MOVEZ VENTURES

THE INTELLIGENCE INFRASTRUCTURE THESIS

Investing in the compute, power, physical systems, and applications that make artificial intelligence possible.

Joshua Hollis
Founder, Movez Ventures
July 2026

DISCUSSION DRAFT
This paper presents an original investment framework informed by publicly available research. It is not an offer to sell securities, investment advice, or a prediction that any specific technological timeline will occur.

Table of Contents

Executive Summary

Upon examination of the present artificial-intelligence boom, it is tempting to describe the opportunity as another software cycle. Movez Ventures rejects that narrow interpretation. Artificial intelligence is becoming an industrial system: software capability is inseparable from chips, memory, advanced packaging, networking, data centers, electricity, cooling, security, robotics, and the precision components that convert computation into physical work.

Leopold Aschenbrenner's Situational Awareness provides a provocative starting point. Its central economic claim is not merely that models may become more capable; it is that continued scaling could produce an extraordinary mobilization of capital and industrial capacity. The exact AGI timetable remains uncertain and should not be treated as fact. Yet the observable direction - larger clusters, stronger demand for accelerators, rising electricity loads, and increasing attention to supply-chain security - already supports an investable infrastructure thesis.[1]

> **We do not need to know which model wins to identify what every serious model, robot, and AI-native business will require.**

Movez Ventures therefore invests across three connected pillars: Intelligence Infrastructure, Embodied Intelligence, and Applied Intelligence. Intelligence Infrastructure includes accelerators, memory, semiconductor manufacturing, interconnects, data-center systems, power, cooling, and infrastructure software. Embodied Intelligence includes robots, actuators, bearings, reducers, motors, sensors, and manufacturing tools. Applied Intelligence includes defensible AI products and fintech platforms whose economics improve through automation, better underwriting, lower transaction costs, or new machine-native workflows.

The thesis is intentionally broader than a conventional 'AI software fund' but narrower than a generalist technology strategy. Every investment must connect to a measurable constraint, enabling layer, or durable application within the AI economy. Our primary question is simple: if intelligence becomes cheaper and more abundant, which scarce inputs, systems, and trusted interfaces become more valuable?

Sources: [1], [2], [3], [4]

1. The Central Argument

Throughout history, general-purpose technologies have required complementary infrastructure. Steam power depended on coal, machine tools, and railways. Electrification depended on generation, transmission, motors, and factories redesigned around distributed power. The internet depended on semiconductors, fiber, data centers, cloud software, and payment rails. Artificial intelligence follows the same pattern, although its infrastructure is unusually concentrated, capital intensive, and globally interdependent.

Aschenbrenner organizes his argument around orders of magnitude: more compute, better algorithms, and fewer practical limitations may create large capability gains. His most aggressive predictions are contestable, particularly the timing of AGI and an intelligence explosion. For an investor, however, the strongest portion of the argument is the connection between capability scaling and physical demand. Even without accepting a specific 2027 timeline, larger and more frequently used models increase requirements for compute, energy, networking, storage, cooling, security, and capital.[1]

This distinction is important. A venture thesis should not depend upon a single forecast being perfectly correct. The Movez thesis is constructed to remain attractive across several futures: rapid frontier-model progress, slower but steady enterprise adoption, greater use of specialized models, and widespread robotics. In each case, demand shifts toward efficient compute, reliable power, orchestration, physical automation, and trusted financial infrastructure.

The bottleneck principle

When a technology improves quickly, value often accumulates at the inputs that cannot scale as quickly. Model weights can be copied; a permitted power interconnection cannot. Software can be deployed globally; an advanced packaging line, a transformer, a precision reducer, or a high-purity bearing-steel supply chain requires time, capital, quality control, and manufacturing knowledge. Movez Ventures seeks ownership in these bottlenecks before scarcity is fully reflected in private-market valuations.

