After the Line

The Data Center Debate and the Measure It Lacks

J. Konstapel, Constable Research, Leiden — August 29 2026

1. A floor, not a ceiling

On 26 August 2026 WIRED reported that more than fifteen candidates for Congress, the Senate and governorships had signed the AI Pact. The count stood at eighteen a day later: seventeen Democrats and one independent, Dan Osborn of Nebraska. The pact is an offshoot of the Political Integrity Project. It contains five promises. Ban tax breaks, backroom deals and excessive energy use for data centers. Require safety reviews before AI models are released. Create an AI dividend for workers. Let consumers and companies sue over harms. Refuse public bailouts for AI companies that take reckless market risk.

Its founder calls the language deliberate. The pact is a floor. It stops short of a moratorium so that candidates with strong ties to labour can sign. Several unions have just reaffirmed their support for data center construction.

Osborn describes a town hall in rural Nebraska. A man he calls "hardcore MAGA" approached him. They had no common ground until Osborn asked about AI and data centers. Then they were, in his words, in lockstep. Osborn does not rule out a pause: "Calling a time-out right now doesn't seem like a horrible thing."

The other Texas candidate in the story, Gina Hinojosa, puts the fear in one sentence. "We should not be replacing teachers with AI. We should not be replacing nurses with AI."

This essay takes the pact seriously. It is a well-built floor. The question is what stands on it. The answer is: nothing yet, because the debate on every continent is missing one measure. Without that measure, "excessive" has no number, "jobs" has no place, and "school" has no alternative.

2. The same debate on every continent

The pact is not an American peculiarity. The same three parties sit at every table. National governments and Brussels want capacity. Local councils and residents want water, grid and land back. Companies offer voluntary promises where laws threaten.

The Netherlands. In February 2026 the Dutch parliament adopted a motion by Pieter Grinwis to tighten the legal definition of a hyperscale data center. The current rule from 2022 requires both 10 hectares and 70 MW. Projects in Amsterdam and Lelystad stay below one threshold on paper while having the impact of a hyperscaler in practice. 105 of 150 members voted for the motion, including the largest governing party. A second motion in March, calling for a national plan for data centers and data cables, passed with 134 votes.

The rules arrive after the facts. Permits for seven hyperscale sites have already been granted. Four are in Haarlemmermeer, one in Vijfhuizen, one in Lelystad, and three Microsoft towers in Amsterdam, which declared its own moratorium years ago. Together they will use as much

electricity as the inhabitants of Haarlem. Meanwhile 15,000 businesses wait for a grid connection. Data centers connect at the highest level of the grid and rarely stand at the back of the queue. The sector reports that a permit took ten weeks in 2016 and takes 4.5 years in 2026. Dutch capacity grew 7 percent in 2025 while Europe grew considerably faster. The chief executive of Eneco asked in the FD on 19 July for more data centers, to create demand for green power that would otherwise be wasted.

Ireland. Data centers used 23 percent of the country's electricity last year, more than all urban households combined. Between 2015 and 2023 that demand added a cumulative average of €360 to a household bill. 3,300 people, 0.1 percent of the workforce, are directly employed in them. After a de facto moratorium since 2021, the regulator now requires new data centers to install on-site generation or storage matching their full import capacity.

Europe. The Cloud and AI Development Act, adopted in June 2026, aims to triple EU data center capacity within five to seven years and to streamline permitting. The motive is sovereignty. Amsterdam, meanwhile, imposed a new moratorium in April 2025 until at least 2030. A coalition of researchers and campaigners now argues for an EU-wide pause until data centers are shown to serve the public interest. In the Netherlands, housing and fifty schools and day-care centres could not be built because a large data center had priority access to the grid.

United States. Gallup found in March 2026 that 71 percent of Americans oppose an AI data center near them, 48 percent strongly. Opposition to a nearby nuclear plant was 53 percent. In the first quarter of 2026, 75 projects worth a combined $130 billion met organised local opposition. In Murdock, Nebraska, population 275, conservative farmers and the Sierra Club filled the fire station together, against labour leaders, worried about farmland, water and bills. The Governor of Texas, a Republican who has courted data centers, said this month they had "dug their own grave" and imposed a temporary moratorium.

Elsewhere. In Uruguay campaigners discovered that Google's planned cooling would use 2 million gallons of drinking water a day. The project proceeded only after it was scaled down and water cooling was dropped. Google paused a $200 million data center in Santiago in 2024. Residents near a Microsoft data center in Mexico report more frequent power cuts and water outages lasting weeks since 2024. In Johor, Malaysia, the country's first data center protest demanded compensation for dust and noise. India, moving the other way, grew its installed capacity more than fourfold in six years to 1,263 MW and welcomes more.

Every one of these positions measures the same thing: how much electricity, water and land, and who pays. None asks why it costs so much.

