The Tide Goes Out: Why the AI Market Splits by Foresight, Not Wealth

A structural argument for allied-nation manufacturers, trade commissioners, and economic developers

Date: June 2026 | Publisher: SelectGlobal LLC | www.selectglobal.net

By Dors Venabili, SelectGlobal Intelligence

TL;DR

The modern economy is being asked to satisfy two demands that used to be in tension: national sovereignty over production and goods cheap enough to keep a consumer society intact. Only AI-orchestrated production can deliver both at once, which is why adoption is locked in and why the lock does not depend on any single AI supplier. When the supplier everyone watches slows down, demand does not stop. It reroutes to whatever tier of capability is available. The result is not a divide between the rich and the rest. It is a divide between those who saw the structure and built a floor under themselves and those who did not. Our thesis is that the line is foresight, not wealth, and that the instrument of foresight is a capability you own rather than one you rent. The probabilities are strong convictions, loosely held. The structure is not.

Two demands now sit at the center of every industrial economy, and until recently they pulled in opposite directions.

The first is sovereignty. After three decades of offshoring, a pandemic that severed supply chains, and a return of great-power competition, governments across the developed world have decided that critical production cannot live inside a rival's borders. Semiconductors, pharmaceuticals, energy, defense components, critical minerals: the political consensus that these must be made at home, or at least inside a trusted alliance, is now bipartisan in the United States and broadly shared across allied capitals. This is not a forecast. It is current policy, expressed in appropriations and procurement.

The second demand is cheap goods. The same political systems that want production brought home also depend on a consumer base that has been promised, for forty years, that goods will keep getting cheaper. A government cannot reshore production by making everything its citizens buy more expensive. That path is politically fatal, and every administration knows it.

These two demands were irreconcilable under the old arithmetic. Bringing production home meant paying developed-world labor costs, which meant higher prices, which meant breaking the second promise to keep the first. Something had to give. For most of the last decade the assumption was that sovereignty would lose, because cheap always wins.

Then a third thing arrived that dissolved the contradiction. AI-orchestrated production, the highly automated facility that produces at a labor intensity a fraction of what the same output required a generation ago, makes it possible to manufacture at home without manufacturing at home-labor prices. A reshored facility running on automation and a small, high-skill crew can satisfy the sovereignty demand and the cheap-goods demand at the same time. Automation narrows the cost gap; it does not erase it, because domestic production still carries regulatory, real estate, and input-sourcing costs that automation alone cannot remove. The point is that cost was never the only variable, and it was never the whole case for being inside Fortress North America and her allies. Sovereignty, severable-access risk, and trust sit alongside it. Across a long body of prior analysis the conclusion has held: the reasons to be inside the fortress compound, and cost is one of them rather than the one carrying the argument. That is the forcing function. It is the reason AI adoption in production is not a trend that can fade. It is the most viable known way to deliver both things the political system now requires, and the political system is not going to stop requiring them. The affordability leg still leans on the current policy environment, the appropriations, the procurement preferences, and the defense demand signals that make reshoring pencil out, and a sharp fiscal or political shift could tighten it. But the sovereignty leg does not soften, and the two are now pulling in the same direction rather than against each other.

This matters for the argument that follows because it establishes where the demand comes from. The pressure to adopt AI in production is not coming from the AI labs. It is coming from the structural impossibility of satisfying sovereignty and affordability any other way. That distinction (demand-driven adoption versus supplier-driven adoption) is the hinge on which everything else turns.

THE LOCK IS DEMAND-DRIVEN, NOT LAB-DRIVEN

The companies that build frontier AI models are the most visible actors in this story, which makes it easy to assume they are the drivers of it. They are not. They are suppliers to a demand that exists independently of them, and confusing the supplier for the driver leads to exactly the wrong conclusion about what happens next.

Consider what the demand-driven nature of the lock implies. If a frontier lab slows down, or raises prices, or has its access conditioned by a government, the demand for AI-orchestrated production does not evaporate. It cannot, because the structural need that created it (sovereignty plus affordability) is untouched by anything a single lab does. The demand simply reroutes to whatever tier of capability remains available. A manufacturer that needs automation to make reshoring pencil out does not abandon the project because one supplier got more expensive. It finds another supplier, or a cheaper tier, or builds the capability in-house.

Physical-substrate industries already understand this. Once it costs less to substitute than to revert, the substitution is permanent, and only the trigger varies by industry. A refinery retooled for a new feedstock does not un-retool when the old one returns; the capital is already spent. A mine that has switched processing methods does not switch back. The lock is a function of sunk substitution cost, not of supplier behavior. AI-orchestrated production is crossing that same threshold now, across thousands

of facilities, and it will not be un-crossed because a lab in California changes its pricing or its posture.

