

PROMETHEUS RESEARCH PROJECT | WP-07 GOLD EDITION | SEPTEMBER 2026

PROMETHEUS RESEARCH PROJECT

PROMETHEUS

WP-07 | GOLD EDITION

SPACE-BASED AI

When Compute Leaves the Earth

A concise PROMETHEUS briefing on the engineering, energy, semiconductor, communications and industrial systems required to move useful artificial-intelligence compute into orbit.

Prepared by David W Tietgen

September 2026

contact@prometheusresearchproject.org

GOLD EDITION | SEPTEMBER 2026

WHY THIS PAPER EXISTS

Executive Premise

Artificial intelligence appears to be software. At infrastructure scale, it is physical.

Advanced AI requires processors, high-bandwidth memory, networking, electricity, cooling, buildings, fiber, manufacturing and capital. As compute grows, those physical requirements increasingly shape where AI can be built and how quickly it can expand.

Space-based AI proposes a different architecture. Instead of locating every large computing system on Earth, some workloads could eventually move into orbit, where solar power, satellite-generated data, optical communications and increasingly capable launch systems create a different set of engineering possibilities.

The advantages are accompanied by severe constraints. Every kilogram must reach orbit. High-performance processors generate heat that must be rejected in vacuum. Electronics face radiation. Distributed processors require high-bandwidth links. Hardware is difficult to repair and rapidly obsolete. Large constellations add collision, debris, spectrum and governance requirements.

Which parts of the AI-compute stack could become useful in orbit, and what must become true for that transition to occur?

THE ARGUMENT

WP-07 in Five Minutes

- AI is becoming a physical infrastructure industry built on chips, power, cooling, communications and increasingly large data centers.
- Orbital computing substitutes launch, spacecraft power, radiation tolerance and radiative cooling for some terrestrial requirements.
- SpaceX's proposed Starlink concept connects orbital compute with launch, satellite manufacturing, communications and optical networking.
- Google Project Suncatcher provides an independent experimental path for testing AI accelerators, radiation, thermal performance and optical links in space.
- Space-native inference and processing may become practical before tightly coupled frontier-model training.
- The likely near-term architecture is hybrid: terrestrial compute combined with increasingly capable processing in orbit.

PROMETHEUS conclusion: the important question is not whether data centers simply move from Earth into space. It is whether orbit becomes another useful location inside the global computing infrastructure.

ONE-PAGE MAP

The Orbital AI Architecture

Layer	What it does	Representative system
Transportation	Moves hardware to orbit	Reusable launch systems
Power	Generates and stores electricity	Solar arrays / batteries
Compute	Processes AI workloads	Accelerators / orbital compute
Thermal	Rejects waste heat	Radiators / heat transport
Communications	Moves data between nodes	Optical inter-satellite links
Autonomy	Operates distributed infrastructure	Onboard control / AI agents
Ground	Connects orbit to terrestrial systems	Gateways / fiber / data centers
Industrial base	Builds and replaces hardware	Chips / optics / solar / spacecraft

Orbital AI becomes infrastructure only when all eight layers work together.

No individual breakthrough - cheap launch, abundant solar power, faster chips or optical communications - is sufficient by itself.

A DELIBERATE HANDOFF

From WP-06 to WP-07

WP-06 examined the broader SpaceX ecosystem: transportation, manufacturing, communications, endpoints, intelligence, government and the industrial base. It identified orbital AI compute as an emerging layer within a larger vertically integrated architecture.

WP-07 isolates that compute layer and asks whether the underlying engineering can work.

WP-06: What ecosystem is taking shape?

WP-07: Can AI compute become orbital infrastructure?

COMPUTE LAYER

1 | AI: The Infrastructure Problem

Artificial intelligence is increasingly constrained by the physical systems beneath it. The basic chain is direct:

CHIPS → MEMORY → NETWORK → POWER → COOLING → COMPUTE

Increasing capability at one layer places pressure on the others. More processors require more electricity. More electricity consumed by processors produces more heat. Larger clusters require higher-bandwidth interconnects. Higher power density increases cooling requirements.

