BREAKWATER

CAPITAL MARKETS BRIEFING

The Sovereign AI Premium

Why the defining capital-allocation decision of the AI age is not how much to spend, but where to retain control.

AUTHORED BY

Mark Hayes, Partner and Head of Breakwater Capital Markets

VALUATION OPERATING SYSTEM - ADVISORY SOLUTION

A Breakwater Capital Markets briefing

CONTENTS

EXECUTIVE SUMMARY

AI is no longer a technology decision. It is a capital, sovereignty, and control decision.

Artificial intelligence is a determinant of capital intensity, market access, operating resilience, regulatory exposure, and corporate power. Public companies entering their next strategic-planning cycle must ask a more consequential question than how much they should invest in AI. They must decide which parts of the emerging AI stack they need to control, where they can safely rely on external providers, how sovereignty requirements will reshape their products and geographic architecture, and which investments can produce durable economic advantage rather than temporary technological relevance.

Breakwater Capital Markets calls this the **Sovereign AI Premium**. It is not a valuation premium awarded for an AI strategy or attaching artificial intelligence to an existing product. It is the premium that may accrue to companies capable of converting AI into measurable growth, operating leverage, and stronger customer economics while retaining sufficient control over critical data, infrastructure, intellectual property, market access, and operating continuity.

Its opposite is the **AI Dependency Discount**. This discount arises when a company's AI strategy creates uncontrolled capital intensity, excessive reliance on a small number of vendors, exposure to shifting export controls, fragmented regulatory obligations, uncertain intellectual-property rights, rising infrastructure costs, or assets that may become obsolete before earning an adequate return.

Investors will increasingly distinguish among companies based on the quality of their position within the emerging AI political economy. Companies that combine differentiated data, disciplined capital allocation, geographic adaptability, infrastructure access, credible governance, and freedom to change suppliers should be more likely to command a premium. Companies that spend heavily without explaining how their investments improve returns, strengthen competitive advantage, or reduce strategic vulnerability may instead face greater skepticism, a higher cost of capital, and a lower valuation multiple.

> *The next strategic-planning cycle should treat AI not as a technology initiative but as a portfolio of economic, operating, and geopolitical decisions.*

SECTION 01

AI Has Moved From the Technology Budget to the Corporate Balance Sheet

The first phase of corporate AI adoption was largely experimental. Companies funded pilots, purchased software licenses, created innovation teams, and tested productivity tools. The financial commitments were often modest relative to enterprise capital budgets, and the consequences of failure were limited. Many projects could be financed through operating expenses and governed within the technology organization. That phase is ending.

The next generation of AI adoption will require decisions involving cloud capacity, data centers, electricity, cooling, advanced semiconductors, proprietary data, cybersecurity, workforce redesign, product architecture, intellectual property, long-term vendor contracts, and geographically distributed infrastructure. These are not ordinary software decisions. They affect fixed assets, depreciation, working capital, margins, supply-chain resilience, market entry, and the strategic flexibility of the enterprise.

The extraordinary investment cycle among hyperscale technology companies illustrates the direction of travel. A recent company has said that it expects capital expenditure of between billions of dollars in 2026, with spending directed toward AI computing capacity, frontier-model development, cloud demand, consumer services, advertising products, and other strategic opportunities. Management has also emphasized that the expanding asset base will increase depreciation and place pressure on the income statement. The lesson for other public companies is not that they should imitate the company's scale. It is that AI strategy has become inseparable from infrastructure economics, utilization, financing, depreciation, and the credibility of management's return assumptions.

Most companies will not build frontier models or hyperscale data centers. They will nevertheless inherit the economics of this investment cycle through cloud pricing, model-access fees, specialized software costs, energy availability, semiconductor constraints, data-residency requirements, and the bargaining power of infrastructure providers. A company can remain asset-light in accounting terms while becoming strategically dependent on capital-intensive systems owned by others. Boards must therefore ask not only how much AI capital the company intends to deploy, but whose capital, infrastructure, and regulatory permissions the business ultimately depends upon.

A company that rents almost its entire AI stack may avoid near-term capital expenditure while accepting greater pricing risk, supplier concentration, and reduced control over future product economics. A company that attempts to own too much of the stack may protect itself from dependence while destroying returns through underutilized assets and rapid technological obsolescence. The capital-allocation problem is not ownership versus outsourcing in the abstract. It is determining where ownership creates economic or strategic advantage, where contractual access is sufficient, and where redundancy is worth paying for.

SECTION 02

Geopolitics Is Becoming Part of the Technology Architecture

Public companies have traditionally treated geopolitics as an external risk affecting tariffs, sanctions, currencies, supply chains, and market access. Artificial intelligence changes this relationship because government policy can now reach directly into the architecture of a company's products and operations.

Access to advanced semiconductors, computing capacity, model technology, cloud infrastructure, technical talent, and cross-border data flows is increasingly conditioned by national-security policy. Current United States export rules distinguish among countries, end users, corporate ownership structures, equipment classifications, and approved entities. Certain authorizations introduced in July 2026 provide differentiated treatment to named American AI companies and specific entities in the United Arab Emirates, while other destinations and end users remain subject to tighter restrictions. The commercial availability of AI infrastructure is therefore determined not only by technological capability and price, but also by diplomatic relationships, jurisdiction, ownership, end use, and regulatory approval.

This creates a new category of corporate exposure. A company may design a global product around a particular semiconductor, model, or cloud service only to discover that the architecture cannot be deployed in every market. A supplier may be commercially available in one jurisdiction and restricted in another. A national government may require local computing capacity, domestic data storage, security review, local model adaptation, a joint venture, or local ownership as a condition of market access.