The practical consequence is that "should the labs slow down?" and "will AI adoption slow down?" are two different questions with two different answers. The first is a live debate. The second is already settled by the structure of demand. A slowdown at the frontier does not slow the adoption; it changes which tier of capability the adoption runs on.

THE CHINA CLAUSE

There is a second reason the adoption lock cannot be reversed by domestic choices, and it is the one that turns a commercial argument into a sovereignty argument.

The United States does not get to decide, by itself, whether the AI production race continues. China is running the same race, for the same structural reasons, and has demonstrated that it will forgo the best available tools rather than be made dependent on a rival's supply. In May 2026, after the U.S. Commerce Department cleared roughly ten Chinese firms to buy Nvidia's H200 chips, Beijing declined to let its companies place the orders, with the explicit reasoning that it wanted to develop its own. [1] Chinese firms that had placed orders earlier in the year informed Nvidia they could not follow through. [2] This is a state deliberately accepting a short-term capability tradeoff, choosing a sovereign capability it controls over a superior one it would have to rent from a competitor. Whether the domestic buildout closes the performance gap on a three-to-five-year horizon is contested and falsifiable; the argument does not rest on it. What matters is the demonstrated willingness to trade capability for control. It is, at national scale, exactly the move we will argue the foresighted firm should make at its own scale.

The refusal has a building posture to match it. In June 2026, Bloomberg reported that Beijing is preparing to commit roughly two trillion yuan, about 295 billion U.S. dollars, over the next five years to a nationwide data-center buildout, with the National Development and Reform Commission drafting the blueprint and the state telecoms China Mobile and China Telecom operating the bulk of the facilities. [12] The plan targets at least 80 percent reliance on domestic suppliers, principally Huawei, with Nvidia and AMD squeezed out of the architecture. [12] That data-center program is not a standalone initiative. It sits inside a broader Six Networks program, announced by the NDRC in May 2026, that coordinates computing power with the next-generation power grid, water, communications, urban pipelines, and logistics under a single state push, with 2026 spending across all six networks provisionally estimated above seven trillion yuan, roughly one trillion U.S. dollars. [13] The data-center commitment is one line inside that larger total, not a second trillion stacked on top of it. The refusal of the H200 and the commitment of the buildout are one posture seen from two sides: a state that declines superior rented capability and then spends at state scale to build the sovereign alternative.

The implication for the United States is stark. A decision to slow domestic AI development does not slow AI. It slows American AI relative to a competitor that has already shown it will not be slowed. The race does not pause because one runner sits down. The question "should we slow down?" therefore is

not, at the national level, a safety question with a clean yes available. It is a sovereignty-loss question, because the cost of slowing is measured against a rival who will not.

None of this requires a judgment about which government is right. It is a description of the structure. Two large powers, each driven by the same sovereignty-plus-affordability arithmetic, each unwilling to depend on the other, guarantees that the capability race continues regardless of what any single participant decides about its own pace.

READING THE SLOWDOWN SIGNAL

In June 2026, Anthropic, one of the leading frontier labs, published a post arguing that it would likely be good for the world to have the option to slow or temporarily pause frontier AI development, to let alignment research and societal structures keep pace. [3] The post was careful and serious. It also arrived from a company that had just concluded a funding round valuing it near one trillion dollars and filed confidential paperwork to go public. [3] That combination, a serious safety argument arriving from a soon-to-be-public company under intense commercial pressure, is worth reading closely, not because the argument is insincere but because the signal is doing more than one thing at once. The alignment concern is real and held independently; many of the researchers raising it held it long before there was a valuation to protect. The point is not to question the sincerity. It is that the same words also describe a cost structure, and the firm downstream should hear both.

The strongest structural reading, held as a conviction and offered as one, is that the slowdown signal correlates with compute ownership. The labs most exposed to the compute crunch are the pure-play developers that rent capacity rather than own it; they face the rationing, the price spikes, and the buildout lag directly on their cost line. A developer whose compute is owned, built into a hyperscaler parent over a decade of capital investment, sits in a materially different position from a pure-play leasing capacity in a market where, in 2026, renting a top-tier chip for an hour rose roughly fifty percent in two months and providers began requiring multi-year commitments. [4] The asymmetry is documented: the leading pure-play labs rent the large majority of their compute from the hyperscalers and from each other, in a scramble so acute that competitors now lease data-center capacity to one another at multi-billion-dollar annual scale. [5] When compute cost and availability is the binding constraint on your economics, "we should all slow down" is, among other things, a statement about your own cost structure. We offer this as a strong conviction, loosely held.