What we believe - and what we do not assume

- We believe AI usage and computational intensity will continue to rise, but we do not require a precise AGI date.
- We believe data centers increasingly function as integrated AI factories in which chips, networking, software, cooling, and power must be co-designed.[2]
- We believe physical AI will expand demand for precision motion components and industrial materials.
- We do not assume every AI application will produce venture-scale returns; distribution, proprietary data, margins, and workflow control remain decisive.
- We do not treat macro excitement as a substitute for technical diligence, customer proof, or disciplined entry price.

Sources: [1], [2]

2. From Capability Scaling to Industrial Mobilization

Situational Awareness describes a progression from model scaling to a 'trillion-dollar cluster.' Whether or not the largest estimates materialize on schedule, the thesis captures a real change in corporate planning: leading technology companies increasingly treat AI capacity as strategic infrastructure rather than ordinary information technology.[1]

The economic mechanism is circular. Better models create useful products; useful products generate revenue and strategic urgency; revenue and urgency justify larger infrastructure budgets; larger budgets fund more compute and research. This flywheel does not guarantee uninterrupted progress, but it explains why AI investment can remain elevated even when individual applications disappoint.

Evidence already visible

NVIDIA now describes itself as a 'data center scale AI infrastructure company' and explains that modern systems require extreme co-design across chips, networking, systems, software, and algorithms. It states that hundreds of thousands of GPUs can be connected to function as a single giant computer.[2] TSMC reports robust AI-related demand and continued growth in advanced technologies and CoWoS advanced packaging, demonstrating that the constraint is not simply chip design but the ability to manufacture and package complex systems at scale.[5]

The U.S. Department of Energy reports that hyperscale connection requests can reach 300 to 1,000 megawatts or more, with lead times that stretch local grids. This is industrial planning on the scale of power plants and heavy manufacturing, not conventional software deployment.[3]

Investment implication

The venture opportunity is not limited to companies that compete directly with incumbent semiconductor firms. It includes tools that increase utilization, reduce energy per useful token, improve interconnect performance, accelerate permitting, manage power, reduce water use, inspect components, automate facilities, and finance infrastructure. A modest improvement at the system level can create substantial value when multiplied across large clusters.

Constraint	Illustrative opportunities	Why venture can win
Compute	accelerators, inference chips, compilers	architectural specialization and rapid product cycles
Packaging & memory	chiplets, HBM tooling, thermal interfaces	capacity and yield become system constraints
Networking	optics, switches, interconnect software	cluster performance depends on communication
Power & cooling	grid software, generation, storage, liquid cooling	local scarcity and deployment speed matter
Operations	scheduling, observability, security	utilization and reliability determine economics

Sources: [1], [2], [3], [5]

3. Compute, Memory, Packaging, and Networking

The basic unit of AI economics is often described as a token, but tokens are not produced by software alone. Each useful token represents work performed across a stack: processors execute operations, high-bandwidth memory moves model parameters and intermediate values, packaging connects dies, networks coordinate nodes, storage supplies data, and software schedules the system. The cost, speed, and reliability of this stack determine which products can reach market.

Accelerated compute

General-purpose processors remain important, but modern training and inference depend heavily on parallel accelerators. NVIDIA's filing describes a full-stack platform of GPUs, CPUs, DPUs, interconnects, networking, software libraries, and application frameworks. This integration creates both a powerful incumbent advantage and openings for companies that solve specialized workloads, reduce cost per token, or make heterogeneous hardware easier to use.[2]

Memory and advanced packaging

As models and context windows grow, memory capacity and bandwidth become central. Advanced packaging enables multiple compute and memory components to operate as one system. TSMC states that CoWoS has experienced strong growth from surging AI demand. This supports investments in packaging equipment, materials, inspection, thermal management, substrate technology, chiplets, and design software.[5]

Networking becomes computation

At cluster scale, a processor that waits for data is an expensive idle asset. NVIDIA states that the data center has become the new unit of computing, with networking integral to the platform.[2] Movez therefore views optical interconnects, silicon photonics, switching, congestion control, network telemetry, and distributed-training software as core AI investments rather than adjacent communications bets.

