The Electricity Under AI Labour

A build note on the AI Labour Power Budget, an energy-equivalent scenario tool for comparing configured AI-worker profiles with annual electricity denominators.

2026.07.13 15:29 DENNIS HEDEGREEN BUILD V1.0
HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-ELECTRICITY-UNDER-AI-LABOUR/

AI labour is usually counted in agents, tasks, jobs, and productivity.

The physical layer arrives later, if it arrives at all.

That is backwards.

If companies, governments, and institutions try to run AI systems as labour at scale, the question is not only what the software can do. It is also what the scenario requires underneath it: electricity, data centres, cooling, grid connections, land, transformers, cables, water, and political choices.

The AI Labour Power Budget does not model that entire system. It does something smaller and more inspectable.

It takes a configured AI-worker energy profile and compares it with an annual electricity denominator.

The result is an energy-equivalent scenario.

It is not a count of real workers. It is not a deployment forecast. It is not a claim that a country has spare electricity or grid capacity. It says nothing about whether the configured profiles would perform useful work.

The tool exists to keep one question visible:

How much annual electricity sits underneath this version of the AI-labour claim?

Four steps, one boundary

The public tool has four steps.

`Country` selects the jurisdiction and the annual electricity metric used as the denominator.

`Worker` configures the profile: watts, runtime in hours per week, and an optional infrastructure-overhead factor.

`Budget` places electricity share on one side and configured profile count on the other.

`Report` carries the scenario, calculation basis, source evidence, data status, and reference assumptions into a downloadable PDF.

That sequence matters. It stops the large profile number from arriving before the reader can see what produced it.

The profile itself is deliberately narrow. Its annual energy is calculated as:

```
watts × hours per week × 365/7 × overhead
```

That produces annual kilowatt-hours or megawatt-hours for one configured profile. The tool can then read the equation in either direction:

- choose a share of the selected annual electricity denominator and calculate the equivalent number of configured profiles
- choose a target number of configured profiles and calculate the annual electricity share that target would require

No productivity value enters the equation. No human-equivalence factor enters it. No employment number enters it.

That is not an omission to hide. It is the model boundary.

The first bug was in the public reading

The early tool could calculate a large profile number, but the electricity share underneath it did not land with enough force.

A reader could see the output and silently complete the sentence as:

> This country can support this many AI workers.

The calculation had not established that.

“Can support” carries claims about available generation, grid connections, peak load, buildout, permits, cooling, system reliability, and political allocation. An annual-energy comparison cannot settle any of them.

The tool needed a stricter sentence:

> Assigning a stated share of the selected annual electricity denominator to this scenario would provide an equivalent yearly energy budget for a stated number of configured AI-worker profiles.

That sentence is less exciting.

It is also true to the calculation.

The interface became clearer when Step 3 stopped behaving like a one-way calculator and started behaving like an equation. Electricity share can drive the profile count. The profile target can drive the required share. The reader can pressure both sides instead of receiving one polished answer.

