

The Impact of Tariffs on the AI Data Center Buildout

Balancing Supply Chain Security and AI Infrastructure Leadership

By Evan Brown, Michael Gary, Kate Koren, and Philip Luck

MAY 2026

THE ISSUE

The United States is on track to invest more than **\$2.7 trillion** in data center infrastructure by 2030, with semiconductors representing approximately 54 cents of every dollar spent. The Trump administration has embraced two goals that are fundamentally in tension: an aggressive push to build out U.S. AI infrastructure, and broad use of Section 232 tariff authority to restrict the semiconductor and metal imports the buildout requires. A January 2025 proclamation that applied a 25 percent tariff on a specific subset of advanced semiconductors carved out imports for U.S. data center construction but left open the possibility of wider levies on all semiconductors and derivative products. A 100 percent tariff on all semiconductors and products containing them would likely impose an additional \$1.4 trillion burden on the buildout. While such a maximalist approach to semiconductor tariffs is not the expected path for the administration, even more moderate tariff scenarios would function as a tax on the United States AI ambitions. This brief examines how semiconductor and metal tariffs interact with data center economics, assesses cost implications across multiple tariff scenarios, and offers policy recommendations to resolve the tension between supply chain security and AI infrastructure leadership.

UNDERSTANDING DATA CENTER COSTS

Building a modern data center requires two distinct categories of capital expenditure: the physical infrastructure needed to supply reliable power, cooling, and structural support, and the information technology (IT) hardware—servers, storage, and networking equipment—that enables AI model training and inference.

Physical infrastructure costs scale with a data center's megawatt (MW) capacity. Traditional data centers require 5-10 MW, while advanced hyperscale facilities are now built to support 200 MW. According to [SemiAnalysis](#), the non-IT costs of constructing a data center, including cooling, building, and power infrastructure, amount to approx-

imately \$10 million per MW, implying a facility's skeleton can cost anywhere from \$50 million to \$2 billion, depending on capacity.

IT hardware costs dwarf non-IT costs. Consulting firm McKinsey & Company [estimates](#) servers, storage, and networking equipment represent approximately 52 percent, 12 percent, and 1.5 percent of total data center capital expenditures, respectively. Within the server category—the largest single cost component—semiconductors [account](#) for roughly 81 percent of value in traditional data centers and up to 87 percent in facilities optimized for AI workloads. Since nearly 80 percent of projected data center capital costs are earmarked for AI-oriented data centers, chips represent at least 45 percent of total server costs. The storage


cost of data centers is similarly chip-intensive at approximately **80 percent** semiconductors, representing another 9 percent of the total cost.

In aggregate, for every dollar spent on data center infrastructure, approximately 54 cents goes toward semiconductors. As data centers grow more graphics processing unit (GPU)-dense to train increasingly complex AI models, chips' share of overall costs will likely rise even further. Nvidia, which provides chips to a host of industries from gaming to data centers, revealed that data center chips now **account** for almost 90 percent of the company's chip sales—a 52 percentage point **increase** in its share in under five years. As demand for data center IT reaches new heights, it is likely that chip shortages, particularly for memory chips, will **result** in price hikes that further increase the cost of the buildout.

Translating these shares into dollar terms, U.S. data centers can be expected to deploy over \$1.4 trillion in semiconductors by 2030—most of which will be **imported**. Only about **12 percent** of global semiconductor fabrication capacity is found in the United States, and efforts to build out domestic leading-edge capacity, while promising, may not be fully realized until the **end of the decade**. Importantly, many of these chips enter the United States already embedded in other products.

Figure 1: The Data Center Bill of Materials Is Dominated by Semiconductors

Percentage share of total data center capital expenditures




Note: Calculations made using predicted data center workloads.

Source: Adam Barth, Chhavi Arora, Gayatri Shenai, Jesse Noffsinger, and Pankaj Sachdeva, "The Data Center Balance: How US States Can Navigate the Opportunities and Challenges," McKinsey & Company, August 8, 2025, <https://www.mckinsey.com/industries/public-sector/our-insights/the-data-center-balance-how-us-states-can-navigate-the-opportunities-and-challenges>; SemiAnalysis, "Datacenter Industry Model," <https://semianalysis.com/datacenter-industry-model/>.

