# Purchasing Time: The Hidden Economics of AI Infrastructure

Navigating the widening gap between computing demand, energy timelines, and financial risk.

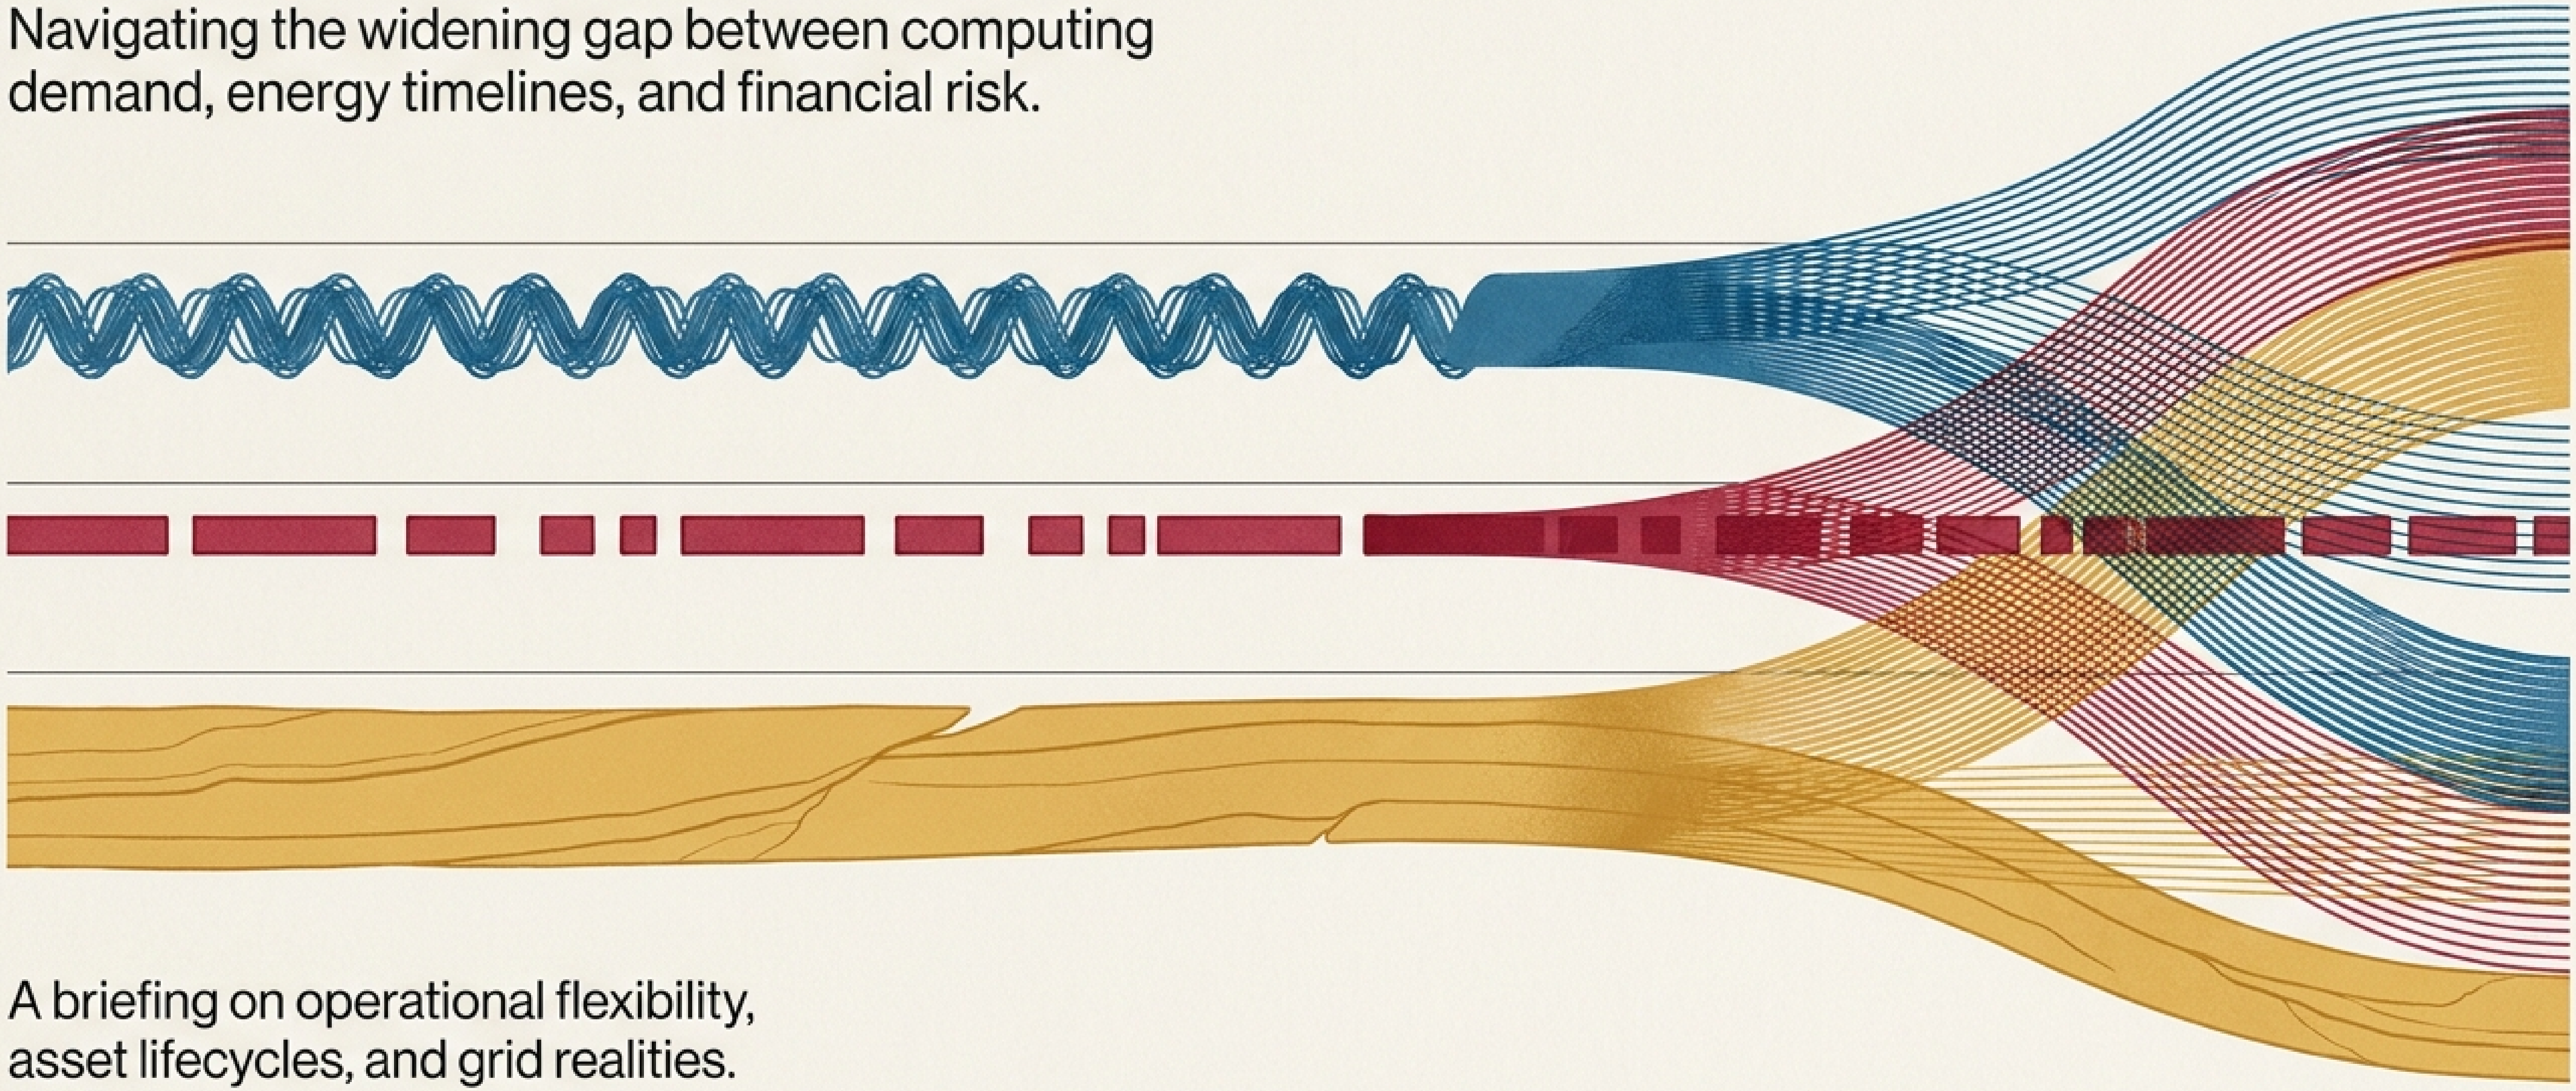

A briefing on operational flexibility, asset lifecycles, and grid realities.

# The Time Dimension in AI Infrastructure

## The Scheduling Gap

A model team ships in months. A chip buyer upgrades in years. A utility expands a substation over decades. AI companies are not just paying for power; they are paying premiums to align these disjointed schedules.

## The Limits of Flexibility

Intermittent compute is a myth. A patient user waiting for cheap, off-peak power does not make hardware cheap; idle time destroys the lifetime output of depreciating GPUs.

