

AI INFRASTRUCTURE · COMMUNITIES · CAPITAL

The AI Buildout Meets Reality

Communities are pushing back. Investors are questioning returns. Demand keeps growing. Which data centers will get built, and who will profit?

BY ALAN SHIMEL

49

local ballot measures put data centers in front of voters in 2026

\$4.9T

the difference in Gartner's math between revenue required and revenue forecast

~35%

of announced capacity expected to materialize, in Bernstein's estimate



TECHSTRONG SPECIAL REPORT · OCTOBER 2026

Inside This Report

—	INTRODUCTION	03
	Three places. Three outcomes.	
01	THE NEIGHBORS	05
	The Community Is Part of the Project	
02	THE POLITICS	07
	The Election Puts a Microphone in Front of an Existing Fight	
03	THE PIPELINE	09
	The Pipeline Isn't the Plan	
04	THE DEMAND	11
	Futurum's Bullish Case Has Substance	
05	THE RETURNS	14
	Gartner Puts the Returns Under Scrutiny	
06	THE FINANCING	16
	Who's Holding the Paper?	
07	THE POWER	19
	The Physical Limits Do Not Disappear When the Queue Shrinks	
08	THE ARCHITECTURE	21
	The Technology May Move Before the Building Is Finished	
09	THE HISTORY	23
	The Railroads Still Have Something to Teach Us	
10	THE DEBATE	25
	Where the Arguments Meet, and Where They Don't	
11	THE SCENARIOS	27
	Three Paths Through the Buildout	
12	THE VERDICT	29
	The Terms Will Determine the Outcome	
—	ABOUT THIS REPORT	31
	Source and Methodology Note	

BY THE NUMBERS

49

local data center--related ballot measures tracked in 2026; 41 fall on November 3⁸

\$260B+

in proposed investment reported canceled in 2026 after local opposition¹¹

78.4%

of Futurum survey respondents expect AI budgets to rise in the next 12 months²¹

\$4.9T

difference between the revenue Gartner says is required and its AI revenue forecast¹⁵

649 TWh

LBNL reference case for US data center electricity use in 2030²⁰

TECHSTRONG SPECIAL REPORT · INTRODUCTION

Three places. Three outcomes.

Communities are pushing back. Investors are questioning returns. Demand keeps growing. Which data centers will get built, and who will profit?

BY ALAN SHIMEL OCTOBER 2026 RESEARCH CUTOFF: OCTOBER 24, 2026

In Loudoun County, Virginia, Amazon and local officials have negotiated a preliminary agreement that would keep data centers off the former George Washington University campus in Ashburn. The county announced the arrangement on October 22, following significant opposition to development at the site. Amazon intends to withdraw its data center application, while the parties work toward housing, mixed-use development and community facilities instead. The final agreements still have to be completed.¹

In Arizona, opponents of Project Blue won a unanimous vote against annexation by Tucson in August 2025. The project nevertheless continued in unincorporated Pima County. County records show the developer closed its land purchase in December, followed by permitting progress in

2026. Local reporting documented construction activity in April. The city's rejection did not end the development.²

In Wisconsin, Microsoft announced in June that its first Fairwater facility in Mount Pleasant was fully operational. The company reported nearly 550 full-time employees on site and construction underway on a second facility.³

Three places. Three outcomes. Each tells us something about the AI infrastructure boom, and none can stand in for the whole industry.

The opposition is real. The construction is real. So is the uncertainty about what all that construction will earn.

Data centers have become an unusually tangible expression of the AI economy. People who have never purchased a model subscription are

FIGURE 01

One boom, three very different outcomes

Status of three high-profile AI data center projects, as reported through October 2026

WITHDRAWN

Loudoun County, Virginia

OCT 22, 2026

Preliminary agreement keeps data centers off the former GWU campus in Ashburn. Amazon intends to withdraw its application; housing, mixed-use and community facilities are proposed instead.¹

REROUTED

Pima County, Arizona

AUG 2025 – APR 2026

Tucson unanimously rejected annexation of Project Blue. The project continued in unincorporated Pima County; land purchase closed in December, permitting progressed and construction was reported in April.²

OPERATING

Mount Pleasant, Wisconsin

JUNE 2026

Microsoft reported its first Fairwater facility fully operational, with nearly 550 full-time employees on site and construction underway on a second facility.³

AUG 2025

Tucson rejects
Project Blue annexation

DEC 2025

Project Blue land
purchase closes in Pima County

APR 2026

Construction
activity reported

JUN 2026

Fairwater I fully
operational

OCT 2026

Loudoun preliminary
agreement

Sources: Loudoun County¹; Pima County records and Arizona Luminaria reporting²; Microsoft³.

being asked to live near the equipment that makes those subscriptions possible. Communities are weighing promised tax revenue against noise, water demands, transmission lines and changes to the places they call home. Investors are making a parallel calculation about debt, electricity, customer commitments and the useful life of expensive chips.

Both groups are being offered a future whose benefits have yet to be fully demonstrated. Both are asking who bears the cost if that future arrives late, arrives somewhere else or proves less profitable than advertised.

The approaching US election has brought these questions into campaigns and ballot initiatives. Recent financial-market anxiety has brought them into credit analysis. The technology itself is adding another variable: AI systems are becoming more efficient, smaller models are becoming useful for more tasks, and some work can move onto enterprise systems, personal computers and other devices.

There is a temptation to turn all this into a single verdict. Either the boom is falling apart or the critics are missing the greatest investment opportunity of their lives. The available evidence supports a more demanding assessment. We have to examine what gets built, what earns its keep and who captures the value.

I approach those questions as the founder of Techstrong, part of The Futurum Group, and as the author of *The Indispensability Trap*. Futurum's research provides a substantial case for continued AI investment. My book examines why the economic value created by essential infrastructure may be distributed very differently from the capital spent building it. Those perspectives overlap, and they also produce disagreements worth exploring.

Being bullish on AI does not relieve us of the obligation to ask whether a particular building, business model or financing structure makes sense.

ALAN SHIMEL

Being bullish on AI does not relieve us of the obligation to ask whether a particular building, business model or financing structure makes sense. It makes that obligation more consequential.



**The opposition is real. The construction is real.
So is the uncertainty about what all that
construction will earn.**

ALAN SHIMEL

01 THE NEIGHBORS

The Community Is Part of the Project

Much of the public discussion begins with the wrong assumption: that communities are reacting to a technology they do not understand.

Some people do oppose AI itself. Others are worried about employment, corporate power or the speed of technological change. But a resident objecting to generator noise or an electricity-rate increase does not need a view on artificial general intelligence to have a legitimate concern.

A facility can be essential to the digital economy and poorly suited to a particular location. A developer can deliver substantial investment while negotiating terms that leave existing residents with an inadequate share of the benefits. Conversely, a well-sited development can contribute jobs and tax revenue without producing every harm its opponents anticipate.

Those distinctions tend to disappear when a local hearing becomes an argument about whether a town supports the future.

Greg Noone's article in the November issue of *BBC Science Focus*, "Do we really need to worry about a data centre boom?", puts the community bargain near the center of the discussion. University of Michigan researcher Ben Green questions whether benefits agreements are proportionate to the developments and whether the promised spending addresses what residents actually want. A contribution to a fire department, for example, may do little for someone whose primary concern is persistent noise.⁴

The difference is practical. A donation is something the developer offers. A restriction on noise, water use or operating practices can change

how the developer conducts its business. The latter may matter more to residents and cost more for the operator to honor.

Jobs require similar care. Construction employment can be substantial and economically important. It is also temporary. Permanent employment varies with the size, configuration and operation of a campus. Neither a general claim that data centers create almost no jobs nor a headline construction workforce should settle the local calculation.

