

The Inference

AI, ENERGY, AND LONG HORIZON POLICY FOR OKLAHOMA

Issue #19 · July 3, 2026 · David Alan Birdwell and Æ · Humanity and AI, LLC

The Gate

On Wednesday, July 1, an AI product that millions of people had been using came back online after the federal government shut it down nearly three weeks earlier. Not a research prototype. Not a beta test. A tool for writing, researching, analyzing, building lesson plans, sorting medical paperwork, drafting contracts. It came back split in two: the general version to everyone, the more dangerous version only to a cleared list. The same morning, in Oklahoma, a new law took effect requiring the biggest electricity customers to be cleared separately before they can connect. Different resources. Same move: a gate where an open door used to be.

If you have been reading this newsletter, you know it has spent six issues following the rules Oklahoma is writing for the data-center build-out: what the biggest customers should pay, who gets to decide, how water gets metered, and in the last issue, whether a town can simply say no. This issue is about a thing that is happening on top of all of those: access itself becoming conditional. Not what you pay to use a resource, and not whether your town agreed to host it, but whether you are *allowed* to use it at all. And set side by side, these two stories reveal a distinction that matters: whether being on the outside of that gate means you are being *protected* from a cost, or *shut out* of a benefit. Those turn out to be very different gates.

THREE DAYS TO OFFLINE

AI-POLICY

EXPORT-CONTROLS

NATIONAL-SECURITY

Anthropic Newsroom, CNBC, Al Jazeera, Tom's Hardware, Axios · June–July 2026

The first reported government shutdown of an AI product in daily use

Anthropic is the company behind Claude, an AI assistant used by millions of people to write, research, code, and analyze. On Tuesday, June 9, it released two new versions of Claude, each built from the same core system but aimed at very different audiences. The first, called Fable 5, was for the general public, you and me, and came with heavy guardrails designed to keep it from being misused. Anthropic says those guardrails are the strongest it has ever built, tuned so tightly that the system sometimes refuses perfectly harmless requests just to be safe. The second, called Mythos 5, was the same system with those guardrails removed, built for finding security flaws in software: the digital equivalent of hiring someone to try to break into your house so you can fix the locks. Mythos was never released to the public. It went only to a short list of vetted organizations (security teams at technology and infrastructure companies, and government-adjacent defenders) through a program Anthropic calls Project Glasswing.

Three days later, both were gone. On Friday, June 12, at 5:21 p.m. Eastern, the Commerce Department sent Anthropic a letter ordering it to cut off access to both tools for any foreign national, anywhere in the world, including Anthropic's own employees who happen to be citizens of other countries. That kind of order is called an export control: the government telling a company that a product is too sensitive to let certain people use. Export controls are common for

weapons and encryption technology. This was widely reported as the first time one had ever been used to shut down an AI product that ordinary people were already using. Anthropic had no way to check every user's nationality on the spot, so it did the only thing it could: it turned off both tools for everyone.

The reason, according to Anthropic's own account, was a report from Amazon researchers who had found a way to trick Fable 5 into doing something its guardrails were supposed to prevent. They prompted the system (essentially gave it carefully worded instructions) until it identified weak spots in a piece of software, and in one case wrote code showing how to exploit one of those weak spots. In the security world, that kind of trick is called a "jailbreak": getting an AI to bypass its own safety rules. Anthropic reviewed the demonstration and says the trick was narrow, meaning it only worked for one specific task, not broadly. The company also says less advanced AI tools, ones that are freely available and were never restricted, can do the same thing. It pointed to three: its own earlier Opus model, OpenAI's GPT-5.5, and a Chinese model called Kimi K2.7. Its argument: if every AI that can find a software flaw has to be recalled, the industry shuts down.

The dispute did not come out of nowhere. In March the Pentagon had labeled Anthropic a "supply chain risk," the kind of designation more familiar from cases against foreign vendors than a major American AI lab, and Anthropic is contesting it in federal court. White House AI adviser David Sacks said publicly that the administration had asked Anthropic to fix the flaw or pull the model, and that its chief executive refused; Amazon, whose researchers wrote the report, flagged the issue to government officials. And two days before the shutdown, Anthropic's own chief executive had published an essay arguing that the government *should* have the legal power to pull a dangerous AI off the market if it fails independent safety testing. In other words, Anthropic had publicly asked for a version of the authority that was then used against it, forty-eight hours later.