This cost structure presents two key implications for Section 232 tariffs on metals and semiconductors. First, tariffs on metals increase costs for metal-intensive inputs within data center infrastructure, unless such items are excluded. Second, and more consequentially, tariffs on

Figure 2: Semiconductors Will Cost Over \$1.2 Trillion by 2030

Total data center capital expenditures by category



Note: Calculations made using predicted data center workloads.

Source: Adam Barth, Chhavi Arora, Gayatri Shenai, Jesse Noffsinger, and Pankaj Sachdeva, "The Data Center Balance: How US States Can Navigate the Opportunities and Challenges," McKinsey & Company, August 8, 2025, <https://www.mckinsey.com/industries/public-sector/our-insights/the-data-center-balance-how-us-states-can-navigate-the-opportunities-and-challenges>; SemiAnalysis, "Datacenter Industry Model," <https://semianalysis.com/datacenter-industry-model/>.

semiconductors without an exclusion for data center infrastructure would target the largest and least-substitutable portion of the cost stack.

SECTION 232 METAL TARIFFS

The United States has already applied Section 232 tariffs to several industrial metals foundational to data center construction. In 2018, the Trump administration **imposed** tariffs of 25 percent on steel and 10 percent on aluminum on national security grounds. Upon returning to office in 2025, the administration equalized both tariffs at 25 percent and extended coverage to a wide range of derivative products, including structural components, fasteners, and other construction materials. In April 2025, these duties were doubled to **50 percent**. Separately, a Section 232 investigation into copper **culminated** in a new 50 percent tariff effective

on August 1, 2025, with refined copper exempt until a scheduled review in January 2027.

These three tariffs raised the price of an array of upstream inputs used across the data center stack. Steel and aluminum are core components of building shells, server racks, containment systems, power housings, and cooling equipment. Copper is required for power distribution, uninterruptible power supplies (UPS), cabling, and switchgear. Because original equipment manufacturers source globally for cost efficiency and volume, U.S. tariffs translate into higher input costs for domestic suppliers and, ultimately, for data center developers.

The United States currently **imports** approximately 26 percent of the steel it consumes, most prominently from Canada, Mexico, South Korea, and Brazil. It also imports 50 percent of the **aluminum** it consumes, with Canada

Table 1: Data Center Components Exposed to Steel and Aluminum Derivative Tariffs

Component	Percent of Total Capital Expenditure (CAPEX)	Tariff Status	Effective Tariff Rate
Building shell and containment structures	5.1%	Steel derivatives	50%
Computer room air conditioners/computer room air handlers	3.8%	Steel derivatives	25%
Pumps and valves	1.9%	Steel and aluminum derivatives	50%
Raised floor, drop ceiling, and other metal building components	0.5%	Steel and aluminum derivatives	50%
Fire protection systems	0.5%	Aluminum derivatives	25%
GPUs and accelerator boards	48.1%	Excluded	None
Server racks	2.0%	Excluded	None
High-speed networking (interconnects and switches)	7.0%	Not listed	None
UPS systems	5.5%	Not listed	None
Diesel backup generators	5.2%	Not listed	None
Chillers and cooling towers	4.3%	Not listed	None
Switchgear assemblies	3.6%	Not listed	None
Optical fiber cables	3.0%	Not listed	None
PDU's	2.8%	Not listed	None
Security and monitoring systems	1.4%	Not listed	None

Note: AI data center bill of materials (BOM) mapped to April 2026 Section 232 derivative annexes.

Source: SemiAnalysis, "Datacenter Industry Model," <https://semianalysis.com/datacenter-industry-model/>; White House, "Annexes I-A, I-B, II, III, and IV" (Section 232 Tariff Schedules for Steel, Aluminum, and Copper), April 2026, <https://www.whitehouse.gov/wp-content/uploads/2026/04/Metals-ANNEXES-I-A-I-B-II-III-IV.pdf>.