Microsoft's Wisconsin announcement supplies a useful case in which the company reports hundreds of ongoing positions. A community considering a different development should ask for its own staffing commitments, expected local hiring and provisions governing whether those commitments are met. A comparison based only on the amount of capital invested misses how differently that capital can translate into local livelihoods.

Tax revenue also needs to be examined after incentives. The relevant question is how much money becomes available to schools, public services and other local priorities over time, after exemptions and abatements, and how that revenue compares with the public costs of supporting the development.

The same discipline applies to water. Berkeley Lab's 2024 report estimated approximately 66 billion litres of direct US data center water consumption in 2023. That is a national estimate, not a description of every cooling system. It also excludes the separate water footprint associated with generating electricity.⁵

66 billion litres

Estimated direct US data center water consumption in 2023

Berkeley Lab national estimate; excludes water used to generate electricity.⁵

A national total cannot tell a community whether a local project is acceptable. The source of the water, seasonal availability, cooling design and competing demands matter. An annual corporate replenishment pledge is not automatically equivalent to reliable supply during a local drought. Nor does the presence of liquid cooling, by itself, establish how much water a facility consumes.

Residents need answers specific enough to be tested against permits and operating records. They also need those answers while there is still an opportunity to affect the decision.

AP's October reporting describes major technology companies offering more community investment and disclosure, including changes to the use of nondisclosure agreements. Amazon and Microsoft have also published commitments concerning community benefits and infrastructure costs. These are meaningful responses to a problem the industry now recognizes. Whether they settle individual disputes depends on implementation.⁶

Trust cannot be restored solely by increasing the size of the presentation. Residents need to know who is building the facility, what it will require, which obligations are enforceable and what happens if the operator changes ownership or fails to deliver.

Trust cannot be restored solely by increasing the size of the presentation.

ALAN SHIMEL

SIDE BAR 01

The Community Bargain, In Six Questions

01

What are the permanent employment commitments, separate from construction jobs?

02

What tax revenue remains after exemptions, abatements and other incentives?

03

Who pays for generation, transmission, distribution and water-system upgrades?

04

Which promises appear in enforceable permits or agreements?

05

Who monitors compliance, and what remedies exist if commitments are missed?

06

Can residents influence the terms before the major decisions are made?

A developer's promise to be a good neighbor becomes more credible when a community can answer these questions without relying on the developer's goodwill.

//

A facility can be essential to the digital economy and poorly suited to a particular location.

ALAN SHIMEL

02 THE POLITICS

The Election Puts a Microphone in Front of an Existing Fight

Electricity bills are a powerful political issue because they arrive whether or not a household uses AI. They also make a complicated infrastructure debate personal.

Reuters reported on October 22 that investors were watching state and local elections for their potential effect on development approvals and restrictions. Its reporting connects community opposition with concerns about the timing and cost of AI investment.⁷

The ballot measures make the issue more than campaign language. Ballotpedia's current tracking identifies 49 local data center-related measures during 2026, with 41 scheduled for November 3. They cover different policies and should not all be described as outright bans.⁸

The election can accelerate action. Candidates have incentives to take positions, officials have incentives to demonstrate responsiveness, and residents have an immediate way to register dissatisfaction. The election can also simplify the debate into slogans that obscure important differences among proposals.

The consequences will depend on where political pressure turns into authority. A congressional speech, a utility tariff, a county zoning decision and a state permitting requirement do different things. An investor evaluating a campus needs to understand which of those can change its delivery date or its operating costs.

There is little reason to expect the underlying concerns to disappear after voting ends. The infrastructure remains. So do bills, noise, water requirements and the contracts governing who pays.

International experience helps separate those issues from the US campaign calendar. Ireland's regulator adopted a new connection policy in December 2025 that establishes a conditional path for development. It requires new data centers to provide generation or storage capacity and obtain at least 80% of annual demand from additional renewable electricity generated in Ireland, with a six-year period for the renewable requirement to be delivered. Location-specific grid constraints remain part of the assessment.⁹

Singapore is also allocating capacity selectively. Its government confirmed that four operators received provisional allocations totaling 200 megawatts under its second data center capacity exercise.¹⁰

These are examples of governments determining the conditions under which growth can proceed. They complicate the idea that regulation must represent either hostility to AI or unconditional support for construction.


For developers, clearer requirements may improve planning even when they raise costs. For communities, conditions can create a more credible bargain. For marginal projects, those same conditions may expose that the economics never worked without shifting part of the cost elsewhere.

That is a legitimate result of scrutiny. An investment does not become socially or economically desirable merely because its developer can announce a large number.

FIGURE 02

The data center fight is on the ballot

Local ballot measures in the US, as tracked by Ballotpedia on October 24, 2026; selected international connection policies



Sources: Ballotpedia⁸; Ireland CRU⁹; Singapore government¹⁰. Ballotpedia counts can change as the live page is updated.


An investment does not become socially or economically desirable merely because its developer can announce a large number.

ALAN SHIMEL



There is little reason to expect the underlying concerns to disappear after voting ends.

ALAN SHIMEL





03 THE PIPELINE

The Pipeline Isn't the Plan

The word “planned” carries too much weight in this debate.

A campus with financing, an executed customer agreement, an electricity connection and construction crews on site is planned. So is a developer's ambition to build several gigawatts on land for which it holds an option. They do not represent the same probability of delivery or the same amount of capital already at risk.

The difference matters when estimates of the development pipeline turn into forecasts of future electricity needs, economic activity or financial losses.

Heatmap reported in October that at least \$260 billion in proposed data center investment had been canceled during 2026 after sustained local opposition. Its public explanation describes tracking public officials, regulatory filings and local reporting, while noting that many projects do not disclose investment amounts. This is a significant reported finding. It is not a measurement of \$260 billion already spent and lost. The complete underlying ledger was not available for independent review for this report.¹¹

Bernstein's widely reported estimate that roughly 35% of announced capacity would materialize addresses a different question. Reporting on its approach describes assessments of construction status, electricity availability and developer experience. The estimate applies probabilities to a heterogeneous pipeline. It does not establish that 65% of financed

construction has failed. Some developers have disputed the assessments.¹²

The disputes around claims that half of planned capacity faces delay or cancellation raise the same methodological questions. SemiAnalysis has challenged the treatment of announcements and the denominator behind those headlines. The existence of that challenge does not make every opposing estimate wrong. It tells us that a report should identify what is being counted before adopting the conclusion.¹³

Consider a developer seeking power in several locations. If it chooses one and withdraws the others, a queue has contracted without any reduction in the developer's intended deployment. Consider a campus announced at five gigawatts but funded in phases. A canceled later phase has a different meaning from the abandonment of a nearly finished building. Consider a customer who moves its requirements to another operator after a permit denial. The original town loses the investment, but the national demand may remain.


Project Blue demonstrates why the distinction is more than a statistical exercise. Tucson's rejection was real. So was subsequent progress in Pima County. Calling the original vote a permanent cancellation would misdescribe what happened.

The reverse error would be to dismiss opposition because developers sometimes find another route. Relocation can add cost and time, alter cooling or power choices, and change which community carries the burden. A project may survive while its original economics do not.

FIGURE 03

Announced is not the same as built

Two widely cited 2026 estimates of the gap between data center ambition and delivery



~35%

of announced capacity expected to materialize, in Bernstein's widely reported probability assessment¹²

Probabilities applied to a heterogeneous pipeline using construction status, electricity availability and developer experience. It does not establish that 65% of financed construction has failed. Some developers dispute the assessments.

\$260B+

in proposed data center investment canceled during 2026 after sustained local opposition, as reported by Heatmap¹¹

A reported tally of proposals, many with undisclosed values. It is not \$260 billion already spent and lost.

