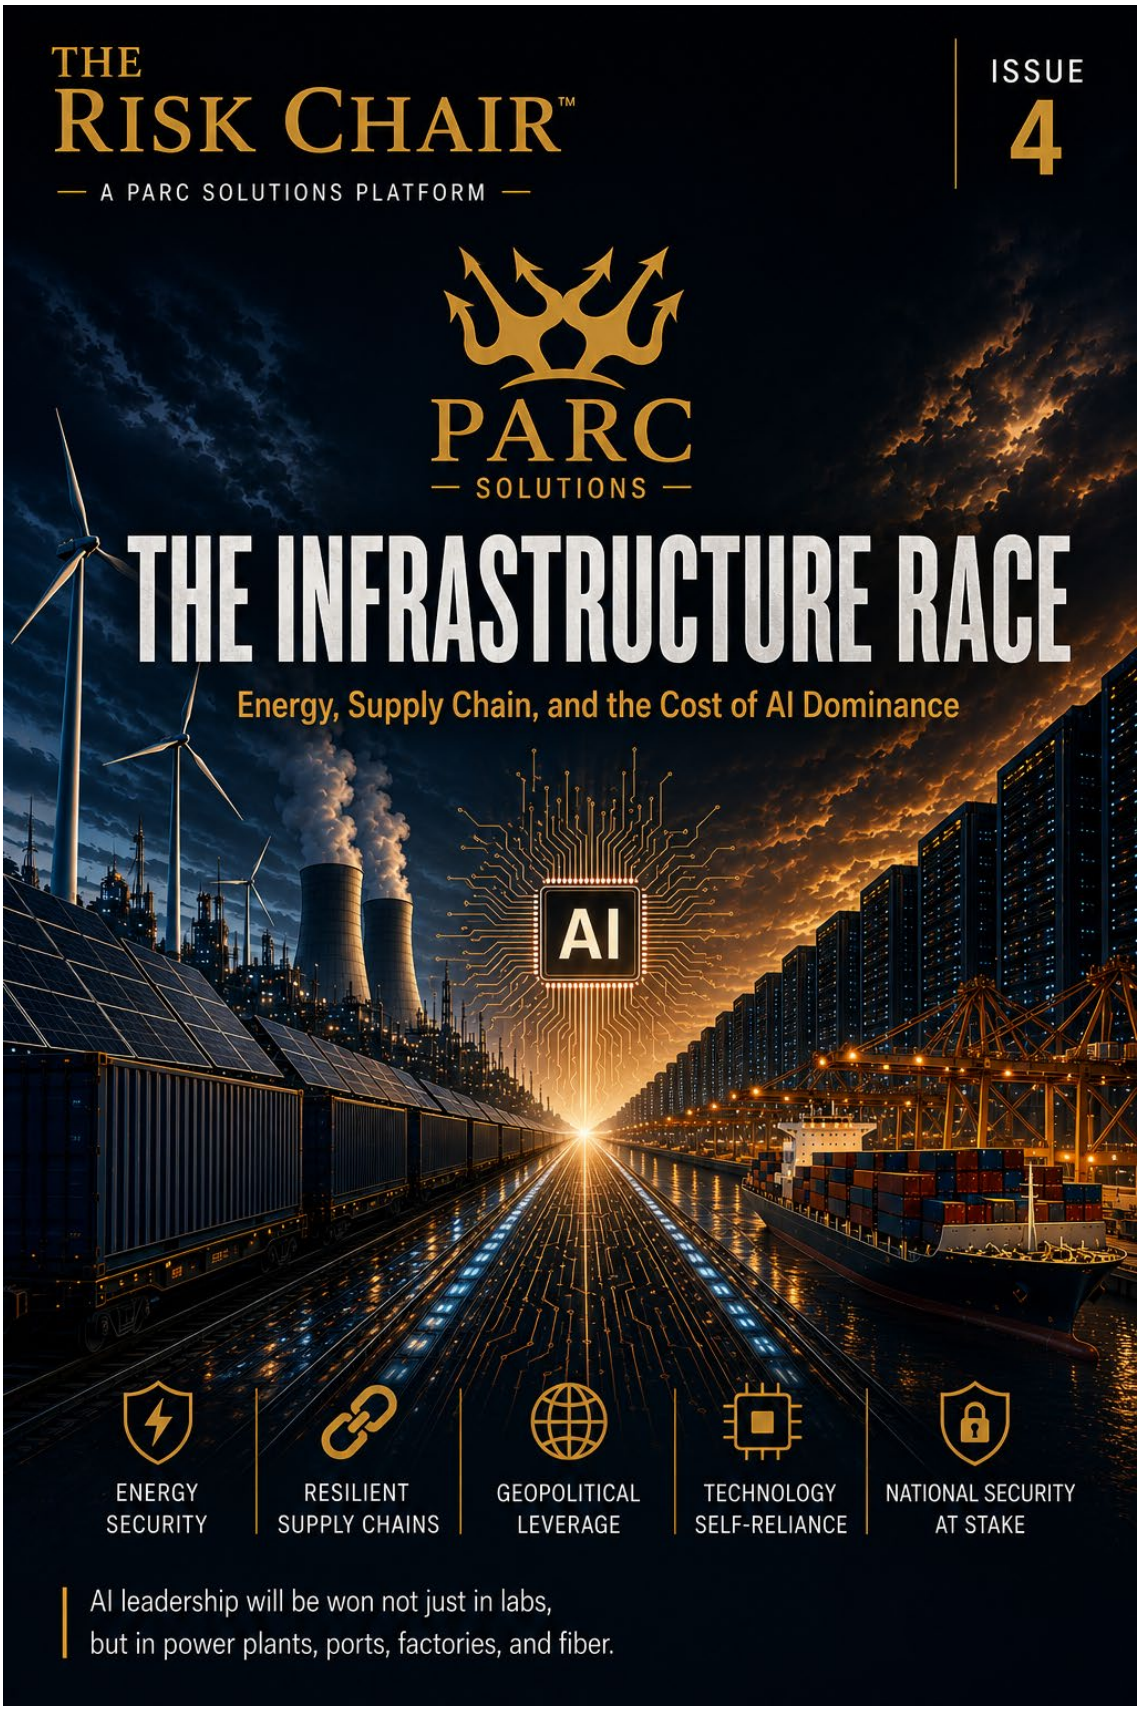
THE
RISK CHAIR™
— A PARC SOLUTIONS PLATFORM —
ISSUE
4
PARC
— SOLUTIONS —
THE INFRASTRUCTURE RACE
Energy, Supply Chain, and the Cost of AI Dominance
AI
ENERGY SECURITY
RESILIENT SUPPLY CHAINS
GEOPOLITICAL LEVERAGE
TECHNOLOGY SELF-RELIANCE
NATIONAL SECURITY AT STAKE
AI leadership will be won not just in labs,
but in power plants, ports, factories, and fiber.