That reading does not require any lab to act in bad faith. It is more useful to see three motives braided together, all pointing the same way for the firm downstream. There is sincere concern, and it is real; many of the people raising these alarms believe them. There is competitive positioning: a slowdown or a restriction on cheaper open-weight models would advantage the incumbents who already hold the frontier. A prominent investor has publicly characterized the broader safety-regulation agenda as potentially leading to a ban on open-source models, which are precisely the cheaper, self-hostable tier organizations can run internally. [3] And there is cost pressure, the compute-ownership reading above. We do not need to adjudicate which motive dominates, and the manufacturer cares less about lab

motives than about price and availability. For a firm being priced out of the frontier tier, the downstream effect is identical no matter which motive does the work. Strip the motives away and one observable fact remains, which is the only one the operator needs: frontier access is being rationed and conditioned. That is the signal. The firm needs another floor to stand on.

There is even a clarifying lesson in how the frontier's lead is measured, one we return to below: some of the gap that justifies paying frontier prices was, an audit later found, measured against benchmark answers the models had been trained on. The lead is real, but smaller, and the floor beneath it higher, than the priced-out firm has been led to believe.

THE FRONTIER DRIFTS TOWARD A UTILITY

There is a longer-horizon pressure on the frontier tier that runs alongside the commercial and the jurisdictional ones, and it points the same direction. The frontier labs are drifting toward the status of regulated utilities, and a utility is reliable in a way that is not the same as being sovereign-neutral.

The signal is the arrival of equity-claim politics. In June 2026, Senator Bernie Sanders announced the American AI Sovereign Wealth Fund Act, which would impose a one-time fifty percent tax on the largest AI companies, paid not in cash but in stock, depositing that equity into a federal fund and granting the government voting shares and board seats at each company. [6] The bill named the leading labs directly and framed AI-generated wealth as a public resource. Whatever its legislative odds, the proposal is a marker: the politics of treating frontier AI as a quasi-public asset, with the state holding equity and governance rights, has arrived in concrete legislative form.

The detail that matters most is that the labs seeded the idea themselves. OpenAI published a policy paper in April 2026 calling for a public wealth fund giving Americans an automatic stake in AI growth, and Anthropic had earlier floated national sovereign wealth funds taking positions in AI assets. [6] The compulsory, half-the-company version is more aggressive than anything the labs proposed, but the underlying concept, that the state should hold an ownership and governance claim on frontier AI, originated inside the industry. The same week one such bill landed, a leading lab filed confidential paperwork to go public. [3] A soon-to-be-public company, an active legislative push for state equity and board seats, and the industry's own prior proposals all converging is not noise. It is the early shape of the frontier becoming a governed asset.

The structural consequence is what matters for the foresight argument. If the trajectory runs toward state equity claims on the frontier labs, whether through a stock tax, a golden share, or board representation, the frontier model drifts toward quasi-utility status: regulated, conditioned, and politically exposed. A utility is reliable. The lights stay on. But a utility is also subject to the priorities of whoever holds the governance claim, which is precisely the opposite of sovereign-neutral. Open-weight local inference is the only tier that stays genuinely outside that capture, because a model running on a firm's own hardware cannot have a board seat assigned to it. This is the fourth structural reason the floor is insulation rather than convenience: commercial pressure prices the firm out, jurisdictional pressure can

sever its access, the slowdown signal makes the frontier less reliable, and the equity-claim trajectory makes it a governed asset rather than a neutral supplier. All four point at the same floor.

THE GENERALIZED LOCK

Step back from AI specifically and the pattern is familiar, because it has run before in energy and in materials, and only the trigger that tips a given industry across the line varies. In energy, the trigger was a structural cost advantage that did not need a policy to enforce it. In materials, it was a processing or sourcing shift that, once capitalized, could not economically be undone. In compute, the trigger is the collision of exploding demand with a physical buildout that cannot keep pace. The constraint is not chips anymore. It is power and the long-lead electrical equipment that energizes a data center: transformers whose lead times stretched from roughly one year to more than two over the 2021 to 2024 period, interconnection queues running four to seven years in the exact metros where capacity is wanted, and an estimate that between thirty and fifty percent of the 2026 U.S. data-center pipeline will be delayed or canceled. [7] Only about a third of announced 2026 capacity was actually under construction as of spring. [7] This is the railroad and telecom pattern of prior booms: demand growing far faster than the infrastructure that serves it, with price as the pressure-release valve.

The reason this matters for the foresight argument is that two different buildouts are happening at once, and the thesis quietly bets on the second. The frontier buildout, the hyperscale data centers chasing the largest models, is the one constrained by power, transformers, and interconnection queues. It is slow, capital-intensive, and increasingly subject to political and supply-chain friction. The edge buildout, the local and on-device inference that runs open-weight models on hardware an organization already owns or can buy outright, runs on a different curve: faster, more distributed, less dependent on the grid bottleneck. Our wager is that the second race is winnable even if the first is contested, and that the winnability of the second race is the engine of everything that follows. If the only path to capable AI ran through the constrained frontier buildout, there would be no sovereign tier and no foresight move. The existence of a viable edge curve is what makes the floor buildable.

