nearly four in ten, 38.7%, reallocated money from another portion of IT. Another 23.1% took money from a non-IT business unit. Only 17.2% paused or reduced AI initiatives.

Faced with an overrun, organizations were more than twice as likely to cut another IT category as to reduce the AI program causing it.

Among the 297 respondents taking money from IT, 60.9% reduced contractors and consultants. Nearly half, 49.5%, cut legacy-infrastructure modernization. Another 36.4% reduced non-AI software and SaaS, while 27.9% reduced internal headcount.

Within the affected software budgets, IT operations and monitoring lost funding at 47.2% of organizations. Productivity and collaboration tools were cut by 46.3%. Business intelligence and analytics were reduced by 38.9%. Communications, ERP, CRM and even cybersecurity also surrendered dollars.

Many companies are paying for AI by cutting elsewhere in their technology budgets.

Cuts to legacy modernization, clean data, observability, cybersecurity, integration and experienced people can leave AI projects without the support they need. These capabilities make production AI reliable and useful.

Management teams appear more willing to seek additional money, postpone reconciliation and cut other priorities than to slow AI. Those decisions make AI spending harder to interrupt before its return is established.

WHERE THE MONEY COMES FROM

AI overruns draw on the rest of the business.

Overlapping survey responses, not additive cash flows. Bar length represents the share of each stated respondent base.