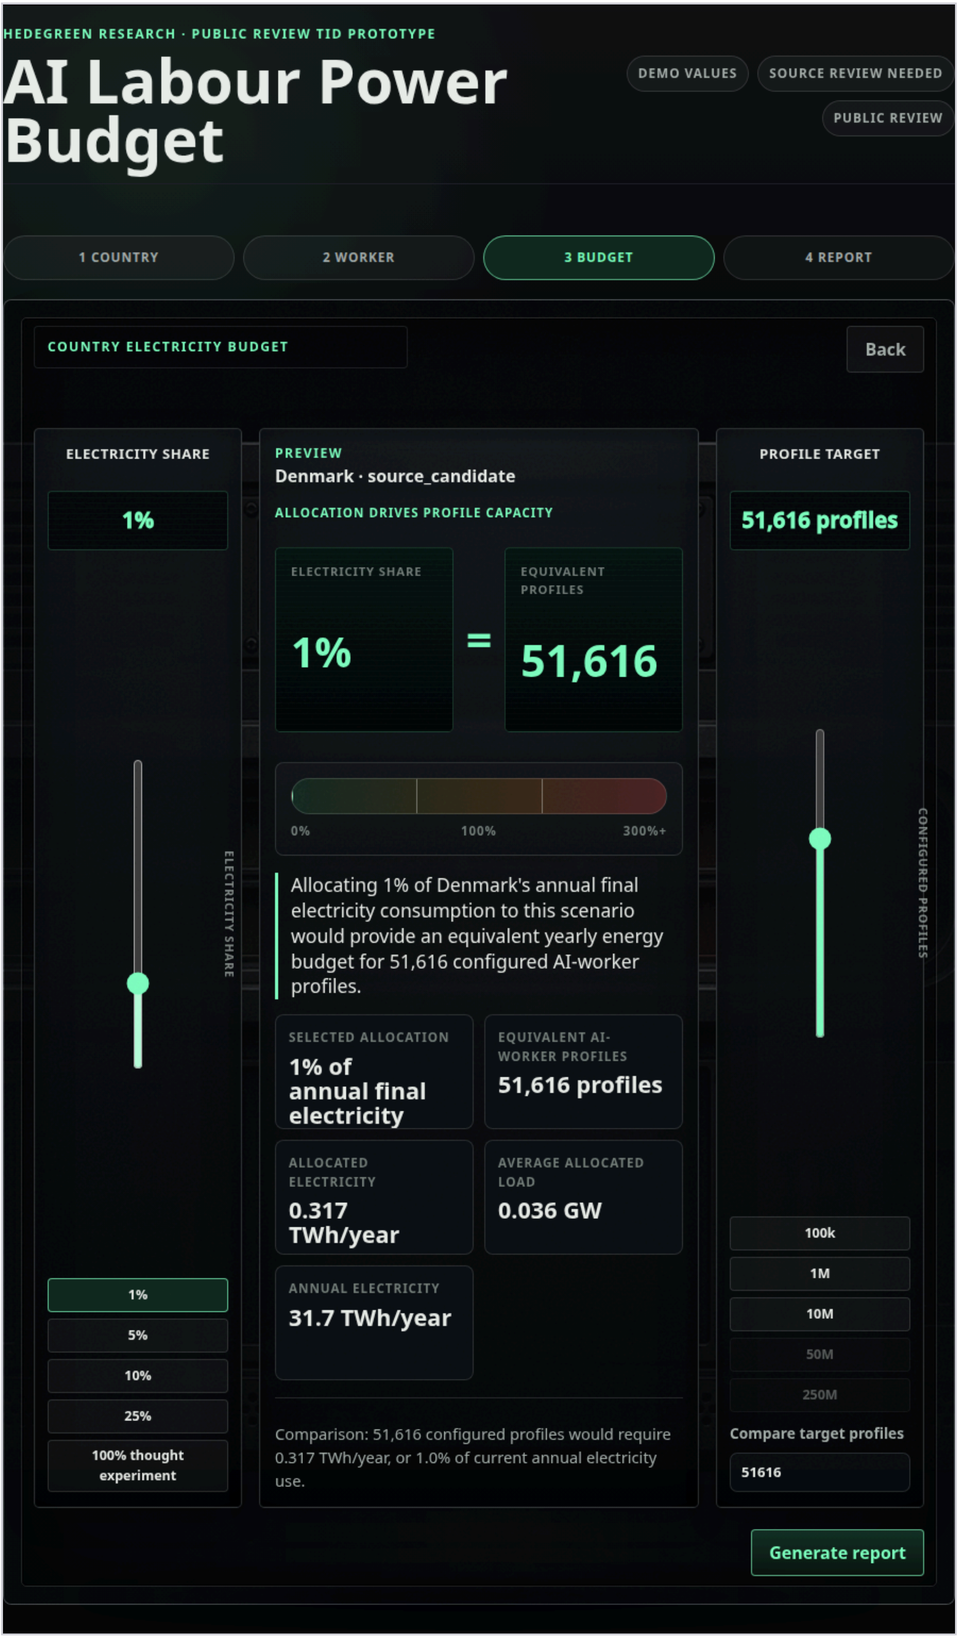

Step 3 keeps both directions of the scenario in one instrument: electricity share can drive the energy-equivalent profile count, or a profile target can drive the required annual energy. The Denmark value shown is a source candidate, not evidence of spare grid capacity.

This was the real design problem. The server surface could not merely look serious. The control, number, unit, and caveat had to remain attached to the same physical-looking instrument.

Otherwise the machine would become theatre.

Above 100 percent is not an allocation

The equation creates a specific language problem when a target passes 100 percent of the selected annual electricity denominator.

At 1 percent, 10 percent, or 100 percent, the tool can describe a scenario allocation. The percentages are not recommendations. They are positions inside the thought experiment.

Above 100 percent, “allocation” stops making sense.

A country cannot allocate 212 percent of its current annual total inside the normal meaning of the word. The interface therefore changes mode. It describes a **required-energy multiple** and labels the result as a **beyond-total-electricity thought experiment**.

That is not a mathematical change. The calculation continues.

The claim changes.

