INSTITUTE FOR AMERICAN
MANUFACTURING & TECHNOLOGY

Powering the Buildout

Energy Demand Projections for Data Centers, Reshored Manufacturing, and Defense Production

A White Paper from the Atlas and Aegis Institute at IAMT

Written by: Patrick J. Wolf, PhD

February 2026

IAMT Policy Series

Table of Contents

Executive Summary 3
I. The Convergence: Three Buildouts, One Grid 4
AI Data Center Infrastructure 4
Reshored Manufacturing 4
Defense Production Expansion 4
The Cumulative Challenge 5
II. Grid Capacity: Current State and Structural Limitations 6
Generation Capacity 6
Natural Gas as Transition Fuel 6
Transmission Constraints 7
Interconnection Queues 7
III. Permitting as the Binding Constraint 8
Federal Environmental Review 8
State and Local Permitting 8
The Timeline Mismatch 8
What Failure Looks Like 9
IV. Energy as Industrial Policy: The Case for Integration 10
The Cost of Misalignment 10
An Integrated Framework 10
V. Regional Analysis: Where the Constraints Are Most Acute 11
Northern Virginia and the PJM Interconnection 11
Central Texas 11
Central Arizona 11
Ohio River Valley 12
VI. Recommendations 13
For Government 13
For Industry 13
For Policy and Research Institutions 14
Conclusion: Energy Is the Foundation of Industrial Strategy 15

Executive Summary

The United States is simultaneously pursuing three of the most energy-intensive industrial buildouts in its modern history: the construction of artificial intelligence data center infrastructure at unprecedented scale, the reshoring of advanced manufacturing capacity in semiconductors and other strategic sectors, and the expansion of defense production to meet the requirements of great power competition. Each of these initiatives is central to the nation's economic competitiveness and national security. Each is also constrained by the same binding resource: electric power.

This white paper, produced jointly by the Atlas Institute and the Aegis Institute at the Institute for American Manufacturing and Technology, models the aggregate energy demand generated by these three converging buildouts, assesses the capacity of the existing electric grid to absorb this demand, and identifies the structural barriers—permitting timelines, interconnection queues, generation capacity shortfalls, and transmission constraints—that threaten to make energy the rate-limiting factor in America's industrial strategy.

The central finding is that current energy planning does not account for the cumulative demand of these initiatives. Grid capacity projections, generation planning, and transmission investment are proceeding on trajectories that predate the AI infrastructure boom, the CHIPS Act fabrication buildout, and the defense production expansion. Unless energy policy is aligned with industrial policy, the nation's ambitions in all three domains will be constrained not by a lack of capital, technology, or strategic intent, but by an inability to deliver the electrons required to power them.

The paper argues that energy is the connective tissue of industrial policy. Manufacturing, compute, and defense production are distinct policy domains, but they share a common dependence on reliable, abundant, and affordable electric power. A coherent industrial strategy requires an energy policy framework that treats power as a strategic input to national competitiveness rather than a commodity governed solely by market dynamics and regulatory process.

I. The Convergence: Three Buildouts, One Grid

The United States is in the early stages of an industrial expansion that, if executed, will represent the most significant increase in domestic manufacturing and compute infrastructure in a generation. Three distinct but simultaneous buildouts are driving this expansion, each supported by substantial public and private investment, and each generating energy demand at a scale that challenges the existing power infrastructure.

AI Data Center Infrastructure

The capital expenditure commitments of major technology companies for AI data center construction have reached extraordinary levels. Hyperscale cloud providers and AI companies committed over three hundred billion dollars in combined AI infrastructure spending in 2025 alone, with additional hundreds of billions planned through the end of the decade. These facilities are among the most power-intensive structures in the built environment. A single large-scale AI training cluster can consume one hundred megawatts or more of continuous power, and individual hyperscale campuses now request one gigawatt or more—equivalent to the power produced by half of the Hoover Dam. The aggregate demand from planned data center construction is projected to add tens of gigawatts of load to the U.S. electric grid over the next decade. In Texas alone, ERCOT projects that data center demand could reach seventy-eight gigawatts by 2031, and the grid operator is tracking over two hundred gigawatts of large-load interconnection requests—more than double the state's record peak demand of eighty-five point five gigawatts.