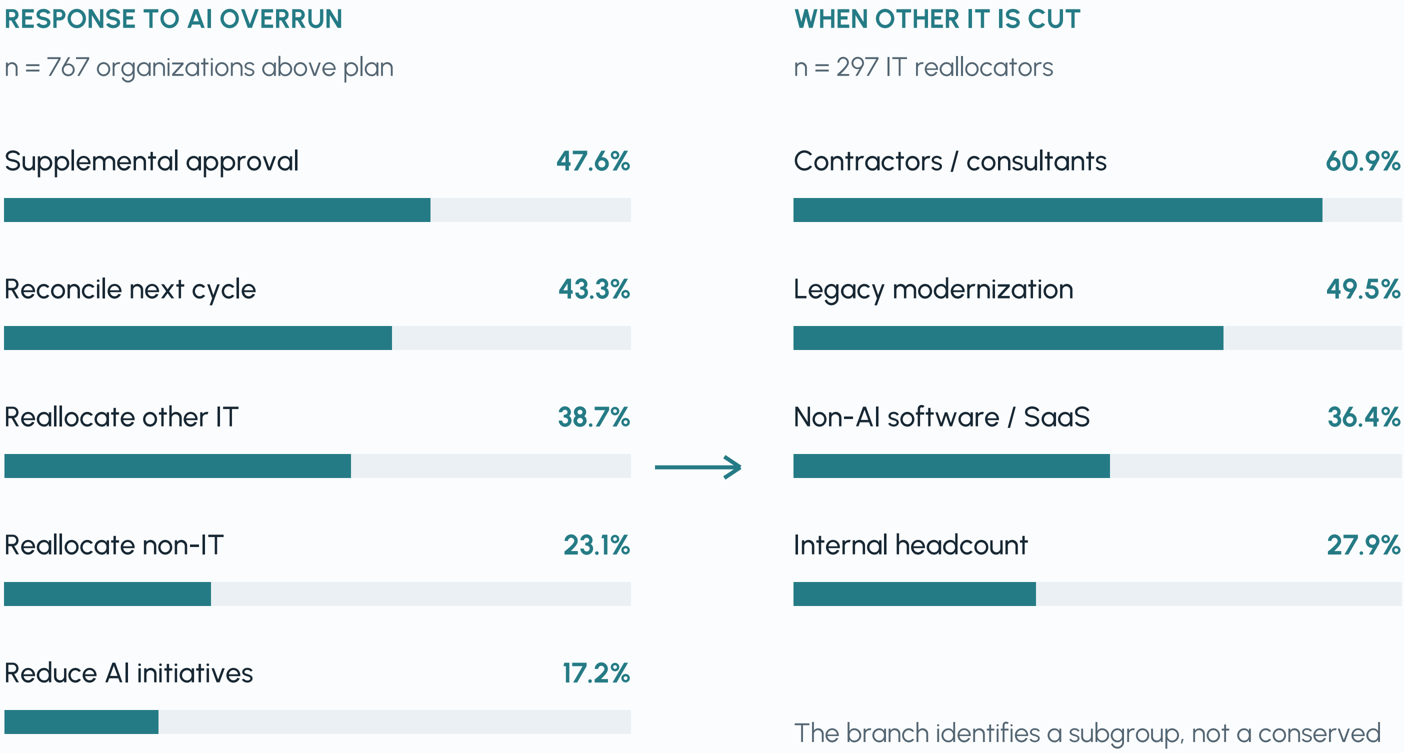

The branch identifies a subgroup, not a conserved flow of dollars.

Where the Money Comes From When AI Exceeds Budget

Among organizations exceeding their planned AI spending, 47.6% sought supplemental approval, 43.3% accepted the overrun for later reconciliation and 38.7% moved money from another IT category. Only 17.2% reduced their AI initiatives.

When IT absorbed the cut, contractors, modernization, software and internal headcount were the primary sources. Some AI spending therefore replaces other technology spending rather than adding to it. Organizations risk cutting the capabilities they need to put AI into production.

Source: Futurum CIO and Technology Buyers Decision Maker Panel, July 2026. Multiple responses can overlap within and across categories.

THE OPERATING TEST

Adoption Is Rising Faster Than Proven ROI

Enterprise AI adoption continues to grow.

From July 2025 to July 2026, the share of organizations using AI to enhance workforce productivity increased from 67.5% to 75.1%. AI-driven analytics rose from 55.4% to 67.8%. The share using AI to transform the business model increased from 40.9% to 48.8%.

The workforce effects are also becoming more explicit. Organizations using AI to limit future headcount increased from 21.3% to 34%. Those using it to reduce headcount strategically increased from 14.3% to 20%. The percentage not using AI at all fell from 10.1% to 4.1%.

Widespread adoption has yet to establish the returns.

Futurum's AI Platforms Decision Maker Survey, 1H 2026 measured adoption maturity among 820 respondents. Of those, 10.2% remained at Awareness and 21.7% at Experimentation. The largest group, 30.5%, described itself as being at Optimization. Another 24.3% had reached Standardization. Only 13.3% reported Transformation.

Most organizations have moved beyond basic experimentation, but only about one in eight says AI has transformed the organization.

Production use is broad and increasingly multi-model. Among 736 organizations with generative models in production, 63.9% used OpenAI GPT-4 or 4o, 62.1% used Azure OpenAI and 53.9% used Gemini. Roughly a quarter used Llama, small language models, Claude or domain and sovereign models. DeepSeek was already in production at 19.8% of respondents, while Mistral reached 15.6%. Multiple responses were allowed, reflecting model portfolios rather than winner-take-all selection.

For AI developed in-house, 35.5% reported ROI in pilots that had not yet scaled. Another 29.5% had adopted AI but had not realized ROI. Only 15.8% reported sustained ROI at scale.

For purchased AI, 40.3% reported pilot ROI without scale. Another 33% had adopted the technology without realizing ROI. Only 12.6% reported sustained ROI at scale.

Pilot results have not yet translated into sustained returns for most respondents.

AI spending exceeds plan at almost half of organizations. Sustained scaled ROI remains in the mid-teens or lower.

Futurum's trend data shows gradual improvement in ROI. Infrastructure investment still runs well ahead of the evidence that applications can pay for it.

THE BUSINESS-MODEL GAP

Spending exceeds plan. Scaled ROI remains uncommon.

Different questions and survey populations. These measures are not a conversion funnel or matched-cohort comparison.