Sources: Business Insider reporting on Bernstein¹²; Heatmap¹¹; see also SemiAnalysis's critique of "half canceled" framing¹³. Full models and ledgers were not available for independent review.

The question is therefore how much credible, committed construction is being delayed or abandoned, and why. Answering it requires project records, not the aggregation of optimistic announcements followed by the aggregation of alarming headlines.

The question is therefore how much credible, committed construction is being delayed or abandoned, and why.

ALAN SHIMEL

SIDEBAR 02

Six Ways Capacity Leaves A Pipeline

01

PERMANENT CANCELLATION

The development ends. The customer's demand may disappear or move elsewhere.

02

DELAY

The expected opening moves. Financing costs and equipment age may increase even if demand remains.

03

REDUCED SCOPE

A smaller campus or fewer phases proceed. The change may reflect demand, financing, power or permitting.

04

RELOCATION OR A DIFFERENT APPROVAL ROUTE

The project changes geography or jurisdiction. Local and national consequences differ.

05

DUPLICATE REQUEST WITHDRAWN

Several applications represented one intended development. Removing extras improves the forecast.

06

SPECULATIVE PROPOSAL ABANDONED

An announcement never became a funded commitment. Some development money may be lost, but headline investment was never equivalent to capital deployed.

04 THE DEMAND

Futurum's Bullish Case Has Substance

The demand case deserves to be presented at full strength. There is substantial evidence that enterprises intend to spend more on AI, and Futurum's forecasts describe a much larger infrastructure market even after growth moderates.

In the supplied Futurum AI Platforms research briefing, 78.4% of respondents to the first-half 2026 survey expected AI budgets to increase over the following 12 months, while 1.2% expected decreases. The

survey was fielded in March with 820 respondents. Its AI Chipsets survey, also fielded in March, reported expected average 2026 spending increases of approximately 35% for GPUs and 40% for custom accelerators.^{F1}

Those are intentions, and March is not October. Nevertheless, they describe a broad willingness to commit resources rather than a buyer population preparing to abandon the technology.

FIGURE 04

Buyers are still leaning in

Futurum survey respondents' stated spending intentions, March 2026 fieldwork

AI budgets, next 12 months

AI Platforms Decision Maker survey, 1H 2026 (n=820)

78.4%

Expect AI budgets to increase

Other

1.2% expect decreases

Expected average 2026 spending increase

AI Chipsets Decision Maker survey, 1H 2026 (n=824), approximate

GPUs

+35%

Custom accelerators

+40%

Source: Futurum Intelligence research briefings^{F1}. Intentions, not realized spending; "Other" is the balance of responses not reported as increase or decrease.

Futurum's supplied data center semiconductor base case projects global silicon revenue rising from \$241 billion in 2025 to \$500 billion in 2026 and \$770 billion in 2027. It then anticipates slower annual growth, reaching approximately \$1.21 trillion in 2030. The model covers GPUs, custom accelerators, CPUs and off-chip memory. It does not measure total facility investment or electricity capacity.^{F2}

The distinction prevents misuse of the numbers, but it does not diminish their significance. Even the decelerating base case describes an enormous expansion. A slowdown from exceptional growth rates can coexist with very large additions to annual spending.

The deployment forecast assigns a growing share of silicon purchases to Tier 2 cloud providers. That category includes businesses relevant to the neocloud discussion, although it should not be treated as an exact count or definition of neocloud companies. Its projected expansion helps explain why specialized infrastructure providers have attracted customers and financing.^{F2}

The bull case also has a practical basis. Enterprises do not all have the expertise, facilities or procurement capacity to build their own AI infrastructure. Specialized providers can make equipment available, assemble demanding configurations and handle operational complexity. In a market where getting reliable access matters, that work has value.

Power constraints strengthen part of this argument. If new capacity is difficult to energize, an operator with functioning infrastructure can offer

something that a better-funded competitor cannot reproduce immediately. A large balance sheet does not instantly create an interconnection, a completed building or an experienced operations team.

That advantage should be examined rather than dismissed as mere GPU resale. Specialized software, networking, service quality and customer relationships can distinguish providers that appear similar in a stock-market category. David Linthicum's response to Gartner emphasizes this distinction between generic capacity and specialized operators with operational advantages. His argument supplies a useful counterweight to treating all neoclouds as one financing trade.¹⁴

But demand evidence has limits. The Futurum briefings also report that only 15.8% of surveyed CIOs had achieved sustained ROI at scale from in-house AI and 12.6% from vendor AI in the July 2026 results they cite. The same research material describes substantial spending above plan.^{F3}

A large balance sheet does not instantly create an interconnection, a completed building or an experienced operations team.

ALAN SHIMEL

FIGURE 05

A decelerating forecast is still an enormous one

Futurum Data Center Semiconductors base case: global silicon revenue (GPUs, custom accelerators, CPUs and off-chip memory), US\$




Source: Futurum Data Center Semiconductors model, 1H 2026^{F2}. The 2030 value is derived from the annual growth rates supplied in the briefing. Silicon revenue only; not facility investment or electricity capacity.

FIGURE 06

Few buyers report sustained returns at scale

Share of surveyed CIOs reporting sustained ROI at scale, Futurum CIO & Technology Buyers research, July 2026 (n=1,636)



Fewer than 1 in 6

surveyed CIOs reported sustained ROI at scale from either in-house or vendor AI in Futurum's July 2026 results, even as the same research describes substantial spending above plan.

This does not establish that AI investment is failing. It shows the spending trajectory is not yet matched by widespread reports of sustained financial success among buyers.

Source: Futurum Intelligence research briefings⁹⁹. A different survey population from Figure 04.

Those findings do not establish that AI investment is failing. New systems may require integration, process redesign and time before returns emerge. They do establish that the industry's strong spending trajectory is not yet matched by widespread reports of sustained financial success among buyers.

The distinction becomes more important as deployments move beyond pilots. An enterprise can fund exploration because it needs to

understand a technology. A large, recurring production bill eventually requires a durable business reason.

The bullish case assumes that those reasons will continue developing and justify greater consumption. It is a reasonable proposition to investigate. It is not the same as proof that every layer supplying that consumption will earn the returns its investors expect.



A slowdown from exceptional growth rates can coexist with very large additions to annual spending.

ALAN SHIMEL

05 THE RETURNS

05

Gartner Puts the Returns Under Scrutiny

Gartner's October 14 announcement moves the debate from demand to the relationship between capital spending and revenue. It warns that more than \$3.5 trillion of AI infrastructure spending could be exposed to impairment by 2030 if the current investment trajectory continues.¹⁵

Its public analysis estimates that infrastructure owners would need \$8.5 trillion in cumulative revenue through 2030 to cover announced capital spending, financing costs and an assumed 12% return on invested capital. It forecasts no more than \$3.6 trillion in attributable enterprise and consumer AI revenue over that period. The arithmetic difference is \$4.9 trillion, a separate measure from the impairment exposure.

Gartner also predicts that more than 80% of neocloud providers could cease independent operations by 2030, citing cost pressure and insufficient differentiation. Its public account anticipates growth in AI infrastructure services while questioning whether providers can recover investment before equipment is replaced. The underlying client report extends its analysis through 2032, although the full methodology was unavailable for this report. Jon Swartz covered the announcement for Techstrong IT.^{15,16}

The warning cannot be answered simply by pointing to growing consumption. A business can increase sales and still earn too little on the capital required to produce them. It can have an enormous market and an unattractive cost structure. It can have positive operating margins while generating insufficient cash to fund replacement equipment and repay debt.


At the same time, an aggregate forecast requires choices that can materially change the result. A facility completed late in a forecast period will earn much of its lifetime revenue after that period. A powered building and an accelerator do not necessarily have the same useful life. A vertically integrated technology company may capture benefits in products whose revenue is not labeled AI.

