

PROMETHEUS RESEARCH SERIES

WHITE PAPER No. 3

COMPUTE FOLLOWS ENERGY

The Geography of the Intelligence Age

David W Tietgen

August 2026

Executive Summary

Artificial intelligence is beginning to expose a weakness in the way modern infrastructure has been planned.

For more than a century, industrial systems generally moved energy toward machines. Coal moved by rail. Oil moved through pipelines and ships. Natural gas moved through pipelines. Electricity moved across transmission lines.

Computing introduces another possibility.

Information can move.

Some computation can move with it.

That does not mean electricity transmission becomes unnecessary. It means the Intelligence Age has acquired an additional infrastructure choice.

At times, the rational answer will be to move more electricity toward the computer. At other times, the rational answer may be to move a flexible computational workload toward available electricity. Often the correct answer will combine both.

The International Energy Agency reports that data centers are unusually concentrated geographically. Nearly half of U.S. data-center capacity is concentrated in five regional clusters, making local grid effects much larger than the sector's national electricity share might suggest. [1]

That concentration is becoming harder to sustain as individual AI campuses grow from tens of megawatts toward hundreds of megawatts and, in announced or planned cases, gigawatt scale.

Meanwhile, the Federal Energy Regulatory Commission is changing the regulatory conversation. Its June 2026 actions on large-load interconnection explicitly recognize that some large loads can change consumption rapidly and that flexible service may allow faster connection while avoiding unnecessary grid upgrades. [2]

Industry practice is moving in the same direction. Google reported in March 2026 that it had incorporated 1 GW of data-center demand-response capacity into long-term utility contracts, using the ability to shift or reduce selected machine-learning workloads. [3]

These developments point toward a new infrastructure principle:

When electricity is scarce in one place and computation is flexible, the system should at least ask whether moving the computation is cheaper than moving the electrons.

That question is the foundation of WP-03.

But it comes with limits. Data movement consumes energy. Fiber capacity is finite. Latency matters. Massive datasets are expensive to replicate. Some AI workloads are tightly synchronized. Some applications must respond in milliseconds. Cybersecurity, data sovereignty, hardware availability, water, land, workforce, and network resilience constrain geography.

So PROMETHEUS does not propose that fiber replaces transmission.

It proposes something more useful:

Move electrons when electrons are the better thing to move. Move information when information is the better thing to move. Move computation when computation is the better thing to move. Design the system so it can choose.

This is the geography of the Intelligence Age.

MOVEMENT I - GEOGRAPHY RETURNS

1. The Cloud Was Never Weightless

The word cloud encouraged a useful abstraction and a dangerous illusion.

To the user, computing appeared locationless. A file existed 'in the cloud.' An application ran 'online.' A model answered from somewhere beyond the screen.

But every calculation occurs in a physical place.

It occupies a processor. The processor sits in a rack. The rack sits in a building. The building occupies land. It receives electricity, rejects heat, uses communications networks, depends on workers and suppliers, and exists under a legal jurisdiction.

AI makes that physicality impossible to ignore.

Once a campus demands hundreds of megawatts, location becomes a first-order engineering decision rather than a real-estate afterthought.

PROMETHEUS therefore rejects the idea that the cloud erased geography.

The cloud hid geography. AI is bringing it back.

2. Why Data Centers Cluster

Data centers cluster for rational reasons.

Fiber networks already exist in certain corridors. Customers and enterprises are nearby. Cloud ecosystems develop around established regions. Skilled labor accumulates. Utilities become familiar with large digital loads. Land and permitting practices become understood. Network effects reinforce the original choice.

Northern Virginia is the clearest U.S. example, but the pattern is broader. The IEA notes that nearly half of U.S. data-center capacity is concentrated in five regional clusters. [1]

Clustering improves connectivity and operating familiarity.

But clustering also concentrates electrical demand, land use, water concerns, transmission requirements, and community impact.

The feature that created efficiency in the first generation of cloud infrastructure can therefore create scarcity in the next.

3. The Local Grid Becomes the Constraint

National electricity statistics can be misleading.

