The World After AI: intelligence gets permission to act

The next big thing is not a smarter chatbot. It is the permission layer that lets machine intelligence move money, matter, evidence, biology, power, and liability without breaking the world.

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

post-AI world / next big thing after AI

Horizon

2030 to 2040

Issued

2026-06-17

Method

Wide cast, adversarial gate, public resolution criteria.

Board summary

The cross-cutting read

The post-AI world is not defined by intelligence becoming scarce. It is defined by intelligence becoming cheap enough that the bottleneck moves to permission, coupling, and verification. The winning layer is wherever an AI system tries to act on the world and hits an inelastic gate: a payment rail that can authorize a machine, a legal record that can assign blame, a lab protocol that another machine can replay, a materials dataset that proves lifetime reliability, a synthesis screen that lets biology leave the model, a causal atlas that proves a drug target in human tissue, a grid front end that can absorb AI-load dynamics, or a heat sink that lets the compute campus exist. The phrase "after AI" therefore does not mean AI fades. It means AI becomes the cheap input, and the scarce assets are the rails that make action safe, admissible, insurable, and physically possible.

At a glance

#	Claim	Binding constraint	Case	Call	Resolves
P1	Payment networks, identity providers, agent runtime platforms, and compliance layers that can bind an AI action to a...	Verifiable agent identity, delegated permission tokens, spending and action limits, per-action audit logs...	76%	49%	2031-12-31
P2	The durable post-AI compliance market is not explainability theater. It is tamper-evident event recording that lets...	Tamper-evident action logs, incident reports, model and tool version provenance, authority trace, retention...	76%	51%	2031-12-31
P3	The next scientific platform is not the AI scientist in isolation. It is the shared, replayable lab stack where...	Machine-readable protocol, sample lineage, instrument-control standards, data and knowledge management...	76%	48%	2031-12-31
P4	AI will make plausible materials abundant; the rent moves to accelerated lifetime testing, reliability statistics, and...	Accelerated lifetime testing, reliability statistics, standardized qualification data, critically evaluated...	72%	42%	2035-12-31
P5	When AI can design biological sequences, the chokepoint shifts to who is allowed to synthesize them, order them, or...	AI-resilient sequence and function screening, customer verification, benchtop synthesizer access controls...	74%	46%	2030-12-31
P6	As AI target generation gets cheap, pharma pays for proprietary maps that show how human cells and organoids actually...	Human perturbation atlases in organoids, primary cells, disease-relevant tissue, spatial multi-omics, and...	68%	34%	2033-12-31
P7	AI campuses become grid actors. The valuable infrastructure is the power-electronics, battery, model, and...	Dynamic load models, EMT-grade studies, ride-through commitments, grid-forming inverters, battery front ends...	68%	38%	2033-12-31
P8	The best AI real estate is not only cheap land with power. It is a site with permitted heat rejection, reclaimed water...	Legal and physical heat rejection: district-heat offtake, non-potable or reclaimed cooling water, thermal...	72%	41%	2032-12-31

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

P1

Agent authority rails become the transaction network after AI.

Domain: agentic commerce / identity / payments 2031-12-31

Structural case	Our call, dated	Resolves
76%	49%	2031-12-31

Once agents can plan and execute, the binding input is no longer task intelligence. It is legally attributable authority: who allowed the agent to act, what the agent was allowed to spend or change, whether the action can be reversed, and who carries loss if it exceeds scope.

The boom	Payment networks, identity providers, agent runtime platforms, and compliance layers that can bind an AI action to a delegated legal authority capture the first durable post-AI rent.
Why it is not priced yet	Agent apps are visible and payment pilots are now visible, but the market still mostly prices front-end agent experiences. The deeper rail is credentialed authority with controls, settlement, auditability, and dispute handling across vendors.
Where the price sits today	Mastercard launched Agent Pay for Machines in June 2026, with credentialing, permissioning, and settlement for machine-driven transactions. Visa Intelligent Commerce similarly frames trusted AI-initiated transactions around credentials, controls, authentication, and protections. The IMF notes that agentic payments create a tension between probabilistic AI behavior and deterministic payment infrastructure.
The binding constraint	Verifiable agent identity, delegated permission tokens, spending and action limits, per-action audit logs, dispute rules, and settlement rails that work at machine speed.
What we are watching	Count of payment, identity, or cloud platforms operating agent credential networks; transaction volume through agent-specific credentials; procurement or regulatory text requiring agent identity and delegated authority for high-value actions.
What would prove us wrong	Kill if by 2031 agents still transact at scale through ordinary user credentials, API keys, or human-initiated payment flows with no separate agent identity, permission, or liability layer.
How we tried to break it	Existing OAuth, card tokenization, and enterprise SSO could stretch far enough, especially inside closed ecosystems. The call survives because open commercial agents need cross-platform settlement, merchant trust, revocation, chargeback evidence, and authorization semantics that ordinary human credentials do not provide.

