After AI: where the constraint moves when intelligence leaves the screen

The next boom is not another chatbot layer. The rent moves to power, bodies, labs, devices, factories, and liability rails, because those are the inputs that do not scale at software speed.

Frame

When a system scales, the money moves to the input that cannot scale with it. This board names that input, the date it starts to bite, and the line that would break the call.

Area

what comes next after AI

Horizon

2028 to 2031

Issued

2026-06-17

Method

Wide cast, adversarial gate, public resolution criteria.

Board summary

The cross-cutting read

The shared shift is that AI stops being mainly a model market and becomes a physical deployment market. The obvious layers, GPUs, frontier labs, generic agents, and benchmark wins, are no longer where the edge lives. The edge moves into the inputs that cannot answer a software-speed demand curve: firm power and siting, robot commissioning and safety validation, experimental throughput, always-on edge thermals, biomanufacturing scale-up, and auditable authority for agents that can act. The commercial wedge for Vati is to identify these binding constraints before they change capex, procurement, site selection, portfolio exposure, compliance budgets, or partnership timing.

At a glance

#	Claim	Binding constraint	Case	Call	Resolves
P1	Behind-the-meter firm power, geothermal-capable campuses, and data-center sites with power rights become strategic AI...	Contiguous land, fiber proximity, behind-the-meter firm generation rights, interconnection optionality...	82%	52%	2028-12-31
P2	Robot commissioning software, safety validation, simulation assets, and task libraries capture more durable value than...	Verified task data, sim-to-real validation, workcell commissioning, safety certification, and field support...	78%	46%	2028-12-31
P3	Self-driving labs, assay capacity, robotic testbeds, metadata infrastructure, and provenance systems become the scarce...	Robotic experimental throughput with standardized metadata, reliable assays, calibrated uncertainty, and...	76%	44%	2029-12-31
P4	Wearables, glasses, phones, cars, drones, and robots with local NPUs become the main surface where AI feels useful.	Low-power NPUs, sensor fusion, local memory, privacy-preserving orchestration, and thermal envelopes for...	70%	43%	2028-12-31
P5	Pilot fermentation, downstream processing, QA, feedstock contracts, and industrial bioprocess talent become more...	Pilot and commercial-scale fermentation, downstream processing, strain robustness, feedstock flexibility, QA...	72%	41%	2030-12-31
P6	Agent security, identity, execution controls, audit logs, insurance, and governance middleware become the enterprise...	Identity, permissioning, tool-access control, audit trails, reversible execution, human authorization...	74%	48%	2028-06-30

Case is the strength of the structural thesis. Call is the probability on the exact dated clause.

P1

The AI frontier moves from model access to firm-power siting.

Domain: AI infrastructure / energy

2028-12-31

Structural case	Our call, dated	Resolves
82%	52%	2028-12-31

Data-center electricity demand is rising faster than grid equipment, interconnection, and local political consent can adjust. Rhodium cites LBL projections that US data centers could reach 7 to 12 percent of US electricity demand by 2028, while pv magazine reported four-year waits for power transformers in May 2026. Once GPU access is not the only scarcity, the gating variable becomes where a buyer can get clean, firm, permitted electricity without waiting in the queue.