Table 2: Estimated Metal Tariff Cost Impact Under Current and Hypothetical Scenarios

Tariff Scenarios for Steel, Aluminum, and Copper	Additional Cost \$1 Billion Facility	Additional Cost \$2.7 Trillion Buildout
Current policy	\$48 million	\$130 billion
Hypothetical: derivative lists expanded to include networking components and power systems (50 percent rate)	+\$157 million	+\$424 billion
Hypothetical: exclusions for GPUs and server racks removed (50 percent rate)	+\$251 million	+\$676 billion

Source: CSIS.

contributing by far the largest share.

Section 232 metal tariffs apply in two layers. The first layer covers steel, aluminum, and copper articles directly. The second layer extends to derivative articles—finished goods containing substantial metal content—but only when explicitly listed by the harmonized tariff schedule (HTS) code in an annex. In April 2026, the administration **amended** the scope and calculation of the Section 232 metal tariffs. GPU boards and accelerator cards were explicitly excluded, as were server racks. The major power infrastructure components—large diesel backup generators, UPS systems, switchgear, cooling systems, and power distribution units (PDUs) are not included in the derivative products list.

The current derivative metal tariffs cover construction materials and some cooling infrastructure while exempting (explicitly or implicitly) the IT hardware stack. GPU boards, server racks, and networking components, which account for roughly 64 percent of total facility CAPEX, are either specifically excluded or absent from the derivative product lists. The primary exposure falls on structural steel and cooling units, with more limited exposure in pump parts, cooling coils, insulated conductors, and transformer components.

STRUCTURAL SUPPLY CONSTRAINTS

The Section 232 tariffs on metals and their derivatives have not addressed the underlying structural constraints on U.S. production, which present high cost and logistical burdens. Supply chains for many of the key nonsemiconductor inputs to data centers are global, with steel and aluminum derivative products that are subject to tariffs ubiquitous throughout data center infrastructure. For example, **industry reports** suggest that metals tariffs will have severe repercussions for the sourcing of specialized cooling machinery:

Even “Made in USA” cooling systems typically contain 30–60 percent imported content, including specialized compressors, which are covered by the current tariffs.

The United States is also **reliant** on imports for pumps and valves, with only 55 percent of demand met by domestic manufacturing. Increased costs of subcomponents should be expected to increase the cost of the final product, with worst-case scenarios adding hundreds of billions of dollars in additional costs to the buildout. While tariffs could be expected to shift specialized manufacturing to the United States in the long term, building new or restarting idle facilities will require extensive capital investment, significant improvements in U.S. cost structures, and, most crucially, time.

Because U.S. demand for steel, aluminum, and their innumerable key derivative products will continue to exceed domestic supply for the foreseeable future, the United States will remain dependent on imported metals regardless of tariff levels in the interim. The buildout will proceed on the back of imports, but at a higher cost. Yet despite the significant challenge posed by metal tariffs, the cost uplift pales in comparison to the potential impact of a broad Section 232 semiconductor tariff, which would target the highest-cost and least-substitutable input.

SEMICONDUCTOR TARIFF SCENARIOS

The Section 232 investigation into semiconductors initiated in April 2025 **encompasses** leading-edge chips, legacy chips, semiconductor manufacturing equipment, and derivative products, including all products containing embedded semiconductors. If the Department of Commerce determines these imports threaten national security, tariffs may be applied at any rate across any subset of these categories.

Table 3: Semiconductor Tariff Cost Scenarios for the U.S. AI Data Center Buildout

Scenario	Description	Added Cost at Varying Tariff Rates		
		100 Percent Tariff	50 Percent Tariff	25 Percent Tariff
A: All semiconductors and derivative products	Levies applied across all chips and chip-embedded products, regardless of origin or end use.	\$1.44 trillion	\$720 billion	\$360 billion
B: Derivative products only (chips exempt)	Exempts bare semiconductors; tariffs are applied to products with embedded chips.	\$1.15 trillion	\$576 billion	\$288 billion
C: Taiwan exempted	Removes Taiwan-sourced chips (est. 44.2 percent of U.S. logic imports) from the tariff base.	\$379 billion	\$190 billion	\$95 billion
D: Firms with a U.S. manufacturing presence are exempt	Exempts chip designers and fabricators with committed U.S. operations.	Minimal effect	Minimal effect	Minimal effect

Note: Figures represent additional cost burden above the \$2.7 trillion baseline, assuming 60 percent of data center chips are imported. Scenario D is treated as cost-neutral given broad exemption coverage among major technology firms.