## The Lifecycle Risk

Power plants and transmission lines operate for decades. The economic competitiveness of a GPU changes rapidly. Securitization and PPAs do not make renewal risk disappear.

# The Timeline Collision

A developer cannot get a new data center connected on the schedule the business needs. Schedules across the tech and utility sectors are fundamentally mismatched.

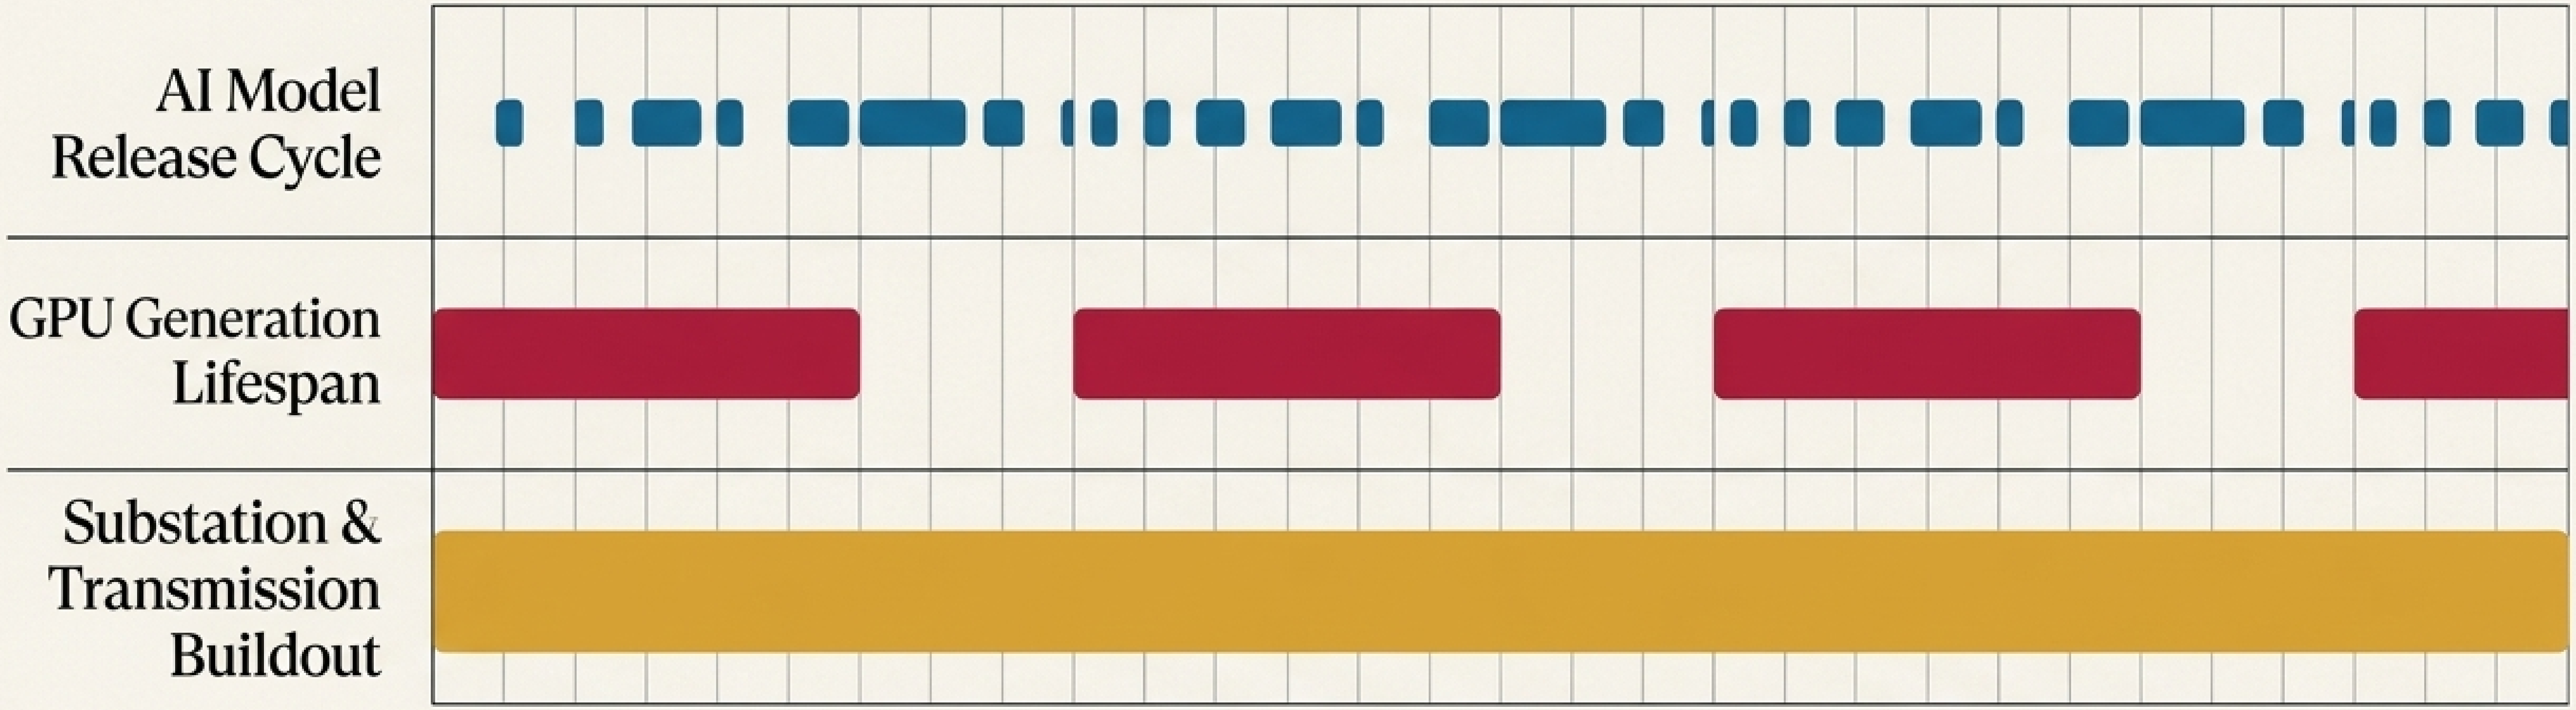


Denominator Callout

Paying more for an earlier connection is economically sensible. The comparison includes revenue earned sooner, less idle equipment, and the risk that hardware becomes obsolete while waiting in the queue.