The boom	Behind-the-meter firm power, geothermal-capable campuses, and data-center sites with power rights become strategic AI assets.
Why it is not priced yet	The market sees power constraints, but still talks about AI infrastructure as a chip and cloud-capex race. The less priced layer is site selection as an energy-development problem: fiber plus land plus behind-the-meter generation plus permitting, especially where next-generation geothermal can bypass interconnection delays.
Where the price sits today	Transformer shortage and grid congestion are visible. The residual edge is not the fact that power is tight; it is the claim that firm-power site rights become a primary AI platform asset, and that geothermal or other clean firm behind-the-meter resources get valued as AI infrastructure rather than as ordinary generation.
The binding constraint	Contiguous land, fiber proximity, behind-the-meter firm generation rights, interconnection optionality, cooling, and local permitting bundled into a deployable data-center campus.
What we are watching	Track hyperscaler and data-center developer announcements that name on-site firm power, geothermal, or interconnection bypass as the reason for site selection; count 100 MW plus campuses with direct power-development partnerships; track transformer lead times and local moratoria.
What would prove us wrong	Kill if by end 2028 fewer than two hyperscaler-scale campuses publicly secure behind-the-meter firm clean generation as a core siting advantage, or if transformer and interconnection delays normalize below roughly 24 months in the main US AI data-center markets.
How we tried to break it	Already priced at the headline power level, not at the site-rights level. The obvious constraint is power. The promoted claim is narrower: the next strategic asset is not generic megawatts, but the combined package of firm power, land, fiber, and permission that lets AI load come online faster than grid-dependent sites.

Why we are making the call

The first-order AI power story is now consensus. The second-order call survives because it names where the rent migrates after that consensus: not just to generators, but to sites and developers that can collapse the time from model demand to energized capacity.

If the call is right

Who is exposed

Hyperscaler infrastructure teams, data-center developers, power developers, large AI labs with reserved compute needs, and investors underwriting AI infrastructure.

Action now

Map sites by firm-power time-to-energize, not just land cost and fiber; secure options on campuses where geothermal, gas with carbon capture, nuclear restart, or other firm power can be contracted behind the meter.

Decision it changes

Data-center site selection, PPA strategy, power-development partnerships, capex phasing, and portfolio exposure to grid-dependent versus power-secured data-center assets.

ROI / risk logic

A site that energizes 12 to 24 months earlier can be worth more than a cheaper site with stranded shells and delayed transformers. The asymmetry is time: idle GPUs and delayed leases burn capital while power-secured campuses monetize demand.

Rent lands in developers and landowners with power-secured campuses, geothermal developers such as Fervo-style EGS operators, and utilities or IPPs that can deliver firm interconnection alternatives. It does not land evenly across data-center real estate.

Who gains

Geothermal and clean firm power developers: They turn data-center load into a bankable offtake and bypass part of the grid queue.

Data-center developers with power-secured land: They can sell time-to-energize, not square footage.

Hyperscalers with flexible inference siting: Inference workloads can move toward power instead of clustering only near legacy hubs.

Who loses

Grid-dependent campus projects in congested markets: They face transformer, interconnection, and local consent delays that make announced capacity less real.

AI infrastructure investors underwriting shell capacity only: They risk owning buildings that cannot be energized on the underwriting timeline.

What reprices

Power-secured land options, behind-the-meter PPAs, geothermal development rights, and data-center lease premiums should reprice upward relative to ordinary shells.

The next constraint it creates

The next constraint moves to drilling capacity, high-voltage equipment, water and cooling permits, and local political consent for large loads.

Earliest sign it has begun

A hyperscaler or top data-center REIT announcing a 100 MW plus campus whose stated differentiator is behind-the-meter clean firm power rather than cheap land.

P2

Physical AI's bottleneck is certified deployment, not robot bodies.

Domain: robotics / industrial automation

2028-12-31

Structural case	Our call, dated	Resolves
78%	46%	2028-12-31

Humanoid and industrial robot pilots are moving toward early commercial deployment, especially in automotive and logistics, but every physical deployment has to cross a hard boundary: task data, safety, uptime, workcell integration, maintenance, and liability. NVIDIA is explicitly pushing world models, Isaac Sim, Isaac Lab, Cosmos, and GR00T as simulation and validation infrastructure. IDTechEx's humanoid materials and robotics coverage highlights tactile sensors, training data, software, simulation, and component bottlenecks. The body is visible; the deployable task stack is the constraint.