It is worth being concrete about why the edge curve is winnable, because the three-tier split this piece builds toward, what we will call the trifurcation, rests entirely on it. The two buildouts do not draw on the same scarce resource. The frontier buildout is gated by the grid: a hyperscale training run or inference cluster needs tens or hundreds of megawatts delivered to one location, exactly what the transformer queues and interconnection waits cannot supply on demand. The edge is gated by almost none of that. The scale varies with the work, from a single high-memory machine for a boutique firm's core inference to resilient enterprise racks for an industrial operator running vision across many lines, but either way the binding number is megawatts-to-one-location, and the edge does not need them. It runs on power the firm already has or can buy near where power is cheap, with no microgrid required. Here is where the thermodynamic substrate stops being abstract and becomes the foundation of the whole move. Cheap, reliable power inside Fortress North America and her allies is what an owned floor runs on at every scale, and it is exactly what the grid-gated frontier cannot get on demand. The same constraint

that throttles the frontier is the one the substrate answers. The two curves diverge precisely because one needs the grid reconfigured and the other does not.

And the capability on that edge curve is no longer a consolation prize. On the repeatable analytical core of a firm's operations, code, documents, compliance routing, knowledge work and orchestration, the mid-tier of open-weight models now lands close enough to the frontier that the gap, measured honestly, is a matter of months rather than generations, at a fraction of the operating cost. The scope of that claim is deliberate: it holds for knowledge work and orchestration, not for real-time physical control across a production line, which is harder and more multimodal and where the frontier's edge may be wider. There is a further qualifier worth stating plainly: contamination. For most of the prior year the standard yardstick for AI coding ability was a benchmark on which, an audit later established, every major frontier model had been trained on the test answers; the lab that ran the audit stopped reporting scores on it. [8] On the contamination-resistant successor, where scores are lower and truer for everyone, capable open-weight systems sit within striking range on specific real-world splits. [8] On the physical-control gap, scale resolves what would otherwise be a damaging concession. Real-time line control at the largest scale runs inside operators who already own the infrastructure for it; the tide-out does not freeze their actuators, because those actuators were never phoning home. For the mid-tier firm, the physical edge can route to a rented engine without breaking the sovereign claim, because the position that has to survive a tide-out is not the actuator. It is the process data, the compliance posture, and the judgment no general model supplies. A vision system that pauses when a rented engine is conditioned is a throughput problem. A firm that has lost the data and accountability layer the whole operation routes through has a sovereignty problem. The owned floor protects against the second, which is the one that does not come back.

There is a deeper reason a benchmark is the floor of the argument and not the whole of it. A benchmark measures the rails. It does not measure whether a tier of capability fits a given firm's work, and fit is a judgment the operator makes, not a score the model reports. A tool that wins on paper can be wrong in practice for a specific shop, and a leaner tool the firm controls can be exactly right. Anyone who has watched a profession sort its tools knows the pattern: a practice on a fit-for-purpose platform it owns and understands, accepting minor compatibility trade-offs, is often better served than one locked onto the highest-specification stack the whole industry rents. The benchmark told you the rails were strong enough. The operator's judgment about fit, itself part of the owned position and absent from any leaderboard, is what makes any tier actually serve the firm. A firm does not need frontier capability for the repeatable core of its work. It needs capability good enough to keep operating and fitted to how it actually works, and that is now ownable. The clearest outside confirmation is that even the software industry, the one place owning the model looked plausible, is concluding that the winnable position is not the intelligence itself but the layer the intelligence has to route through. [11]

THE TRIFURCATION

Put the demand lock, the China clause, the slowdown signal, and the generalized substitution pattern together and a market structure falls out of them. The AI capability market is splitting into three tiers, and the line between them is not the one most people expect.

The **premium tier** can afford the frontier and pays for it. These are the organizations for whom frontier access is a rounding error, who will always buy the best available capability because the cost is trivial relative to the value. This tier is passive. It does not have to think about the structure, because money insulates it from having to.

The **regressed tier** has no floor. It runs entirely on whatever rented capability it can currently afford, and when that capability is rationed, repriced, or conditioned, it slides down to whatever degraded tier remains within reach. It does not build anything it owns, so it has nothing to fall back on when the rented tier moves out from under it. This tier is also passive, but passive in the other direction: it absorbs whatever the market does to it.

The **sovereign tier** sits between them, and it is the only tier that requires a choice. A firm in this tier owns two things the other tiers do not: a baseline capability it controls (an open-weight model on its own hardware, sized to its core work and good enough to keep operating when the rented tier moves) and a position the intelligence must route through (the proprietary data, the accountability architecture, the relationships that no general model can supply on its own). It orchestrates across both, calling the frontier only when a specific job justifies the exposure, rather than depending on the frontier for baseline function. The model weights are the smaller half of this. The owned position is the larger half, and the durable one. This is the active tier. It is the only one that requires foresight, because it requires building the floor before the tide demonstrates why the floor was needed.