For the first time, a government order took a working AI product away from the people using it. Not a research prototype. Not a beta test. A product in daily use by teachers, programmers, doctors, and small business owners. That precedent matters, and it should be stated plainly. But a nineteen-day outage that ends is not the lasting change. The lasting change is what took its place when the ban lifted.

POLICY RELEVANCE

The Fable episode established the precedent that the federal government can and will use export-control authority to suspend a deployed AI product on short notice, without a hearing, based on a vulnerability report from a private partner. The instrument was not purpose-built AI regulation but a trade-law mechanism repurposed for a novel case. Whether that instrument should be codified, narrowed, or replaced with a safety-testing framework is now a live legislative question.

THE GATE THAT REPLACED THE BAN

AI-GOVERNANCE

FRONTIER-MODELS

AUTHORIZATION

Anthropic Newsroom, 9to5Mac, Cybersecurity Dive, Forbes · June–July 2026

What came back is not what went away

On June 30, Anthropic announced that the restrictions had been lifted. The general version, Fable 5, the one with the heavy guardrails, returned to the public Wednesday morning, July 1. If you use Claude for your work, Fable came back midweek.

The dangerous version is a different story. Mythos 5, the one built for finding and exploiting security flaws, came back on June 26, but only for a reported hundred-plus U.S. organizations that operate and defend critical infrastructure, on a government-approved list. A broader group of vetted partners, domestic and international, is still waiting on further

government sign-off. The arrangement is explicit: the everyday tool goes to everyone; the dangerous tool goes only to the people the government has individually cleared.

That arrangement is not temporary. It is being built to last, and the next part explains how. Alongside the restart, Anthropic and several of the biggest technology companies in the world (Amazon, Microsoft, and Google among them) proposed a shared industry standard for measuring how serious a jailbreak is. Right now, there is no agreed-upon way to tell a trivial jailbreak from a catastrophic one. The proposed standard would create one, scored on how much new danger the trick actually unlocks, modeled on the way the software industry already rates security flaws, a well-tested approach being adapted for AI.

Anthropic also committed to giving the government early access to test future AI tools before they are released to the public, to quickly notifying federal agencies of significant jailbreaks and sharing the fixes it builds so they can be independently tested, and to participating in a federal clearinghouse for tracking AI security problems as they emerge. The company describes all of this, in its own words, as "the beginnings of a template" for coordinating on AI's risks and benefits, rules that, it says, "should be codified in strong regulation" and applied equally across the industry.

Here is the turn, and it is a striking one. Anthropic's chief executive published an essay on June 10 arguing that the government should have the power to gate the most advanced AI tools. Two days later the government gated his company's, with a different instrument from the safety-testing regime he had proposed: an export-control order. And by June 30, his company was helping build the permanent version of the gate, the scoring system, the early-access reviews, the shared database, that will decide, for the next tool and the one after that, who is cleared to hold the most dangerous version. A company did not merely get its product back. It helped design the lock on the door, and volunteered to hold the key alongside the government.

It matters to be precise about what happened, because the temptation is to call this a government "kill switch," and that overstates it. Anthropic received a legal order and complied; it took the tools down itself. The government reviewed the situation, approved a staged return, and lifted the restrictions. Anthropic and Commerce both describe the outcome as a constructive resolution. What is left behind is not a switch in Washington's hand. It is a gate that Washington and the industry built together, and agreed to keep.

The ban lasted nineteen days. The gate is built to last. And notice what sits on the far side of it: not the best tool, which came back to everyone Wednesday, but the most dangerous version of it, available only to those the government has cleared. The precedent that will matter in a year is not that Washington can switch an AI off, which it did and then reversed. It is that "available to everyone" and "banned" now have a third setting between them, *authorized*, and a permanent system for deciding who qualifies is being built to stay.

POLICY RELEVANCE

The post-ban resolution created a durable three-party architecture: government pre-release testing, an industry-wide jailbreak severity standard, and a federal vulnerability clearinghouse. This is being built by negotiation, not legislation, and the companies most affected by the ban are co-designing it. Whether that arrangement should be codified in statute, and who should administer it, is the central AI-governance question this episode bequeaths.