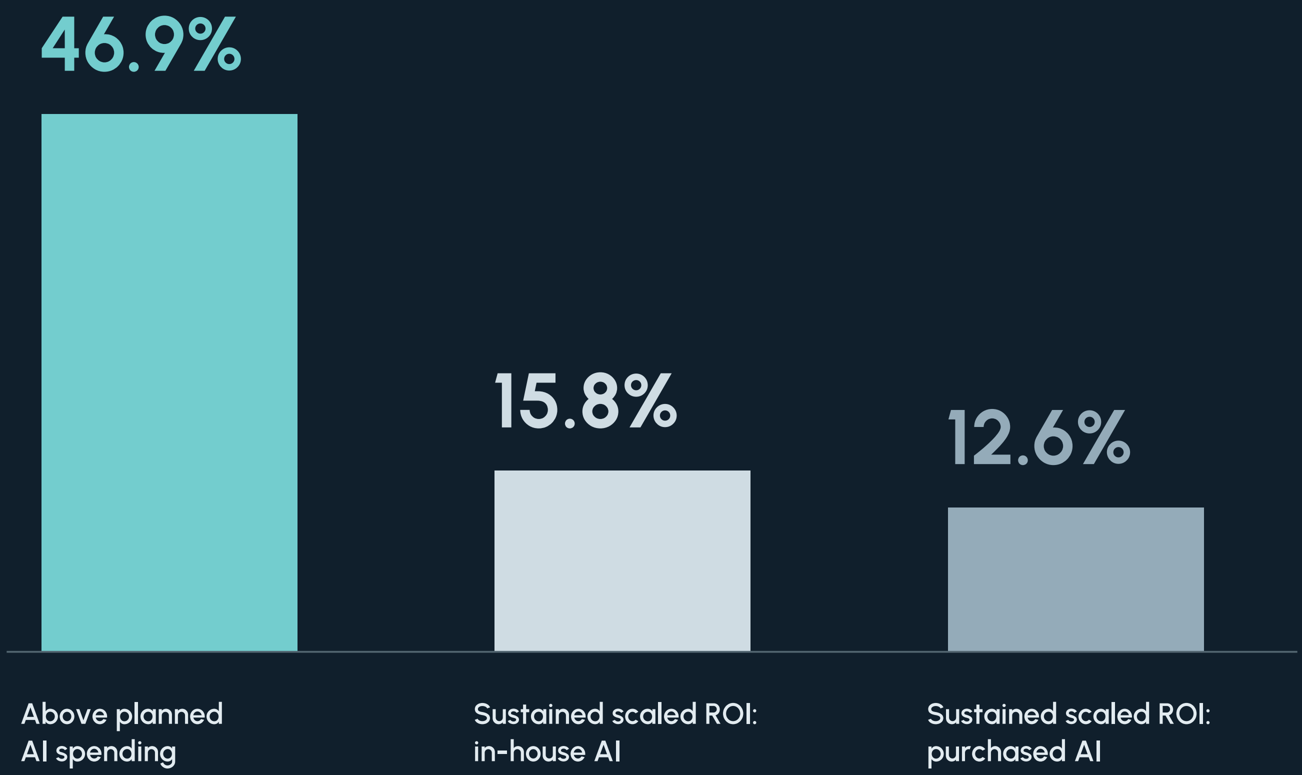

The Three Numbers That Define the Business-Model Gap

46.9% of organizations report spending above their AI plan.

15.8% report sustained ROI at scale from internally developed AI.

12.6% report sustained ROI at scale from purchased AI.

The surveys measure different populations and cannot establish eventual success or failure. Infrastructure is being financed against expectations of scaled adoption, while most enterprises have yet to report sustained ROI at scale.

Sources: spending, Futurum CIO and Technology Buyers Decision Maker Panel, July 2026 (n = 1,636); ROI, Futurum AI Platforms Decision Maker Survey, 1H 2026. Different questions and survey populations.

THE NETWORK OF EXPOSURE

Who Pays Until the Applications Catch Up?

The first stage of the AI buildout was financed largely from the cash flows of some of the strongest companies in the world. Microsoft, Alphabet, Amazon and Meta could spend aggressively without depending on speculative financing.

The next stage is larger. It extends beyond hyperscaler balance sheets into bonds, project finance, private credit, leases, special-purpose vehicles, utility investment and long-term capacity contracts.

Reuters reported in September 2026 that gross debt issuance from hyperscalers could reach a record $420 billion in 2027, up 60% from 2026 estimates. AI-related borrowers were paying spreads around 115 basis points, compared with 78 basis points for the broader investment-grade market.

Investors were not predicting imminent default by the largest technology companies. Their concerns involved the sheer amount and unpredictability of issuance, the lack of visibility into returns and growing concentration in the same corporate names through multiple financing structures.