The boom	Robot commissioning software, safety validation, simulation assets, and task libraries capture more durable value than many undifferentiated humanoid bodies.
Why it is not priced yet	Capital chases the robot OEM narrative. Buyers pay for productive hours, uptime, and safety sign-off. The pre-consensus layer is that task deployment and validation become separately priced software/services, not just bundled OEM support.
Where the price sits today	Robotics platform hype is visible in private valuations and public AI narratives. The narrower deployment-layer thesis is less directly priced because the value sits inside integration contracts, simulation tools, safety cases, and task libraries rather than a clean public instrument.
The binding constraint	Verified task data, sim-to-real validation, workcell commissioning, safety certification, and field support for physical AI in factories and warehouses.
What we are watching	Track robot OEMs or large integrators selling task libraries, simulation validation, or commissioning layers as separate line items; track public claims of 40 percent plus cuts in commissioning time; track safety certification language in physical-AI deployments.
What would prove us wrong	Kill if by end 2028 humanoid or mobile manipulation deployments scale mainly through turnkey robot hardware with little separate pricing for task validation, commissioning software, or safety case tooling.
How we tried to break it	The existing FUTURE_MAP already called auto-commissioning and robot components as bottlenecks. This call survives by shifting the buyer object: not 'robotics booms', but 'deployment assurance becomes the product buyers pay for'. It is a commercial packaging call layered on the technical constraint.

Why we are making the call

Robots do not sell as intelligence; they sell as safe, repeatable labor hours. The scarce layer is the proof that a robot can do the task in a customer's messy environment without breaking the line, hurting someone, or requiring heroic integration labor.

If the call is right

Who is exposed

Manufacturing COOs, logistics operators, industrial automation integrators, robot OEMs, and investors in physical AI companies.

Action now

Inventory repeatable tasks by commissioning burden and safety risk; structure pilots around validated productive hours and time-to-commission, not robot count or demo quality.

Decision it changes

Automation capex, integrator selection, pilot design, safety certification budgets, and whether to invest in robot OEMs or deployment-layer tooling.

ROI / risk logic

A robot that takes six months of integration labor to deliver a narrow task has poor ROI even if the body is cheap. Deployment tooling pays by converting demos into billable productive hours faster.

Rent flows to simulation and validation platforms, robot integrators with reusable task libraries, and OEMs that can prove uptime and safety. Commodity bodies get squeezed unless they own the deployment layer.

Who gains

NVIDIA Isaac/Cosmos-style simulation ecosystems: They sit in the validation path between robot policy training and real-world deployment.

Industrial integrators with reusable task libraries: They turn one-off commissioning labor into repeatable software-enabled deployment.

Manufacturers with clean process data and standardized workcells: They become easier customers and capture automation ROI earlier.

Who loses

Undifferentiated humanoid OEMs: Hardware demos do not create durable margin if deployment assurance is owned elsewhere.

Custom-only automation integrators: Reusable task libraries and simulation reduce the value of bespoke labor.

What reprices

Physical-AI valuation should migrate from unit shipments to deployment software attach rate, commissioning margin, and validated task-hour economics.

The next constraint it creates

The next constraint becomes high-quality real-world task data, tactile sensing, safety certification labor, and maintenance networks.

Earliest sign it has begun

A major automotive or logistics deployment where the press release names the validation, simulation, or task-library layer as the reason the rollout scaled.

P3

Autonomous science shifts the bottleneck from model discovery to experimental throughput.

Domain: autonomous science / biotech / materials

2029-12-31

Structural case	Our call, dated	Resolves
76%	44%	2029-12-31

AI can propose experiments faster than labs can run, measure, standardize, and trust them. DOE's 2026 robotics and automation testbed call for autonomous scientific discovery points to the institutionalization of this bottleneck, with expected awards of 3 to 5 million dollars per year. The binding input is not another model; it is reliable experiment generation, instrument uptime, assay reproducibility, and machine-readable provenance.