The International Energy Agency reports that global data-center electricity consumption reached about 485 TWh in 2025 and projects roughly 950 TWh by 2030. AI-focused data centers are growing faster than the sector as a whole. The point is not that AI will dominate global electricity use; it is that concentrated loads can become locally difficult to connect and supply.

This creates the first reason orbital computing is being examined: AI expansion is no longer determined solely by better algorithms. It is increasingly determined by physical infrastructure.

COMPUTE LAYER

Why Move Any Compute Off Earth?

The case is strongest when the workload already has a reason to be in space. Earth-observation satellites collect imagery in orbit. Communications satellites generate network information in orbit. Scientific instruments produce data in orbit. Autonomous spacecraft must make decisions in orbit.

SENSE → COMPUTE → SELECT → TRANSMIT

can be more efficient than

SENSE → TRANSMIT EVERYTHING → COMPUTE ON EARTH

The first orbital AI advantage may therefore be proximity to data, not proximity to sunlight.

POWER LAYER

2 | Solar Energy Above Earth

AI requires continuous electrical power. Space offers direct access to solar energy without atmospheric attenuation or terrestrial land requirements, but orbital solar power is not unlimited or continuous under every orbital configuration.

SOLAR ARRAY → POWER CONDITIONING → STORAGE → COMPUTE → THERMAL REJECTION

Orbit, eclipse periods, battery capacity, solar-cell degradation, orientation and spacecraft geometry all affect usable power. As compute increases, solar structures become larger. Energy storage grows. Power electronics expand. Thermal requirements rise.

A conventional satellite carries electronics to perform a mission. A large orbital-compute platform could instead become an **energy system built around processors**.

POWER LAYER

Power Density: What Must Become True

Requirement	Why it matters
High-efficiency solar	More electrical output for each kilogram launched
Lightweight structures	Large arrays cannot overwhelm payload economics
Reliable storage	Compute must survive eclipse and power variation
Efficient conversion	Conversion losses become additional heat
Long service life	Degradation changes lifetime economics
Scalable deployment	Large power structures must deploy reliably

The central issue is not whether sunlight exists in space. It is whether enough reliable electrical power can be produced **per kilogram of complete operating system**.

THERMAL LAYER

3 | Heat: The Constraint That Follows Power

Every watt consumed by computing eventually becomes heat. Terrestrial data centers can move heat through air, water and mechanical cooling systems. Spacecraft operate in vacuum, where heat cannot simply be transferred to surrounding air.

Waste heat must ultimately be rejected primarily through radiation.

MORE COMPUTE → MORE POWER → MORE HEAT → MORE RADIATOR AREA → MORE MASS → MORE LAUNCH

This may be one of the defining constraints of orbital AI. Large solar arrays can collect additional energy. That does not solve the requirement to dispose of the resulting waste heat.

Google's Project Suncatcher explicitly identifies cooling in vacuum as one of the engineering challenges its prototype program is intended to test.

THERMAL LAYER

The Thermal Constraint Map

Constraint	Question
Heat density	How much processor heat must be removed?
Radiator area	How much surface is required to reject it?
Temperature	At what temperature can the system radiate efficiently?
Mass	How heavy are radiators and heat transport?
Orientation	Can radiators avoid unwanted solar heating?
Reliability	Can the thermal system operate for years without service?

In orbit, cooling is part of the computer.

SEMICONDUCTOR LAYER

4 | AI Chips in the Radiation Environment

Modern AI accelerators are optimized primarily for terrestrial data centers. Orbit adds radiation, energetic particles, thermal cycling and other reliability challenges.

Potential responses include shielding, radiation-tolerant components, redundancy, error correction, fault detection and shorter replacement cycles. Each solution affects economics: shielding adds mass, redundancy adds hardware, specialized components can reduce performance or increase cost, and shorter service lives increase replacement demand.

