




CATEGORY ROLLUP

Space-based compute & power - 2026-W40

Week of September 28, 2026

This week represents a structural pivot in the trajectory of the space industry: the narrative has officially transitioned from the era of "edge inference" on small, low-power satellites to the era of "Orbital Supercompu

AI News Evolved · cognitoriverdelta.ai

CATEGORY ROLLUP

Space-based compute & power - 2026-W40

Week of September 28, 2026 · 11 min read

Space-based compute & power

Launch & the space economy

Space-based compute & power · week 2026-W40: Sep 27 - Oct 03, 2026 · 1 subtopic(s) covered · 2219 words · expanded

Overview

This week represents a structural pivot in the trajectory of the space industry: the narrative has officially transitioned from the era of "edge inference" on small, low-power satellites to the era of "Orbital Supercomputing." While previous satellite paradigms focused on local processing to reduce the bandwidth needed to downlink imagery, the current industry focus has shifted toward the creation of "SI Factories"—massive, orbital data centers designed to produce Super Intelligence (SI) *in situ*. We are witnessing the emergence of a new industrial stack that integrates heavy-lift launch capacity, massive-scale solar energy production, and high-end AI silicon into a single, cohesive ecosystem.

The primary engine for this transition is the recent successful deployment of Starship Flight 14, which achieved orbit and deployed 26 Starlink Version 3 (V3) satellites. This is not merely a connectivity upgrade; the V3 constellation, with its massive throughput capabilities, acts as the high-bandwidth "nervous system" required to support the data-intensive workloads of an orbital supercomputer. When this connectivity is paired with the ambitious 200 GW annual solar production target cited by Elon Musk and Google's "Project Suncatcher"—an initiative testing high-end AI chips in orbit—a clear strategic direction emerges. The industry is no longer attempting to make space "useful" for Earth-bound applications; it is attempting to move the most resource-intensive components of the digital economy *into* space to bypass the escalating terrestrial constraints on energy, water, and land.

However, this pivot has introduced a fundamental technical tension. The industry is currently split between two competing implementation philosophies. One camp argues that we can utilize existing, high-end Commercial Off-The-Shelf (COTS) hardware, such as NVIDIA or Google TPUs, by relying on advanced software-level hardening to mitigate environmental risks. The opposing camp, led by innovators like Terafab, argues that the unique physics of space—specifically the relentless bombardment of radiation and the extreme difficulty of thermal management in a vacuum—require a fundamental departure from terrestrial design. They contend that the industry needs an entirely new class of specialized, "orbital-native" silicon designed from the ground up to operate under these specific physical constraints.

Space-based AI compute and power (orbital datacenters)

The technical landscape of orbital computing moved from conceptual frameworks to hardware-in-orbit validation this week, marking a critical step in the transition from general-purpose satellite hardware to dedicated, high-power compute architectures.

The Hardware Validation Phase: COTS vs. Custom Silicon The announcement of Google's "Project Suncatcher" serves as a landmark experiment in the "COTS-to-Orbit" strategy. By utilizing SpaceX to send high-end AI chips into orbit, Google is seeking empirical data on whether the silicon required for modern LLMs and multi-modal models can survive the transition from controlled terrestrial environments to the vacuum and radiation of space. This testing is designed to resolve the "Radiation vs. COTS" dilemma. Current high-performance AI chips are often built on advanced 3nm–5nm nodes, which are highly susceptible to bit-flips and Single Event Effects (SEEs) caused by cosmic radiation. Conversely, traditional radiation-hardened chips, while resilient, operate on much older, slower architectures that lack the FLOP/watt necessary for modern AI. Google's results will likely dictate whether the industry's future relies on "software-level hardening"—using redundancy and error correction to manage COTS hardware—or if it must pivot toward the specialized hardware approach championed by Terafab.

Terafab's strategy highlights why this distinction matters. Farzad Mesbahi reports that Terafab is developing two distinct chip families: one optimized for terrestrial applications like robotics and vehicles (including Tesla's Optimus), and a specialized line engineered specifically for orbital compute. This dual-track approach acknowledges that space is not merely a different location, but a different physical environment. An orbital chip must be designed around the Stefan-Boltzmann law to facilitate radiative cooling, as the lack of convective cooling (air or

water) in a vacuum makes heat dissipation a primary design constraint rather than a secondary one.