The boom	Self-driving labs, assay capacity, robotic testbeds, metadata infrastructure, and provenance systems become the scarce layer under AI-for-science.
Why it is not priced yet	AI-for-science narratives still reward model papers and discovery claims. The more durable bottleneck is boring: the labs, assays, protocols, and reference data that make closed-loop discovery trustworthy.
Where the price sits today	Self-driving lab companies and lab automation vendors are funded, but the market has not cleanly priced experimental throughput as the scarce asset under model-driven discovery. Much of the value sits inside R&D budgets, national labs, and pharma partnerships.
The binding constraint	Robotic experimental throughput with standardized metadata, reliable assays, calibrated uncertainty, and provenance that lets models learn from results without poisoning the loop.
What we are watching	Track DOE and national-lab testbed awards, pharma/materials partnerships that buy autonomous lab capacity, assay throughput per researcher, and publications where the bottleneck is experiment generation rather than model inference.
What would prove us wrong	Kill if by 2029 model-only AI discovery companies repeatedly produce commercially validated materials or therapies without materially expanding wet-lab or physical-test throughput.
How we tried to break it	This is not a claim that autonomous labs are unknown. The promoted edge is that experimental throughput, provenance, and assay quality become the decision-grade moat once proposal generation gets cheap.
Why we are making the call	Discovery systems learn from measurements. If measurements are slow, noisy, unstandardized, or untrusted, model quality cannot compound. The valuable asset becomes the loop, not the idea generator.

If the call is right

Who is exposed

Pharma R&D leaders, materials companies, national labs, AI-for-science founders, and investors underwriting discovery platforms.

Action now

Audit which discovery programs are experiment-limited; secure access to robotic testbeds, assay vendors, and metadata/provenance infrastructure before model-generated proposal volume spikes.

Decision it changes

R&D budget allocation between foundation models, lab automation, assay capacity, data infrastructure, and strategic partnerships.

ROI / risk logic

A better model has low ROI if every validation loop takes months. Shortening the experiment loop compounds across the entire discovery funnel and creates proprietary datasets competitors cannot scrape.

Rent lands in autonomous lab operators, instrument integration platforms, assay providers, and organizations with high-quality closed-loop experimental datasets.

Who gains

DOE and national lab testbed ecosystems: They become shared infrastructure for autonomous discovery workflows.

Autonomous lab operators and assay platforms: They own the throughput and data rights needed to make AI-for-science real.

R&D teams with strong metadata discipline: Their experiments become reusable training signal instead of isolated results.

Who loses

Model-only discovery startups: They can propose more than they can validate, making their output easy to discount.

Manual CROs without automation or data standardization: They provide services but not compounding machine-readable learning loops.

What reprices

Autonomous lab capacity, assay contracts, and proprietary experimental datasets should command higher strategic value than generic AI-for-science claims.

The next constraint it creates

The next constraint becomes sample preparation, protocol ontologies, instrument interoperability, and rights to high-quality negative results.

Earliest sign it has begun

A major pharma, chemicals, or materials company publicly saying its AI discovery bottleneck is lab throughput or assay availability, not model quality.

P4

The consumer AI interface moves to always-on edge devices, gated by thermals, battery, sensors, and privacy.

Domain: edge AI / consumer hardware

2028-12-31

Structural case	Our call, dated	Resolves
70%	43%	2028-12-31

The next interface cannot rely on sending every private, continuous, multimodal signal to the cloud. Mordor estimates the edge AI hardware market at 30.74 billion dollars in 2026 and 68.73 billion dollars by 2031. Qualcomm's 2026 Snapdragon Wear Elite announcement explicitly frames wearables as personal AI devices with on-device NPUs, sensor fusion, always-on context, and partners including Google, Motorola, and Samsung. The hard part is not fitting a toy model on device; it is making the device always available without killing battery, privacy, thermals, or trust.

