



DATA CENTERS



AI INFRASTRUCTURE



POWER & COOLING



CONNECTIVITY



POLICY & MARKETS

THE NEW COMPUTE EPOCH

POWERING INTELLIGENCE. BUILDING THE FUTURE.

The Infrastructure Backbone of
the AI Era and What Comes Next



UNPRECEDENTED POWER DEMAND

AI is driving a historic surge in power requirements, reshaping grids, markets, and policy.



THE DATA CENTER ARMS RACE

New infrastructure paradigms are emerging at a scale and speed the world has never seen.



CAPITAL, GEOGRAPHY, STRATEGY

Winners will be defined by access to power, strategic locations, and patient capital.



RISKS, CONSTRAINTS, CONSEQUENCES

From grid reliability to geopolitics, the challenges ahead are complex and interdependent.

**\$500B+**

Projected data center infrastructure spend by 2030

Source: Goldman Sachs

**2.5–3.5x**

Increase in U.S. data center power demand by 2030

Source: EPRI, McKinsey

**Global Impact**

A structural shift with far-reaching economic, energy, and societal consequences

Source: McKinsey

INFRASTRUCTURE.
INTELLIGENCE.
IMPACT.

CRITICALSTACK.COM

01 | EXECUTIVE OVERVIEW

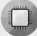







THE COMPUTE STACK HAS DECOUPLED

Silicon still matters. But software, orchestration, memory systems, and inference optimization now create independent performance cycles across the AI infrastructure stack.

For decades, performance, cost decline, and adoption were paced primarily by silicon scaling. That era is over. Dennard scaling faded by ~2012 and Moore's Law has slowed materially. Today, AI performance advances on a second, increasingly independent clock driven by model architecture, data efficiency, compilers, kernels, runtimes, and system design.

This decoupling changes where value is created and captured. The winners will be those who control scarce physical constraints and convert them into higher utilization, lower total cost of ownership, and better intelligence at scale.

FROM A SINGLE CLOCK TO MULTIPLE, INDEPENDENT CYCLES

OLD COMPUTE CYCLE	NEW COMPUTE EPOCH
 CPU / NODE SCALING Performance lifted by transistor shrinks.	 FULL-STACK OPTIMIZATION Software, models, and systems unlocked on every hardware gen.
 MONOLITHIC COMPUTE General-purpose processors dominated workloads.	 HETEROGENEOUS SYSTEMS GPUs, TPUs, ASICs, CPUs, DPUs, HBM, and advanced networking.
 HARDWARE-PACED PERFORMANCE Upgrades came with each silicon generation.	 INDEPENDENT SOFTWARE CYCLES Compilers, kernels, runtimes, and models move faster than silicon.
 PROCESSOR-CENTRIC VALUE CAPTURE Value accrued to the fastest chip and its roadmap.	 INFRASTRUCTURE-CENTRIC VALUE CAPTURE Value flows to bottlenecks and utilization leverage.

KEY INSIGHT

“
 The new compute race is no longer won by hardware alone. It is won by whoever converts constrained infrastructure into usable intelligence most efficiently.”

— CRITICAL STACK RESEARCH

BY THE NUMBERS

 **280x**
 Decline in GPT-3.5-class inference cost between Nov 2022 and Oct 2024
 Source: Stanford AI Index 2025

 **3.0x**
 Annual improvement in pre-training compute efficiency
 Source: Epoch AI (2025)

 **\$500B+**
 Projected AI infrastructure spend over the next several years
 Source: Company filings, analyst estimates

 **2–5 YR**
 Grid interconnection delays in many major markets
 Source: EPRI, McKinsey (2024)



KEY TAKEAWAY


Compute is no longer constrained by transistor scaling alone. The new bottleneck is coordinated infrastructure efficiency across silicon, software, networking, memory, and power.

02 | SILICON CONSTRAINTS














SILICON IS NO LONGER ON AN EASY CURVE.

Dennard scaling is gone. Moore's Law has slowed. And the economics of advanced silicon are now defined by physics, complexity, and capital intensity.

Transistor scaling delivered decades of exponential gains. That era is over. Today, performance improvements are harder, more expensive, and increasingly dependent on system-level innovation, advanced packaging, and new materials. Silicon remains foundational—but it is no longer the only or even primary driver of compute progress.