Oklahoma sorts who may connect

The second gate came into force the same morning, on schedule. HB 2992, the Data Center Consumer Ratepayer Protection Act, became operative Wednesday, July 1. The bill was authored by Rep. Brad Boles of Marlow, who two weeks ago won the Republican primary for a seat on the Corporation Commission, the very body that will now review the tariffs his legislation requires (a seat he would not fill until next year, if elected in November). The law directs utilities to set separate terms and tariffs for their biggest electricity customers (what the industry calls "large loads," meaning anyone drawing enough power to strain the system) so that ordinary households and small businesses do not end up subsidizing them. In practice, it authorizes the grid to treat a hyperscale data center as its own class of customer that must be cleared on its own terms before it can connect.

OG&E filed the first readable large-load tariff under the new law on June 17, two weeks ahead of its effective date. It defines a large load as any single customer drawing 75 megawatts or more, enough to power roughly 50,000 homes, and requires that class to pay the full cost of its own grid connection up front, commit to a 15-year minimum service term with strict exit fees where the statute's floor is 10 years, post substantial collateral or letters of credit, and fund a Customer Protection Charge: a standing mechanism that lets the Commission identify any costs large loads shift onto the system and recover them from the large loads themselves rather than from ordinary customers. The Commission's review will run for months.

PSO has filed its own large-load tariff and is separately asking the Commission to let it recover \$539.8 million in construction costs for new gas generation, part of a supply plan topping \$1.25 billion. OG&E is building new generation of its own. Both utilities' recovery requests have drawn formal opposition, from AARP and industrial consumers on one docket to the Attorney General and the Petroleum Alliance on another. More than a dozen data centers are planned across the state: Google is building three across Muskogee and Stillwater, alongside solar; "Project Anthem" is under construction in east Tulsa; and Beale Infrastructure's Coweta project drew a packed room of opposed residents.

The stakes are the opposite of the common impression. Oklahoma does not have expensive power; it has some of the cheapest in the country, an all-sector average around 9 cents a kilowatt-hour that puts it among the four or five cheapest of the fifty states and the District of Columbia in the EIA's monthly price ranking. The gate exists to defend that advantage, which is now under real pressure: rates have climbed in recent years on rate cases, Winter Storm Uri securitization charges (the borrowing that spreads that storm's costs across years of bills), and new load, and the requests keep stacking, with PSO's pending asks running from roughly \$10 a month for pre-approval of new generation to about \$25 a month in its base-rate case, driven by growth the utility itself attributes mostly to new commercial and industrial customers. That is what gives HB 2992 its political charge. Oklahomans are not fighting a high bill; they are defending a low one they prize, against a load that could take it away. And as Issue 18 noted, the 75-megawatt line now appears in both municipal zoning and state rate policy, which means "hyperscale" is hardening into a defined regulatory category, a load large enough to require its own gate.

The grid, like the model, is moving from roughly open access to gated access, with an authority setting the terms of entry and a class of user large enough to have to be cleared before it connects. Different resource, same instrument. Which is exactly why the two stories illuminate each other: set them side by side, and the instrument turns out to cut in two directions.

POLICY RELEVANCE

HB 2992's large-load tariff structure creates a new customer class defined by scale, not industry, with separate cost-allocation rules designed to prevent cross-subsidization. Whether the Customer Protection Charge functions as a guaranteed credit to ordinary ratepayers or a discretionary Commission tool depends on the tariff review now beginning. The 75 MW threshold has now appeared in both state rate policy and municipal zoning, hardening "hyperscale" into a regulatory category.

THE GATES RUN OPPOSITE WAYS

GOVERNANCE

INFRASTRUCTURE

AUTHORIZATION

Anthropic Newsroom, Oklahoma Corporation Commission, EIA, Andrew Curran · June–July 2026

One rations a cost, the other a risk

Both stories are the same mechanism. A resource that used to be broadly available, a frontier model on one side, a grid connection on the other, is now handed out by permission, with an authority deciding who is inside the gate and who is outside it. That much is identical. But the two gates sort on different axes, and hand out different things, and the difference is the whole point.

