



AI INFRASTRUCTURE

Owning the Physical Economy Behind Artificial Intelligence

A W.G. Edwards Investments thematic research publication

AUGUST 2026

Artificial intelligence is digital. The infrastructure behind it is not.

POWER

GRID

DATA CENTERS

COOLING / ELECTRICAL

NETWORKING

SEMICONDUCTORS

COMPUTE

The physical AI stack

This publication is informed by the analytical framework developed in The Compute Economy - Updated Model, August 2026.

THE INVESTMENT OPPORTUNITY MAY EXTEND FAR BEYOND THE COMPANIES BUILDING AI MODELS

Some describe artificial intelligence as a software revolution. We believe that description is incomplete.

Every artificial-intelligence workload ultimately depends on a physical system capable of producing, transmitting, housing, cooling, connecting and operating enormous amounts of compute.

That system begins with electricity and extends through the electric grid, data centers, cooling systems, networking equipment, semiconductor manufacturing and ultimately the processors on which AI workloads run.

In other words, intelligence may be digital. Compute is physical.

That distinction forms the foundation of the W.G. Edwards Investments AI Infrastructure framework.

Our research suggests that the continued expansion of artificial intelligence may create a multi-year capital-investment cycle extending across a much broader group of industries than technology alone.

Who owns the infrastructure that AI cannot scale without?

FROM SOFTWARE REVOLUTION TO CAPITAL CYCLE

The first phase of the modern AI investment cycle was understandably focused on models, accelerators and hyperscale computing platforms. Those technologies remain essential.

But as AI systems become larger, more widely deployed and increasingly embedded into business processes, the physical requirements behind those systems become more visible.

A data center cannot operate without power. Power cannot reach the facility without transmission and interconnection capacity. Higher-density computing requires increasingly sophisticated electrical distribution and cooling. Large AI clusters require high-speed networking. Advanced processors require leading-edge semiconductor fabrication. And leading-edge semiconductor fabrication depends on some of the most sophisticated industrial equipment ever produced.

Weakness at any one of these layers can constrain the productive capacity of the entire system. That is why we increasingly view artificial intelligence not simply as a technology cycle, but as a physical infrastructure cycle.



The physical AI stack

THE BOTTLENECK MATTERS

One of the central principles behind our research is that economic value frequently accumulates where demand encounters supply that is difficult to expand quickly.

Not every part of the AI ecosystem has the same supply elasticity. Software can often scale relatively quickly. Physical infrastructure usually cannot.

Power plants require years to develop. Transmission projects can require long permitting and construction timelines. Data centers require land, power availability and specialized infrastructure. Transformers, switchgear and cooling equipment have manufacturing constraints. Advanced semiconductor plants cost tens of billions of

dollars. And some of the equipment required to manufacture leading-edge chips is produced by only a handful of companies globally.

This creates a fundamentally different investment dynamic. The question is not simply whether AI demand grows. The question is which layer becomes the limiting constraint as that demand grows.

Those constraints can change over time. Semiconductors may constrain supply during one phase. Power and interconnection may constrain it during another. Cooling, electrical equipment, networking bandwidth or advanced semiconductor manufacturing may become more important later.

The opportunity therefore is not based on identifying one permanent AI winner. It is based on understanding the entire physical system and recognizing where scarcity, earnings sensitivity and valuation intersect.

THE MISSING LAYER: SEMICONDUCTOR MANUFACTURING

Our original AI infrastructure framework emphasized power generation, electrical equipment, cooling, networking, data centers and AI processors. Further research identified an important structural gap.

Owning companies that design advanced AI processors provides exposure to compute demand. But those processors cannot exist without the manufacturing infrastructure required to produce them.

That manufacturing system introduces three particularly important layers: lithography; process control and metrology; and leading-edge fabrication and advanced packaging.

For that reason, semiconductor manufacturing has become a dedicated component of our AI Infrastructure framework.

ASML - Lithography

Advanced semiconductor production depends on extreme ultraviolet lithography. ASML occupies a highly specialized position within that manufacturing ecosystem.

As transistor geometries become smaller and semiconductor designs become more complex, lithography becomes increasingly important to producing leading-edge chips economically and at scale. High-NA EUV represents the next evolution of that manufacturing process.

The investment significance is straightforward: before the semiconductor industry can manufacture more advanced processors, it must possess the equipment capable of producing them.

KLA - Process Control

Manufacturing increasingly complex semiconductors requires far more than simply printing smaller circuits. Manufacturers must constantly inspect wafers, identify defects, measure processes and improve yields.

As complexity increases, the economic value of process control may increase with it. KLA provides inspection, metrology and process-control technology used throughout advanced semiconductor manufacturing.

The company's position is therefore somewhat different from a chip designer. It may benefit not simply from the number of chips produced, but from the increasing difficulty of producing those chips successfully.

