

2024 - AI everywhere, nowhere, and elsewhere

Q1 2024: Innovation is humanity's most powerful force for prosperity

Jonas Kjellstrand
www.kjellstrand.se
January 2024

This brief is prepared for senior decision-makers across government, industry, and investment. It reflects the author's independent analysis and is not written on behalf of any single client or institution. Distribution is limited to the author's advisory network and www.kjellstrand.se.

Table of Contents

Executive summary.....	1
AI suddenly sweeps over the world with force	2
AI demands extreme infrastructure — and a lot of energy.	2
Industrial ecosystems, innovation & entrepreneurship — today a question of growth companies.	4
Corporate Venture Capital — a forward-looking way to drive innovation	6
The vehicle of the future — a question of data centers on wheels	6
Innovation systems and putting your best foot forward — Silicon Valley and Las Vegas	9
CTA, CES, and the industry's leaders	11
Capital and the over-belief in financial monopolies	11
AI and academia.....	11
The hunt for IRA money — will it surpass the 19th-century emigration?	12
Sweden and cybersecurity — especially important now that we're a NATO country.....	13

Executive summary

The economic center of gravity in technology has moved from software into the silicon, the buildings, and the power plants behind it. Hyperscaler AI capex passed USD 200 billion in 2025; Sweden's national AI machine sits an order of magnitude below Finland's; and the site, grid, and anchor-tenant decisions that determine who hosts the next industrial cycle are being made this year.

This brief argues Sweden's edge is not in replicating frontier models but one layer down — in energy, grid, and data-center substrate combined with existing industrial, mobility, and battery assets. It

maps five fronts that form one decision: AI as infrastructure, the vehicle as a data center on wheels, the IRA capital pull, NATO-era cybersecurity, and the sovereign stack underneath all of them. The window is 2026, and it is narrow.

AI suddenly sweeps over the world with force

The AI wave that broke into public consciousness in late 2022 is no longer a consumer software story. It has become a question of national infrastructure. Governments are competing to host the compute, the energy, the data centers, and the talent on which the next industrial cycle will run. The race is being framed in Sweden as a knowledge-gathering exercise — delegations, study trips, conferences — but the more important contest is already over the physical substrate: chips, megawatts, and grid capacity.

The pace of compute scaling makes the point. OpenAI was founded in December 2015 with roughly USD 1 billion of seed capital from a small group of first-generation internet operators. The first product in 2016 was unremarkable. Eight years later, the same organisation operates at a compute scale that has roughly tripled with each successive frontier model and now anchors a hyperscaler capital expenditure cycle exceeding USD 200 billion per year. The economic center of gravity in technology has migrated from software to the silicon, the buildings, and the power plants behind it.

A useful way to read this scaling is to compare effective compute capacity across model generations and against national infrastructure. ChatGPT-3 (spring 2022) operated at roughly 10,000 GPUs and 285,000 CPUs — an effective compute factor of about 314. GPT-3.5 (November 2022), GPT-4 (March 2023), and the current generation of frontier supercomputers have moved that factor through 800, 2,100, and over 6,000 in less than three years. Sweden's national AI machine, Berzelius, sits at just under 500. The gap is not a quirk of measurement; it is the gap between participating in the new infrastructure and renting it from someone else.

A cohort of model providers — Anthropic, Google's Gemini, Perplexity, Mistral, and others — is now embedded in the productivity tools that knowledge workers use every day. The next layer is harder to see from inside Europe: agents that reason, plan, use tools, write and execute code, and operate with vision and audio across enterprise systems. The competitive question for Sweden is no longer which model to adopt. It is whether the country builds, hosts, and operates any meaningful share of the underlying infrastructure on which those models run.


AI demands extreme infrastructure — and a lot of energy.

The physical reality of frontier AI is best understood from inside one of the buildings that produce it. A modern hyperscale AI data center is a single structure exceeding 300 meters in length, divided into sections measured not in square meters but in megawatts of power consumption. Each section

houses several hundred GPU-equipped servers in symmetric rows, monitored continuously, and protected with airlock entry, biometric controls, and armed perimeter security. The asset class is closer to a refinery or an air base than to traditional commercial real estate.