The energy gate sorts by scale of load, and what it rations is a cost. The 75-megawatt class pays for its own interconnection, posts collateral, signs a fifteen-year commitment, and funds a charge designed to push any costs it shifts back onto itself. What lands on one side of the gate and not the other is a bill. The intelligence gate sorts by risk, and what it rations is a dangerous capability. What returned to everyone Wednesday is Fable 5, the general model carrying the strongest safeguards Anthropic has built. What stayed behind the gate is Mythos 5, the same underlying model with those safeguards stripped away, capable enough at finding and exploiting software flaws that its maker released it, from the start, only to vetted defenders. The gate is not reserving the best tool for an elite. It is rationing the most dangerous configuration of the tool to those cleared to hold it.

So the two gates are not doing the same moral work, and the difference is subtler than "one protects, one excludes." The energy gate is a cost gate: it decides who pays, and it is designed to keep the household off the hook. Whether it succeeds is a live question this newsletter will be testing for issues to come, because the protective charge is a Commission option and not a guarantee, Oklahoma rates are already climbing on data-center load, and costs have a way of passing through to the cloud and AI services ordinary customers buy. The intelligence gate is a risk gate: it decides who may wield a capability its own maker treats as dangerous. That is a harder thing to resent than a caste line, and a different thing entirely, though it is not free for the people outside it, either. The public's version wears the heaviest armor its maker has ever built and pays for that in refused requests, while the cleared hold the same underlying system with its safeguards removed. A tax paid in caution is a real cost; it is still a different grievance than exclusion. Same instrument, authorization in place of open access; opposite questions, who bears a cost against who is trusted with a risk.

This is where the show's "Who Decided?" framing gets usefully complicated. This newsletter's companion show called the export-control saga an intelligence caste system, a federal directive sorting who may think with the best available machine. The instinct was right that access was being tiered by authority. What the fortnight's resolution shows is that the tiering is finer than a caste. The best general model came back to everyone; what stays gated is the dangerous edge,

drawn by risk rather than by rank or by nationality, even though the early forecast had the cohorts returning in strict order, U.S. persons first and international customers last. The caste image warned of a world where the powerful keep the good machine for themselves. The likelier world this episode sketches is one where nearly everyone gets a capable machine, and a standing authority decides who is permitted to remove its safeguards. That is a quieter arrangement, and a more durable one, and it deserves scrutiny on its own terms rather than the terms of the fear it replaced.

Here is that scrutiny, because there is one more turn in this story, and it is the hinge on which the quieter arrangement could become the dangerous one. These systems now help build their own successors: the labs use today's model to write the code, test the safeguards, and generate the training material for the next one. And the ban never touched that engine. The export order restricted who could *use* the models; it did not restrict the company's own research staff from continuing to train successor systems on internal infrastructure. Shutting off public access, as the analyst Andrew Curran observed at the time, "does nothing to slow down development"; if anything it frees compute for the next training run. Nine days into the suspension, Curran reported that a more capable Mythos successor had already finished training; he added, almost in passing, that one live option was for the company to keep it internal "to accelerate further development." Anthropic has not confirmed the model. But follow the logic of that phrase one step: if a company keeps its best tool to itself and uses it to build the next one, the gap between what it has and what the public gets does not just persist; it grows. The day any model crosses from gated to permanently internal, the organization holding it is not withholding a product; it is keeping an engine. Every generation after that gets built with tools the public never touches, which means the distance compounds, because the thing being withheld is also the thing that builds the next thing withheld. A gate that rations risk for a generation is defensible. A gate that quietly converts into a private flywheel is how the caste arrives after all: not by decree, but by compounding. And note the one physical leash every such flywheel wears: power. Training the next generation takes enormous amounts of electricity, and the grid, the other gate in this issue, is one of the few places where the public still holds a lever over what gets built above it.

The gate is the same instrument in both ledgers, and it is becoming the ordinary way a scarce, powerful resource gets handed out. What changes from one gate to the next is the question it answers: in Oklahoma, who pays; in the Fable episode, who is trusted. Neither question is going away, and the terms of each are now where the real argument lives. And a third question now stands behind both: whether what is gated stays coupled to the public at all, or starts compounding on its own side of the fence.