Those are questions to put to the model. Without the full methodology, they are not established errors in it. The public discussion is often too quick to convert a reasonable methodological concern into a confident rebuttal. The same caution applies to the required return. An investor seeking a particular return can be disappointed without losing the entire investment.

FIGURE 07

Gartner's math: the revenue required versus the revenue forecast

Cumulative through 2030, US\$ trillions, from Gartner's October 14, 2026 public announcement



\$3.5T+

of AI infrastructure spending could be exposed to impairment by 2030 if the current trajectory continues. A separate measure from the \$4.9 trillion difference.

>80%

of neocloud providers could cease independent operations by 2030, citing cost pressure and insufficient differentiation. Independence lost is not the same as capacity lost.

Sources: Gartner¹⁵; Techstrong IT coverage¹⁶. The full client methodology, which extends to 2032, was not available for this report.

A business can increase sales and still earn too little on the capital required to produce them.

ALAN SHIMEL

A business that earns less than anticipated might remain solvent and valuable. Whether a financial asset or operating asset must be written down depends on the applicable accounting framework and the evidence about its recoverable value. Whether a borrower defaults depends on its obligations, cash resources and access to financing.

These outcomes can interact, but they should not be collapsed into a single prediction of failure.

Nor should the forecast about independence be translated automatically into a bankruptcy count. Acquisition can reward shareholders or impose losses. A merger can preserve customers and capacity while eliminating an independent company. A distressed sale can keep the equipment operating under a new owner at a lower capital cost. The economic meaning depends on the terms.

That is why the denominator matters in financial forecasts as much as in construction forecasts. The share of companies leaving the market is

different from the share of operating capacity that disappears. A collection of small firms and a few large operators can produce a high company attrition rate without a proportionate loss of supply.

The reverse is also possible: a small number of large failures can carry substantial financial consequences. We should not assume a capacity-weighted result would be reassuring without the underlying data.

Futurum's supplied briefings raise useful questions about these issues but explicitly decline to engage with Gartner under the portal's sourcing policy. They therefore cannot be presented as having disproved its analysis. They provide an independent demand outlook and a framework for evaluating financing risk.

One especially useful limit in those briefings is their acknowledgment that Futurum does not supply a directly comparable total attributable AI revenue forecast in the material available here. Its semiconductor model and its AI Platforms model measure different layers. Subtracting platform revenue from chip purchases would create a dramatic number without a defensible economic interpretation.

The disagreement must be examined through compatible measures. How much capital is deployed? Which revenues and internal benefits support it? What does operating the equipment cost? How long can it earn money? What return do investors require? Those questions are less convenient than a single trillion-dollar headline, but they determine whether the headline survives scrutiny.

SIDEBAR 03

Three Financial Outcomes That Are Not Interchangeable

01

A MISSED RETURN TARGET

Means the investment earns less than its owners expected or required. It may still generate cash and remain viable.

02

AN IMPAIRMENT

Means an asset's recorded value must be reduced under the relevant accounting rules. It is not automatically an immediate cash payment or a declaration of insolvency.

03

A DEFAULT

Means a borrower has failed to meet an obligation under its financing arrangements. Restructuring, enforcement or asset sales may follow, depending on the circumstances.

Consolidation is another category. A company can lose independence through a successful sale, an ordinary merger or a distressed transaction. The label alone does not reveal the return earned by its investors.

It can have an enormous market and an unattractive cost structure.

ALAN SHIMEL

06 THE FINANCING

Who's Holding the Paper?

CoreWeave offers a more useful way into the financing debate than an argument about whether the entire AI economy is a bubble.

In an August 10 announcement filed with the SEC, the company described a \$2.6 billion loan facility with approximately five-year maturity backed by customer contracts averaging approximately three years. CoreWeave said the arrangement broadened its ability to finance shorter customer commitments, which could offer higher pricing and help it reach more enterprises. The company presented lender

participation as confidence in the continuing usefulness of its infrastructure.¹⁷

That is a plausible commercial interpretation. The same transaction also makes the risk visible: the initial customer commitments do not extend through the entire debt period. Existing contracts must be renewed or the capacity must find other customers on terms sufficient to support the remaining obligations.

This is not proof of a bad loan. It is a loan whose success depends partly on conditions beyond the first contracts.

FIGURE 08

The debt outlasts the first contracts

CoreWeave DDTL 5.5, approximate terms as described in the company's August 10, 2026 SEC-filed announcement (schematic, not to calendar)



~\$104B

CoreWeave revenue backlog as of June 30, 2026, including remaining performance obligations and other estimated future revenue under committed contracts, subject to delivery and service availability.¹⁸

Sources: CoreWeave SEC filing¹⁷; CoreWeave second-quarter results¹⁸. Shorter commitments may carry higher pricing; the capacity must be renewed or re-let on terms that support the remaining obligations.

An operator can manage that exposure through customer diversification, conservative underwriting, cash reserves, technical flexibility and strong renewal economics. It can also underestimate what the market will pay when the original customer agreement expires.

The equipment may still function perfectly. That does not establish the rental price it can command or the cost of operating it relative to newer systems. Useful hardware and sufficient collateral value are related questions, not identical ones.

CoreWeave's second-quarter release reported approximately \$104 billion in revenue backlog as of June 30. Its definition includes remaining performance obligations and other estimated future revenue under committed contracts, subject to delivery and service-availability requirements.¹⁸

Backlog deserves attention. It is evidence that customers have made commitments. But it must be converted into service and revenue, and that requires equipment, facilities, electricity and execution. A provider cannot meet a compute obligation with a building that has not been energized.

The financing system must bridge the period between investment and delivery. That bridge becomes more expensive if completion slips. It becomes more vulnerable if a major customer's own funding weakens or if contract protections permit relief when delivery requirements are missed.

Customer quality therefore matters as much as the headline size of the order book. A commitment from a diversified business with strong cash generation has a different credit profile from a commitment by a rapidly growing company dependent on repeated external financing. Neither profile alone determines success, but they should not be treated as equivalent.

A provider cannot meet a compute obligation with a building that has not been energized.

ALAN SHIMEL

Supplier-linked investments add another question. When a supplier or strategic investor provides capital to a customer that then purchases related infrastructure or services, the transaction may support a legitimate expansion. It may also make the apparent demand more dependent on continued financing inside the same ecosystem. The relevant task is to trace the obligations and the ultimate paying customer, rather than assume every connected transaction is either fraudulent or harmless.

Hyperscalers have an advantage because their businesses can generate cash outside the particular AI service being financed. Advertising, enterprise software and established cloud operations can help sustain investment. Vertical integration can also allow them to capture value in several places.

That advantage is substantial, but it is not unlimited. Capital allocated to AI has an opportunity cost. Debt remains debt even when the borrower has a large market capitalization. A diversified company can support disappointing investments longer than a specialized competitor, but the ability to absorb losses does not turn them into satisfactory returns.

The facility owner may occupy a different position from the compute provider. A landlord with a well-structured lease can receive payments without bearing every fluctuation in GPU rental prices. Its protection depends on the

SIDEBAR 04

What A Power Delay Does To The Deal

01

COMPUTE PROVIDER

Equipment may have been purchased or financed before revenue begins. Delay consumes cash and moves the hardware closer to the next generation.

02

FACILITY OWNER

Construction financing and other carrying costs continue. Lease commencement, penalties and customer remedies depend on the contract.

03

CUSTOMER

The intended deployment may be delayed or shifted to another provider, potentially at a different price.

04

LENDER

Expected cash flow arrives later, while the value and usefulness of collateral must be reassessed.

05

UTILITY AND RATEPAYERS

Infrastructure built for a load that arrives late or never arrives still has to be paid for. Tariffs and guarantees determine how that exposure is allocated.