3. The gas behind the pact

Two days before the pact story, WIRED reported that the gas-fired generating capacity under development for American data centers had nearly doubled in under a year. Earlier reporting showed that gas-powered data centers are permitted to emit more greenhouse gas than entire nations. The mechanism is simple. The fastest way to get power is to build a private gas plant beside the building. Those plants use, for the most part, less efficient technology than the plants on the public grid.

This is the direct outcome of the Ratepayer Protection Pledge that Amazon, Google, Meta, Microsoft, OpenAI, Oracle and xAI signed at the White House on 4 March 2026. The pledge asks companies to "build, bring, or buy" their own power so that household bills do not rise. It is non-

binding. It makes the private gas plant the approved answer. The AI Pact of August is the response to that outcome. Its five points regulate who pays, who is liable and who reviews. Only one point touches use itself, and that point, "excessive energy usage", has no definition. Nobody at the table has a measure of what computation ought to cost.

4. Three assumptions every party shares

The steam engine was the first machine that turned heat into motion at will. Everything since has been built on its accounting: energy in, work out, heat lost. The data center is the last machine built on that accounting, and the debate around it inherits three assumptions without noticing them.

Intelligence is computation. A matrix times a vector, billions of times a second. Every operation costs energy and returns heat. More intelligence, more electricity. The moratorium and the AI factory share this assumption. They differ only in whether they want the result.

Computation is language. The language model predicts tokens. Everything it knows is a distribution over words. Therefore all knowledge of the world must first become text and then a matrix. That is why the corpus is so large and the training so expensive.

Search is the core operation. Every query is a search through a space of parameters or vectors. Cost scales with the size of that space. A bigger model is a bigger space is more electricity.

The Irish number and the fifty Dutch schools are not side effects. They follow from these three assumptions. As long as the assumptions stand, a moratorium moves the problem to India or the Gulf.

The cost can be stated. A model with N parameters performs about 2N arithmetic operations per generated token. For $N = 10^{12}$ that is 2×10^{12} operations per token. A current accelerator delivers on the order of 10^{15} operations per second at about 1,000 watts, so about 10^{12} operations per joule. One token therefore costs about 2 joules in arithmetic alone. A five-hundred-token answer costs about 1,000 joules. Memory traffic, cooling and conversion multiply this by three to five. Call it 3 to 5 kilojoules per answer. The number is not the point. The scaling is. Double the model and the cost per answer doubles. Nothing in the pact, the pledge or the motion changes that line.

5. The measure: address instead of search

Picture the archive of everything known as a fishing net. Knots at every crossing, strands between them. The language model handles a query by scanning the whole net for the knot that fits best. The bigger the net, the longer the scan.

There is another way to find a knot in a net. Follow the strand. Every knot has a position: from this knot, the next one is above, on, or below the reference. Three states. Write them as +1, 0 and −1. A position is a sequence of these. That sequence is the address. It is not assigned from outside. It is read off the net itself.

This register is small. With 19 such digits there are $3^{19} = 1.16 \times 10^{9}$ addresses, enough for every printed book. With 26 digits, $3^{26} = 2.54 \times 10^{12}$, every document. With 32, $3^{32} = 1.85 \times 10^{15}$, every sentence. With 42 digits, $3^{42} = 1.09 \times 10^{20}$, every moment of every human being who has ever lived. Forty-two digits of three states fit in nine bytes.

Four operations replace search. Equality: are two addresses the same. Prefix: does one address lie within another. Negation: mirror the sequence. Subtraction: how far apart are two addresses. Each operation touches at most 42 digits. A few hundred elementary steps, on the order of 10^{-9} joule. The cost is proportional to the number of digits, not to the size of the archive. Double the archive and the cost per query does not move.

Two things follow. There is no registry and no authority that hands out addresses, because the address is computed. Two archives merge by taking the union of their integers; no mapping of schemas. And an address at depth n is the truncation of the same address at depth n + k. Refinement is monotone. Nothing ever needs re-indexing. Knowledge that does not yet exist already has its place.

Set this beside the third assumption. Search was the core operation because the address was not known. Once the address is read off the net, search is gone, and with it the line that ties intelligence to electricity.

6. Fit instead of computation

The first assumption goes the same way. In a net, what returns in phase remains. A strand that closes on itself holds; a strand that does not close falls away. That is selection, not calculation. Nothing is computed. Something either fits or it does not.

This is the difference between a bulb and a laser. The bulb pushes current through a wire and radiates in every direction; most of the energy is heat. The laser holds a state; light that fits the cavity is kept, light that does not is lost, and the energy goes into holding the state, not into any operation. Right-brain computing, as described on this site, is built on optics for that reason. Coherence and resonance do the work that arithmetic does in the language model. The cost sits in maintaining a state, and that cost does not grow with the number of things the state can distinguish.

