

PROMETHEUS RESEARCH SERIES

WHITE PAPER No. 2

POWERING INTELLIGENCE

The New Energy Architecture for AI

David W Tietgen

August 2026

GOLD PROOF - SUBSTANCE LOCKED / PRODUCTION REVIEW

Executive Summary

Evidence convention: numbered references support current facts and quantitative claims; unnumbered design rules and campus examples are PROMETHEUS analysis unless otherwise stated.

Artificial intelligence is creating an energy problem unlike anything the digital economy has encountered before.

Not because AI is the first industry to consume enormous amounts of electricity. It isn't.

Steel mills, aluminum smelters, chemical plants, refineries, transportation systems, and cities have all required immense quantities of energy.

AI is different because of the combination: scale, concentration, speed, reliability, growth - and one additional characteristic that may ultimately prove more important than all the others: some of its demand can think.

That last statement requires explanation.

A steel mill cannot normally decide that a portion of today's production should instantly move 1,500 miles because electricity is cheaper somewhere else. A city cannot postpone half of tomorrow morning's residential demand until midnight. But certain computational workloads can be delayed. Some can accelerate. Some can move.

That gives the Intelligence Age an energy capability previous industrial revolutions largely did not possess.

Before we exploit that capability, however, we need electricity. A great deal of it.

The International Energy Agency projects global data-center electricity consumption to more than double to roughly 945 terawatt-hours by 2030. AI is the most important driver of that increase.^[1]

The United States is simultaneously experiencing renewed electricity-demand growth. EIA expects developers to add a record 86 GW of utility-scale generation and storage during 2026 if planned projects are completed, including approximately 43.4 GW of solar and 24 GW of battery storage.^{[2] [2] [2]}

Those are enormous numbers. Yet they do not automatically solve the AI energy problem.

A megawatt can exist in the wrong place. It can exist at the wrong time. It can be trapped behind transmission congestion. It can be variable when the campus requires firmness. It can be inexpensive but unavailable for five years. It can require a transformer that has not been manufactured. Or it can arrive perfectly - only for another component of the infrastructure to become the bottleneck.

The Intelligence Age does not merely need more electricity. It needs a new architecture for producing, storing, delivering, and controlling electricity.

That architecture will not be built around one winner.

Not nuclear alone. Not natural gas alone. Not solar. Not wind. Not batteries. Not geothermal. Not the existing grid.

The system will be layered.

Firm generation. Variable generation. Short-duration storage. Long-duration resources. Grid capacity. Dedicated generation. Microgrids. Efficiency. Cooling. Computational flexibility.^[8]

Each resource performs a different job.

The revolutionary step comes when those resources stop operating independently. Energy becomes coordinated with computation. Then energy becomes programmable.

That is the subject of WP-02.

MOVEMENT I - THE NEW POWER PROBLEM

1. The Electricity Reversal

For decades, electricity demand in many advanced economies grew relatively slowly. The digital economy expanded inside that environment.

Servers became more efficient. Computing became more powerful. Software businesses could scale dramatically without their executives thinking like utility planners.

AI is reversing that relationship.

The physical infrastructure beneath advanced computing is becoming too large to ignore.

The IEA expects data-center electricity consumption to rise at roughly 15 percent annually between 2024 and 2030 in its base case - more than four times the growth rate of electricity consumption from other sectors. Accelerated servers, driven mainly by AI adoption, are projected to grow much faster still.^[3]

A technology company planning a gigawatt-scale campus is therefore no longer merely choosing servers. It is making decisions about the power system.

Where does the electricity come from? How quickly can it arrive? What happens during a grid emergency? What portion must be firm? How much storage is necessary? Can generation be built beside the campus? Can future energy technologies be added later?

These are industrial questions.

AI has entered the energy business whether technology companies intended to or not.

2. A Gigawatt Changes the Conversation

Consider a 1-GW AI campus.

If it operated continuously at that electrical load for a year, its energy consumption would be 8.76 TWh.

Actual operation would vary. But the scale is what matters.

At 50 MW, a facility is a large customer. At 500 MW, it is a major infrastructure project. At 1 GW, it begins to resemble a power-system participant.

This changes the conversation with the utility.

We need hundreds of megawatts - potentially a gigawatt or more - and we need to know when you can deliver it.

That is a very different customer.

In June 2026, FERC ordered all six regional grid operators under its jurisdiction to justify or reform the rules governing how data centers and other large loads connect to the transmission system. FERC explicitly framed the issue around faster large-load integration and speed-to-power.^[4]

That phrase deserves to become part of the vocabulary of the Intelligence Age.

3. The Cheapest Electricity May Be Too Late

Imagine two energy options.

Option A: electricity cost lower; available 2032.

Option B: electricity cost higher; available 2028.

Which is cheaper?

The answer cannot be determined from the electricity price alone.

