22 JULY 2026

From intelligence ambition to demonstrable value: the data center debate needs a yardstick

The Dutch data center debate has been caught between two extremes for years. On the one hand, data centers symbolize high electricity consumption, grid congestion, space consumption and water consumption. On the other hand, they are presented as indispensable infrastructure for our digital economy.

In his article [The data center debate: from right-versus-mistake to intelligence-ambition](https://rubenswindow.substack.com/p/het-datacenter-debat-van-goed-versus), Ruben Nieuwenhuis makes an important attempt to get out of these trenches. His central question is no longer: do we want data centers or not? The relevant question is: what intelligence do we want to produce with our digital infrastructure, who benefits from it and under whose control does it happen?

That is a necessary shift. And Ruben does not stop at an abstraction: he translates the ambition into three concrete questions — what do we train ourselves, who buys it, and who comes here to build — and finally, he asks for an annual computing power monitor that measures value per kilowatt hour, so that the myths do not recur every two years.
That is precisely what I want to build on. Because the call for a monitor is justified, but a monitor is only a yardstick when you know what you measure, how you add it up and how you verify it. Without that architecture, "value per kilowatt hour" remains an intention. Concepts such as sustainable AI, sovereign cloud, European compute and circular digital infrastructure sound attractive, but only become meaningful when they are translated into demonstrable performance.
The next step in the debate is therefore not a new concept alongside Ruben's ambition, but the foundation underneath:
From intelligence ambition to measurable and demonstrable compute value.
Ruben names the instrument. This piece is about how to build it.

A data center is not a goal, but infrastructure

The public debate often treats "the data center" as one category. In reality, these are very different infrastructures and applications. Ruben makes that distinction sharply with the image of warehouse versus factory; It is worth extending that distinction here to the yardstick, because each category requires different indicators.
A traditional enterprise or colocation data center supports payment transactions, healthcare systems, government services, business applications and digital communication, among other things. A hyperscale cloud location serves an international market with storage and cloud services. An AI data center uses GPUs and other accelerators to train, refine, and apply models.
A distinction must also be made within AI infrastructure.
Large-scale training requires large amounts of concentrated computing power and in many cases can take place in locations where renewable electricity is widely available. Inference – the actual use of models – often requires capacity closer to users, data and business processes. This is relevant for latency, business continuity, privacy and jurisdiction.
The logical end picture, therefore, is not one all-encompassing AI gigafactory, but a layered compute architecture:

- large-scale training capacity in energy-enabled locations;
- national and regional infrastructure for inference and fine-tuning;
- edge capacity close to machinery, hospitals, governments and industrial processes;
- a measurement and evidence layer that provides insight into performance across this infrastructure.

That last layer is precisely the subject of this piece. The first three layers are about where computing power is; the fourth is about whether we can demonstrate what that computing power is worth. Without that fourth layer, we still judge a layered architecture with the instruments of a single box of blocks.

Europe is now firmly committed to AI Factories and AI Gigafactories. The European Commission has announced a €20 billion facility within InvestAI for up to five AI Gigafactories, in addition to the wider network of European AI Factories.

However, the question is not only how many GPUs Europe can install. The decisive question is which European applications, companies, public services and knowledge clusters will run sustainably on it — and whether wecan turn that value into a demonstrable figure instead of an assumption.

From electricity consumption to value per kilowatt hour

Dutch data centers collectively received about 5.1 billion kilowatt hours of electricity from the public grid in 2024. This corresponded to 4.6 percent of Dutch electricity consumption.

That is a relevant fact, but not a judgment in itself.

A kilowatt hour used to sell advertising profiles more efficiently has a different social value than a kilowatt hour that is used to train a medical model, optimize an electricity grid or control an industrial installation in a more energy-efficient way.

At the same time, social value must not become a license for unlimited energy use. An application can be strategically important and still run inefficiently. That is why costs and revenues must be brought together in the same yardstick.

Not only:

- how many megawatts are connected;
- how much electricity is consumed;
- how much water is taken;
- how many square meters are used.

But also:

- How much usable compute is provided;
- which workloads are being run with it;
- how much economic and social value is created;
- how much energy or CO_2 is avoided elsewhere;
- what proportion of the capacity is available to European organisations;
- which legal and operational control data and models fall under;
- how much residual heat is actually used;
- what flexibility is provided to the energy system.

The central indicator should therefore not only be Power Usage Effectiveness, but increasingly also:

Useful Compute Value per kilowatt-hour

Let me define that concept sharply, because its strength lies in its precision. Useful Compute Value is not a single financial number and does not pretend to be. It is a composite measure that shows per unit of energy how much *productive* computing power is supplied, what value is created with it, what environmental impact there is in return and under whose control this happens. In other words, it is a normalised dashboard of technical, economic, environmental and sovereignty indicators – not a price tag, but a profile.