THE INFRASTRUCTURE RACE

Energy, Supply Chain, and the Cost of AI Dominance

Pawneet Abramowski
CEO & Founder, PARC Solutions · Independent Board Director, Titan Acquisition Corp. (NASDAQ: TACHU) · Former FBI Intelligence Analyst
Published July 2026 · The Risk Chair™ · parcsolutionsllc.com

About This Paper

This is the fourth publication from The Risk Chair™. The Decoupling Doctrine (Issue 01) mapped the strategic reorientation of American economic and national security posture. What the Boardroom Does Not Yet Know About Artificial Intelligence (Issue 02) addressed the governance gap inside institutions. The Governance Deficit (Issue 03) connected Basel III and AI governance as a single compliance obligation. This paper goes to the physical layer beneath all three: the energy, supply chain, and geopolitical infrastructure that determines who can actually sustain AI dominance at scale. It is written for board directors, chief risk officers, chief compliance officers, general counsel, and the senior executives responsible for institutions with credit, counterparty, underwriting, or operational exposure to the AI infrastructure buildout. The risk exposures described in this paper are not confined to financial institutions alone. Legal firms advising on data center transactions, logistics and infrastructure companies moving minerals and equipment through contested corridors, insurers underwriting power purchase agreements and marine risk, technology vendors supplying AI hardware through regulated supply chains, and real estate developers building in water-stressed markets all carry embedded Power, Processing, and Passage risk. The governance argument applies wherever the infrastructure race touches an organization's operations, contracts, or counterparty relationships.

FOREWORD

When a Theoretical Risk Becomes an Operational Reality

On February 28, 2026, Israel and the United States launched airstrikes against Iran. Within hours, the Islamic Revolutionary Guard Corps closed the Strait of Hormuz. Tanker traffic through the world's most consequential maritime chokepoint dropped to nearly nothing. Brent crude reached $126 per barrel at peak. The largest disruption to global energy supply since the 1970s oil crisis unfolded in 48 hours.

I have been in this space long enough to recognize the moment when a theoretical risk becomes an operational reality. The Strait of Hormuz situation was not a surprise to anyone tracking the strategic architecture this series has been mapping since Issue 01. The Decoupling Doctrine laid out the logic: supply chain sovereignty is not a geopolitical abstraction. It is a physical condition, measured in chokepoints, in mineral deposits, in fabrication capacity, and in the energy required to run it all. The February 28 strikes did not create a new vulnerability. They revealed one that was already there.

What I found striking, watching the situation develop over the months that followed, was not the disruption itself. It was the pattern. The MOU signed on June 17 brought temporary relief. The market priced in normalization. Then on July 7, Iran attacked a Qatari LNG tanker, struck two other vessels, and the US bombed more than 80 Iranian targets overnight. Brent surged again. Trump declared the ceasefire over. The MOU is now on life support, kept functioning by periodic diplomatic injections that hold for days or weeks before the next violation and the next escalation. This pattern, ceasefire, violation, strikes, renewed pause, repeat, is not a bug in the negotiating process. It is the operating environment that financial institutions with exposure to this corridor have to plan for through the remainder of 2026 and beyond.

The AI race made this particular chokepoint more consequential than it would otherwise have been. The LNG that powers the data centers being built to run AI moves through that strait. The petrochemical feedstocks that go into semiconductor fabrication move through it. The physical infrastructure layer of the AI race runs through geography that is actively contested. That is the argument this paper makes, and the Hormuz situation is its live proof point.

The framework this paper uses has three physical chokepoints and a governance translation layer. Power, Processing, and Passage are the three conditions that determine whether AI infrastructure gets built and sustained. The Governance Layer is where those conditions translate into what a board, a risk committee, or a senior executive actually does about them. That structure emerged from the same analytical lens I brought to the Decoupling Doctrine: the operational-transformation perspective that asks not just what is happening at the policy level, but what it means for the institutions that have to navigate it in real time. The boards and executives reading this paper are not in the business of geopolitical prediction. They are in the business of managing the exposure their institutions carry, much of which is now tied, directly or indirectly, to who wins the infrastructure race.

Pawneet Abramowski | CEO & Founder, PARC Solutions
Publisher, The Risk Chair™ · July 2026

EXECUTIVE SUMMARY

The Infrastructure Race: Three Arguments

The AI race will not be won in a model architecture. It will be won in the ground beneath it.

The first three issues of this series addressed the strategic decoupling underway at the geopolitical level, the governance gap inside institutions, and the compliance deficit between Basel III and AI governance. This paper goes to the physical layer beneath all three: the energy, supply chain, and geopolitical infrastructure that determines who can actually sustain AI dominance at scale. The implications for credit risk, counterparty exposure, underwriting, and board-level oversight are direct and underappreciated.

The AI ecosystem is splitting. BCG's 2026 analysis documents a clean bifurcation between the US and Chinese stacks. Stargate, the $500 billion infrastructure joint venture announced at the White House in January 2025, is the defining US response: winning through physical scale. China is building a parallel stack, deliberately and at pace. For institutions financing or insuring infrastructure in either ecosystem, the choice of stack now determines legal operating jurisdiction and geopolitical exposure. Sitting between the two is no longer an option for most corporations or nations.