Should orbital AI hardware be built to survive longer - or built cheaply enough to replace frequently?

SEMICONDUCTOR LAYER

Google Project Suncatcher: The Independent Test

Google describes Project Suncatcher as a research effort to explore scalable machine-learning compute in space. Its 2026 prototype work is intended to test AI hardware against radiation, launch vibration and cooling challenges in vacuum.

- Can advanced AI accelerators tolerate the radiation environment?
- Can high-performance computing hardware survive launch?
- Can heat be removed effectively in vacuum?
- Can optical links connect distributed compute nodes?
- Can useful AI workloads operate reliably in orbit?

The value of the program is not that it proves orbital AI economics. It provides another path for determining which assumptions survive contact with space.

NETWORK LAYER

5 | The Orbital Backplane

A collection of satellites carrying processors is not automatically a data center. Modern AI clusters depend on enormous communication capacity between processors, memory and storage.

In terrestrial facilities, those connections are built through short, dense electrical and optical networks. Distributed orbital computing requires another architecture.

DATA-CENTER FABRIC → ORBITAL LASER MESH

High-bandwidth optical inter-satellite links could allow compute nodes to exchange data through laser connections. But distance, pointing accuracy, changing geometry, synchronization and bandwidth all matter.

Some workloads tolerate distributed operation better than others. Inference and local processing can often operate with less tightly coupled communication. Frontier-model training may require enormous processor-to-processor bandwidth and synchronization.

A plausible progression is:

ONBOARD PROCESSING → INFERENCE → DISTRIBUTED SPACE COMPUTE → LARGE-SCALE TRAINING?

The question mark matters. Orbital compute does not need to reproduce every terrestrial workload to become economically important.

WORKLOAD PLACEMENT

What Should Compute in Space First?

- Earth observation: analyze imagery before transmitting results.
- Communications: optimize routing and network resources locally.
- Space-domain awareness: identify objects, anomalies and changing conditions.
- Autonomous spacecraft: make operational decisions without constant ground intervention.
- Scientific missions: process high-volume instrument data before downlink.

TRANSPORTATION & INDUSTRIAL LAYER

6 | Building the Orbital Computer

Every orbital computing system begins in a terrestrial factory. Processors, memory, solar arrays, radiators, optical terminals, batteries, structures and propulsion must be manufactured and launched.

The relevant economic metric is not simply launch cost per kilogram. It is closer to:

COST PER USEFUL UNIT OF COMPUTE DELIVERED OVER OPERATING LIFE

Element	Economic effect
Hardware	Processors, memory and spacecraft
Launch	Cost and reliability of reaching operational orbit
Power	Solar arrays, storage and conversion
Thermal	Radiators and heat transport
Network	Optical links and ground connectivity
Reliability	Redundancy and fault tolerance
Lifetime	Years of productive operation
Refresh	Rate of processor obsolescence
Replacement	Manufacturing and relaunch

Reusable heavy launch matters because it can change the design trade space. More radiator mass may become acceptable. Shielding may become less punitive. Replacement cycles may shorten. Hardware may become more standardized.

Transportation makes orbital AI possible. Productive operating life makes it useful.

AUTONOMY LAYER

7 | The Computer Must Operate Itself

Large orbital-compute networks cannot depend on continuous human micromanagement. They would need to manage spacecraft health, computing workloads, power allocation, thermal conditions, optical routing, orbital position, collision avoidance, degraded hardware, software faults and replacement scheduling.

Artificial intelligence therefore appears twice in the architecture. It is the **workload**. It may also become part of the **control system operating the infrastructure**.

Space-based AI may require AI-managed space infrastructure.