THE NEW PHYSICS OF COMPUTE

	1. MOORE ERA ~1970s–2000s	2. MULTICORE ERA ~2000s–2015	3. AI INFRASTRUCTURE ERA ~2015–BEYOND
	SCALING WAS SIMPLE	SCALING HIT LIMITS	SYSTEMS DRIVE PERFORMANCE
	 Smaller transistors, lower cost, higher performance.	 Frequencies stalled. Multicore and parallelism became necessary.	 Heterogeneous systems, software, and data movement unlock the next gains.
 TRANSISTOR SCALING	Rapid (Moore's Law)	Slowing (Power & leakage constraints)	Incremental (Atomic & material limits)
 POWER DENSITY	Manageable	Rising (Thermal wall)	Binding Constraint (Power is the limit)
 ECONOMICS (COST / TRANSISTOR)	Falling Exponentially	Flattening	Rising (Complexity & capex)
 VALUE MIGRATION	Silicon 	Architectures & Software 	System-Level: Packaging, Memory, Networking, Power, Software    


THE IMPLICATION

Compute performance is now gated by a system of constraints—power, thermal, packaging density, memory bandwidth, and economics—not just transistor scaling.

WHERE VALUE IS MIGRATING

LITHOGRAPHY

Precision tools are the new choke point. Capacity, not just technology, is scarce.


HBM & MEMORY

Bandwidth per watt is now a primary differentiator. HBM is the new performance lever.


ADVANCED PACKAGING

Chiplets, 2.5D/3D, and CoWoS drive density and performance. Packaging is becoming strategic.


NETWORKING

Scale-out performance depends on bandwidth, latency, and efficiency.


POWER-ENABLED SITES

Power availability, grid access, and cooling define where compute can scale.



SILICON REALITIES



TSMC CAPEX ESCALATION
\$32B → \$40B+
 2023 to 2025E

Leading-edge fabrication requires unprecedented and compounding capital.

Source: TSMC, Company Filings



ASML HIGH-NA EUV
\$380M
 Per system

High-NA EUV systems are ~2.5x the cost of standard EUV, with limited annual supply in the near term.

Source: ASML, SemiAnalysis



ADVANCED PACKAGING
 VALUE SHARE
~20%

of leading-edge chip cost by 2027E

Up from ~7% in 2022 as 2.5D/3D packaging becomes mission critical.

Source: Yole Group



EXPORT CONTROLS
 EXPANDING
3x

Increase in restricted items since 2022

Broader scope, more countries, more layers of the stack.

Source: CSIS, BIS

03 | SOFTWARE PERFORMANCE LAYER


SOFTWARE IS NOW ITS OWN SCALING ENGINE.

Performance gains in AI are increasingly driven by software innovation, not just silicon. Compilers, kernels, runtimes, and model optimizations now deliver independent performance cycles that compound faster than hardware advances.

From training efficiency to inference throughput, the software stack has become a force multiplier—extracting more from the same hardware, more quickly and at lower cost.

THE SECOND PERFORMANCE CLOCK

Independent software cycles compound within each hardware generation.


WHAT DRIVES SOFTWARE CYCLES

- Compilers & Graph Optimization
- Kernels & Algorithmic Innovation
- Runtime & Memory Management
- Inference Orchestration & Serving
- Model Design & Efficiency
- Quantization & Compression

SOFTWARE EFFICIENCY SIGNALS

280x

Decline in GPT-3.5-class inference cost
 Nov 2022 → Oct 2024

Driven by software optimizations across the stack.

Source: Stanford AI Index 2025

43%

Training speedup with PyTorch 2.x 'torch.compile' on common LLM workloads.

Source: PyTorch Blog (2024)

2-4x

Higher inference throughput using optimized serving engines like vLLM.

Source: vLLM Benchmark (2024)

20x

Memory efficiency improvement with FlashAttention vs. naive attention.

Source: Tri Dao et al. (2023)

1.5-3x

End-to-end cost reduction from quantization and model compression techniques.

Source: MLCCommons (2024)

WHERE SOFTWARE CREATES VALUE

HIGHER CLUSTER UTILIZATION 85% UTILIZATION Better scheduling and resource management extract more from existing infrastructure.	LOWER COST PER TOKEN Software efficiencies drive structural declines in cost across the model lifecycle.	MEMORY EFFICIENCY Advanced algorithms reduce HBM pressure and expand effective context length.	HIGHER INFERENCE THROUGHPUT Optimized serving stacks increase tokens/sec and improve latency at scale.	DEPLOYMENT PORTABILITY Better abstractions make models easier to deploy across diverse environments.	SYSTEM-WIDE ORCHESTRATION Cross-layer coordination maximizes performance, reliability, and cost efficiency.
--	--	--	--	--	---


KEY TAKEAWAY

Software has become an independent scaling engine. It compounds faster than hardware, delivers outsized economic impact, and is now a primary determinant of AI infrastructure ROI.