One company has recently warned investors that divergent approaches to AI regulation, data flows, sanctions, export controls, geopolitical instability, and national market interventions may affect operating strategies, costs, market access, product continuity, and profitability. Its disclosures specifically connect shifting export-control policies with accelerating sovereignty initiatives among international governments and customers. The world is unlikely to divide cleanly into two self-contained technology blocs. The more plausible outcome is a fragmented system of overlapping restrictions, preferred suppliers, localization mandates, regulatory regimes, and political bargains.

A global company may rely on American semiconductors, Chinese telecommunications equipment, European regulation, Gulf-financed data centers, Indian digital infrastructure, and locally developed applications. Its AI architecture will not sit outside geopolitics. It will embody the compromises and dependencies created by that environment.

> *AI geopolitics is becoming product architecture.*

SECTION 03

Sovereignty Is Becoming a Customer Requirement

Sovereign AI is frequently defined as a government's desire to develop a domestic foundation model. That definition is too narrow for corporate strategy. Sovereign AI is the ability of a country, institution, or enterprise to employ artificial intelligence while maintaining acceptable control over its data, infrastructure, intellectual property, security, technical dependencies, and decision-making authority.

Sovereignty does not require complete self-sufficiency. Few countries or companies could achieve it. It requires sufficient technical capability, contractual freedom, and operating resilience to understand critical systems, determine where sensitive information travels, maintain service through disruption, and change suppliers without unacceptable cost.

This demand will increasingly influence corporate procurement. Customers will ask where inference occurs, whether data are used to train external models, which personnel or subcontractors can access their information, what laws govern that information, whether systems can be deployed in a private environment, whether outputs can be audited, and what happens when the vendor relationship ends.

These requirements should not be treated only as compliance costs. In many industries, they will become sources of differentiation. A cloud provider that can offer nationally compliant infrastructure may gain access to government, financial-services, healthcare, energy, and defense-related customers. An industrial company that can operate AI systems locally rather than transmitting sensitive information across borders may win business that competitors cannot serve. A software company that supports multiple clouds, models, languages, and data-residency regimes may become a strategic orchestration layer. A professional-services company that can combine technology with regulatory implementation and organizational transformation may capture substantial value without owning the underlying model.

The opportunity lies in identifying where sovereignty is becoming part of the customer value proposition rather than merely a legal obligation.

REGULATORY MARKER · EUROPEAN UNION

Obligations for providers of general-purpose AI models entered into application on August 2, 2025. They include technical-documentation requirements, information for downstream providers, copyright policies, and summaries of training content. The European Commission's enforcement powers apply from August 2, 2026. Additional transparency requirements for certain AI systems and AI-generated content also apply from that date.

Companies integrating, adapting, or distributing AI systems must understand not only their own obligations but also the documentation, controls, and contractual commitments available from their suppliers. The companies that convert regulatory complexity into reusable product capabilities can strengthen their market positions. Those that respond through bespoke manual processes in every jurisdiction may accumulate cost without creating advantage.

Sovereignty should therefore be designed into products, contracts, data architecture, and infrastructure from the beginning.

SECTION 04

The AI Stack Is Becoming a Map of Corporate Power

Every public company should enter its strategic-planning process with a clear map of the AI stack upon which its future economics will depend.

At the base of the stack are electricity, land, cooling, connectivity, semiconductors, data centers, and cloud infrastructure. Above them sit data platforms, cybersecurity systems, foundation models, development tools, orchestration software, applications, autonomous agents, workflow integrations, and customer interfaces. Surrounding the entire stack are financing, technical talent, intellectual-property rights, standards, regulation, and geopolitical permissions. Few companies should attempt to own all these layers. Every company must identify which layers determine its competitive advantage and which could become unacceptable points of dependence.

Companies generally occupy one or more of five positions.

INFRASTRUCTURE BUILDERS	Supply power, chips, telecommunications networks, data centers, or computing capacity.
PLATFORM ORCHESTRATORS	Allow customers to use multiple models and infrastructure environments through a common control layer.
DOMAIN OWNERS	Possess proprietary data, workflows, customer relationships, and specialized expertise.
SOVEREIGN ENABLERS	Help governments and regulated enterprises deploy AI within national or sector-specific requirements.
AI USERS	Most companies: purchasing increasingly standardized capabilities and applying them to existing operations.

No position is inherently superior. The strategic failure is occupying one position while allocating capital as though the company occupied another.

A company that lacks the scale and expertise to operate infrastructure should not invest as though it were a hyperscaler. A company possessing unique industry data should not allow that data to become an undifferentiated input into another firm's platform without determining how value will be shared. A business serving highly regulated or strategically sensitive customers should not design products that can function only through a single foreign cloud or model provider.

The planning process should identify the layers at which the company intends to differentiate, the layers at which it will partner, and the layers at which it requires credible alternatives. This is the AI version of make-versus-buy, but the consequences are greater. A conventional outsourcing decision can often be reversed through a new contract. An AI architecture can embed data structures, development tools, security controls, employee capabilities, customer integrations, and operating routines so deeply that changing providers becomes a multiyear transformation.

The initial commercial discount offered by a supplier may conceal a long-term strategic premium paid by the customer.

SECTION 05

The Sovereignty Paradox

Many governments and enterprises are responding to geopolitical uncertainty by demanding local data centers, domestic cloud regions, local-language models, national infrastructure, and technology transfer. These requirements may reduce certain risks. They may keep sensitive information inside a jurisdiction, support local employment, and strengthen domestic technical capability. Yet physical localization does not by itself create practical autonomy.