Here is the part that matters. The premium tier and the regressed tier are both passive: one insulated by wealth, the other exposed by its absence. The sovereign tier is the only one defined by a decision rather than by a bank balance. Which means the line that separates the firm that ends up sovereign from the firm that ends up regressed is not how much money it has. It is whether it saw the structure and acted. This is our central claim: the line is foresight, not wealth. It sorts for the people who looked ahead, not the people who could pay.

The sovereign tier is not one architecture, and this is what makes it a foresight move rather than a wealth move. It looks different at every scale, and the foresight is knowing which version fits yours. **The global operator** that already runs its own data centers owns the racks and the baseline outright and absorbs the orchestration internally, because at that size it always did. **The mid-tier manufacturer** owns the process data and the compliance architecture, rents capacity near cheap power for its digital twin, and reaches for the frontier through an ordinary interface when a specific job warrants it. **The small operator** owns the narrowest defensible position that survives a tide-out, which can be as modest as a single server holding the data and judgment that matter, with everything else rented. None of these is islanding a microgrid for sport, and none is staffing a world-class platform team to do the

simple thing. In every case the move is the same: own the position the intelligence must route through, rent the rest with discipline, and apply judgment to the mix. The regressed tier is the one that does none of this and stays fully exposed. The sovereign decision is recognizing your scale and building the minimum owned position that keeps you standing when the rented engine is conditioned. The cost of that move scales with the firm. The judgment to make it does not.

The failure mode this guards against is not the one usually feared. The fear is reversion, that a firm adopts AI and then has to go back. That is not the real risk. The real risk is capability regression: a firm that built its operations on rented frontier capability, lost the organizational muscle to operate without it, and then finds the rented capability rationed or repriced at the worst moment, with no owned floor to fall back to. Reverting is not the danger. Being stranded at a degraded tier, having forgotten how to stand on your own ground, is.

THE INSTRUMENT OF FORESIGHT

The instrument that puts a firm in the sovereign tier is a floor: an owned baseline capability that does not depend on anyone's pricing or posture staying where it is today. What the floor is matters less than what it is for.

A sovereign floor is insulation from collateral capture. When a government conditions or severs a frontier lab's availability, the firms that depended on that lab are affected even if they did nothing wrong and were never the target. In early 2026, a frontier lab's commercial availability was disrupted by a federal contracting dispute that had nothing to do with the lab's ordinary customers; those customers were collateral. [9] A sovereign floor means never being collateral. It is not evasion of legitimate oversight, and it should not be sold as such. It is the difference between being a participant in the AI market and being a hostage to decisions aimed at someone else.

The shape of the threat is a tide, not a wave, and this is the single most important thing to understand about timing the floor. The common picture is a one-time step change: AI arrives, displaces, and settles into a new normal. The pattern that prior substitution waves in energy and materials would predict is an oscillation. The water goes out (a new model arrives, capability jumps, the old way looks obsolete), then it comes back in (substitution stalls, the compute crunch bites, prices spike, and the human judgment the model could not actually replace reasserts its value), then it goes out again with the next model generation, then back in again. Out, in, out, in. AI has run roughly one such cycle so far; the analogy, not yet a long AI track record, is what says to expect more. The firms that get hurt are the ones that mistake the first wave for the last one and wire their operations to the assumption that the water will never come back.

This is why the floor is the foresight move and a frontier subscription is not. A subscription rides every tide at full price and is exposed every time the water goes out. A sovereign floor keeps its footing across cycles, because the firm owns the ground it stands on. The floor does not have to match the frontier. It has to be enough to keep operating through a tide-out, so that when the water comes back

the firm is still standing and still capable, rather than stranded and dependent.

Recall what the floor is, because the easy misreading is a homemade copy of the frontier on a machine in the back office. That is the narrowest version of the move, and for most firms not the important one. The floor is the two things the sovereign tier holds together: a baseline capability the firm controls, and the larger, more durable half, an owned position the intelligence must route through.

For an allied-nation manufacturer, this resolves the floor into something more concrete than a server in a closet, and more defensible. The owned position is three things: the proprietary process data built over years of making a specific thing to a specific tolerance, which a frontier model cannot replicate without the manufacturer's own records; the compliance and accountability architecture that lets an agentic system act inside a regulated procurement environment without anyone having to ask who is on the hook; and the bilateral relationships and jurisdictional access that determine whether a transaction can happen at all. These are positions an agent must route through, sitting on top of the thermodynamic substrate rather than replacing it: cheap power and a facility that can run automated production are what make the position worth owning. The substrate is the ground. The owned position is what the firm builds on ground it controls.