05 | POWER CONSTRAINT

POWER IS THE NEW BOTTLENECK.

AI infrastructure growth is no longer limited by chips or algorithms. It is constrained by the real-world availability of electricity.

From interconnection queues to transmission buildouts, power has become the critical path. Without sufficient, reliable, and affordable power, compute capacity cannot be deployed—no matter how much capital or silicon is available.

ELECTRICITY IS THE NEW SCARCITY	POWER LEAD TIMES ARE LENGTHENING RAPIDLY	AI WORKLOADS ARE DRIVING EXPLOSIVE GROWTH	THE GAP IS WIDENING FAST
--	---	--	---------------------------------

GRID REALITY

INTERCONNECTION WAIT TIMES
2-5 YEARS

For large (>100MW) projects in major markets.

Source: Berkeley Lab (2024)


U.S. QUEUED CAPACITY
2,600+ GW

More than double current U.S. generation capacity.

Source: Lawrence Berkeley National Lab (2024)


U.S. DATA CENTER POWER DEMAND BY 2030
2.5-7.3x

Growth by 2030 vs. 2023 (EPRI high-demand range).

Source: EPRI (Dec 2024)


UTILITY CAPEX ACCELERATION
\$400B+

Annual U.S. utility capex by 2030 (up from \$150B in early 2020s).

Source: BCG (2024)


TRANSFORMER LEAD TIMES
50-130 WEEKS


For high-voltage transformers (major constraint globally).

Source: Grid Strategies (2024)

THE POWER GAP
AI & DATA CENTER POWER DEMAND VS. GRID EXPANSION

— AI & Data Center Power Demand (High Growth Scenario)

- - - Grid Capacity Additions (Committed + Likely)


~100-120 GW
 AI + Data Center Demand by 2030

THE POWER GAP
~25-35 GW
 Grid Capacity Additions by 2030

Hyperscale AI buildout accelerates

Inference drives 24/7 power demand

Interconnection queues exceed 3 years

Transmission buildouts lag demand

Power becomes the binding constraint

Source: SemiAnalysis, EPRI, BCG, Lawrence Berkeley National Laboratory, Company Filings

POWER SOLUTIONS EMERGING

Natural gas, dual-fuel, and microgrids provide immediate capacity.

Nuclear offers 24/7, carbon-free baseload at scale.

Fast-to-build capacity to back intermittent demand spikes.

Smoothing demand curves and supporting grid reliability.

Higher density compute requires advanced thermal solutions.

Transmission, HVDC, and grid upgrades are critical to close the gap.

Locating data centers near power sources reduces constraints.

KEY TAKEAWAY

Compute can be bought.
 Power must be built.

The AI race may ultimately be won by
whoever secures power first.



06 | GEOGRAPHY & STRATEGIC SHIFTS

INFRASTRUCTURE IS MOVING TO WHERE POWER, POLICY, AND PROXIMITY ALIGN.

AI infrastructure is no longer distributed evenly. It is concentrating in regions with the right combination of power availability, policy support, grid access, and connectivity.

Geography is now a competitive advantage.

EMERGING INFRASTRUCTURE HOTSPOTS

Power availability, grid readiness, and strategic advantages

- Power Abundance
- Policy & Permitting Advantage
- Hyperscaler Expansion
- Sovereign AI Initiatives
- Connectivity Corridors

CANADA

- Clean power
- Abundant land
- Low latency to U.S.

U.S. SUN BELT

- Low-cost power
 - Fastest grid growth
 - Pro-business policy
 - Large land parcels
- Major expansion:
TX, OK, LA, AR, TN, AZ

NORDICS

- Renewable power
- Cool climate
- Stable regulation
- Strong fiber

MIDDLE EAST

- Massive power buildout
- Capital availability
- Strategic location
- Sovereign AI ambitions


SOUTHEAST ASIA

- Rising power capacity
- Growing digital demand
- Manufacturing proximity
- Policy incentives

— Major Subsea Fiber Routes - - - Emerging Fiber Routes ● Landing Stations ○ Key Data Center Corridors

TREND: IS FOLLOWING THE ADVANTAGE


Announced data center investments (2024–2030)