A disciplined way to underwrite compute startups

- Measure performance on commercially relevant workloads, not isolated benchmarks.
- Evaluate total cost of ownership, including software migration, networking, memory, cooling, and developer labor.
- Determine whether the company owns defensible intellectual property or merely integrates commodity components.
- Test manufacturing feasibility, supply access, yield assumptions, and customer qualification timelines.
- Prefer products that improve cost per useful outcome rather than raw peak performance alone.

Sources: [2], [5], [6]

4. Power and the New AI Factory

The clearest bridge between Aschenbrenner's thesis and observable public data is electricity. The International Energy Agency projects that global data-center electricity consumption will roughly double to about 945 terawatt-hours by 2030 in its base case, with AI the most important driver of growth. The U.S. Department of Energy reports that data centers consumed about 4.4 percent of U.S. electricity in 2023 and could consume approximately 6.7 to 12 percent by 2028.[4][7]

These forecasts are uncertain, but they establish the scale of the planning problem. Power is not a fungible national commodity at the moment a project needs it. A data center requires deliverable capacity at a specific site, interconnection approval, transformers, switchgear, cooling, backup systems, and often long-term contracts. Delays in any component can strand billions of dollars of compute.

Our power thesis

- Generation: geothermal, nuclear, natural gas, renewables, and other firm or complementary sources where deployment economics are credible.
- Grid equipment: transformers, switchgear, substations, power electronics, and monitoring systems.
- Flexibility: storage, demand response, workload shifting, and software that aligns compute with power availability.
- Cooling and water: liquid cooling, heat exchangers, thermal materials, water treatment, and heat reuse.
- Development infrastructure: permitting, interconnection analytics, site selection, construction automation, and financing tools.

The objective is not to make a directional bet on one energy source. It is to invest in technologies that increase the speed, reliability, and efficiency with which power becomes usable AI capacity. The IEA expects renewables to meet a substantial share of incremental supply while natural gas, coal, and later nuclear also contribute, which implies a diversified and region-specific buildout.[8]

> **In the AI economy, a megawatt delivered on time may be more valuable than a theoretical megawatt delivered years later.**

Sources: [3], [4], [7], [8]

5. Security, Sovereignty, and Resilience

Situational Awareness devotes significant attention to laboratory security, geopolitical competition, and the possibility that advanced AI becomes a national-security priority. Movez does not adopt every geopolitical conclusion in the series. Nevertheless, the underlying investment point is sound: as model weights, training methods, chips, and clusters become more valuable, customers will pay more for security, provenance, resilience, and sovereign control.[1]

Security is an infrastructure layer

AI security includes conventional cybersecurity, but extends further: protecting model weights, securing data pipelines, verifying hardware and firmware, controlling privileged access, monitoring clusters, preventing model exfiltration, and documenting how systems were trained and deployed. Startups that embed security into AI infrastructure may enjoy stronger retention than point solutions added after deployment.

Sovereign AI and supply-chain concentration

The semiconductor and materials supply chains are geographically concentrated. TSMC's planned U.S. expansion, including advanced semiconductor manufacturing and packaging, reflects the strategic importance of domestic capacity.[9] USGS data also show that nonfuel minerals underpin manufacturing and that production of many commodities remains concentrated across a limited number of countries.[10]

What Movez will seek

- hardware-rooted security and confidential computing
- model and data provenance, testing, and governance
- secure orchestration for regulated and sovereign environments
- domestic or allied manufacturing technologies that relieve strategic bottlenecks
- dual-use technologies with clear compliance, governance, and responsible-customer frameworks

This area requires judgment. National-security relevance can create durable demand, but it can also produce long sales cycles, export restrictions, customer concentration, and ethical risk. Movez will evaluate governance and end use alongside market size.

Sources: [1], [9], [10], [11]

6. Robotics and Embodied Intelligence

Artificial intelligence becomes economically transformative when it can act. Robots convert perception, planning, and learned policies into motion. This makes robotics a synthesis of AI software, sensors, embedded compute, motors, actuators, reducers, bearings, power systems, materials, manufacturing, and safety engineering.