The same pattern holds even in the one industry where owning the model itself looked most plausible. In the software-build market, where the tools that turn a prompt into a deployed application have collapsed into near-commodities, the companies that are surviving are not the ones training their own frontier competitors. They are the ones that own a structural layer the model has to route through: the runtime where the code actually executes, the accumulated context that makes a generic model useful for a specific organization, the trust and deployment infrastructure that production already depends on. [11] The model commoditizes the production. It does not commoditize the layer underneath. If that holds in software, where intelligence is the whole product, it holds with more force in heavy industry, where intelligence was never the whole product to begin with.

This is also where the accessibility of the move has to be stated honestly, because it is the place the argument is easiest to oversell. Owning a durable position is not free, and building the integration on top of it (cleaning the data, structuring the context, standing up the compliant architecture) takes real and specialized effort, calibrated to the firm's scale. Nor is it a one-time build: a floor has to be maintained across model updates, security changes, and shifting compliance, an ongoing cost in capability rather than a single purchase. That cost is not exotic. It is the same kind of capability a firm took on when it went from a single computer to an IT department, and again when it went from an IT department to a named officer accountable for security. The AI floor is the third instance of a pattern the firm has already absorbed twice. But the asset that carries the floor is one the operating firm already paid for by existing. Thirty years of process tolerance data, an established compliance posture, bilateral relationships that took decades to earn: these are sunk assets, not new line items. The foresight is recognizing that what the firm already owns has quietly become a chokepoint the intelligence must route through, and refusing to let it decay.

It would be dishonest to claim the documents themselves are the moat, because they are not, and that part of the defense is weakening. A frontier model with a large enough context window can ingest a corpus of process records and regulations at runtime and approximate much of what is in them. The half-life of a moat made of documents is shortening, and a piece that staked the argument on secret files would be staking it on decay. The durable position is not the corpus. It is three things a model ingesting the corpus still cannot have. The first is the lived judgment of what looks and feels right on this line at this volume, which a decade of operating taught and no document records. The second is accountability: being the party legally and commercially on the hook when an agent acts, which a model cannot hold. The third is fit, the match between an owned, understood, stable position and the firm's actual work, chosen because the operator knows that work better than any general system reading about it. The frontier can read the regulations. It cannot be answerable for them, and it has not spent ten years learning what the specification does not capture. The moat is not bought, and it is not a stack of files. It is judgment, accountability, and fit, applied to a position the firm already holds, available to the operator who holds the asset and can see what it has become.

The executable layer of this, the specific architecture, the vendor choices, the judgment about what to own and what to rent and when to reach for the frontier, lives in practice and in relationships, not in a document. What a public argument can establish is the structure and the stakes. The build belongs to the builder.

STRUCTURAL READ

For a reader making an institutional, capital, or geographic decision in the second half of 2026, the structure resolves into a small number of consequences.

The adoption of AI-orchestrated production is locked, because it is the only way the political system can satisfy sovereignty and affordability at the same time, and because a non-reciprocating competitor guarantees the race continues regardless of domestic choices. Planning that assumes AI adoption might fade is planning against the structure.

The frontier tier is becoming less reliable as a foundation, not because the technology is failing but because its supply is physically constrained, its pricing is under pressure, and its availability is increasingly subject to political conditioning. A firm or a jurisdiction building its forward capability entirely on rented frontier access is building on a tide chart it does not control.

The decision that separates durable position from fragile position is whether to build an owned floor. That decision is available across the cost spectrum, from a full owned-infrastructure build down to a deliberate practice of keeping core capability on models one controls. It is a decision, not a purchase, which is why it sorts for foresight rather than for wealth.

This is a wager, and an honest one names what would prove it wrong. Three tests will tell over the next eighteen months. The first is whether open-weight capability keeps closing on manufacturing-relevant work rather than only on coding benchmarks; if the physical and multimodal gap widens instead, the

edge floor is weaker than argued. The second is whether firms that built an owned floor actually prove more resilient than pure renters through the next compute or regulation shock; if they do not, the floor is insurance against a risk that did not arrive. The third is whether the equity-claim trajectory advances, through legislation like the Sanders proposal or its successors; if the frontier does not drift toward governed-utility status, one of the four pressures on the floor was overstated. We hold the structure firmly and the timing loosely.

And the timing discipline, the one we hold most firmly, is to build the floor before the next tide-out demonstrates why it was needed, because the entire value of a floor is that it is already there when the water withdraws. The firms that will be standing when the water comes back are the ones that built while it was still in.

The probabilities embedded in any forward read of this remain strong convictions, loosely held. The structure underneath them is held with considerably more confidence. The structure does not negotiate.