Three arguments anchor the paper.

> *First: The compute buildout has become an energy procurement problem before it is anything else. America's investor-owned utilities have announced $1.4 trillion in capital spending through 2030. Lead times for large power transformers have stretched to five years, leaving 40% of planned US data center capacity facing delays. The hyperscalers have responded by building nuclear capacity directly: 13 announced projects committing more than 9.8 gigawatts. Meanwhile, domestic community opposition stalled over 75 data center projects worth $130 billion in Q1 2026 alone and is driving capital to Nordic sites in Europe. The energy layer is not a supporting detail. It is the binding constraint.*

> *Second: The processing layer, chip fabrication and critical mineral inputs, remains a strategic vulnerability despite two years of reshoring effort. China controls 40 to 90 percent of global processing capacity for the minerals in every advanced chip. US enforcement has moved from broad export controls to surgical secondary sanctions: a Bangkok firm was flagged in May 2026 for diverting Nvidia servers to Chinese entities including Alibaba, establishing that the US views the physical location of AI hardware as a national security boundary. The hyperscalers are spending $600 billion on infrastructure this year, with 75 percent targeting AI, consuming up to 70 percent of globally produced memory chips and forcing reallocation away from automotive and industrial sectors.*

Third: The Strait of Hormuz is not recovering. It is on life support. The pattern since February 28 has been consistent: ceasefire, violation, US strikes, renewed pause, repeat. On July 7, Iran attacked a Qatari LNG tanker and two other vessels. The US bombed 80 Iranian targets overnight. Trump declared the ceasefire over. The MOU signed June 17 is effectively voided. Brent crude surged above $78 at peak. The oil market's tendency to price in normalization before normalization is secured has created systematic mis-pricing of this risk, and it has now repeated twice in six weeks. One-fifth of global oil and LNG, including the energy that powers the data centers being built to run AI, transits a waterway that could remain below 50 percent of pre-war traffic levels for many months.

Power. Processing. Passage.

Three physical chokepoints. One governance obligation. The board that understands these as a unified risk category, not three separate departmental concerns, is the board that will govern this transition rather than be governed by it.

1

POWER

Energy generation, grid capacity, and storage. The compute buildout is an energy procurement problem before it is anything else. 40% of planned US data center capacity faces delays due to transformer shortages alone.

2

PROCESSING

Chip fabrication and critical mineral inputs. China controls 40-90% of processing capacity for the minerals in every advanced chip. The reshoring effort is real. The gap it is closing is larger than the timeline for closing it.

3

PASSAGE

The geopolitical chokepoints that move physical inputs and finished capacity. The Strait of Hormuz has been on life support since February 28. The pattern repeats: ceasefire, violation, strikes, renewed pause.

PART I

Power

The Grid Cannot Keep Up

The energy problem predates the AI buildout. AI has made it structural. America's investor-owned utilities have collectively announced $1.4 trillion in capital spending through 2030, up 27 percent from the prior year and effectively double what was invested in the previous decade. Duke Energy: $102 billion. Southern Company: $81 billion. The driver is not electrification alone. It is data center demand, which has rewritten utility capital plans faster than regulators or ratepayers have absorbed.

Global data center electricity consumption stood at 415 terawatt-hours in 2024. The IEA projects 945 by 2030, more than double in six years. US data center demand alone is forecast to increase 160 percent. Goldman Sachs has identified energy availability as the binding AI infrastructure constraint, displacing chip supply. Nvidia has confirmed it slowed cluster expansion not because of GPU shortages but because the power was not there.

The physics are simple and the math is relentless. A single H100 GPU draws 700 watts at full load. An AI rack draws 80 to 140 kilowatts. A 10,000-GPU cluster demands continuous power the grid in most cities cannot deliver. Sam Altman has estimated 0.34 watt-hours per ChatGPT query. At 2.5 billion daily requests, that is a full-time power plant's output. Every day.

The $653 Billion Midstream Blind Spot

The constraint nobody is talking about loudly enough is not generation. It is the midstream equipment that moves power from the grid to the facility. US imports of data center equipment, transformers, switchgear, and midstream battery processing hardware surged past $653 billion, more than doubling since 2020. Lead times for large power transformers now run five years. Roughly 40 percent of planned US data center capacity is facing delays not from GPU shortages or permitting problems but because the physical equipment required to power the facilities cannot be procured on the timeline the buildout demands. This is an industrial policy gap, and it sits inside the credit profile of every institution financing a data center.

US dominance in frontier models and chip design does not solve this problem. True infrastructure sovereignty requires domestic production of the grid components that feed the data centers, not just the chips inside them. No amount of AI investment accelerates a transformer procurement queue.

AEP Ohio has paused all new data center interconnections due to insufficient power infrastructure. That is a live data point, not a scenario. It is what grid constraint looks like when it materializes in a specific market, and it belongs in every project finance stress test for data center lending.

The Nuclear Turn

Every major hyperscaler has reached the same conclusion independently: renewables cannot deliver the reliability AI requires. Solar and wind are intermittent. Data centers are not. Storage infrastructure at the required density does not exist. The answer, for all of them, is nuclear.

Thirteen announced projects now commit more than 9.8 gigawatts of nuclear capacity to AI infrastructure. Microsoft secured a 20-year PPA for the restart of Three-Mile Island Unit 1. Google contracted Kairos Power for 500 megawatts of small modular reactor capacity. Amazon invested $700 million in X-energy. Meta leads the field with up to 6.6 gigawatts across TerraPower, Oklo, Vistra, and Constellation. Eighteen months ago, nuclear was politically complicated. Today it is strategically necessary.

Stargate, the $500 billion joint venture between OpenAI, SoftBank, Oracle, and MGX, is the US infrastructure bet made tangible. Seven sites under active construction across Texas, New Mexico, Ohio, and Wisconsin, targeting more than 9 gigawatts of planned capacity. That is comparable to the peak power demand of New York City. The Abilene flagship is partially operational. The Michigan campus is projected at $16 billion. JPMorgan Chase provided a $2.3 billion project finance loan for the Abilene facilities. The governance implications for institutions financing similar buildouts are addressed in Part IV.