If a multi-billion-dollar computing campus sits partially unused for four years waiting for Option A, the lower-cost electricity may produce the higher-cost intelligence.

This introduces a variable traditional energy procurement often underweights: time.

PROMETHEUS calls the resulting metric SPEED-TO-POWER.

Energy economics for AI therefore need to consider price, availability date, reliability, capacity factor, transmission, fuel, expansion potential, capital cost, and the economic value of delayed computation.

The cheapest megawatt-hour is not necessarily the cheapest system.

4. Capacity Is Not Energy

FIGURE 1 - POWER, ENERGY, AND FIRM CAPACITY

Three different questions: How fast can power be delivered? How much energy is available over time? How much dependable capacity can be counted on when the campus needs it?

One of the easiest mistakes in energy discussions is treating all megawatts as equivalent. They aren't.

A 1-GW solar installation may be capable of producing approximately 1 GW under favorable conditions. That does not mean it supplies 1 GW continuously.

A nuclear plant behaves differently. A gas turbine behaves differently. A battery is different again because it is primarily storing energy produced elsewhere.

The distinction is fundamental: power measures the rate at which energy is delivered; energy measures how much is delivered over time; and firm capacity describes dependable capability under defined conditions.

A 1-GW AI campus needs all three concepts. Ignoring the distinction produces bad architecture.

5. The Grid Problem

There is enough electricity in the United States to operate enormous amounts of computing. That statement can be true while a particular AI campus still cannot obtain the electricity it needs.

Why? Because electricity is geographically constrained.

Generation must connect through transmission. Transmission must have available capacity. Substations must exist. Transformers must be installed. Interconnection studies must be completed. Reliability requirements must be met.

The grid therefore creates an unusual condition: national energy abundance can coexist with local power scarcity.

This is one reason FERC's 2026 large-load proceedings matter.

The problem is not simply building more generators. It is integrating enormous new loads into a physical network that was not designed around their speed of arrival.

Never confuse electricity that exists with electricity that can reach the campus when it is needed.

MOVEMENT II - BUILD A PORTFOLIO, NOT A BET

6. Stop Asking Which Technology Wins

Energy debates often begin with the wrong question: Which technology should power AI?

Nuclear? Natural gas? Solar? Wind? Geothermal? Batteries?

The question assumes there must be one winner. There doesn't.

A gigawatt-scale intelligence campus has multiple energy problems. It needs electricity continuously, instantaneous stability, protection against outages, the ability to survive hours of constraint, economical energy during normal operation, expansion capacity, and increasingly power quickly.

Different technologies solve different portions of those requirements.

Which resource performs which job best?

That produces an architecture rather than an ideology.

7. The Grid - Diversity at Scale

The grid remains extraordinarily valuable.

A strong regional grid connects many generators, customers, fuel types, and locations. It allows diversity that would be extremely expensive for a single campus to reproduce independently.

PROMETHEUS therefore does not advocate abandoning the grid. It advocates changing the campus's relationship with it.

Instead of GRID -> PASSIVE CUSTOMER, think GRID <-> ACTIVE CAMPUS.

The campus imports electricity, but it may also possess generation, storage, flexible demand, sophisticated controls, and possibly export capability.

The grid remains the foundation. The campus becomes more capable.

8. Natural Gas - The Speed Layer

Natural gas belongs in a realistic near-term discussion because it solves several difficult problems well.

It is dispatchable, commercially mature, capable of producing large amounts of firm power, and in suitable locations may be developed faster than some alternatives.

That makes gas especially relevant to speed-to-power.

But gas has weaknesses: pipeline capacity, turbine availability, fuel-price volatility, emissions, methane leakage, and the risk that a long-lived facility outlives the assumptions that justified its original fuel.

PROMETHEUS therefore gives natural gas a specific role: bridge and firmness, not necessarily destination.

The principle is the same one applied throughout this paper: evaluate the resource that actually exists, identify the job it can perform, account for its full costs and constraints, and preserve the ability to replace it when a better option is ready.

PROMETHEUS therefore gives coal a narrower role than gas or nuclear: an existing-system and transition resource that must be counted honestly, especially where delaying retirement protects reliability while replacement capacity is built. That is different from making new coal construction the default architecture for AI.

For an AI campus, existing coal has real engineering strengths: dispatchable generation, on-site fuel inventories at many plants, established transmission connections, and large existing units that can contribute firm power. It also carries substantial disadvantages: high carbon emissions, conventional air pollutants, water and ash-management requirements, aging equipment, and weak economics for building new conventional coal plants in many U.S. markets.

The United States is moving in a different long-term direction: coal-fired generation is expected to decline as plants retire and other resources expand. But the decline is no longer a simple straight line. EIA reported that only 2.6 GW of U.S. coal capacity retired in 2025, far below what operators had originally planned, after multiple retirements were delayed or cancelled. Another 6.4 GW was scheduled to retire in 2026, although EIA warned that further delays were possible.