TSMC - Advanced Manufacturing

Taiwan Semiconductor Manufacturing Company represents another critical layer. TSMC operates at the frontier of semiconductor fabrication and advanced packaging.

Its importance reflects a basic reality of the AI infrastructure system: designing an advanced processor is not the same as manufacturing one.

Leading-edge manufacturing requires enormous capital investment, sophisticated process knowledge, advanced equipment and accumulated production experience. These capabilities are difficult to reproduce quickly.

THE FULL AI INFRASTRUCTURE STACK

Power and Grid

AI infrastructure requires extraordinary quantities of dependable electricity. Companies exposed to generation, grid equipment, transmission construction and electrical distribution may therefore participate in the physical buildout required to support additional compute.

Electrical Distribution and Cooling

Power entering a data center must still be converted, distributed and managed. High-density AI computing generates substantial thermal loads. As rack density rises, cooling architecture and electrical infrastructure may become a larger portion of overall data-center capital spending.

Networking

Thousands of processors operating independently do not create an efficient AI cluster. Those processors must communicate at extremely high speeds. As AI clusters become larger, network bandwidth and latency become increasingly important.

Compute

AI accelerators and custom silicon perform the underlying computation required for training, inference and increasingly sophisticated reasoning workloads. Our framework intentionally places compute inside a larger system rather than treating processors as the complete AI investment opportunity.

Data Centers

Data centers convert land, power, cooling, networking and computing equipment into functioning infrastructure. As AI workloads increase in density, power availability, interconnection, geography and physical design can become increasingly important determinants of value.

Natural Gas and Generation Infrastructure

The rapid addition of electricity demand can affect upstream energy infrastructure. Natural gas remains an important fuel source for dispatchable generation in many regions, while nuclear and other dependable generation may become increasingly valuable where data-center operators prioritize reliability.

Materials

The physical AI buildout requires substantial quantities of industrial materials. Copper is particularly important across electrical equipment, transmission, transformers, data centers, cooling systems, networking and semiconductor manufacturing.

AN ILLUSTRATIVE PORTFOLIO ARCHITECTURE

To translate the research into an investable framework, W.G. Edwards Investments maintains an illustrative AI Infrastructure Model Allocation.

The largest allocations remain concentrated in areas where we believe physical capacity can require significant time and capital to expand.

Importantly, diversification across these infrastructure categories should not be confused with broad portfolio diversification. These businesses remain connected to a common investment theme and could become highly correlated if AI capital spending, economic growth or equity valuations weaken simultaneously.

The allocation is therefore best understood as a thematic research model, not as a complete portfolio appropriate for every investor.

ECONOMIC FUNCTION	ILLUSTRATIVE WEIGHT
Electrical / Cooling / Grid	29%
Compute / Networking	23%
Semiconductor Manufacturing	10%
Generation / Utilities	9%
Hyperscale Platforms	8%
Data Centers	7%
Gas Infrastructure	7%
Materials	3%
Strategic Optionality	4%

THREE DIFFERENT TYPES OF EXPOSURE

Structural Bottlenecks

Businesses associated with infrastructure or manufacturing capabilities that may be difficult to replace or expand rapidly. Examples can include electrical systems, cooling infrastructure, AI accelerators, advanced networking and semiconductor manufacturing.

Enabling Infrastructure

Businesses that provide capacity required for the physical buildout but may operate in markets where scarcity is more regional, regulatory or construction-driven. Examples include grid construction, generation, natural-gas infrastructure and data-center capacity.

Demand and Materials

Companies that participate through hyperscale AI demand or through the materials required to build the infrastructure. Their earnings exposure to AI may be less direct but can still become meaningful if capital spending remains elevated.

Strategic Optionality

Companies whose potential AI infrastructure role is significant but considerably more uncertain. SpaceX represents this type of exposure in our current research framework. Its combination of launch capability, satellite connectivity, artificial intelligence activities and potential future orbital-compute infrastructure represents a differentiated long-duration possibility, but technological, execution, financing, regulatory and valuation uncertainty remains substantially greater than for established infrastructure businesses.

WHAT COULD GO WRONG?

AI Capital Spending Could Slow

If AI monetization fails to justify current expenditure levels, hyperscaler capital budgets could decline. That would affect companies across the infrastructure chain.

Efficiency Could Improve Faster Than Demand

AI processors, software and model architectures may become substantially more efficient. If efficiency gains reduce aggregate physical infrastructure requirements faster than new workloads emerge, projected capacity needs could prove too high.

The Industry Could Overbuild

Periods of scarcity often attract capital. New semiconductor capacity, data centers, generation assets and electrical-equipment manufacturing could eventually reduce shortages. Scarcity is rarely permanent.