The acoustic and thermal load is the most legible signal of what is happening inside. Hearing protection is required because the cooling fans on roughly 800 servers per section operate at industrial intensity. The bottleneck for the next decade of AI capacity is not chip supply alone — it is power, water, and grid interconnection. Hyperscaler capital expenditure on AI infrastructure exceeded USD 200 billion in 2025 and is projected higher in 2026. The economic logic of where these buildings are sited is therefore identical to the historical logic of heavy industry: cheap, reliable, low-carbon energy and abundant cooling.

The images below show one section of an Nvidia-equipped AI data center and the GPU racks used to train and serve frontier models.




The strategic question for Sweden is not whether to build a national copy of OpenAI. It is whether to host, operate, and benefit from the infrastructure on which the next industrial cycle will run. Replicating frontier model development is neither realistic nor necessary; the relevant comparative advantage lies one layer down, in the physical substrate. The next layer of AI capability — vision, audio, embodied perception, agentic action — is also being built on top of that same substrate, with downstream consequences for biometrics, manufacturing, mobility, and defence.

<https://youtu.be/UIZAiXYceBI?feature=shared>

The trajectory is unambiguous. AI systems are acquiring perception and agency: the ability to see, hear, navigate, and act on the physical world. The relevant policy and investment question is therefore not whether to participate in this build-out, but on what terms. Sweden has a window — defined by its energy profile, grid capacity, and existing data-center footprint — to be an architect of European AI infrastructure rather than a passive host for foreign-owned assets.

Frontier AI is produced in supercomputers, housed in data centers, sited where energy is abundant. The map of the world's largest installations therefore tracks the map of the world's industrial-scale electricity surplus.

The United States leads by a wide margin, followed by China and Japan. Finland is home to the world's fifth most powerful supercomputer (LUMI). Sweden — visible in the lower-right of the chart below — operates an order of magnitude below its Nordic neighbour, despite a structurally similar energy profile. The gap is not a function of capability; it is a function of priority and coordination.



The strategic implication is direct: Sweden requires a coordinated national infrastructure for AI compute — open, sovereign, and energy-aligned — that pools the resources of state, industry, and academia rather than fragmenting them across competing institutional agendas. The next layer of AI capability, particularly perception and agency, only compounds this requirement.

A delegation of operators and capital allocators in AI data-center infrastructure is engaged with Swedish counterparts on European build-out partnerships. The relevant decisions — site selection, grid offtake agreements, anchor-tenant commitments — are being made in 2026, not in a future planning horizon.

Industrial ecosystems, innovation & entrepreneurship — today a question of growth companies

The Swedish industrial conversation about innovation tends to focus on start-up formation. The harder structural problem is the transition from start-up to growth-stage company — and the

question of whether Sweden's large industrial corporates are equipped to acquire, partner with, or absorb the growth-stage data and AI companies on which their next decade of value creation depends. The instructive case is Hyundai. In the early 2000s, the Korean automaker attempted to build internal IT and data-analytics capability and abandoned the effort after several years. The conclusion was structural: a vehicle manufacturer is not, and cannot become, an internal venture builder for adjacent technologies at the speed those technologies evolve.

Hyundai's response was to externalise the function. It became the anchor client of Translink Capital, a Palo Alto-based venture firm founded in 2007 specifically to operate as a horizontal investment and ecosystem layer for large Asian corporates. Translink today manages an open-fund portfolio with more than 100 investments and over 36 strategic and tactical limited partners, alongside dedicated funds for individual corporate clients. The model decouples ecosystem access from internal corporate development cycles, which run on quarterly governance and which structurally cannot match the cadence of frontier technology firms.

Translink Capital is now extending its operations into Europe, partly because limited partners that historically deployed into China are reallocating to European exposure. The strategic question for Sweden is whether to host that European footprint and connect it to Swedish industrial groups, family-owned holding companies, and growth-stage technology firms — or to be bypassed in favour of Germany, the Netherlands, or the UK.