Why we are making the call

The post-AI economy starts when agents can buy, bind, reserve, meter, settle, and contract. The winning primitive is not a model. It is permissioned action.

If the call is right

Who is exposed

Payment networks, issuer processors, merchant acquirers, cloud agent platforms, identity vendors, and enterprise risk teams that expect agents to move money or make commitments.

Action now

Inventory where agents could initiate payments or commitments, then require scoped credentials, explicit consent, revocation, logging, and dispute evidence before any production rollout.

Decision it changes

Payment and identity roadmap priorities shift from human checkout UX toward agent-specific credentials, programmable limits, merchant trust signals, and liability handling.

ROI / risk logic

Early rails win standards gravity. If the agent economy turns into high-volume microtransactions, the infrastructure provider that owns credentialing and settlement gets transaction take-rate and compliance lock-in before apps commoditize.

Value lands at Mastercard, Visa, Stripe, Adyen, Cloudflare, identity providers, agent runtime vendors, and compliance startups that can bridge agent authorization with real settlement and dispute processes.

Who gains

Mastercard and Visa: They can extend existing merchant acceptance, tokenization, fraud controls, and settlement trust into agent-specific commerce.

Cloudflare, Stripe, Adyen, and identity vendors: They sit at the control plane where agents authenticate, authorize, and transact across business systems.

Agent runtime and policy-control startups: They can become the layer that enforces scopes, approvals, and audit trails before money or state changes happen.

Who loses

Closed single-app agent stores: They lose leverage if merchants and enterprises demand interoperable credentials and dispute rules.

Unlogged consumer-agent apps: They cannot clear enterprise, insurance, or payment risk reviews once agents handle money.

Merchants without agent-readable inventory and policy interfaces: They become invisible to automated buyers that require machine-readable terms and trusted checkout.

What reprices

Network fees, merchant acceptance contracts, fraud-liability terms, stablecoin settlement services, and enterprise identity spend reprice toward agent-specific controls.

The next constraint it creates

After authority rails exist, the next bottleneck is agent reputation: portable evidence that a given agent or agent vendor performs within policy over time.

Earliest sign it has begun

Watch for two major card networks, processors, or cloud providers reporting production agent credential volume, not just pilots, and for procurement language requiring distinct credentials for autonomous software actors.

P2

Liability moves to the evidence recorder.

Domain: liability / insurance / audit evidence

2031-12-31

Structural case	Our call, dated	Resolves
76%	51%	2031-12-31

Autonomous systems create losses before institutions agree who is responsible. Courts and regulators cannot assign blame from model descriptions alone; they need logs that preserve version, input, authority, tool call, output, human override, and incident context.

The boom	The durable post-AI compliance market is not explainability theater. It is tamper-evident event recording that lets insurers, courts, regulators, and boards reconstruct what an autonomous system did.
Why it is not priced yet	AI governance is crowded, but most spend is still framed as policy, evaluation, red teaming, or certification. The granular rent moves to admissible causal trace, because every serious incident becomes a dispute over what happened.
Where the price sits today	The EU AI Act has record-keeping, post-market monitoring, and serious-incident reporting obligations. NHTSA already requires identified manufacturers and operators to report certain ADS and ADAS crashes. ISO/IEC 42001 defines an AI management system that auditors can certify, but event reconstruction remains underbuilt.
The binding constraint	Tamper-evident action logs, incident reports, model and tool version provenance, authority trace, retention rules, and replayable forensic records.
What we are watching	Number of G7 or EU-regulated sectors requiring AI action logs or incident records for legal defense, incident reporting, market access, or insurance coverage; policy wording requiring AI event records.
What would prove us wrong	Kill if by 2031 serious AI incidents are still litigated mainly through ordinary negligence discovery and logs remain voluntary internal telemetry with no sector-specific effect on insurance, legal defense, or market access.
How we tried to break it	Regulators may keep requirements flexible and vendors may resist discoverable logs. The call survives because high-stakes AI systems need a blame surface, and existing sectors such as automated driving already show the shape of mandatory incident reporting.
Why we are making the call	After AI acts, the scarce asset is not an explanation of a model. It is admissible causality.