The United States can possess ample generation while a specific substation, transmission corridor, or utility territory lacks the capacity to connect another 500 MW load.

This is the distinction WP-02 established between electricity that exists and electricity that can reach a campus.

WP-03 adds another response.

If the grid cannot economically deliver another block of power to Location A, perhaps the next block of compute should not be built at Location A.

That is not always possible. But at gigawatt scale, it must become part of the siting calculation.

4. The New Siting Equation

The old data-center siting equation emphasized land, taxes, fiber, and electricity price.

The AI-era equation is larger:

deliverable power; speed-to-power; fiber capacity; network diversity; latency; water and cooling; land; skilled labor; construction capacity; generation options; storage; regulatory environment; community acceptance; natural-hazard exposure; security; and room to expand.

These variables interact.

Cheap land without fiber is not cheap.

Cheap electricity without transmission is not available.
Abundant power without cooling may be unusable.
Excellent power and fiber without skilled labor can delay construction and operation.
Low latency without expansion capacity may create a dead end.
PROMETHEUS therefore treats site selection as systems engineering.

MOVEMENT II - FIBER IS THE OTHER GRID

5. What We Mean by Fiber

WP-01 introduced fiber-optic infrastructure. WP-03 needs to define it plainly.

Fiber-optic cable carries information as pulses of light through extremely thin strands of glass or plastic. Compared with older copper communications systems, fiber can move vastly more data over long distances with low signal loss.

But 'fiber' is not one thing.

A campus needs local fiber within buildings, high-capacity links between buildings, metro connections, long-haul routes to other cities, connections to cloud regions and internet exchanges, and ideally physically diverse paths so that one construction accident or equipment failure does not isolate the site.

The fiber network is therefore not merely the internet connection.

For distributed AI, fiber becomes the transportation system for computation.

6. Bandwidth Is Capacity

Bandwidth measures how much information a communications link can carry over time.

AI increases bandwidth requirements in several ways.

Training datasets can be enormous. Model checkpoints are large. Distributed inference can generate substantial east-west traffic. Replication between campuses consumes network capacity. Moving a workload may require moving not just instructions but data, model weights, state, and results.

This creates an important correction to the phrase 'move the computation.'

Sometimes moving the computation means sending a relatively small job to a remote location where the necessary model and data already exist.

Sometimes it means moving terabytes or petabytes of supporting data.

Those are entirely different infrastructure problems.

7. Latency Is Distance Made Visible

Bandwidth answers how much data can move.

Latency answers how long movement takes.

Light travels extraordinarily fast, but not instantaneously. Fiber routes are longer than straight-line distance, and signals pass through optical equipment, routers, switches, security systems, and software stacks.

For a nightly training job, tens of milliseconds may be irrelevant.

For an interactive AI assistant, latency matters more.

For industrial control, autonomous systems, financial trading, or safety-critical applications, it can matter enormously.

Therefore not all compute can follow energy equally.

The farther computation moves from the user or machine it serves, the more latency becomes part of the energy decision.

8. The Fiber Constraint

Moving computation away from a constrained grid does not eliminate infrastructure.

It substitutes one infrastructure burden for another.

A region with abundant electricity but inadequate long-haul communications may not be a practical AI location.

That is why energy maps and fiber maps should increasingly be studied together.

The most valuable future AI sites may not simply be where electricity is cheapest.

They may be where abundant electricity intersects high-capacity, diverse communications.

MOVEMENT III - CLASSIFY THE COMPUTATION

9. Not All Compute Is the Same

The phrase 'AI workload' hides important differences.

Some workloads are rigid. Some are flexible.

Some require specialized hardware. Some can run on several accelerator types.

Some must stay near a dataset. Some can operate anywhere.

Some have deadlines measured in milliseconds. Others have deadlines measured in hours or days.

WP-02 used this distinction to make demand flexible in time.

WP-03 extends it into geography.

10. Four Geographic Classes of Compute

PROMETHEUS proposes four practical classes.

CLASS A - LOCATION-BOUND. Workloads that must remain near a user, machine, regulated dataset, specialized facility, or unique hardware environment.

