KPMG

LATENCY TO LEADERSHIP

ORBITAL COMPUTE, SPACE-AS-A-SERVICE , AND THE FUTURE OF AI INFRASTRUCTURE

The New Frontier

An Introduction to the Orbital Compute Paradigm

We stand at a defining crossroads: will space remain merely a footnote in history, or will it ignite the next wave of global digital transformation by emerging as the backbone of our modern infrastructure? The scope has outgrown rockets and satellite launches. Now, space is central to building resilient digital networks, securing data sovereignty, and unlocking unprecedented opportunities for enterprise growth across the planet. In this new era, nations and organizations compete for strategic dominance not just through technology, but through bold, visionary investments that shape the future. This decade demands we choose. Will we relegate space to the margins, or seize its potential to power the world's next great leap?

The space economy has entered a new phase.

What once centered on launch capability and satellite deployment now extends into digital infrastructure, data sovereignty, and enterprise value creation at planetary scale. Space has become a domain where strategic advantage is designed, engineered, and defended through long-term investment decisions.

AI is reshaping infrastructure priorities.

As artificial intelligence accelerates demand for compute, energy availability, resilience, latency, and governance rise to the forefront of leadership agendas. These forces are redefining how governments and enterprises determine where intelligence belongs and how it operates.

Orbital compute changes the model.

Placing compute, storage, and analytics closer to space-based data introduces an environment where physics, economics, and policy converge. Leadership in this domain emerges through early recognition of constraints and deliberate design choices that convert them into advantage.

Orbit becomes strategic infrastructure.

Winning organizations align orbital platforms to long-term growth, national priorities, and enterprise resilience. They apply the same rigor to space infrastructure that they apply to capital markets, cybersecurity, and critical supply chains.

This marks the first chapter.

Across two parts, this series examines how artificial intelligence, infrastructure, and physics are reshaping competitive advantage beyond Earth, and how leaders respond with intent and precision.

The future belongs to decisive leaders.

Brian Miske
Americas Space Lead
KPMG LLP

INTRODUCTION

The next frontier of competitive advantage isn't just in orbit. It's in mastering the physics that will define the orbital compute market. As the Space industry races to deploy infrastructure powered by artificial intelligence (AI) beyond Earth, it's time to move past the allure of "free cooling" and low latency and confront the real engineering and economic paradox at play. We're not simply launching data centers; we're architecting complex thermal ecosystems where the absence of convection means every watt of heat must be radiated away, shifting the economic equation from operational expense to bold, up-front engineering investment.

The true differentiator will be proprietary thermal intellectual property (IP), not just faster chips or cheaper launches. For leaders, this means reframing orbital compute as a thermodynamic innovation challenge, where value creation hinges on the ability to design, patent, and optimize advanced heat rejection systems. If you don't own the heat dissipation, then you don't own the uptime. The organizations that recognize this shift, and invest accordingly, will set the pace for the space economy, transforming thermal management from a hidden constraint into a strategic asset.

Authored by

Brian Miske, Americas Space Lead, KPMG LLP

Contributors

Jonathon Gill, Global Head of Industrial Markets & Aerospace, KPMG UK
Joe Cassidy, EMA Head of Technology, Media & Telecoms, KPMG UK
Chad Seiler, Head of TMT, KPMG LLP
Chris Brown, Head of Strategy, KPMG Ireland
Neil Rae, Head of Space, KPMG UK
Angie Gildea, Partner, KPMG LLP
David Pessah, Executive Director, Ignition, KPMG LLP
Jozy Assa-Jacobson, Director, Ignition, KPMG LLP

Cover Art produced in Adobe Firefly

"An Earth-orbiting satellite silhouetted against the planet's horizon at sunrise, with illuminated city networks visible below and the Milky Way spanning deep space above…"