Source: DC Byte, SemiAnalysis

SUBSEA CONNECTIVITY EXPANDING

New & planned subsea cable system capacity



Source: TeleGeography (2024)

POWER AVAILABILITY IS THE DIFFERENTIATOR

Estimated power cost for large-scale data centers (USD/MWh)



Source: BloombergNEF (2024)

THE SHIFT IS GLOBAL

GLOBAL DATA CENTER CAPACITY (GW)

83

2024

PROJECTED (GW) BY 2030

219

+164%

NEW CAPACITY OUTSIDE TRADITIONAL MARKETS

>60%

By 2030

Source: SemiAnalysis, DC Byte, McKinsey

STRATEGIC SHIFTS UNDERWAY


SOVEREIGN AI INVESTMENT SURGE

\$100B+

Public and sovereign funds committed to national AI infrastructure buildouts.

Source: CNBC, Global SWF



HYPERSCALER GEOGRAPHIC DIVERSIFICATION

70%+

Of new capacity announcements in the past 24 months are outside traditional tier-1 markets.

Source: DC Byte (2024)



EDGE & REGIONAL INFERENCE GROWTH

3–5x

More inference compute deployments at the edge by 2030.

Source: Gartner (2024)



INDUSTRIAL POLICY COMPETITION

>30

Countries with active programs to attract AI/data center investment.

Source: OECD, BCG (2024)



CO-LOCATION WITH GENERATION

2x+

Increase in planned data center projects with on-site or adjacent power generation.

Source: Wood Mackenzie (2024)

WHY GEOGRAPHY MATTERS AGAIN

ABUNDANT POWER Access to low-cost, reliable, and scalable electricity is the ultimate constraint breaker.	POLICY CERTAINTY Fast permitting, incentives, and stable regulation reduce execution risk and speed to online.	FIBER PROXIMITY Access to dense subsea networks and IXPs lowers latency and boosts performance.	COOLER CLIMATES Lower cooling requirements improve efficiency and reduce operating costs.
INDUSTRIAL LAND Large, affordable land parcels enable hyperscaler buildouts at speed and scale.	GEOPOLITICAL STABILITY Political alignment and security reduce long-term risk and support sovereign priorities.		

KEY TAKEAWAY

AI infrastructure is becoming a geopolitical asset.
 The next decade will be won by those who control the right geography.

07 | INVESTMENT IMPLICATIONS

WHERE DURABLE VALUE WILL ACCRUE.

AI infrastructure is a full-stack industrial economy. Value accrues to the layers with scarcity, leverage, and structural importance—not to the most headlines.

THE AI INFRASTRUCTURE INVESTMENT THESIS



SCARCITY
Limited supply relative to explosive demand growth.



LEVERAGE
High utilization leverage or pricing power.



ECOSYSTEM CONTROL
Critical to others' ability to build and scale.



DURABILITY
Structural tailwinds that persist across cycles and geos.

CAPITAL IS MOVING TOWARD



HBM & MEMORY
~\$160B
HBM market by 2028 (>\$40B in 2024)

Source: TrendForce (Apr 2025)



NETWORKING & OPTICS
~\$55B

DC networking market by 2028 (2.5x vs. 2024)

Source: Dell'Oro (Mar 2025)



POWER INFRASTRUCTURE
\$500B+

U.S. grid capex through 2030 (vs. ~\$300B prior plan)

Source: EPRI (Dec 2024)



AI DEPLOYMENT TOOLING
~\$25B+

AI infra software TAM by 2028 (3x vs. 2024)

Source: BCG (Apr 2025)



INFERENCE OPTIMIZATION
>70%

Of compute spend by 2028 will be inference

Source: SemiAnalysis (Apr 2025)



SOVEREIGN & NATIONAL AI
>30

Countries with announced sovereign compute initiatives

Source: OECD.AI (Apr 2025)



CO-LOCATION WITH POWER
2-4x

More deployments planned near generation by 2030 vs. today

Source: McKinsey (Apr 2025)