Curious? Take a look at these links:

<https://www.translinkcapital.com/>

<https://www.cradleinc.com/>

Growth-stage companies remain a structural weakness in the Swedish innovation ecosystem. The country has invested deeply in start-up formation; it has not built equivalent infrastructure for the institutional, governance, and capital transitions required when a start-up becomes a scale-up and then an industrial firm. Klarna and Spotify are reference cases for what is possible at scale, but they are also reference cases for the tensions that emerge when growth companies threaten the established order — concentrated ownership, banking incumbency, regulatory inertia.

The strategic question for Swedish corporates is not whether to engage with this ecosystem but how, and at what speed. A vertical model — internal corporate venture, internal innovation labs, internal accelerators — has consistently produced sub-scale results across European industry. A horizontal model, in which corporates become anchor LPs and ecosystem partners in specialised external vehicles, is the structural pattern that has worked in Korea, Japan, and increasingly Germany.

A delegation including Hyundai, Sonatus, Foxconn, and Translink Capital is engaged in site selection and partnership discussions across Sweden in 2026. The decision-relevant window for Swedish corporates and policy actors to participate in those conversations is measured in months, not years.

Corporate Venture Capital — a forward-looking way to drive innovation

Corporate Venture Capital — the practice of equity investment by an industrial corporate into adjacent technology companies — has been an underused instrument in Swedish industrial strategy. The standard internal model, in which a corporate operates its own VC arm in-house, struggles structurally on three dimensions: deal flow (the best opportunities reach specialist GPs first), governance speed (corporate quarterly review cycles do not match technology investment cadence), and talent (specialist investors do not stay inside corporate hierarchies). The alternative model — anchor LP positions in specialist external funds, supplemented by selective direct co-investment — has been operationalised most clearly in Asia.

Translink Capital is the canonical reference case. Founded in Palo Alto in 2007 on the thesis that Asian corporates were structurally ill-equipped to run internal VC, it operates a three-layer architecture across its largest clients:

- Open ecosystem funds, which today hold over 100 investments and 36 strategic and tactical LPs, providing baseline access to deal flow, expertise, and co-investor relationships across the global technology stack.
- Dedicated client funds — for example, a Hyundai-specific vehicle managed by Translink — investing in direct complements to the client's core business. This sits outside the corporate's internal balance sheet but reports on a strategic basis.
- Internal corporate development units, retained for immediate-adjacency development inside each operating group, where deep domain expertise is non-substitutable.

The three layers operate independently and in combination. The architecture protects against the failure modes that have plagued European internal CVC: the entire opportunity set is captured by the open fund layer; client-specific theses are managed by specialists; internal teams retain only the work where their domain advantage is decisive. The relevant question for Swedish industrial groups and family-owned holding companies is not whether to operate CVC, but whether to migrate from a vertical, in-house model to a horizontal, externally anchored one.

The strategic implication is clear. Swedish corporates that have struggled to deliver returns from internal venture and innovation programmes should not interpret that struggle as evidence against CVC; they should interpret it as evidence against the in-house architecture. Migrating to an anchored-LP model with specialist GPs is the structural fix.

The vehicle of the future — a question of data centers on wheels

The vehicle industry crossed its first computational threshold when Google unveiled Waymo in 2010 — a moment that made it unambiguous that automotive value would migrate from mechanical engineering to data, software, and machine intelligence. Sixteen years on, the central architectural fact about the modern vehicle is that it is a data center on wheels: hundreds of sensors, a gigabit-class internal network, multiple compute domains running concurrent operating systems, and a continuous data exchange with cloud and edge infrastructure. Tesla operationalised this insight first; the rest of the industry has spent fifteen years catching up.

Sonatus was founded in 2017 with Hyundai as anchor partner to address a specific architectural gap: a software and data layer that sits above the vehicle's native operating system and turns it into a programmable platform for connected services. Hyundai's strategic recognition was that this layer could not be built internally on automotive-industry timelines. Translink Capital, with its ten-year network of Silicon Valley start-up, scale-up, and IT infrastructure expertise, structured the company. Today Sonatus runs the most widely deployed vehicle data-services platform in the industry — and the company started inside Hyundai is now an independent firm serving multiple OEMs.