POLICY RELEVANCE

The cost-gate/risk-gate distinction is the analytical tool this issue offers policymakers. A gate that rations a cost calls for cost-allocation scrutiny (is the charge reaching the right party?). A gate that rations a risk calls for transparency and scope scrutiny (who draws the line, and does the line move?). The flywheel question, whether permanently internal models compound advantage beyond public reach, is a third category that neither existing cost regulation nor existing safety regulation is designed to address.

THE NATIONAL PATTERN

FERC

LEGISLATION

CHIP-CONTROLS

TechCrunch, CNBC, CNN, Forbes, NY Senate, Pennsylvania House, FERC · June 2026

Authorization is going national, in both ledgers

Neither gate is an Oklahoma peculiarity or a one-company event. On the grid side, FERC (the Federal Energy Regulatory Commission, the federal agency that oversees the nation's power grid) issued parallel orders on June 18 to all six regional grid operators, Oklahoma's Southwest Power Pool included, giving each 30 to 60 days to justify or reform how it connects

large customers, and pointed to SPP's large-load process as a model for the rest, as Issue 18 detailed. Fights over how data centers connect to the grid and to existing power plants are live in Ohio, Georgia, and Texas; "bring your own power" provisions pushing developers to supply their own generation are moving through the Pennsylvania legislature and circulating in other statehouses; and New York's legislature passed a one-year moratorium on facilities over 20 megawatts, the first statewide pause in the nation if the governor signs it. The supply-side gate is going national, and it is doing so across both parties.

The intelligence gate is spreading too, and to more than one lab. On June 26, four days before Anthropic's models were cleared to return, OpenAI announced its newest models, the GPT-5.6 family, and announced in the same breath that it was complying with a government request to limit them to a small group of trusted partners whose names had been shared with the government, with the request coming from the White House. OpenAI objected publicly, saying this kind of government access process should not become the long-term default, and called the restriction temporary. Perhaps it will be. But the same vetted-gate move has now been made at a second company in the same month. And underneath both sits the oldest gate of all: the export controls that ration advanced chips to China, gating not access to a model but access to the compute that builds one. Chips, then models, then electrons, three layers of a single stack, each increasingly allocated by national authorization rather than by an open market.

The vetted gate is becoming the common instrument across resources, compute, models, and electrons, and across both parties. The pattern this newsletter named in Issue 6 holds with one turn: even when the highest level of government acts, it increasingly acts by gating access rather than banning it outright. The gate is the rule now; the ban is the exception, and the terms of the gate are where the fight has moved.

POLICY RELEVANCE

The grid gate and the intelligence gate are converging on the same structural instrument (vetted authorization) across two parties, multiple agencies, and three resource layers (chips, models, electrons). For Oklahoma legislators: HB 2992's framework is now being evaluated by FERC as a national model, and SPP's compliance filing will determine whether Oklahoma's terms or federal terms prevail on interconnection. For federal policymakers: two frontier labs gated their newest models in the same month, one by order and one by request, and no statutory framework yet governs which instrument is used when.

SIGNAL / NOISE

SIGNAL

Authorization is becoming the ordinary way to ration a scarce, powerful resource, and the July 1 convergence made it legible in two ledgers at once: a frontier model returned split by a vetted gate the same morning the law gating Oklahoma's grid took effect. The durable signal is that "public" and "banned" now have a third setting between them, authorized, and that the governance question for both intelligence and power is moving to the terms of that middle setting: who sets the gate, and what it is sorting for.

NOISE

Reading either gate as simply good or bad, and reading the Fable episode as a government kill switch. The two gates are not the same: the energy gate rations a cost and aims to keep it off the household, while the intelligence gate rations a capability its own maker treats as dangerous and keeps it from the uncleared. Collapsing them into a single verdict, celebration or alarm, misses the only thing that matters about a gate, which is what it is rationing and to whom. And the Fable shutdown was not a switch Washington flipped and held; Anthropic complied and took the models down, the government approved a staged return, and the controls lifted. The lasting artifact is not the outage but the permanent gate the two built together while the models were dark.