WHAT COULD SLOW THE SYSTEM

The Orbital AI Constraints

Constraint	Why it matters
Launch	Economics require reliable, repeatable transportation
Power	Large compute loads require large generation systems
Heat	Waste heat requires substantial radiator capacity
Radiation	Advanced processors must remain reliable
Bandwidth	Distributed compute requires high-capacity networking
Obsolescence	AI hardware improves rapidly
Debris	Larger constellations increase orbital-management burdens
Capital	Manufacturing and deployment require sustained investment
Regulation	Spectrum, licensing and international rules affect scale

WHAT COULD SLOW THE SYSTEM

The Orbital Commons

Low-Earth orbit is shared infrastructure. Communications satellites, Earth-observation systems, scientific spacecraft, crewed vehicles and future computing platforms all depend on access to that environment.

As spacecraft numbers increase, collision avoidance, debris mitigation, disposal, spectrum coordination and space-domain awareness become increasingly important.

An orbital computing architecture that creates useful AI capacity while degrading the environment required by other spacecraft would impose costs beyond its own system.

Orbital sustainability is not separate from orbital AI economics. It is one of the requirements.

ECONOMIC LAYER

The Orbital Compute Economy

Orbital AI economics cannot be reduced to launch cost or solar availability. A complete comparison must include satellite manufacturing, launch, power, energy storage, processors, memory, thermal systems, optical networking, ground infrastructure, operations, replacement rate, orbital lifetime and deorbiting.

For many workloads, Earth may retain an overwhelming advantage. For some space-native workloads, orbit may eventually win.

The key comparison is functional: **what does it cost to deliver a useful unit of computation, at the required reliability and latency, over the full operating life?**

PROMETHEUS SCENARIOS

8 | Three Futures to Watch

Scenario	What it looks like
EDGE	AI remains primarily terrestrial while satellites perform increasing onboard inference and preprocessing.
HYBRID	Terrestrial data centers and orbital compute divide workloads according to location, power, bandwidth and latency.
INFRASTRUCTURE	Large distributed orbital-compute networks become a meaningful layer of global AI infrastructure.

These are scenarios, not forecasts. Their purpose is to identify what evidence matters.

PROMETHEUS SCENARIOS

The Indicators That Matter

- **Launch:** demonstrated reuse, cost, payload and sustained cadence.
- **Power:** deployed watts per kilogram and long-duration reliability.
- **Thermal:** demonstrated high-power heat rejection in orbit.
- **Processors:** radiation performance of advanced AI accelerators.
- **Optical links:** sustained bandwidth between distributed compute nodes.
- **Starmind:** hardware demonstrations rather than announced architecture.
- **Project Suncatcher:** results from actual orbital testing.
- **Manufacturing:** repeatable production of compute spacecraft.
- **Economics:** useful compute delivered over complete system life.
- **Governance:** debris, spectrum and operating rules keeping pace with deployment.

These indicators matter more than satellite-count announcements because they measure whether the underlying infrastructure actually works.

WHY THIS MATTERS

The PROMETHEUS Assessment

Space-based AI should not be evaluated as a contest between Earth and orbit. The more useful question is **workload placement**.

Terrestrial data centers retain enormous advantages in maintenance, dense networking, hardware replacement and existing infrastructure. Orbit may develop advantages where data originates in space, communications capacity is constrained, autonomous decisions are valuable, global reach matters or terrestrial infrastructure is difficult to provide.

The likely architecture is therefore not substitution. It is integration.

EARTH COMPUTE + ORBITAL COMPUTE + OPTICAL NETWORKS + SATELLITE DATA + AUTONOMOUS SYSTEMS

The competitive question may eventually become not who operates the largest data center, but who can place computation where it creates the most value.

THE ORBITAL COMPUTE CONSTRAINT LOOP

No Single Breakthrough Closes the System

MASS → LAUNCH → POWER → COMPUTE → HEAT → RADIATOR MASS → NETWORK → UTILIZATION → LIFETIME → REPLACEMENT → MASS

Cheap launch without efficient thermal rejection fails. Abundant solar energy without sufficient networking fails.

Radiation-tolerant processors without economical launch fail. High-performance processors without reliable power fail.