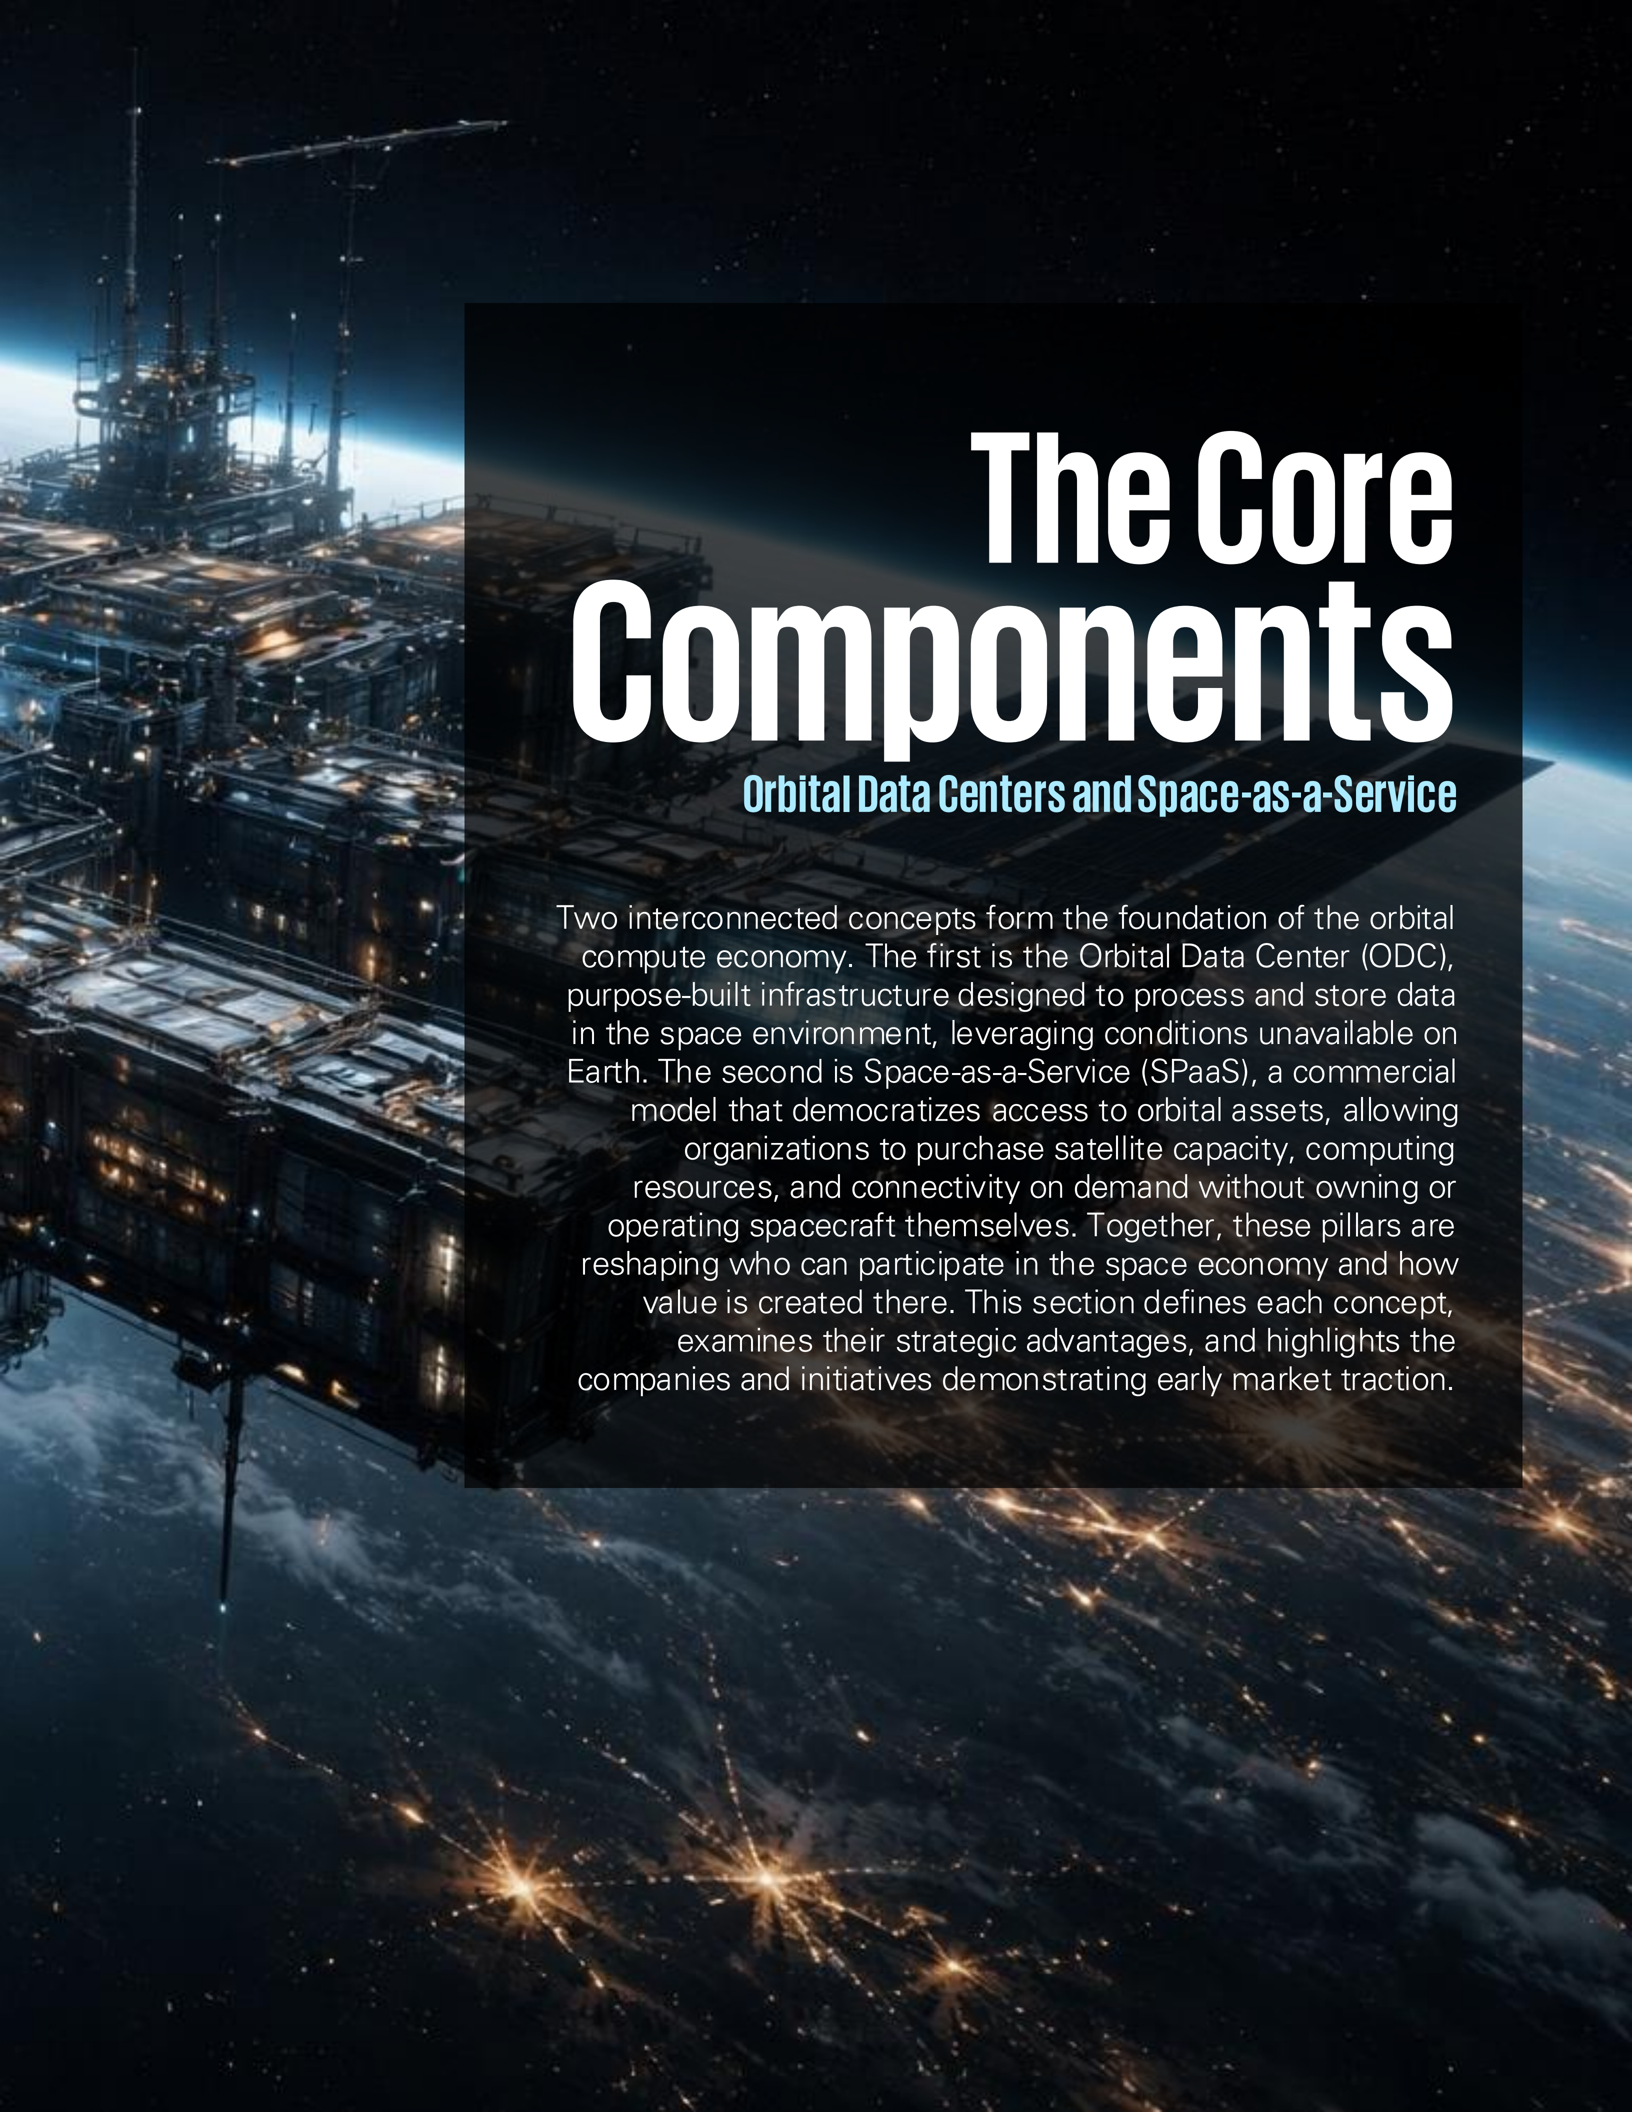

The Core Components

Orbital Data Centers and Space-as-a-Service

Two interconnected concepts form the foundation of the orbital compute economy. The first is the Orbital Data Center (ODC), purpose-built infrastructure designed to process and store data in the space environment, leveraging conditions unavailable on Earth. The second is Space-as-a-Service (SPaaS), a commercial model that democratizes access to orbital assets, allowing organizations to purchase satellite capacity, computing resources, and connectivity on demand without owning or operating spacecraft themselves. Together, these pillars are reshaping who can participate in the space economy and how value is created there. This section defines each concept, examines their strategic advantages, and highlights the companies and initiatives demonstrating early market traction.

SPaaS, LPUs versus CPUs in a new business model set to reshape global infrastructure

The cloud is leaving Earth, and the companies that prepare for orbital computing today will own the competitive edge tomorrow. A new paradigm is emerging, moving compute from “cloud on Earth” to “cloud in orbit.” Orbital data centers (ODCs) collocate compute, storage, and analytics with the rapidly expanding constellation of space-based sensors, enabling low-latency inference, bandwidth-efficient “data thinning,” and resilient operations decoupled from terrestrial infrastructure. Early exemplars, HPE’s Spaceborne Computer missions, OrbitsEdge micro-datacenter concepts, and Axiom Space’s ODC nodes, demonstrate increasing technical feasibility and market intent.[1,2,3,4]

Concurrently, **Space-as-a-Service (SPaaS)** is extending the cloud model into orbit, offering hosted payloads, on-orbit compute, storage, and connectivity on subscription or utility-like terms. Loft Orbital, Azure Orbital Ground Station, and allied “-as-a-Service” frameworks illustrate a maturing ecosystem that lowers capital intensity and accelerates speed-to-orbit.[5, 6, 7, 8]

Orbital data centers

What if the cloud moved into orbit? Imagine a future where data centers aren't anchored to Earth, they're orbiting above it. **ODCs** and **SPaaS** represent a radical shift in architecture, business models, and governance. These platforms promise differentiated energy solutions, improved latency performance, and entirely new forms of value creation.

ODCs are compute infrastructures deployed in **low Earth orbit (LEO), medium Earth orbit, geostationary equatorial orbit, or even cislunar space**, designed to manage data processing, storage, and AI inference close to the source. That source includes **Earth observation systems, intersatellite telemetry, and space-based sensors.** By processing data in orbit, ODCs dramatically reduce round-trip latency and minimize downlink demands, unlocking efficiencies that terrestrial cloud models simply cannot match.[9, 10]

Evidence of traction

Spaceborne Computer (HPE–National Aeronautics and Space Administration [NASA], International Space Station [ISS]): Demonstrated the viability of commercial-off-the-shelf (COTS) servers in orbit using software-based fault management to mitigate radiation without traditional rad-hard hardware cycles.[1]

OrbitsEdge SatFrame and micro-datacenter: A colocation concept for in-orbit analytics and data thinning, partnered with HPE.[2]

Axiom Space ODCs: Announced initial ODC nodes in LEO for secure, scalable AI/machine learning (ML) and cloud-enabled processing for defense and commercial use; prototype launched to ISS with Red Hat Device Edge.[4, 11]

Industry analysis and start-ups: Cutter Consortium's 2025 overview, Sophia Space, Starcloud/Lumen Orbit, and ongoing studies model business cases and architectures.[12, 13, 14]

Key benefits

Latency reduction

Eliminates intermittent downlink bottlenecks and enables near-real-time exploitation (e.g., disaster response; intelligence, surveillance, and reconnaissance [ISR]).[9, 10]

Operational resilience

Decouples mission-critical analytics from terrestrial outages (cyber, physical), leveraging intersatellite laser links and distributed mesh architectures.[16, 17]

Bandwidth efficiency

On-orbit "data thinning" reduces volumes transmitted to ground stations constrained by spectrum and visibility.[9, 15]

Energy opportunity

Continuous solar availability and radiative heat rejection via space-grade radiators (not "free cooling" in vacuum) offer distinct thermodynamic and sustainability trade-offs; academic and industry work is probing feasibility and constraints.[18, 19]

There's no air in space, so heat can't just float away like it does on Earth. Instead, spacecraft use special tools like radiators and heat pipes to move heat out, because "free cooling" doesn't exist in a vacuum. While space offers a radiative sink, there is no convective cooling in vacuum. Thermal management depends on radiators, conduction paths, and heat pipes, not ambient "free cooling." Successful platforms (e.g., ISS) dissipate heat via shaded radiators pointed to deep space; terrestrial data center "free cooling" analogies do not directly apply.[19]

"Orbit expands the cloud's control plane. When latency, spectrum, and sovereignty dominate the business case, you move inference up, not away."