The result now means that the configured target would require annual energy equal to more than the full selected denominator. It does not mean the jurisdiction possesses, can build, or should allocate that energy.

Public calculators often fail in exactly this place. The formula keeps working, so the language is allowed to keep working too.

It should not.

Data status belongs beside the number

The tool uses three explicit data modes:

- `demo` for unverified interface placeholders
- `source_candidate_review` for candidate records that can be inspected and used in review scenarios without being promoted
- `promoted_runtime` for records that have passed an explicit promotion gate

The current EU, U.S., and World data used by the public-review instrument remain in `source_candidate_review` mode. None of it has been promoted as runtime data.

The distinction matters even though the candidate values carry official-source evidence and have passed source-value checks.

The EU candidate uses Eurostat's 2023 annual final consumption of electricity.

The U.S. candidate uses 2024 EIA retail electricity sales to ultimate customers, which excludes direct use.

The World candidate uses the United Nations Statistics Division's 2023 total final consumption of electricity.

Those are not identical metrics. The interface and report keep their denominator labels separate instead of flattening them into a generic claim about “national electricity”.

Promotion would also not transform a value into universal truth about a grid. It would mean that a reviewed record had been accepted for the tool's declared runtime purpose. The annual-energy boundary would remain.

This is why source status cannot live only in a repository note or at the bottom of the page.

It is part of the result.

The PDF has to survive outside the interface

The report is the most portable part of the tool, which also makes it the easiest part to misuse.

A screenshot may lose its footer. A number may be copied into a presentation. A PDF may be forwarded to someone who never sees the interactive explanation.

So the report cannot export only the large profile count.

It carries:

- the selected jurisdiction and denominator metric
- the configured watt, runtime, and overhead profile
- the annual energy per configured profile
- the calculation basis
- the scenario share or required-energy multiple
- the source name, year, retrieval date, URL, and status
- the `candidate not promoted` boundary where it applies
- the energy-equivalent scenario warning
- the assumptions behind reference comparisons