Some institutional investors were approaching exposure limits after combining parent-company debt with related data-center vehicles. Double-A borrowers sometimes paid spreads closer to those associated with triple-B credits. The market was still open, but it was beginning to charge more for uncertainty.

The greatest vulnerability is unlikely to begin with Microsoft or Google missing a bond payment. It lies farther out in the network.

Data-center developers build against customer commitments. Neocloud providers purchase accelerators against contracts with a few model companies. Utilities add generation and transmission for enormous new loads. Suppliers expand factories. Communities grant incentives and plan around property-tax revenue. Private lenders finance assets whose residual value depends on utilization and technological life.

The system can remain healthy as long as the largest buyers continue spending and their customers keep using more compute. If either assumption weakens, the impact spreads beyond the company making the original decision.

Unlike subprime mortgage borrowers, the largest AI buyers have substantial cash flow. Many facilities have contracted demand, and the equipment provides real computing capacity. Concentrated exposure, opaque related structures and dependence on continued expansion still deserve scrutiny, even when the anchor credits are strong.

AI may not be too big to fail today. It is becoming too interconnected to fail cleanly.

Too interconnected to fail cleanly.

CAPACITY BEFORE CERTAINTY

The buildout has become the business. The return comes next.

Technology success does not guarantee that value accrues to the layer where the largest capital was placed.

AI-GENERATED EDITORIAL ILLUSTRATION

USEFUL INTELLIGENCE PER DOLLAR

The Efficiency Wildcard

Where AI runs and how much capital each useful result requires could determine which infrastructure investments pay off.

The current buildout reflects a frontier assumption: better AI requires larger models, more accelerators, more memory, more power and larger centralized clusters. That assumption has produced the modern AI factory.

Small models are improving. Open-weight models make capable systems available without the economics of a closed API. Distillation and quantization reduce memory and compute requirements. Custom accelerators improve performance for specific workloads. Better interconnects, memory systems and software increase utilization. Inference is moving toward PCs, phones, vehicles, robots, enterprise servers and edge devices.

DeepSeek showed how algorithmic and systems efficiency can challenge assumptions about the capital required to reach competitive capability.

Efficiency does not necessarily reduce total demand. William Stanley Jevons observed that greater efficiency in the use of coal could increase total coal consumption by making more uses economical. AI could follow the same pattern. If the cost of an inference falls 90% while usage grows twentyfold, total inference spending still rises.

The largest forecasts assume cheaper intelligence will be embedded into more applications, agents, devices and workflows. New uses would absorb the capacity freed by efficiency gains.

Jevons paradox does not guarantee that demand will grow faster than efficiency. Where the additional work runs also affects infrastructure demand.

If efficient models allow more work to run on custom silicon, enterprise systems and local devices, AI use can explode while demand for some centralized facilities grows more slowly than anticipated. The world could consume far more AI without buying the precise mix of infrastructure now being financed.

AI may become more indispensable because it requires less centralized infrastructure to use.

Technology success does not guarantee that value accrues to the layer where the largest capital was placed.

Jevons Paradox and the AI Buildout

Falling compute costs can increase total demand by making AI economical in more places. If usage grows fast enough, improved efficiency could help fill the infrastructure being built.

Jevons paradox does not say demand must always rise faster than efficiency or determine where the work will run. Cheaper inference could fill hyperscale data centers, spread across millions of enterprise servers and devices, or do both. Investors need to know whether that usage will support the centralized infrastructure, pricing and financing model being built today.

THE GEOGRAPHY OF COMPUTE

America, China and Competing Infrastructure Models

The AI race pits competing infrastructure systems against one another as well as models.

Futurum estimates that the United States accounted for $109.5 billion of data-center semiconductor revenue in 2025, compared with $36.3 billion for China. The U.S. represented 45.4% of the global market. China represented 15.1%.

Under Futurum's base case, the U.S. reaches $477.9 billion by 2030 while China reaches $115.2 billion. Export controls and restricted access to the most advanced accelerators are among the reasons China's projected share declines.