The IEA estimates that coal supplies roughly 15 percent of the electricity consumed by U.S. data centers today. Globally, coal is the largest single source of electricity physically consumed by data centers, at about 30 percent. In the IEA base case, natural gas and coal together meet more than 40 percent of the additional electricity demand from data centers through 2030.^[5]

Coal cannot be omitted from a serious 2026 energy architecture simply because relatively little new U.S. coal capacity is being proposed. It remains part of the electricity physically serving data centers today, and rising demand can change how quickly existing plants retire.

9. Coal - The Existing Firm Resource

10. A Bridge Must Lead Somewhere

If a campus uses gas to become operational sooner, its architecture should anticipate what happens later.

Perhaps additional grid capacity arrives. Perhaps nuclear becomes available. Perhaps geothermal matures locally. Perhaps storage economics improve. Perhaps processor efficiency dramatically lowers electrical demand.

The campus should not need to tear itself apart to accommodate those developments.

This is why PROMETHEUS emphasizes interfaces: electrical buses, switchgear, microgrid controls, generation interconnections, storage connections, cooling architecture, and software control.

Design those correctly and today's generator becomes a replaceable component. Design them poorly and today's generator becomes tomorrow's trap.

11. Nuclear - The Firmness Layer

Nuclear's attraction to AI is straightforward. AI campuses value electricity that is available continuously.

Nuclear plants can provide large quantities of firm power over long operating lives.

The challenge is equally clear. Large conventional nuclear projects can require substantial capital, complex regulation, extensive construction, and long schedules.

That creates an uncomfortable mismatch. AI moves quickly. Traditional nuclear construction does not.

So nuclear may solve the long-duration energy problem while struggling with the immediate speed-to-power problem.

That is why advanced nuclear matters.

12. Modular Nuclear

The most interesting aspect of advanced nuclear may not be that reactors become smaller. It may be that nuclear generation becomes incremental.

Imagine a campus built in stages: Compute Block 1, Block 2, Block 3, Block 4.

Now imagine energy capacity expanding in corresponding stages: Reactor Module 1, Module 2, Module 3.

The architecture begins to resemble the load.

That is potentially powerful.

But we must be disciplined about what is real today.

DOE reported several advanced reactor zero-power criticality milestones during summer 2026. Those are meaningful technical demonstrations, but they are not equivalent to commercially operating power reactors supplying gigawatt-scale AI campuses.^[7]

PROMETHEUS therefore classifies advanced modular nuclear as strategically important, technically advancing, commercially emerging - not solved.

13. Nuclear Should Not Operate Alone

Now combine nuclear with storage.

A reactor provides sustained energy. A battery provides rapid response. The grid provides regional diversity. Flexible computation provides demand response. The microgrid coordinates them.

This is much more interesting than arguing whether nuclear is better than batteries. They perform different jobs.

Conceptually: NUCLEAR -> MICROGRID <-> BATTERY <-> GRID -> COMPUTE.

The value lies in the combination.

14. Geothermal - The Quiet Candidate

Geothermal deserves more attention than it usually receives in AI discussions.

Its traditional limitation is geography.

Advanced geothermal attempts to broaden that map through improved drilling, engineered reservoirs, and techniques partly derived from oil-and-gas experience.^[13]

If that effort succeeds economically, geothermal could offer AI campuses something extremely attractive: firm electricity without continuous fuel delivery.

It also has a comparatively small surface footprint.

But advanced geothermal must prove drilling economics, repeatability, reservoir performance, development schedules, and financing at scale.^[13]

So geothermal belongs in the architecture as an emerging firm resource whose importance could grow substantially.

15. Solar - Abundance Has Value

Solar does not provide firm round-the-clock electricity by itself. That fact should not obscure what solar can do extremely well: produce enormous quantities of electricity during favorable daylight conditions.

U.S. developers plan approximately 43.4 GW of new utility-scale solar capacity during 2026.^[2]

For PROMETHEUS, the important question is not whether solar can independently power a 24/7 AI campus. It is what the campus should do when solar creates several hours of abundant electricity.

Charge batteries. Advance flexible AI model-training workloads. Perform batch computation. Reduce fuel consumption. Possibly export power.

The campus should recognize energy abundance and exploit it.

16. Wind - Different Time, Same Opportunity

Wind offers another source of variable electricity. Its production pattern often differs from solar.

Solar may create daytime abundance. Wind may produce strongly overnight. The two can therefore complement each other.

But neither should be confused with guaranteed firm capacity.

PROMETHEUS treats wind and solar as energy resources whose value increases when paired with flexibility.

Storage, flexible compute, geographic diversity, and transmission all increase their usefulness.

The architecture determines how valuable the resource becomes.

17. Batteries - The Electrical Shock Absorber

Battery deployment is accelerating rapidly.

For AI campuses, batteries can respond almost instantly, smooth rapid load changes, bridge short interruptions, shift electricity from one part of the day to another, reduce peak grid withdrawals, support microgrid islanding, and allow slower generators time to respond.