If the call is right

Who is exposed

AI assurance firms, cyber insurers, regulated enterprises, autonomous vehicle and robotics operators, healthcare AI vendors, and enterprise legal teams.

Action now

Define an AI event-record schema before deployment: model version, prompt/input class, tools called, permissions used, output, human approval, rollback path, incident marker, and retention period.

Decision it changes

AI compliance budgets shift from one-off model evaluations to persistent logging, forensic replay, evidence retention, and independent assurance.

ROI / risk logic

A firm with usable incident evidence can defend claims, obtain coverage, and fix systems faster. A firm without it pays through exclusions, discovery costs, downtime, and regulatory penalties.

Value lands at audit firms, observability vendors, security logs, GRC platforms, and specialist AI forensics providers that can make event records acceptable to insurers and regulators.

Who gains

Audit and assurance providers: They can certify whether controls and logs actually exist and operate.

Security observability and data-retention vendors: They already own log pipelines and can extend them to AI event provenance.

Insurers with AI-specific underwriting data: They can price coverage instead of issuing blanket exclusions.

Who loses

Model vendors with opaque tool execution: They become harder to procure in regulated settings.

Enterprises that treat AI logs as optional telemetry: They will lack evidence when a failure becomes a claim or investigation.

Generic GRC vendors without runtime data: Static policy workflows do not reconstruct autonomous action.

What reprices

Cyber, E&O, product liability, and D&O insurance terms reprice around AI event evidence and governance warranties.

The next constraint it creates

Once event records exist, the next constraint is standardizing which logs are admissible without exposing trade secrets, personal data, or security-sensitive prompts.

Earliest sign it has begun

Watch insurance questionnaires and regulated-sector procurement templates for mandatory AI action-log, incident replay, and retention language.

P3

Autonomous labs bottleneck on standards before autonomy.

Domain: autonomous science / lab infrastructure

2031-12-31

Structural case	Our call, dated	Resolves
76%	48%	2031-12-31

AI can propose experiments faster than labs can execute them, but a more subtle bottleneck appears before raw robot capacity: every autonomous lab has bespoke instrument commands, sample IDs, metadata, protocols, safety constraints, and data schemas. Without standards, each lab is an island.

The boom	The next scientific platform is not the AI scientist in isolation. It is the shared, replayable lab stack where protocols, samples, instruments, data, and model versions can move across machines.
Why it is not priced yet	Autonomous labs are now a visible theme. The underpriced layer is not lab robots; it is machine-readable experiment memory, interoperability, and audit replay. That is where a shared scientific execution market becomes possible.
Where the price sits today	NIST says there is no standardized autonomous experimentation ecosystem for materials R&D and calls this a major obstacle to industrial adoption. DOE's Genesis Mission connects supercomputers, experimental facilities, AI systems, and datasets. NSF has funded BioFoundries to create high-throughput measurement infrastructure.
The binding constraint	Machine-readable protocol, sample lineage, instrument-control standards, data and knowledge management standards, model integration, audit replay, and safety-constrained orchestration.
What we are watching	Public procurements or funding programs above $25 million requiring named autonomous-lab standards; industrial contracts requiring standards for sample handling, instrument communication, data management, and model integration.
What would prove us wrong	Kill if by 2031 major autonomous-lab deployments remain dominated by bespoke vendor stacks with no named protocol, sample, data, or instrument standards required in public procurements or industrial contracts.
How we tried to break it	Vertical platforms may win by bundling hardware, software, and data, making open standards irrelevant. The call survives because public labs, industrial consortia, and multi-vendor instrument ecosystems need interoperability to avoid lock-in and obsolete custom integration.
Why we are making the call	When AI makes ideas cheap, the next scarce layer is trusted physical execution that another machine can repeat.