Chad Seiler, Head of TMT, KPMG LLP

What is SPaaS?

SPaaS applies cloud business models to orbital infrastructure

Compute-as-a-Service in orbit, Hosted-Payload-as-a-Service, and Connectivity-as-a-Service via ground/space integration.[20]

Market examples

Loft orbital (hosted payloads)

Offering standardized buses, plug-and-play integration, and multitenant missions, recent €170 million financing to scale.[5, 6]

SPaaS taxonomy (academic)

University of Luxembourg's framework clarifies "-as-a-Service" offerings for space, mirroring cloud virtualization concepts.[20]

Azure orbital ground station and cloud access

Bridging satellite communications to cloud services and democratizing ground segment access.[8]

Industry overviews

Kratos and VivaTech summarize the ecosystem shift from ownership to subscription, expanding access and speed-to-orbit.[21, 22]

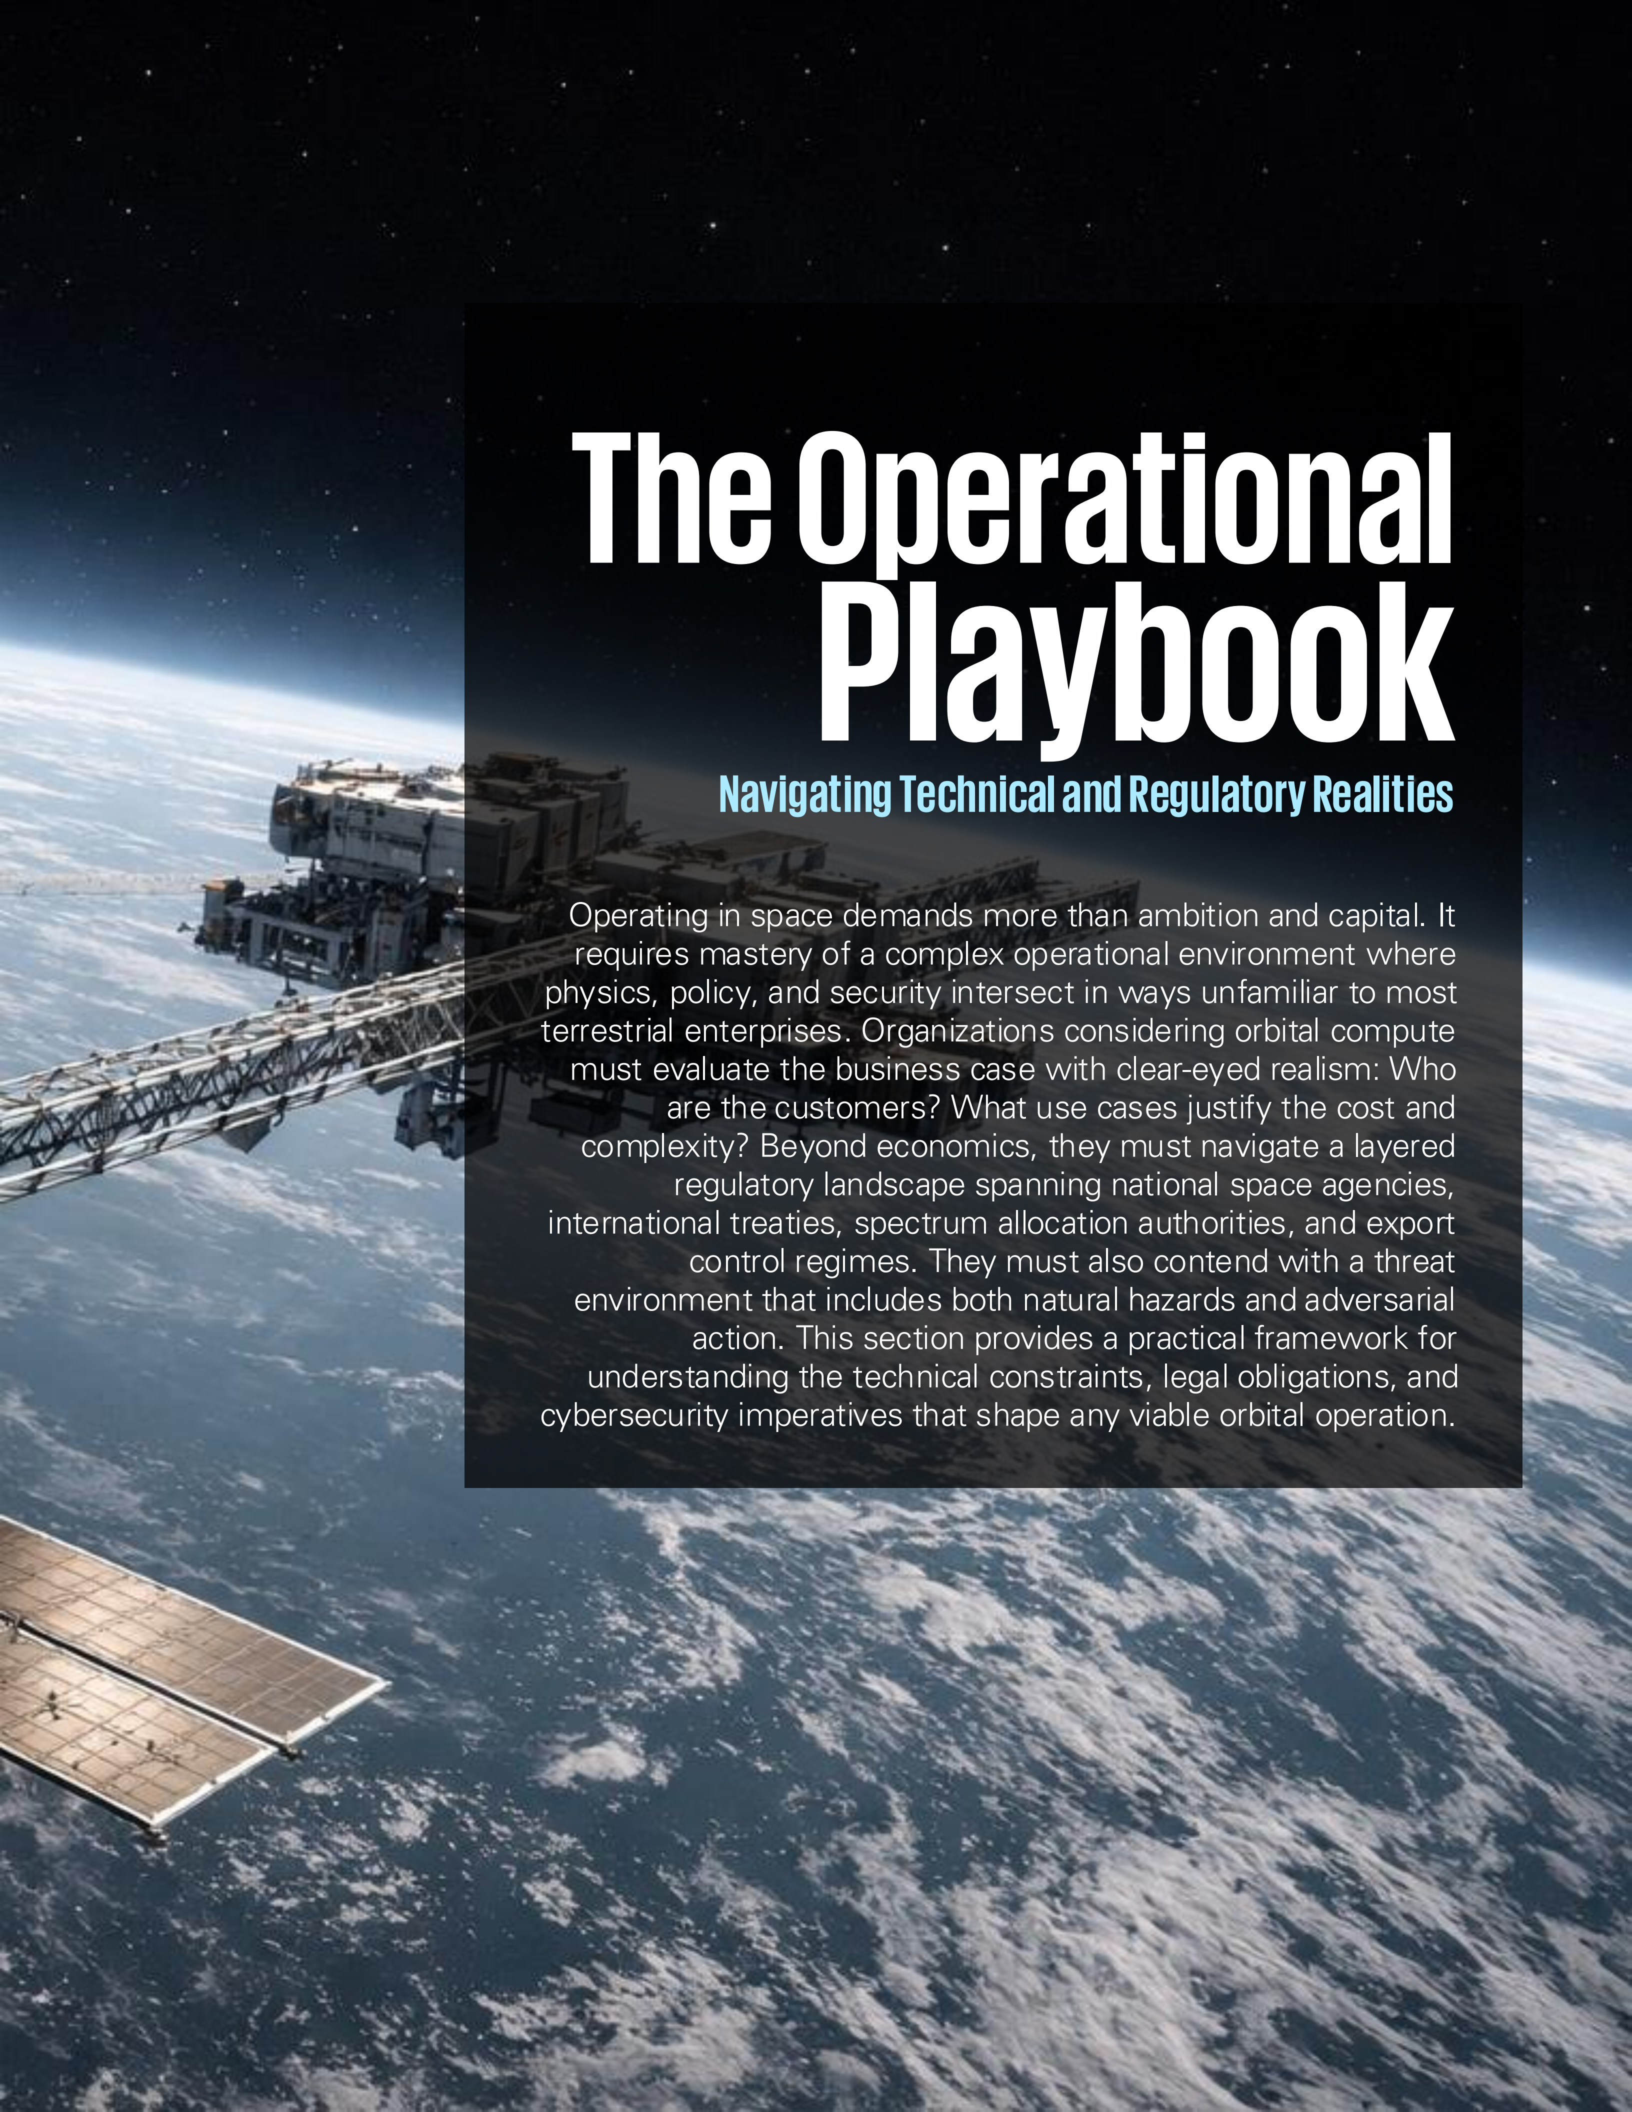

The Operational Playbook

Navigating Technical and Regulatory Realities

Operating in space demands more than ambition and capital. It requires mastery of a complex operational environment where physics, policy, and security intersect in ways unfamiliar to most terrestrial enterprises. Organizations considering orbital compute must evaluate the business case with clear-eyed realism: Who are the customers? What use cases justify the cost and complexity? Beyond economics, they must navigate a layered regulatory landscape spanning national space agencies, international treaties, spectrum allocation authorities, and export control regimes. They must also contend with a threat environment that includes both natural hazards and adversarial action. This section provides a practical framework for understanding the technical constraints, legal obligations, and cybersecurity imperatives that shape any viable orbital operation.

Downstream impact

SPaaS lowers time-to-value, spreading capital across standardized platforms, enabling **software-defined missions** (virtual payloads) and in-orbit analytics that compress iteration cycles for government and commercial customers.[5]

LPUs versus CPUs

Choosing the right brain for orbit under space constraints

Central processing unit (CPU) (general-purpose)

Excellent for sequential and mixed workloads but **not optimized** for massive parallelism or energy-efficient inference under tight SWaP-C (size, weight, power, cost) constraints typical in space platforms.[23] Computers designed to operate in space need to be **small and light and use very little power** due to payload capacity and energy limitations. CPUs aren't the best for that because they use more energy and aren't super-fast at doing many jobs at once.

Language processing unit (LPU) (low-power/light processing unit)

LPUs prioritize **parallelism for inference**, **energy efficiency**, and **fault tolerance**, critical in orbital environments with constrained power budgets and radiation exposure. Inside an LPU, there are special parts that make math faster (matrix engines), quick memory so they don't waste time waiting, and extra tools called accelerators that help with hard jobs. These architectural features include specialized matrix engines, low-latency memory hierarchies, efficient data paths, and integrated accelerators.

Radiation considerations

The space radiation environment exposes electronics to **single-event effects (SEEs), total ionizing dose (TID), and displacement damage**, requiring either rad-hard parts or **COTS devices using screening plus architectural mitigation** (e.g., error-correcting code [ECC], triple modular redundancy [TMR], watchdogs, checkpoint-rollback). NASA/ European Space Agency (ESA) tutorials and NASA Electronic Parts and Packaging guidance codify assurance and mitigation strategies.[25, 26, 27]

Established approaches to feasibility

HPE Spaceborne

Software-based derating and fault management showed COTS viability for space compute.[1]

Microchip RT PolarFire system-on-chip (SoC) (RISC-V)