That makes the battery less like a fuel source and more like the electrical shock absorber.

18. Four Hours Is Not Four Days

Storage discussions become confused when duration is ignored.

Suppose a 1-GW campus has a four-hour battery. Its nominal stored-energy requirement at full output is 4 GWh.

That is enormous. But after four hours, the fundamental problem returns unless another resource takes over.

Now imagine trying to maintain 1 GW for four days: 96 GWh.

The scale changes dramatically.

Using a four-hour battery to solve a four-day problem is poor system design.

19. The Storage Ladder

FIGURE 2 - THE STORAGE LADDER

Milliseconds-minutes: power quality / UPS | Minutes-hours: peak reduction / bridging | Hours-overnight: energy shifting | Multi-day: prolonged shortages

PROMETHEUS organizes storage by job.

Milliseconds to minutes: power quality, ride-through, UPS, rapid fluctuations.

Minutes to hours: peak reduction, solar shifting, generator bridging, grid response.

Hours to overnight: energy shifting, extended grid constraints, renewable balancing.

Multi-day: rare prolonged shortages, fuel disruptions, weather events, extended grid failure.

Different technologies may dominate different layers. Lithium-ion does not have to win all of them.

20. Long-Duration Storage

The search for longer-duration storage includes flow batteries, sodium-based systems, iron-based chemistries, thermal storage, compressed air, pumped hydro, hydrogen, and other electrochemical or mechanical approaches.^[14]

Each comes with different efficiency, geography, capital cost, maturity, safety, and operational characteristics.

PROMETHEUS does not need to choose the winner today.

Duration defines the problem. Economics and reliability select the technology.

MOVEMENT III - THE POWER PLANT WE DO NOT BUILD

21. Efficiency Becomes Capacity

Imagine improving total campus electrical efficiency by 5 percent.

At a 1-GW facility, 5 percent equals 50 MW.

No new generator was constructed. No transmission line was added. No fuel was burned. But the campus now needs 50 MW less electricity to produce the same useful result.

This is why efficiency becomes an infrastructure resource at scale.

Percentages become power plants.

PROMETHEUS calls this virtual capacity. It is not literally generation, but capacity that does not have to be built can have enormous economic value.

22. The Processor Is Part of the Energy System

Processor efficiency matters far beyond the server rack.

Suppose a future accelerator produces twice as much useful computation per watt.

The consequence could mean fewer generators, less transmission, smaller substations, less cooling, fewer batteries, less water consumption, less land, and lower capital requirements.

The boundary between semiconductor engineering and energy engineering begins to disappear.

The processor becomes an energy technology.

23. Cooling Is Also an Energy Technology

The same applies to cooling.

Every watt entering computing equipment eventually becomes heat. That heat has to go somewhere.

Cooling architecture therefore determines how much useful computation can be produced from the electricity entering the site.

Liquid cooling becomes especially important as rack power densities rise.

But the water discussion requires precision.

A cooling system can circulate enormous quantities of water without consuming the entire circulated volume. Closed-loop systems repeatedly recirculate coolant.^[10]

Fresh-water consumption depends on the heat-rejection method, evaporation, blowdown, local climate, treatment requirements, and system design.

Water withdrawal is not the same as water consumption.

A PROMETHEUS campus should minimize unnecessary fresh-water consumption through appropriate combinations of closed-loop cooling, reclaimed municipal or industrial water, dry or hybrid heat rejection where practical, and intelligent thermal management.^[10]

24. What If Heat Is a Product?

Now go one step further.

The campus produces enormous quantities of heat. We normally call it waste heat. But is it always waste?^[11]

Potential uses may include district heating, industrial processes, greenhouses, water heating, or other nearby thermal loads.

The economics depend heavily on temperature, distance, climate, and the existence of a customer for the heat.

A data center in the desert and one beside a northern city should therefore not necessarily use identical thermal architectures.

Before treating an output as waste, ask whether another part of the system can use it as an input.

MOVEMENT IV - ENERGY BECOMES PROGRAMMABLE

25. The Microgrid Changes Everything

Until now, we have discussed resources individually: grid, gas, nuclear, geothermal, solar, wind, batteries, efficiency, cooling.

The microgrid is where they become a system.

A sophisticated campus microgrid can coordinate generation, storage, grid imports, critical loads, noncritical loads, and island operation.

But PROMETHEUS adds one more variable: computation itself.

26. Not Every Calculation Is Equally Urgent

Some AI computation is time-sensitive. A user expects an inference response now. A safety-critical industrial system cannot wait three hours for cheaper electricity.

But other work may have scheduling flexibility.

Examples include portions of AI model training, fine-tuning, evaluation, synthetic-data generation, batch analytics, and other non-real-time computational jobs.

The exact flexibility depends on the workload.

But the principle is real: energy planning acquires a load that can sometimes negotiate with time.