If the call is right

Who is exposed

National labs, pharma R&D operations, materials companies, analytical-instrument vendors, lab automation vendors, and cloud scientific-workflow platforms.

Action now

Require every autonomous-lab pilot to export protocols, sample state, instrument state, data provenance, model version, and safety constraints in a machine-readable schema.

Decision it changes

R&D infrastructure budgets shift from isolated robot workcells toward interoperable lab operating systems, data schemas, and integration standards.

ROI / risk logic

A standards-first lab stack avoids repeated custom integration, lets experiments move across sites, improves reproducibility, and creates a market for experiment capacity.

Value lands at instrument vendors and lab orchestration platforms that become standard-compliant by default, plus national lab networks that offer trusted experiment capacity.

Who gains

NIST-aligned standards consortia and national lab networks: They can define the interoperability layer buyers require.

Thermo Fisher, Danaher, Agilent, Waters, Bruker, and automation vendors: They can sell instruments and control software into closed-loop experiment networks.

Lab orchestration software providers: They become the system of record for experiment execution.

Who loses

Closed bespoke self-driving-lab stacks: They become procurement risks when buyers require interoperability and replay.

Pure model-discovery startups: They lack the physical validation capacity that buyers increasingly demand.

Manual CRO workflows without machine-readable provenance: They cannot plug into agentic research loops at scale.

What reprices

Instrument control APIs, lab information management systems, scientific cloud workflows, and national-lab user facility access reprice upward.

The next constraint it creates

Once standards exist, the next bottleneck is scarce calibrated metrology and instrument uptime, especially inline characterization.

Earliest sign it has begun

Watch DOE, NSF, NIST, pharma, and materials RFPs for named requirements around sample management, instrument communication, data management, and model integration standards.

P4

Materials discovery becomes a qualification-data business.

Domain: materials / qualification / reliability 2035-12-31

Structural case	Our call, dated	Resolves
72%	42%	2035-12-31

The quantity of AI-suggested materials rises faster than the ability to qualify them for batteries, aircraft, chips, catalysts, grids, or medical devices. For high-consequence products, discovery is not the gate; lifetime, manufacturability, and reliability evidence are.

The boom	AI will make plausible materials abundant; the rent moves to accelerated lifetime testing, reliability statistics, and qualification datasets that let a material enter a product.
Why it is not priced yet	The market celebrates AI materials discovery. It underprices the slow part: duty-cycle testing, accelerated aging, standards, sample history, and critically evaluated datasets that product engineers trust.
Where the price sits today	NIST's Materials Genome Initiative emphasizes essential data exchange protocols, quality of materials data and models, and metrologies needed for accelerated materials development. That is the quiet infrastructure beneath AI discovery claims.
The binding constraint	Accelerated lifetime testing, reliability statistics, standardized qualification data, critically evaluated materials datasets, and duty-cycle-specific testbeds.
What we are watching	Count of AI-discovered or AI-optimized materials missing original commercial design-in targets by at least 36 months due to qualification, reliability, or lifetime evidence; funding or revenue for shared qualification testbeds and data consortia.
What would prove us wrong	Kill if by 2035 three or more AI-discovered materials reach high-volume design-in within 36 months of disclosure without public qualification delays or shared qualification-data businesses becoming material.
How we tried to break it	Materials qualification has always been slow, so this could be obvious rather than alpha. The call survives because AI increases candidate volume and shifts the scarce input from ideas to trusted evidence, creating a new data business rather than merely continuing old slowness.

Why we are making the call

The post-AI world gets flooded with plausible designs. The valuable actor is the one that can prove which designs survive reality.

If the call is right

Who is exposed

Battery, aerospace, semiconductor, catalyst, energy equipment, and advanced-manufacturing firms evaluating AI-generated material candidates.

Action now

Build or buy qualification-data capacity before model-generated candidate volume overwhelms test pipelines: accelerated aging, inline metrology, sample lineage, and field-failure feedback.

Decision it changes

R&D budgets shift from model-only discovery tools toward reliability testbeds, reference datasets, and standards participation.

ROI / risk logic

The cost of waiting is a stranded candidate portfolio: many plausible materials and no credible path to product insertion. Qualification data can shorten design-in cycles and prevent expensive late failures.

