

The Water Problem That Isn’t:

The Upstream Binding Constraint And Strategies on Mitigation

Paul R. Leonard

The argument over data centers and water is becoming increasingly contentious. Communities are being asked to accommodate rapidly growing computational infrastructure while confronting rising water costs, aging treatment systems, constrained electricity supplies, and the infrastructure investment required to support both.

It is tempting to frame this as a new problem created by artificial intelligence. It isn't. AI and data centers are making an older physical relationship much harder to ignore: **water and energy are inseparable infrastructure problems.**

The planet is not running out of water. Earth contains an essentially closed hydrological system with an enormous inventory of water distributed among oceans, ice, groundwater, surface water, and the atmosphere-- 332.5 million cubic miles of H_2O, give or take. The problem is that the water available to humans at a particular location, at a particular time, in a particular physical state and at a required level of purity is not necessarily where we need it.

We already know how to change that.

We know how to desalinate seawater. We know how to purify contaminated water. We know how to reclaim wastewater, pump water across great distances, move it uphill, store it, heat it, cool it, and separate it from dissolved materials. None of those technologies is particularly mysterious.

What limits their deployment is largely the cost of doing it.

That cost is fundamentally energetic.

We don't pay for water itself

We pay for the energy, infrastructure, and technology required to alter water's location, state, or quality.

Desalination is a good example. The oceans contain an effectively inexhaustible supply of water for human purposes, but seawater is not freshwater. Converting it requires energy and infrastructure. Moving the resulting water inland or uphill requires more energy. Treating wastewater requires energy. Pumping, cooling, heating, storing, and distributing water all require energy.

The same logic applies to many of the resources modern civilization treats as scarce.

Energy allows us to extract and refine minerals, manufacture materials, cultivate and process food, heat and cool buildings, transport goods, purify water, and operate the infrastructure through which all of those things become useful to humans.

This is why Thermoeconomics treats energy as a **meta-constraint**. Energy is not simply one resource competing with other resources. It is a principal means by which humans alter the physical state, location, and usability of those resources.

When energy is expensive or constrained, many things appear scarce.

When energy becomes abundant, some of those apparent scarcities become engineering problems.

AI is exposing the Energy-Water Nexus

Data centers make this relationship unusually visible because they concentrate enormous computational demand into a relatively small physical footprint.

A data center requires electricity, cooling, communications, backup capacity, land, transmission and distribution infrastructure, and often substantial water resources. The energy system supporting the data center can also have its own water requirements, depending on how electricity is generated.

Consequently, asking how much water a data center consumes is only part of the question.

A more useful question is:

What total physical infrastructure is required to deliver useful computation at this location?

The answer can include generation, transmission, cooling, water treatment, water supply, wastewater handling, and the infrastructure required to support all of them.

This is where the issue becomes contentious. A large computational load can compete with existing human demand not only for electricity and water, but for the capital and infrastructure required to make both available.

The resulting political argument is usually framed as a conflict between communities and data centers.

At a deeper level, it is a conflict over the allocation of constrained physical capacity.

We are also making energy the hard way

The energy side of the problem deserves the same scrutiny.

Modern electricity generation relies heavily on two broad physical strategies: capture an environmental energy flux or exploit an energy gradient.

Solar and wind depend on relatively diffuse environmental fluxes. Their aggregate resource may be enormous in theory, but their power density and geographic distribution impose engineering consequences. Large amounts of land and infrastructure may be required to collect relatively modest continuous power.

Conventional thermal generation takes a different route. Coal, natural gas, geothermal and nuclear sources can provide much higher energy density, but electricity is generally produced by using thermal energy to establish a temperature gradient and then running a heat engine. That gradient is primarily dependent upon the limits of our infrastructure materials. Rankine-cycle architectures, including those used with conventional nuclear generation and proposed small modular reactors, necessarily reject a substantial fraction of the original thermal energy rather than converting it into electricity.

The issue is not that these technologies are useless. They have built modern civilization.

The issue is that our generation architecture carries substantial thermodynamic and infrastructural consequences before electricity ever reaches the load.

For a data center, the chain can then continue:

primary energy → electrical generation → transmission → computation → heat → cooling → water and infrastructure

We have built an extraordinarily sophisticated civilization around managing this chain.

The question is whether we are required to keep building it this way as other options emerge.

What if the gradient is not the resource?

This is where a very different energy concept becomes interesting.

At our company, we refer to the alternative we are investigating as **temporal-driven energy generation**, in contrast to conventional flux-driven or gradient-driven generation.

The distinction begins with the material architecture rather than the external energy source.