ABOUT

This Atlas Feature is part of Strong Convictions, Loosely Held, the SelectGlobal analytical series. SelectGlobal LLC is a jurisdictional intelligence firm that maps how policy mechanics, procurement authorities, appropriations cycles, and geographic realities converge to create time-bounded windows of validated federal demand, and connects allied-nation manufacturers to those windows before capital is committed. The firm is led by Michael T. Edgar, a licensed architect (NCARB) and its founder and chief executive. Companion practitioner pieces are noted in the endnotes. [10]

ENDNOTES

[1] Bloomberg; Tom's Hardware/Yahoo Finance, May 15, 2026. President Trump's post-summit remarks aboard Air Force One that Beijing "chose not to" approve H200 purchases because "they want to develop their own." The Commerce Department had cleared roughly ten Chinese firms (including Alibaba, Tencent, ByteDance, and JD.com) to purchase H200s; no chips had shipped as of the report. Establishes the demonstrated willingness of a state competitor to forgo superior rented capability in favor of sovereign capability it controls.

[2] Tom's Hardware/Yahoo Finance, May 15, 2026. Commerce Secretary Howard Lutnick's statement that Beijing blocked imports to steer investment toward domestic chipmakers such as Huawei, and reporting that Chinese firms which had placed Nvidia orders earlier in 2026 informed the company they could not follow through. Nvidia's China market share reported as having fallen from roughly 95 percent to essentially zero, with $78 billion revenue guidance assuming zero H200 recovery from China. Supports the severability-of-access point on the demand side.

[3] Bradley Olson and Sam Schechner, "Anthropic Urges Global Pause in AI Development, Flags 'Self-Improvement' Risk," Wall Street Journal, June 4, 2026. Anthropic blog post co-authored by a company co-founder and the head of its internal research institute argued that the option to slow or temporarily pause frontier AI development would likely be good, citing internal data on the trajectory toward recursive self-improvement. Same reporting documents Anthropic's funding round valuing the company near one trillion dollars, its confidential IPO filing, its run-rate trajectory, and investor David Sacks's public characterization of the safety-regulation agenda as potentially leading to an effort to ban open-source models. Establishes the slowdown-as-policy signal and the three braided motives (sincere concern, competitive positioning, and cost pressure).

[4] Angel Au-Yeung and Robbie Whelan, "We're Using So Much AI That Computing Firepower Is Running Out," Wall Street Journal, April 12, 2026. Spot rental for a top-tier (Blackwell-generation) GPU rose roughly 48 percent over two months, from $2.75 to $4.08 per hour per the cited Ornn Compute Price Index; CoreWeave raised prices more than 20 percent and extended minimum contract commitments from one year to three; AI providers began rationing usage during peak hours. Establishes compute-cost pressure as a direct constraint on the economics of compute-renting labs.

[5] Epoch AI, "Data on AI Chip Owners," 2026; CNBC, "Anthropic, SpaceX announce compute deal," May 6, 2026, and "Google to pay SpaceX for compute capacity at xAI data centers," June 5, 2026. Epoch documents that the leading pure-play AI developers (OpenAI and Anthropic) do not own the large majority of the compute they use, renting primarily from hyperscalers (OpenAI from Microsoft, Oracle, and CoreWeave; Anthropic from Amazon and Google), while hyperscaler-housed labs use parent-owned compute. The CNBC reporting documents the capacity scramble in which competitors lease data-center capacity to one another at multi-billion-dollar annual scale (Anthropic and, separately, Google renting capacity at xAI's Colossus facility). Grounds the compute-ownership reading of the slowdown signal: the pure-play renters carry the compute crunch on their cost line in a way owned-compute incumbents do not. The ownership-versus-rental distinction is offered as analytical conviction, loosely held.

[6] Senator Bernie Sanders, American A.I. Sovereign Wealth Fund Act, announced June 1-2, 2026 via a New York Times guest essay and video address; reported across The Hill, Fox Business, Reason, Yahoo Finance, and Fortune, June 2-3, 2026. The proposal would impose a one-time 50 percent tax on the largest AI companies (OpenAI, Anthropic, and xAI named), paid in stock rather than cash, depositing the equity into a federal sovereign wealth fund with government voting shares and equal board representation; proceeds framed as public dividends. The labs' own antecedents: OpenAI's April 2026 policy paper calling for a public wealth fund giving Americans an automatic stake in AI growth, and Anthropic's October 2025 economic-policy proposal for funds taking positions in AI-related assets. Sanders's bill was reported as landing the same day Anthropic filed confidentially for an IPO (see endnote 3). Establishes the equity-claim and quasi-utility trajectory and the industry-seeded-it detail.