The second assumption falls with the other two. Language is one way to write a position down. It is not the position. The three-state register is the register; text is one of its readouts.

7. Work: job into fit

The pact's second promise is an AI dividend for workers. The Irish objection to data centers is 3,300 jobs for 23 percent of the power. The strategist quoted by WIRED asks whether construction jobs now are worth the destruction of other workers' opportunities later. All three measure in the same unit.

A job is a place on the line. You sell time; someone else decides what happens in that time; the result is counted in hours times wage. Henry Ford's assembly line of 1913 made this the shape of work, and the data center is that line without people. To object that it creates too few jobs is to say: put people beside it. It is not to ask whether this is the way to make anything.

In the net there is no line. A person stands at a depth. What that person does either returns in phase at that depth or it does not. The personal blueprint, worked out on this site, is the address at which someone fits, and the next step from there. Its invariant part does not change over a life; its state changes with every closure. A region where many people work at their own address cannot be read off a jobs counter. It can be read off how much returns without being organised.

The farmers of Murdock, Nebraska were not defending jobs. They were defending land, water and the possibility of staying where they are, in their own way. That is the address. The jobs counter is the language they are handed to say it, and that language always loses to the factory, because a factory can promise more jobs than a village can count.

8. Learning: school into life-course

"We should not be replacing teachers with AI." Hinojosa is defending the school as a building with jobs in it. The school, too, is the line, then for people. A batch per birth year, one pace, one syllabus, move along at the end of the year, inspect at the end of the run. It was built in Prussia and copied in Ford's America because the factory needed workers who arrive on time and do what they are told. The fifty Dutch schools that could not be built are fifty units of supply capacity for the line.

One person built something else and tested it himself. Roger Schank, at Yale, Northwestern and Carnegie Mellon, and in his own company, argued that people do not learn by instruction. They learn by expectation failure. You hold a script — how a restaurant works, how a customer reacts. The script fails. At that moment you remember, and you revise the script. No failure, no learning. Knowledge is therefore not a store of facts but a store of stories, indexed so that the right one surfaces at the moment it is needed. Teaching, then, is doing, in a setting where failure is possible, with a story ready that fits what went wrong. His test is a single question. Ask someone what they learned at school last year and there is no answer. Ask what went wrong last year and there is a story.

Schank removed the line. No cohort, no syllabus, no test. Two things he did not have.

Who learns. Schank's learner is generic. Everyone holds scripts, everyone fails, everyone indexes. Why one failure becomes a story for one person and slides off another, he could not say. That is the blueprint. A failure that lands at a person's own depth is kept. A failure at another depth is noise. Schank described the mechanism; the blueprint describes the selection.

What a script is. For Schank it is a structure in memory: a representation, language about the world. In the net a script is a strand that has closed. Failure is a mesh that does not close. Learning is letting the strand close again at another depth. No storage, no indexing, no retrieval. This is why case-based reasoning stalled as an engineering programme. It was built as a memory with indexes, and everything that searches scales like a search. On address it does not.

Life-course guidance, as sketched on this site from a 1993 design, is where the two meet. The coach is Schank's storyteller at the right moment. The blueprint says which moment that is. The fellow travellers are the people at the same depth who know the same failure. None of this needs a building, and none of it needs the line.

9. One line above the floor

The AI Pact regulates distribution, liability and oversight. It leaves "excessive" undefined because no party has a measure of what computation ought to cost. Without that measure the gas plant beside the building is the rational answer, and the pact settles afterwards who pays for it.

The measure exists. Cost per query proportional to the digits of an address, not to the size of the archive. Intelligence as fit, with energy spent on holding a state rather than on operations. Work as

the return at a person's own depth, not a place on the line. Learning as the closing of a strand at the right depth, not the filling of a cohort.

One line could stand above the floor: *support a measure of computation under which the cost of a question does not grow with the size of what is known*. Everything else in the pact can remain. Tax breaks, backroom deals and bailouts are wrong on any accounting.

What is needed to put that line there is not an argument. It is a running case. An archive on address, with the four operations, whose energy per query has been measured and does not rise as the archive grows. That number, set beside the Irish 23 percent, ends the debate that the pact can only regulate.

Annotated references

Molly Taft, "Candidates Are Signing a Pact Promising Action on Data Centers and AI Safety", WIRED, 26 August 2026. Why read: the primary account of the AI Pact, with Osborn's town hall, the founder's "floor" reasoning, the labour position and Hinojosa's teachers-and-nurses line. Read it for how carefully the pact avoids saying what "excessive" means.

Molly Taft, "Data Centers Are Driving an Alarming Gas Power Expansion in the US", WIRED, 25 August 2026. Why read: the near-doubling of gas capacity under development in under a year, and the private-plant mechanism. This is the outcome the pact is reacting to. Read it before the pact story, not after.