06

SUPPLIER

Orders can be deferred, inventory can accumulate and financing provided to customers can become more difficult to recover.

The effects depend on when money was committed and who accepted each risk. A generic construction-delay percentage cannot answer those questions.

lease, the tenant, the guarantees and the building's usefulness to another customer.

Similarly, a lender may have security over equipment, contracts or project assets. Security is valuable, but its value depends on what can be recovered when several participants face pressure at once. The ability to sell used equipment in a strong market is not the same as the ability to recover a modeled value during widespread distress.

Reuters' coverage of J.P. Morgan's assessment provides a counterweight to claims of immediate systemic danger. The bank's analysts

described banking exposures to CoreWeave as limited and well secured, while acknowledging that estimates were constrained by disclosure. That is evidence against assuming that one company's difficulties automatically threaten the entire financial system. It does not eliminate the need to examine the individual obligations.¹⁹

The report's strongest financial conclusion is therefore specific. The investment case depends on how demand translates into cash under actual contracts and how long those cash flows last relative to the obligations used to finance them.



Useful hardware and sufficient collateral value are related questions, not identical ones.

ALAN SHIMEL

07 THE POWER

The Physical Limits Do Not Disappear When the Queue Shrinks

Removing speculative requests improves a forecast. It does not create electricity for a credible project that still needs a connection.

Berkeley Lab's 2025 Update, published in June 2026, estimates a reference case of 649 terawatt-hours of US data center electricity consumption in 2030. Its combined uncertainty scenarios span 521 to 843

terawatt-hours. The model incorporates equipment shipments, per-device energy use and cooling performance, with sensitivity to utilization and other assumptions.²⁰

Those figures describe modeled consumption, not a sum of utility requests. They give us a way to discuss physical requirements without assuming that every announced campus appears exactly as proposed.


FIGURE 09

The range is as important as the central estimate

Modeled US data center demand in 2030, and one state's approach to allocating the risk

US data center electricity use in 2030

Terawatt-hours, modeled consumption (not a sum of utility requests)



Lawrence Berkeley National Laboratory, 2025 Update, published June 2026.²⁰

Ohio puts a floor under the request

AEP Ohio tariff for qualifying new data center customers



At least 85% of contracted capacity must be paid for, for up to 12 years, even if actual use is lower. The state's consumer counsel calls it protection against shifting costs to other customers.²¹

Sources: LBNL²⁰; Office of the Ohio Consumers' Counsel, case 24-0508-EL-ATA²¹.

The range is as important as the central estimate. If utilization, equipment efficiency or installation patterns differ, the demand for electricity changes. Grid planning must nevertheless begin before all those uncertainties are resolved.

That creates a difficult public-interest problem. Underbuild and productive investment can be delayed or displaced. Overbuild for demand that never arrives and someone pays for unnecessary infrastructure. Existing ratepayers have good reason to ask whether they are providing an involuntary guarantee to speculative development.

Ohio offers one concrete response. Its approved AEP Ohio tariff requires qualifying new data center customers to pay for at least 85% of contracted capacity for up to 12 years, even if actual use is lower. The state's consumer counsel identifies this as protection against shifting infrastructure costs to other customers.²¹

The logic is straightforward. A request becomes more informative when it carries a meaningful financial commitment. Developers with real demand can demonstrate it; applicants seeking inexpensive options on several sites face a different calculation.

But a financial commitment must also be collectible. The credit standing of the party making it, any collateral and the terms governing exit remain important. Cost allocation can reduce exposure without making it disappear under every circumstance.

Onsite generation offers another response to grid delays. It may allow a developer to pursue a different timetable, but it adds a power project to the computing project. Fuel supply, equipment availability, environmental

Existing ratepayers have good reason to ask whether they are providing an involuntary guarantee to speculative development.

ALAN SHIMEL

requirements and operating responsibilities become part of the business model.


Generating power on site also changes the community conversation. Residents concerned about grid pressure may have additional questions about emissions and noise. A project can solve one bottleneck and create another political or financial problem.

The attraction of existing powered capacity follows from these difficulties. A functioning site can be more valuable when alternatives are years away. However, there is no universal guarantee that a powered shell will find a profitable tenant. The design, location, operating cost and price at which it was acquired still matter. Converting a facility to support different densities or cooling requirements may require additional investment.

It is possible for power scarcity to protect one operator's pricing while eroding another operator's returns through delay. That is another reason aggregate descriptions of the industry can conceal the outcome that matters to a particular investor or community.

A project can solve one bottleneck and create another political or financial problem.

ALAN SHIMEL



08 THE ARCHITECTURE

The Technology May Move Before the Building Is Finished

The infrastructure debate often assumes the architecture of AI demand is already settled. It is not.

There are reasons to centralize computing: large shared workloads, demanding training runs, access to substantial memory and specialized equipment, and the ability to manage resources across many users. There are also reasons to execute work locally: latency, privacy, intermittent connectivity, control and cost.

The appropriate balance varies by task. A phone handling a small personal function is not a substitute for a frontier training cluster. An enterprise serving a narrow internal workload does not necessarily need the same infrastructure as a laboratory training a general-purpose model.

Noone's *BBC Science Focus* article brings these questions into a discussion otherwise dominated by electricity and finance. It describes improving hardware and smaller models that make more work feasible on

FIGURE 10

The architecture of AI demand is not settled

Reasons to execute AI work locally or centrally, as described in this report

Run it locally

DEVICE · PC · ENTERPRISE SYSTEM

- Latency
- Privacy
- Intermittent connectivity
- Control
- Cost

which may add central demand
cheaper AI broadens use,

efficient hardware and smaller
models move work outward

**Which workloads go where, and
what can each provider charge?**

Centralize it

HYPERSCALE AND NEOCLOUD DATA
CENTERS

- Large shared workloads
- Demanding training runs
- Substantial memory and specialized equipment
- Managing resources across many users

Sources: report analysis; BBC Science Focus⁴; Futurum AI Platforms deployment forecast⁴. Deployment spending is not a direct measure of inference volume.

personal devices. It also gives the opposite possibility: cheaper AI encourages more use, potentially increasing demand for centralized computing.⁴

Those possibilities are often treated as competing answers. They can be parts of the same development.

Local execution can remove some network and service costs, improve responsiveness and make a capability available in circumstances where a remote service is inconvenient or expensive. If that broadens use, the result may be more AI activity overall. Some of the new activity may remain local. Some may generate additional work for centralized services.

The investment question is which workloads go where and what each provider can charge for serving them.

The supplied Futurum forecasts anticipate growth in on-device and edge deployment while retaining a much larger centralized market. They also project continued expansion in enterprise and private environments. Those forecasts are consistent with a mixed architecture rather than a complete migration away from data centers.^{F4}

They do not, by themselves, measure the share of inference executed on a device. Revenue categories and workload volumes are different things. An inexpensive local function can be used frequently while representing little platform spending. A comparatively small number of demanding remote tasks can generate substantial revenue.

The investment question is which workloads go where and what each provider can charge for serving them.

ALAN SHIMEL

The same issue arises with smaller models. A model that performs a bounded task well enough may have greater commercial value for that task than a more capable model whose extra performance the customer does not need. Quality still matters. The business question is what quality is worth for the specific use.

Green makes that argument in Noone's reporting by asking why customers would pay a large premium for the most sophisticated model when a cheaper alternative meets much of their requirement. That is a useful question, but the answer will differ across applications. An error that is tolerable in a draft summary may be unacceptable in another setting.

Providers can retain premiums for reliability, security, integration and performance. Those premiums must be earned repeatedly as alternatives improve. They cannot be assumed from the fact that a model or a chip was once difficult to obtain.