# Precision in Scale: Capacity vs. Consumption

The IEA projects global data-center electricity consumption will reach ~950 TWh by 2030 (up from 485 TWh in 2025). This represents just ~3% of global electricity demand.

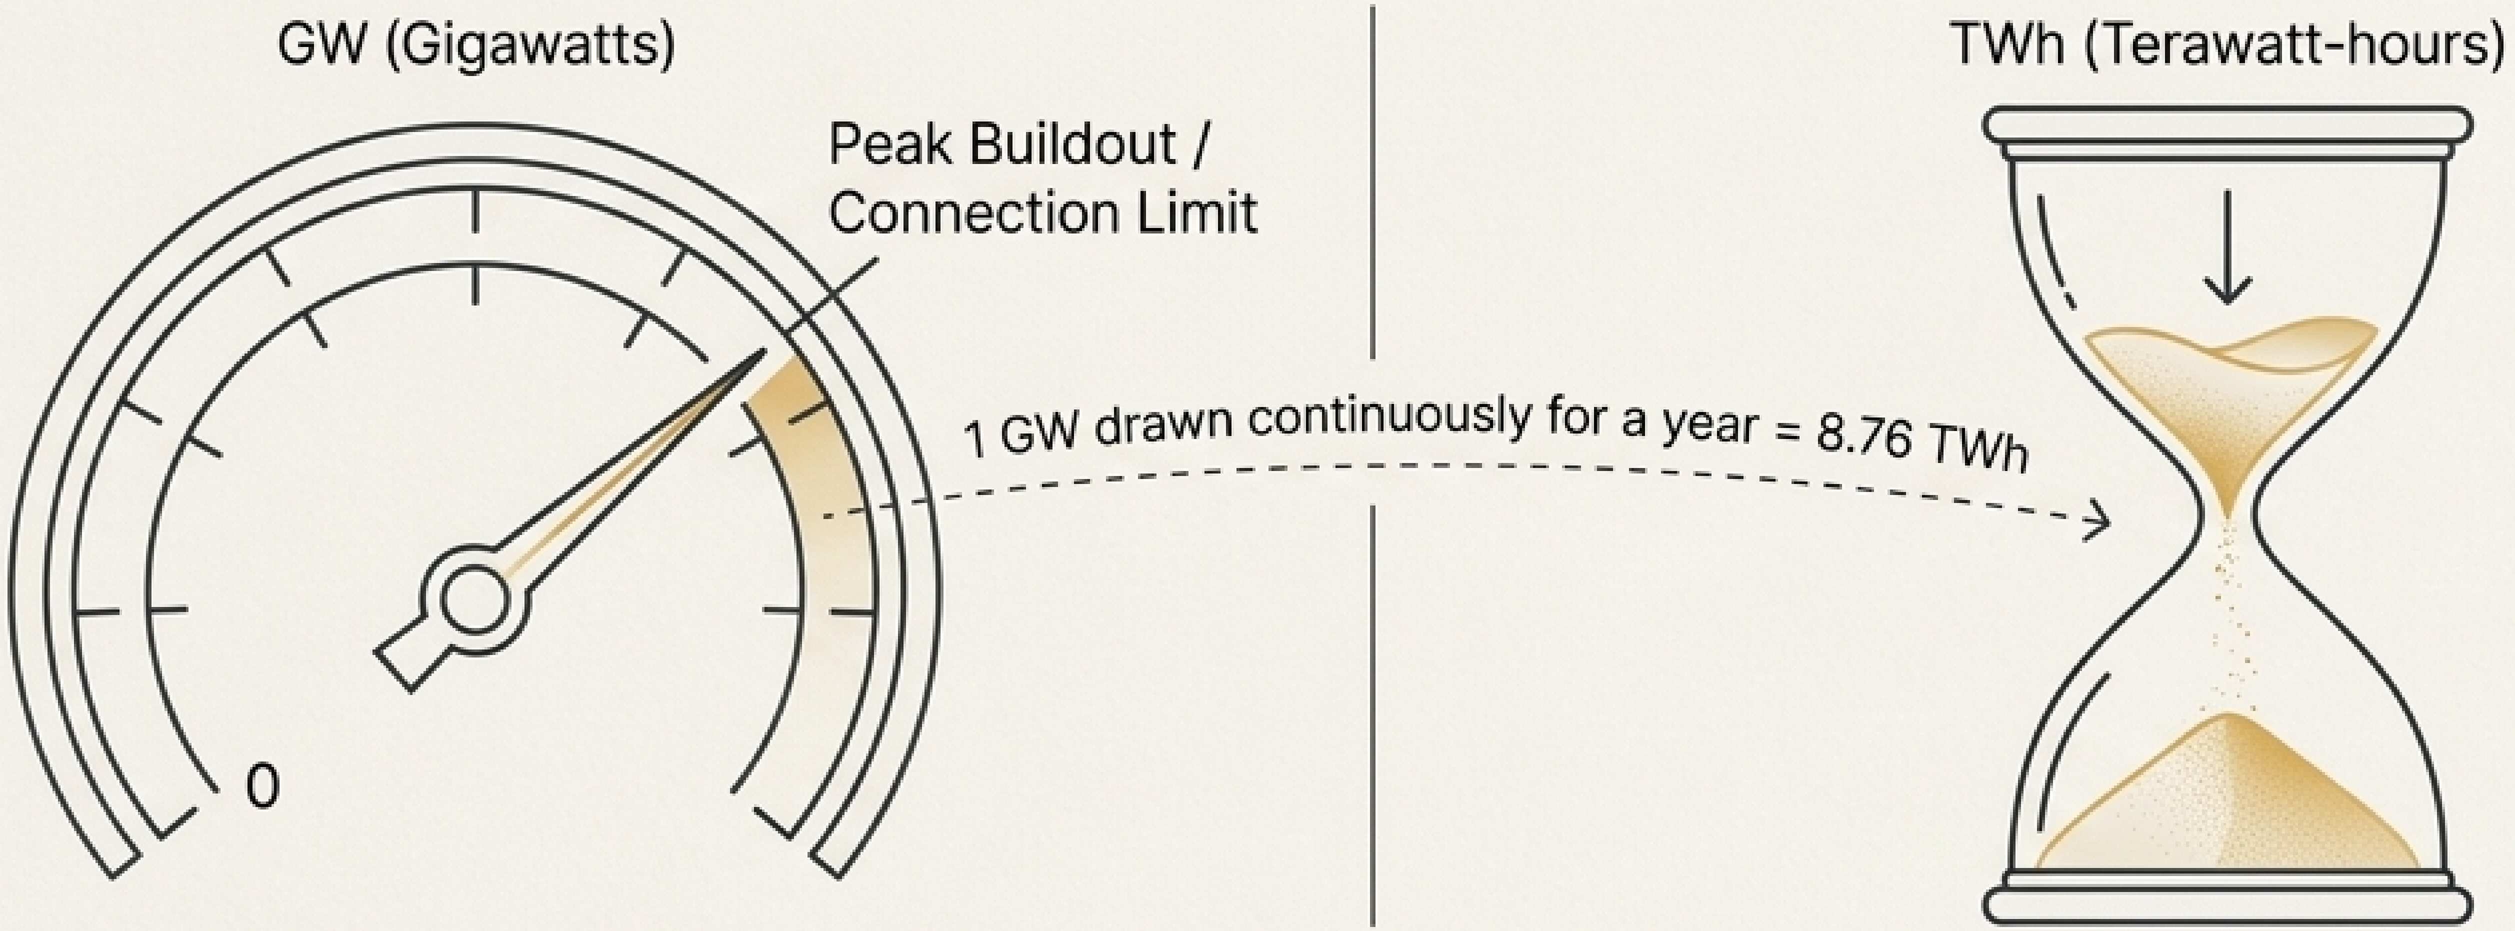


America is not running out of electricity. A modest global share coexists with severe local constraints. Permitting, generation availability, transmission, and equipment supply prevent large new loads from being delivered reliably to specific locations by developer deadlines.

# The Idle GPU Economic Trap

If an operator waits for cheap electricity, depreciation continues while machines sit idle. This reduces the total work the equipment delivers over its useful life, driving up the fixed cost per completed task.

Hardware Depreciation + Facilities + Networking & Financing + Electricity = Total Cost of Output

When podcasts or reports claim "Electricity is 60–80% of operating costs," we must ask: Operating costs of what? A facility operator's day-to-day expenses? Or a GPU service including hardware depreciation? Without a clear denominator, the percentage cannot establish how much cheaper the final AI service becomes when power prices fall.

# The Limits of Operational Flexibility

A 2025 field demonstration proved a 256-GPU cluster could reduce power by 25% for three hours. But small, bounded demonstrations do not prove large campuses can economically follow the weather all year.

| OPERATIONAL CHOICE | POTENTIAL BENEFIT | HIDDEN CONSTRAINTS |
|---|---|---|
| Run entire job in another region | Use available power | Data movement, jurisdiction limits, latency. |
| Curtail at peak hours and catch up later | Relieve grid stress | Delays, recovery overhead, requires spare capacity to catch up. |
| Remain idle awaiting cheap power | Closely follow generation | Lost lifetime output, higher fixed cost per task. |

# Designing for Asynchrony

Calling a product asynchronous does not magically give infrastructure operators scheduling freedom.

An AI **agent service** must **actively** distinguish between “respond now” and “deliver by morning.” True flexibility requires software designed with reliable checkpoints and recovery overhead so jobs can be safely interrupted and resumed without data loss.