Value lands at NIST-aligned data consortia, testing labs, instrument vendors, reliability analytics providers, and incumbents with long field data histories.

Who gains

Testing labs and standards bodies: They can turn qualification into a paid gate for AI-generated materials.

Instrument and metrology vendors: They sell the measurement infrastructure needed to validate lifetime and reliability.

Incumbent materials suppliers with field data: Their historical reliability evidence becomes a moat against model-only entrants.

Who loses

AI materials startups with no qualification path: They may generate candidates faster than customers can trust or adopt them.

Product teams that accept model confidence as qualification: They face late-stage design failures, warranty risk, and certification delays.

Commodity data vendors: Uncurated datasets do not clear product engineering or regulatory thresholds.

What reprices

Reliability datasets, accelerated aging facilities, materials testing contracts, and instrument suites reprice upward; generic AI discovery software reprices downward if it lacks validation access.

The next constraint it creates

After qualification data, the next bottleneck is manufacturing process transfer: whether the material can be made consistently at production scale.

Earliest sign it has begun

Watch public postmortems of AI-materials programs for phrases like lifetime testing, qualification, design-in delay, field reliability, and accelerated aging.

P5

Programmable-biology access becomes a licensed screening layer.

Domain: biosecurity / synthetic biology / access control

2030-12-31

Structural case	Our call, dated	Resolves
74%	46%	2030-12-31

AI biodesign expands the space of possible DNA and protein sequences faster than legacy homology screening can classify them. The scarce layer becomes functional screening, customer legitimacy, recordkeeping, and access control for synthesis providers and benchtop instruments.

The boom	When AI can design biological sequences, the chokepoint shifts to who is allowed to synthesize them, order them, or operate benchtop synthesis equipment.
Why it is not priced yet	Biosecurity is visible in policy circles, but the commercial implication is underpriced: screening becomes an API, compliance service, and access-control business for programmable biology.
Where the price sits today	NIST explicitly warns that AI could design novel DNA sequences undetectable by current sequence screening tools and is developing standards, databases, tools, and capacities for emerging sequences of concern. NIH policy, as issued in 2024 for 2025 implementation, required NIH-funded procurement from providers adhering to the OSTP screening framework.
The binding constraint	AI-resilient sequence and function screening, customer verification, benchtop synthesizer access controls, provider attestations, audit records, and secure screening APIs.
What we are watching	Binding US, EU, or OECD screening rules that go beyond simple sequence-identity matching; number of providers selling AI-resilient screening compliance as a paid API or enterprise product.
What would prove us wrong	Kill if by 2030 screening remains mostly voluntary and identity-based, with no broader functional or AI-resilient requirement and no paid compliance API market.
How we tried to break it	Open-source design tools, local synthesis, and weak international enforcement could route around screening. The call survives because federal funding, institutional procurement, and large commercial providers create a compliance chokepoint even if fringe bypass remains possible.
Why we are making the call	The next big biological constraint after AI is not design. It is permission to turn designs into molecules.

If the call is right

Who is exposed

DNA synthesis providers, biofoundries, benchtop synthesizer makers, universities, pharma, cloud labs, and institutional biosafety officers.

Action now

Map every synthesis procurement path and require provider attestations, customer screening, record retention, and AI-resilient sequence/function screening before expanding automated biology workflows.

Decision it changes

Biofoundry and synthetic-biology platform budgets shift from raw synthesis speed toward compliance screening, access control, and auditable procurement.

ROI / risk logic

Institutions that clear compliant procurement can keep research flowing while others face funding restrictions, halted orders, or reputational damage after a misuse incident.

Value lands at Twist, IDT, Ginkgo-style biofoundries, screening software providers, benchtop synthesizer vendors with access controls, and compliance APIs that become procurement defaults.

Who gains

Large synthesis providers with mature screening: They become preferred vendors for federally funded and institutional research.

Biosecurity screening API vendors: They can sell compliance as a service to smaller providers and benchtop platforms.

Institutions with strong biosafety procurement: They maintain access to synthetic biology tools under tighter rules.

Who loses

Low-cost synthesis providers without screening: They are excluded from funded institutional demand.

Open benchtop synthesis devices with weak access controls: They become procurement and liability risks.