Reshored Manufacturing

The CHIPS and Science Act, combined with broader industrial policy initiatives, is catalyzing the construction of semiconductor fabrication plants, battery manufacturing facilities, and other advanced manufacturing installations across the United States. Semiconductor fabrication is exceptionally energy intensive. A modern leading-edge fabrication plant consumes approximately one hundred to one hundred fifty megawatts of continuous power, with additional demand for ultra-pure water systems, HVAC, and supporting infrastructure. The planned fabrication facilities in Arizona, Ohio, Texas, and New York collectively represent a significant new load on regional grids that were not designed to accommodate them.

Defense Production Expansion

The expansion of defense manufacturing capacity driven by the need to replenish stockpiles, support allies, and prepare for potential high-intensity conflict, adds a further layer of energy demand. Munitions production, missile manufacturing, shipbuilding, and the expansion of production lines for aircraft and ground vehicles all require energy, much of it in forms—high-temperature industrial heat, specialized

power profiles—those present distinct challenges for grid accommodation. Defense production facilities are often located in regions where grid capacity is already constrained, and the classified nature of defense programs can complicate the energy planning and permitting process.

The Cumulative Challenge

Taken individually, each of these buildouts represents a significant but potentially manageable increase in energy demand. Taken together, they represent a structural shift in the nation's power consumption profile that current energy planning does not adequately account for. Grid operators, regional transmission organizations, state public utility commissions, and federal energy agencies are planning for load growth based on historical trends and approved projections that predate the AI data center boom and the CHIPS Act construction wave. The gap between projected and actual demand growth is widening, and the consequences of this gap will be felt across all three policy domains.

II. Grid Capacity: Current State and Structural Limitations

The U.S. electric grid is the product of over a century of incremental development, designed to serve a load profile that has been relatively stable in recent decades. Total U.S. electricity generation has been roughly flat since the mid-2000s, with efficiency gains and the decline of energy-intensive manufacturing offsetting modest demand growth. Grid planning, transmission investment, and generation capacity additions have reflected this period of stability.

The infrastructure buildouts now underway are disrupting this equilibrium. Regions that are attracting data center and manufacturing investment such as northern Virginia, central Texas, the Ohio River Valley, central Arizona, are experiencing demand growth that exceeds historical planning assumptions by substantial margins. Grid operators in these regions are reporting that available capacity is being consumed faster than new generation and transmission can be built.

Generation Capacity

The U.S. generating fleet is undergoing a simultaneous transition in fuel mix and an expansion in total capacity. Coal plant retirements are accelerating, driven by economics and policy. Natural gas generation remains the dominant source of dispatchable power. Renewable generation, wind and solar, is growing rapidly but presents challenges of intermittency and geographic concentration that complicate its role in serving the continuous, high-reliability loads that data centers and manufacturing require.

Nuclear generation, which provides approximately twenty percent of U.S. electricity on a carbon-free, dispatchable basis, is attracting renewed interest from technology companies seeking to power data centers. Several major technology firms have announced agreements to purchase power from existing or new nuclear facilities. However, the timeline for new nuclear capacity, whether conventional large-scale plants or small modular reactors, extends well beyond the timelines for the data center and manufacturing facilities that need the power. U.S. Nuclear Regulatory Commission (NRC) licensing for new reactor designs takes years, and only in a limited capacity has Small Modular Reactor (SMR) design has entered commercial-scale production. SMRs are a critical element of the long-term generation mix but cannot solve the near-term power gap facing industrial buildout in this decade.

Natural Gas as Transition Fuel

Natural gas remains the dominant source of dispatchable power and will inevitably serve as the primary near-term generation source for new industrial and data center load. This reality must be confronted directly. Gas turbine procurement lead times now exceed three years due to surging global demand. Gas pipeline infrastructure faces its own permitting constraints. The same National Environmental Policy Act (NEPA) and state-level review processes that delay transmission lines also delay the pipelines needed to

deliver fuel to new generating stations. Long-term policy uncertainty around decarbonization targets further complicates investment in new gas generation. Natural gas is a near-term necessity, not a complete solution, and policy must address both its essential role in the current decade and the generation transition that follows.