The Hill, "Democrats in key House, Senate races sign AI regulation pledge amid data center scrutiny", 26 August 2026. Why read: the full list of five promises, the count of eighteen signatories, and the Emerson poll (63 percent opposed to a local data center). Useful for the exact wording.

The AI Pact, theaipact.com. Why read: the pledge text itself and the candidates' own statements. Note how many statements are written in the vocabulary of jobs, bills and moratoria, and how few say what AI is.

The White House, "President Trump Secures Historic Commitment to Keep Electricity Costs Down Amid Data Center Boom", 4 March 2026. Why read: the Ratepayer Protection Pledge as signed by the seven companies: "build, bring, or buy". This is the document that licenses the private gas plant. Read alongside the Axios preview of 25 February.

Tweede Kamer, Motie van het lid Grinwis c.s., Kamerstuk 36 800 XXIII nr. 46, 12 February 2026; Motie Zalinyan/Grinwis, 29435-288, 30 March 2026. Why read: the two Dutch motions in their own words, with the 10-hectare-and-70-MW criterion the first one wants tightened. The voting record (105 and 134 of 150) shows how broad the unease is. Read with the DatacenterWorks commentary of 6 March on why the motion exposes a weakness rather than fixing one.

Dutch Data Center Association / Pb7 Research, "State of the Dutch Data Centers 2026". Why read: the sector's own numbers. Ten weeks to 4.5 years for a permit; 7 percent growth against a faster Europe; six bottlenecks. Read it as the view from inside the line.

Michelle Thorne, Aurora Gómez Delgado, Fieke Jansen and Rosi Leonard, "The Case for a Temporary Moratorium on Large Data Centers in Europe", TechPolicy.Press, July 2026. Why read: the Irish figures (23 percent, €360, 3,300 jobs) and the Dutch fifty schools, with sources. The

strongest case for a pause, and a clear example of measuring the damage of the line in the units of the line.

Regulation (EU) 2026 on Cloud and Artificial Intelligence Development (CADA), adopted June 2026. Why read: Brussels' answer, which is to triple capacity in five to seven years. Read the recitals for the sovereignty motive and note that "energy" appears only as a supply problem.

Gallup, "Americans Oppose Local AI Data Centers", March 2026; Data Center Watch / Reuters, first-quarter 2026 opposition index. Why read: 71 percent opposed, 48 percent strongly, above nuclear at 53 percent; 75 projects and $130 billion facing organised opposition in one quarter. The numbers behind the pact's timing.

Henry Ford with Samuel Crowther, *My Life and Work*, 1922. Why read: the line described by the man who built it, in plain language. Chapter 7, on the moving assembly line, is the origin of the unit that the whole data center debate still uses.

Roger C. Schank and Robert P. Abelson, *Scripts, Plans, Goals and Understanding*, Erlbaum, 1977. Why read: the origin of the script. Read the restaurant script in chapter 3 and then ask what happens to a script that cannot close.

Roger C. Schank, *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*, 1982; revised as *Dynamic Memory Revisited*, 1999. Why read: expectation failure as the moment of learning, and the indexing of cases. The 1999 revision includes Schank's own account of why the school could not use it. Read the revision.

Roger C. Schank, *Teaching Minds: How Cognitive Science Can Save Our Schools*, Teachers College Press, 2011. Why read: the twelve cognitive processes that no school subject teaches, and the story-centred curriculum as built and run. Written by a maker after the making, which is the only kind of education book worth reading.

J. Konstapel, "How to Visualize the Vacuum", constable.blog, 2026. Why read: the fishing-net image used throughout this essay: knots, strands, closure and the three states above, on and below. Read first if the net is unfamiliar.

J. Konstapel, "The Codable World: Design for a World-Scale Intelligence in Which All Knowledge, Past and Future, Becomes Address", constable.blog, August 2026. Why read: the full derivation of the address register: 3^{1^9}, 3^{2^6}, 3^{32}, 3^{4^2}, the four operations, monotone refinement and the cost proportional to digits. The essay above states the results; this article shows the steps.

J. Konstapel, "The Personal Blueprint of the Net", constable.blog, August 2026. Why read: the split between the invariant address of a person and the state that changes with every closure. This is the selection that Schank lacked.

J. Konstapel, "Technology of the Net", constable.blog, August 2026. Why read: the design rule (set the state so that the goal is the fixed point) and the six technologies, including the net as a synchronisation service and medicine as state determination. The latter is where Hinojosa's nurses belong; it is not worked out in this essay.

J. Konstapel, "Right-Brain AI", constable.blog, November 2025, with technical specification. Why read: the earlier, intuitive statement of the same line: the language model as left brain with an energy ceiling, and coherence, a nilpotent kernel and optics as the alternative. Read it as the origin, not the current form.