AI-biology startups ignoring biosecurity workflow: They may design faster than they can legally or reputationally synthesize.

What reprices

Synthesis procurement contracts, provider attestations, benchtop instrument controls, and screening API subscriptions reprice upward.

The next constraint it creates

Once screening is standardized, the next bottleneck is international coverage: whether lower-control jurisdictions become leakage points.

Earliest sign it has begun

Watch NIH, DOE, EU, OECD, and provider procurement language for requirements around AI-designed sequences, customer verification, record retention, and benchtop synthesis equipment.

P6

Human causal perturbation atlases become the drug-discovery scarce asset.

Domain: biology / pharma / causal data

2033-12-31

Structural case	Our call, dated	Resolves
68%	34%	2033-12-31

Biology is not text. Model-generated targets need human-relevant causal evidence from perturbation atlases, organoids, primary cells, spatial multi-omics, and longitudinal phenotypes. The rent moves from hypothesis generation to human causal data.

The boom	As AI target generation gets cheap, pharma pays for proprietary maps that show how human cells and organoids actually respond to perturbations.
Why it is not priced yet	Single-cell atlases and virtual cell models are visible, but most AI-biology narratives still emphasize model architecture and target discovery. The underpriced layer is proprietary human perturbation data that ties mechanism to pharmacology.
Where the price sits today	The Human Cell Atlas first draft appeared across Nature Portfolio in 2024, Arc's Virtual Cell Atlas includes observational and perturbational data from hundreds of millions of cells, and Recursion's Roche/Genentech collaboration explicitly uses Maps of Biology and single-cell perturbation data.
The binding constraint	Human perturbation atlases in organoids, primary cells, disease-relevant tissue, spatial multi-omics, and longitudinal phenotypes linked to drug response.
What we are watching	Top-20 pharma or mega-biotech deals above $500 million that explicitly center proprietary human perturbation, spatial, or organoid atlases; FDA or EMA approval packages citing such platforms as material mechanism or pharmacology evidence.
What would prove us wrong	Kill if by 2033 fewer than two such major deals are disclosed and approval packages continue relying on conventional animal and in-vitro evidence without atlas-derived human causal evidence.
How we tried to break it	Pharma may build atlases internally, undercounting deal evidence, and regulators may keep platform evidence confidential. The call survives because AI lowers target-generation cost and makes proprietary human causal data the scarce differentiator.
Why we are making the call	After AI can suggest biology, the bottleneck is proving biology behaves that way in human systems.

If the call is right

Who is exposed

Pharma strategy teams, AI-biology startups, organoid companies, spatial-omics platforms, single-cell technology vendors, and biotech investors.

Action now

Treat causal perturbation data as the asset, not a support function: secure tissue access, assay rights, metadata standards, and longitudinal phenotype linkage.

Decision it changes

Business-development teams prioritize partnerships and acquisitions around human perturbation atlases rather than model architecture alone.

ROI / risk logic

A proprietary causal atlas can reduce wrong-target spend, improve translational confidence, and create licensing leverage even if frontier biology models commoditize.

Value lands at Recursion-like maps of biology, Arc/CZI-style atlas infrastructure, organoid and spatial-omics platforms, and pharma groups that own disease-specific perturbation data.

Who gains

Single-cell and spatial-omics platforms: They provide the measurement layer for causal human biology.

Organoid and primary-cell data owners: They supply human-relevant perturbation systems models cannot infer from text alone.

Pharma companies with proprietary disease atlases: They can test targets against internal causal maps before competitors see them.

Who loses

Model-only AI drug discovery vendors: Their hypotheses become cheap without exclusive human causal validation.

Animal-model-heavy programs in poor-translating diseases: They lose credibility as human perturbation evidence improves.

Undifferentiated public atlas mirrors: Open data helps the field but does not create proprietary decision advantage.

What reprices

Disease-specific atlas rights, tissue access agreements, spatial-omics instruments, organoid CRO capacity, and perturbation-screening contracts reprice upward.

The next constraint it creates

After causal atlases, the next bottleneck is clinical linkage: longitudinal outcomes and responder/non-responder data tied back to perturbation signatures.