NVIDIA describes physical AI as an end-to-end platform spanning data-center infrastructure, models, simulation, embedded computing, and software for automation and robotics.[2] This framing is useful because it shows why robotics demand can benefit multiple layers at once. A capable robot must be trained and simulated in data centers, deployed on embedded hardware, and constructed from reliable physical components.

Why robotics may now cross a threshold

- Improved foundation models can reduce the amount of task-specific programming.
- Simulation and synthetic data can accelerate training and testing.
- Better sensors and embedded accelerators improve perception and control.
- Labor scarcity and reshoring create economic pressure for automation.
- Manufacturing scale can lower component costs and improve reliability over time.

The opportunity is not limited to humanoids. Movez will consider industrial arms, mobile manipulation, warehouse systems, medical robots, agricultural automation, inspection platforms, autonomous machines, and component suppliers. Humanoids may create a large market, but purpose-built machines can reach positive unit economics sooner.

The underwriting standard

Robotics companies should be judged by task economics, reliability, duty cycle, safety, service requirements, and deployment time - not by demonstration videos. A technically impressive prototype may fail as a venture investment if it requires excessive teleoperation, maintenance, or customer integration. Conversely, a seemingly unglamorous component can become extremely valuable if it is qualified into many platforms and difficult to replace.

Sources: [2], [12], [13]

7. Bearings, Motion Systems, and Material Inputs

Upon examination of a robot's mechanical system, bearings appear ordinary only because their function is fundamental. They support loads, constrain motion, reduce friction, and help joints rotate accurately over repeated cycles. SKF states that bearings in humanoid robots directly influence precision, efficiency, noise, and service life from shoulder joints to wrist articulation.[12] NSK's technical research describes bearings used in robot reducers and servomotors, including the need for precise control of rotation, stopping, and positioning under harsh operating conditions.[13]

Why bearings are an investable AI theme

A robot is only as useful as its uptime and repeatability. Small improvements in friction, backlash, weight, contamination resistance, acoustic performance, or bearing life can improve battery endurance, payload, precision, maintenance cost, and customer trust. Specialized thin-section bearings also contribute to smaller, lighter, and lower-torque robotic systems.[14]

The investment opportunity extends beyond finished bearings. It includes high-purity steels, ceramics, coatings, lubrication, heat treatment, precision grinding, metrology, condition monitoring, recycling, and manufacturing automation. These businesses may be less visible than robot brands, yet they can sell across many platforms and industries.

Materials are part of the robotics thesis

Bearing performance depends on material quality. SKF identifies 100Cr6 - approximately 1 percent carbon and 1.5 percent chromium - as the standard steel for many bearing rings and washers.[15] NSK explains that through-hardened chromium steel is commonly used for rings and rolling elements, while chromium-molybdenum and nickel-chromium-molybdenum steels may be used where shock resistance is required.[16] This connects robotics growth to steel cleanliness, chromium, nickel, molybdenum, ceramics, polymers, lubricants, and the industrial processes that transform these inputs into precision components.

Robotic motors can also rely on permanent magnets and power electronics. U.S. Department of Energy materials planning identifies neodymium, praseodymium, and dysprosium as important inputs for permanent magnets in electric motors and other advanced equipment.[17] The precise bill of materials varies by robot architecture, but the broader conclusion is clear: scaling physical AI requires resilient material and manufacturing supply chains, not merely better algorithms.

Robot subsystem	Critical components	Relevant materials / processes
Joint	bearings, reducer, encoder, seals	bearing steel, chromium, nickel, ceramics, lubricants, precision grinding
Actuation	motor, magnets, power electronics	copper, electrical steel, rare earths or alternative magnet materials, semiconductors
Structure	frames, links, fasteners	aluminum, steel, titanium, composites, casting and machining
Perception	cameras, lidar, force and torque sensors	semiconductors, optics, specialty glass, wiring
Energy	battery, power distribution, thermal system	lithium-ion materials, copper, cooling fluids, thermal interfaces

Sources: [12], [13], [14], [15], [16], [17]

8. Applied Intelligence and Fintech

Infrastructure is the foundation of the thesis, but infrastructure alone does not capture all value. Applications turn falling intelligence costs into customer outcomes. Movez will invest selectively in AI-native products where the company owns distribution, workflow, proprietary data, trust, or a structural cost advantage.