A company may legally own a data center while relying on foreign chips, software, maintenance, and security updates. It may operate a locally branded model but depend on an external cloud provider for training and inference. It may store data domestically while lacking the internal expertise required to audit the system that processes them.

Formal control and practical autonomy are not the same. This is the sovereignty paradox. A company or country can own the asset while remaining dependent on the knowledge, infrastructure, and permissions needed to operate it.

> *The practical test of sovereignty is not where the equipment is located. It is whether the organization can understand, secure, alter, and ultimately replace the system.*

SECTION 06

The Sovereign AI Premium

There is no automatic valuation premium for AI exposure. Public markets may temporarily reward a compelling narrative, but durable value continues to depend on expected cash flows, growth, margins, returns on invested capital, competitive durability, and risk.

The Sovereign AI Premium is the additional value investors may assign to a company that captures the economic benefits of AI without surrendering control over the factors determining its future economics. The premium consists of five elements.

01 **Growth** — the ability to create new products, improve customer outcomes, enter sovereign or regulated markets, and expand addressable demand.

02 **Operating leverage** — the ability to automate workflows, increase asset utilization, raise employee capacity, improve decision quality, and reduce the marginal cost of serving customers.

03 **Durability** — ownership or control of differentiated data, distribution, infrastructure, intellectual property, customer relationships, or industry workflows that cannot be readily replicated.

04 **Resilience** — the ability to continue operating through supplier disruption, geopolitical restrictions, cyber incidents, regulatory changes, and shortages of infrastructure or technical capacity.

05 **Credibility** — management's ability to explain where AI is producing value, how investments are governed, what assumptions support the investment case, and what evidence will determine whether further capital is released.

Against this potential premium must be charged the AI Dependency Discount. The discount reflects excessive capital intensity, weak utilization, uncertain monetization, supplier concentration, regulatory fragmentation, cyber exposure, intellectual-property uncertainty, infrastructure shortages, and the possibility that rapidly improving technology will render expensive assets obsolete.

The strategic objective should not be to maximize AI expenditure. It should be to maximize the difference between the Sovereign AI Premium and the AI Dependency Discount. The relationship can be expressed as a management framework rather than a literal valuation formula.

THE FRAMEWORK

Sovereign AI Premium = Growth + Operating Leverage + Durability + Resilience + Credibility

AI Dependency Discount = Capital Intensity + Concentration + Regulatory Friction + Obsolescence Risk + Strategic Opacity

The same investment can produce very different valuation consequences depending on the company making it. A data-center investment may create substantial value for a business with contracted demand, favorable power access, low-cost financing, high expected utilization, and infrastructure expertise. The same investment may destroy value for a company with uncertain workloads, no operating advantage, and a high cost of capital.

Developing a proprietary model may be rational for a company possessing unique data, specialized technical requirements, sufficient scale, and an opportunity to create defensible intellectual property. It may be an expensive act of corporate vanity for a company that could obtain comparable capabilities from several competing providers.

Investors will not award a lasting premium for technological ambition divorced from economic logic.

SECTION 07

The Decisions Most Likely to Drive a Valuation Premium

The market will eventually treat basic AI adoption as ordinary. A lasting premium will require evidence that AI is improving the quality, duration, and resilience of earnings.

01 **Differentiated growth.** Companies must demonstrate that AI produces products, pricing power, retention, customer outcomes, or access to markets that competitors cannot easily reproduce.

02 **Operating leverage.** Productivity must translate into additional output, lower expense, accelerated growth, improved service, or greater organizational capacity. Hours theoretically saved have little valuation significance unless they alter financial performance.

03 **Capital discipline.** Investors will distinguish between companies purchasing useful capability and those accumulating infrastructure, vendor contracts, or acquisitions without sufficient demand or economic return.

04 **Control over scarce or differentiated assets.** Proprietary data, customer distribution, embedded workflows, power access, infrastructure locations, regulatory permissions, and specialized talent may prove more durable than ownership of a particular model.

05 **Regulatory and geographic adaptability.** Companies able to operate across sovereign environments without reconstructing their entire platform should address more demand with lower incremental complexity.

06 **Resilience.** An operationally tested ability to maintain service through supplier disruption, cyber incidents, regulatory intervention, and geopolitical restrictions can reduce the risk attached to future cash flows.

07 **Disclosure credibility.** Markets will reward management teams that articulate their investment logic, measure results honestly, abandon weak projects, and distinguish strategic necessity from technological fashion.

> *The Sovereign AI Premium will accrue not to the companies that speak most confidently about artificial intelligence, but to those that make their AI economics legible.*

SECTION 08

Capital Allocation Should Begin With the Control Question

The first question in the capital-allocation process should be: what must the company control to preserve its ability to compete? Control does not always require full ownership. It can be achieved through intellectual-property rights, contractual protections, portable data, multiple suppliers, reserved capacity, local deployment, joint ventures, minority investments, or the internal expertise required to replace a provider. The answer will vary by industry.

A bank may require strong control over customer data, model governance, explainability, and operating resilience but may not need to own data centers. A manufacturer may need control over process data, plant continuity, robotics integration, and digital twins while purchasing general-purpose models. A pharmaceutical company may need proprietary control over scientific data and validated applications but can rent much of its computing capacity. A media company may consider audience relationships and content rights more strategically important than model ownership.