Transmission Constraints

Transmission infrastructure, the high-voltage lines that move power from generation sources to load centers, is among the most constrained elements of the energy system. The permitting and construction of new transmission lines is a multi-year, often multi-decade process involving federal, state, and local approvals, environmental review, land acquisition, and public opposition. The current transmission network was designed for a different era of generation and load distribution. Connecting new generating resources to the grid, and delivering power to new large-load customers, requires transmission capacity that in many cases does not exist and cannot be built on the timelines that industrial customers require.

Interconnection Queues

New generation projects must apply for interconnection to the grid through regional transmission organizations. The interconnection queue process has become a major bottleneck. As of recent reporting, there were over two thousand gigawatts of proposed generation projects waiting in interconnection queues across the United States, with average wait times exceeding four years. The queue is dominated by renewable energy and storage projects, but the congestion affects all new generation, including the gas and nuclear projects that may be best suited to serve the continuous loads of industrial and compute customers.

III. Permitting as the Binding Constraint

The single greatest structural barrier to aligning energy supply with industrial demand is the permitting process. The construction of new generation, transmission, and distribution infrastructure in the United States requires navigating a layered system of federal, state, and local approvals that was designed for an era of stable demand and incremental infrastructure additions. This system is fundamentally mismatched with the pace and scale of the industrial buildout now underway.

Federal Environmental Review

The National Environmental Policy Act requires environmental impact statements for major federal actions, including the permitting of energy infrastructure on federal lands and projects that require federal approvals. NEPA reviews for energy projects routinely take three to five years or longer, and complex projects can take a decade or more. Recent legislative efforts to reform NEPA and accelerate permitting have made incremental progress, but the fundamental structure of the review process remains intact. For the energy projects needed to power the current industrial buildout, NEPA timelines represent a constraint that no amount of private capital can overcome.

State and Local Permitting

State-level sitting authorities, public utility commission proceedings, local zoning approvals, and community opposition add additional layers of delay to energy infrastructure projects. Transmission lines, generation facilities, and substations all require state and local approvals that can extend project timelines by years. The geographic distribution of industrial buildout, which is mainly concentrated in a relatively small number of states, means that the permitting capacity of these states is being tested by a volume of applications that exceeds their administrative capacity to process.

The Timeline Mismatch

The fundamental problem is one of timeline mismatch. A data center can be designed and constructed in eighteen to twenty-four months. A semiconductor fabrication plant takes three to five years. The energy infrastructure required to serve these facilities, such as new generation, transmission upgrades, substation construction, takes five to ten years or more to permit and build. Average permitting timelines for interstate transmission lines range from seven to ten years, and individual projects have exceeded fifteen years from proposal to energization. The Gateway West transmission project in Wyoming and Idaho, for example, required over a decade of permitting before construction could proceed. This means that industrial facilities are being designed and constructed faster than the energy infrastructure they depend on, creating a structural timing gap that has real operational consequences.

What Failure Looks Like

The consequences of timeline mismatch are not theoretical. Consider a fabrication plant completed on schedule at a cost of twenty billion dollars. The grid upgrades required to deliver reliable power were submitted to the utility concurrently, but interconnection studies took eighteen months, transmission construction required four years, and permitting added two more. The facility opens but operates at reduced throughput for two years because full power is not available. The output shortfall delays chip supply to defense and commercial customers. Unit costs rise because fixed capital is underutilized. The strategic case for domestic fabrication which is already challenged by higher operating costs is further weakened. The facility eventually reaches full capacity, but the two-year delay has cost billions in foregone production and eroded confidence in domestic manufacturing economics. This scenario is not hypothetical. Elements of it are already visible in multiple CHIPS Act-funded projects encountering infrastructure gaps between facility completion and power availability.

IV. Energy as Industrial Policy: The Case for Integration

The United States does not currently treat energy policy and industrial policy as integrated disciplines. Energy policy is primarily the domain of the Department of Energy, the Federal Energy Regulatory Commission, state public utility commissions, and regional transmission organizations. Industrial policy is primarily the domain of the Department of Commerce, the Department of War (Formerly the Department of Defense), and the relevant congressional committees. These institutions do not share planning frameworks, demand projections, or decision-making processes. The result is that energy supply and industrial demand are planned in parallel but not in coordination.