CLASS B - LATENCY-SENSITIVE. Workloads that can move only within a limited geographic radius because response time matters.

CLASS C - REGIONALLY FLEXIBLE. Workloads that can move among multiple campuses within a country or broad region if data, hardware, and network capacity are available.

CLASS D - HIGHLY FLEXIBLE. Batch processing, portions of training, evaluation, synthetic-data generation, and other work that may tolerate substantial temporal and geographic movement.

The exact classification will vary by application.

But the architecture improves the moment the system stops treating every megawatt of compute as equally immovable.

11. Training Is Not One Workload

Even AI model training should not be treated as a single category.

A tightly synchronized frontier training run may require thousands of accelerators communicating continuously and may be difficult to move after it begins.

But preparation, data processing, evaluation, checkpoint analysis, fine-tuning, experimentation, and many batch tasks may have more flexibility.

This is why PROMETHEUS uses the phrase 'flexible AI model-training workloads' carefully.

It does not mean that a giant training run can casually jump across the continent every afternoon.

It means portions of the broader training pipeline may be schedulable or locatable according to infrastructure conditions.

12. Inference Has Its Own Geography

Inference - using a trained model to produce answers or actions - often has different constraints.

Interactive inference benefits from proximity to users.

But not all inference is interactive.

Batch inference, document processing, media analysis, forecasting, and background services may tolerate more distance.

A May 2026 Microsoft Research paper modeled geographically distributed AI data centers and found that flexible inference routing can create substantial operational value when compute scheduling, local generation, batteries, grid interaction, latency, and carbon constraints are optimized together. The work is research rather than proof of universal commercial performance, but the architecture is directly relevant. [4]

The lesson is not that inference should always move.

It is that inference should be classified.

MOVEMENT IV - RETURN TO THE 150-MW PROBLEM

13. The Problem We Already Know

WP-01 introduced a temporary 150-MW deficit at a 1-GW AI campus.

WP-02 returned to the same deficit and showed how batteries, flexible training, cooling optimization, and shifted computation could respond together.

WP-03 asks one narrower question:

What does 'shifted computation' actually require?

Suppose 30 MW of suitable work can be moved from Campus A to Campus B for two hours.

At Campus A, local demand falls by 30 MW.

But Campus B must possess at least 30 MW of available electrical and computational headroom.

The network must have capacity.

The model and data must be accessible.

The hardware must be compatible.

The workload must tolerate the added latency.

The receiving campus must not create a new local grid problem.

Only when all of those conditions are satisfied has geography actually solved the 30-MW problem.

14. The 30-MW Transfer

This is where the PROMETHEUS concept becomes concrete.

Campus A experiences a temporary energy constraint.

The orchestration system identifies 30 MW of Class C or Class D workloads.

It checks Campus B, Campus C, and Campus D.

Campus B has spare compute but its grid is also constrained.

Campus C has available power but insufficient accelerator capacity.

Campus D has available power, compatible accelerators, adequate network capacity, and acceptable latency.

The workload moves to Campus D.

Campus A's shortage falls from 150 MW to 120 MW before another generator starts.

This is not theoretical magic. It is resource scheduling across geography.

Google already reports using demand response that can shift selected machine-learning workloads across times and locations, and in 2026 said it had incorporated 1 GW of such demand-response capacity into U.S. utility contracts. [3]

15. What Was Actually Moved?

The electrons did not travel from Campus D to Campus A.

The computation moved to where electrons were available.

That distinction is the heart of WP-03.

Electricity transmission solves a physical energy-location mismatch by moving electrical energy.

Compute mobility can solve some mismatches by moving the demand.

Neither method is universally superior.

The value comes from having both.

MOVEMENT V - THE CONTINENTAL AI CAMPUS

16. Stop Thinking of the Campus as One Place

The next conceptual step is larger.

What if an AI campus is not one campus?

What if the true infrastructure unit is a network of campuses operating as one computational system?

One site may have excellent nuclear power.

Another may have abundant wind.

Another may have solar and batteries.

Another may sit near users and provide low-latency inference.

Another may possess water and land suitable for large training systems.

Fiber connects them.

Software coordinates them.