Radiation-tolerant SoC field-programmable gate arrays (FPGAs) offering lower power and “zero configuration upsets” versus static random-access memory FPGA reliance on scrubbing.[28]

NASA Goddard Space Flight Center review of AI onboard

Advocates COTS accelerators screened for TID/SEE with system-level fault tolerance to unlock feasible AI inference in orbit.[23]

Edge AI trends (industry)

LPUs/neural processing units (NPUs) on microcontroller units and specialized application-specific integrated circuits push high-efficiency inference, principles that translate to orbital constraints.[29]

Why LPUs will likely dominate ODCs

- **Autonomy and real-time analytics** require parallel inference close to sensors.[8]
- **Power/thermal constraints** prioritize high performance per watt and simplified thermal paths; LPUs offer better **giga operations per second per watt** than general CPUs.[28]
- **Radiation tolerance strategies** favor simpler, lower-power fabrics combined with error-resilient software, more tractable in LPUs/SoC-FPGAs than in large, hot CPUs.[23]

"Across discussions with senior leaders in both the public and private sectors, a consistent picture of near-term orbital compute is emerging. The strongest commercial pull today is driven by bandwidth economics: for data-heavy sensing missions, processing in orbit to reduce downlink loads is viewed as the most immediate source of value.

At the same time, operators expect early deployments to be modest in scale – lightweight compute nodes embedded within inter-satellite networks – reflecting practical constraints around power, radiation exposure and thermal stability. These dynamics also reinforce the need for clearly scoped, high-value workloads, alongside careful consideration of governance and interception exposure once traffic leaves terrestrial fibre.

Taken together, these signals point to a measured but meaningful trajectory: orbital compute becomes compelling when it strengthens what existing cloud infrastructure already does well – extending capability, improving resilience and enhancing the overall performance of digital infrastructure."

Chris Brown, Head of Strategy, KPMG Ireland

Key considerations

What is the business case for orbital compute versus terrestrial cloud?

Latency and timeliness: Missions (e.g., wildfire response, maritime monitoring, ISR) benefit from on-orbit processing / exploitation / dissemination (PED) within minutes versus hours/days given ground pass windows and downlink constraints. Empirical studies of LEO downlink architectures and EO delivery latency support the case.[10, 15]

Bandwidth economics: This is where speed is king. Spectrum limits and cloudy optical ground paths create structural bottlenecks; in-orbit analytics reduce volume and ground-segment spend.[15, 30]

Resilience and security: Orbital compute can sustain operations during terrestrial outages or contested environments; cyber events (e.g., KA-SAT) and outages highlight ground-segment fragility.[16]

Energy and sustainability: Continuous solar, radiative heat sinks, and land-neutral siting are attractive; yet engineering realities (thermal radiators, storage for eclipse, launch emissions, debris risk) must be factored. Academic and industry assessments are ongoing.[18]

Key customer segments

Government and defense: Mesh networks aligned with the Space Development Agency (SDA), ISR, threat tracking, and autonomous satellite control, explicit ODC use cases identified by Axiom Space and industry analysts.[4, 12]

Commercial EO and geospatial analytics: Near-real-time insights for agriculture, insurance, logistics, and disaster response when bandwidth is scarce.[10]

Industrial Internet of Things and maritime/aviation: Spaceborne data fusion and edge analytics for global asset monitoring under intermittent connectivity.[7]

Space manufacturing and research and development: Hosted compute for materials, pharma crystallization, and microgravity processes on commercial stations.[11]

What regulatory and cybersecurity frameworks apply (ITAR, FCC, GDPR in space)?