The application filter

- The product completes a valuable workflow rather than merely generating content.
- The company can improve through proprietary feedback, data, or integration depth.
- Gross margins remain attractive after inference and human-review costs.
- The customer can quantify revenue gain, labor savings, risk reduction, or faster cycle time.
- The product has a credible path to defend against foundation-model vendors and incumbents.

Why fintech belongs in the strategy

Financial services are information-intensive, regulated, and rich in repetitive decisions. AI can improve fraud detection, underwriting, compliance, customer service, treasury, capital markets, accounting, and wealth management. At the same time, financial institutions require auditability, security, explainability, and controlled deployment. This creates opportunities for companies that combine AI capability with domain-specific data, workflows, and regulatory trust.

Movez will distinguish between fintech that simply adds an AI interface and fintech whose unit economics or risk performance fundamentally improve through intelligence. Attractive areas include machine-native payments and identity, AI-assisted underwriting with strong governance, fraud and cybersecurity, capital-markets infrastructure, programmable treasury, compliance automation, and vertical financial operating systems.

Applications create demand for infrastructure

Applied and infrastructure investments reinforce each other. More useful applications generate more tokens, transactions, simulations, and robotic actions. That demand improves utilization of compute and data centers, while better infrastructure lowers the cost of creating new applications. Movez intends to use this feedback loop as a sourcing advantage: understanding bottlenecks helps identify applications that become newly possible, and understanding applications reveals the next bottlenecks.

Sources: [2], [18]

9. Movez Ventures Investment Framework

Movez Ventures organizes the opportunity into three pillars. The pillars are not separate funds; they are a method for maintaining coherence while investing across a complex value chain.

Pillar	What it includes	Core question
Intelligence Infrastructure	compute, chips, memory, packaging, networking, data centers, power, cooling, security, orchestration	What must exist for intelligence to be trained, deployed, and trusted at scale?
Embodied Intelligence	robots, actuators, motors, bearings, reducers, sensors, manufacturing tools, critical materials	What converts intelligence into safe, economical, repeatable physical work?
Applied Intelligence	vertical AI, developer products, enterprise systems, fintech, financial infrastructure	Where does abundant intelligence create durable customer value and defensible distribution?

Stage and company profile

The intended emphasis is early-stage investing, from pre-seed through Series A, with selective later-stage participation where Movez has differentiated access or technical conviction. Ideal companies combine a large market with a narrow initial wedge, technically credible founders, clear customer pain, and a path toward capital-efficient proof before scaling.

The Movez operating advantage

Movez Ventures intends to be more than a source of capital. The firm will support product definition, full-stack development, AI implementation, go-to-market testing, recruiting, operational workflows, and fundraising preparation. This operating model is particularly useful at the earliest stages, when a founder's speed of iteration can matter more than organizational scale.

What we will avoid

- undifferentiated model wrappers without workflow ownership
- capital-intensive projects with no credible financing path or milestone structure
- hardware companies whose cost and qualification timelines are hidden behind optimistic prototypes
- robotics businesses dependent on indefinite teleoperation or uneconomic service labor
- fintech products that ignore licensing, compliance, fraud, or balance-sheet risk
- companies whose only thesis is that AI market growth will rescue weak fundamentals

Sources: [1], [2], [3]

10. Portfolio Construction and Diligence

A coherent thesis still requires portfolio discipline. Infrastructure and robotics companies often need more capital and time than software companies. Movez should therefore size initial checks, reserves, and ownership targets according to technical and financing risk rather than applying one template to every company.