The Infrastructure and Power Scale: Moving to the Gigawatt Era The scale of proposed orbital infrastructure is undergoing a massive expansion, moving from the kilowatt scale of traditional satellites to the gigawatt scale of industrial-grade data centers. A central component of this discussion is the "AI1" satellite rack. As highlighted by Phil Bicil, a single AI1 unit has the potential to provide 200–250 kW of compute. To visualize this scale, a single AI1 unit would possess a power profile comparable to a small terrestrial industrial facility, but contained within a single orbital chassis.

Meeting this unprecedented power demand requires a radical reimagining of the solar economy. Elon Musk has indicated that SpaceX and Tesla are targeting 200 GW of annual solar production in space. Crucially, this strategy is not about solving the historical hurdle of Space-Based Solar Power (SBSP)—the complex task of beaming energy back to Earth via microwaves or lasers—but about creating an "energy-compute loop" in orbit. By consuming the massive amounts of solar energy generated *in situ*, orbital data centers can bypass the atmospheric attenuation and conversion losses associated with power-beaming, instead converting raw solar irradiance directly into digital intelligence.

The Connectivity Backbone: The V3 "Nervous System" The viability of these orbital "SI Factories" is entirely dependent on the data plumbing provided by next-generation satellite constellations. The specifications for Starlink Version 3 (V3) represent a massive leap in throughput: 1 Tbps downlink, 160 Gbps uplink, and 2.4 Tbps optical inter-satellite link capacity. If the AI1 unit represents the "brain" of the orbital datacenter, the V3 constellation serves as the "synapses," enabling the rapid movement of data between orbital compute nodes and terrestrial users.

This connectivity is essential for the "digital twin" architecture proposed by Phil Bicil. In this model, the high-level "mind" and long-term memory of an agent—such as a Tesla Optimus robot—reside in orbit, while the immediate, low-latency "reflexes" are handled by edge hardware on the robot itself. This creates a symbiotic relationship between terrestrial edge computing and orbital supercomputing, facilitated by the high-bandwidth V3 backbone.

Cross-cutting themes

The developments this week confirm that "Space-based compute" and "Space-based power" are no longer distinct sectors; they have converged into a single, integrated techno-economic system.

1. **The Energy-Compute Symbiosis:** The fundamental driver for relocating compute to space is the terrestrial energy bottleneck. As AI workloads grow, the demand for grid capacity and water for cooling is hitting unsustainable levels on Earth. Space offers an unlimited, high-density solar supply and the ability to dissipate heat into the void via radiation. Consequently, the two industries are locked in a symbiotic relationship: the success of orbital datacenters is predicated on the successful scaling of space-based solar, and the economic justification for massive solar arrays in space is the insatiable demand for AI compute.
2. **Launch Capacity as the Ultimate Bottleneck:** Both the power and compute subtopics are tethered to the progress of the Starship program. The sheer mass required to deploy megawatt-scale solar arrays, combined with the heavy, radiator-laden chassis required for orbital datacenters, makes Starship the most critical piece of infrastructure in this category. The orbital compute market is essentially a downstream beneficiary of the heavy-lift launch market.
3. **The Shift from "Data Downlink" to "Intelligence Uplink":** There is a profound paradigm shift occurring in how we value satellite data. Historically, the value proposition of satellites was "Earth Observation"—sensing data and sending it *down* to Earth. The new paradigm is the production of intelligence. Space is becoming a factory that generates processed, high-level intelligence (Super Intelligence) which is then beamed down to Earth, shifting the value from raw data to refined cognitive output.

Where sources agree

- **SpaceX is the indispensable gatekeeper:** There is total consensus that SpaceX's Starship program and the Starlink V3 infrastructure are the foundational "enablers" that make an orbital datacenter market even theoretically possible.
- **Solar efficiency is a unique advantage:** All sources agree that space-based solar energy provides a superior and more consistent power profile than terrestrial solar due to the lack of atmospheric interference and the constant availability of sunlight.

- **AI is the primary economic driver:** The push into orbit is not being driven by traditional telecommunications or incremental sensor improvements, but by the exponential growth of AI and the specific pursuit of Super Intelligence (SI).
- **The hardware hurdle is a fundamental reality:** There is no dispute that operating high-performance silicon in a high-radiation, vacuum environment is a massive engineering challenge that will require either significant advances in specialized hardware or highly sophisticated software mitigation.