Operating in orbit goes beyond engineering. Compliance and resilience are also crucial. From export controls like [International Traffic in Arms Regulations] (ITAR) and Export Administration Regulations (EAR) to spectrum licensing under the Federal Communications Commission (FCC) and International Telecommunication Union (ITU), and data protection rules similar to the General Data Protection Regulation (GDPR), organizations must navigate complex regulatory frameworks while implementing zero-trust cybersecurity architectures to safeguard orbital data centers against power constraints, radiation, and evolving threats:

- **Export controls (ITAR/EAR):** Recent US rulemakings (October 2024) adjust controls and propose **license exception commercial space activities (CSA)**. affecting payloads, software, and data flows. Compliance planning must precede multinational development.[31, 32, 33]
- **Licensing (FCC/ITU):** Frequency, link budgets, and spectrum coordination remain gating factors for ODC communications. NASA ground systems guidance highlights frequency selection and licensing.[17]
- **Data protection and privacy (GDPR-like):** Jurisdiction and sovereignty complicate personal data handling when sensing and processing occur off-territory; organizations must map legal obligations across registration states and customer domiciles (see Legal section below). (Sources below in Governance.)
- **Cyber resilience:** Cybersecurity and Infrastructure Security Agency (CISA) recommendations (2024) and White House industry workshops (2025) urge **zero-trust architectures across space segments**, secure uplink/downlink, and Earth-independent endpoint detection and response (EDR) for on-orbit systems, directly relevant to ODCs.[34, 35]
- **KPMG perspective:** Mature operational technology (OT)/space-system cybersecurity practices, including supply-chain assurance and information technology/OT convergence, should be integrated from design (not retrofit).[36]

Technical perspectives

What workloads justify orbital deployment? Not every workload belongs in orbit. Only those where speed, proximity, and efficiency outweigh Earth-based processing. AI inference on Earth observation data, space domain awareness, multisensor fusion across constellations, and edge analytics for intersatellite networks are prime candidates, along with select quantum experiments that demand the unique conditions of space:

- **AI inference on EO data** (object detection, segmentation, anomaly detection), **space domain awareness, rapid multisensor fusion** across constellations, and selected **quantum experiments** (e.g., space quantum links) requiring proximity to space environment.[9]
- **Edge analytics for intersatellite networks**, reducing backhaul; demonstration architectures (ORBITCAST) project less than a five-minute EO delivery via LEO constellations.[10]

What compute architecture is best (LPUs versus CPUs versus GPUs versus TPUs)?

Choosing the right compute architecture for orbit is about **balancing power, performance, and resilience.** Low-power LPUs and radiation-tolerant FPGAs excel for inference-heavy workloads, while screened GPUs and heterogeneous nodes, combining CPUs with accelerators, deliver higher performance for AI/ML under strict space constraints:

- **LPUs/SoC-FPGAs** for **inference-heavy, low-power** workloads with radiation tolerance strategies.[28]
- **COTS GPUs/accelerators** screened for TID/SEE plus fault-tolerant software when higher performance is essential; Spaceborne Computer proofs inform risk posture.[1, 2, 3]
- **Heterogeneous nodes:** CPU control and LPU/FPGA/NPU accelerators for AI/ML, with ECC memory, lockstep cores, and watchdog telemetry.[23]

> **The TMT play is hybrid by default: policy-aware schedulers on top, zero-trust below, sovereign partitions throughout.”**
>
> Chad Seiler, Head of TMT, KPMG LLP

How do we manage cooling and energy (solar versus nuclear versus beamed power)?

In the vacuum of space, getting rid of heat is a major challenge, so engineers design spacecraft with special radiators and heat pipes that radiate thermal energy away, a process carefully managed based on the craft's orientation and exposure to sunlight. Power is typically generated by large solar arrays that store energy in batteries for when the spacecraft is in shadow, but for high-power or deep-space missions, nuclear options provide a more consistent energy source. While concepts like beaming solar power from space to Earth are still a long-term vision, the research is already influencing the design of new orbital power systems:

- **Thermal:** Radiators, heat pipes, and conductive structures; no ambient convective cooling. Design must balance radiator area, pointing, and eclipse profiles.[19]
- **Power:**
    - **Solar arrays** with energy storage, heritage on ISS and small satellite electric power system; design to survive eclipses and peak loads[19,27]
    - **Nuclear options** (radioisotope thermoelectric generators/fission) for high-duty cycles or deep space, recognized by the United Nations Office for Outer Space Affairs (UNOOSA) principles and active US work (JETSON program)[38, 39]
    - **Beamed power (space-based solar power)** remains long-horizon but under study; not near-term for ODC baseload to Earth, yet components (space solar collection) could inform orbital energy architectures.[40]

The choice of these critical systems involves significant trade-offs between proven heritage and future potential. Given these options, the chief financial officer and chief technology officer must ask: "How do we balance the high up-front investment and risk of developing advanced nuclear or beamed power technologies against the long-term operational limitations and replacement costs of relying on established solar power systems?"

Operational considerations

What is the redundancy and failover strategy?

To ensure ODCs remain operational, they are designed with redundancy in mind, using a constellation of multiple interconnected nodes so that if one fails, then the others can seamlessly take over its tasks. The system is also built to contain faults using radiation-aware software that can automatically detect and correct errors, allowing for rapid recovery without depending on commands from Earth. Furthermore, these space-based centers are tightly integrated with ground stations, creating a hybrid network that guarantees continuous data coverage and prioritizes the most critical information:

- **Constellation-based redundancy:** Multiple ODC nodes with **optical crosslinks** and distributed tasking enable graceful degradation and rolling upgrades.[12]
- **Fault containment:** SEE-aware software (checkpointing, ECC, TMR), watchdogs, and **Earth-independent EDR** reduce mean time to recovery.[4, 27]
- **Ground segment integration:** Hybrid architectures (Ground-Segment-as-a-Service [GSaaS] and ODCs) ensure coverage continuity and priority routing; NASA's ground systems guidance emphasizes end-to-end compatibility testing.[17]

Creating a resilient ODC network requires a complex and deeply integrated system of both space and ground assets. Therefore, management teams must ask, "What is the total cost of ownership for this comprehensive redundancy, and how do we ensure our operational teams have the skills to manage both the orbital and ground segments effectively as they evolve?"

Data sovereignty and jurisdiction, whose laws govern orbital data centers?

Determining which country's laws apply to data in ODCs is complicated, as the Outer Space Treaty gives legal control to the nation that registered the space object, creating a situation where multiple jurisdictions, the registrar, the operator, and the customer, may all have a claim. When personal information is involved, stringent data privacy laws like GDPR must be followed, which demand specific contractual agreements and technical protections such as encryption to handle cross-border data flows. This legal puzzle means that navigating data sovereignty requires a careful blend of international treaty law, national regulations, and private contracts to ensure compliance.

There is a strong consideration to take on a proactive legal posture, combining strategic corporate structuring, robust contracts, and a clear regulatory compliance plan, that is designed to create a defensible legal shield that protects the company from the inherent ambiguities and financial risks of operating in the legally gray area of outer space:

- **Registration and jurisdiction:** Under the **Outer Space Treaty** and **Registration Convention**, a space object is subject to jurisdiction and control of the **State of registry**; payload data governance may then implicate the registry, operator domicile, and customer jurisdiction simultaneously. (Legal overviews in NASA/CISA materials and export-control briefings provide context; apply counsel for transaction-specific interpretations.)[41]
- **Cross-border data flows:** GDPR-class rules can attach when personal data is involved; contractual terms (processing agreements), technical safeguards (encryption, anonymization), and **SCC** analogues must be reconciled with space-specific jurisdiction. (See Regulatory sources above.)[31]

The intersection of space law and data privacy creates significant legal and financial risks. Therefore, management teams must ask, "How do we develop a unified compliance strategy that is robust enough to satisfy the potentially conflicting legal demands of multiple countries, and what is our financial exposure if a jurisdictional dispute arises?"

Implications: From “cloud on Earth” to “cloud in orbit”

The move to place data centers in orbit represents a fundamental paradigm shift, transforming everything from hardware design and business models to the complex rules of cybersecurity and international law. This evolution is driven by the growing limitations of Earth-based data centers and the emergence of new SPaaS financial models, positioning orbital computing to handle specialized workloads where it offers a distinct advantage. As this new ecosystem develops, it brings forth critical board-level challenges that must be managed proactively.

For telecommunications, media, and technology leaders, orbital compute is not about abandoning hyperscale; it's about *placing* the right workload at the right layer, ground or orbit, based on **latency, bandwidth economics, resilience, and regulatory posture.** Think of orbit as another *availability zone* with unique energy and sovereignty properties. The winners will instrument **policy-aware schedulers** to route inference close to sensors while enforcing **zero-trust** and jurisdictional controls across Earth/space links. In capital terms, this is an **operating-expense-friendly augmentation** of your cloud footprint, leveraging **SPaaS** and hosted payloads to compress time-to-value without carrying space hardware on your balance sheet.

ODCs and SPaaS transform compute architecture (heterogeneous, radiation-aware), business models (subscription / utility for orbital infrastructure), and governance (export controls, licensing, cyber resilience, and data jurisdiction). Market signals, from **capital expense surges in terrestrial data centers** and power constraints to **SPaaS financing and ODC prototypes**, suggest orbital compute will complement (not replace) hyperscale cloud, absorbing **space-native workloads** where latency, bandwidth, and resilience dominate:[5, 42]

- KPMG 2024–2025 thought leadership underscores the pressure on energy ecosystems and the need for **collaborative data-center strategies**, a frame transferable to orbital deployments where **energy, regulatory, and cybersecurity** must be engineered coherently across vendors and states.[43]

Market Shift: TMT

For TMT leaders, orbital compute extends the cloud into a new strategic layer, enabling the right workloads to run where latency, resilience, and sovereignty deliver the most value. This is portfolio thinking applied to infrastructure.