27. Return to the 150-MW Problem

WP-01 introduced a temporary 150-MW deficit at a 1-GW AI campus. We return to that same problem here - not to repeat it, but to examine what the energy architecture must actually coordinate.

The illustrative response remains:

50 MW - battery discharge

40 MW - flexible AI model-training workloads temporarily deferred

30 MW - reduced cooling and electrical overhead

30 MW - suitable computation shifted to another campus

TOTAL - 150 MW

In WP-01, the example demonstrated that a shortage does not always require an equal amount of new generation. In WP-02, the lesson is more specific: the energy system must be designed in advance so these resources can respond together.

The battery must have sufficient charge and duration. Cooling controls must be capable of reducing auxiliary demand without threatening equipment. Flexible AI workloads must already be classified by priority. Communications must be capable of moving suitable computation. And the microgrid must coordinate the response without compromising critical loads.

The arithmetic is illustrative. The architecture is the point.

28. Energy Orchestration

Now imagine the microgrid knows electricity price, grid constraints, battery state, generator availability, weather, solar production, wind production, cooling demand, fuel availability, and maintenance status.

The compute system knows job urgency, deadlines, hardware requirements, network requirements, data location, latency, and whether a workload can move.

The two systems communicate.

ENERGY ORCHESTRATION <-> COMPUTE ORCHESTRATION.

That is the central technical idea of WP-02.

The energy system no longer sees one enormous fixed load. It sees categories of demand. And the compute system no longer sees electricity as an infinite background utility. It sees energy conditions.

29. Predict Tomorrow Before Tomorrow Arrives

Suppose tomorrow afternoon will bring extreme heat. Regional demand will rise. Solar production will decline toward evening. Grid reserves may tighten.

A conventional campus waits. PROMETHEUS prepares.

Tonight: batteries charge strategically; flexible computation advances; maintenance schedules adjust; fuel reserves are checked.

Tomorrow morning: solar begins rising; additional flexible workloads accelerate.

By afternoon: batteries are ready; firm generation is positioned; noncritical computation can reduce; remote campuses are evaluated; critical workloads remain protected.

The system has not created energy. It has created options.

And options are one of the most valuable forms of resilience.

MOVEMENT V - THE 1-GW PROMETHEUS ENERGY CAMPUS

30. Start With the Load

We can now return to our blank sheet.

We need to supply a 1-GW intelligence campus.

Do not begin by choosing a technology. Begin by classifying the load.

How much is absolutely critical? Firm but schedulable? Interruptible briefly? Movable in time? Movable geographically?

Then examine the site: grid capacity, generation options, fuels, renewable resources, climate, water, fiber-optic communications capacity, workforce, and expansion.

Only then do we build the energy portfolio.

31. The Layered Energy Stack

FIGURE 3 - THE 1-GW PROMETHEUS ENERGY STACK
GRID + FIRM GENERATION + VARIABLE ENERGY + STORAGE + FLEXIBLE COMPUTE + ORCHESTRATION

A PROMETHEUS campus could contain six functional layers.

Layer 1 - Grid: regional diversity, market access, backup, import/export capability.

Layer 2 - Firm generation: natural gas, nuclear, geothermal, hydro where available, other proven dispatchable resources.

Layer 3 - Variable energy: solar, wind, other geographically appropriate resources.

Layer 4 - Storage: fast batteries, multi-hour batteries, longer-duration resources where justified.

Layer 5 - Flexible compute: time-shiftable workloads, geographically movable workloads, efficiency, cooling optimization.

Layer 6 - Orchestration: the control system coordinating everything above it.

That is the PROMETHEUS energy stack.

32. There Is No Universal Fuel Mix

PROMETHEUS should not prescribe a false-precision national fuel recipe.

A Texas campus and a Washington State campus should not have identical portfolios. Neither should Arizona and Pennsylvania.

Resources, climate, grid conditions, fuel, water, regulation, and construction schedules differ.

Standardize the architecture. Localize the portfolio.

33. Build for N+Options

Traditional critical infrastructure often uses redundancy: N+1, 2N, backup generators, multiple electrical feeds.

PROMETHEUS retains redundancy but adds another idea: N + OPTIONS.

If one pathway fails, the campus should possess multiple different responses.

Grid constraint? Battery. Generation. Flexible compute. Another campus.

Generator failure? Grid. Storage. Reduced noncritical demand.

Extreme heat? Cooling strategy changes. Workloads shift.

Energy abundance? Charge storage. Accelerate computation.

Resilience becomes more than duplicate equipment. It becomes operational choice.

34. The Energy Campus Should Be Upgradeable

A campus built in 2028 may still operate in 2050. The processors inside it certainly will not be the same.

Neither should we assume the energy technologies will be.

The campus should anticipate new battery chemistries, new reactor technologies, better geothermal, better power electronics, better cooling, better processors, better communications, and technologies not yet identified.