Physical construction is even more concentrated. BloombergNEF reported approximately 23.1 gigawatts of data-center IT capacity under construction globally, with 15.9 gigawatts in the United States. APAC as a whole had approximately 3.2 gigawatts under construction, while EMEA had 2.9 gigawatts.

Beijing is preparing a reported 2 trillion yuan, or roughly $295 billion, five-year plan for a national network of computing hubs. Goldman Sachs has projected that leading Chinese internet companies will invest more than $70 billion in AI-related infrastructure, although that remains far below comparable U.S. hyperscaler spending.

The American buildout is led by hyperscalers and financed through corporate cash flow, public debt, private capital and project structures. Government helps with permits, land, power, chips and exports, but private companies decide where much of the compute is built and who receives access.

China combines corporate spending with state planning, state-owned telecommunications providers, regional computing hubs and domestic accelerator development. Export controls limit its access to the leading edge, but they also create pressure to use scarce hardware more efficiently and accelerate alternatives such as Huawei Ascend.

Europe is pursuing another path. The European Union is mobilizing more than €30 billion in public and private investment for AI gigafactories, motivated by sovereignty and the fear of dependence on American or Chinese infrastructure.

The Gulf is using sovereign wealth and energy resources. The UAE has started the first phase of a planned multi-gigawatt AI campus, while Saudi Arabia is targeting several gigawatts of capacity through state-backed initiatives.

Semiconductor revenue, public programs, construction costs and gigawatts measure different things and cannot be combined into a single ranking. Across the major regions, governments and companies are treating compute as strategic infrastructure.

Spending alone will not determine leadership. Useful intelligence per dollar, chip and watt will count.

The United States leads in spending without yet establishing an economic model that will sustain it.

SEMICONDUCTOR MARKET COMPARISON

The U.S. leads the spending race.

Data-center semiconductor vendor revenue, US$ billions. 2030 is Futurum's base forecast, not a realized outcome.