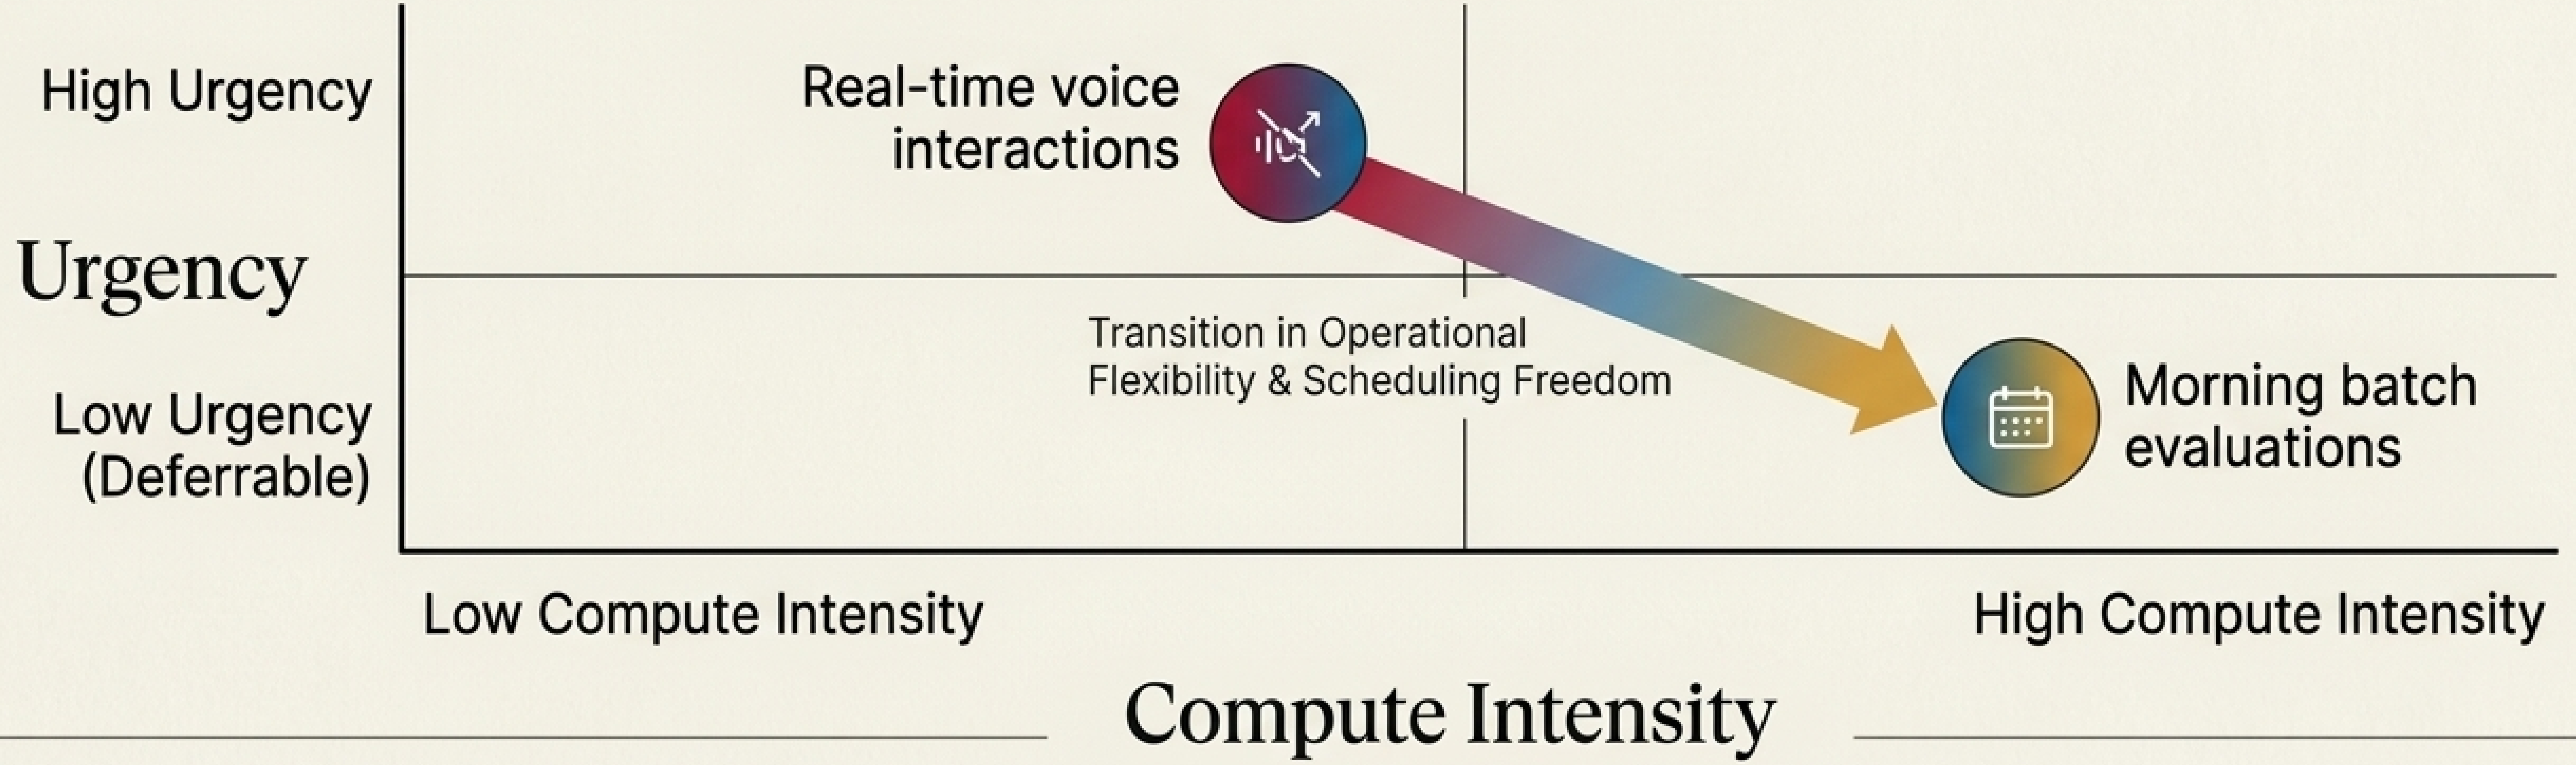

# The Tripartite Incentive Clash

Based on dynamics in the Chinese market, these interests overlap but are not identical. A location can obtain construction investment without actually gaining downstream tech companies.

Cheap land and abundant wind do not automatically produce round-the-clock reliability. Even with a direct renewable connection, someone must pay for storage, grid balancing, and backup generation when the wind stops. A dedicated line does not provide free reliability.

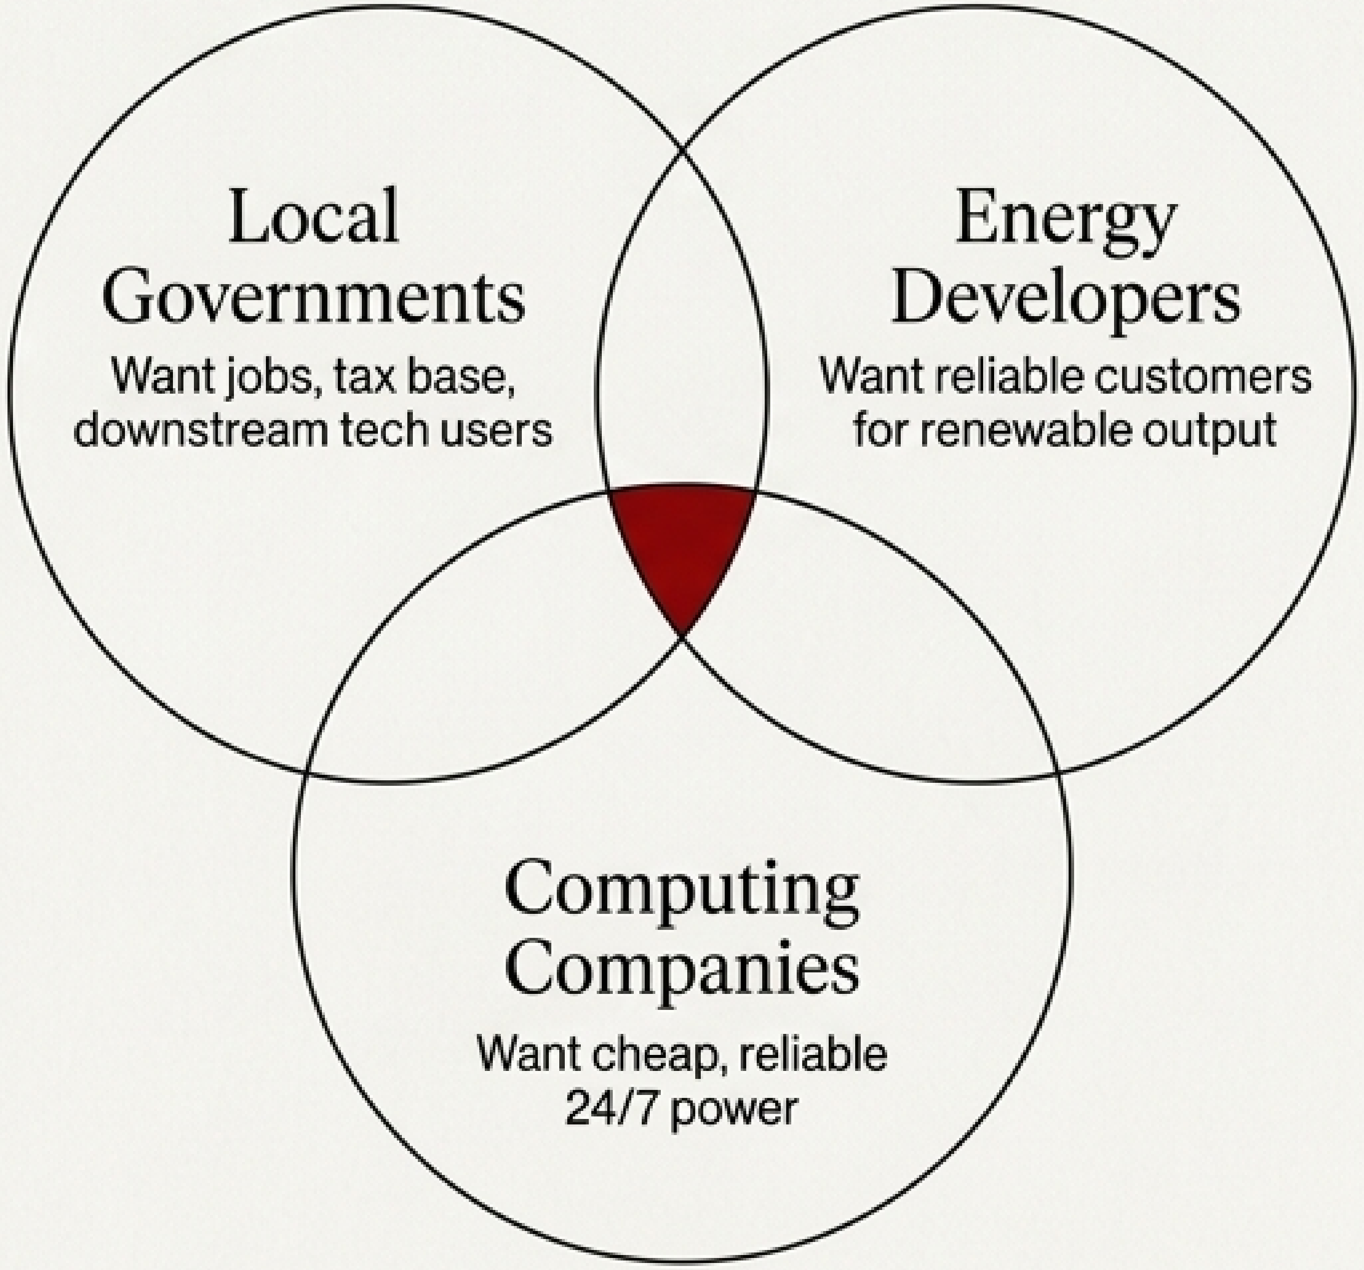

# The Timing Question of Clean Energy

**Denominator Callout**

Procuring annual renewable certificates does not establish that clean generation was available in the grid during every hour the company operated. Google made this distinction explicit in 2020: annual matching is easy; carbon-free energy every hour is an entirely different engineering challenge. A low-cost annual certificate tells us nothing about the actual cost of reliable, hourly clean supply. Delivering generation at the required time and place is the true constraint.