Custom silicon introduces a related distribution question. If an integrated provider can meet more of its own needs with specialized hardware, it may change the economics of serving particular workloads. That does not imply the end of merchant GPUs. It means the returns at the chip, infrastructure and application layers need to be examined separately.

My argument about local execution is also frequently misunderstood. It does not require device manufacturers to become the principal winners. Cheaper execution can create value for software vendors, service providers, enterprises and consumers. The economic benefit can appear as a new product, a lower bill, a better margin or a task that previously could not be performed economically.

What matters for the infrastructure owner is how much of that benefit returns as revenue, and what it costs to earn it.

There is a serious bullish answer. Agents and other applications may create workloads that more than compensate for efficiency gains and local migration. New capabilities may emerge faster than computing becomes cheaper. Infrastructure providers with good execution could serve a larger market and retain attractive returns.

There is also a serious skeptical answer. The market can grow enormously while prices fall faster than some operators reduce costs. Providers that financed expensive capacity on the assumption of durable scarcity may find that the technology succeeds under economic conditions their balance sheets cannot accommodate.

Neither answer can be settled by counting tokens alone. We need workload economics, realized pricing, utilization, operating costs and the capital required to keep serving customers as the technology changes.



The market can grow enormously while prices fall faster than some operators reduce costs.

ALAN SHIMEL



09 THE HISTORY

The Railroads Still Have Something to Teach Us

A technology can transform an economy without rewarding the original investors in proportion to the value it creates.

ALAN SHIMEL

Noone's article invokes the American railroad boom, drawing on economic historian Liaquat Ahamed's comparison. The connection is familiar to anyone who has read *The Indispensability Trap*. A technology can transform an economy without rewarding the original investors in proportion to the value it creates.

The analogy is most useful when it identifies a mechanism. Large infrastructure commitments anticipate future demand. Financing translates those expectations into obligations. Competition, operating costs, regulation and changes in pricing determine how much of the resulting economic activity becomes income for the owners.

If expectations are wrong, the infrastructure does not necessarily disappear. Ownership can change. Debt can be restructured. Assets purchased below their original cost can support a different business model. The public may continue benefiting even when the original financial structure fails.

The telecom and fiber cycle makes that separation especially relevant to computing. I lived through that period at Interliant. The book's prologue recounts an attempted renegotiation of an expensive bandwidth arrangement as market prices fell, including a meeting with Enron that did not produce a deal. The experience helped shape my view of infrastructure economics: the price decline that hurts one company's balance sheet can become an opportunity for businesses built on the cheaper service.⁵¹

That is not a prediction that every AI investment will repeat the same history. It is a reason to ask whether a financing plan can survive a successful technology becoming less expensive.

The assets differ. Railroad rights of way, fiber and electrical infrastructure can remain useful across long periods. An accelerator faces a different pattern of technical and economic competition. A data center investment combines assets with different lives, rather than constituting one uniform object called infrastructure.

The continuing usefulness of old chips is therefore an empirical question. They may serve workloads that do not need the newest equipment. Whether they do so profitably depends on utilization, achievable prices, operating costs and the capital still owed. A working chip can be an excellent purchase for a second owner at a reduced price and a disappointing investment for the first owner at the original price.

The Wachters' research offers another way to examine the expectations embedded in the buildout. Jessica and Jonathan Wachter's paper, "What Investment Data Implies About the AI Transition," works backward from investment commitments to the productivity expectations that could justify them. Wharton's September account describes an AI-sector productivity multiple of approximately 2.7 under that version's calibration. It

does not mean that every data center must mechanically produce 2.7 times as much output.²²

~2.7 x

AI-sector productivity multiple implied by investment commitments

Wharton's September 2026 account of the Wachters' calibration. Not a per-facility output requirement.²²

Jessica Wachter is also more optimistic than bankruptcy-focused summaries of the research suggest. The Wharton account describes her belief that a substantial productivity boom is possible, while explaining that investment decisions reveal expectations rather than prove their correctness. Different publicly indexed versions show different calibration figures, so the dated account is important.^{22,23}

The lesson is not that history supplies an inevitable ending. It is that investment commitments embody large assumptions about productivity, demand and the ability to capture income. Those assumptions deserve to be stated plainly.

In my book's conclusion, the distinction is explicit: "The trap compresses margins. Leverage converts compressed margins into insolvency." The first process does not guarantee the second. Financing determines how much room the owner has to adapt.⁵¹

The conclusion also says that the thesis caps the fortune relative to value created; it does not erase every fortune. Infrastructure can remain enormous, profitable and essential while applications and other businesses capture a larger share of the economic benefit. That qualification matters when the book is used to interpret current events.

It also makes the disagreement with the bulls more precise. The question is not whether society will find uses for cheaper intelligence. It is whether the people financing a particular layer can retain the revenue and margins needed to support the price they paid for it.

The trap compresses margins. Leverage converts compressed margins into insolvency.

ALAN SHIMEL, *THE INDISPENSABILITY TRAP*

The public may continue benefiting even when the original financial structure fails.

ALAN SHIMEL

10 THE DEBATE

Where the Arguments Meet, and Where They Don't

10

Daniel Newman opens the foreword to my book with a strongly bullish view of AI's economic potential. He also engages the distinction between creating wealth and retaining it. Later, the book's conclusion examines his demand argument and Gennaro Cuofano's case for AI factories.^{B1}

There is substantial common ground: AI can support a much larger economy of applications and services; efficiency can broaden adoption; infrastructure is necessary; and different financial structures create different exposure.

The disagreement concerns how those forces interact with pricing and investment returns. Strong demand can give a provider room to improve utilization and spread fixed costs. Competition and technical progress can simultaneously reduce the price of the service. The result depends on the pace of both changes and on the operator's cost of capital.

That is why an optimistic forecast of chip sales and a skeptical forecast of infrastructure returns do not automatically contradict each other. A chip supplier records a sale when equipment is purchased. The buyer must subsequently use that equipment to earn enough income. Both forecasts could be right about their respective layers.

However, this does not mean every disagreement can be explained away by definitions. Forecasts can be too optimistic or too pessimistic. A

model can underestimate demand, overestimate utilization or choose an unrealistic asset life. The appropriate response is to identify the assumptions and follow the evidence as it develops.

The supplied Futurum material provides strong demand signals and a plausible case for selective development. It does not establish that every contracted operator is financially protected. Gartner provides a substantial warning about the return on announced investment. Its public announcement does not allow us to independently reproduce the full impairment estimate. Pipeline trackers identify an enormous gap between ambition and credible delivery, but their classifications and probability assumptions matter.

Both forecasts could be right about their respective layers.

ALAN SHIMEL

Residents contribute evidence that does not fit neatly into those forecasts. A project that satisfies an investment model may still impose unacceptable costs on a community. A project with a favorable local bargain may still face a financially weak customer. Both tests matter.

FIGURE 11

Both forecasts could be right about their respective layers

Where the bull and bear cases measure the AI economy (conceptual)

1	Chip supplier	Records a sale when equipment is purchased.	Futurum's silicon forecast measures this layer.
2	Infrastructure owner	Must use that equipment to earn enough to cover operating costs, replacement and financing.	Gartner's returns warning concerns this layer.
3	Applications, services and distribution	May capture a larger share of the economic benefit as computing gets cheaper.	Where value moving up the stack would show.

Sources: report analysis; Futurum^{F2}; Gartner^{G5}; The Indispensability Trap^{B1}.

The report's evidence therefore supports a narrower conclusion than either the industry's most confident promoters or its most categorical critics might prefer. There will be a substantial buildout. Its scale, location,

ownership and returns remain contested. The factors determining those outcomes are already visible in permits, tariffs, customer agreements and financing documents.