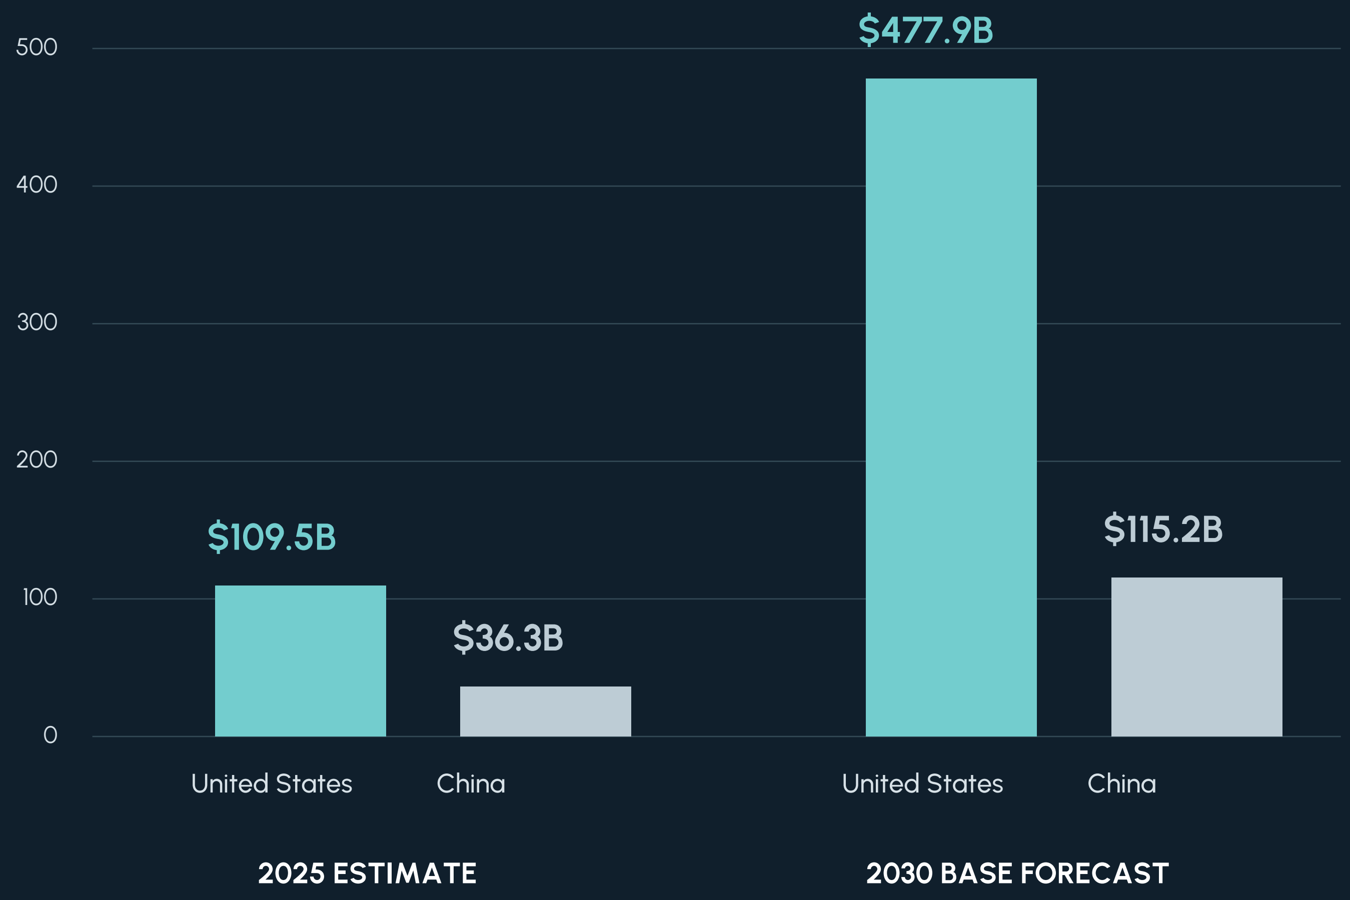

Spending alone will not determine leadership. Useful intelligence per dollar, chip and watt will count.

Source: Futurum Data Center Semiconductors Market Model, 1H 2026, Regional Forecast. Semiconductor revenue must not be combined with construction spending or gigawatts in one ranking.

THE COST OF SLOWING DOWN

Too Economically Important to Slow

Washington's AI strategy began with technological leadership, competition with China, national security and control of the standards governing the next computing platform, well before information-processing equipment passed housing in a quarterly table.

The scale of the buildout has given Washington economic reasons to accelerate it as well.

The White House released America's AI Action Plan on July 23, 2025. Its three pillars were accelerating innovation, building American AI infrastructure and leading in international diplomacy and security. The plan called for faster data-center and semiconductor-fab permitting, fewer regulatory barriers, workforce development and full-stack AI exports.

A companion executive order on data-center infrastructure directed agencies to identify financial support that could include loans, guarantees, grants, tax incentives and offtake agreements. It also sought expedited environmental reviews for qualifying projects.

The administration has opened Department of Energy land to data centers and power generation, encouraged new nuclear and natural-gas capacity and promoted American AI systems abroad. Its AI export order describes AI as foundational to future economic growth, national security and global competitiveness. It calls for exporting packages that include hardware, models, software, applications, standards, storage, cloud services and networking.

Federal policy now supports AI permitting, power supply, financing and exports.

A material slowdown would no longer affect only a handful of laboratories. It could reduce business investment, construction, semiconductor demand, electricity projects, manufacturing orders, corporate earnings and market valuations. It could expose weaker data-center developers and financing vehicles. It could be portrayed as surrendering ground to China.

The same administration has acknowledged some externalities. Its Ratepayer Protection Pledge asks leading technology companies to build, bring or buy generation and cover the grid upgrades required by their facilities rather than shifting those costs to households.

The pledge addresses household costs while the broader policy continues to accelerate construction.

The economic benefits of AI investment appear in today's GDP, earnings, construction and jobs. Many of the risks described by AI researchers are uncertain, probabilistic or deferred. Even when those risks are serious, the benefits of slowing down are harder to measure than the immediate costs.

There is no evidence that Washington is knowingly trading human safety for quarterly growth. But the economic cost of slowing AI has increased dramatically, making strong precautionary action harder to sustain politically, particularly when competitors may continue building.

AI has become too economically important for Washington to approach it as a technology experiment. Policy will favor acceleration with guardrails because the alternative now carries consequences across the economy.