This institutional separation was tolerable during a period of stable demand and incremental industrial change. It is not tolerable during a period of rapid, policy-driven industrial expansion. When the federal government commits tens of billions of dollars to semiconductor fabrication, it is implicitly committing to the energy infrastructure required to power those fabrication plants. When the Department of Defense expands munitions production, it is implicitly committing to the grid capacity required to serve those production facilities. When the market builds hundreds of data centers to power AI, it is implicitly assuming that the grid will accommodate the load. None of these implicit commitments are currently being explicitly planned for in a coordinated manner.

The Cost of Misalignment

The consequences of this misalignment are beginning to materialize. Data center developers are being told by utilities that power will not be available for three to five years. Semiconductor fabrication plants are investing in on-site generation as a hedge against grid uncertainty. Defense production facilities are discovering that grid upgrades required to support capacity expansion will take longer than the production ramp-up they are designed to support. In each case, the underlying industrial investment is being delayed, constrained, or made more expensive by energy infrastructure that was not planned in coordination with industrial demand.

An Integrated Framework

The Atlas Institute proposes that energy planning be formally integrated into the industrial policy process. This means that every major industrial policy initiative—whether it involves semiconductor fabrication, defense production, AI infrastructure, or other strategic manufacturing—should include an energy impact assessment that identifies the power requirements of the initiative, the available grid capacity, the infrastructure investments required, and the permitting timelines involved. This assessment should be conducted at the project planning stage, not after construction has begun, and should be used to inform both the industrial investment decision and the energy infrastructure planning process.

V. Regional Analysis: Where the Constraints Are Most Acute

The energy-industrial demand convergence is not uniformly distributed across the United States. Certain regions are experiencing disproportionate pressure due to the concentration of data center development, manufacturing investment, or both. Understanding the regional dimension of this challenge is essential to targeted policy intervention.

Northern Virginia and the PJM Interconnection

Northern Virginia hosts the largest concentration of data center capacity in the world. Dominion Energy, the principal utility serving the region, has reported that data center load requests exceed its available generation and transmission capacity. The PJM Interconnection, the regional transmission organization serving the mid-Atlantic and parts of the Midwest, has implemented new rules for large-load interconnection requests in response to the unprecedented volume of demand. The region faces a fundamental question of whether sufficient generation can be built, permitted, and connected to serve the continued growth of data center capacity, or whether the industry will be forced to disperse to other regions.

Central Texas

Texas is attracting both data center and manufacturing investment, including Samsung's semiconductor fabrication expansion. The ERCOT grid, which operates independently from the two major national interconnections, has experienced capacity challenges during extreme weather events and faces growing baseload demand from industrial customers. The combination of data center growth, manufacturing expansion, and the inherent constraints of an island grid creates a concentration of energy risk that warrants careful planning.

Central Arizona

TSMC's fabrication campus in Phoenix represents a major new industrial load in a region with constrained water and energy resources. The desert Southwest faces long-term challenges related to water availability, cooling requirements, and the integration of solar generation with industrial baseload demand. Ensuring that the energy infrastructure to support semiconductor fabrication is in place and reliable is essential to the success of what is arguably the most strategically important manufacturing investment in the CHIPS Act portfolio.

Ohio River Valley

Intel's massive fabrication investment in New Albany, Ohio, part of the $28 billion "Ohio One" campus, represents one of the largest single new industrial loads in the region. The project includes dedicated infrastructure such as a 500 MW substation built by AEP Ohio, though Intel's phased rollout has delayed full operations for advanced nodes to around 2031. Combined with surging data center development in Central Ohio from operators including Google, Meta, Amazon, and Microsoft, the area is seeing rapid baseload growth. Ohio ranks among the top U.S. states for data center energy consumption, with recent estimates at approximately 7.5 TWh annually.

The PJM Interconnection, which serves much of the Ohio River Valley, faces acute pressure in this region. American Electric Power (AEP Ohio) has reported large-load interconnection requests, primarily from data centers, in the 13 GW range through the early 2030s. Adjusted forecasts and new tariffs have reduced some projections amid debates over potential over-forecasting to justify grid upgrades. PJM's overall load is projected to rise dramatically, potentially adding tens of GW regionally by 2030 to 2035, with data centers as a primary driver. This strains a system transitioning from coal retirements to natural gas, renewables, and emerging storage.