The board and management team should identify a limited number of control points that directly determine economic advantage or existential risk. Capital should then be concentrated around them.

> *Spreading investment evenly across the AI stack is likely to produce mediocrity. Strategic control is created through concentration.*

SECTION 09

Companies Must Choose Which Infrastructure Risks to Carry

AI infrastructure creates three principal exposures. The appropriate capital structure must balance all three.

UTILIZATION RISK

Arises when a company acquires more capacity than it can economically employ.

ACCESS RISK

Arises when the company depends on infrastructure that may become unavailable, restricted, too expensive, or allocated to more powerful customers.

OBSOLESCENCE RISK

Arises because semiconductors, models, and technical architectures can improve faster than traditional capital assets.

Companies with large, predictable, and differentiated workloads may justify owning or reserving dedicated infrastructure. Companies with volatile or experimental demand should generally preserve variable costs until the use cases mature. Most enterprises should combine contracted capacity for critical workloads with competitive external markets for noncritical demand. The objective is not merely to minimize current cost. It is to preserve options.

Long-term contracts can secure capacity and reduce pricing uncertainty, but they can also create take-or-pay exposure. Building infrastructure can provide control but increase depreciation and utilization pressure. Relying entirely on external capacity can preserve flexibility but expose critical operations to shortages and repricing.

The optimal architecture will often be a portfolio. Critical workloads may require reserved or sovereign capacity. Differentiated applications may justify specialized infrastructure. Routine productivity tools may remain fully external and interchangeable.

Capital allocation should reflect the strategic importance and economic maturity of each workload.

SECTION 10

AI Investments Should Be Managed as a Portfolio of Real Options

The uncertainty surrounding AI technology, regulation, and market structure makes conventional single-point business cases inadequate for many investments. Companies should instead manage AI capital as a portfolio of real options.

Core productivity investments should face strict near-term economic tests. Management should identify the process being changed, the expected improvement in cost, quality, capacity, or cycle time, the cost of implementation, and the time required to capture the benefit. Productivity claims should be adjusted for integration costs, human review, model errors, control requirements, software expenses, and the possibility that saved employee time is not converted into lower cost or additional output.

Customer-facing investments should be evaluated according to revenue quality rather than adoption alone. The relevant measures are willingness to pay, retention, usage intensity, cross-selling, customer acquisition, gross margin, and whether the capability strengthens or commoditizes the broader product.

Sovereign-market investments should be evaluated as market-access options. A local data environment, regulatory certification, regional model, or domestic partnership may appear inefficient when assessed against a single contract. It may become valuable when it opens access to governments, regulated industries, and national champions across an entire market.

Strategic-infrastructure investments should be assessed against the value of capacity assurance, control, and future bargaining power as well as immediate financial return. These projects require especially rigorous governance because the language of strategic necessity can easily be used to excuse weak economics.

Frontier investments should be staged. The company should purchase learning before purchasing scale. Small investments, partnerships, pilot facilities, minority positions, and contractual options can preserve access to emerging capabilities without forcing the balance sheet to absorb the full risk of an immature architecture.

> *Every significant AI investment should have a defined next decision rather than only an imagined final destination.*

SECTION 11

The AI Capital Allocation Scorecards

A common framework so that AI expenditure competes for capital on comparable terms — and so the assumptions behind judgment are visible.

AI investments should compete for capital under a common framework rather than advance through separate technology, business-unit, or geopolitical narratives. The purpose of the following scorecards is to make different types of AI expenditure comparable while recognizing that financial return alone does not capture the value of resilience, sovereign-market access, strategic control, and future optionality.

Each proposal should receive three primary scores. The first assesses economic value and capital efficiency. The second assesses sovereignty, resilience, and geopolitical exposure. The third assesses the likelihood that the investment will create a durable competitive advantage and support a valuation premium. Together, the primary scorecards produce a composite score of 100 points.

The framework should not replace judgment. It should expose the assumptions behind that judgment, reveal hidden trade-offs, and establish the conditions under which additional capital will be released, withheld, or withdrawn.

Scorecard One: Economic Value and Capital Efficiency

MAXIMUM 40 POINTS

CRITERION	WT.	BOARD-LEVEL QUESTION	SCORING STANDARD
Revenue creation	10	Will the investment create measurable revenue, pricing power, retention, or addressable-market expansion?	1 — an unquantified opportunity. 3 — an identified customer problem and plausible commercial case. 5 — contracted demand, demonstrated willingness to pay, or strong evidence from scaled deployment.
Margin and productivity	8	Will the investment lower unit costs, increase output, shorten cycles, improve utilization, or strengthen service quality?	1 — theoretical time savings. 3 — validated operational improvement. 5 — evidence that improvement will translate into earnings, capacity, or cash flow.
Return on invested capital	8	Is the expected return attractive relative to the company's cost of capital and competing uses of funds?	1 — returns below the hurdle rate or no credible estimate. 3 — returns near the hurdle rate. 5 — strong risk-adjusted returns supported by realistic assumptions.
Time to value	5	How quickly can the company reach deployment and measurable financial impact?	1 — a multiyear program with little interim value. 3 — measurable value within eighteen to twenty-four months. 5 — material value within twelve months or a sequence of near-term milestones.
Utilization confidence	5	Is demand sufficient to justify the capacity, contract, acquisition, or infrastructure commitment?	1 — speculative demand. 3 — credible forecasts or validated internal demand. 5 — contracted, recurring, or demonstrably constrained demand.
Capital flexibility	4	Can the investment be staged, resized, transferred, repurposed, or terminated?	1 — a large irreversible commitment. 3 — partial staging or contractual flexibility. 5 — modular deployment with clear decision gates and limited stranded-cost risk.