The boom	Wearables, glasses, phones, cars, drones, and robots with local NPUs become the main surface where AI feels useful.
Why it is not priced yet	The market understands on-device AI in broad terms. The narrower call is that the winning interface is not a standalone AI gadget, but an always-on sensor-rich device layer where privacy, power, and context are the differentiators.
Where the price sits today	Chip and device companies are announcing edge AI features. The residual edge is the constraint stack: battery life, local memory, secure sensor fusion, and developer access to context without leaking user trust.
The binding constraint	Low-power NPUs, sensor fusion, local memory, privacy-preserving orchestration, and thermal envelopes for always-on personal AI.
What we are watching	Track commercial devices that run billion-parameter local models, all-day context capture, or private agentic features on wearables/glasses/phones; track battery complaints, thermal throttling, and local AI developer APIs.
What would prove us wrong	Kill if by end 2028 the dominant consumer AI usage remains cloud-chat inside phones and browsers, with wearables and glasses failing to show persistent local context as a major usage mode.
How we tried to break it	Standalone AI gadgets have already disappointed once, and phones can absorb many features. That is why the call is not 'new gadget category wins'. The call is that always-on local context becomes the scarce interface layer across devices, whether branded as watch, glasses, earbuds, car, or phone.
Why we are making the call	The next UI for AI is context. Context lives near sensors, bodies, homes, and vehicles. Cloud-only agents lack the continuous private state needed to be useful, while device-only systems fail unless power and thermals work.

If the call is right

Who is exposed

Device OEMs, chip vendors, mobile OS teams, consumer AI app developers, regulated enterprise buyers, and privacy/security teams.

Action now

Design around local context and power budgets; build features that improve when personal data stays on device instead of depending on generic cloud memory.

Decision it changes

Hardware roadmap, NPU vendor selection, app architecture, privacy posture, and whether to build for cloud-first chat or local context surfaces.

ROI / risk logic

Always-on local context can raise retention and trust, while cloud-only interfaces compete on commoditized model access. Battery and privacy are the cost of admission.

Rent lands in chip vendors with efficient NPUs, OEMs with integrated sensor stacks, OS owners with trusted permission models, and apps that turn local context into recurring utility.

Who gains

Qualcomm, Apple, Samsung, and other integrated edge-AI silicon/device players: They control the NPU, sensors, battery envelope, and developer surfaces.

Privacy-preserving local agent platforms: They can access sensitive context without forcing everything into cloud logs.

Apps with daily personal context loops: Health, mobility, work, and communication features improve with persistent local state.

Who loses

Cloud-only personal assistant apps: They lack trusted continuous context and face rising privacy friction.

Low-end device makers without NPU or sensor-fusion capability: They cannot support the local features users come to expect.

What reprices

NPU-rich device BOMs, sensor vendors, secure local memory, and OS-level permission surfaces should reprice relative to generic chatbot app layers.

The next constraint it creates

The next constraint becomes developer access to local context without privacy collapse, plus energy-efficient multimodal inference.

Earliest sign it has begun

A Google, Samsung, Apple, or Meta product launch where the lead AI feature runs persistently on device and is marketed around private context rather than a bigger cloud model.

P5

Biomanufacturing's bottleneck is scale-up, not AI organism design.

Domain: industrial biotech / biomanufacturing

2030-12-31

Structural case	Our call, dated	Resolves
72%	41%	2030-12-31

AI and metabolic engineering can generate more candidate organisms and pathways, but industrial deployment still runs through strain robustness, fermentation, downstream purification, feedstock economics, QA, and price parity. Nature Communications noted in 2026 that large-scale commercialization remains limited despite systems metabolic engineering progress, highlighting the persistent gap between lab success and industrial deployment. The model proposes; the plant decides.