Where sources disagree or differ

- **The "COTS vs. Custom" Hardware Debate:** A significant divide exists regarding the implementation of silicon. Steven Mark Ryan suggests that existing high-end hardware (NVIDIA, Google TPUs) may be sufficient for initial deployments. In contrast, Farzad Mesbahi (Terafab) argues that the physical realities of space necessitate a proprietary, orbital-native class of chips designed specifically for the thermal and radiative environment.
- **Economic Valuation and TAM:** Estimates for the total addressable market (TAM) diverge wildly. ARK Invest presents a hyper-bullish \$28.5 trillion figure. Conversely, Steven Mark Ryan offers a more granular, hypothetical model where 1 GW of AI compute generates \$50 billion in annual recurring revenue, which presents a different scale of economic modeling.
- **Timeline of Consolidation:** There is disagreement on when the industry will see a "Great Consolidation" of assets. Peter H. Diamandis, citing Cathie Wood, suggests a potential Tesla-SpaceX merger could occur as early as this year to consolidate energy and compute. Steven Mark Ryan suggests a more conservative timeline, pointing toward late 2026 or early 2027.
- **The Scale of Orbital Dominance:** Phil Bici presents a "total dominance" model, predicting that 99% of all compute will eventually reside in space. Other analysts view orbital compute more as a specialized high-performance niche rather than a replacement for the entirety of terrestrial computing.

Numbers and claims to verify

- **\$28.5 trillion TAM:** This figure from ARK Invest is an extreme outlier that requires rigorous validation of the assumptions regarding AI's long-term impact on global GDP.
- **200–250 kW per AI1 unit:** This power target for the AI1 satellite rack must be verified against the realistic mass and surface-area limits for solar arrays and radiators on a single satellite bus.
- **Starlink V3 technical specs:** The specific throughput figures (1 Tbps downlink / 2.4 Tbps optical link) should be cross-referenced with official SpaceX technical documentation to ensure accuracy.
- **Cost-parity projections:** Jo Bhakdi's specific projections—\$68 billion per GW in the medium term, falling to \$43–\$50 billion per GW by 2035—require tracking against actual Starship launch costs and orbital deployment efficiencies.

Investment and strategic implications

- **Vertical Integration is the primary competitive moat:** The most potent strategic position is to control the entire "intelligence-production" stack. This means integrating launch (SpaceX), energy (Tesla), and compute (AI chipmakers). Companies that can manage the transition from a solar photon to a digital bit through a single integrated chain will capture the most value.
- **A fundamental shift in CapEx modeling:** For the "Big Tech" players, capital expenditure is moving away from terrestrial real estate and local power-grid interconnections. Instead, CapEx will increasingly be allocated to orbital launch manifests and the deployment of massive satellite constellations.
- **The rise of "Intelligence Factories" as a new asset class:** Orbital datacenters represent a different type of industrial asset than terrestrial ones. Unlike traditional datacenters, which are largely depreciating real estate, an orbital datacenter is a high-velocity, rapidly replenished constellation of machines designed for the continuous production of intelligence.
- **Regulatory freedom as a strategic advantage:** As Jo Bhakdi notes, the ability to operate in space provides a "regulatory moat." Orbital providers can bypass the terrestrial constraints that slow down ground-based expansion, such as water rights, land use permits, and the lengthy interconnection queues of local power grids.

What to watch next week

- **Starship Flight 15 developments:** Any updates regarding the reliability, turnaround time, or payload capacity of Starship will serve as a direct proxy for the feasibility timelines of the orbital datacenter market.
- **Project Suncatcher progress:** Watch for any data or official commentary from Google regarding the performance of their orbital AI chip tests, specifically regarding bit-flip rates and thermal stability.
- **Terafab prototype announcements:** Look for any updates on Terafab's specialized orbital-specific chip architecture and whether they can demonstrate a viable solution to the radiative cooling challenge.
- **Consolidation signals:** Monitor for any regulatory filings or strategic moves between SpaceX and Tesla that might substantiate the merger/consolidation theories discussed by analysts.