A proposal should generally receive at least 24 points before advancing beyond a controlled pilot unless it is required for legal compliance, operational continuity, or access to a strategically important market.

An investment should not receive a high productivity score merely because employees report saving time. Management must demonstrate how that time becomes lower expense, additional output, accelerated growth, improved customer outcomes, or increased organizational capacity.

Scorecard Two: Sovereignty, Resilience, and Geopolitical Exposure

MAXIMUM 30 POINTS

CRITERION	WT.	BOARD-LEVEL QUESTION	SCORING STANDARD
Data sovereignty and rights	7	Does the company retain sufficient control over sensitive data, training inputs, generated outputs, and derivative intellectual property?	1 — unclear rights or extensive vendor use. 3 — adequate contractual protections. 5 — strong legal, technical, and operating control, including portability, deletion, and audit rights.
Supplier concentration	6	Would the failure, restriction, or repricing of one provider materially disrupt the business?	1 — an irreplaceable dependency. 3 — contractual alternatives that have not been operationally tested. 5 — genuine portability, multiple qualified suppliers, or internally controlled continuity.
Geographic and regulatory portability	5	Can the system operate across the company's most important jurisdictions?	1 — an architecture that is prohibited or commercially impractical in major markets. 3 — material regional adaptation. 5 — a modular architecture capable of meeting different localization, privacy, security, and governance requirements.
Infrastructure access	4	Does the company have reliable access to the power, chips, compute, cloud capacity, connectivity, and technical support it requires?	1 — uncertain or geopolitically exposed access. 3 — adequate capacity with some concentration. 5 — secured capacity, credible alternatives, and protection against material interruption.
Cybersecurity and operational resilience	4	Can the company identify, contain, and recover from model, data, supplier, or infrastructure failures?	1 — untested controls and unclear accountability. 3 — documented controls and incident procedures. 5 — tested recovery, continuous monitoring, independent assurance, and executive ownership.
Exit feasibility	4	Could the company replace the provider or architecture without unacceptable cost, delay, data loss, or disruption?	1 — prohibitive switching costs. 3 — a documented but untested exit plan. 5 — tested migration, documented interfaces, retained expertise, and enforceable transition support.

An investment supporting a critical business process should generally receive at least 18 points.

> **VETO CONDITIONS**
>
> Certain weaknesses should operate as veto conditions rather than merely reducing the score: unclear ownership of strategically important data; an inability to comply with applicable law; dependence on a supplier that may be legally unavailable in a core market; no incident-response process for a critical system; or contractual terms that make reasonable migration impossible.

Scorecard Three: Strategic Advantage and Valuation Premium

MAXIMUM 30 POINTS

CRITERION	WT.	BOARD-LEVEL QUESTION	SCORING STANDARD
Differentiation and defensibility	7	Does the investment create an advantage based on proprietary data, workflows, distribution, infrastructure, expertise, or customer integration?	1 — a capability readily available to competitors. 3 — temporary differentiation. 5 — an advantage that strengthens through data, use, integration, or scale.
Strategic control point	5	Does the investment strengthen a layer of the stack that materially determines the company's economics or market position?	1 — ownership of a nonessential capability. 3 — influence over an important dependency. 5 — control central to pricing, customer access, cost, resilience, or innovation.
Sovereign and regulated-market access	4	Will the investment enable the company to serve governments, regulated industries, or markets with localization requirements?	1 — no incremental access. 3 — a credible route into new customers or jurisdictions. 5 — a reusable capability unlocking material demand across several markets.
Option value	4	Does the investment preserve access to a future technology, market, capability, or partnership without requiring a premature commitment?	1 — limited future flexibility. 3 — meaningful learning or access rights. 5 — a low-cost option on a potentially significant position.
Management credibility	5	Can management explain the economics, milestones, risks, financial effects, and conditions for further funding?	1 — promotional language and weak measurement. 3 — defined milestones and partial financial linkage. 5 — transparent economics, accountable ownership, credible metrics, and clear stop-or-scale decisions.
Earnings durability	5	Will the investment improve the predictability, resilience, or quality of future cash flows?	1 — increased volatility or dependence. 3 — modest risk reduction. 5 — material improvement in retention, continuity, regulatory access, or protection from disruption.

Investments scoring below 18 points are unlikely to support a durable valuation premium even when they provide an attractive tactical return.

Such projects may still be approved as productivity measures, but they should not be presented to investors as transformational strategic assets.

Composite Investment Score

DIMENSION	MAXIMUM SCORE
Economic value and capital efficiency	**40**
Sovereignty, resilience, and geopolitical exposure	**30**
Strategic advantage and valuation premium	**30**
Total	**100**

80–100	**Priority investment.** May be suitable for scaled deployment, subject to ordinary financial and risk approvals.
65–79	**Conditional investment.** Management may advance the proposal in stages, but additional funding should depend on explicit commercial, technical, regulatory, or operating milestones.
50–64	**Option-stage investment.** The company may fund a controlled pilot, minority investment, capacity reservation, or learning program, but should avoid large irreversible commitments.
Below 50	**Reject, redesign, or defer.** Exceptions should be limited to legally required investments, essential risk remediation, or circumstances in which the cost of inaction demonstrably exceeds the expected loss.

A high aggregate score must not override a critical failure within one dimension. A project with excellent revenue potential but unacceptable data rights or no viable path to compliance should not proceed merely because its commercial score is strong.