Illustrative construction principles

- Build a concentrated but diversified portfolio across the three pillars, avoiding excessive exposure to one technical architecture or customer.
- Reserve meaningful capital for companies that achieve technical validation and commercial pull.
- Syndicate capital-intensive deals with investors experienced in hardware, energy, industrials, or government markets.
- Use milestone-based financing where manufacturing, certification, or infrastructure deployment risk is high.
- Track look-through exposure to GPU vendors, foundries, cloud providers, power markets, rare materials, and government policy.

Diligence by category

Category	Key diligence questions
AI infrastructure	Does performance persist at production scale? What are utilization, migration, and integration costs?
Semiconductors / hardware	Can it be manufactured at target yield and cost? Are supply and packaging secured?
Energy / data centers	Is power deliverable, permitted, financed, and contracted? What is the interconnection timeline?
Robotics	What is the demonstrated duty cycle, failure rate, maintenance burden, and customer payback?
Components / materials	What qualification creates switching costs? Can quality scale without yield loss?
Fintech	Who bears regulatory, fraud, credit, liquidity, and model risk? Is the AI decision auditable?

For LPs, the critical point is that thesis breadth will not become style drift. Every investment memo should identify the specific AI-economy constraint or application, the evidence that demand is durable, the source of company-level defensibility, and the milestones that would falsify the investment case.

11. Risks and Disconfirming Evidence

A serious investment thesis must explain how it could be wrong. Situational Awareness is intentionally forceful, and some of its predictions may prove early, exaggerated, or incorrect. Movez should benefit from the framework without becoming captive to it.

Principal risks

- Capability progress slows because scaling produces diminishing returns or data constraints.
- Algorithmic efficiency reduces infrastructure demand faster than usage expands it.
- Overbuilding creates excess data-center, chip, or power capacity and compresses returns.
- Incumbents capture most value through vertical integration, leaving startups with weak margins.
- Energy, permitting, export controls, or supply-chain disruptions delay deployments.
- Robotics adoption remains limited by reliability, safety, customer integration, or total cost.
- AI applications commoditize as foundation models improve and distribution consolidates.
- Valuations price in success before companies demonstrate durable economics.

How the thesis adapts

If model scaling slows, efficiency, inference, edge systems, and applied automation may become more important than frontier training. If compute becomes abundant, value may shift toward power, distribution, trusted applications, and physical deployment. If humanoids develop slowly, industrial and purpose-built robotics can still create component demand. If one accelerator architecture dominates, software, networking, cooling, security, and power remain necessary.

Signals that would cause us to reduce exposure

- multi-year stagnation in commercially useful model capability despite rising spend
- sustained collapse in accelerator utilization or data-center pricing without offsetting application growth
- customer evidence that AI projects fail to produce measurable returns at scale
- material deterioration in robot deployment economics or reliability
- a structural reduction in power demand per unit of AI value that usage growth does not offset
- private-market valuations that eliminate plausible venture returns

The purpose of these conditions is not to predict pessimism. It is to preserve intellectual honesty. Movez Ventures seeks conviction with the ability to update, not certainty unsupported by evidence.

Sources: [1], [4], [7]

Conclusion: Investing in What Intelligence Requires

Upon examination of the AI economy, the most important conclusion is that intelligence is not weightless. It is produced by machines, powered by electricity, connected by networks, cooled by physical systems, secured by institutions, and increasingly expressed through robots. Even the most advanced software remains dependent on industrial capacity.

Situational Awareness is valuable because it forces the reader to consider scale. Its most dramatic timelines should be treated as hypotheses, not certainties. Yet its infrastructure insight is already visible in corporate filings, government power forecasts, semiconductor investment, and robotics engineering. The world is preparing for far more computation, and that preparation creates opportunities across the entire enabling stack.[1][2][3]

Movez Ventures will invest in the companies that make intelligence possible, economical, secure, and useful. This includes accelerators and infrastructure software, but also transformers and cooling systems; robotics models, but also reducers, motors, bearings, and the steels and materials from which those components are made; AI applications, but only where technology becomes a defensible business.