Valuations Matter

A strong industry does not automatically produce a strong investment. Even companies with excellent competitive positions can generate disappointing shareholder returns if investors pay too high a price.

Technology Can Change the Architecture

New processors, cooling technologies, network architectures or computing approaches could shift economic value between layers. Today's bottleneck may not remain tomorrow's bottleneck.

Geopolitical and Regulatory Risks Are Significant

Export controls, trade restrictions, political tensions, geographic concentration, environmental rules and permitting constraints can materially affect semiconductor, power and data-center development.

THE SCENARIO FRAMEWORK

BEAR CASE

AI returns on investment disappoint. Capital spending slows. Semiconductor demand becomes more cyclical. Projects are delayed and valuation multiples contract. Under this scenario, balance-sheet quality, diversification and valuation discipline become more important.

BASE CASE

AI adoption continues to broaden. Inference, reasoning and agent-based workloads increase utilization. Investment in power, networking, data centers and semiconductor manufacturing remains elevated but does not accelerate indefinitely. Under this scenario, diversified exposure to the physical infrastructure stack may continue to participate in the buildout.

BULL CASE

Machine-generated workloads compound faster than efficiency gains. Power, networking, semiconductor manufacturing and data-center capacity become increasingly constrained. Infrastructure investment accelerates. Under this scenario, businesses positioned at difficult-to-expand bottlenecks could experience greater earnings sensitivity to the AI capital cycle.

These scenarios are analytical frameworks, not predictions. Actual outcomes may differ materially.

FROM ELECTRONS TO INTELLIGENCE

Artificial intelligence ultimately sits at the end of an enormous physical production chain.

Electricity must be generated. The grid must deliver it. Data centers must house the equipment. Electrical systems must distribute the power. Cooling systems must remove the heat. Networks must connect the processors. Semiconductor equipment must manufacture the chips. Foundries must fabricate them. Only then can compute be produced.

The investment implications of artificial intelligence therefore extend much further than software.

We believe one of the defining capital-allocation questions of the coming decade may be: where does rapidly expanding demand for intelligence collide with slowly expanding physical capacity?

That intersection is where scarcity can emerge. Scarcity can create pricing power. And pricing power, when combined with strong competitive positions and disciplined valuations, can create investment opportunity.

The objective of the W.G. Edwards Investments AI Infrastructure framework is not to identify one permanent group of AI winners. It is to understand the physical system that makes artificial intelligence possible - and to follow the bottlenecks as that system evolves.

INVESTMENT CONCLUSION

Artificial intelligence may ultimately become one of the largest technological transitions of the modern economy. But the infrastructure required to support that transition is decidedly physical.

We therefore believe investors evaluating the AI opportunity should consider not only the companies developing intelligence, but also the businesses responsible for producing the electricity, equipment, facilities, networks and semiconductors required to deliver it.

Artificial intelligence is digital. The infrastructure behind it is not.

That physical infrastructure may represent one of the most important capital-investment cycles of the decade.

POWER

GRID

DATA CENTERS

COOLING / ELECTRICAL

NETWORKING

SEMICONDUCTORS

COMPUTE

IMPORTANT DISCLOSURES

This publication is provided for informational and discussion purposes only and does not constitute individualized investment advice, an offer to sell, or a solicitation to purchase any security.

Any companies, securities, portfolio weights or allocations discussed are illustrative components of the W.G. Edwards Investments research framework and should not be interpreted as recommendations appropriate for every investor. Actual portfolios may differ materially, and holdings, allocations, estimates and opinions may change without notice.

The AI Infrastructure Model is a concentrated thematic investment framework and may experience materially greater volatility than a broadly diversified portfolio. Securities associated with a common investment theme may become highly correlated during periods of market stress.

Forward-looking statements, forecasts, scenarios, growth estimates and analytical conclusions involve substantial uncertainty. They reflect assumptions and judgments made by W.G. Edwards Investments and are not guarantees of future results. Actual results may differ materially because of economic conditions, technological developments, competition, regulation, interest rates, capital spending, semiconductor cycles, energy availability, geopolitical developments, export controls and other factors.

Investing involves risk, including possible loss of principal. International securities may involve additional currency, political and regulatory risks. Semiconductor, technology, energy, utility, real-estate and infrastructure companies each involve industry-specific risks.

Past performance is not indicative of future results.

Third-party information is obtained from sources believed to be reliable, but its accuracy or completeness cannot be guaranteed.

Investors should consider their individual objectives, risk tolerance, time horizon, liquidity requirements, diversification needs, tax circumstances, existing holdings and account restrictions before implementing any investment strategy.

This publication is informed by the analytical framework developed in W.G. Edwards Investments - The Compute Economy - Updated Model, August 2026.

COMPLIANCE APPROVAL CODE: PM-02122028-5834770