The score is intended to support judgment, not average away existential risks.

The AI Infrastructure Scorecard

Physical and cloud infrastructure investments require an additional assessment because they involve long-duration commitments, material capital intensity, and significant obsolescence risk.

CRITERION	WT.	EVIDENCE REQUIRED
Contracted or highly predictable demand	15	Customer commitments, internal workload forecasts, utilization history, and downside cases
Power and connectivity security	10	Power contracts, grid access, network redundancy, generation mix, and expansion rights
Expected utilization	15	Base, upside, and downside utilization by year, including ramp assumptions
Unit economics	15	Revenue or avoided cost per unit of compute, energy cost, maintenance, depreciation, and margin
Obsolescence protection	10	Upgrade pathways, modularity, equipment life, residual value, and redeployment options
Geographic and regulatory suitability	10	Data residency, export controls, permits, political risk, and customer eligibility
Financing and balance-sheet effect	10	Cost of capital, lease or ownership structure, free-cash-flow effect, covenants, and ratings implications
Strategic control value	10	Capacity assurance, customer access, bargaining power, intellectual property, or differentiated service
Exit or repurposing value	5	Alternative users, equipment resale, contract transfer, and site repurposing

Infrastructure proposals should generally receive at least 70 points before full-scale approval. Every proposal should include a downside case covering slower AI adoption, lower utilization, falling compute prices, higher energy costs, regulatory restrictions, and the arrival of materially more efficient technology.

Management should compare ownership with at least three alternatives: long-term reserved capacity, joint investment, and fully variable external capacity. The preferred structure should produce the strongest risk-adjusted economics while preserving sufficient control over critical workloads.

The AI Partnership and Vendor Scorecard

Every strategically important AI supplier should be assessed as a potential long-term operating dependency rather than as an ordinary software vendor.

CRITERION	WT.	EVIDENCE REQUIRED
Technical capability and roadmap	15	Current performance, reliability, roadmap, scalability, and compatibility
Financial durability	10	Capitalization, profitability, funding dependency, customer concentration, and long-term support capacity
Jurisdiction and geopolitical exposure	15	Ownership, headquarters, sanctions exposure, export controls, supply-chain dependencies, and government relationships
Data and intellectual-property protections	15	Training rights, data-use restrictions, output ownership, confidentiality, indemnification, and deletion rights
Portability and interoperability	15	Open interfaces, standard formats, migration tools, model choice, and compatibility with alternative infrastructure
Security and assurance	10	Certifications, audit rights, incident history, subcontractors, and vulnerability management
Commercial economics	10	Current and projected pricing, minimum purchases, volume commitments, and repricing rights
Transition support and exit rights	10	Notice, migration assistance, continuity obligations, escrow, data return, and termination rights

Providers supporting critical operations should generally receive at least 75 points and undergo periodic reassessment.

A high score does not eliminate the need for alternatives. The strongest provider can still create concentration risk when the company lacks the technical and contractual ability to operate without it.

The Sovereign-Market Entry Scorecard

Investments in local cloud regions, data environments, regulatory certifications, joint ventures, and market-specific models should be evaluated as market-access decisions.

CRITERION	WT.	EVIDENCE REQUIRED
Addressable revenue and growth	20	Market size, customer pipeline, pricing, share assumptions, and growth outlook
Sovereignty requirement	10	Legal mandates, procurement preferences, customer requirements, and likely policy durability
Reusability	15	Ability to serve multiple customers, sectors, or neighboring markets through the same investment
Local partner quality	10	Governance, ownership, technical capability, political exposure, reputation, and alignment
Capital and operating requirements	10	Investment, staffing, compliance cost, infrastructure, and ongoing support
Profit repatriation and ownership rights	10	Foreign-ownership rules, taxes, currency exposure, cash transfer, and intellectual-property protection
Regulatory and geopolitical stability	10	Policy continuity, international relationships, sanctions risk, and rule of law
Strategic importance	10	Relevance to major customers, regional influence, supply chains, talent, and expansion
Exit feasibility	5	Contract termination, asset transfer, data migration, and stranded-cost exposure

A proposal should not be justified solely by the size of a national AI market. Management must demonstrate which customers the company can serve, why localization is necessary, how the investment will be reused, and what margin remains after local infrastructure and compliance costs.

Sovereign-market investments are most attractive when one capability can unlock several sectors or adjacent markets. They are least attractive when the company must construct a bespoke environment for a small and politically volatile revenue pool.

The Quarterly Capital Review

Approved AI investments should return to the capital committee at predefined intervals. The review should compare actual performance with the assumptions contained in the original scorecard.

Management should report capital committed, capital deployed, operating expense, utilization, revenue, margin effects, productivity realized, implementation costs, customer adoption, supplier concentration, regulatory developments, incidents, and changes in exit costs. Projects should then be classified as scale, continue, redesign, contain, or exit.

SCALE	Requires evidence that demand, economics, security, and operating performance are at least consistent with the original case.
CONTINUE	Used when learning remains valuable but the evidence does not support expansion.
REDESIGN	Applies when the strategic objective remains valid but the architecture, supplier, product, or commercial model is underperforming.
CONTAIN	Prevents new commitments while preserving essential operations or contractual value.
EXIT	Made when the core assumptions have failed or a superior alternative has emerged.

SECTION 12

Geographic Architecture Has Become a Capital-Allocation Decision

Global companies have historically sought common technology architectures to reduce cost and complexity. AI sovereignty may make complete standardization impossible. The strategic challenge is determining where a common global core remains feasible and where regional or national variation is required.