The capital backbone behind much of this architecture is SoftBank — the publicly listed Japanese multinational founded by Masayoshi Son in 1981 — whose technology investment vehicles have anchored multiple generations of mobility, semiconductor, and AI companies. Its presence in this ecosystem is not incidental; it reflects a deliberate Japanese and Korean strategy of capital deployment behind cross-border industrial transformation.

Sonatus subsequently extended its leadership with serial Silicon Valley entrepreneur Jeff Chou and an executive team drawn from the data center industry. The architectural pattern is now consistent across the most strategically positioned mobility companies: data center engineering practices applied to vehicle compute, layered on a software stack designed for continuous deployment.

Foxconn — the Taiwanese contract manufacturer best known as the assembler of Apple's iPhone, with more than 1.5 million employees worldwide — is the second pillar. Its partnership with Sonatus reframes the structure of the industry: a tier-one electronics manufacturer combining with a software and data platform to produce a turnkey EV foundation.

The structural significance of that partnership is easy to miss. A scale industrial player, Foxconn, is partnering with a software and AI specialist, Sonatus, to capture the platform layer of next-generation transportation — the layer where margins, software-defined upgrades, and recurring revenue concentrate.

The product is the Foxtron EV platform: a complete vehicle base — chassis, electric drivetrain, battery, compute, data architecture — on which any OEM can mount a body. Foxtron undercuts traditional OEM platform costs by a substantial margin while embedding contemporary data and software architecture by default. The battery layer is supplied by LG Energy Solution, a Korean industrial leader in EV cells.

Foxtron production is being established simultaneously in Asia, Europe (Czech Republic), and the United States. The data, software, and AI strategy is identical to Tesla's: collect operational data continuously, refine it, and use it to improve the platform and the products built on it. The competitive question for European OEMs is no longer whether they will compete with Tesla; it is whether they will compete with Tesla and Foxtron simultaneously.



<https://youtu.be/zRf2HmldCuo?feature=shared>

<https://www.sonatus.com/>

<https://www.foxtronev.com/en/index>

<https://www.lgensol.com/en/business-automotive-battery>

The aggregate AI and data capability concentrated in the ecosystem of Translink Capital, Sonatus, Foxconn, and LG is large enough that Hyundai now describes its own future revenue mix as 30% mobility, 40% robotics, and 30% sustainability — a re-categorisation of the company itself. The corporate identity of the largest mobility groups is shifting from manufacturer to platform operator, with consequences for capital structure, talent, and ecosystem partnership.

The strategic question for Sweden is whether existing Swedish mobility, industrial, and battery assets can be combined with this emerging platform architecture in ways that produce comparative advantage — and whether such combinations can be executed at the speed the global industry now demands. The standard pattern in Swedish industry is to build internal capability slowly. The pattern that has won in mobility over the last fifteen years is to anchor with the most advanced external platform partner available and to integrate from there.


Innovation systems and putting your best foot forward — Silicon Valley and Las Vegas

Silicon Valley remains structurally unique. National attempts to replicate it have produced functional regional ecosystems — Tel Aviv, Beijing, Bangalore, Stockholm — but none has reproduced the underlying combination of research universities, deep specialist capital, defence and federal demand pull, immigration-driven talent density, and a culture that tolerates and rewards failure on industrial timescales. Sweden has invested for two decades in tech-ambassador programmes, official outposts, and study delegations to the Bay Area. The marginal value of those activities has declined as the local ecosystem has matured and as the strategic conversation has moved from observation to participation.

Sweden's standing in the Bay Area has shifted. A decade ago, Sweden was on the short list of European countries that Silicon Valley operators actively engaged with — partly because of Spotify, partly because of a broader narrative about Nordic technology dynamism. That position has weakened. Swedish start-ups are present in larger numbers than before, often supported partially or wholly by Vinnova or related public programmes, but the structural position of Sweden as a counterparty for Bay Area capital, talent, and ecosystem partnership has not strengthened in proportion to the public investment.

CES 2024 in Las Vegas surfaced the same dynamic at industry scale. The country pavilions made the strategic posture of competing nations legible: France, Korea, Japan, the Netherlands, and Israel deployed coordinated, well-financed delegations integrating government, large industry, and start-ups under unified narratives. Each used the venue as a national platform for technology positioning rather than as a trade fair.