The $130 Billion Domestic Governance Vulnerability

The federal strategy requires rapid infrastructure scaling. Local governments are not cooperating. In Q1 2026, municipal blocks and zoning delays stalled over 75 data center projects worth $130 billion in the US. Google walked away from a $1 billion Indiana site. Amazon was voted down in Arizona over power and water constraints. Communities are pushing back on water consumption, power strain, and the gap between what developers promise and what neighbors actually receive.

Capital is responding by leaving. Bitzero and others have shifted billions to pre-permitted, gigawatt-scale Nordic sites in Norway and Finland, where power is abundant, cooling is natural, and nobody votes down a data center. The federal strategy and the local operating environment are pulling in opposite directions. That misalignment is a credit risk variable in every US data center project finance deal being written today.

Most project finance credit models have not priced community opposition and zoning risk for data center construction. They should. An institution financing a data center in a market where water rights are contested, utility capacity is constrained, or local approval has been challenged is financing a project whose completion depends on variables no technology investment thesis can control. The contrast between US community opposition and Nordic acceptance is not simply a story about environmental attitudes or geography. It is a governance story, and it has direct implications for the risk profile of AI infrastructure investments across jurisdictions. US communities opposing data centers are not primarily opposed to the technology. They are responding to three compounding failures of the development process. The first is health and environmental opacity: most communities receive corporate disclosure reports rather than

independent environmental impact assessments, and the data center industry's standard response to water consumption questions has not met the credibility threshold communities now require. The second is the broken promise pattern: the economic benefit commitments — jobs, tax revenue, local investment — have frequently underdelivered relative to the infrastructure stress that materialized, and communities have learned this through experience. Google's Indiana withdrawal after public commitment and Amazon's Arizona rejection both happened after the development process was already underway. The third is the equity dimension: the communities most affected by data center energy and water demand are often not the communities that benefit from the digital services the facilities provide. This is an organized, legally sophisticated opposition movement that understands zoning law and permitting processes well enough to stop $130 billion in projects in a single quarter. Nordic communities have opposition too. A Norwegian Parliament member has called data center electricity demands "hugely controversial." The Narvik facility has reignited national debate over electricity consumption and national control. Finland is moving data center electricity from the lower industrial tax category to the standard tax bracket on July 1, 2026, and proposed legislation would require large facilities to obtain matching new power generation capacity before connecting to the national grid. What is different in the Nordics is the social contract governments have built around the industry. Norway's National Data Centre Strategy, launched in January 2026, requires facilities to be energy-efficient, run on renewable energy, and capitalize on energy recovery and reuse of excess heat. Microsoft's facility in Espoo, Finland will pipe waste heat to up to 100,000 homes, eliminating the need for fossil-fueled heating systems and helping the city meet its carbon neutrality target. This is the biggest heat recovery technology project anywhere in the world. Communities in Finland and Norway are accepting data centers they might not actively welcome because the regulatory compact gives them a tangible benefit, cheaper, cleaner home heating in direct exchange for absorbing the infrastructure. The governance lesson for financial institutions is specific. Regulatory certainty is a credit variable. Projects in jurisdictions with defined social contract frameworks — mandatory waste heat recovery, renewable energy requirements, community benefit assessments, and national strategies that establish what data centers owe communities in exchange for access to power and water carry structurally lower community opposition and litigation risk than projects in the US regulatory patchwork. An institution financing AI infrastructure in the US without assessing community relations history, water rights certainty, and permitting environment is underwriting a variable that its Nordic counterparts have addressed at the regulatory level. That is not an endorsement of one model over another. It is a risk differentiation that belongs in the credit analysis.

Water deserves far more than a separate mention. It is the fourth infrastructure chokepoint the AI race has produced, and the least visible in current risk frameworks. Modern hyperscale data centers consume between one and five million gallons of water daily for cooling. A typical facility uses the daily water equivalent of 1,000 households. Projections show water used for AI cooling may increase by 870 percent as more facilities come online. In March 2026, Ecolab paid 29 times forward EBITDA to acquire CoolIT Systems, a liquid cooling company owned by KKR. That multiple is not paid for a sleepy industrial business. It is paid when the smartest money in

industrial M&A believes a category is on the cusp of structural transformation. Every hyperscale data center being designed for 2026 to 2028 is liquid-cooled by default. Water is now an infrastructure procurement problem before it is a sustainability footnote. In Arizona, Nevada, and the Colorado River basin, data centers compete with farmers and residents for a resource that is already overallocated. Virginia, which has the world's highest concentration of data centers, is moving toward restricting the most water-intensive cooling methods. Google's expansion in The Dalles, Oregon has faced years of community resistance. Water rights and access have evolved from due diligence checkboxes to deal-breaking issues requiring the same sophisticated analysis traditionally reserved for power procurement. An often-overlooked dimension of President Trump's interest in acquiring Greenland lies here: the island is the world's largest inhabited natural cooling reservoir and freshwater source, making it potentially transformative for hyperscale AI data centers where access to low ambient temperatures and abundant water could dramatically reduce cooling costs. Bill Gates, Sam Altman, and Jeff Bezos have reportedly explored Greenland for exactly this reason. The Nordic countries have built their competitive advantage around this insight. Microsoft has committed $6.2 billion in Norway through Nscale, $3.2 billion in Sweden, and $3 billion in Denmark. Oslo, Stockholm, and Helsinki have emerged as the three primary development corridors, with Sweden accounting for 44 to 45 percent of Nordic investment share and power capacity. These facilities use naturally cold fjord water for passive cooling without drawing on scarce freshwater supplies. China has solved the problem differently: its undersea data center near Shanghai, now fully operational with $226 million invested, uses surrounding ocean water as a passive heat sink and achieves a Power Usage Effectiveness below 1.15. No freshwater consumed. This is the same strategic logic China applied in the DRC, finding a chokepoint others depend on and building around it before the dependency becomes visible to competitors. For financial institutions, water access has joined grid connection and mineral supply chains as a primary credit variable in AI infrastructure financing. Projects in water-stressed markets without secured water rights, community consent, and permitting certainty carry a risk profile that most credit models have not yet priced.

