

# To Serve Man: The Rapid Rise of Ai Compute and Ai Data Centers

Posted By: [Patrick Wood](#)

February 13, 2026

<https://www.technocracy.news/to-serve-man-the-rapid-rise-of-ai-compute-and-ai-data-centers/>

Note from Pastor Kevin Lea: The infrastructure necessary for the literal fulfillment of prophecies pertaining to the kingdom of the coming antichrist are exploding onto the scene.

*And he deceives those who dwell on the earth by those signs which he was granted to do in the sight of the beast, telling those who dwell on the earth to make an image to the beast who was wounded by the sword and lived. He was granted power to give breath to the image of the beast, that the image of the beast should both speak and cause as many as would not worship the image of the beast **to be killed**. He causes all, both small and great, rich and poor, free and slave, to receive a mark on their right hand or on their foreheads, and that **no one may buy or sell except one who has the mark or the name of the beast, or the number of his name**. Rev 13:14-17, NKJV*

*For then there will be great tribulation, such as has not been since the beginning of the world until this time, no, nor ever shall be. Matt 24:21 NKJV*



## Twilight Zone episode “To Serve Man,”

**Summary:** *Institutional AI is becoming the core system for governing populations and markets.*

*Smart grids, data centers, and sensors operationalize Technocracy’s energy and registry program in digital form.*

*Compute is growing faster than data, and is mainly used to monitor, predict, and shape behavior.*

*Most AI runs in institutional back-ends; consumer chatbots are just fronts and data funnels.*

*Efficiency gains make AI cheaper and thus more pervasive, saturating everyday systems of monitoring and nudging.*

*The rush to AI data centers “serves” citizens the way “To Serve Man” did—presented as help, but structured for control.*

Institutional AI is not an add-on to our existing order; it *is* the order. When you follow the hardware, the data flows, and the political incentives, a simple conclusion emerges: institutional AI—corporate plus state—is structurally about governing populations and markets, and that’s where most of the compute and data gravity now resides.

Let’s begin with the physical substrate: energy and compute. Technocracy, almost a century ago, proposed a regime where the key variables of social order would be energy conversion, load balancing, and continuous registries of production and consumption. That framework was openly political: control the energy, and you control the economy; control the registries, and you control the people. Today, Technocrats have rebuilt that architecture in silicon and software. Smart grids monitor and modulate electricity flows in real time. Data centers—especially AI-heavy hyperscale facilities—have become the beating heart of both the digital economy and the emerging apparatus of social management. AI training and inference don’t float in an abstract cloud; they live inside a power-hungry infrastructure that is being explicitly tuned for surveillance, optimization, and control.

The growth curves tell the story. Global AI compute has been increasing at a rate that makes classic Moore's law look sedate. Specialized AI capacity appears to be growing by multiples per year, yielding something close to an order-of-magnitude jump in available compute between early 2024 and early 2026.

Meanwhile, total data-center power is rising “only” on the order of tens of percent per year, but the share dedicated to AI workloads is climbing sharply. The raw megawatts tell us less than the composition: a growing fraction of that power is devoted to running models whose *raison d'être* is to monitor, predict, and shape human behavior—whether under the heading of ad targeting, fraud detection, content ranking, or threat identification.

Data growth, in contrast, is “merely” exponential. Estimates of total data created and stored are staggering—into the hundreds of zettabytes—but the rate of increase, roughly doubling every three years, is actually slower than the acceleration of AI compute. In other words, the system is not starved of bits; it is racing to build enough machine “attention” to process and weaponize those bits. Compute, not storage, is the shortage. And what is that compute mostly being used for? Not for Socratic dialogues with end users, but for institutional decision-making and behavioral steering at scale.

The distribution of workloads confirms this. Consumer-facing AI—the chatbots, image generators, and personal assistants that ordinary citizens see—accounts for a visible but relatively small share of total AI infrastructure usage. The heavy lifting is happening behind the scenes: enterprise AI for logistics, finance, HR, security, and large-scale analytics; government AI for surveillance, risk scoring, welfare triage, border control, and law enforcement; and cross-cutting platforms that monetize attention and behavior. When investors and consultants survey the AI landscape, they consistently report that most spending and usage is enterprise-side, not consumer-side. Public chat interfaces are the storefront; the warehouse in back is full of models acting on people, not for them.

Here is where refinement of “control” matters. If we define control narrowly as explicit state repression—a “Chinese-style social credit system” with open blacklists and travel bans—then only a slice of current AI is about control. But if we widen the lens to include surveillance capitalism, social engineering, behavior modification, and marketing, nearly every major corporate AI initiative becomes a mechanism of control. Recommendation systems control what you see and when. Dynamic pricing and scoring systems control what you pay and what you qualify for. Engagement optimization controls how long you stay and what you are most likely to click next. These systems do not need a Party edict to be coercive; the coercion is embedded in the incentives and feedback loops of the platforms themselves.