HOW THE WAGER CAN RESOLVE

Four Ways the Wager Can Resolve

Four outcomes are plausible, and different parts of the market may experience more than one.

The buildout succeeds

Enterprise ROI catches up with investment. Inference expands across business and consumer applications. Falling costs unlock enough new demand to absorb the capacity. Productivity becomes visible in macroeconomic data. American AI infrastructure becomes a strategic asset comparable to the cloud, the electric grid or the internet backbone.

In this scenario, the largest concerns shift from financial viability to concentration, labor displacement, electricity demand and control of infrastructure. America may discover that it did not build too much, but struggled to build enough.

The returns arrive late

AI becomes indispensable, but the payback period stretches beyond current expectations. Hyperscalers continue investing because they cannot risk falling behind. Free cash flow remains under pressure. Some capacity sits underused before demand catches up. Weaker developers, neoclouds and model companies consolidate or restructure.

The buildout proves broadly correct in retrospect, but not every investor survives the wait. Public policy becomes more entangled as communities, utilities and strategic suppliers seek protection from a temporary gap between investment and revenue.

AI succeeds but the infrastructure model changes

Efficient models, open weights, custom silicon and edge inference make AI ubiquitous. Enterprises generate substantial value. But more workloads run on smaller clusters, private infrastructure, PCs, phones, vehicles and robots. Some centralized facilities lose pricing power or utilization. The technology delivers while parts of the stack built to support it disappoint.

Railroads reached places their original investors never profited from. Fiber transformed communications after fortunes were destroyed building excess capacity. Public value and private return can travel on different timetables.

Material disappointment

Utilization falls short, token prices collapse faster than demand expands and model providers cannot support their commitments. Data-center projects are delayed or canceled. Semiconductor, memory, networking and power-equipment orders fall. Leveraged operators face distress. Utilities and communities confront infrastructure built for loads that arrive late or not at all.

The impact would reach beyond technology stocks. It would pull against business investment, construction, manufacturing and electricity development. Markets would reprice an AI complex that now occupies a meaningful share of earnings expectations and corporate borrowing.

Washington would then have to choose between allowing a painful adjustment and supporting infrastructure it already considers essential to national security.

A SCENARIO MAP

AI success and infrastructure success are not identical.

Conceptual positions, not probabilities. Multiple outcomes can occur in different parts of the market.

CENTRALIZED INFRASTRUCTURE DEMAND

HIGHER

Returns arrive late High capacity commitments; value realization lags.	**Buildout succeeds** Scaled value and centralized demand grow together.
Material disappointment Weak value realization and disappointing demand.	**AI succeeds; infrastructure changes** Value grows while work shifts across the stack.

LOWER

AI ADOPTION & VALUE REALIZATION: LOWER → HIGHER

AI Can Succeed Even If Infrastructure Investors Lose

Railroads, electric utilities, telecom networks and fiber all created enormous social and economic value. They also produced bankruptcies, restructurings and long periods of poor investor returns.

Infrastructure can be correctly identified as essential and incorrectly financed, priced or timed. AI may follow the same pattern. Models and applications could transform the economy while particular data centers, neoclouds, frontier laboratories or chip commitments fail to earn their cost of capital. A useful technology does not rescue every company that builds ahead of it.

STRESS-TESTING THE THESIS

The Strongest Case Against This Report

Consumer spending remains the largest support for the American economy. AI investment does not approach the scale of household consumption. "Pillar" here describes dependence across several systems, not dominance of GDP.

Information-processing equipment is not all AI. The housing crossover is an opening marker; AI-specific semiconductor data, data-center construction, electricity forecasts, credit issuance and corporate earnings support the broader argument.

Housing weakness has little to do with data centers. Mortgage rates, affordability, zoning and local supply restrictions explain far more than competition from AI. Two lines on a chart do not establish direct causation. Computing investment has risen rapidly while housing has declined, changing the economy's investment mix.

Strong adoption establishes AI's importance, but not sustained ROI. It does not prove that eventual revenue will support every capital commitment.

Falling prices may unlock demand far beyond current forecasts. That prospect supports the bull case, which depends on the elasticity, location and profitability of future demand.