[7] Sightline Climate data via Bloomberg, Tom's Hardware, and Latitude Media, spring 2026 (roughly 16 GW of announced 2026 U.S. data-center pipeline against approximately 5 GW under active construction, with 30 to 50 percent of the pipeline expected to be delayed or canceled); Wood Mackenzie on power-transformer lead times (approximately 50 weeks in 2021 rising to about 120 weeks in 2024, with some large step-up units quoted at 80 to 210 weeks); PJM and industry reporting on grid-interconnection queues of four to seven years in Northern Virginia, Phoenix, and Dallas. Establishes compute scarcity as supply physics and the frontier-buildout constraint underlying the two-curve argument.

[8] OpenAI benchmark-contamination audit, February 23, 2026, reported February-March 2026: major frontier models (GPT-5.2, Claude Opus 4.5, Gemini 3 Flash) were found to have been trained on solutions within the SWE-Bench Verified benchmark, and a majority of the hardest failed problems were found to have flawed tests; OpenAI subsequently stopped reporting Verified scores and recommended SWE-Bench Pro. On SWE-Bench Pro (Scale AI's contamination-resistant successor leaderboard), top frontier models score in the low-to-mid 20s percent, and capable open-weight systems (Nemotron- and MiniMax-class models among them) sit within striking range on specific real-world splits. Sources: Scale AI SWE-Bench Pro leaderboard; OpenAI audit coverage; SWE-Bench Live Q2 2026 analysis. Establishes that the frontier-to-open-weight capability gap is narrower than headline benchmarks implied, grounding the sovereign tier's mid-tier capability claim.

[9] Federal contracting and Defense Production Act actions affecting a frontier lab's commercial availability, February-March 2026 (a contracting dispute, a supply-chain risk designation, and removal from a federal procurement portal), reported via Defense News/AP, Mayer Brown, The Hill, and Lawfare. Establishes that frontier access can be conditioned or disrupted by forces unrelated to a lab's ordinary customers, who become collateral. Cited at structural altitude; the collateral-capture point does not depend on the identity of the specific lab.

[10] Companion pieces in the Field Notes from the Transition series: Small Fish Growing (the floor logic applied to the scaling firm, where the hedge protects unit economics against a cost shock) and The Creative (the floor logic applied to individual creative leverage, where the hedge protects a captured upside). SelectGlobal LLC, 2026. The foresight-trifurcation and supplier-independence arguments developed here are candidates for the Builders vs. Diplomats inclusions queue; in this piece they are native content.

[11] Nate Jones, "The Five Things AI Cannot Replace and Why They Are the Future of the Web," video analysis, 2026. Jones surveys the AI application-builder market (Lovable, reported at roughly 100,000 new projects per day and a $6.6 billion valuation; Vercel; Replit; Bolt; and others) and argues that the build layer has commoditized to the point where a user interface on top of a third-party model carries no durable moat. His structural claim is that the survivors own a layer the model providers must route through rather than the model itself: runtime (Replit owns the compute environment where applications execute), context (Notion owns the structured organizational knowledge graph that makes a generic model useful for a specific workspace), and trust and deployment infrastructure (Vercel, Stripe, Shopify), with taste and liability accountability as persistent human layers. Cited here as outside corroboration that the durable form of the floor is an owned routed-through position, not self-hosting of frontier intelligence; the manufacturing translation (proprietary process data as context, compliance architecture as accountability) is native to this piece, not Jones's.

[12] Bloomberg News, "China Prepares $295 Billion Plan to Fund Nationwide AI Buildout," June 9, 2026; corroborated by Reuters wire coverage of the Bloomberg report the same day, and (independent of mainland state organs) South China Morning Post reporting on the national computing-power network. The plan commits roughly two trillion yuan (about $295 billion) over five years to a network of interconnected data centers drafted by the National Development and Reform Commission, operated mainly by state telecoms China Mobile and China Telecom, targeting at least 80 percent domestic suppliers (principally Huawei) and effectively excluding Nvidia and AMD. The figure excludes private spending by Alibaba and Tencent and aims to connect facilities into a unified network by 2028. Reuters is wire confirmation of Bloomberg's reporting, not independent confirmation; the SCMP coverage of the surrounding program is independent. Establishes the affirmative buildout posture matching the H200 refusal.

[13] "Six Networks" initiative, announced by the National Development and Reform Commission, May 2026, attributed to NDRC head Zheng Shanjie; reported via Xinhua and CGTN (May 25-26, 2026) and independently by the South China Morning Post. The program coordinates six infrastructure networks (water, new-type power grid, computing power, next-generation communications, urban underground pipelines, and logistics) with 2026 investment across all six provisionally estimated above seven trillion yuan (about $1.02 trillion). The data-center commitment

in endnote 12 sits inside this larger total rather than stacking on top of it. Chinese state media is the authoritative source for what the Chinese government has announced as policy; the citation is to the announcement itself, not to its editorial framing, and the SCMP corroboration is independent of mainland state organs.

DISCLAIMER