PART II

Processing

The Chip Layer and the Reshoring Gap

The CHIPS Act has produced real results. TSMC's Arizona plant is in production. Intel, Samsung, and Micron have announced or begun domestic facilities. The January 2026 Section 232 executive order on processed minerals, the February Critical Minerals Ministerial with eleven new bilateral agreements, and the CFIUS and COINS Act perimeter represent a coherent strategy. But strategy and supply chain reality are on different timelines.

EUV lithography tools come only from ASML, manufactured in the Netherlands with components sourced across three continents. Lead times for new tools run two years or longer. Even with fabrication investment fully funded, the ramp to meaningful domestic production of the most advanced chips takes time the AI deployment velocity does not allow. TSMC's Arizona plant is producing. It is producing at a fraction of the scale of the Taiwan facilities. The near-term dependency has not changed.

The Bangkok Precedent: Enforcement as the New Export Control

In May 2026, US authorities flagged a Bangkok firm tied to Thailand's national AI initiative for allegedly diverting billions of dollars' worth of Nvidia-powered servers to Chinese entities including Alibaba. This is the shift that matters: Washington has moved from broad export controls to surgical enforcement of diversion networks. Where AI hardware physically sits is now a national security question, and getting the answer wrong can trigger secondary sanctions.

Southeast Asian nations are now navigating a compliance dilemma they did not ask for: hosting the wrong hardware could draw US enforcement attention. Institutions with trade finance or counterparty exposure in Thailand, Malaysia, Vietnam, or the UAE need to assess whether their clients are exposed to diversion risk under the enforcement framework the Bangkok case established. The risk is not theoretical. It is being prosecuted.

The AI systems financial institutions are deploying run on GPU clusters that source chips through a fabrication ecosystem subject to US export controls. Most third-party risk assessments do not ask where a vendor's hardware was procured or what supply chain it traversed to get there. They need to.

The Mineral Layer: DRC and the Refining Chokepoint

The DRC produces approximately 70 to 75 percent of global cobalt. China controls the refining stage between Congolese mines and Western manufacturers. When this administration took office, Chinese firms controlled roughly 80 percent of Congo's mining output. CMOC's Tenke

Fungurume mine, once US-owned, is now the world's second-largest cobalt source. China paid $2.65 billion for it while Western investors were citing conflict risk as reasons to stay away.

The DRC implemented cobalt export quotas in late 2025 capping hydroxide exports at 96,600 tonnes annually for 2026 and 2027, less than half of 2024 production of approximately 204,000 tonnes. Chinese-affiliated operators carry a structural advantage in this quota environment because their integrated supply chains allow them to move material through processing stages before export, reducing the customs complexity that Western miners, who predominantly export raw or intermediate material, must navigate.

The US response has been substantial. A December 2025 strategic minerals partnership with the DRC offered security support in exchange for preferential mineral access. Virtus Minerals finalized its acquisition of Chemaf SA in April 2026, describing itself as the first US-owned operator to return to the DRC in over ten years, with combined operations targeting 75,000 tonnes of copper cathodes and 25,000 tonnes of cobalt hydroxide annually. The Department of Defense's Office of Strategic Capital signed a conditional $725 million loan commitment with Energy Fuels to expand domestic rare earth processing. China responded by signing a new mining cooperation agreement with the DRC in March 2026 and extending duty-free access to Congolese mineral exports starting May 2026.

Mining is one thing. Refining is another. China's dominance is at the refining stage, and building alternative capacity requires capital, technology, regulatory approval, and years of operational development that no policy commitment can compress. The institutions financing AI infrastructure today are financing a buildout whose mineral supply chain still runs, in the near term, through the ecosystem US policy is working to displace.

Supply Chain Cannibalization: When AI Crowds Out Everything Else

The top five hyperscalers are projected to spend over $600 billion on infrastructure in 2026, with 75 percent targeting AI. Data centers are consuming up to 70 percent of all globally produced memory chips. Samsung, SK Hynix, and Micron are reallocating fabrication capacity away from automotive and industrial customers to meet AI demand. A car manufacturer that cannot source chips is experiencing the downstream consequence of capital allocation decisions made in Cupertino.

Controlling AI infrastructure is not just a national security capability. It is leverage over the global manufacturing supply chain for every industry that depends on the same chips AI is consuming. Financial institutions with lending exposure to automotive, industrial, or consumer electronics companies need to understand that their counterparties' supply chains are being restructured by decisions made in Cupertino and Mountain View, not by their borrowers.

PART III

Passage

The Strait on Life Support

The Strait of Hormuz is 34 kilometers wide at its narrowest point. Before February 28, roughly 20 percent of global seaborne oil trade, 20 percent of global LNG, and a meaningful share of the LNG powering the data centers being built to run AI transited it daily. For five months, almost none of it has moved at pre-war levels.

When Israel and the United States launched airstrikes on February 28, killing Supreme Leader Khamenei and targeting nuclear and missile infrastructure, Iran went straight for the chokepoint. The IRGC closed the strait, attacked merchant shipping, and laid mines. Tanker traffic dropped to nearly nothing. Brent hit $126 at peak. The IEA activated emergency measures. The Philippines began a four-day work week to reduce fuel consumption. Analysts started running scenarios with oil at $200.

What followed was not recovery. The June 17 Islamabad MOU brought some traffic back to the strait and markets priced in normalization. Then on July 7, Iran attacked a Qatari LNG tanker and two other vessels with missiles and drones. The US bombed 80 Iranian targets overnight. Trump declared the ceasefire over at the NATO summit in Turkey. Treasury revoked Iran's oil sale license effective July 17. Brent surged above $78. The MOU is not dead. It is on life support, kept functioning by periodic interventions that hold for days or weeks before the next violation and the next escalation.