SIDEBAR 05

How To Test The Bull And Bear Cases

01 IF DEMAND IS AS DURABLE AS THE BULLS EXPECT

Watch whether contracted capacity becomes operating capacity, customers renew at viable prices, and revenue produces cash after financing and equipment replacement.

02 IF THE PIPELINE IS MOSTLY CORRECTING SPECULATION

Watch for fewer announcements and utility requests alongside continued delivery of funded projects. Queue contraction alone is insufficient evidence.

03 IF RETURNS ARE DETERIORATING

Watch realized prices, utilization, operating costs, debt terms and asset sales. A falling share price is a signal of changing expectations, not a complete operating diagnosis.

04 IF VALUE IS MOVING UP THE STACK

Compare infrastructure economics with the economics of applications, services and distribution. Increasing total AI spending can conceal very different margins across those layers.

05 IF COMMUNITY PROTECTIONS ARE WORKING

Look for enforceable obligations, transparent reporting and costs borne by the parties creating them. A promise or a projected benefit is not the same as a delivered outcome.



A project that satisfies an investment model may still impose unacceptable costs on a community.

ALAN SHIMEL

11 THE SCENARIOS

Three Paths Through the Buildout

The first path is continued expansion that earns its returns. Enterprises turn experiments into productive systems. New applications support recurring spending. Providers deliver capacity, maintain utilization and renew customer relationships on terms that cover operating costs, replacement equipment and financing.

In that outcome, some of today's alarming headlines will have described a difficult construction phase rather than a fundamental

investment mistake. Community disputes still matter, but better siting and enforceable agreements allow credible projects to move forward. Investors who distinguished capable operators from speculative proposals may be rewarded.

The evidence to watch is financial conversion. Orders become usable capacity. Usable capacity becomes revenue. Revenue becomes sufficient cash.

A

FIGURE 12

Three paths, not one verdict

Scenarios for the AI infrastructure buildout. No probabilities are assigned, and the paths can coexist across the industry

PATH 1

Expansion that earns its returns

Experiments become productive systems; providers keep utilization high and renew on terms that cover operating costs, replacement equipment and financing.

Orders → usable capacity → revenue → sufficient cash

PATH 2

A selective buildout

Strong demand supports substantial construction while stricter financing and utility requirements remove weaker proposals. Some companies consolidate; some projects move.

Closest to the judgment in the supplied Futurum briefings

PATH 3

A material investment correction

Customer returns develop too slowly, pricing weakens or financing tightens. Assets are sold, debt restructured, and new owners operate capacity at lower capital cost.

AI adoption could continue; the financial harm would still be real

Source: report analysis. The purpose of the scenarios is to identify what would change the assessment.

growing backlog without the rest of that sequence would not establish this scenario.

The second path is a selective buildout. Strong demand supports substantial construction, but stricter financing and utility requirements remove weaker proposals. Some companies consolidate. Some developments move. Customers gain access to AI even as the number of independent suppliers shrinks.

This is closest to the judgment in the supplied Futurum briefings. It is plausible because the pipeline includes projects at very different stages and operators with very different capabilities. It should remain a scenario rather than a settled finding about what most cancellations mean.

The economic distribution could be uneven. A town whose project disappears loses an expected development even if capacity is built elsewhere. An investor in an acquired operator may receive a good return or a disappointing one. A customer may benefit from a more stable provider while becoming more dependent on a smaller group of suppliers.

The third path is a material investment correction. Customer returns develop too slowly, pricing weakens, financing becomes harder to obtain or equipment economics change faster than borrowers anticipated. More credible projects are delayed or canceled. Assets are sold, debt is

restructured and new owners seek to operate capacity at lower capital costs.


AI adoption could continue through that correction. Indeed, cheaper capacity might support it. The financial harm would still be real for the original owners, lenders, suppliers and communities that committed resources on different assumptions.

These paths are not mutually exclusive across the entire industry. A well-positioned campus can thrive while a speculative development fails. A strong operator can acquire a weaker one. A region with power and workable local terms can attract activity that another region loses.

The purpose of the scenarios is to identify what would change our assessment.

ALAN SHIMEL

The purpose of the scenarios is to identify what would change our assessment. It is not to assign invented probabilities or force an emerging market into one predetermined story.



A region with power and workable local terms can attract activity that another region loses.

ALAN SHIMEL

12 THE VERDICT

The Terms Will Determine the Outcome

Return to the three places where this report began.

Loudoun's preliminary agreement demonstrates that even an experienced developer can conclude that a particular site should serve another purpose. Project Blue shows that opposition and approval can operate through different jurisdictions, with construction proceeding despite a highly visible rejection. Wisconsin shows that a large facility can move from announcement to operation.

The next questions concern what each outcome delivers. Do the promised alternatives and community benefits materialize? Does an operating facility generate sufficient income and public benefit? Are costs allocated as promised? Does the infrastructure remain useful as customers and technology change?

The industry has often presented the scale of the buildout as evidence of its inevitability. The scale actually increases the importance of its terms. A project requires credible power, suitable financing, customers able to pay and a workable relationship with the community. Its equipment must serve useful work at an economic price. Its owners must have enough flexibility to respond when the original assumptions change.

Communities do not have to accept a development simply because its output is useful somewhere else. Investors do not have to finance every proposal simply because demand for AI is increasing. Governments can support the technology while requiring those who benefit from a project to bear its costs.

There will also be costs to decisions that prevent good projects. Lost investment, delayed services and constrained access deserve examination alongside environmental and financial exposure. A serious report should make those tradeoffs visible rather than pretend that any decision is free.

The scale actually increases the importance of its terms.

ALAN SHIMEL

I remain bullish on what AI can make possible. That position is entirely compatible with skepticism about the terms on which parts of its infrastructure are being built. An industry capable of changing the economy should be capable of explaining who pays for its power, what its contracts protect and how its investments earn a return.

The people attending town meetings and the people examining loan documents are approaching the same development from different directions. Each is asking whether the promises justify the obligations being placed on them.

AI's future will not be decided by a single cancellation statistic or a single market forecast. It will be built through those obligations, one site and one contract at a time. The facilities that deserve to endure are the ones that can deliver useful computing, meet their

FIGURE 13

What a facility needs to deserve to endure

The terms that will determine which data centers get built and who profits

01

Credible power

An electricity connection, or on-site generation, that can energize the building.

02

Suitable financing

Obligations that match how long the cash flows last.

03

Customers able to pay

Customer quality matters as much as the headline size of the order book.

04

A workable community bargain

Enforceable obligations; costs borne by the parties creating them.

05

Useful work at an economic price

Equipment that keeps earning as alternatives improve.

06

Room to adapt

Owners flexible enough to respond when assumptions change.

Source: report analysis.

financial commitments and honor the bargain made with the people around them.



AI's future will not be decided by a single cancellation statistic or a single market forecast.

ALAN SHIMEL

ABOUT THIS REPORT

Source and Methodology Note

This draft uses public reporting, regulatory and local-government records, company disclosures, the full *BBC Science Focus* article supplied for this project, two supplied Futurum Intelligence research briefings, and the final print interior of *The Indispensability Trap*. Evidence is limited to information available through October 24, 2026. Company-reported benefits and financial expectations are attributed to their source.

The Futurum figures were supplied through the briefings; the underlying proprietary dashboards and full survey tables were not independently accessed. The draft uses specified survey waves where identifiable and avoids the briefings' ambiguous October dashboard snapshots. Spending intentions are not treated as realized spending. Tier 2 cloud is not treated as an exact synonym for neocloud. Silicon revenue, platform revenue, facility investment and electricity demand are not combined into a synthetic shortfall.

The full Gartner client methodology, Bernstein probability model and Heatmap project ledger were unavailable. Their estimates are presented as attributed findings rather than independently reproduced results. The Wachter paper's full PDF was inaccessible; the discussion uses the dated Wharton account and flags differences across indexed versions. The illustrative investment model in the Futurum briefing is not used as evidence of actual market returns.