Earliest sign it has begun

Watch pharma BD language for acquisitions or collaborations where the named asset is a human perturbation atlas, organoid atlas, spatial atlas, or virtual-cell training dataset.

P7

Grid stability replaces megawatts as the AI load gate.

Domain: energy / data centers / grid stability

2033-12-31

Structural case	Our call, dated	Resolves
68%	38%	2033-12-31

Data centers are not only large loads; AI workloads can create fast-changing power behavior through power electronics and operational controls. As capacity planning catches up, the constraint migrates from raw megawatts to dynamic models, ride-through, grid support, and interconnection obligations.

The boom	AI campuses become grid actors. The valuable infrastructure is the power-electronics, battery, model, and interconnection stack that lets a 100 MW load behave safely on the grid.
Why it is not priced yet	Power scarcity is already priced. The underpriced second-order constraint is stability: whether large AI loads can connect without creating voltage, frequency, oscillation, or protection problems.
Where the price sits today	IEA projects data-center electricity consumption to double to around 945 TWh by 2030. LBNL shows US interconnection queues remain huge and slow. NERC and DOE/PNNL are already studying emerging large-load risks and electromagnetic transient models for data centers.
The binding constraint	Dynamic load models, EMT-grade studies, ride-through commitments, grid-forming inverters, battery front ends, MVDC architectures, and utility tariff obligations for large AI loads.
What we are watching	Number of US RTO/ISO regions or NERC/FERC-linked reliability processes requiring large data-center loads to provide dynamic load models and ride-through or grid-support commitments; GW of AI-campus capacity publicly using BESS, grid-forming inverter, or MVDC front ends as interconnection enablers.
What would prove us wrong	Kill if by 2033 grid operators integrate large AI loads through ordinary substation and capacity studies with no special dynamic-load, EMT-model, ride-through, or grid-support requirements.
How we tried to break it	Utilities may solve this through conventional planning, local generation, and static tariffs. The call survives because NERC, DOE, and PNNL are already pointing at large-load dynamics and model gaps, not just energy supply.
Why we are making the call	The post-AI physical world has to couple intelligence factories to AC-grid physics. That coupling becomes a business.

If the call is right

Who is exposed

Hyperscalers, data-center developers, utilities, battery integrators, power-electronics vendors, grid planners, and infrastructure funds.

Action now

Add dynamic-load modeling, ride-through design, battery smoothing, and grid-support capabilities to site selection before committing to land or power contracts.

Decision it changes

Data-center capex moves from lowest-cost megawatts toward interconnection-ready power architecture and utility-approved operating behavior.

ROI / risk logic

A site that can connect quickly and satisfy stability requirements beats a cheaper stranded site. The timing advantage can be worth more than headline power price.

Value lands at battery integrators, power-electronics firms, electrical engineering consultants, utilities with mature large-load processes, and campuses that can provide grid services rather than merely consume power.

Who gains

Battery and grid-forming inverter vendors: They can turn fast AI-load swings into acceptable grid behavior.

Utilities and RTOs with transparent large-load rules: They attract serious AI campuses that can clear interconnection reviews.

Power-engineering firms with EMT and large-load expertise: They become gating consultants for AI-campus permits and studies.

Who loses

Data-center developers treating power as a commodity input: They get delayed when grid stability requirements harden.

Regions with opaque interconnection and weak grid models: They become harder to finance for large AI campuses.

Generation projects without deliverability or load-integration strategy: They do not solve the stability or connection bottleneck by themselves.

What reprices

Large-load interconnection studies, BESS front-end contracts, MV switchgear, power-electronics vendors, and utility tariff concessions reprice upward.

The next constraint it creates

After grid stability, the next bottleneck is who pays for shared upgrades and whether AI loads can monetize flexibility without shifting costs to ratepayers.

Earliest sign it has begun

Watch NERC, FERC, ERCOT, PJM, and utility filings for mandatory dynamic models, ride-through rules, or BESS/grid-support requirements for data centers above 100 MW.

P8

AI campuses become heat-sink and water-rights assets.

Domain: data-center real estate / cooling / water / heat

2032-12-31

Structural case	Our call, dated	Resolves
72%	41%	2032-12-31

AI campuses turn electricity into heat. As power gets solved site by site, the next physical gate becomes where the entropy and cooling burden go: water access, non-potable supply, thermal discharge, district heat, and local acceptance.