This choice for a profile instead of one number is deliberate. A single key figure invites you to optimize on the number instead of on reality; a coherent profile forces the showing of trade-offs. A data center that excels in energy efficiency but runs structurally underutilized hardware will not get away with it in a profile. And neither is an application with high social value that is carried out in an energy-wasteful way.

PUE alone is not enough

PUE remains a useful indicator of the ratio between the total energy consumption of a data center and the energy consumption of the IT equipment. But PUE does not tell what the IT equipment actually produces.

A concrete example makes that clear. Take two data centers with an identical PUE of 1.2. One delivers near-full GPU power to a medical imaging model that supports thousands of diagnoses daily; The residual heat goes into a nearby heat network and the training tasks shift along with the supply of wind power. The other runs structurally underused servers for an oversized model that is hardly used, on gray electricity, without heat utilization. They are indistinguishable on PUE. On Useful Compute Value, they differ like night and day. A low PUE can even go hand in hand with inefficient software, oversized models or hardware that is hardly used — a good report mark for the packaging says nothing about the contents.

For AI infrastructure, we therefore need to measure at least on five levels.

1. Efficiency of the facility

This includes the existing indicators for energy, cooling and water:

- PUE
- Water Usage Effectiveness
- Share of renewable energy
- Cooling effectiveness
- Actual reused residual heat
- Energy use of ancillary installations

The European reporting obligations for data centers from 500 kW installed IT capacity include energy use, water use, renewable energy, cooling and residual heat. The

European database publishes this information at an aggregated level. This is the only layer for which there is already a legally enshrined reporting flow — and as I will discuss later, compliance with it is not yet a done deal.

2. IT infrastructure efficiency

Next, it is necessary to measure how efficiently the existing hardware is used:

- average CPU and GPU utilization rate
- idle time
- Energy consumption per compute task
- Energy consumption per training run
- Energy per inference request or per million tokens
- Ratio of reserved to actual capacity used
- Technical lifespan and reuse of hardware

This shows whether installed computing power is actually productive. This is the layer that makes the sharpest difference between "connected load" and "delivered compute", and it is no coincidence that it is the layer that is completely missing from the public debate. We count megawatts at the gate; We rarely count whether those megawatts inside actually do work.

3. Value of the workload

Not every workload needs to be assessed by a government in terms of content. However, clients and operators can provide insight into which categories are supported:

- Care
- Energy
- water management;
- Mobility
- Science
-
- Defence and security
- Industry
- Government
- General commercial applications

This makes it possible to link infrastructure policies to national and European priorities without the need for the government to approve individual algorithms.

This is where the sharpest objection to my entire argument comes in, because it is guaranteed to come: isn't measuring "value of the workload" just a disguised way to let the government designate winners? And won't such a classification immediately become gameable — won't everyone soon call their workload "care" or "energy"?

Both objections are justified and both are solvable, provided you design the instrument well. The yardstick does not indicate individual applications right or wrong; He makes categories transparent, so that policy and purchasing can steer at the portfolio level without assessing one algorithm. It is the difference between a quality mark and a censor. And gameability is not overcome with nicer definitions but with verification: a workload classification that is not linked to independently verifiable technical data (which models, which data streams, which customers) is a marketing label, not a measurement value. Therefore, value classification in this framework never stands alone, but always alongside the hard, verifiable levels 1, 2, and 4. Value without proof does not count.

4. Contribution to the energy system

Data centers should not be treated exclusively as fixed electricity demand. AI workloads, in particular, offer opportunities to shift computational tasks in time and location.
Relevant are therefore indicators such as:

- Available flexible power
- Participation in congestion management
- Interruptibility during peak load
- Use of local storage
- Alignment with wind and solar energy production
- avoided peak load
- Grid capacity fed in or avoided

Flexibility must be demonstrable. Not the technical possibility, but the actual system service provided should count. This distinction is not academic. The grid is partly *administratively* full: contracted capacity that is not fully used cannot be reissued by the grid operator. A data center that demonstrably returns unused space or shifts peak load thus de facto expands the grid without the need for a cable. It is precisely that service provided — not its promise — that belongs in the yardstick.