The boom	Pilot fermentation, downstream processing, QA, feedstock contracts, and industrial bioprocess talent become more valuable than many AI-designed strains.
Why it is not priced yet	AI-bio narratives reward discovery speed and design quality. The underpriced layer is the boring industrial bridge from bench to tons: COGS, contamination, yield under real feedstocks, purification, regulatory proof, and buyer qualification.
Where the price sits today	Biomanufacturing capacity and CDMO services are visible, but scale-up bottlenecks are not cleanly priced as the scarce input under AI-designed biology. Many private AI-bio valuations still lean on discovery rather than manufacturability.
The binding constraint	Pilot and commercial-scale fermentation, downstream processing, strain robustness, feedstock flexibility, QA systems, and process-development talent.
What we are watching	Track AI-designed or engineered bio-products that fail or delay on COGS and scale-up; track pilot fermentation capacity, downstream bottlenecks, and offtake contracts tied to price parity rather than sustainability premium.
What would prove us wrong	Kill if by 2030 multiple AI-designed industrial bio-products reach commodity-relevant scale and price parity without scarce pilot capacity, downstream processing, or process-development labor becoming a public bottleneck.
How we tried to break it	This is a known pain inside biotech, but not yet the headline investment story. The call survives because AI increases the number of designs competing for the same scale-up bridge, making the bridge more scarce, not less.
Why we are making the call	A thousand plausible organisms do not equal one profitable plant. The constraint migrates from search to industrialization once AI makes search cheaper.

If the call is right

Who is exposed

Industrial biotech founders, chemicals strategics, food and materials companies, climate-tech investors, and CDMOs.

Action now

Underwrite AI-bio companies on scale-up evidence, COGS path, downstream process, feedstock contracts, and pilot capacity access before valuing design claims.

Decision it changes

Investment diligence, partnership timing, pilot plant capex, CDMO contracting, and whether to buy design IP or process know-how.

ROI / risk logic

Early access to pilot and downstream capacity can turn designs into qualified products while competitors wait in the scale-up queue. Avoiding failed scale-up saves years of burn.

Rent lands in companies with process-development talent, qualified pilot capacity, downstream purification know-how, and feedstock-secured plants. Design-only IP captures less unless tied to manufacturability.

Who gains

Bioprocess CDMOs and pilot fermentation operators: They own the scarce bridge from strain to commercial product.

Strategics with existing fermentation and downstream assets: They can industrialize AI-generated designs faster than pure software-led entrants.

AI-bio teams with process development built in: They can select designs for manufacturability instead of novelty.

Who loses

Design-only AI-bio startups: They face valuation compression when customers ask for yield, cost, and qualified volume.

Commodity bio-products relying only on green-premium demand: Nature's review notes commercial viability ultimately needs price parity or a tangible product benefit.

What reprices

Pilot fermentation slots, downstream processing assets, process-development teams, and scale-up data should reprice above generic AI-bio discovery claims.

The next constraint it creates

The next constraint becomes low-cost feedstock, contamination control, waste handling, and buyer qualification at industrial volumes.

Earliest sign it has begun

A well-funded AI-bio or synbio company delaying launch because pilot capacity, downstream recovery, or COGS fails despite successful lab design.

P6

Agentic AI's scarce layer becomes authority, auditability, and rollback.

Domain: enterprise AI / security / governance

2028-06-30

Structural case	Our call, dated	Resolves
74%	48%	2028-06-30

Agents that only answer questions are easy to pilot. Agents that execute code, move money, change records, message customers, operate devices, or trigger workflows create liability. Berkeley's Agentic AI Risk Management Profile highlights risks from system configuration, tool access, and real-world interaction, including unauthorized privilege escalation and unsupervised erroneous tasks. As agents gain authority, the binding constraint moves from reasoning quality to controlled action.