Overbuilding can create long-term public value even when investors lose. Excess fiber eventually made bandwidth cheap and abundant. The country may benefit from surplus compute capacity without every financing decision having been sound.

The largest hyperscalers can withstand enormous spending and should not be confused with fragile dot-com startups. Risk grows as their commitments extend into utilities, suppliers, data-center developers, model laboratories, private lenders and local economies.

Much of the hardware is imported, reducing spending's direct contribution to domestic GDP. The investment, construction, cloud revenue, power demand and financial exposure remain centered in the United States.

The construction boom is geographically concentrated. Virginia, Texas, Ohio, Arizona and a limited number of other markets capture much of it. The semiconductor supply chain, corporate credit, stock market, electricity equipment and federal policies affected by AI extend much farther.

Productivity may arrive faster than skeptics expect. I hope it does.

THE ECONOMIC CONSEQUENCES ARE HERE

The Indispensability Trap Comes for the Economy

In The Indispensability Trap, I argue that becoming indispensable does not guarantee that the company, layer or business model providing the capability captures the value.

AI may become indispensable. I believe it will. That does not mean every frontier lab, infrastructure provider, financing structure or centralized-computing model becomes indispensable with it.

America is building an intelligence utility before settling who ultimately pays for it, how returns will be earned, where workloads will run, how much centralized capacity will be required and who absorbs the losses if demand develops differently than expected.

The government faces a related problem. How aggressively can it regulate a technology that increasingly supports investment, energy development, market earnings and national-security strategy? What happens when prudent friction imposed for safety collides with an economic system rewarded for speed?

Technology companies must earn enough from AI services to cover their infrastructure commitments before those commitments outgrow their ability to carry them.

Enterprises are shifting money from modernization, software, expertise and headcount into AI before most can show sustained scaled ROI. They risk making AI more central to the budget while leaving the organization less prepared to use it.

The financial risk now extends beyond the companies making the original investments.

AI may eventually produce enough productivity, scientific discovery and economic value to justify everything now being built and more. For now, America depends on an investment cycle whose full returns have yet to arrive.

The business model is still being proven. The economic consequences are already here.

NOTES FOR THIS EDITION

What the figures measure.

Housing crossover

Exhibit 01 uses 30 quarterly observations per series, Q1 2019 through Q2 2026. Both are real annualized investment series in chained 2017 dollars. The chart does not measure AI-only investment or data-center construction. The last points are $752.114B and $748.029B.[1 2]

Semiconductor scenarios

Exhibit 02 compares the $241B 2025 estimate with three 2030 scenario endpoints; intervening annual paths are not shown. The market covers GPUs, custom accelerators, CPUs and off-chip memory, not the full data-center stack.

Budget reallocation

Exhibit 03 uses two explicit bases: 767 organizations exceeding AI plans and 297 reallocating other IT spending. Choices overlap; bar lengths are response rates, not shares of conserved dollars or mutually exclusive groups.

Spending versus return

Exhibit 04 compares spending responses from the July 2026 CIO and Technology Buyers panel with ROI responses from the 1H 2026 AI Platforms survey. The measures use separate questions and populations; they do not track a single cohort from spending to returns.

Regional comparison and scenarios

Exhibit 05 compares like-for-like semiconductor revenue estimates. Exhibit 06 is an editorial framework, not a quantitative forecast. “Returns arrive late” represents a timing mismatch between infrastructure commitments and realized application value.

Research attribution

Semiconductor estimates and forecasts are attributed to Futurum's Data Center Semiconductors Market Model, 1H 2026. Enterprise spending and adoption measures are attributed to the Futurum surveys named beside the figures. Forecasts and scenarios are not realized outcomes.

Artwork and typography

Editorial imagery was generated with GPT Image 2.5 Flare. It is illustrative and does not depict identified facilities. The report is set in Urbanist.

1. https://fred.stlouisfed.org/data/Y034RX1Q020SBEA.txt
2. https://fred.stlouisfed.org/data/PRFIC1.txt

THE QUESTION THAT REMAINS

Who ultimately pays for the intelligence utility?

Techstrong
a Futurum company

A Techstrong Special Report
By Alan Shimel
September 2026