Companies that create an entirely separate stack for every jurisdiction will suffer from duplication, inconsistent controls, and weak returns. Companies that insist upon one universal architecture may lose access to markets with different localization, privacy, cybersecurity, or model-governance requirements. The solution is a modular global platform capable of accepting sovereign variations.

Sensitive data may remain local while aggregated insights move globally. Models may operate in different clouds while applications share common interfaces. High-risk decisions may require national governance even when lower-risk productivity tools remain globally standardized. A company may use a common model library while maintaining country-specific policies, training data, access controls, and audit records.

This architecture is more expensive to design initially. It can nevertheless create leverage because sovereignty capabilities can be reused across jurisdictions and customers. Compliance becomes a product capability rather than an accumulation of exceptions.

Companies should be selective about which markets justify this investment. Not every jurisdiction warrants a customized platform. Decisions should reflect market size, growth, strategic importance, customer concentration, regulatory direction, political durability, and the cost of maintaining the local architecture. Geographic complexity should be funded where it creates defensible access, not where it merely accommodates marginal revenue.

SECTION 13

Partnerships Should Be Evaluated as Geopolitical Commitments

AI partnerships are often presented as technology selections. In practice, they can become long-term strategic commitments. A model provider influences product capabilities and economics. A cloud provider influences architecture, security, deployment geography, and operating cost. A semiconductor provider influences capacity and performance. An implementation partner influences internal knowledge and the company's future ability to operate independently. A government-backed infrastructure partner may influence both market access and political relationships.

Companies should evaluate strategically important providers using criteria normally applied to capital projects, acquisitions, and joint ventures. Management must understand ownership, jurisdiction, financial strength, regulatory exposure, supply-chain dependencies, security practices, intellectual-property rights, subcontractors, and the consequences of political change.

Concentration should be measured at the level of critical business processes rather than total vendor spending. A supplier representing a modest share of the technology budget can still constitute an existential dependency if it controls a system required for revenue generation, customer access, or regulatory compliance.

Multicloud or multimodel claims offer little protection when applications, data, and personnel remain deeply tied to one environment. Genuine portability requires documented interfaces, internal capability, contractual rights, and tested migration processes.

> *Redundancy that has never been tested is not resilience.*

SECTION 14

Mergers and Acquisitions Must Be Reconsidered Through the AI Stack

AI will alter corporate-development priorities. Many companies will be tempted to acquire model developers or highly visible AI applications. Some may be strategically attractive, but public companies should resist paying frontier valuations for capabilities that may become commoditized.

The more durable acquisition targets may be located elsewhere in the stack. Proprietary datasets, vertical software, specialist engineering talent, cyber capabilities, power access, data-center locations, workflow platforms, local licenses, regulatory certifications, and trusted distribution channels may retain strategic value even as models improve.

The central M&A question should be whether an acquisition strengthens a control point that management has already identified as important. An acquisition that adds AI branding without changing growth, margins, market access, resilience, data advantage, or customer economics is unlikely to produce durable value.

Minority investments and corporate venture positions can be particularly useful where the technology remains uncertain. They can provide information, commercial access, partnership rights, and strategic optionality without requiring the company to select a permanent architecture prematurely. The price of optionality should be compared with the cost of being forced to act later from a position of dependence.

SECTION 15

Governance Must Connect Technology, Geopolitics, and Finance

AI governance is too important to be confined to a technical committee and too technically complex to be absorbed into a generic enterprise-risk process. The board does not need to select models or design infrastructure. It must ensure that management has connected AI decisions to corporate strategy, capital allocation, risk appetite, operating resilience, and disclosure.

The board should receive an integrated view of where AI is expected to create revenue, margin, productivity, or strategic advantage; which data and infrastructure are critical; where the company is dependent on individual providers; how geopolitical scenarios could affect deployment; and how capital will be increased, withheld, or withdrawn as evidence develops.

Cybersecurity governance offers a useful precedent. SEC rules require public companies to describe their processes for assessing and managing material cybersecurity risks, management's role in those processes, and the board's oversight. As AI becomes integrated into critical operations, the distinction between AI governance and cybersecurity governance will narrow. AI systems can create new attack surfaces, expose sensitive information, automate consequential decisions, and amplify the effects of compromised suppliers or flawed data.

The NIST AI Risk Management Framework and its generative-AI profile provide voluntary structures for identifying, assessing, governing, and managing such risks. NIST is also developing more specific guidance concerning AI in critical infrastructure. The value of such frameworks is not formal compliance. It is the creation of repeatable disciplines across procurement, design, deployment, monitoring, and incident response.

Governance must also preserve speed. A system that requires every experiment to pass through a central committee will either inhibit learning or drive employees toward unapproved tools. Companies need clear risk tiers, standard environments, approved providers, rapid pathways for low-risk experimentation, and stronger review for systems affecting customers, safety, employment, regulated decisions, or critical infrastructure.

> *Good governance is not the opposite of innovation. It is what allows innovation to scale.*

SECTION 16

Investor Communication Is a Key Driver of the Valuation Outcome

As references to AI become nearly universal in earnings calls, annual reports, and investor presentations, their informational value declines.

The SEC has warned that companies making AI-related claims need a reasonable basis for those claims and should provide company-specific rather than boilerplate discussions of material risks. It has also cautioned against "AI washing": making inaccurate or misleading representations about how a company uses artificial intelligence or what its systems can do. The investor-relations challenge is not to demonstrate enthusiasm. It is to demonstrate command.