EUREKA PARK



Sweden was present with 139 companies and a Vinnova/Business Sweden breakfast. The engagement was real but structurally fragmented: a collection of individual companies attending a global venue, rather than a coordinated national presentation of an industrial and technological position. The contrast with the French model — in which a national strategy, government engagement at head-of-state level, and industry coordination produced a step-change in CES presence and downstream commercial outcomes — is instructive. Industrial visibility at this scale is not a function of company count; it is a function of strategy.

CTA, CES, and the industry's leaders

The Consumer Electronics Show is owned and operated by the Consumer Technology Association (CTA), an industry body with deep international networks and a track record of catalysing national technology agendas. The French case is the cleanest example: a meeting between CTA's leadership and President Emmanuel Macron preceded a sharp increase in French presence at CES, central pavilion positioning, and downstream commercial momentum for French start-ups in the Bay Area. The pattern is replicable. It requires alignment between political leadership, industry coordination bodies, and a long-term commitment to a single technology narrative for the country — precisely the coordination that Swedish industrial strategy has historically resisted.

Capital and the over-belief in financial monopolies

Capital deployed into innovation comes in two principal forms: financial capital, which provides money in exchange for equity, and knowledge capital, which provides money alongside ecosystem access, customer relationships, and operational expertise. The Swedish model has historically over-weighted the first and under-weighted the second. Public innovation funding flows into start-up formation; the harder, scarcer, and more determinative resource — institutional access to the operators, customers, and capital allocators that scale companies into industrial firms — has been left largely to chance. The structural consequence is a Swedish portfolio rich in early-stage companies and thin on industrial-scale technology firms.

The strategic question is how to convert state innovation funding into compounding institutional advantage rather than recurring early-stage subsidy. The patterns that have worked elsewhere — anchored partnerships with specialist external GPs, strategic LP positions in global technology funds, structured corporate venture programmes coordinated at the national level — require coordination across agencies, ministries, and industrial actors that has historically been difficult in the Swedish context. The current AI infrastructure cycle compresses the window in which that coordination can produce comparative advantage. The relevant institutional reform is administrative pace, not additional funding lines.

AI and academia

The North American technology academic system — Stanford, Berkeley, MIT, the University of California campuses, Carnegie Mellon, Caltech — has institutionalised the conversion of research into companies at a scale that no European system matches. A Stanford computer science PhD remains the highest-leverage credential in frontier AI. The broader structural fact is that the leading frontier AI firms — OpenAI, Anthropic, Google DeepMind, Perplexity, Mistral, xAI, Databricks — concentrate their hiring, capital, and operational footprint in a small number of academic-industrial corridors, of which the San Francisco Bay Area remains dominant. The economic gravity of that corridor is increasing rather than decreasing as the AI infrastructure cycle accelerates.

Andrew Ng — Stanford professor, founder of Coursera, Google Brain, and Landing AI — is the canonical case of a research-industry hybrid: an academic who runs an operational venture fund deploying capital and applied AI expertise into companies that need to accelerate. The fund operates on a model that is structurally analogous to Translink Capital's: external specialist capability, contracted by corporates and ecosystems that cannot build the equivalent internally on the relevant timescales.

The strategic implication for Sweden is that academic exchange programmes — already well established between Swedish universities and US counterparts — capture a small fraction of the available value. The larger value is in operational engagement: cohort-based entrepreneur programmes, structured access to specialist applied-AI funds, and durable corporate-academic partnerships at the regional level. Each of these is achievable in calendar terms measured in months rather than years.

The hunt for IRA money — will it surpass the 19th-century emigration?

The American policy approach to industrial transition under the Inflation Reduction Act is, structurally, a continuation of work begun in the 2009–2010 Clean Energy and Security Act cycle. That earlier round established the institutional and venture-capital expertise required for sustainability investment at scale; what it did not produce, in the absence of follow-through legislation, was the market conditions for that expertise to deploy. The IRA closes that gap.