The physical architecture must outlive the technologies initially installed within it.

Build the interfaces before you know all the machines that will eventually connect to them.

MOVEMENT VI - THE ECONOMICS OF INTELLIGENCE

35. Levelized Cost Is Not Enough

Energy analysis frequently compares technologies using cost per megawatt-hour. That remains useful.

But AI infrastructure requires a larger calculation.

Suppose one resource costs \$45/MWh and another \$70/MWh. The first looks cheaper.

But what if it requires five additional years of transmission development? What if the second can operate three years earlier?

What is three years of lost computational output worth? What is the capital cost of idle servers? What is the opportunity cost of delayed AI services? What is the cost of unreliability? What is the value of future expandability?

The energy decision therefore becomes a system-capital decision.

PROMETHEUS proposes evaluating total cost of delivered, reliable computational energy, and ultimately cost per unit of useful intelligence.

36. A More Expensive Generator Can Produce Cheaper Intelligence

This sounds contradictory. It isn't.

Imagine a generation option with higher electricity cost but faster construction, higher reliability, lower transmission requirements, better expansion capability, and easier integration with storage.

Its electricity may cost more. But the campus may begin productive operation earlier and experience fewer constraints.

The resulting intelligence may be cheaper.

This is why energy procurement and computing economics can no longer be separated. The correct unit of optimization is the system.

37. The Grid Should Value Flexibility

If a campus can reduce load during grid stress, that flexibility has system value.

If it can charge batteries during abundance, that has value.

If it can avoid demanding its full interconnection capacity during a constrained period, that may reduce the need for some upgrades.

The customer begins to become a grid resource.

38. Bring Your Own Power

We may therefore see a new development model.

Instead of a data-center developer arriving with only a load request - We need 1 GW - the developer may increasingly arrive with: We need 1 GW; here is dedicated generation; here is storage; here is our flexibility profile; here is our grid connection; here is our expansion plan.

That changes the utility conversation.

The campus becomes partially responsible for creating the infrastructure required to support itself.

This will not work everywhere. Nor should it replace sound regional grid planning. But at very large scales, it may become increasingly common.

MOVEMENT VII - THE TECHNOLOGIES BEYOND TODAY

39. Advanced Nuclear

Advanced nuclear could become one of the most consequential energy developments for AI.

Could. That word matters.

The attractive characteristics are obvious: firm generation, high energy density, small land requirements, potential modularity, and long operating life.

The unresolved questions include commercial cost, manufacturing, licensing, fuel availability, construction schedules, and repeatability.

PROMETHEUS should follow the evidence rather than the enthusiasm.

If modular reactors demonstrate reliable, economical serial deployment, they could fit the architecture beautifully.

Until then: prepare the interface. Do not assume the outcome.

40. Advanced Geothermal

Geothermal deserves the same discipline.

Improved drilling could dramatically expand viable resources. Oil-and-gas expertise may transfer. Firm generation near load could be extremely valuable.

But commercial performance across diverse geology must be demonstrated.

So again: prepare the architecture. If geothermal wins locally, use it.

41. Long-Duration Storage

Long-duration storage could eventually change how much firm generation a campus needs.

But economics, round-trip efficiency, materials, cycle life, and duration matter enormously.

A technology that is excellent for twelve hours may be poor for seven days.

There will probably not be one universal storage winner. That's fine. PROMETHEUS does not require one.

42. Hydrogen and Synthetic Fuels

Hydrogen receives substantial attention because it can store energy and potentially provide fuel for power generation.

But every conversion has costs. Electricity creates hydrogen. Hydrogen is stored and transported. Then perhaps converted back into electricity. Efficiency losses accumulate. Infrastructure is required.

That does not make hydrogen useless. It means hydrogen should be used where its characteristics justify those costs.

If multi-day or seasonal storage becomes the problem, hydrogen may deserve consideration. If the problem is a two-hour evening peak, there may be much simpler answers.

43. Space-Based Energy and Computation

We should address this because it represents the outer boundary of the architecture.

Space offers extraordinary solar availability. That has produced concepts involving space-based solar power and, increasingly, computation performed in orbit.

But significant challenges remain: launch cost, thermal rejection, radiation, maintenance, hardware replacement, communications, orbital infrastructure, and economics.

PROMETHEUS therefore classifies orbital energy/computation as frontier - not a solution upon which today's terrestrial 1-GW campus should depend.

But conceptually it reinforces an important idea: energy geography can shape compute geography.

WP-03 will follow that idea much further.

MOVEMENT VIII - THE PROMETHEUS ENERGY DOCTRINE

44. What Should Be Built Now

If we were making a real infrastructure decision today, PROMETHEUS would favor technologies according to maturity rather than excitement.