This shift presents challenges, including transmission bottlenecks such as reliance on proposed lines from Indiana that may not energize until 2031, permitting delays for new generation, and risks of higher rates from over-forecasting. At the same time, it creates opportunities for dispatchable gas additions like the Trumbull Energy Center, on-site or near-site power solutions, and leveraging Ohio's natural gas abundance to provide reliable, competitive supply that attracts further manufacturing and compute investment.

VI. Recommendations

For Government

1. Establish a formal energy impact assessment requirement for all major industrial policy initiatives, ensuring that power requirements, grid capacity, and infrastructure timelines are evaluated at the project planning stage.
2. Create an interagency coordination mechanism linking the Department of Energy, Department of Commerce, Department of War (Formerly Department of Defense), and relevant regulatory agencies to align energy supply planning with industrial demand projections.
3. Accelerate permitting reform for energy generation and transmission infrastructure, with particular attention to projects that serve nationally significant industrial facilities such as semiconductor fabrication plants and defense production installations.
4. Support the development of advanced nuclear generation, including small modular reactors, as a dispatchable, carbon-free power source suited to the continuous, high-reliability loads of data centers and manufacturing.
5. Direct regional transmission organizations to incorporate industrial policy demand projections into their long-term transmission planning processes, moving beyond historical trend-based forecasting.

For Industry

1. Engage early and systematically with utilities and grid operators in site selection processes, incorporating energy availability and infrastructure timelines as first-order criteria rather than secondary considerations.
2. Invest in on-site and near-site generation where grid capacity is constrained, including long-term power purchase agreements for dedicated generation assets.
3. Collaborate with other large-load industrial customers in shared service territories to coordinate demand growth and infrastructure investment with utilities and regulators.

For Policy and Research Institutions

1. Develop integrated energy-industrial demand models that project cumulative power requirements across data centers, reshored manufacturing, and defense production on a regional basis.
2. Monitor and analyze the effectiveness of permitting reform efforts in reducing infrastructure deployment timelines for energy projects serving industrial customers.
3. Publish regular assessments of regional grid capacity relative to committed and projected industrial load to provide a public information baseline for policy and investment decisions.

Conclusion: Energy Is the Foundation of Industrial Strategy

The United States cannot build what it cannot power. The convergence of AI data center construction, reshored manufacturing, and defense production expansion is generating energy demand at a scale and pace that the existing grid and energy planning infrastructure are not prepared to accommodate. The investments being made in semiconductor fabrication, compute infrastructure, and defense industrial capacity are at risk of being stranded not by a lack of technology or capital, but by an inability to deliver the electric power required to operate them.

Energy is not a secondary consideration in industrial policy. It is the foundation. A national strategy that commits to rebuilding domestic manufacturing, achieving compute sovereignty, and maintaining defense readiness must simultaneously commit to the energy infrastructure required to power these objectives. This requires integrating energy planning with industrial planning, accelerating the deployment of generation and transmission infrastructure, and reforming the permitting processes that currently impose multi-year delays on critical energy projects.

The Atlas Institute and the Aegis Institute will continue to produce analysis that connects energy, manufacturing, and technology policy into the integrated framework that effective industrial strategy demands. The nation's ability to build, power, and govern its industrial future depends on treating these as dimensions of a single challenge rather than separate policy domains.

About the Institute for American Manufacturing and Technology

The Institute for American Manufacturing and Technology (IAMT) is a policy research organization dedicated to strengthening America's manufacturing capacity and technological competitiveness. The Atlas Institute focuses on energy systems, power generation, grid reliability, and the infrastructure required to support industrial growth.

About the Author: ***Patrick J. Wolf, PhD***, is a strategist and advisor working across technology, manufacturing, energy, and economic policy. He advises organizations in both the public and private sectors on strategy, industrial competitiveness, and the economic implications of technological change. Trained in industrial engineering and holding an MBA and a PhD in strategic media and communications, his work focuses on political economy, industrial strategy, and the institutional decisions that shape national competitiveness, manufacturing capacity, and long-term technological leadership.

© 2026 Institute for American Manufacturing and Technology. All rights reserved.