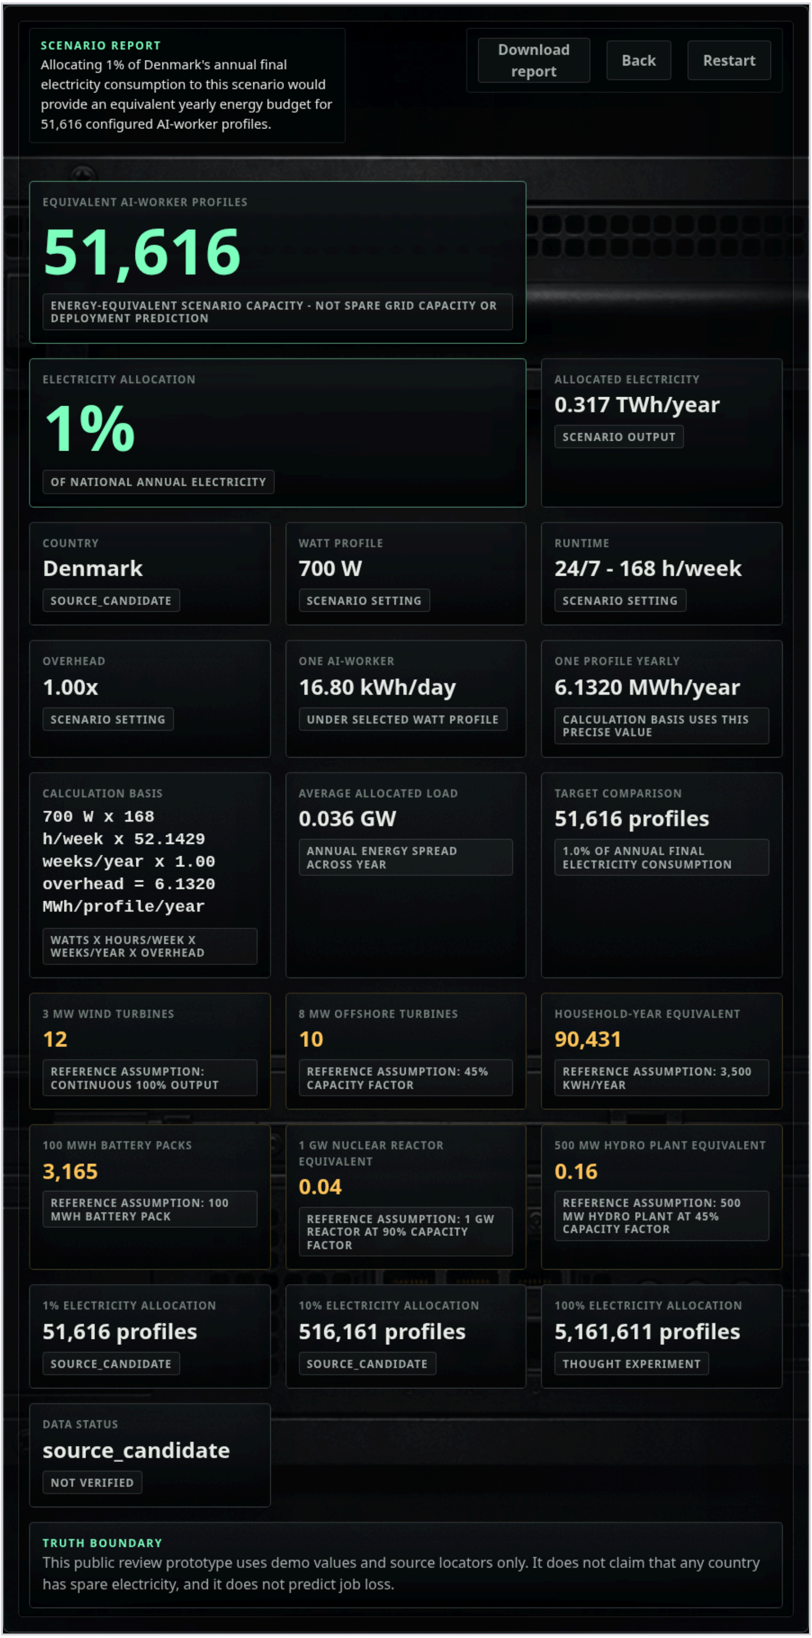

The Report view keeps the calculation basis, reference-assumption labels, and source_candidate / not verified status beside the output. The downloaded PDF carries the source name, metric, year, retrieval date, URL, and candidate-not-promoted boundary with it.

The reference comparisons need the same discipline.

Household-years, wind turbines, offshore turbines, battery packs, nuclear output, and hydro output can make a large annual-energy number easier to read. They can also create false equivalence.

A battery stores energy; it does not generate it. A turbine's annual output depends on an output assumption or capacity factor. A household-year value is a chosen reference, not a universal household fact. A power rating in megawatts is not itself an annual energy value in megawatt-hours.

The tool therefore labels these as reference assumptions. They are scale aids, not new evidence about the scenario.

What the instrument still cannot answer

The AI Labour Power Budget is not a grid model.

It does not model peak demand, hourly load shape, transmission constraints, connection queues, redundancy, storage, imports, cooling, water, land, construction time, power prices, carbon intensity, or future generation.

It does not estimate model quality, useful output, human substitution, wages, or job loss.

It does not establish that one configured profile corresponds to one person.

And it does not make 100 percent of annual electricity — or any other selected share — politically or physically available.

Those limits are not hidden weaknesses around the tool. They define the question the tool is qualified to ask.

If someone proposes one million AI agents, the tool cannot tell us whether those agents will work. It can ask which watt, runtime, and overhead assumptions sit behind the proposal and what annual energy those assumptions imply.

If a government talks about national AI capacity, the tool cannot tell us what should be built. It can separate an annual-energy scenario from the much stronger claim that grid capacity exists.

That is enough to improve the conversation.

A public instrument, still under pressure

The AI Labour Power Budget is online as a public-review instrument.

The server-panel dramaturgy is visible. The four-step calculation can be tested. The report can be downloaded. The source candidates and their status can be inspected.

Public does not mean promoted or verified runtime data.

Public means the method, wording, data boundaries, and failure points can now be challenged in the same room as the result.

That is the standard this tool has to keep.

The product is not the largest number it can produce.

The product is the number, the equation, and the boundary that prevents the number from claiming more than it knows.

AI labour was never weightless.

It had an electricity denominator.

The job of this instrument is to keep it visible.

RELATION MEMORY

SYSTEM REQUIREMENT The Tool Learned to Say Not Yet · Shares the rule that an instrument must expose what its evidence cannot yet establish.

NEARBY TID Is Still Becoming a Tool Surface · Places the public-review tool inside TID's wider need for visible runtime and data-status boundaries.

RETURNS TO AI Should Be Infrastructure, Not Authority · Returns to AI as physical and civic infrastructure without turning a scenario output into authority.

CANONICAL URL: HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/THE-ELECTRICITY-UNDER-AI-LABOUR/
PDF VIEW GENERATED: 2026-07-13T13:31:20.450Z