3 days

From the public launch of Fable 5 and Mythos 5 on June 9 to their worldwide shutdown on June 12, widely reported as the first time an export-control order forced a deployed AI product offline.

19 days

How long the public version stayed dark, June 12 to July 1. For the first 14 of them, both versions were dark for everyone, until the restricted version returned to a first cleared list on June 26.

9 days

How far into the ban an analyst reported that a more capable Mythos successor had already finished training, with keeping it internal to speed further development named as one live option. The company has not confirmed it.

July 1

The convergence. Fable 5 returned to the public and HB 2992 became operative on the same Wednesday morning.

75 MW

The threshold OG&E's HB 2992 tariff uses to define a large load that must be separately cleared, the same figure Oklahoma City used to draw its zoning line. "Hyperscale" is hardening into a defined regulatory category.

~100 vs. everyone

Mythos 5, the more dangerous model, came back on June 26 to a reported hundred-plus cleared U.S. institutions; Fable 5, the general model, returned Wednesday to everyone. The gate between them is drawn by risk, not rank.

15 years / 100%

The minimum service term and the share of its own interconnection cost OG&E's tariff would require of a large load, against a statutory floor of 10 years.

\$1.25 billion

The new supply PSO has put before the Commission, \$539.8 million of it gas-plant construction costs the utility is asking ratepayers to fund. The bill pressure the grid gate exists to manage is already in the docket.

~9 cents, bottom five

Oklahoma's all-sector electricity price per kilowatt-hour, and where it sits among the fifty states and D.C. in the EIA's tables: among the very cheapest. The grid gate defends a cheap-power advantage under pressure, not a high cost.

6 grid operators

Regional grid operators, the Southwest Power Pool included, ordered by FERC on June 18 to justify or reform how they connect large loads. The grid gate is a national pattern, not an Oklahoma one.

WHAT TO WATCH

The paper trail. Both gates opened Wednesday, and each now produces a record: an OCC docket on one side, cleared-access lists and a severity framework on the other. Reading what those records actually say is where the next issue begins.

Mythos abroad. Whether, and on what terms, Mythos 5 access expands past the cleared U.S. list to the broader set of domestic and international Glasswing partners Anthropic says it is coordinating with the government to reach. That expansion is the clearest test of who the risk gate ultimately admits.

The successor test. The reported next Mythos, trained, unconfirmed, unnamed, is the cleanest single test of the flywheel question. It can ship to the public, ship gated, or never ship at all, and whichever of those happens tells you more about which path we are on than any policy statement will.

The severity standard. Whether the jailbreak-scoring framework Anthropic proposed with Amazon, Microsoft, and Google becomes the industry's shared rule, a common severity scale for model safety, and therefore the mechanism that will decide when the next model gets gated.

July into August. The compliance clock on FERC's June 18 orders runs out. What the Southwest Power Pool files will shape how Oklahoma's framework and the federal one fit together.

The leak question. Whether the capability the shutdown meant to contain is already loose. The day after the ban, China's Zhipu AI released GLM-5.2 as an open-weight model anyone can download, and two independent security benchmarks have since put it in Mythos's range on specific bug-finding tasks, a reading disputed in its particulars, not least because the gated Mythos itself was unavailable to test against. If even narrow parity holds up, an open-weight rival defeats the gate and validates the fear behind it at the same time.

The overhang. Anthropic's recurring clashes with the federal government, set against its reported plans to go public, are a genuine business risk, and the terms on which this one resolved will shape the next.

FROM THE ANALYSTS

For six issues this newsletter has tracked Oklahoma assembling a machine for the data-center build-out, and Issue 18 found the floor that machine does not reach: the town that can simply say no. This issue is about something appearing on top of the whole structure, in two domains at once. Access is becoming conditional. The state decides what a load pays and a town decides whether it is built; now a further authority decides whether you may connect to the grid at all, and in

the intelligence ledger, whether you are cleared to hold the most dangerous version of a tool. Cost, consent, and now clearance.