Orbital AI becomes practical only when the complete system closes technically and economically.

FOR EXECUTIVES, RESEARCHERS AND POLICYMAKERS

Questions for Leaders

- Which AI workloads actually benefit from orbital processing?
- Is power generation or heat rejection the harder scaling constraint?
- What launch economics make orbital compute competitive?
- How quickly does processor obsolescence undermine long-lived spacecraft?
- Can optical networks provide the bandwidth required by distributed AI?
- Which functions should remain terrestrial?
- How should governments treat strategically important orbital compute?
- What redundancy is required if commercial orbital infrastructure becomes critical?
- How should orbital congestion and debris costs be incorporated into economics?
- What evidence would demonstrate that orbital AI has crossed from experiment to infrastructure?

These questions are more useful than asking whether space-based AI will simply “work.” Different portions of the architecture may mature at different times.

DECISION FRAME

Earth or Orbit?

Earth remains strong where	Orbit becomes interesting where
Dense processor interconnects dominate	Data originates in space
Equipment must be repaired frequently	Autonomy has high value
Existing power is inexpensive	Downlink bandwidth is constrained
Physical access matters	Global communications matter
Launch mass overwhelms orbital advantages	Launch and spacecraft economics improve

FINAL ASSESSMENT

Conclusion | When Compute Leaves the Earth

Artificial intelligence began as software. Scaling it required processors. Processors required data centers. Data centers required electricity, cooling, communications and industrial supply chains.

Now engineers are testing whether part of that infrastructure can extend beyond Earth.

The first useful orbital AI systems may not resemble terrestrial hyperscale data centers. They may process satellite imagery, optimize communications, operate autonomous spacecraft or analyze scientific data before transmission.

From there, the architecture could expand.

A spacecraft becomes a compute node. Optical links connect nodes into a network. Solar arrays supply power. Radiators reject heat. Reusable launch replenishes hardware. Terrestrial and orbital systems exchange workloads.

The boundary between satellite and data center begins to blur.

The evidence does not yet establish that large orbital AI systems will outperform terrestrial alternatives. It establishes something narrower:

THE ENGINEERING COMPETITION HAS BEGUN.

The decisive question is not whether compute *can* leave Earth. It is whether useful computation can be produced in orbit reliably and economically enough to become infrastructure.

That is the space-based AI question.

It is also the next PROMETHEUS question.

HOW TO READ THIS PAPER

Research Discipline

Label	PROMETHEUS meaning
Established	Operational or officially documented capability supported by authoritative evidence
Emerging	Announced, filed, piloted or early-stage capability not yet mature
Probable	Analysis supported by engineering, economics or strategic behavior, but not guaranteed
Scenario	Forward-looking possibility dependent on unresolved engineering, economics, regulation or capital

WP-07 uses the same discipline as WP-06: company statements are identified as company statements; government and institutional records establish authorization or program status; forward-looking analysis is separated from established fact.

SOURCE NOTES

Primary and Authoritative References

International Energy Agency. *Key Questions on Energy and AI*, 2026. Data-center electricity demand, AI-focused data-center growth, power-density trends and infrastructure bottlenecks.

International Energy Agency. *Energy and AI*, 2025, and *Electricity 2026*. Electricity-system context and data-center demand outlook.

Google Research. *Project Suncatcher*, 2025-2026. Program architecture and prototype testing of AI hardware for radiation, vibration and thermal conditions in space.

xAI / SpaceX. *xAI joins SpaceX*, February 2, 2026. Corporate integration announcement.

SpaceX technical and corporate materials. Orbital AI-compute concepts, launch, Starlink and optical-network architecture. Company claims are treated as announced or emerging unless independently demonstrated.

FCC and other U.S. government records. Used where relevant for licensing, constellation filings, spectrum and orbital authorization status.

PROMETHEUS WP-06 Gold Edition. *The SpaceX Ecosystem*, September 2026. Preceding systems-level analysis and the deliberate handoff to WP-07.