The boom	The best AI real estate is not only cheap land with power. It is a site with permitted heat rejection, reclaimed water, and a credible offtake for waste heat.
Why it is not priced yet	Power is now consensus. Cooling and water are discussed, but heat-sink rights are not yet priced as a primary siting asset across AI campuses.
Where the price sits today	The European Commission now maintains reporting for energy performance and water footprint of significant data centers. Germany's data-center energy-efficiency rules move toward waste-heat reuse requirements for new sites starting in 2026. Current EU reporting shows energy reuse is still limited, which leaves room for regulatory pressure.
The binding constraint	Legal and physical heat rejection: district-heat offtake, non-potable or reclaimed cooling water, thermal discharge permits, closed-loop cooling, and local consent.
What we are watching	Number of OECD jurisdictions requiring new large data centers to reuse material waste heat, use non-potable or closed-loop cooling, or prove infeasibility; share of AI-campus announcements citing heat sink, WUE, reclaimed water, or district-heat offtake as a siting reason.
What would prove us wrong	Kill if by 2032 rules remain disclosure-only and hyperscale AI campuses continue scaling in water- and heat-stressed regions without heat or water constraints materially changing site selection.
How we tried to break it	Waste heat is low-grade and often uneconomic; many campuses can use dry cooling, tolerate higher server temperatures, or site remotely. The call survives because local politics and cooling physics can bind before global energy numbers look large.
Why we are making the call	After AI gets power, it still has to dump heat. In a post-AI buildout, thermodynamic permission becomes real estate value.

If the call is right

Who is exposed

Data-center developers, hyperscalers, district-energy utilities, municipal planners, water authorities, and infrastructure real-estate investors.

Action now

Screen AI-campus sites for reclaimed-water access, heat offtake, district-energy adjacency, dry-cooling feasibility, thermal permits, and community opposition before land banking.

Decision it changes

Site acquisition and power procurement decisions add heat/water rights as first-order diligence, not late-stage sustainability checks.

ROI / risk logic

A campus with permitted cooling and waste-heat offtake can avoid delays, community fights, retrofits, and future water or heat-reuse mandates.

Value lands at industrial parks with non-potable water, district-heating networks, cold-climate or heat-demand municipalities, and cooling technology vendors that reduce water exposure.

Who gains

District heating operators and municipalities with heat demand: They can turn AI waste heat into negotiated offtake and permitting leverage.

Data-center sites with reclaimed-water rights: They avoid local freshwater conflict and regulatory delays.

Closed-loop and liquid-cooling vendors: They become permitting tools, not just efficiency upgrades.

Who loses

Hot, water-stressed data-center markets without reclaimed supply: They face local resistance and stricter cooling conditions.

Developers that secure power before cooling permits: They can strand land or interconnection rights on thermal and water constraints.

Municipalities giving away heat-sink rights cheaply: They may underprice a scarce permitting asset.

What reprices

Land near district heating, reclaimed-water infrastructure, heat-reuse contracts, water permits, cooling systems, and environmental approvals reprice upward.

The next constraint it creates

After heat and water, the next bottleneck is local political consent: whether communities accept AI campuses as infrastructure or treat them as extractive load.

Earliest sign it has begun

Watch >100 MW AI-campus announcements for WUE targets, reclaimed-water agreements, district-heat offtake, and local moratoria tied to water or heat.

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
Physical AI safety cases become the robot bottleneck	The 2025 industrial robot safety-standard refresh and global robot installation growth suggest a certification bottleneck for human-adjacent robots.	The earlier after-AI board already covered certified deployment as the physical-AI bottleneck, so this pass demoted it to avoid duplication.
Machine edge zones become a new real-estate class	Physical AI may require low-latency inference near factories, ports, hospitals, and grids.	The clause depends on uncertain latency needs and may be weakened by on-device inference; kept below heat/water and grid-stability calls.
AI insurance splits fully from cyber	AI-specific exclusions and endorsements are likely as loss data accumulates.	This is partially contained inside P2 liability evidence and event recorders.
Inline metrology outruns robot synthesis	Autonomous science may bottleneck on calibrated in-situ measurement rather than robot arms.	Strong but nested under P3 autonomous-lab standards and P4 qualification data.

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