5. Digital sovereignty

A data center on Dutch soil is not automatically sovereign infrastructure.
To do this, questions must be answered about:

- Ownership of the infrastructure
- control over operational processes
- Applicable jurisdiction
- Access to data and metadata
- Encryption key management
- Portability of data and models
- Dependence on one cloud or chip provider
- Ability to move workloads to another infrastructure

- Model ownership, fine-tuning results and usage data

There is a dimension to the latter that often remains out of the picture and that fundamentally exacerbates the question of sovereignty. Ruben rightly brings it in through what Satya Nadella calls the *Reverse Information Paradox*: those who have their data and work processes processed by someone else's model pay twice — once with money, and once with their own knowledge that the model absorbs along the way, through prompts, corrections and traces of use. Translated to our strongest sectors: train the greenhouse data, the delta data or the medical images via non-European foundation models, and in ten years' time our domain knowledge will be in those models, while we have a subscription to them.

This is not a fringe phenomenon, but goes to the heart of why sovereignty should be *measurable*. It is not enough to know where the data is; You must be able to demonstrate where what you have learned lands — with you, or with your future competitor. A mature sovereignty indicator therefore not only reveals ownership and jurisdiction, but also whether the IP that arises during use accumulates within its own control. That is the hinge between sovereignty and the layer of evidence later in this piece.

The tightened national cloud policy shows that the government is also taking dependence and extraterritorial legislation into account more emphatically. For core tasks of critical government organizations, the use of cloud services from providers that fall under legislation outside the EU or EEA is not recommended.

But sovereign storage without sufficient sovereign computing power is only a partial solution. Data can then be stored in Europe, while processing, model development and inference still remain dependent on non-European infrastructure. Having your own vault full of raw materials that you then have processed elsewhere is not sovereignty but outsourcing with a European label.

From voluntary promises to performance agreements

Many data center projects promise residual heat, flexibility, sustainable electricity and employment. The problem is not that these contributions are impossible. The problem is that they are often insufficiently made concrete during licensing and political decision-making.

Therefore, permits and public support must increasingly become performance-based.

For example, a location can be judged on:

- Maximum water use
- minimum use of residual heat, provided that a customer and heat network are available
- Demonstrable flexibility during congestion
- Minimal energy efficiency
- reporting on IT capacity utilisation

- Circularity plan for servers, chips and cooling installations
- capacity for European or public workloads
- Periodic, independently verifiable reporting

This does not create a ban on data centers, but a clear social quid pro quo.
A condition such as "residual heat must be used" only works if the demand side is also organised. Without a heat network, customers and a sound business case, even technically usable heat cannot be sold. This requires area development in which energy infrastructure, heat consumption and digital capacity are jointly planned. The original article also points out that supply-side obligations have little effect if the customer is not organised: residual heat rarely stops at the source, but at the networks, the customers and the business case around it. Whoever imposes the obligation must organize the customer. That is precisely why a performance requirement without a yardstick is an empty shell: you can only judge on performance if you can also determine that performance independently.

Modular infrastructure is needed in addition to gigafactories

The European debate is currently strongly dominated by AI Factories and very large GPU clusters. This scale is necessary for training the largest models, but does not solve all issues.
Many European organisations need the following:

- shielded inference capacity
- fine-tuning on your own data
- Predictable capacity
- Low latency
- Local legal control
- Linking with local renewable energy
- Scalability without directly reserving hundreds of megawatts

Therefore, in addition to large-scale central infrastructure, modular and regional capacity is also needed.
A fair objection is appropriate here, because it comes from the corner of the scale economy: is modular, regional capacity not less efficient per euro than one large cluster, and is there no risk of fragmentation? The answer is that efficiency per euro is not the only axis on which you assess infrastructure. Latency, jurisdiction, phased investment and energy siting are equally realistic design criteria, and on those axes a modular layer wins. It is therefore not about replacement but about complementarity: the largest models are trained centrally where power is abundant; The inference, the fine-tuning and the data-sensitive workloads are placed closer to the user. A layered architecture is not an inefficient version of one gigafactory — it is the answer to issues that onegigafactory cannotsolve.

A possible Dutch infrastructure layer consists of modular data center concepts such as SOLOdc for general digital infrastructure and SoloAI for high-power AI workloads. SoloAI is designed as a modular AI infrastructure with approximately 1.5 MW of IT capacity per unit, high rack densities, and liquid cooling. This allows capacity to be added in phases and placed closer to available energy, heat consumers and users.

This does not replace an AI gigafactory. It forms a complementary layer between central training clusters and the end user.

The value of such a modular architecture lies in:

- Shorter realisation times
- Phased investments
- Placement in locations with available energy
- Support for regional inference
- Linking to energy hubs and heat consumers
- Prevent each application from becoming dependent on one central infrastructure

And, last but not least, modular, measurable units lend themselves ideally to the layer of evidence that follows. A defined unit with its own energy, cooling and IT instrumentation is simply easier to make transparent than an amorphous multi-tenant building. Measurability is therefore not just a policy wish; It's also a design choice that you can build into the infrastructure itself.