- **Debris and stewardship:** Every new orbital data center adds to the risk of space debris, demanding responsible lifecycle planning, including deorbiting strategies and ethical stewardship, as a core part of the business model. Each ODC adds mass and area; lifecycle planning (deorbiting, passivation), shielding mass versus collision risk, and **ethical stewardship** are board-level concerns, aligned to UNOOSA guidelines.[38]
- **Cyber escalation:** Because these systems operate autonomously, a single security flaw could have cascading consequences, making it imperative to build them on "zero-trust" and "secure-by-design" principles from the very beginning. Earth-independent compute heightens the need for **trusted autonomy**; misconfigurations or compromises could propagate across constellations. Follow **zero-trust** and **secure-by-design** tenets.[34]
- **Jurisdictional equity:** Companies cannot treat space as a lawless frontier to bypass terrestrial data privacy and sovereignty regulations; they must build transparent and auditable systems that respect the legal rights attached to the data they process. Data processed "above" Earth must respect **privacy and sovereignty**, avoiding regulatory arbitrage; transparency and auditing will be necessary.[31]

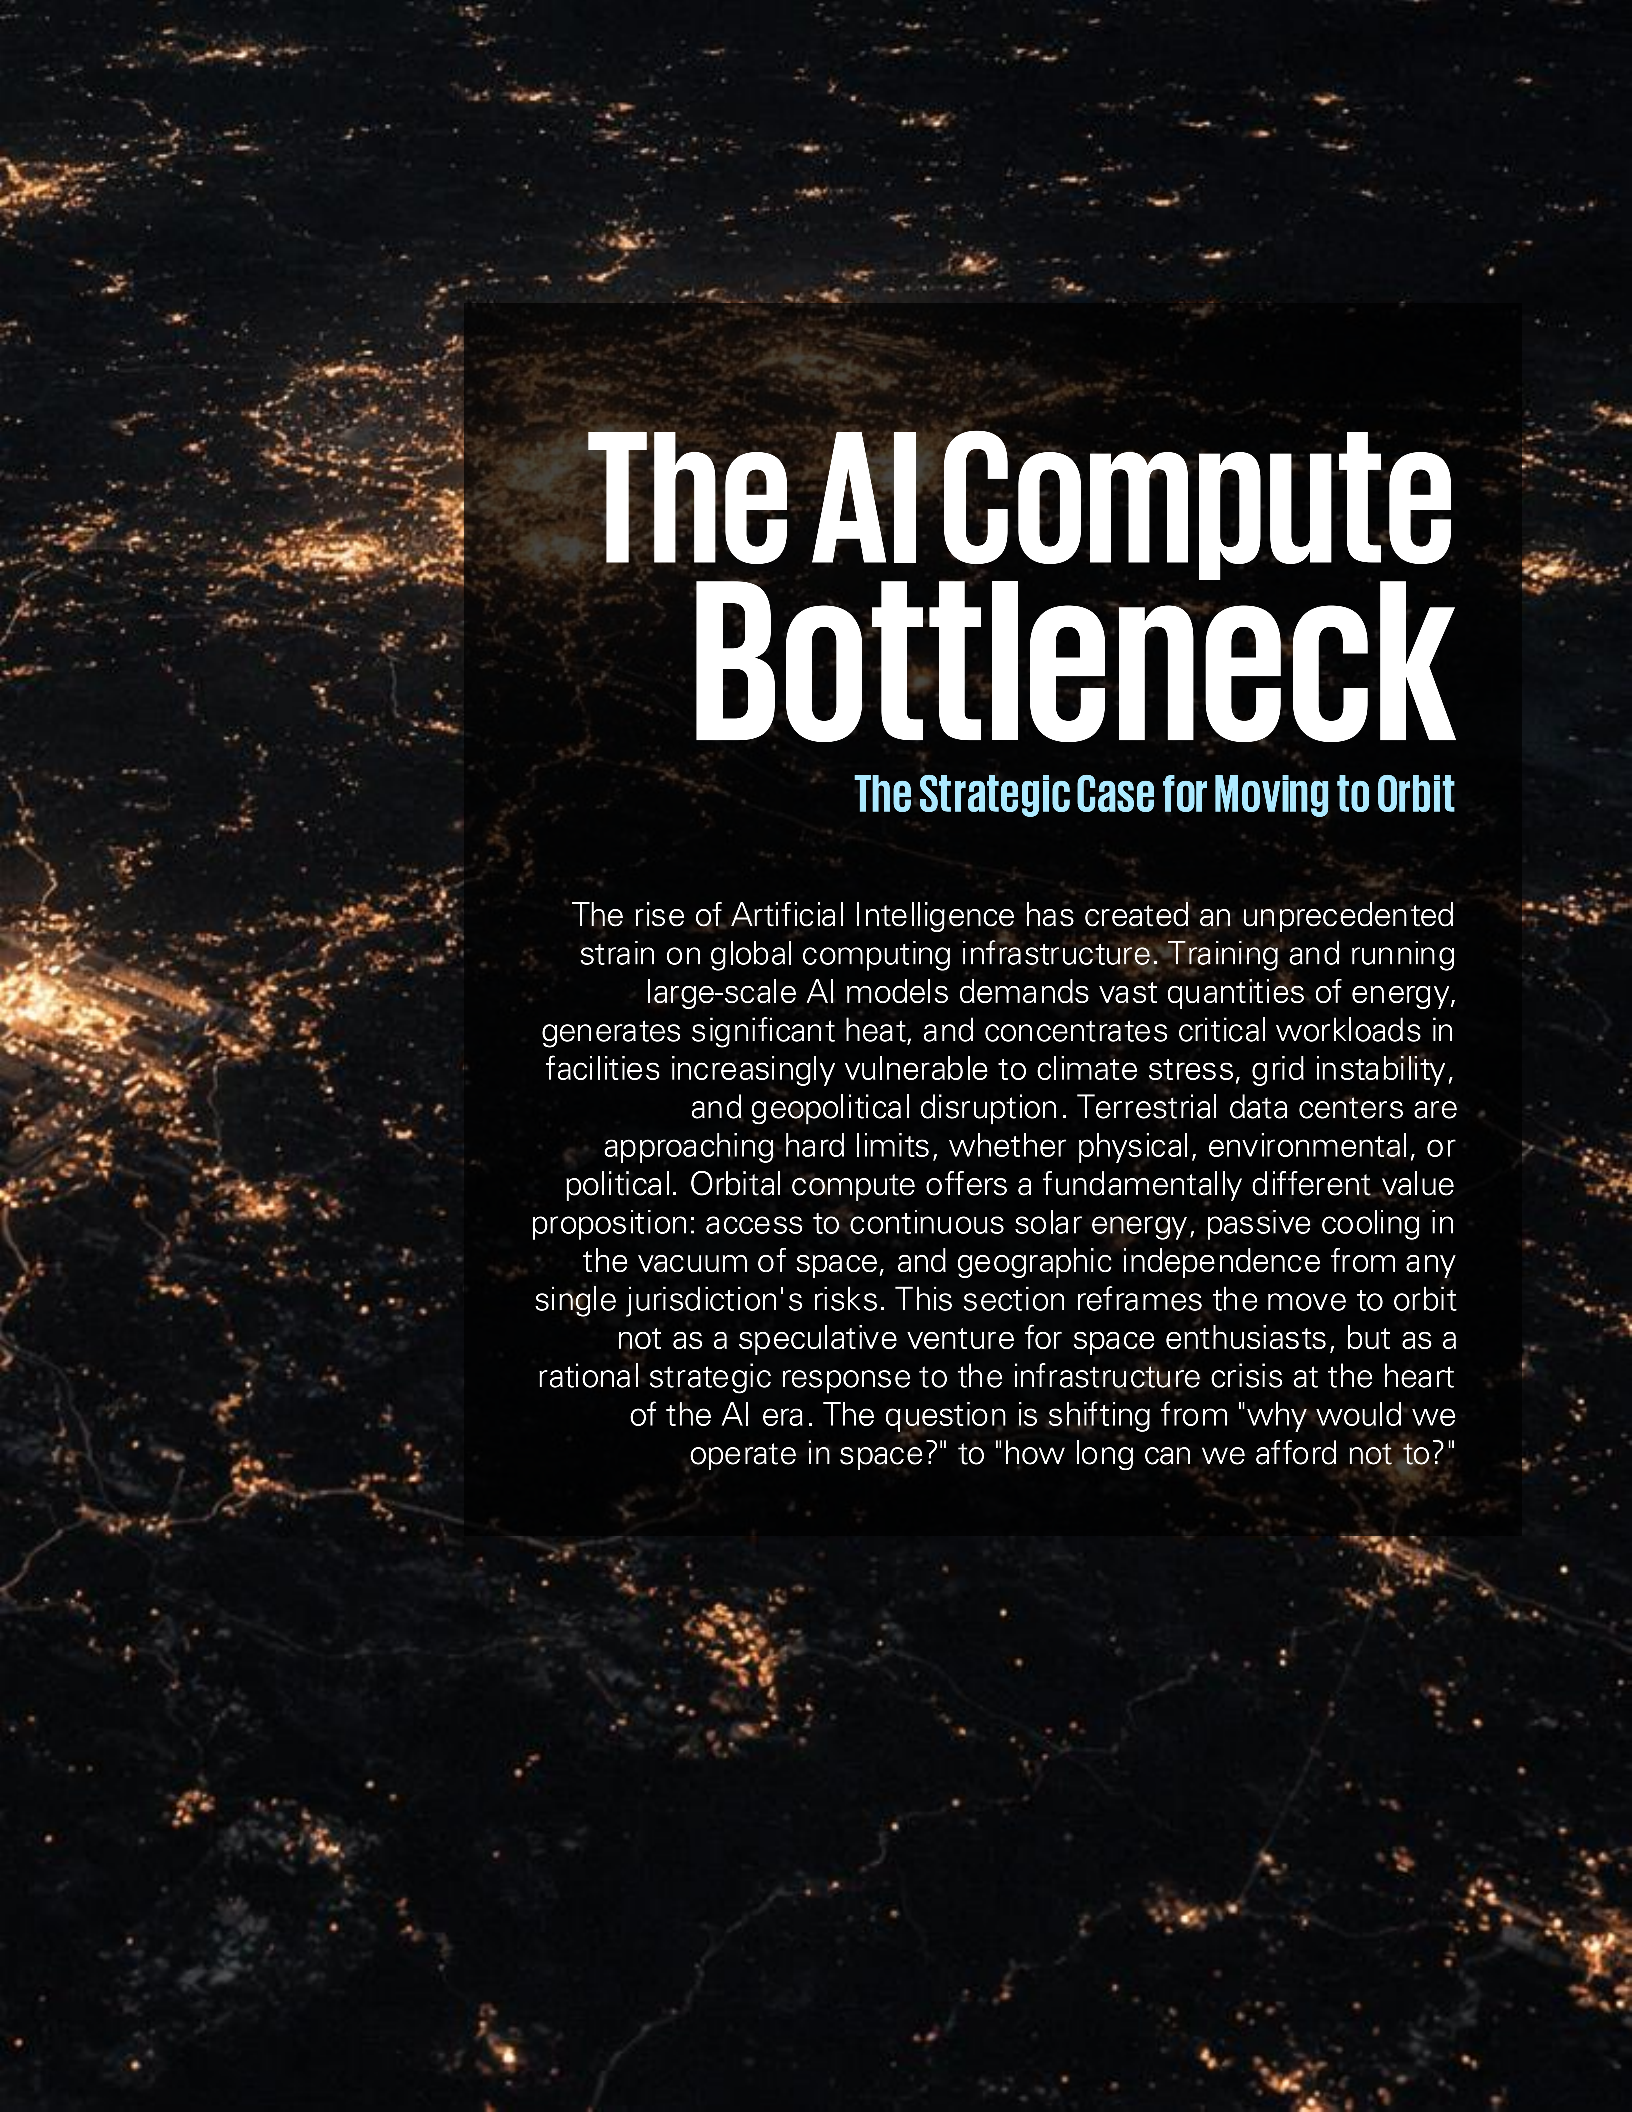

The AI Compute Bottleneck

The Strategic Case for Moving to Orbit

The rise of Artificial Intelligence has created an unprecedented strain on global computing infrastructure. Training and running large-scale AI models demands vast quantities of energy, generates significant heat, and concentrates critical workloads in facilities increasingly vulnerable to climate stress, grid instability, and geopolitical disruption. Terrestrial data centers are approaching hard limits, whether physical, environmental, or political. Orbital compute offers a fundamentally different value proposition: access to continuous solar energy, passive cooling in the vacuum of space, and geographic independence from any single jurisdiction's risks. This section reframes the move to orbit not as a speculative venture for space enthusiasts, but as a rational strategic response to the infrastructure crisis at the heart of the AI era. The question is shifting from "why would we operate in space?" to "how long can we afford not to?"

Orbital Data Centers Timeline

A decade of transformation: from prototype to planetary infrastructure

Scientific Implications

2027 — Proof of Concept Rollouts

Prototype systems with custom AI chips validate real workloads like image analysis at lower energy use, confirming hardware reliability in space.

Enterprises launch pilots for data security and efficiency, building evidence to shift budgets toward off-planet operations.

2028 — Initial Commercial Operations

Scaled satellite networks enable cost-effective training of larger AI models, leveraging massive constellations for distributed compute.

Revenue starts flowing from orbital data centers, drawing mid-sized firms to bypass grid limits and access scalable AI without heavy capex.

2029 — Expansion and Integration

Hybrid Earth-space systems integrate seamlessly, supporting smooth workload shifts with improving orbital connectivity.

Pay-per-use models emerge, letting companies handle AI peaks flexibly and sparking new revenue streams in cloud-orbit hybrids.

2030 — Operational Maturity

Gigawatt-scale power in orbit supports enterprise-level AI training that exceeds terrestrial constraints.

Strategic alliances create shared infrastructure, accelerating AI innovation and opening orbital manufacturing for faster, cheaper hardware cycles.

2031 — Cost Parity Achievement

Enhanced links minimize delays, enabling real-time AI apps from space with performance rivaling ground systems.

Orbital becomes the default for energy-heavy projects, driving VC surges and new jobs in space ops management.

Business Implications

Scientific Implications

2032 – Scalability Breakthroughs

Scientific: Massive expansions handle global demands, with automated maintenance making orbital centers reliable and low-maintenance.