Nature already produces non-equilibrium steady states. We did not invent NESS. What we are learning to do is design NESS into functional systems and purposeful materials.

Cascade has already demonstrated first-principles design of non-equilibrium steady-state architectures in quantum catalysis. Our approach is based on designing material structure and excited-state behavior

around the intended function and environment of use rather than relying solely on empirical materials screening. Early applications have included falsification-oriented formulation design, systematic element substitution, operation at atmospheric pressure and moderate temperatures, and catalytic architectures capable of accumulating high-energy excited states. Those approaches have progressed into early commercial catalytic iterations.

The unresolved question is whether the same design methodology can be extended from quantum catalytic function to **quantum catalytic charge separation and durable, net-positive electrical generation**.

The generation technology has not already been demonstrated at useful scale, and mainstream adoption would not be justified until the claims of net work, durability, scalability, and reproducibility have been established.

The underlying architectural idea, however, is not that a new environmental flux must be found or that an increasingly large gradient must be engineered.

It is that a material system can be deliberately designed into a non-equilibrium state and that useful ordered energy may be extracted during its controlled temporal relaxation. The more time you have to work with ordered energy, the better.

Entropy, or disorder, is not eliminated. Thermodynamics does not get repealed. The proposed architectural difference is that useful work is sought during the engineered relaxation process rather than beginning with the conventional requirement to establish a macroscopic thermal gradient and pass energy through a heat engine. The point is to extend the time that ordered energy is available for harvest, as opposed to extending the upper limits of materials that might expand the temperature gradient.

If that approach ultimately works-- and the design model has already been proven in other applications, the significance would extend well beyond electricity generation.

A formal technical treatment of the NESS benchmarking and the Cascade evidentiary record is forthcoming.

Change the energy boundary and the water boundary changes with it

Suppose energy generation could become substantially less dependent on geographic energy fluxes and gradients.

The implications would not stop at the generator.

Energy could become available closer to where it is needed. Transmission requirements could change. Computational systems could become less dependent on centralized power infrastructure. Desalination could be performed closer to population centers or wherever water is needed. Wastewater could

become a more economically recoverable resource. Water could be pumped, purified, stored, and redistributed with a different cost structure.

None of those capabilities requires a new water technology.

We already have the technologies.

What changes is the cost of deploying them.

This is why the Energy-Water Nexus is more interesting than the current argument over whether AI uses “too much water.” The deeper question is whether we are allowing a particular generation architecture to define the practical limits of what can be done with an abundant physical resource.

The oceans are not going away. The water cycle is not being consumed. What we lack is inexpensive capacity to transform water into the form and location required by human systems.

Energy is what provides much of that capacity.

The data center is a stress test, not the whole problem

AI has simply made the relationship impossible to ignore.

The same physical logic applies to agriculture, housing, manufacturing, hospitals, municipalities, food production, wastewater treatment, remote communities, and ordinary households. Every one of them requires energy to transform physical resources into useful states.

That is why the water question ultimately leads back to a much larger question about energy.

If energy remains constrained by the density and geography of environmental fluxes, or by the thermodynamic cost of conventional heat-engine architectures, then every downstream transformation inherits those constraints.

If a new energy architecture can materially relax those constraints, the consequences could extend far beyond the electricity market.

The most interesting question may therefore not be how to make AI consume less water.

It may be whether we can change the energy architecture that makes water, computation, and so many other necessities compete for the same constrained infrastructure.

We know how to make seawater potable. We know how to reclaim wastewater. We know how to move water. We know how to build computational infrastructure.

The unresolved technological question is whether we can learn to generate useful electrical energy from a fundamentally different physical architecture.

If temporal-driven energy generation proves viable, the consequence would not simply be another source of electricity. It would be a potential change in the physical boundary within which humans can transform the resources already available to us.

The water is already here.

The question is how much energy we need to make it useful.

###

#Thermoeconomics #EnergyWaterNexus #DataCenterInfrastructure #NonEquilibriumSteadyState #QuantumCatalysis

About the Author

Paul R. Leonard is President & CEO of Cascade Catalysts, Inc., a Vancouver-based company developing quantum-effect catalysts and advanced materials for industrial applications.

His broader work examines seemingly disparate problems through a thermoeconomic lens, seeking structural relationships among energy, surplus, human agency, institutional behavior, and technological development.

Leonard is also the author of *Big Quantum Nerd*, a publication exploring structural clarity through thermoeconomics and non-equilibrium systems.

His work draws on applied experience across physics, energy, materials, complex systems, and organizational problem-solving, with an emphasis on identifying structural solutions that expand human maneuverability rather than merely optimizing existing constraints.