Energy conditions influence where suitable work runs.

The campus becomes continental before it becomes orbital.

17. Regional Energy Diversity

America's energy system is geographically diverse.

The Pacific Northwest has major hydro resources. The Great Plains possess exceptional wind resources. The Southwest has exceptional solar resources. Natural gas is abundant in several producing regions. Nuclear capacity is concentrated in other areas. Coal remains important in portions of the Midwest and interior. Reuters' August 2026 mapping of the U.S. power system illustrates how strongly the national energy mix still follows geography. [5]

This diversity is usually treated as a grid-planning challenge.

For distributed AI it can also become an asset.

Different campuses can be optimized around different regional strengths rather than forcing every site into the same energy recipe.

18. Follow the Surplus

Electricity systems experience periods of relative abundance and scarcity.

A windy night in one region may produce surplus energy while another region is constrained.

Midday solar can create abundance in another location.

Hydro conditions vary seasonally.

Generator outages change local conditions.

Transmission congestion can isolate cheap energy from distant demand.

A geographically distributed compute network can respond to some of those differences.

The goal is not to chase every hourly price movement.

The goal is to recognize sustained or predictable infrastructure advantages and use workload flexibility where the benefit exceeds the cost of movement.

19. The Virtual Transmission Line

Consider a simplified case.

Region A has 100 MW of unused compute capacity but scarce electricity.

Region B has 100 MW of unused compute capacity and abundant electricity.

If a suitable 50-MW workload can move from A to B, the effect on Region A resembles adding 50 MW of supply.

No transmission line carried 50 MW between the regions.

Demand moved instead.

PROMETHEUS calls this effect a virtual transmission line.

The phrase must be used carefully. It does not replace physical transmission and it cannot serve immovable loads.

But for flexible digital demand, it describes a real planning effect: communications and compute mobility can sometimes relieve an electrical constraint.

20. Why We Still Need Transmission

This concept should never become an excuse to underbuild the grid.

Most economic activity cannot teleport its demand.

Homes, factories, hospitals, transportation, and many data-center workloads need electricity where they physically exist.

Transmission also provides regional diversity and allows generation to serve many customers.

PROMETHEUS therefore supports a both-and strategy.

Build transmission where transmission creates broad value.

Build fiber where information movement creates broad value.

Build distributed compute where computation can use both.

Then optimize across all three.

MOVEMENT VI - THE NETWORK BECOMES AN ENERGY ASSET

21. Communications Has an Energy Value

A fiber route is normally valued by bandwidth, latency, reliability, and revenue.

WP-03 adds another dimension.

If that route allows 20, 50, or 100 MW of flexible compute to relocate during a local grid constraint, the communications network has indirectly acquired energy-system value.

This does not mean the fiber 'produces' electricity.

It means communications can increase the flexibility with which electrical demand is placed.

That can defer or reduce some peak infrastructure requirements.

22. Network Diversity Is Power Resilience

Once workload mobility becomes part of the energy architecture, communications reliability becomes part of power resilience.

A campus that depends on shifting workloads cannot rely on one fiber route.

Long-haul paths should be physically diverse where possible.

Carrier diversity matters.

Network equipment requires backup power.

Routing must fail gracefully.

Cybersecurity becomes critical because an attacker who can disrupt workload movement may indirectly constrain the campus's energy response.

The energy plan and network plan therefore converge.

23. Data Gravity

One of the strongest limits on compute mobility is data gravity.

Large datasets attract computation because moving the computation to the data can be easier than moving the data to the computation.

If a training corpus occupies petabytes and changes continuously, duplicating it across several campuses can be expensive and slow.

Model weights, checkpoints, caches, and application state create similar constraints.

The solution is architectural rather than rhetorical.

Replicate what is economical to replicate.

Keep location-bound data where it belongs.

Pre-position models and datasets likely to be needed elsewhere.

Move only workloads whose data requirements make movement worthwhile.

24. The Cost of Moving Bits

Moving information consumes electricity too.

Routers, optical amplifiers, switches, transceivers, storage systems, and receiving servers all consume power.