Sources: Author’s calculations based on data from Adam Barth, Chhavi Arora, Gayatri Shenai, Jesse Noffsinger, and Pankaj Sachdeva, “The Data Center Balance: How U.S. States Can Navigate the Opportunities and Challenges,” McKinsey & Company, August 8, 2025, <https://www.mckinsey.com/industries/public-sector/our-insights/the-data-center-balance-how-us-states-can-navigate-the-opportunities-and-challenges>; and SemiAnalysis, “Datacenter Industry Model,” accessed May 1, 2026, <https://semianalysis.com/datacenter-industry-model/>.

The authors’ research analyzed four scenarios that collectively illustrate the range of potential outcomes and calculated the additional cost burden to the \$2.7 trillion buildout at tariff rates of 25, 50, and 100 percent.

Scenario A: Tariffs applied to all finished semiconductors and products containing semiconductors. In the broadest tariff scenario, levies would be applied across all semiconductors and semiconductor products irrespective of origin and end use. On average, the United States imports 60 percent of the semiconductors it uses, implying 60 percent of data center chips would see a cost increase due to tariffs. Putting aside the significant increase in prices of consumer products with embedded semiconductors, tariffs on all imported semiconductors would raise the cost of the data center buildout from an estimated \$1.44 trillion to \$4.14 trillion (53.3 percent).

Scenario B: Semiconductors are exempted from tariffs, but products with embedded semiconductors are not. While the United States can be inferred to import about 15 percent of the world’s chips, trade data shows that

it only imports about 3 percent. This is because roughly 80 percent of the chips the United States imports are embedded in other goods and are counted as other products, with servers accounting for roughly 72 percent of the value of GPU imports in 2025. Between 2019 and 2022, Mexico accounted for roughly 70 percent of server exports to the United States, but Taiwan acquired a growing share of these exports. In 2025, for instance, Mexico exported \$80 billion in servers, while Taiwan exported \$86 billion. Exempting semiconductors but not their derivative products, therefore, only reduces tariff burdens by about 20 percent. This option would likely encourage the onshoring of server construction in the United States, but realizing those gains would require a multiyear ramp-up.

Scenario C: Taiwan is exempted. Some reports have **suggested** Taiwan may dodge semiconductor tariffs. According to the most recent data available, the United States **imports** an estimated 44.2 percent of its logic chips from Taiwan. If this percentage is used as a proxy for Taiwan’s contribution to data center chips specifically (which

could well be an underestimate), tariffs will apply to 15.8 percent of data center chips, assuming no adjustments in purchasing behavior. Given the high switching costs associated with changing chip manufacturers, firms will likely refrain from switching fabricators, especially if they suspect the tariffs may be lifted. In the 100 percent tariff case, the \$1.2 trillion chip portion of the buildout cost would rise by approximately 31.6 percent, raising the cost of U.S. AI infrastructure by \$379.2 billion to \$3.08 trillion by 2030.

Scenario D: Scenarios A-C, where an exemption is made for firms with a U.S. presence. In early conversations around potential semiconductor tariffs, President Donald Trump suggested tariffs would not be applied to firms with current or pledged manufacturing operations in the United States. Experts argue a semiconductor tariff with this exemption means “much of the sector will be exempt,” with leading technology end users like Apple, hyperscale firms like Google, Amazon, and Microsoft, and chip designers and manufacturers like Nvidia and TSMC all committing to U.S. manufacturing operations. From the U.S. data center buildout perspective, semiconductor tariffs with a carve out for firms that commit to invest in U.S. manufacturing would likely add minimal cost.