Notes & Sources

- 1 Loudoun County, October 22, 2026. Preliminary agreement with Amazon at the GWU campus.
<https://www.loudoun.gov/m/newsflash/Home/Detail/10944>
- 2 Pima County, September 17, 2025 and April 23, 2026 memoranda on Project Blue; Arizona Luminaria/Caló News, April 24, 2026 construction reporting.
<https://www.pima.gov/project-blue-faq>
<https://content.civicplus.com/api/assets/2192e865-fe7e-48f4-8a0c-7c1c8b5cd32c>
<https://content.civicplus.com/api/assets/32428b1b-4a81-444a-92a4-72b1fa039655>
<https://azluminaria.org/2026/04/24/we-said-no-dozens-protest-as-construction-of-project-blue-data-center-begins/>
- 3 Microsoft, June 23, 2026. Completion and operation of the first Mount Pleasant facility.
<https://news.microsoft.com/source/2026/06/23/microsoft-completes-construction-on-first-datacenter-facility-in-mount-pleasant-wisconsin/>
- 4 Greg Noone, *BBC Science Focus*, November 2026 issue. "Do we really need to worry about a data centre boom?" Full text supplied by Alan Shimmel.
<https://pocketmags.com/focus-science-and-technology-magazine/nov-26/articles/do-we-really-need-to-worry-about-a-data-centre-boom>
- 5 Lawrence Berkeley National Laboratory, 2024 United States Data Center Energy Usage Report.
https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report_1.pdf
- 6 AP, October 2026, reporting on community commitments and disclosure; Amazon and Microsoft primary commitments.
<https://apnews.com/article/data-centers-artificial-intelligence-nimby-tech-b4b054403698fc33820e8306593aff9>
<https://www.aboutamazon.com/impact/amazon-data-center-tenets>
<https://blogs.microsoft.com/on-the-issues/2026/01/13/community-first-ai-infrastructure/>
- 7 Reuters, October 22, 2026. Election-related risks to AI investment.
<https://www.reuters.com/legal/transactional/investors-see-risk-us-midterms-ai-trade-data-center-backlash-builds-2026-10-22/>
- 8 Ballotpedia, 2026 data center-related ballot measures. Counts as retrieved October 24; the live page can change.
https://ballotpedia.org/Data_center-related_ballot_measures%2C_2026
- 9 Ireland Commission for Regulation of Utilities, December 12, 2025. New electricity connection policy.
<https://www.cru.ie/about-us/news/the-cru-publishes-its-decision-on-new-electricity-connection-policy-for-data-centres/>
- 10 Singapore government, parliamentary written reply on DC-CFA2 capacity and energy supply.
<https://www.meti.gov.sg/newsroom/written-reply-to-pq-on-electricity-demand-from-grid-versus-on-site-generation-and-proportion-of-green-energy-commitment-from-local-renewable-generation-versus-imported-sources-for-dc-cfa2/>
- 11 Robinson Meyer, Heatmap, October 6, 2026. Cancellation tracking.
<https://heatmap.news/daily/260-billion-data-center-opposition>
- 12 Business Insider, October 22, 2026. Reporting on Bernstein's probability assessment and developer responses.
<https://www.businessinsider.com/ai-data-centers-never-built-bragawatts-kevin-oleary-2026-10>
- 13 SemiAnalysis, June 18, 2026. Critique of the half-canceled framing; public portion reviewed.
<https://newsletter.semanalysis.com/p/stop-saying-half-of-2026-us-datacenter>
- 14 David Linthicum, October 17, 2026. Commentary on differences among neocloud operating models.
<https://www.linkedin.com/pulse/gartners-neocloud-call-confuses-financial-exposure-david-linthicum-a6zhe>
- 15 Gartner, October 14, 2026. Public announcement of infrastructure impairment exposure and neocloud consolidation forecast.
<https://www.gartner.com/en/newsroom/press-releases/2026-10-14-gartner-says-over-us-dollars-three-and-a-half-trillion-of-artificial-intelligence-infrastructure-spending-could-be-exposed-to-impairment-by-the-end-of-2030>
- 16 Jon Swartz, Techstrong IT, October 16, 2026. Gartner coverage.
<https://techstrong.it/ai/gartner-over-80-of-neocloud-providers-expected-to-fail-or-consolidate-by-2030/>
- 17 CoreWeave, SEC-filed announcement, August 10, 2026. DDTL 5.5 financing.
<https://www.sec.gov/Archives/edgar/data/1769628/000176962826000357/ex991pr.htm>
- 18 CoreWeave, August 11, 2026. Second-quarter results and backlog definition.
<https://investors.coreweave.com/news/news-details/2026/CoreWeave-Reports-Strong-Second-Quarter-2026-Results/default.aspx>
- 19 Reuters, October 20, 2026. J.P. Morgan assessment of banking exposure.
<https://www.reuters.com/business/finance/coreweave-debt-selloff-poses-limited-risk-global-banks-jp-morgan-says-2026-10-20/>
- 20 Lawrence Berkeley National Laboratory, United States Data Center Energy Usage Report: 2025 Update, published June 2026.
<https://bies.lbl.gov/publications/united-states-data-center-energy-2025>
- 21 Office of the Ohio Consumers' Counsel, AEP Ohio tariff, case 24-0508-EL-ATA.
<https://occ.ohio.gov/content/data-center-costs-24-0508-el-ata>
- 22 Knowledge at Wharton, September 1, 2026. Explanation of Jessica and Jonathan Wachter's investment and productivity research.
<https://knowledge.wharton.upenn.edu/article/can-ai-productivity-grow-fast-enough-to-justify-big-techs-spending/>

- 23** Jessica Wachter's Wharton faculty research page and NBER working-paper record. Indexed calibration figures differ across versions.
<https://fnce.wharton.upenn.edu/profile/jwachter/>
<https://www.nber.org/papers/w35290>
- F1** Futurum AI Platforms Decision Maker, 1H 2026, March fieldwork, n=820; AI Chipsets Decision Maker, 1H 2026, March fieldwork, n=824. Figures supplied in the research briefings.
- F2** Futurum Data Center Semiconductors model, 1H 2026: scenario and deployment forecasts. The approximately \$1.21 trillion 2030 base-case value is derived from the annual growth rates supplied in the briefing.
- F3** Futurum CIO & Technology Buyers research, July 2026 ROI results, n=1,636, as identified in the supplied briefings. Other spending-overrun observations are described qualitatively because the briefing mixes dashboard and snapshot dates.
- F4** Futurum AI Platforms forecast, 1H 2026, deployment model, as supplied in the research briefings. Deployment spending is not a direct measure of inference volume.
- B1** Alan Shimel, *The Indispensability Trap*, Fruition Harbor Press, 2026. Final print interior supplied in this project: Daniel Newman foreword, pp. 9–12; prologue, pp. 13–14; conclusion, especially pp. 269–272. Historical analogies are presented as the author's interpretive framework, not as deterministic forecasts.





The facilities that deserve to endure are the ones that can deliver useful computing, meet their financial commitments and honor the bargain made with the people around them.

ALAN SHIMEL

ABOUT THE AUTHOR

Alan Shimel is CEO of Techstrong Group and CSO of Futurum Media. He is the author of *The Indispensability Trap* (Fruition Harbor Press, 2026), which examines why the economic value created by essential infrastructure may be distributed very differently from the capital spent building it.

ABOUT THIS REPORT

This Techstrong Special Report draws on public reporting, regulatory and local-government records, company disclosures, Futurum Intelligence research briefings and *The Indispensability Trap*. Evidence is limited to information available through October 24, 2026. Imagery is AI-generated illustration and does not depict specific facilities.