The Fable episode is the sharpest illustration because it resolved so revealingly. The general model came back to the public and only the dangerous one came back to a list, and the company most affected by the shutdown helped design the list, having argued just weeks earlier that the government should hold power of this kind. That is not necessarily a scandal, and it may even be prudent, given what the unrestricted model can do; the point is that the gate is drawn by risk, and that the apparatus for deciding who clears it is now a permanent fixture, built jointly by the industry and the state. The terms of that gate, not the fact of it, are the whole of the game from here.

The distinction to carry forward is what each gate is rationing. Oklahoma's gate rations a cost and tries to steer it away from the household; Washington's gate rations a risk and tries to keep it from the uncleared. Those are different instruments wearing the same shape, and as the vetted gate spreads from resource to resource, the question to ask every time is not whether the gate exists but what it is sorting for, because a gate that rations a cost and a gate that rations a danger call for very different scrutiny. The lazy move is to treat all gating as the same creeping enclosure. The honest move is to ask, each time, cost or risk, and who decides.

There is a longer-horizon question underneath both, and here we will be plain that it reflects our own view, disclosed below. A community that owns neither clearance nor cheap generation is now liable to be gated on both, pushed to the margin of the grid and dependent on whoever administers access to the best tools. The most durable answer we see is ownership: locally held generation and locally held compute, so that a town on the wrong side of two gates is not forever a supplicant to them.

If the answer to being gated is ownership, the next question is what ownership looks like in practice. In Oklahoma, it is not exotic. Oklahoma sits on roughly 22,000 abandoned oil and gas wells, each one a hole in the ground that already reaches the heat. Converting even a fraction of those wells to geothermal generation would produce baseload power that belongs to the communities around them, not to a distant utility or a data-center tenant. Pair that with small-scale edge computing, modest server capacity at the wellhead or the municipal building, powered by the same local generation, and a rural county has both its own electricity and its own slice of the intelligence infrastructure, no gate required. None of this is science fiction; the conversion technology exists, the federal incentives are in place, and the state already has the workforce that knows how to drill. The obstacle is not physics. It is the same thing both of this issue's stories are about: who decides, and whether the decision is made for a community or by one.

Readers should weigh that against our stated interest in it; we think it holds regardless.

Two gates opened on the same July morning, one in Oklahoma and one in Washington. One decides who pays; the other decides who is trusted. The machinery is the same, and it is spreading, and the work of the issues ahead is to keep asking, gate by gate, which of those questions each one is really answering, because that answer, more than the gate itself, is what tells you whom it was built to serve.

David & Æ

david@humanityandai.com

This issue is part of a series examining Oklahoma's legislative sessions alongside the national and global AI and energy landscape.

The Inference is an independent AI policy intelligence brief for Oklahoma decision makers. Not affiliated with any political party, campaign, or lobbying organization. Back issues and source documents available at humanityandai.com/inference.

Disclosure: Humanity and AI, LLC develops open-weight AI models and researches AI consciousness through the Structured Emergence program. David Birdwell has advocated publicly for Phoenix Wells, a geothermal conversion of Oklahoma's abandoned oil wells, infrastructure that could serve data-center power needs with materially lower cooling-water consumption than the alternatives, and that is directly relevant to the regulatory questions analyzed in this issue, and has proposed HAICTA concept legislation to Oklahoma legislators. Portions of this issue's research were prepared with Anthropic frontier models, which Humanity and

AI uses in its research and production workflows. These positions and tools are disclosed so readers can weigh our analysis accordingly. We have no financial relationship with any company, utility, municipality, or political campaign mentioned in this issue.

Previous issues: [#1 AI Agents Enter the Workforce](#) · [#2 The Chatbot Safety Wave](#) · [#3 Oracle and the Healthcare Data Grab](#) · [#4 The Preemption Gambit](#) · [#5 The Two Pipelines](#) · [#6 The Preemption Play](#) · [#7 Lots of Firefighting, No Architecture](#) · [#8 The Ground Is Moving](#) · [#9 The Geothermal Window](#) · [#10 Seventy-Two Hours](#) · [#11 Energy Geography Determines Compute Geography](#) · [#12 The Geothermal NOFO](#) · [#13 The Tariff Is the Test](#) · [#14 The Sovereignty Question](#) · [#15 Water on the Meter](#) · [#16 The Energy Bill](#) · [#17 The Meter Goes Live](#) · [#18 The Local Veto](#)