Business: Industries like finance and healthcare adopt widely, while brands leverage "space-powered AI" for sustainable, premium positioning.

2033 – Innovation Acceleration

Scientific: Fully orbital-trained models drive faster AI breakthroughs through uninterrupted, high-scale compute.

Business: Specialized services diversify (e.g., defense, autonomy), fueling GDP growth via new data and tech value chains.

2034 – Mainstream Adoption

Scientific: Over half of large AI compute moves orbital, backed by proven reliability and cheaper launches.

Business: Shared platforms optimize costs industry-wide, with mature risk strategies ensuring steady operations.

2035 – Global Transformation

Scientific: Orbital centers power unlimited-scale AI across sectors, fully integrated into the global compute fabric.

Business: New consulting and integration revenues help firms adapt, while sustainability advantages align with ESG mandates.

2036 – Mature Ecosystem

Scientific: Seamless space-Earth networks deliver high-performance compute universally, with orbital hubs pioneering next-gen AI.

Business: Resilience becomes standard, positioning leaders for uninterrupted growth and brand dominance in tech-forward markets

Business Implications

A practical roadmap

NOW (1-2 years)

Organizations should consider running select data processing tasks directly on satellites for missions where it's most cost-effective, while simultaneously designing hardware that can withstand the harsh radiation and power constraints of space. To accelerate this move and reduce up-front costs, companies should leverage existing “Space-Platform-as-a-Service” providers and ensure they establish a robust legal and cybersecurity compliance program from day one:

1. **Select targeted orbital workloads:** Pilot **on-orbit inference** for EO missions; use GSaaS for integration and prioritize missions with high downlink costs or latency sensitivity.[10, 17]
2. **Architect for radiation and power**: Choose **LPU/SoC-FPGA** accelerators, implement ECC/TMR, and design radiators/battery margins for eclipses; leverage Spaceborne lessons.[1, 28]
3. **Establish compliance program:** Map ITAR/EAR classifications, plan for CSA exceptions, and institute zero-trust cyber controls across space, ground, and link segments.[32, 34]
4. **Adopt SPaaS contracts:** Use hosted payloads (Loft Orbital) and cloud-integrated ground services (Azure Orbital) to reduce cycle time and capital burden.[5, 8]

Treat **NOW** as *evidence generation*: pick **one inference-heavy EO workload** and run it on SPaaS, instrumenting **cost-per-insight**, **GOPS/W**, and **downlink savings** as your key performance indicator triad. In **NEAR**, scale to a **mesh of multitenant nodes** with **optical crosslinks**, codifying **failover playbooks** and **cryptographic data separation** per customer and jurisdiction. By **NEXT**, converge terrestrial and orbital clouds under **policy-aware workload placement** and **SLA-backed orbital Compute-as-a-Service**. Success is not measured by raw performance alone, but by:

1. **Unit economics** (bandwidth and energy arbitrage)
2. **Resilience** (Earth-independent operations)
3. **Compliance by design** (export, spectrum, privacy).

"Energy leaders are grappling with a structural demand shock from AI and data centers that is pushing our terrestrial grid to its limits. Orbital compute isn't science fiction; it's a pragmatic response to this reality. It represents the ultimate 'behind-the-meter' solution, taking the entire compute load, and its immense power requirement, off the terrestrial grid entirely. By collocating processing with dedicated, continuous solar power, we are not just reducing latency; we are creating a new, resilient, and independent infrastructure. For an industry facing intense pressure to ensure reliability and meet unprecedented demand, looking to orbit is a strategic imperative."

Angie Gildea, Partner, KPMG LLP

"In both commercial and military space, true advantage comes from owning the orchestration and control layer—it's the center of gravity for operational dominance."