WHY INDIGENIZATION CANNOT FILL THE GAP

Section 232 tariffs on semiconductors without exclusions would have dramatic cost ramifications for the AI data center buildout because the United States is currently reliant on imports, and efforts to indigenize the semiconductor supply chain have long timelines.

The Semiconductor Industry Association estimates \$630 billion has been committed to developing the U.S. supply chain across 140 projects in 28 states. U.S. semiconductor manufacturing capacity is projected to increase by 203 percent from 2024 to 2032. Yet domestic supply will still not satisfy domestic demand, which stands at 25 percent of global consumption in 2025 and will likely rise throughout the buildout.

EFFORTS TO EXPAND U.S. SEMICONDUCTOR MANUFACTURING CAPACITY


Passed in August 2022, the bipartisan CHIPS and Science Act aims to strengthen the U.S. semicon-

ductor ecosystem by providing incentives for domestic investment, research, and manufacturing. While U.S. chip fabrication shrank from 37 percent of global capacity in 1990 to 10 percent in 2022, the CHIPS Act has stimulated over half a trillion dollars in private sector investments, priming the United States to triple its semiconductor manufacturing capacity from 2022 to 2032. This growth rate implies that the United States could supply nearly 30 percent of all advanced logic chips by 2032. Such a trajectory, however, is not a fait accompli and will depend on future legislation that extends and expands incentives, such as investment and manufacturing tax credits.

The cost disadvantage for U.S. fabs is significant. Taking into account existing U.S. subsidies, a standard mature logic fab costs 10 percent more to build in the United States and 35 percent more to operate compared to an equivalent facility in Taiwan. Much of this discrepancy

Figure 3: U.S. Fabs Costs Are Roughly 30 Percent Higher than Taiwan

Annual cost of goods sold of an example mature node fab with subsidies, percentage compared with Taiwan



Note: Considers wafer starts per month for a node fab for 28-nanometer nodes, with 28,000 starts per month fab capacity. Assumes 85 percent utilization for all regions with subsidies and 5 percent uncertainty margin.

Source: Bill Wiseman, Henry Marcil, Marc de Jong, Raphaela Wagner, Taylor Roundtree, and Teddy Stopford, “Semiconductors Have a Big Opportunity—but Barriers to Scale Remain,” McKinsey & Company, April 21, 2025, <https://www.mckinsey.com/industries/semiconductors/our-insights/semiconductors-have-a-big-opportunity-but-barriers-to-scale-remain>.

ancy is associated with the reality that labor costs in the United States are **two to four times** higher than in Asia, impacting both the construction and maintenance of fabs. With recent U.S. investments **projected** to require nearly 115,000 new employees in the U.S. semiconductor industry by 2030—with roughly 67,000 of those positions at risk of going unfilled—competition could drive labor costs even higher.

TSMC’s CEO Morris Chang has **stated** that the production cost per chip at the company’s Phoenix, Arizona, facility is 50 percent higher than at its Taiwanese counterparts, and that the facility is “noncompetitive in the world market.” On the timeline side, U.S. fabs take over three years on average from permitting to first production, roughly 1.5 years longer than equivalent Taiwanese facilities beginning permitting at the same time.

These barriers are certainly surmountable, but catching up will require long-term structural shifts. Even if the United States triples its domestic chip capacity by 2032 as projected, the critical AI infrastructure buildout of the next five years will be achieved on the back of semiconductor imports.

ating the reshoring timeline. A structure that combines conditional tariff relief with investment accountability would help meet both goals.

- **Establish a five-year exemption window** for semiconductor imports destined for U.S. data center construction. This window would align tariff relief with the critical AI buildout period before domestic fab capacity begins to come online at scale.
- **Define clear domestic capacity milestones** measured in wafer starts per month and, where applicable, tied to CHIPS Act commitments that trigger the phased introduction of tariffs as domestic supply grows. This would help ensure tariffs function as a bridge to domestic production rather than a permanent tax.
- **Require beneficiary firms to report annually** on their sourcing mix and U.S. manufacturing investments as a condition of exemption eligibility, generating the data needed to inform future tariff calibration.