No credible sustainable AI without a layer of evidence

The physical infrastructure is only one part of the solution. An independent layer of measurement and accountability is also needed.

That's where the role of a platform like Zirrow lies: combining data on energy, cooling, water, CO_2, IT load, workloads and compliance in real time.

Such a layer of evidence must go beyond a sustainability dashboard. It must make visible relationships between:

- de workload
- the infrastructure used
- energy consumption
- The source and timing of the electricity
- the amount of useful compute
- the water and CO_2 impact
- the flexibility provided
- the use of residual heat
- the applicable governance and jurisdiction

This can demonstrate whether an AI application not only functions, but also meets agreed performance.

This is relevant for operators, but also for customers. A hospital, bank, ministry or industrial company will soon have to be able to assess which digital infrastructure is behind an AI service and what impact is associated with it.

Measurability thus becomes a purchasing criterion.**

To show what that means: suppose a hospital purchases an AI service for image analysis. Today, it assesses that service on functionality and price. In a world with a layer of evidence, it also assesses where the inference is about, under which law the patient data falls, whether what has been learned sticks within its own environment, what CO_2 impact belongs to each study, and whether the supplier can demonstrate those figures independently. The purchasing question shifts from "does it work and what does it cost?" to "does it work, what does it cost socially, and who keeps control?". That is exactly the shift that the whole debate needs — and it starts not with the legislator, but with the purchasing organization that asks for evidence.

Here my argument touches on an inconvenient truth from Ruben's afterword, and it deserves to be made explicit. A yardstick is never stronger than the data underneath it. The figures with which we are having this debate are based on reports that have only recently become partly public, and although large energy consumers are legally obliged to declare their consumption, this is still not happening everywhere, or only halfway. In other words, even before we get to the richer levels of Useful Compute Value, the most basic layer is already faltering. This is not an argument against a yardstick — it is the strongest argument fororfor. The foundation we are asking for does not need to wait for anything: the legal basis for level 1 already exists, it just needs to be respected and monitored. And for a sector that likes to boast of its sustainable character, openness about its own consumption is the cheapest defence of its own credibility.

Towards a Dutch compute agenda

A mature national compute agenda must bring together technology policy, energy policy and industrial policy.

This requires five choices.

First, make a policy distinction between hyperscale cloud, colocation, AI training, inference, and edge infrastructure. What now falls under one reflex requires different considerations and different indicators.

Second: organize the demand. Governments, healthcare institutions, knowledge institutions and companies can make European compute projects financeable through long-term purchase commitments. The lesson of offshore wind is inexorable here: infrastructure without organized demand stalls, no matter how good the technology is. Supply is necessary, demand is decisive.

Third: develop a national yardstick for useful compute value, in which energy, water, CO_2, circularity, sovereignty and economic output come together. This is the computing power monitor that Ruben also asks for, but with a defined measurement architecture underneath instead of a good intention.
Fourth, plan data centers as part of the energy system. Select locations based on available electricity, grid flexibility, heat draw, connectivity, and water availability.
Fifth: build a layered infrastructure. Combine European training clusters with national and regional inference capabilities and local edge solutions.

Conclusion

The data center debate does not have to be between proponents who justify every energy use and opponents who see every server hall as wasteful.
The right question is indeed what we want to produce with our computing power. But there is also a second question:
Can we demonstrate that this computing power is used efficiently, sustainably, sovereignly and in a socially valuable way?
The Netherlands has the knowledge, digital connectivity, energy infrastructure and strong economic sectors to occupy an important position in the European intelligence economy.
But that only works if we connect three layers:

- **Ambition**: what intelligence do we want to develop and use ourselves?
- **Infrastructure**: where and under whose control do the models run?
- **Prove** what value and impact are actually realized?

Ruben has sharply placed the first layer and mapped the second. The third layer is what this piece is about, and without it, the first two remain unproven. An ambition without evidence is a slogan; infrastructure without evidence is a gamble; And proof without the other two is accounting without purpose. Only when the three come together does a debate arise that goes beyond the meter reading.
Not every kilowatt-hour compute is automatically right or wrong. The value is determined by what we do with it, how efficiently we do it and who ultimately retains control over the intelligence produced. The difference between a slogan and a strategy is whether we can demonstrate that.
That is the debate that the Netherlands must now have.
**This article is a response to: Rubenswindow, [The data center debate: from right-versus-error to intelligence-ambition](https://rubenswindow.substack.com/p/het-datacenter-debat-van-goed-versus), Substack, July 19, 2026.*