So a system that moves a workload to save 30 MW locally must ask what the transfer costs elsewhere.

In many cases, the communications energy may be small compared with the compute energy being relocated.

In other cases - especially when huge datasets must be copied repeatedly - the network and storage cost may erase much of the benefit.

PROMETHEUS therefore requires full-system accounting.

Never call computation movable until you have priced the movement.

MOVEMENT VII - LAND, WATER, LABOR AND COMMUNITY

25. Energy Is Not the Only Geography

An energy-rich location is not automatically an AI location.

WP-01 taught us not to overlook land and skilled labor. WP-03 makes them part of the siting architecture.

A 1-GW campus requires substantial land not only for buildings but potentially for substations, generation, storage, cooling infrastructure, security setbacks, roads, construction staging, and future expansion.

Land also determines proximity to communities and competing uses.

A remote site may have cheap acreage and excellent power but face severe construction-labor and logistics constraints.

26. Skilled Labor Is Infrastructure

Electricians, lineworkers, pipefitters, controls specialists, network engineers, data-center technicians, construction managers, cooling specialists, security teams, and many other trades and professions are required to build and operate large AI campuses.

At gigawatt scale, workforce availability can affect schedule just as surely as a transformer shortage.

Training pipelines, regional universities, apprenticeship programs, housing, transportation, and quality of life therefore become part of infrastructure strategy.

A site that cannot attract or develop people may not be a viable site no matter how attractive its power price appears.

27. Water Changes the Map

Cooling technology can reduce water consumption dramatically, but climate and heat-rejection design still matter.

Closed-loop liquid systems can recirculate coolant. Dry and hybrid systems can reduce fresh-water consumption. Reclaimed water can substitute for potable supplies in suitable designs.

Yet every site must solve heat rejection.

A hot, water-constrained region may favor a different cooling architecture from a cool northern region.

That affects both energy consumption and capital cost.

Water therefore remains a geographic variable even when the objective is to minimize water use.

28. Community Acceptance Is a Capacity Constraint

AI infrastructure does not exist outside society.

Large campuses affect land, tax bases, roads, transmission, water, construction activity, noise, visual character, and electricity planning.

Communities increasingly ask whether data centers create enough local value to justify those impacts.

A technically ideal site that cannot secure durable public acceptance may not be an ideal site at all.

The better development model is to make benefits visible: jobs, infrastructure investment, tax revenue, grid upgrades where appropriate, workforce development, heat reuse where feasible, and transparent water and energy planning.

Social license is not a public-relations accessory.
At infrastructure scale, it becomes schedule risk.

MOVEMENT VIII - SATELLITES AND THE EDGE

29. Fiber Cannot Reach Everything

Terrestrial fiber will remain the backbone of high-capacity AI infrastructure.
But not every user, machine, ship, aircraft, rural facility, military unit, mine, farm, or remote sensor sits near fiber.
Satellite communications extend the geographic reach of intelligence.
Low-Earth-orbit systems can reduce latency relative to traditional geostationary satellite links and provide connectivity where terrestrial infrastructure is weak or absent.
For PROMETHEUS, satellites are not a substitute for the high-capacity fiber links connecting major AI campuses.
They are an extension layer.

30. The Edge Moves Compute Toward the User

WP-03 has emphasized moving computation toward energy.
Edge computing sometimes does the opposite.
It moves computation toward the user or machine.
Why? Latency, privacy, resilience, bandwidth, and autonomy.
An autonomous vehicle cannot send every critical decision to a distant gigawatt campus.
An industrial robot may need local inference.
A remote military or scientific system may need to continue operating when communications fail.
The future architecture therefore moves compute in both directions.
Large flexible workloads can migrate toward energy.
Latency-critical intelligence can migrate toward the edge.

31. The Hierarchy of Compute Geography

PROMETHEUS therefore sees a hierarchy:
DEVICE - immediate local intelligence.
EDGE - nearby low-latency compute.
METRO - regional services and aggregation.
HYPERSCALE CAMPUS - massive training and general-purpose AI infrastructure.
CONTINENTAL NETWORK - geographically coordinated campuses responding to energy and demand.
SATELLITE EXTENSION - connectivity beyond terrestrial reach.
FRONTIER ORBITAL COMPUTE - possible future computation in space.
Each layer solves a different geographic problem.