POLICY RECOMMENDATIONS

Time-Limit Semiconductor Tariffs

and Tie Relief to Domestic Investment Milestones

A blanket semiconductor tariff would impose hundreds of billions of dollars in costs on the buildout without acceler-

Address the Structural Barriers That Tariffs Cannot Solve

Tariffs are a demand-side intervention: they raise the cost of imports to incentivize domestic substitution. But the constraints on U.S. semiconductor and metals production are

Figure 4: The United States’ Permitting, Design, and Construction Phases for Fabs are Slower Than in Peer Countries

Months to complete



Source: Pete Singer, “Building Fabs in the U.S. vs Taiwan: Twice as Long, Twice as Much,” Semiconductor Digest, February 18, 2025, <https://www.semiconductor-digest.com/building-fabs-in-the-u-s-vs-taiwan-twice-as-long-twice-as-much/>; Herbert Blaschitz, “Challenges and Opportunities in Semiconductor Fab Planning and Construction,” Exyte, <https://www.exyte.net/About-Us/Innovation-And-Experts/Opportunities-semiconductor-fab-planning-and-construction>.

fundamentally supply-side, permitting delays, workforce gaps, capital costs, and geological disadvantages. Complementary interventions would help create the conditions for successful and speedier onshoring.

- **Permitting reform:** Semiconductor fab permitting currently adds one to two years to project timelines relative to Asian competitors. Pre-approved environmental review templates for standard fab configurations and dedicated interagency coordination for major projects would accelerate the development of domestic capacity.
- **Workforce development:** Industrial worker shortages are a key driver of U.S. cost disadvantage. Targeted investment in community college semiconductor technician programs, visa provisions for specialized manufacturing skills, and apprenticeship pipelines would address a constraint that tariffs cannot.
- **Allied supply chain coordination:** Coordinating tariff exemption frameworks with Japan, South Korea, the European Union, and Taiwan would allow the United States to maintain supply chain security without imposing costs on imports from trusted partners. A United States-Mexico-Canada Agreement-aligned sourcing preference for metals, combined with allied-nation exemptions for semiconductors, would significantly reduce buildout cost impacts while preserving leverage over adversarial supply chains.

CONCLUSION

The United States has a genuine national security interest in reducing its dependence on semiconductor imports from geopolitical adversaries. However, the administration's ambition to lead global AI infrastructure development

depends on the very global supply chains that its tariff agenda is designed to reduce. Resolving this tension does not require abandoning either goal, but it requires precision in how tariffs are designed and applied and complementary policies to support domestic growth. The administration's restraint on tariffs affecting AI data center inputs to date suggests a recognition that indiscriminate measures are at odds with buildout goals. However, Section 232 tariffs that fall outside these carveouts will still impose real costs on the broader technology ecosystem. Sectors such as domestic research and development and electronics manufacturing, if exposed to higher input costs, risk being disadvantaged relative to foreign competitors. A national security-driven tariff strategy should therefore take a holistic view of the technology stack and avoid shifting vulnerabilities to less visible segments of U.S. competitiveness. ■

Evan Brown is a program manager and research assistant with the Economics Program and Scholl Chair in International Business at the Center for Strategic and International Studies (CSIS) in Washington, D.C. Michael Gary was a former research intern with the Economics Program and Scholl Chair in International Business at CSIS. Kate Koren is deputy director of the Economics Program and Scholl Chair in International Business at CSIS. Philip Luck is director of the Economics Program and Scholl Chair in International Business at CSIS.

The authors would like to thank Justin Hu for his research contributions to this piece.

This brief is made possible by support from the Computer and Communications Industry Association.

CSIS BRIEFS are produced by the Center for Strategic and International Studies (CSIS), a private, tax-exempt institution focusing on international public policy issues. Its research is nonpartisan and nonproprietary. CSIS does not take specific policy positions. Accordingly, all views, positions, and conclusions expressed in this publication should be understood to be solely those of the author(s). © 2026 by the Center for Strategic and International Studies. All rights reserved.

Cover Photo: halbergman/Getty Images