The boom	Agent security, identity, execution controls, audit logs, insurance, and governance middleware become the enterprise adoption gate.
Why it is not priced yet	Enterprise AI buyers talk about agents, but many roadmaps still assume model capability is the adoption gate. The next purchasing gate is operational trust: least privilege, chain of custody, approval routing, audit logs, rollback, and insurance language.
Where the price sits today	AI security and governance are increasingly visible. The narrower edge is action-governance as a separately budgeted execution layer for agents, especially where agents touch money, regulated records, infrastructure, or physical systems.
The binding constraint	Identity, permissioning, tool-access control, audit trails, reversible execution, human authorization policies, and liability evidence for agent actions.
What we are watching	Track enterprise RFPs requiring agent audit logs, least-privilege controls, rollback, or insurance support; track public agent incidents; track vendors selling action-governance rather than generic model monitoring.
What would prove us wrong	Kill if by mid-2028 large enterprises widely deploy multi-step agents with real system authority using mostly prompt-level guardrails and generic logging, without a separate action-governance budget.
How we tried to break it	Governance is not new, and some budget already exists. The promoted claim is that execution authority creates a new control plane: not AI ethics reporting, but operational proof of who authorized what action, which tool ran, and how damage can be reversed.
Why we are making the call	The useful agent is the dangerous agent. Once AI can act, buyers need evidence, not vibes. That evidence layer is monetizable because it gates procurement, insurance, and legal approval.

If the call is right

Who is exposed

CISOs, CIOs, compliance leaders, insurers, enterprise software buyers, agent platform vendors, and legal teams.

Action now

Design agent deployments around authority boundaries first: least privilege, approval ladders, audit logs, rollback, and evidence packets for regulated workflows.

Decision it changes

Enterprise AI procurement, security architecture, insurance underwriting, compliance budgets, and whether agents are allowed to touch systems of record.

ROI / risk logic

A controlled agent can be deployed into high-value workflows that an uncontrolled agent cannot enter. The upside is adoption; the avoided loss is incident cost, regulatory exposure, and procurement blockage.

Rent lands in identity/security vendors, agent orchestration platforms with strong execution controls, audit-log infrastructure, and insurers that can price agent risk.

Who gains

Identity and access-management vendors: Agent authority maps naturally onto permissioning, secrets, and least-privilege workflows.

Agent orchestration platforms with audit and rollback: They become the control plane for letting agents touch real systems.

Insurers and compliance tooling providers: They can turn auditable controls into underwritable risk.

Who loses

Prompt-only agent frameworks: They cannot satisfy buyers who need proof, rollback, and authority boundaries.

Model vendors that ignore enterprise control planes: Capability alone does not clear procurement for regulated action.

What reprices

Agent security budgets, identity-control spend, audit-log infrastructure, and cyber insurance language should reprice as agents get execution rights.

The next constraint it creates

The next constraint becomes standard evidence formats for agent actions and cross-vendor interoperability between identity, workflow, model, and audit systems.

Earliest sign it has begun

A Fortune 500 RFP, insurance policy, or regulator explicitly requiring agent action logs, approval chains, and rollback for production deployment.

Seeds considered

These cleared the supply-side test but did not make the final board, usually because the trade was not clean or the move was already priced.

Seed	Physical case	Why not promoted
World models and digital twins as the physical-AI data layer	Validated simulators and real-to-sim-to-real data could become the scarce training asset for robots, factories, and spatial agents.	Folded into P2 because the buyer-facing decision is deployment assurance. A standalone world-model call risks repeating the physical-AI platform narrative without a sharper buyer action.
Local opposition to data centers becomes the AI deployment governor	Water, noise, rates, and local moratoria could block nominal AI capacity.	Important, but it is a sub-constraint under P1's power-and-siting thesis. Local consent belongs in the site-rights bundle.
AI fine-tuning and benchmark wins as fundraising leverage	Public benchmark wins can provide social proof for a solo founder even when they are not the product.	Useful commercialization surface, not a structural 'what comes after AI' thesis. Keep it in the fundraising story, not the future board.
Transformer and GOES shortages as the data-center bottleneck	Electrical equipment remains a hard constraint on AI capacity.	Already strongly represented in FUTURE_MAP. This board moves one step sideways to site-controlled firm power as the actionable buyer decision.

Each call is dated. The line that would prove it wrong is fixed when the board is issued.