MOVEMENT IX - ORBITAL COMPUTE: FRONTIER, NOT PLAN

32. Why Space Keeps Returning

Space appears repeatedly in PROMETHEUS because it changes the boundaries of communications, energy, and geography.

Solar energy is abundant in orbit. Satellites already move enormous quantities of information. Launch costs have fallen dramatically over the reusable-launch era.

That naturally produces a provocative question:

If energy and communications are available in space, could some computation eventually move there too?

The answer today is: perhaps, for specialized cases - but the barriers remain substantial.

Thermal rejection in vacuum is difficult. Radiation damages electronics. Maintenance and hardware replacement are expensive. Launch mass matters. Communications capacity matters. Orbital debris and regulation matter. Terrestrial accelerators improve rapidly, making space hardware vulnerable to technological obsolescence.

WP-03 therefore treats orbital compute as frontier research, not as an assumption in today's campus plan.

33. Space's Nearer-Term Role

The nearer-term role of space is communications and sensing.

Satellite networks can connect remote users and edge systems to terrestrial AI infrastructure.

Earth-observation satellites can generate data that AI systems process.

Optical inter-satellite links can move information through space-based networks.

Direct-to-device services can extend basic connectivity.

These capabilities expand the geography from which intelligence can receive data and to which it can deliver results.

That is strategically important even if the heavy computation remains on Earth.

MOVEMENT X - THE COMPUTE ORCHESTRATION LAYER

34. The Scheduler Becomes an Infrastructure Operator

Traditional schedulers ask: Which processor is free?

The PROMETHEUS scheduler asks more.

Which processor is free?

Where is it?

What electricity is available there?

What will electricity look like three hours from now?

What is the network condition?

Where is the data?

What latency can the workload tolerate?

What is the battery state?

Is the local grid constrained?

What cooling capacity is available?
What does moving the workload cost?
What happens if the network fails?
This is no longer simple job scheduling.
It is infrastructure orchestration.

35. Energy-Aware Routing

Google's operational demand-response work demonstrates that selected machine-learning loads can already be shifted or reduced to support utilities. In March 2026 the company said it had signed 1 GW of data-center demand response across multiple U.S. utilities. [3]

Microsoft Research's 2026 work similarly models geographically distributed AI data centers as coupled compute-energy systems, jointly optimizing workload placement, local generation, batteries, grid interaction, latency, and carbon constraints. [4]

These are not proof that every future AI system will operate this way.

They are evidence that the boundary between computing operations and energy operations is already dissolving.

36. Predictive Geography

Now add forecasting.

Tomorrow's weather is known imperfectly but usefully.

Solar and wind output can be forecast.

Grid conditions can be forecast.

Electricity prices can be forecast.

Network maintenance is scheduled.

Training deadlines are known.

Campus cooling limits are known.

The system can therefore position workloads before scarcity arrives.

A flexible job does not have to wait for a grid emergency before moving.

It can be placed twelve hours earlier in the region expected to have the best combination of energy, compute, network, and thermal capacity.

Geography becomes predictive.

37. The Compute-Energy Operating System

WP-02 introduced the Energy Operating System.

WP-03 extends it.

The future platform coordinates:

generation; grid imports; storage; cooling; compute priority; workload location; fiber capacity; latency; data placement; hardware availability; and forecast conditions.

The result is a Compute-Energy Operating System.

It does not eliminate physical infrastructure.

It extracts more value from the infrastructure already built.

MOVEMENT XI - WHAT WE SHOULD BUILD

38. Build Energy and Fiber Together

The first design recommendation is simple.

Do not plan a gigawatt AI campus with separate energy and communications teams that meet late in the process.

Map generation, transmission, substations, fiber corridors, internet exchanges, carrier diversity, water, land, workforce, and expansion together.

A site with a powerful grid but weak communications may be inferior to a site with slightly higher energy cost and exceptional network diversity.

A site with cheap fiber but no path to another 500 MW of power may be equally constrained.