Companies seeking a valuation premium should explain how AI affects the economic model. Investors need to understand whether AI is increasing revenue, improving margins, protecting market share, reducing capital intensity, creating regulatory access, or extending the duration of competitive advantage. They need to know whether the benefit is recurring, whether competitors can reproduce it, and whether the company must continue increasing expenditure simply to maintain its position.

Management should distinguish among AI expenditure supporting the existing business, transforming the cost structure, creating new revenue, enabling compliance or market access, and preserving a future option. Combining all five into a single investment figure obscures the economics and invites skepticism.

Credible disclosure should connect capital deployed with operating milestones and financial outcomes. It should explain how management tests productivity, how implementation costs are treated, where capacity is constrained, and when expenditure is expected to affect depreciation, operating expense, free cash flow, and return on invested capital.

Companies should disclose material dependencies with sufficient specificity for investors to assess resilience without revealing sensitive technical information. A credible discussion of concentration, data rights, regulation, governance, and uncertainty may initially appear less promotional. It can nevertheless lower the risk premium applied to the company by reducing ambiguity and distinguishing management from less disciplined competitors.

> *A premium multiple is built partly through growth and partly through the reduction of doubt.*

SECTION 17

The Strategic-Planning Agenda

The upcoming planning process should integrate AI into corporate strategy rather than attach it after business-unit plans have been completed. Every business unit should identify how AI may change customer value, industry structure, sources of differentiation, cost curves, and required capabilities. The enterprise should then determine which opportunities are local experiments and which require common data, infrastructure, governance, talent, or capital.

FINANCIAL	Separate AI expenditure by strategic purpose and economic maturity. Identify the anticipated effect on revenue, operating expense, capital expenditure, depreciation, working capital, and free cash flow. Assumptions concerning model costs, compute pricing, utilization, workforce changes, implementation timelines, and vendor economics should be explicit and capable of being stress-tested.
GEOPOLITICAL	Map where critical data, models, suppliers, infrastructure, and talent are located. Identify the capabilities that could be affected by export controls, sanctions, localization requirements, political conflict, or a deterioration in relations between major powers — and the markets in which sovereign capabilities could create incremental demand or competitive advantage.
TECHNOLOGY	Define the company's intended position across the AI stack: which layers it will control, where it will partner, how portability will be preserved, and which workloads require local, dedicated, or redundant infrastructure.
GOVERNANCE	Define decision rights, risk tiers, supplier approval, model testing, monitoring, incident response, and board oversight.
DISCLOSURE	Ensure that public statements concerning AI are consistent with actual capability, investment, governance, and financial evidence.

These should not be five separate exercises. Together they form the company's AI capital strategy.

SECTION 18

The Broader Political Economy

The need for such a strategy is intensified by the concentration of global AI capacity.

UN Trade and Development has reported that AI research, patents, infrastructure, and private investment remain heavily concentrated, with the United States and China accounting for a substantial share of publications and patents and most major corporate AI research investment occurring outside developing countries. The World Bank frames national AI capacity through four foundations: connectivity, compute, context, and competency. These include reliable energy and digital infrastructure, access to computing resources, locally relevant data, and the skills required to adapt and govern AI systems.

The same framework applies, in modified form, to companies. An enterprise cannot create a defensible AI position through software procurement alone. It needs reliable infrastructure, sufficient computing access, differentiated context in the form of data and workflows, and the institutional competence to deploy and govern the technology.

Companies that possess these foundations can choose among models and providers. Those that lack them may own nominally sovereign systems while remaining dependent on external institutions for every consequential decision.

SECTION 19 · SUMMARY

Where to retain sovereignty, where to purchase access, and where to preserve the right to change course.

Artificial intelligence is often described as a technological revolution. For public companies, it is becoming something broader: a reordering of infrastructure, sovereignty, market access, capital intensity, and corporate power.

The central strategic question is no longer whether companies should adopt AI. They will. The question is what forms of dependence they will accept, which capabilities they will control, and whether their investments will improve the durability and quality of future cash flows. Companies that treat AI as another software-procurement cycle risk surrendering control over critical data, economics, and market access. Companies that treat every component of the stack as strategically essential risk committing vast amounts of capital to assets that become obsolete or remain underutilized. The winners will occupy the disciplined space between these extremes.

They will control the limited number of assets that genuinely differentiate their businesses. They will preserve competition and portability across the remainder of the stack. They will construct products capable of operating within different sovereign environments. They will stage uncertain investments, match infrastructure commitments to measurable demand, and communicate the financial consequences with unusual clarity.

> *The defining capital-allocation decision of the AI age will not be how much to spend. It will be where to retain sovereignty, where to purchase access, and where to preserve the right to change course. That is the foundation of the Sovereign AI Premium.*

BREAKWATER

ABOUT BREAKWATER CAPITAL MARKETS:

WE HAVE A GLOBAL MANDATE.

We bring distinctive judgment and deep market insight to the decisions that shape valuation, command investor confidence, and position companies for enduring leadership.

Our proprietary Valuation Operating System (VOS) transforms how companies build, defend, and accelerate value, integrating market intelligence, compelling equity narratives, and precision investor engagement into a single, powerful engine. The result is sharper perception, stronger confidence, and resilient valuation in an increasingly uncertain and complex global environment.

Trusted by boards, management teams, and investors worldwide, we help leaders anticipate shifts, seize opportunity, create advantage, and excel on the most demanding global stage driving sustained value creation.

Contact Mark Hayes for more information: mark@breakwaterstrategy.com