> *The market's tendency to price in normalization before normalization is secured has now produced systematic mis-pricing of this risk twice in six weeks. An analyst at MST Financial described it directly: the market 'is treating this temporary ceasefire as a permanent deal' and has 'overshot to the downside.' Every institution that marked its Hormuz-exposed positions to market-implied normalization in late June and then repriced them again after the July 7 attacks has now experienced this dynamic twice. The governance question is whether those institutions have updated their scenario frameworks to treat the life-support pattern as the base case, not as an exception.*

What the Hormuz Situation Proves

The Hormuz crisis is not a geopolitical aside to the infrastructure race thesis. It is the proof of it. The three physical chokepoints, Power, Processing, and Passage, are not independent risk categories. They are connected through the same underlying dynamic: the AI buildout has made the physical infrastructure layer of technology competition a national security priority, which means it has made every point of physical vulnerability in that layer a potential target for adversarial leverage.

Iran's response was not random. It went straight for the chokepoint that mattered most to the countries whose support the US needed. That is the lesson. Geographic control over a physical passage is leverage in a world where AI infrastructure, energy supply, and semiconductor supply chains all run through the same narrow waterways. The Middle East is the example. The category of risk is universal.

There is a dimension of this risk that barely registers in most portfolio assessments: fertilizer. The Persian Gulf accounts for 30 to 35 percent of global urea exports and 20 to 30 percent of ammonia. Up to 30 percent of internationally traded fertilizers transit the strait. A prolonged closure does not just affect energy prices. It affects food production, which ripples through agricultural lending, trade finance, and sovereign credit for food-importing nations. Most of those exposures are not classified as Hormuz risk. They should be.

China's Strategic Response: Converting Disruption into Positioning

China's behavior during the crisis deserves attention. While the US has been operationally entangled in the Middle East, China has used the disruption to extend energy relationships with Global South suppliers, position itself as the reliable alternative LNG source, and frame itself as the stable partner while the Western-led order proved fragile. The countries that found Chinese suppliers during the closure will remember who delivered. That positioning advantage compounds, and it is one the decoupling thesis this series has been mapping since Issue 01 identified as the long game.

China receives approximately one-third of its oil through the Strait of Hormuz, which means the disruption-imposed costs on China as well. Beijing's response was to accelerate bilateral energy agreements with alternative Gulf suppliers, reduce its own strait dependency, and use the diplomatic space created by the US-Iran conflict to strengthen its role as the indispensable infrastructure partner for the Global South. The BCG bifurcation analysis finds that third-party nations and corporations are increasingly unable to maintain neutrality between the US and Chinese AI stacks. The Hormuz crisis has demonstrated the same dynamic in energy markets: when the Western-led system fails to deliver energy security, the Chinese alternative gains credibility regardless of its own limitations.

Pax Silica as the Institutional Hedge

Pax Silica, the four-pillar multilateral framework from the December summit now including India, is the diplomatic hedge against chokepoint exposure. The same way the domestic buildout hedges Processing risk and the utility spending program hedges Power, Pax Silica is the Passage hedge. The framework assumes that supply chain diversification and allied coordination can reduce the leverage a single adversary can exercise over a critical passage.

The Hormuz crisis is Pax Silica's first real test. Framework commitments do not immediately become operational resilience. Routing around a disrupted chokepoint requires logistics infrastructure, pricing mechanisms, and alternative supplier relationships that take years to build. The lesson is simple: build them before you need them, not after. Every financial

institution that had counterparty exposure through the Persian Gulf on February 27 is learning that lesson the expensive way.

PART IV

The Governance Layer

The Integration Problem

Power, Processing, and Passage are not geopolitical abstractions. They are risk exposures sitting in loan books, underwriting files, vendor contracts, and counterparty portfolios right now. The question is whether the board knows they are there, and whether anyone inside the institution is responsible for making sure they surface before they become loss events.

Energy risk, supply chain risk, and geopolitical risk are established categories in every major institution's enterprise risk framework. The problem is convergence. The AI infrastructure buildout has made these three categories interdependent in ways existing frameworks were not designed to see. A data center project finance deal, a semiconductor manufacturing loan, and a trade finance facility for a petrochemical company can each be performing independently while all three carry concentrated exposure to the same underlying chokepoint. The concentration is invisible in any single risk category view.

This is the same dynamic Issue 03 identified for Basel III and AI governance: interconnected risk categories managed in parallel frameworks, invisible as a whole at the board level.

Who Owns It

Who owns this? That is the question. Power, Processing, and Passage exposure does not sit cleanly with the CRO, the CGO construct introduced in Issue 02, or a traditional geopolitical risk function. It spans all three, and most institutions have not named an owner for it. The result is a risk category visible in pieces across multiple functions and invisible as a whole at the board level, which is precisely where it needs to be visible.

The fix is the same one Issue 03 described for AI and Basel III: a single accountable owner with cross-functional visibility, presenting a unified picture to the board. The title matters less than the mandate. What matters is that someone is responsible for ensuring the board sees Power, Processing, and Passage as one risk category, not three separate departmental concerns that happen to share an underlying driver.

Four KRIs for Board-Level Tracking

The first KRI is AI infrastructure exposure concentration by chokepoint corridor. Map the underlying logistics and energy dependencies of any data center, semiconductor, mineral processing, or energy infrastructure exposure against the known chokepoint corridors: Hormuz, the Taiwan Strait, the South China Sea, and the domestic grid interconnection bottlenecks that are the US-side equivalent. The Hormuz crisis has produced five months of empirical data on

what disruption looks like and how long it lasts. Scenario models built on pre-war assumptions need to be rebuilt against that data.

The second is the percentage of AI-infrastructure-linked lending with unpriced Power risk. The data center financing boom has moved faster than the underwriting frameworks for it. Loan books with material exposure to data center construction should be reviewed specifically for whether grid interconnection risk, power purchase agreement counterparty quality, transformer procurement timelines, permitting assumptions, and water resource availability have been stress-tested against the delays and constraints now materializing in actual markets. AEP Ohio's interconnection pause and Google's Indiana withdrawal are not anomalies. They are the leading edge of a constraint that will become visible across multiple markets over the next 24 months.