39. Build Multiple Campuses as One System

Where scale justifies it, AI developers should think in portfolios of sites rather than isolated campuses.

Different campuses can specialize.

One may be optimized for frontier training.

One for flexible batch work.

One for low-latency inference.

One for disaster recovery.

One for energy abundance.

The software layer then determines how much specialization is useful without creating operational complexity that outweighs the benefit.

40. Pre-Position What Makes Movement Possible

Workload mobility cannot be invented during an emergency.

Models must already exist at alternate sites.

Data must be replicated or accessible.

Security permissions must be established.

Hardware compatibility must be known.

Network routes must be tested.

Service-level rules must identify what can move.

Energy contracts must recognize flexibility.

The system must rehearse the movement.

Resilience is built before it is needed.

41. Do Not Over-Move

Flexibility can become its own form of waste.

Constantly moving workloads across regions can increase network traffic, operational complexity, storage duplication, and failure risk.

The objective is not maximum movement.

It is optimal movement.

Sometimes the correct decision is to leave a workload where it is and start a generator.

Sometimes it is to discharge a battery.

Sometimes it is to delay the job.

Sometimes it is to move it.

The system must be free to choose the least-cost reliable response.

MOVEMENT XII - THE PROMETHEUS GEOGRAPHY DOCTRINE

42. The New Infrastructure Choice

The Intelligence Age creates three things that can move.

ENERGY can move through wires, pipelines, rail, ships, and fuel systems.

INFORMATION can move through fiber, radio, microwave, and satellite networks.

COMPUTATION can move when software, data, hardware, latency, and policy allow it.

Industrial infrastructure has traditionally optimized the first.

Digital infrastructure optimized the second.

AI requires us to optimize all three together.

43. The Geography Hierarchy

When choosing where computation should run, ask:

1. MUST IT STAY HERE? Latency, regulation, data gravity, hardware, or physical control may bind it.
2. CAN IT MOVE IN TIME? Delay may solve the energy problem without geographic movement.
3. CAN IT MOVE LOCALLY? Another building or metro site may have capacity.
4. CAN IT MOVE REGIONALLY? Another campus may have better energy conditions.
5. WHAT DOES MOVEMENT COST? Include network, storage, replication, security, and operational risk.
6. WOULD MOVING ENERGY BE BETTER? Transmission or local generation may be the durable answer.
7. WHAT OPTION PRESERVES THE MOST FUTURE FLEXIBILITY?

That is the PROMETHEUS geography hierarchy.

44. What We Should Not Do

Do not assume the cloud has no geography.

Do not site AI campuses on electricity price alone.

Do not call a workload flexible without understanding its data and latency requirements.

Do not assume fiber is unlimited.

Do not use compute mobility as an excuse to avoid needed transmission.

Do not build a distributed system that depends on a single network path.

Do not ignore land, water, workforce, or community acceptance.

Do not move workloads merely because movement is technically possible.

Do not confuse orbital possibility with terrestrial readiness.

And do not design the Intelligence Age as if energy, communications, and computing were separate industries.

CONCLUSION - COMPUTE FOLLOWS ENERGY

The first industrial revolution concentrated machines near energy.

Later infrastructure learned to move fuels and electricity over enormous distances.

The digital revolution learned to move information.

The Intelligence Age can combine all three.

This does not abolish geography.

It makes geography programmable.

A gigawatt AI campus will still need land, power, cooling, fiber, workers, equipment, security, and community acceptance.

But it may no longer need every computation to remain where it began.

That changes the economics of scarcity.

A constrained substation does not automatically mean the entire workload must stop.

A surplus of wind or solar hundreds of miles away does not automatically require a new transmission line before it can create value for flexible digital demand.

A remote region with abundant firm power is not automatically too remote if communications, latency, workforce, and data architecture make it usable.

Each of these statements has limits.

Those limits are exactly why orchestration matters.

The future AI system will decide continuously whether the best response is to generate, store, transmit, delay, reduce, or move.

WP-01 established that intelligence is becoming infrastructure.

WP-02 established that energy for intelligence must become a layered, programmable portfolio.