THE NEW AI VALUE STACK

WHAT IT INCLUDES		VALUE DYNAMICS			STRUCTURAL OUTLOOK
		SCARCITY	MARGIN POWER	DURABILITY	
	POWER & ENERGY Generation, transmission, grid access, onsite power, nuclear, gas, storage	●●●●	●●●●	●●●●	↑ STRONG TAILWIND
	NETWORKING FABRIC Switches, routers, optics, DCI, fiber, across and within data centers	●●●●	●●●●	●●●●	↑ STRONG TAILWIND
	MEMORY & BANDWIDTH HBM, DRAM, high-bandwidth interfaces, memory controllers	●●●●	●●●●	●●●●	↗ FAVORABLE STRUCTURAL
	SILICON & PACKAGING GPUs, ASICs, advanced packaging, CoWoS, 2.5D/3D	●●●	●●●	●●●	→ CYCLICAL ADVANTAGE
	INFERENCE SOFTWARE Model optimization, runtimes, compilers, model delivery	●●●	●●●●	●●●●	↑ STRONG TAILWIND
	ORCHESTRATION LAYER Kubernetes, schedulers, observability, fleet mgmt, FinOps	●●●	●●●	●●●	↗ FAVORABLE STRUCTURAL
	INFRASTRUCTURE DEPLOYMENT Data centers, cooling, real estate, construction, EPC, modular	●●●	●●●	●●●	→ CYCLICAL ADVANTAGE
	PLATFORM & CLOUD Hyperscaler clouds, APIs, AI platforms, data platforms	●●●	●●●	●●●	→ MARGIN COMPRESSION RISK
	APPLICATIONS & END MARKETS Vertical AI apps, consumer apps, agents, enterprise solutions	●●	●●	●●	↓ HIGH COMMODITIZATION RISK

●●●● Very High ●●● High ●● Moderate ● Low Outlook: ↑ Positive → Neutral ↓ Negative

QUESTIONS INVESTORS SHOULD ASK

WHO CONTROLS THE BOTTLENECKS? <ul style="list-style-type: none"> • Where is supply tight? • What cannot be easily replicated? • What has long lead times? 	WHO IMPROVES UTILIZATION? <ul style="list-style-type: none"> • Who drives more tokens per GPU/CPU? • Who reduces cost per inference? • Who extends hardware useful life? 	WHO OWNS ECOSYSTEM LEVERAGE? <ul style="list-style-type: none"> • Who is embedded in mission-critical layers? • Who benefits when the ecosystem grows? • Who is the default standard? 	WHO BENEFITS FROM INFERENCE SCALE? <ul style="list-style-type: none"> • Who captures the bulk of inference economics? • Who has pricing power as volume explodes? • Who monetizes data flywheel effects? 	WHO HAS DEFENSIBLE MOATS? <ul style="list-style-type: none"> • What moats are structural vs. cyclical? • Who has switching costs? • Who has regulatory or sovereign edges? 	WHO IS EXPOSED TO COMMODITIZATION? <ul style="list-style-type: none"> • Where are margins structurally compressing? • Where is capacity oversupplied? • What is easiest to replicate?
---	--	---	--	--	---

KEY TAKEAWAY

The next decade of AI value will be created in infrastructure— and captured by those who **control the scarce, the essential, and the enduring.**

08 | SIGNALS, RISKS & WHAT HAPPENS NEXT

THE NEXT DECADE WILL BE DEFINED BY INFRASTRUCTURE EXECUTION.

The AI era is no longer speculative. It is physical, operational, and geopolitical. The bottlenecks are real. The capital is mobilizing. The winners will be determined by execution, access, and ecosystem leverage.

Compute has become infrastructure.

WHAT WE'RE WATCHING



INFERENCE & TOKEN ECONOMICS

Token growth, price per 1M tokens, and unit economics convergence.

Trend: ↻



HBM CAPACITY EXPANSION

New capacity, yield improvements, and bandwidth per dollar.

Trend: ↻



UTILITY INTERCONNECTION TIMELINES

Grid connection queues, transformer lead times, and transmission buildout.

Trend: ↻



SOVEREIGN AI CAPEX

National AI initiatives, sovereign clouds, and strategic compute programs.

Trend: ↻



OPEN-WEIGHT MODEL CONVERGENCE

Model performance, adoption curves, and inference cost declines.

Trend: ↻



POWER-ENABLED LAND PRICING

Prime sites near power, cooling, fiber, and industrial ecosystems.

Trend: ↻



AI NETWORKING UTILIZATION

Switching capacity, east-west traffic, and cluster-scale efficiency.

Trend: ↻

THE NEXT DECADE OF SIGNALS



THE EXECUTION GAP WILL DETERMINE THE LEADERS



FIVE STRATEGIC CONCLUSIONS

01 COMPUTE IS NOW INFRASTRUCTURE.
AI infrastructure is a core economic input—like power, transport, and communication.

POWER MATTERS AS MUCH AS SILICON.
Access to affordable, reliable power is the ultimate competitive advantage.