## Annual Matching (The Illusion)

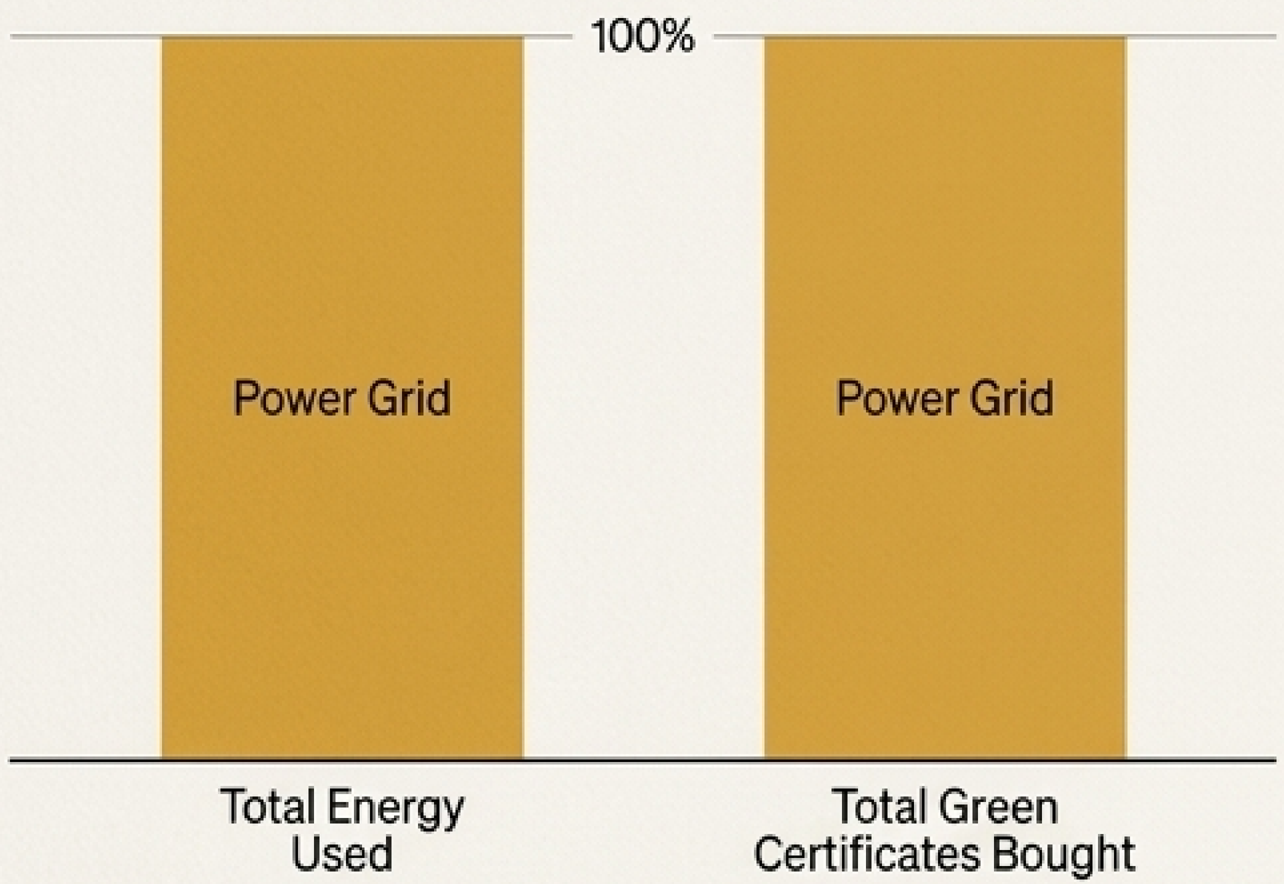


Appears balanced on paper, but masks temporal discrepancies.