WP-03 adds the geographic dimension:

Sometimes the shortest path between energy and intelligence is not a transmission line. It is a fiber-optic cable carrying the computation somewhere else.

That will not always be true.

But the fact that it can be true changes infrastructure planning.

The companies and nations that understand this will not merely build more data centers.

They will build networks of energy, communications, and computation that behave as one machine.

That is the geography of the Intelligence Age.

THE PROMETHEUS GEOGRAPHY RULE

Keep what must remain local.

Delay what can safely wait.

Move what is cheaper to move.

Transmit energy where transmission creates durable value.

Build fiber where information mobility creates durable value.

Place compute where power, communications, cooling, land, labor, and security intersect.

Preserve multiple routes.

Price the movement.

Protect the latency.

Let geography become an option - not a prison.

REFERENCES AND CURRENT EVIDENCE

[1] International Energy Agency (IEA), Energy and AI, 'Understanding the energy-AI nexus' and Executive Summary. Reports strong geographic concentration of data centers; nearly half of U.S. data-center capacity is in five regional clusters. Accessed August 2026. <https://www.iea.org/reports/energy-and-ai/understanding-the-energy-ai-nexus>

[2] Federal Energy Regulatory Commission (FERC), 'FERC Launches Aggressive Targeted Action to Speed Large Load Integration,' June 18, 2026; and Commissioner Rosner remarks. Addresses six regional grid operators, flexible large-load service, speed-to-power, and associated generation. <https://www.ferc.gov/news-events/news/ferc-launches-aggressive-targeted-action-speed-large-load-integration>

[3] Google, 'A new milestone for smart, affordable electricity growth,' March 19, 2026. Reports 1 GW of data-center demand-response capacity integrated into long-term U.S. utility contracts and describes shifting or reducing selected machine-learning workloads. <https://blog.google/innovation-and-ai/infrastructure-and-cloud/global-network/demand-response-data-center-milestone/>

[4] Jue Zhang et al., 'Carbon-Aware Compute-Power Scheduling for AI Data Centers with Microgrid Prosumer Operations,' Microsoft Research / arXiv, May 2026. Models geographically distributed AI data centers with workload routing, local generation, batteries, grid interaction, latency and carbon constraints. <https://www.microsoft.com/en-us/research/publication/carbon-aware-compute-power-scheduling-for-ai-data-centers-with-microgrid-prosumer-operations/>

[5] Reuters, 'Mapping America's mammoth power machine,' August 10, 2026. Provides a current geographic overview of U.S. natural gas, coal, nuclear, hydro, solar, wind, batteries, transmission and emerging data-center demand. <https://www.reuters.com/commentary/reuters-open-interest/mapping-americas-mammoth-power-machine-2026-08-10/>

[6] Google Cloud, 'Supporting power grids with demand response at Google data centers,' October 3, 2023. Describes shifting non-urgent compute tasks to other times and locations and the operational basis for grid demand response. <https://cloud.google.com/blog/products/infrastructure/using-demand-response-to-reduce-data-center-power-consumption>

[7] Federal Energy Regulatory Commission, Docket RM26-4-000, Interconnection of Large Loads to the Interstate Transmission System. Defines and examines large-load interconnection, curtailment, cost allocation and flexible service. Accessed August 2026. <https://www.ferc.gov/rm26-4>

Evidence convention: numbered references support current facts and examples. The 150-MW allocation, 'virtual transmission line,' geography hierarchy, continental-campus architecture, and other design rules are PROMETHEUS analysis unless explicitly sourced.

NEXT IN THE PROMETHEUS RESEARCH SERIES

WHITE PAPER No. 4

THE INTELLIGENT INFRASTRUCTURE STATE

Governance, Security, Capital and the Race to Build the Intelligence Age

WP-04 will move from the physical architecture to the institutions required to build and govern it: permitting, capital, national security, supply chains, cyber resilience, workforce, public-private coordination, community legitimacy, standards, and the strategic competition surrounding intelligent infrastructure.

The question changes again:

If AI, energy, communications and compute are converging into strategic infrastructure, who is responsible for making the whole system work?