> **The next great AI companies will not only create intelligence. They will build the systems through which intelligence enters the economy.**

The firm's standard is therefore simple. We ask which resources remain scarce as intelligence becomes abundant, which systems every credible participant must use, and which founders possess the technical and operational ability to convert those constraints into enduring companies. That is the Intelligence Infrastructure Thesis.

References

[1] Aschenbrenner, Leopold. Situational Awareness: The Decade Ahead. June 2024. https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf

[2] NVIDIA Corporation. Annual Report on Form 10-K, fiscal year ended January 25, 2026, pp. 3-5, 35. https://investor.nvidia.com/files/doc_financials/2026/q4/10K-NVDA.pdf

[3] U.S. Department of Energy, Secretary of Energy Advisory Board. Recommendations on Powering Artificial Intelligence and Data Center Infrastructure. July 30, 2024. https://www.energy.gov/sites/default/files/2024-08/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf

[4] International Energy Agency. Energy and AI: Energy Demand from AI. 2025. https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai

[5] Taiwan Semiconductor Manufacturing Company. 2025 Annual Report; discussion of AI demand and CoWoS advanced packaging. https://investor.tsmc.com/static/annualReports/2025/english/index.html

[6] Taiwan Semiconductor Manufacturing Company. 2025 Annual Report, Business Activities. https://investor.tsmc.com/static/annualReports/2025/english/pdf/2025_tsmc_ar_e_ch5.pdf

[7] U.S. Department of Energy. DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centers. December 20, 2024. https://www.energy.gov/articles/doe-releases-new-report-evaluating-increase-electricity-demand-data-centers

[8] International Energy Agency. Energy and AI: Energy Supply for AI. 2025. https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai

[9] Taiwan Semiconductor Manufacturing Company. TSMC Intends to Expand Its Investment in the United States to US$165 Billion to Power the Future of AI. March 4, 2025. https://pr.tsmc.com/english/news/3210

[10] U.S. Geological Survey. Mineral Commodity Summaries 2026. February 2026. https://pubs.usgs.gov/publication/mcs2026

[11] Aschenbrenner, Leopold. Lock Down the Labs: Security for AGI; The Free World Must Prevail. 2024. https://situational-awareness.ai/lock-down-the-labs/

[12] SKF. Humanoid Robotic Solutions. Accessed July 2026. https://www.skf.com/sg/industries/humanoid-robotics

[13] Ishiguro, Hiroshi. Technical Trend of Industrial Machinery Bearings. NSK Motion & Control No. 30, June 2019, pp. 13-17. https://www.nsk.com/content/dam/nsk/common/tools-resources/rd/pdf/techJournal/etj-0030.pdf

[14] NSK. N Series Thin-Section Ball Bearings. Accessed July 2026. https://www.nsk.com/products/ball-bearings/thin-section-ball-bearings/

[15] SKF. Components and Materials / Bearing Basics. Accessed July 2026. https://www.skf.com/in/products/rolling-bearings/principles-of-rolling-bearing-selection/general-bearing-knowledge/bearing-basics/components-and-materials

[16] NSK. Materials for Bearings: Ensuring Reliability and Quality. https://www.nsk.com/content/dam/nsk/eu/en_gb/documents/bearings-europe/P_TI-0106_EN.pdf

[17] U.S. Department of Energy. FY 2024 Congressional Justification, Crosscutting Activities; critical materials used in permanent magnets and advanced energy systems. https://www.energy.gov/sites/default/files/2023-03/doe-fy2024-budget-volume-2-crosscutting.pdf

[18] NVIDIA. State of AI in Financial Services 2026. https://www.nvidia.com/en-us/industries/finance/ai-financial-services-report/

Method and Attribution Note

This paper is an original Movez Ventures investment framework. It summarizes and critiques selected arguments from Situational Awareness rather than reproducing that work. Current market and energy figures are drawn from the cited primary or institutional sources. Forecasts are inherently uncertain. Company references are illustrative and do not constitute recommendations to buy or sell securities.