SOFTWARE COMPOUNDS HARDWARE VALUE.
The software layer—models, inference, and orchestration—drives utilization and ROI.

GEOGRAPHY IS BECOMING STRATEGIC.
Infrastructure is clustering where power, policy, and proximity align.

EXECUTION, NOT EXPOSURE, WINS.
The AI economy rewards those who build, operate, and scale—consistently.

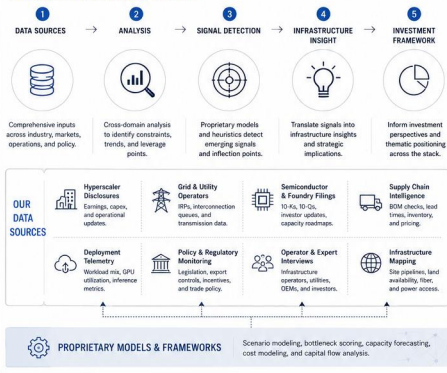
KEY TAKEAWAY

This is not a short cycle.
It is a long-duration restructuring of the digital infrastructure economy.

09 | METHODOLOGY, SOURCES & SIGNALS

METHODOLOGY, SOURCES, AND SIGNALS WE'RE TRACKING.

Critical Stack combines multi-source intelligence, operator insights, and proprietary analysis to identify infrastructure bottlenecks, capital flows, and long-term strategic inflection points shaping the AI economy.

THE CRITICAL STACK RESEARCH ENGINE

REPORT AT A GLANCE

120+
Datasets Analyzed
 Industry, market, operator, and policy datasets.

300+
Companies Tracked
 Across semiconductors, infrastructure, software, utilities, and networks.


25+
Infrastructure Regions Monitored
 North America, Europe, Middle East, Asia, and emerging markets.

200+
Operator & Expert Conversations
 Across hyperscalers, utilities, OEMs, foundries, and investors.

15+
Policy Frameworks Followed
 U.S., EU, China, UK, India, GCC, and more.

7
Signal Categories Continuously Tracked
 Power, cooling, compute, networking, supply chain, policy, and geography.

DISCIPLINED. INDEPENDENT. DATA-DRIVEN.
 Our analysis is independent and uncompromised. We do not accept compensation for coverage or viewpoints.

SIGNALS WE'RE TRACKING


These signals feed our forward-looking models and shape our long-term infrastructure outlook.

THE INFRASTRUCTURE ERA HAS BEGUN.

The AI economy is no longer hypothetical. It is becoming physical, operational, and global. The next decade will be defined by the infrastructure we build—and how we execute. Compute became infrastructure. Execution will determine outcomes.

We track the systems.
You build the future.

OUR MISSION



CLARITY IN COMPLEXITY

We bring clarity to the most complex infrastructure transformations of our time.



INTELLIGENCE THAT MATTERS

We deliver real intelligence on systems, capital flows, and strategic inflection points.



INFRASTRUCTURE FIRST

We focus on the physical and digital foundations powering the AI economy.



ACTIONABLE INSIGHT

We empower investors, operators, and builders with intelligence they can use.



INDEPENDENT. RIGOROUS. GLOBAL.

Our analysis is independent, research-driven, and built for long-term decision-making.



Infrastructure is the multiplier of every technological era. The winners will be those who build, connect, and power what others imagine.

CRITICAL STACK

Intelligence for the
Infrastructure Economy

STAY CONNECTED



VISIT OUR PLATFORM
criticalstack.com



FOLLOW ON LINKEDIN
linkedin.com/company/criticalstack



SUBSCRIBE FOR FUTURE REPORTS
criticalstack.com/subscribe



UPCOMING REPORTS
Quarterly intelligence on the forces reshaping the infrastructure economy.



CONTACT US
research@criticalstack.com

*New signals. Deeper insights.
Every quarter.*

TOPICS WE'RE TRACKING NEXT



POWER MARKETS

Grid capacity, energy markets, and the race to secure abundant, reliable power.



SOVEREIGN COMPUTE

National AI strategies, sovereign clouds, and strategic compute independence.



AI NETWORKING

High-speed interconnects, optical innovation, and the next generation of network fabrics.



COOLING SYSTEMS

Liquid cooling adoption, thermal design, and efficiency breakthroughs at scale.



EDGE INFRASTRUCTURE

Edge compute growth, distributed architectures, and enterprise deployment strategies.



ADVANCED PACKAGING

2.5D/3D integration, chiplet ecosystems, and packaging capacity constraints.