The third KRI is vendor and counterparty dependency mapping for diversion risk. After Bangkok, any institution whose vendors or borrowers are sourcing AI hardware through jurisdictions where US diversion enforcement is active needs that dependency identified. The standard third-party risk assessment does not ask whether a vendor's GPU cluster was procured through a supply chain that could trigger secondary sanctions. It should.

Water Access and Cooling Infrastructure Risk

The fourth KRI is water access and cooling infrastructure risk for data center construction. For any institution with project finance or lending exposure to data center development, the water rights certainty, cooling technology specification, permitting status relative to local water authority requirements, and community relations posture of each project should be assessed as a primary credit variable alongside grid interconnection and power procurement. The Ecolab-CoolIT acquisition at 29 times forward EBITDA signals that the market has already repriced water infrastructure as a strategic asset. Credit models for data center lending that do not reflect that repricing are working from outdated assumptions. Projects in Nordic jurisdictions with regulated social contract frameworks carry structurally lower water and community risk than comparable projects in the US regulatory patchwork. That risk differential belongs in the underwriting.

The Insurance and Underwriting Gap

Insurance has not kept pace. Marine war risk premiums for Hormuz corridor shipping repriced dramatically in February, with some carriers suspending coverage at peak disruption. Property and casualty underwriters are reassessing how to price data center facilities that depend on grid connections with five-year interconnection queues. Political risk and trade credit underwriters are working through what a five-month strait closure means for their extreme-but-recoverable event models.

Stargate is itself a new category of underwriting challenge. JPMorgan Chase provided a $2.3 billion project finance loan for the Abilene facilities. The broader $500 billion commitment depends heavily on project debt secured against signed compute contracts, not committed equity. Most of the remaining capital depends on future financing rounds. A multi-partner, multi-site, power-dependent AI infrastructure project with community opposition risk, grid

constraint risk, and geopolitical supply chain exposure across all three chokepoints described in this paper is a credit risk profile that has not been underwritten at this scale before.

PARC SOLUTIONS PRACTICE IMPLICATIONS

What Institutions Need Now

The exposures described in this paper translate into three specific advisory needs. None of them is speculative. The Hormuz pattern, the transformer queue, the DRC quota system, the Bangkok enforcement action, and the water infrastructure constraint have all made them immediate. And while financial institutions are the primary audience for this series, the governance needs described here extend through the entire AI infrastructure ecosystem: legal firms structuring data center transactions, logistics and insurance providers operating in contested corridors, technology vendors managing hardware procurement through regulated supply chains, and real estate developers navigating the US regulatory patchwork. Power, Processing, and Passage risk does not stop at the bank's balance sheet. It runs through every organization whose operations intersect with the infrastructure race.

The first is AI infrastructure exposure mapping and credit risk advisory. Financial institutions with project finance, lending, or counterparty portfolios touching data center development, semiconductor manufacturing, critical mineral processing, or energy infrastructure need a structured assessment of where those exposures carry embedded Power, Processing, and Passage risk that has not been priced. This is cross-functional analysis that connects sector-based credit categories to the infrastructure dependencies that make them convergent. PARC Solutions brings direct experience inside major financial institutions managing these risk categories to engagements that help clients see the whole picture rather than its parts.

The second is governance framework development for boards and risk committees. The integrated architecture described in Part IV, a single accountable owner, board-level KRI reporting, and stress tests designed for chokepoint scenarios, does not exist as a standard offering in most enterprise risk programs. PARC Solutions builds these frameworks directly with boards and senior risk leadership, drawing on the governance constructs developed across this series to build architecture that reflects how these risks actually interact, not how organizational charts have historically divided them.

The third is third-party risk program enhancement for AI supply chain and diversion exposure. Vendor due diligence frameworks were designed for software vendors, not for AI hardware vendors whose procurement chains run through jurisdictions where US diversion enforcement is active. PARC Solutions updates vendor assessment frameworks to surface these dependencies, building due diligence questions, contractual protections, and monitoring protocols that reflect the supply chain architecture of enterprise AI and the enforcement environment now governing it.

CLOSING

The Race Is Physical

The AI race is a physical race. It requires power, chips, and passage. In that order. At a scale that has no precedent. The institutions financing it are taking positions on all three whether or not they have named those positions in a credit memo.

Four issues. Four altitudes of the same challenge. Decoupling at the geopolitical level. Governance gaps inside institutions. The compliance deficit between deployment velocity and oversight capacity. And now the physical infrastructure layer that determines whether any of it holds. The Governance Layer is where they converge. The board that sees Power, Processing, and Passage as one risk category rather than three departmental concerns is the board that governs this transition rather than reacts to it.

The Hormuz MOU is on life support. Transformer procurement queues stretch five years. The DRC cobalt quota system runs through 2027. The nuclear capacity the hyperscalers have committed to comes online between 2027 and 2030. The largest private infrastructure commitment in history depends on power that does not yet exist at the required scale. Build the governance frameworks now. The gap between the exposure and the loss event is narrowing.

Issue 05 takes this from a different vantage point: family office allocations into AI and supply chain infrastructure, and the diligence gap between aggressive allocation and structured risk awareness. The thread continues.