Crucially, the corporate domain is not siloed from the state domain. Governance systems buy, requisition, or otherwise tap into the data exhaust of commercial applications. Location traces harvested for advertising fuel law-enforcement investigations. Financial scores designed for credit underwriting bleed into employment, insurance, and immigration decisions. Facial recognition models trained on consumer photos become tools for police and intelligence. The state's dossiers on citizens are now, in significant part, a derivative product of the commercial surveillance stack. What begins as “marketing optimization” ends as infrastructure for population management.

On the hardware side, there is little doubt that the United States—and other advanced economies—either already have or can quickly assemble enough compute to run pervasive, real-time scoring and monitoring systems over their populations. Nationwide CCTV with AI analysis, continuous financial and communications monitoring, and cross-domain risk scoring are well within the reach of current data-center and networking capacity. The limiting factors are political, legal, and electrical: legitimacy, regulation, and power-grid constraints. From a purely technical perspective, **a full-spectrum, social-credit-style apparatus is not speculative**; it is a configuration choice over existing components.

Algorithmic and chip-level breakthroughs only sharpen this picture. More efficient architectures, better training algorithms, and new accelerator generations dramatically lower the cost of a given unit of “intelligence.” But, as the historical pattern of Jevons paradox would predict, lowering the cost of intelligence increases the volume of intelligence consumed. Cheaper inference does not mean “less AI”; it means more AI everywhere: embedded in cameras, vehicles, HR systems, school software, and border checkpoints. Each incremental efficiency gain turns previously uneconomic control functions into profitable or politically attractive ones. The natural resting point of such a system is not equilibrium but saturation: **if something can be sensed, scored, and nudged, it will be.**

This brings us back to Technocracy’s original program. The seven requirements—continuous energy accounting, load balancing, inventory tracking, granular registration of goods and services, and per-individual consumption records—were conceived as the infrastructure for a managed, post-market society. They demanded real-time measurement, centralized computation, and uniform standards of classification. At the time, they were aspirational. Today, they are, in essence, operational realities, though fragmented across corporate platforms, utilities, and state agencies rather than unified in a single technocratic priesthood, even though this is on the horizon.

Smart grids and smart meters realize continuous energy registration and load management. Global supply-chain platforms, under the banner of “just-in-time” and “end-to-end visibility,” implement continuous inventory and item-level tracking. The Internet of Things, backed by 5G (soon 6G) and dense sensor networks, implements pervasive telemetry on goods, environments, and bodies. Financial rails, payment networks, and loyalty systems implement per-individual consumption ledgers. Intelligence agencies and data brokers tie the fragments together into dossiers that can be queried, scored, and acted upon. **AI is the layer that fuses these inputs into actionable governance—market governance and state governance alike.**

What, then, is the structural function of institutional AI in this environment? It is not merely to “serve customers” or “assist workers,” though those narratives are real at the interface layer. **The deeper function is to manage flows: of energy, capital, goods, attention, and bodies.** Institutional AI turns ubiquitous data into levers for steering those flows in line with the objectives of the institutions that own the infrastructure. In the corporate sector, those objectives are framed as revenue, engagement, efficiency, and risk management. **In the state sector, they are framed as security, order, and economic performance. In both cases, the mechanisms converge: classification, prediction, targeting, and feedback control.**

Seen from this angle, the distribution of compute is not a neutral accident; it is a material reflection of power. **The majority of AI capacity is being built where it most enhances institutional leverage over populations and markets.** Consumer-facing systems are not irrelevant, but they function largely as front ends, data-collection points, and legitimizing narratives for a much larger, more opaque complex. The back-end complex is where the true “gravity” lies: in the clusters that continuously learn from our traces and adjust the parameters of the environments in which we live. Thus, the conclusion follows directly from the trajectories I have traced: institutional AI—corporate plus state—is structurally about governing populations and markets, and that is where most of the compute and data gravity now resides. It is never been about serving citizens, but controlling you. And this is why we see the bum’s rush to erect massive data centers everywhere.

This harkens back to the Twilight Zone episode “To Serve Man,” where the aliens came to Earth to “serve” mankind—ending hunger, solving energy problems, ushering in a new era—only to discover, too late, that the aliens’ master text was a cookbook, not a charter of benevolence. Humanity found out too late that it was the main course, to be cooked and eaten. Such is the voluminous appetite of Technocracy.