Appendix: Individual perspectives

- **Cathie Wood (ARK Invest):** Focuses on the \$28.5 trillion TAM and the transformative economic potential of a Tesla-SpaceX merger to consolidate the AI, energy, and space-based compute sectors.
- **Elon Musk:** Drives the vision of a 200 GW orbital solar-powered infrastructure to circumvent the energy and battery limitations of Earth.
- **Jo Bhakdi:** Provides a roadmap for cost-parity with Earth-based facilities by 2029, citing specific declines in the cost per gigawatt of compute due to regulatory freedom and increasing orbital efficiency.
- **Phil Bicil:** Proposes a "total dominance" vision where 99% of global compute moves to space, supporting a "digital twin" architecture for terrestrial robotics.
- **Farzad Mesbahi:** Emphasizes the technical necessity of specialized, orbital-native silicon (via Terafab) to manage the unique radiation and thermal constraints of the space environment.
- **Steven Mark Ryan:** Contributes to the economic modeling of the industry (revenue per GW) and suggests that the initial phase of orbital computing may rely on existing high-end hardware like NVIDIA and Google TPUs.

Sources

- 26 V3 Sats = 1.5 Years of Starlink Bandwidth? #SpaceX — Brighter with Herbert, Sep 29
- Cathie Wood on Tesla-SpaceX Merger, \$1M Bitcoin, More AIs Than Humans | EP #296 | Moonshot — Peter H. Diamandis, Sep 29
- Elon & Jensen (Full Talk From Today) — Steven Mark Ryan, Sep 29
- Elon Musk Said The Unthinkable in Front of Trump — Farzad Mesbahi, Sep 29
- Google Sends AI Chips Into Orbit on SpaceX Rocket for Space Data Center Test — HokaNews (via Google News), Oct 03
- Google Sends AI Chips to Space, Testing Orbital Data Centers — Aktualita.co (via Google News), Oct 03
- Google launches Project Suncatcher, a step towards AI data centers in space — Hacker News, Oct 03
- Robotics Have Advanced Beyond Our Expectations | Big Ideas 2026 Mid-Year Review — ARK Invest (Cathie Wood), Sep 29
- Sophia Space Aligns With Redwire to Accelerate Orbital Data Center Infrastructure — TipRanks (via Google News), Sep 30
- Sophia Space Inc. And Redwire Corporation Team Up To Advance Orbital Data Center Infrastru — marketscreener.com (via Google News), Sep 30
- Sophia Space and Redwire Team Up to Advance Orbital Data Center Infrastructure — PR Newswire (via Google News), Sep 30
- Space Computing: NVIDIA Chips on SpaceX Rockets! #shorts — Steven Mark Ryan, Oct 03
- SpaceX Tesla Futures Tied to Starlink and Starlink — Randy Kirk, Oct 04
- SpaceX To Launch Prototype Laser Grid For AI Data Centers In Space — HotHardware (via Google News), Sep 29
- Terafab is planning on making two broad families of chips — Farzad Mesbahi, Oct 03
- They Said It Would Never Work - SpaceX just turned on the 10 trillion dollars printing mac — BestInTESLA, Sep 29
- Wood FIRM On Tesla & SpaceX Coming Together — Steven Mark Ryan, Sep 30

Related reports

CATEGORY ROLLUP

COMBINED REPORT

Launch & the space economy - 2026-W40

This week marks the definitive transition of the SpaceX narrative from one of "experimental testing" to one of "operational utility." The successful orbital debut of Starship during Flight 14 was not merely a milestone i

Week of September 28, 2026 · Launch & the space economy · Space-based compute & power

Transition to Physical and Orbital AI

How AI is crossing from software into the physical world and into orbit - robots, factories, energy and space-based compute, read together.

September 25, 2026 · Space-based compute & power · Datacenters & cloud

DAILY AI BRIEFING

Multi-Trillion-Dollar Capital Expenditure Expansion

The global technology and investment landscape is undergoing a structural pivot toward physical AI, autonomous commercial transport, and massive orbital compute infrastructure, driven by multi-trillion-dollar capital expenditure plans across Elon Musk's corporate ecosystem. Industry analyses from le

September 26, 2026 · Space-based compute & power · Autonomous vehicles

Everything here is free, and reader-supported.

Join free on Patreon to get the weekly digest by email, or become a member to get each morning's briefing by email before it reaches the site, request custom analysis, and get early access to new research tools. What members get

♥ [Join or support on Patreon](#)

Informational analysis synthesized by AI from sourced, dated material, curated by a human. Treat specific claims as unverified until checked. Not financial advice.

[About](#) · [How this is made](#) · [Corrections](#)