## Hourly Matching (The Reality)

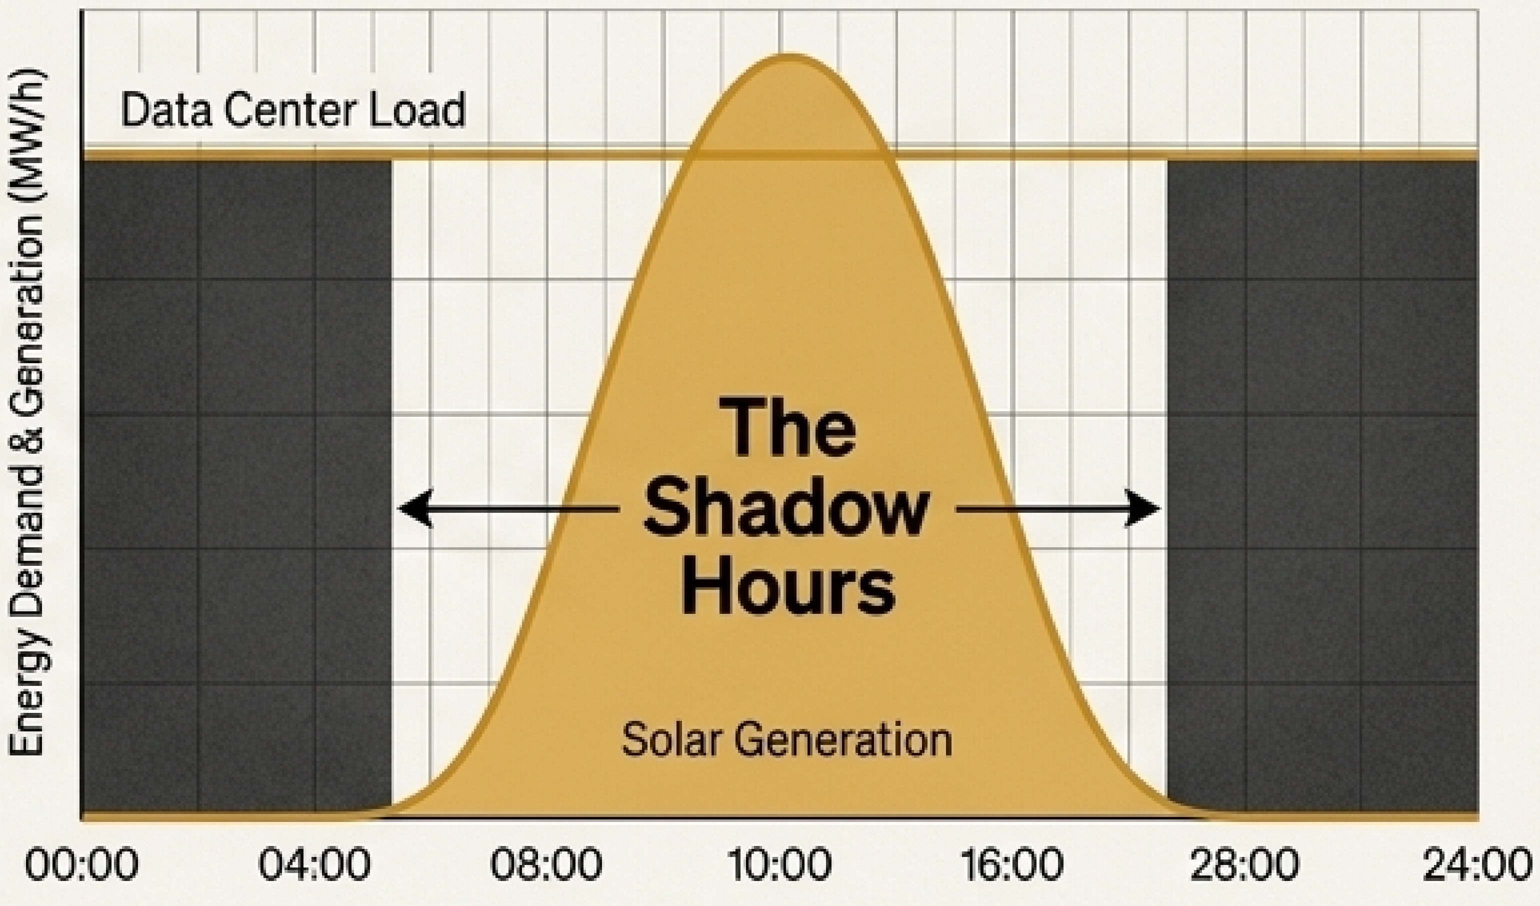


Illustrates the physical reality of 24/7 power requirements and the clean energy deficit at night.

# The Asset Lifecycle Mismatch

Putting different obligations into one project does not automatically align their financial realities. A power plant can operate for decades; the economic competitiveness of a GPU changes in months.

Securitization dictates who holds the rights to cash flow, but it does not make renewal risk, equipment suitability, or long-term electricity-price exposure vanish.

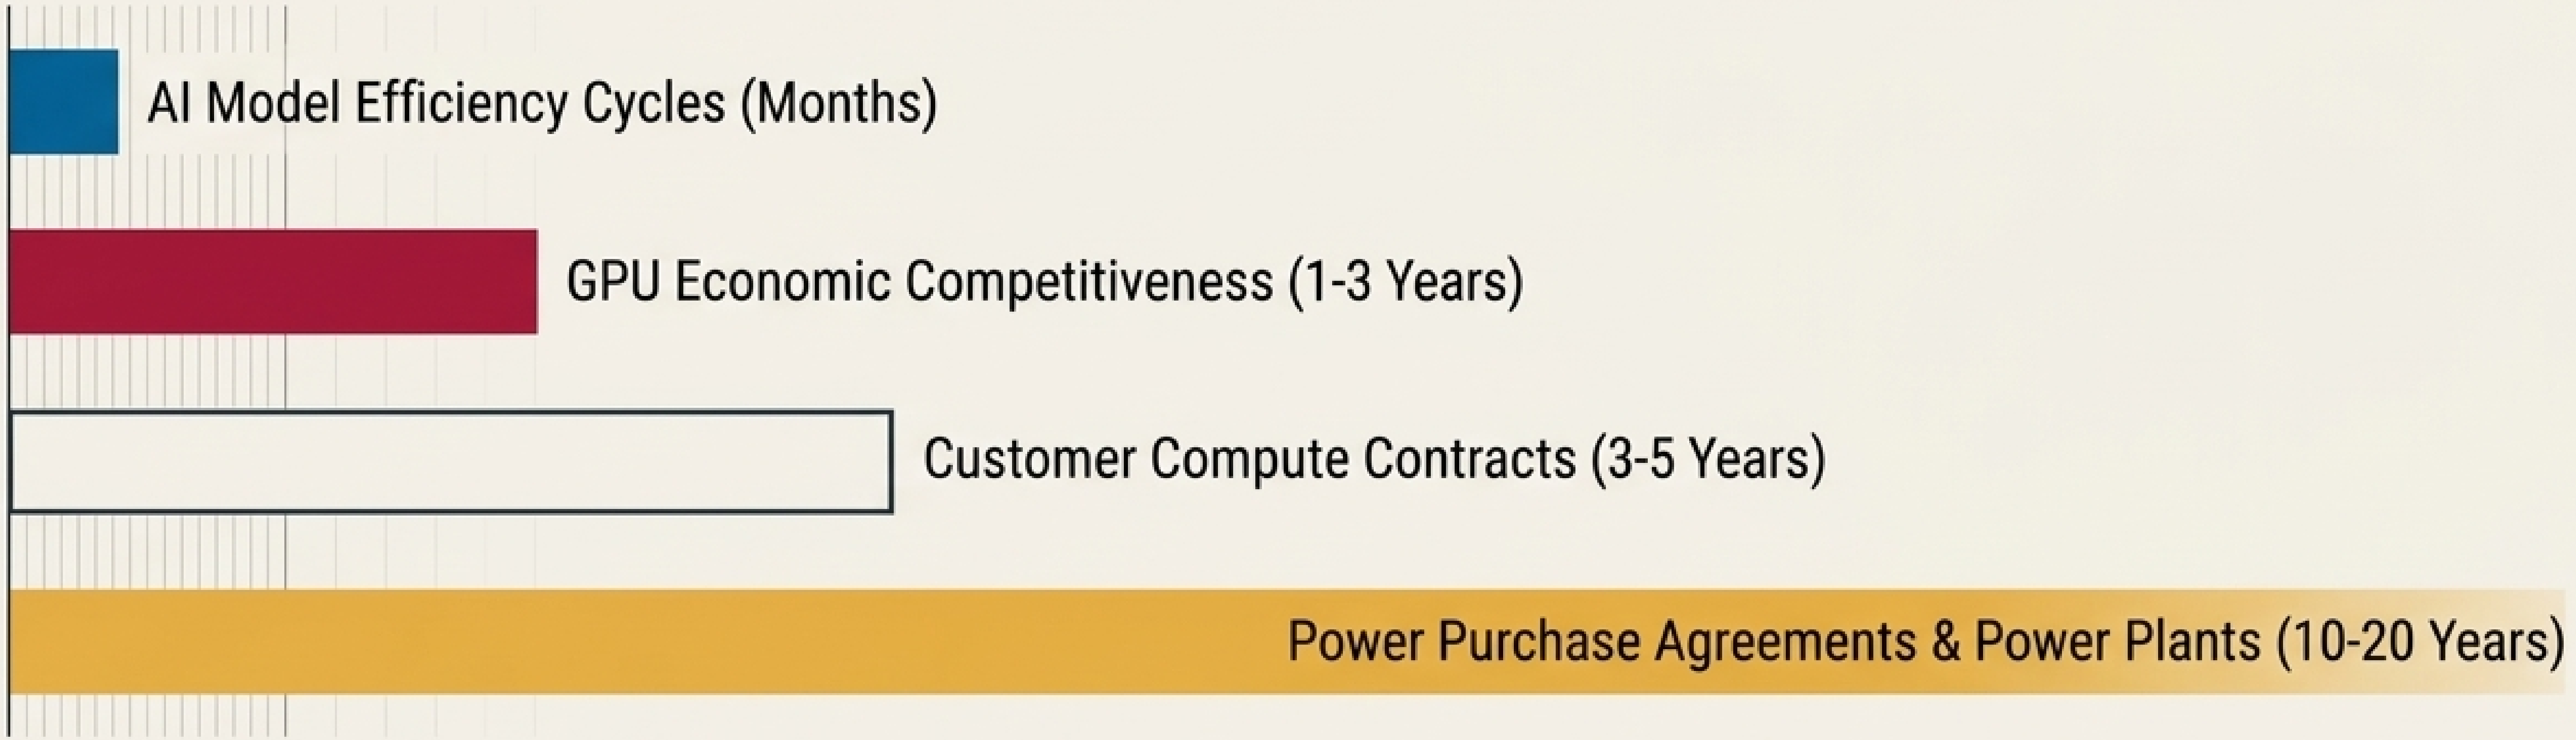

# The Queue vs. The Contract

We must separate the asset holding the building/electrical infrastructure from the business buying GPUs and selling computation. Their revenue and risk profiles are fundamentally different.