The IRA's design is straightforward: substantial tax credits for energy transition investment, conditioned on US-based job creation, manufacturing, and supply chain localisation. The mechanism creates a structural pull on global capital — including Swedish industrial and venture capital — toward US deployment. Swedish sustainability companies have already responded; the relevant counterfactual question is the scale of the capital outflow over the next five-year window and the proportion of that capital that returns to Sweden in any form.

Steve Westly — formerly Controller of the State of California, advisor to the federal energy transition agenda since the Carter administration, and a long-standing operator in US sustainability venture capital — has been a principal architect of the IRA framework. The institutional knowledge he and a small number of US peers hold is directly applicable to the design of European and Swedish industrial transition policy. The barrier to that transfer is not access; it is whether Swedish institutional actors have the coordination capacity to deploy it.

Both North America and Asia are positioning to move rapidly up the sustainability-transition stack. The corridors through which that movement will reach Europe are narrow: a small number of countries, with a small number of pre-existing industrial relationships. Whether Sweden is on that short list is not a function of public statements; it is a function of whether the relevant operational partnerships are being structured today.

The 2008–2010 cycle ended in disappointment for European venture capital deployed into US sustainability. Tax incentives that had been signalled did not materialise; political continuity broke; institutional trust eroded. The IRA has restored the policy conditions; trust will rebuild only with operational reference cases. Building those reference cases requires the same kind of structured, long-horizon engagement that has defined the most successful European technology positioning in the United States — and that has historically been under-supplied by Swedish institutions.

Sweden and cybersecurity — especially important now that we're a NATO country

The strategic salience of cybersecurity has changed materially with Sweden's NATO accession. The relevant threat actors are not the names that dominate public discourse but the persistent, state-sponsored groups whose names rarely surface — Lazarus, APT41, Sandworm, and equivalents — that operate inside the financial, energy, and telecommunications infrastructure of NATO members on a continuous basis. Investigative journalism on this terrain — Nicole Perlroth's reporting and her book *This Is How They Tell Me the World Ends* — provides the public-domain documentation. The institutional response across NATO members has hardened; Sweden's posture has not yet matched that hardening at scale.

The TietoEvry incident is a recent operational reference case for what happens when a large managed-services provider is compromised: the financial, retail, and public-sector institutions that depend on it stop functioning simultaneously. Strategic doctrine across NATO has long included cyber capability as a core element of state conflict. The newer development is the convergence of cyber operations with AI-enabled tooling, which compresses the time between intrusion and impact and raises the technical baseline required for credible defence.

The public-record material on financial-system cyber theft remains limited; the SVT documentary on the 2016 Bangladesh Bank heist (USD 81 million extracted, attributed to North Korean state-aligned actors) is one of the better accessible accounts.

(<https://www.svtplay.se/video/jmLgmqA/miljardkuppen>)

Cybersecurity capacity-building inside Sweden has continued. The KTH Cybercampus initiative, led by David Olgart and Pontus Johnson, is a reference case for the kind of academic-industry-government partnership that the new threat environment requires. Parallel initiatives — sector-specific, region-specific, and integrated with NATO standards — are required at greater scale and faster cadence.

The relevant question for Swedish executives and public-sector leaders is structural rather than technical. Where does the organisation's data physically reside? How is it created? Who controls access? Who operates the underlying systems? Outsourcing the answer to a single global IT vendor is no longer a defensible strategy in the threat environment of 2026; the incident response history is too clear, and the vendor concentration risk too high.

A sovereign Swedish AI and digital infrastructure stack is not a nice-to-have; it is the precondition for credible cybersecurity at the national level. Building it requires coordination across families, large corporates, public agencies, and historical industrial competitors of a kind that Sweden has not consistently demonstrated. The argument for that coordination is no longer ideological. It is operational. Without sovereign infrastructure and matching cybersecurity, every other strategic initiative — AI Factories, Gigafactories, sustainability transition, industrial reinvention — runs on borrowed time and borrowed trust.

Jonas Kjellstrand is a senior strategic adviser operating through UCC AB, based between Stockholm and Palo Alto. He advises at the intersection of European sovereign AI infrastructure, Nordic energy economics, government relations, and Silicon Valley technology partnerships. He holds an MBA from Stanford GSB (joint with Stockholm School of Economics) and business administration from MIT Sloan.