Pawneet Abramowski | CEO & Founder, PARC Solutions
Publisher, The Risk Chair™ · A PARC Solutions Platform · July 2026

APPENDIX A

Acronym Glossary

Acronym	Definition
AI	Artificial Intelligence — Advanced computational systems enabling machine learning, automation, and decision support.
AML/CFT	Anti-Money Laundering / Countering the Financing of Terrorism — Regulatory framework requiring detection and reporting of illicit financial activity.
ASML	Advanced Semiconductor Materials Lithography — Dutch company and sole global manufacturer of EUV lithography systems critical for advanced chip production.
BCG	Boston Consulting Group — Global management consultancy whose 2026 AI bifurcation analysis informs Part I of this paper.
CFIUS	Committee on Foreign Investment in the United States — Interagency body reviewing foreign investments for national security risks.
CHIPS Act	Creating Helpful Incentives to Produce Semiconductors Act — US legislation funding domestic semiconductor manufacturing.
CIPS	Cross-Border Interbank Payment System — China's alternative to SWIFT for international payments, expanding global reach.
CMOC	China Molybdenum Co. Ltd — Chinese mining company operating the Tenke Fungurume cobalt mine in the DRC.
COINS Act	Outbound Investment National Security Act — Legislation governing US outbound investments in sensitive sectors including semiconductors and AI.
DFC	US International Development Finance Corporation — US government agency providing financing for development projects in lower- and middle-income countries.
DoD	Department of Defense — US department responsible for military operations and national defense.
DRC	Democratic Republic of the Congo — Source of approximately 70-75% of global cobalt production.
ERCOT	Electric Reliability Council of Texas — Independent grid operator managing Texas's power system.
ESG	Environmental, Social, and Governance — Framework for assessing non-financial institutional risk and impact.
EUV	Extreme Ultraviolet Lithography — Advanced semiconductor manufacturing process used in the most advanced chip fabrication, monopolized by ASML.
GENIUS Act	Guiding and Establishing National Innovation for US Stablecoins Act — Legislation creating a regulatory framework for dollar-backed stablecoins.
GPU	Graphics Processing Unit — Specialized processor now central to AI model training and inference workloads.

IEA	International Energy Agency — Intergovernmental organization tracking global energy supply and demand.
IRGC	Islamic Revolutionary Guard Corps — Iran's elite military force responsible for closing and enforcing the Strait of Hormuz.
KRI	Key Risk Indicator — Metric used to signal emerging risk conditions for board-level monitoring.
LNG	Liquefied Natural Gas — Cryogenically cooled natural gas transported by tanker; approximately 20% of global supply transits the Strait of Hormuz.
MGX	Abu Dhabi AI and advanced technology investment firm — Fourth equity backer in the Stargate joint venture.
MOU	Memorandum of Understanding — Non-binding agreement between parties; the June 17 Islamabad MOU between the US and Iran established a 60-day negotiating window.
NDAA	National Defense Authorization Act — Annual US legislation governing defense spending and national security policy.
OCC	Office of the Comptroller of the Currency — US regulator supervising national banks.
PJM	PJM Interconnection — Regional transmission organization managing the electric grid across 13 US states.
SMR	Small Modular Reactor — Advanced nuclear reactor design offering flexible capacity for data center power supply.
SPR	Strategic Petroleum Reserve — US emergency stockpile of crude oil.
TSMC	Taiwan Semiconductor Manufacturing Company — World's largest semiconductor foundry.
WTI	West Texas Intermediate — US benchmark crude oil price.

APPENDIX B

Selected Bibliography

Primary sources consulted in the preparation of this publication. All sources reflect publicly available materials as of July 2026.

BCG. 'The Great AI Bifurcation: US and Chinese Ecosystems Diverge.' 2026 Analysis.

Berenberg. Schmieding, Holger. Research Note on Hormuz and Oil Markets. July 2026.

Bloomberg. 'How High Could Oil Prices Get with Strait of Hormuz Closure?' March 2026.

Built In. 'What Is the Stargate Project?' Updated July 2026.

CNBC. 'Oil Prices Rise After Attacks on Tankers in Strait of Hormuz.' July 7, 2026.

CNBC. 'Oil Prices Jump After Trump Threatens to Bomb Iran and Reimpose Naval Blockade.' July 8, 2026.

CNBC. 'OpenAI's First Data Center in $500 Billion Stargate Project is Open in Texas.' September 24, 2025.

Data Center Knowledge. 'Stargate Update: AI's Biggest Data Center Buildout Meets Reality.' July 2026.

Epoch AI. 'OpenAI Stargate: Where the US Sites Stand.' April 17, 2026.

Alliance for Prosperous America. Data Center Equipment Import Surge Analysis. 2026.

Al Jazeera. 'Oil Surges as US Strikes Iran, Reversing Return to Pre-War Prices.' July 8, 2026.

IEA. World Energy Outlook 2025. Data center demand projections.

ING. Commodities Strategy Note on Strait of Hormuz Oil Market Pricing. July 2026.

IntuitionLabs. 'OpenAI's Stargate Project: A Guide to the AI Infrastructure.' February 2026.

MST Financial. Kavonic, Saul. Hormuz Risk Assessment Note. July 2026.

NBC News. 'Oil Rises as Attacks Renew Safety Fears in Strait of Hormuz.' July 7, 2026.

OpenAI. 'OpenAI, Oracle, and SoftBank Expand Stargate with Five New AI Data Center Sites.' September 22, 2025.

PowerLines / Tech-Insider. 'US Utility $1.4 Trillion Capital Spending Plan.' April 14, 2026.

US Department of Defense, Office of Strategic Capital. Energy Fuels $725M Loan Commitment. June 2026.

US Department of State. Critical Minerals Ministerial Outcomes. February 2026.

US Department of State. DRC Strategic Minerals Partnership. December 2025.

US Treasury / Commerce. Bangkok Transshipment Enforcement Action. May 2026.

Value Add VC. 'The SoftBank-OpenAI Stargate Deal: $500B in AI Infrastructure.' June 2026.

Virtus Minerals / Chemaf SA Acquisition Announcement. April 2026.

Wikipedia. '2026 Strait of Hormuz Crisis.' Updated July 8, 2026.

White House. Executive Order on AI National Policy Framework. December 11, 2025.

White House. Section 232 Critical Minerals Executive Order. January 15, 2026.

White House. Executive Order on Federal Reserve Payment Account Access Evaluation. May 19, 2026.

Prior Risk Chair series research: The Decoupling Doctrine (Issue 01), What the Boardroom Does Not Yet Know About AI (Issue 02), The Governance Deficit (Issue 03).