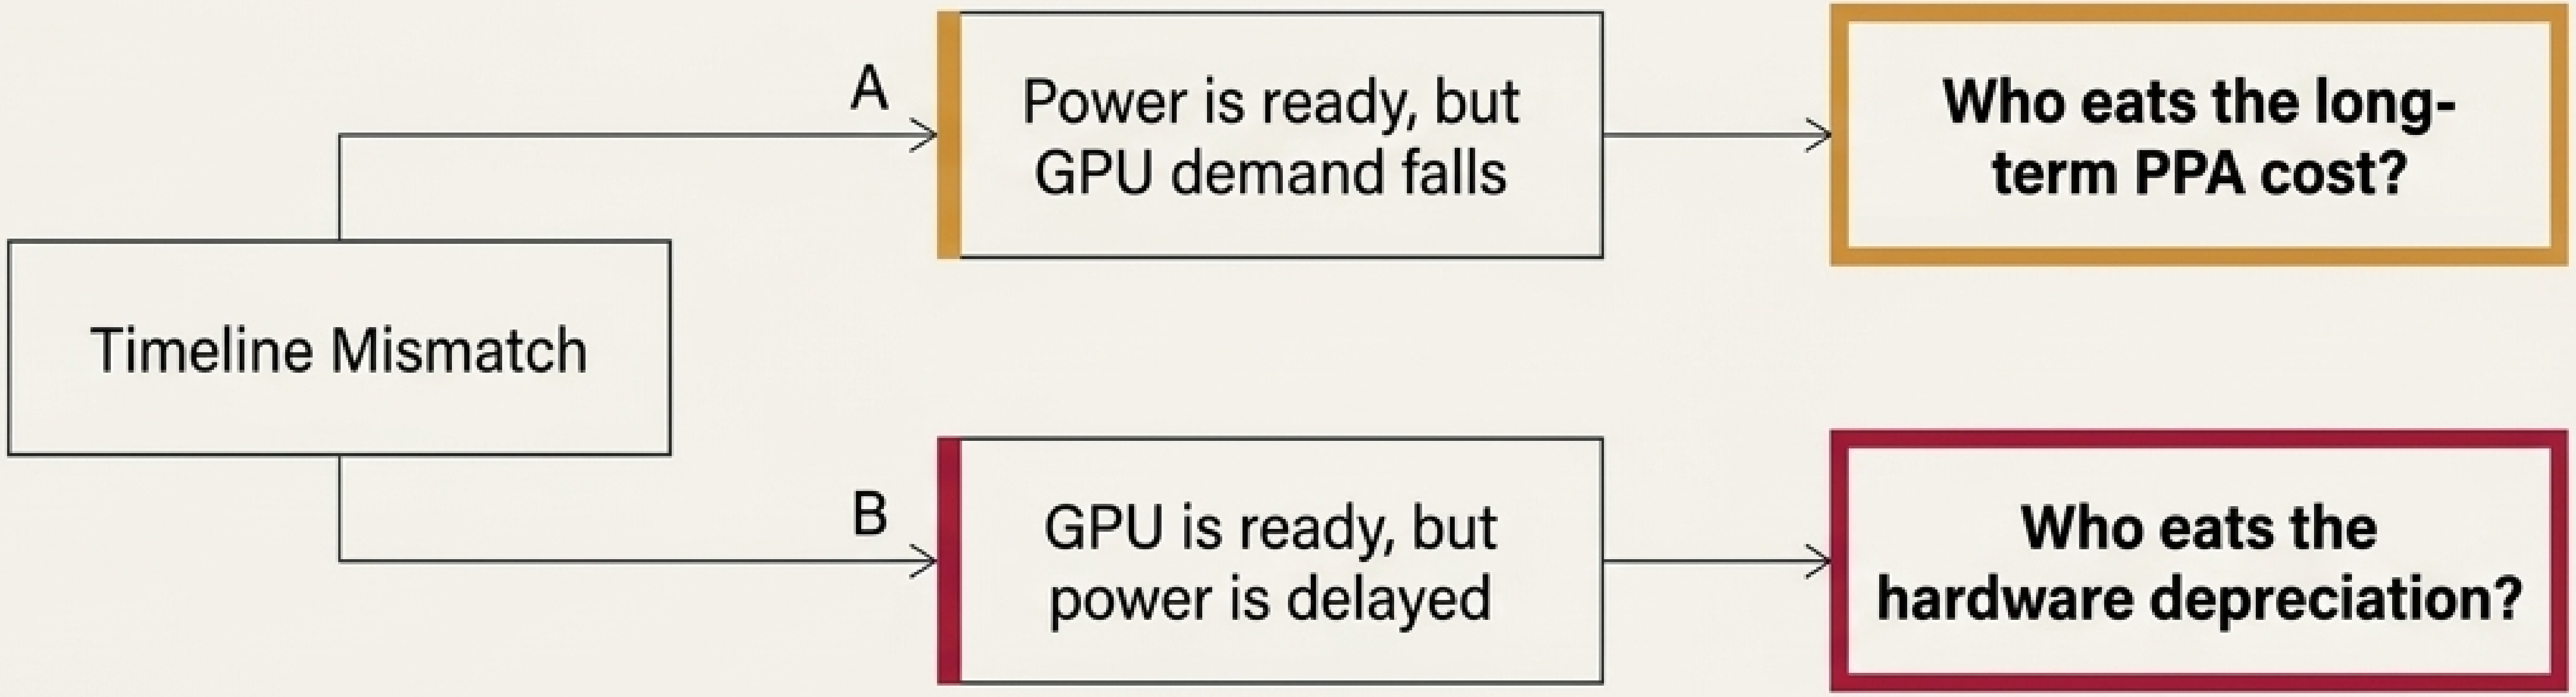


**Denominator Callout**

Advance capacity auctions (like PJM's) procure availability for a specified future delivery period. A capacity price increase is not an immediate per-kWh retail energy price shock hitting household bills today.

# Strategic Diagnostics for Infrastructure Investment

- [ ] Whose project is this asset holding up in the queue?
- [ ] How much will the customer pay to receive power earlier?
- [ ] How long is the revenue contractually supported?
- [ ] Who takes the loss if expected downstream demand fails to arrive?

---

**An industry needing an asset does not establish its owner's return. Scarcity produces higher revenue, but it also produces expensive, rigid construction obligations whose returns remain constrained by time.**