AVAILABLE NOW: grid power; natural gas where appropriate; solar; wind; utility-scale batteries; existing nuclear; hydroelectric resources; modern microgrids; advanced cooling; demand management. ^[8]

SCALING RAPIDLY: hybrid renewable/storage systems; large campus microgrids; liquid-cooled high-density computing; grid-interactive storage; flexible large-load arrangements. ^[8]

EMERGING: advanced geothermal; commercial small modular reactors and microreactors; some long-duration storage technologies; sophisticated energy-aware workload scheduling. ^[13]

FRONTIER: large-scale orbital computation; space-based energy architectures; other technologies not yet commercially demonstrated.

The categories will change. That is the point. The architecture should survive the change.

45. What We Should Not Do

Do not wait for one miracle technology.

Do not assume the grid alone can always keep pace.

Do not abandon the grid unnecessarily.

Do not build permanent dependence on gas simply because gas may be fast.

Do not assume advanced nuclear is commercially ready merely because demonstrations are encouraging.

Do not count variable generation as firm generation.

Do not make batteries solve every duration.

Do not waste potable water when appropriate reclaimed alternatives exist.

Do not treat cooling as an afterthought.

Do not treat efficiency as a minor operating improvement.

Do not assume every computation must run immediately.

Do not design thirty-year infrastructure around five-year technological assumptions.

Do not confuse an energy technology with an energy architecture.

46. The PROMETHEUS Energy Hierarchy

FIGURE 4 - THE PROMETHEUS ENERGY HIERARCHY
REDUCE -> MOVE IN TIME -> MOVE IN SPACE -> STORE -> FIRM -> VARIABLE ENERGY -> ORCHESTRATE

When a campus faces an energy requirement, ask the questions in this order.

1. Can we reduce it? Better processors, better software, better cooling, better utilization.
2. Can we move it in time? Schedule flexible computation around energy availability.
3. Can we move it in space? Run suitable computation somewhere energy is more available.
4. Should we store energy? Choose storage according to duration and economics.
5. What must be firm? Build or contract the generation and grid capacity necessary for the remainder.
6. Where can variable energy reduce system cost? Solar, wind, other resources.
7. How do we orchestrate everything? Continuously.

That final question is the difference between a portfolio and a system.

47. The Energy Operating System

The ultimate PROMETHEUS energy architecture is not defined by its fuel. It is defined by its ability to choose.

At every moment, the system evaluates: grid? generator? battery? reduce load? advance computation? delay computation? move computation? charge? discharge? recover heat? import? export? prepare for tomorrow?

The system becomes adaptive.

Not because energy itself has become intelligent. Because the infrastructure controlling it has.

That is the point where energy infrastructure becomes intelligent energy infrastructure.

CONCLUSION - POWERING INTELLIGENCE

The first electrical grids were built to power lights. Then motors. Factories. Cities. Air conditioning. Telecommunications. Computers. Each new use changed the grid.

Artificial intelligence may do so again.

But this time something unusual is happening. The electricity system is becoming more computational at precisely the same moment computation is becoming more energy-intensive. The two systems are converging.

That creates risk. If energy infrastructure cannot keep pace, processors sit idle. If grids are expanded inefficiently, consumers may bear unnecessary costs. If developers choose technologies solely for speed, they may create long-lived dependencies. If they wait for perfect future technologies, the Intelligence Age will not wait with them.

So we need another approach.

Build what works. Preserve options. Layer technologies according to function. Use the grid. Strengthen it. Add generation where necessary. Use gas where its speed and firmness justify it. Use nuclear where its firmness and longevity justify it. Develop advanced nuclear. Develop geothermal. Exploit solar and wind abundance. Use batteries for the problems batteries solve well. Develop longer-duration storage for longer problems. Improve processors. Improve cooling. Recycle water where practical. Recover useful heat where economics permit. And make computation participate in the energy system.

Then connect everything through orchestration.

The result is not an energy source. It is an energy architecture.

And that distinction may determine who can build intelligence at scale.

How intelligently can we use every watt, every generator, every battery, every grid connection, every hour, and every movable computation available to us?

That is the future of power.

Not simply more electricity.

More intelligent electricity.

THE PROMETHEUS ENERGY RULE

Generate what must be generated.

Store what is economical to store.

Move what is cheaper to move.

Delay what can safely wait.

Exploit abundance when it appears.

Protect what cannot fail.

Preserve the ability to change.

REFERENCES AND CURRENT EVIDENCE

Sources below support the quantitative, technological, and current-status claims in this Gold Proof. PROMETHEUS analysis, design rules, illustrative campus allocations, and forward-looking judgments are identified in the manuscript as analysis rather than established fact.

- [1] International Energy Agency (IEA). Energy and AI: Executive Summary. Data-center electricity consumption is projected to more than double to about 945 TWh by 2030; AI is identified as the most important driver of the increase. Accessed August 2026. <https://www.iea.org/reports/energy-and-ai/executive-summary>
- [2] U.S. Energy Information Administration (EIA). “New U.S. electric generating capacity expected to reach a record high in 2026.” February 20, 2026. Reports 86 GW of planned utility-scale additions, including 43.4 GW solar and 24 GW battery storage. <https://www.eia.gov/todayinEnergy/detail.php?id=67205>
- [3] International Energy Agency (IEA). Energy and AI: Energy Demand from AI. Discusses rapid data-center electricity-demand growth and the unusually concentrated nature of data-center loads. Accessed August 2026. <https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai>
- [4] Federal Energy Regulatory Commission (FERC). “FERC Launches Aggressive Targeted Action to Speed Large Load Integration.” June 18, 2026. Orders six regional grid operators to justify or reform rules governing interconnection of data centers and other large loads. <https://www.ferc.gov/news-events/news/ferc-launches-aggressive-targeted-action-speed-large-load-integration>
- [5] International Energy Agency (IEA). Energy and AI: Energy Supply for AI. Estimates the current U.S. data-center electricity mix at over 40% natural gas, 24% renewables, about 20% nuclear, and about 15% coal; evaluates supply growth through 2035. Accessed August 2026. <https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai>
- [6] U.S. Energy Information Administration (EIA). “U.S. coal-fired generating capacity retired in 2025 was the least in 15 years.” April 13, 2026. Reports 2.6 GW retired in 2025 versus 8.5 GW initially planned, with multiple delays and cancellations. <https://www.eia.gov/todayinenergy/detail.php?id=67427>

- [7] U.S. Department of Energy (DOE). Reactor Pilot Program announcements, June-July 2026. DOE reported four distinct advanced reactor zero-power fueled criticality demonstrations by early July 2026. These demonstrations did not constitute commercial electricity generation. <https://www.energy.gov/articles/departement-energy-celebrates-fourth-criticality-ahead-july-4th-goal>
- [8] U.S. Department of Energy, Office of Electricity. “Microgrids, Large Electric Loads & Grid Support: How to Leverage Microgrids to Support Utilities and Large Load Customers.” June 3, 2026. Describes microgrids as a promising option for faster, reliable data-center buildout relative to conventional grid expansion alone. <https://www.energy.gov/oe/articles/microgrids-large-electric-loads-grid-support-how-leverage-microgrids-support-utilities>
- [9] U.S. Department of Energy. Long-Duration Energy Storage. DOE defines long-duration energy storage as systems capable of delivering electricity for 10 or more hours and supports demonstrations across multiple technology families. <https://www.energy.gov/cmei/oced/long-duration-energy-storage>
- [10] U.S. Department of Energy, Federal Energy Management Program. “Cooling Water Efficiency Opportunities for Federal Data Centers.” Describes recirculating chilled-water loops, direct liquid cooling, cooling-tower water management, and opportunities for water reuse. <https://www.energy.gov/cmei/femp/cooling-water-efficiency-opportunities-federal-data-centers>
- [11] U.S. Department of Energy, Federal Energy Management Program. “Technology Changes, but Energy Efficiency Principles Remain Steadfast in Data Center Design.” Discusses liquid cooling, water-use efficiency, and reuse of data-center waste heat. <https://www.energy.gov/cmei/femp/articles/technology-changes-energy-efficiency-principles-remain-steadfast-data-center>
- [12] Xin Lu and Qianwen Xu. “Grid Integration of Gigawatt-Scale AI Data Centers under Connect-and-Manage.” arXiv, May 2026. Research model separates frontier training, batch training, and inference loads and evaluates battery-supported workload flexibility. This is research evidence, not proof of universal commercial performance. <https://arxiv.org/abs/2605.14109>
- [13] U.S. Department of Energy, Office of Geothermal. “Enhanced Geothermal Systems Demonstration Projects.” DOE announced up to \$171.5 million in February 2026 for next-generation geothermal field tests and drilling intended to reduce commercialization risk. <https://www.energy.gov/hgeo/geothermal/enhanced-geothermal-systems-demonstration-projects>
- [14] U.S. Department of Energy, Office of Electricity. “New Report Showcases How Innovation Can Fast Track Affordable Energy Storage.” Reviews electrochemical, chemical, mechanical, and thermal long-duration storage families, including flow, sodium, zinc, hydrogen, compressed air, pumped hydro, and thermal systems. <https://www.energy.gov/oe/articles/new-report-showcases-how-innovation-can-fast-track-affordable-energy-storage>

NEXT - PROMETHEUS WHITE PAPER No. 3

COMPUTE FOLLOWS ENERGY

The Geography of the Intelligence Age

WP-03 begins where WP-02 deliberately stops.

If energy varies by time and geography, and some computation can vary by time and geography, then perhaps the future AI network should not simply move electricity toward computers.

Sometimes it should move the computation toward the energy.

That leads us into fiber-optic communications, latency, workload mobility, geographically distributed AI campuses, power-aware computing, satellite communications, regional energy arbitrage - and eventually the frontier question of orbital computation.