Joe Cassidy, EMA Head of Technology, Media & Telecommunications, KPMG UK

NEAR (3-5 years)

Looking ahead three to five years, the focus will shift to scaling up by deploying an interconnected mesh of orbital data centers and integrating next-generation energy solutions like advanced solar or small nuclear reactors. Concurrently, it will be essential to codify data governance by implementing cryptographic separation for customer data and establishing firm contractual mechanisms to address the complex legal challenges of sovereignty in space:

1. **Scale ODC mesh nodes:** Deploy multitenant orbital compute with **optical crosslinks**, interoperate with SDA standards, and formalize failover playbooks.[12]

2. **Integrate energy innovations:** Evaluate **high-efficiency solar**, advanced storage, and **small nuclear pilots** where mission profiles demand continuous power.[38, 39]

3. **Codify data governance:** Implement cryptographic separation of customer data, jurisdictional mapping, and contractual mechanisms addressing registry/sovereignty.[41]

NEXT (5-10 years)

Within 5 to 10 years, the goal is to offer on-demand orbital computing as a fully operational service with guaranteed performance, allowing for seamless workload management between ground and space data centers based on cost, latency, and legal requirements. To support this mature ecosystem, companies will need to collaborate with government and industry partners to establish standardized security protocols and harmonized data regulations for space:

1. **Operationalize SPaaS platforms:** Offer **Compute-as-a-Service in orbit** with PED backed by service-level agreement (SLA), sovereign partitions, and threat-resilient operations, aligned to defense and industrial anchors.[4]

2. **Converge terrestrial–orbital cloud:** Abstract workload placement across ground and orbit using policy-aware schedulers that optimize **latency, energy, and regulatory constraints** (industry trajectories noted in Azure Space).[8]

3. **Collaborate on norms:** Participate in cross-industry forums (CISA/White House workshops) to advance **zero-trust** in space and harmonize export/data regimes for orbital compute.[34, 35]

The future in 5 to 10 years could be very different, but the largest differences between today and the future are:

Aspect	NOW (1–2 years)	NEXT (5–10 years)
Primary goal	Prove it works (technical validation)	Make it a global business (market integration)
Focus	Hardware, survival, pilot projects	Software, abstraction, ecosystem, policy
Business model	Cautious, reduce up-front cost, use hosted platforms	Offer guaranteed services (SLA-backed), build a platform
Key question	"Can we build a reliable computer for space?"	"How do we build a global, standardized orbital cloud market?"

Offsetting terrestrial AI compute with low-Earth-orbit data centers

The first installment of this series charted the paradigm shift from terrestrial to orbital data centers (ODCs), revealing how compute is moving from "cloud on Earth" to "cloud in orbit." In Part 2, we examine **a bold alternative: low-Earth-orbit (LEO) data centers as a strategic solution to the artificial intelligence (AI) compute bottleneck.** By harnessing continuous solar power, radiative cooling, and proximity to space-based data sources, orbital compute offers enterprises a transformative path to resilience, efficiency, and sustainability.

Beyond technical novelty, this shift represents a paradigm change in digital infrastructure, unlocking near-instantaneous AI inference, reducing carbon emissions, and enabling a hybrid value chain that integrates terrestrial and orbital clouds. **The imperative is clear: move AI workloads to orbital compute now, because it is critical to overcoming energy constraints, accelerating insight while cutting costs, and securing compliance and resilience at global scale before terrestrial bottlenecks harden into competitive disadvantages.**

We examine the business case, energy advantages, and operational models that position orbital compute as a cornerstone of advantage in the age of AI. Organizations that act decisively will not only shape emerging standards and governance but also lead to a future where agility, compliance, and innovation define digital leadership. The race to orbit is on and the window for transformation is closing.

The AI compute bottleneck and the orbital solution

AI workloads are growing exponentially, with global data center energy consumption projected to reach 1,000 terawatt-hours (TWh) annually by 2026, more than double the 2022 figure and equivalent to the electricity usage of Japan.[47] Hyperscale data centers are straining against grid limitations, land scarcity, and cooling challenges, while the carbon footprint of AI training and inference is under increasing scrutiny.[48] As organizations race to deploy generative AI and real-time analytics, the bottleneck is energy, sustainability, and resilience. LEO compute offers a compelling offset: continuous solar power, radiative cooling, and proximity to space-based data sources. Those who move first to orbital compute will shape the future of digital infrastructure.

> “We stand at a crossroads where the very architecture of intelligence is being reimagined, not merely as a matter of computational power, but as a question of civilization’s capacity to adapt, secure, and sustain its digital future.”
>
> Brian Miske, KPMG LLP

The exponential rise of AI workloads has exposed the limits of terrestrial infrastructure, revealing a bottleneck not only in compute, but also in energy, resilience, and the sovereignty of data itself. As history teaches, those who recognize inflection points and act decisively can shape the destiny of entire industries; today, the leap to orbital compute is such a moment, promising a new era where the boundaries of possibility are defined not by geography, but by vision.

Yet the urgency is unmistakable. The constraints of land, energy, and regulatory complexity are intensifying, and the window for transformation is rapidly closing. In the spirit of bold innovation, organizations must seize the opportunity to migrate workloads to LEO, leveraging continuous solar power, radiative cooling, and proximity to space-based data sources. Those who move first will not only accelerate insights and slash costs, but also set the standards, capture new markets, and lead the charge into a digital ecosystem where agility, compliance, and resilience are the new currency of global leadership.

The value chain: From terrestrial to orbital compute

The possibilities of shifting from terrestrial to orbital compute are nothing short of revolutionary, inviting organizations to imagine a future where digital infrastructure is no longer constrained by land, energy, or bandwidth. By harnessing the power of LEO, enterprises can unlock near-instantaneous AI inference, dramatically improve bandwidth efficiency, and achieve operational resilience that transcends the limitations of traditional data centers. This new value chain, from real-time disaster response to autonomous satellite control, offers a hybrid future where seamless integration between ground and space enables smarter, faster, and more secure decision-making across industries.

Yet the urgency to act is unmistakable. As terrestrial data centers face mounting pressures from energy consumption, land scarcity, and regulatory complexity, the window for transformation is rapidly closing. Organizations that move now to identify, migrate, and optimize workloads for orbital processing will not only slash costs and accelerate insights, but also position themselves as leaders in a digital ecosystem where agility, compliance, and resilience are the new currency of success.

Terrestrial data centers

Current state and constraints

Terrestrial data centers face mounting constraints. In 2024, US data centers consumed over 90 TWh of electricity, representing 2 percent of total US power demand.[47] Land acquisition for hyperscale sites is increasingly competitive, with some regions reporting a 40 percent increase in land costs since 2021.[48] Cooling remains a critical challenge: Up to 40 percent of data center energy is spent on thermal management, and water usage for cooling can exceed 5 million gallons per day at large sites.[49] Grid instability and rolling blackouts in major markets have forced operators to invest in backup generation and microgrid solutions, raising costs and complexity.[50] Regulatory pressures, ranging from carbon taxes to data sovereignty laws, add further risk, while geopolitical tensions threaten supply chains and operational continuity.[51]

The transition point

Workload selection and migration

Not all workloads are suited for orbital processing. The transition begins with identifying AI tasks that benefit from proximity to space-based sensors, real-time inference, and bandwidth efficiency. Earth observation, disaster response, maritime monitoring, and global asset tracking are prime candidates.[52] Policy-aware workload scheduling, balancing cost, latency, energy, and compliance, enables organizations to migrate select AI inference and edge analytics to LEO data centers. Studies show that on-orbit processing can reduce data transmission volumes by up to 90 percent, slashing downlink costs and accelerating time-to-insight.[12]

On-orbit processing

Capabilities and advantages

LEO data centers offer unique capabilities. By collocating compute with satellites and sensors, organizations achieve near-instantaneous AI/machine learning inference, enabling applications such as wildfire detection, crop monitoring, and autonomous satellite control.[53] Bandwidth efficiency is dramatically improved: “data thinning” in orbit means only actionable insights are transmitted to ground stations, reducing spectrum congestion and operational costs.[2] Operational resilience is enhanced, orbital compute is decoupled from terrestrial outages, cyber threats, and physical disruptions. Autonomous mesh architectures and intersatellite laser links provide failover and continuity, supporting mission-critical operations even in contested environments.[4]

"Our terrestrial grid was designed for a different century; it wasn't built to handle the concentrated, 24/7 power demands of AI, nor the increasing frequency of extreme weather events. Orbital compute offers a new layer of resilience by being fundamentally decoupled from these vulnerabilities. It's an infrastructure that is immune to terrestrial blackouts, grid congestion, and local natural disasters. For mission-critical digital operations, that level of reliability isn't just a feature. It's the entire business case."

Angie Gildea, Partner, KPMG LLP

“In an energy transition that requires a ‘mosaic of solutions,’ orbital data centers represent a powerful and unexpected new tile. This represents new class of energy assets and fundamentally redefines what we consider energy infrastructure.”

Angie Gildea, Partner, KPMG LLP

Integration across the value chain

The hybrid future of digital infrastructure is here, where organizations can seamlessly migrate workloads between terrestrial and orbital clouds to unlock unprecedented flexibility, efficiency, and strategic control. By leveraging both environments, enterprises can optimize every task for cost, speed, energy use, and regulatory compliance, transforming orbital compute from a niche experiment into a core pillar of AI, analytics, and secure data processing.

This new paradigm demands more than technical integration. It requires a sophisticated approach to data sovereignty, export controls, and cybersecurity, with ODCs navigating complex international rules and deploying zero-trust architectures to safeguard assets across domains. The organizations that master this hybrid value chain will not only achieve operational excellence but also set the standard for trust, agility, and innovation in the global digital economy.

The future is hybrid. Seamless migration between terrestrial and orbital clouds enables organizations to optimize workload placement based on cost, latency, energy, and regulatory requirements.[54] Orbital compute backed by service-level agreement (SLA) becomes a core component of enterprise digital strategy, supporting AI, analytics, and secure data processing. Data sovereignty and compliance are managed across domains: ODCs must navigate export controls, spectrum licensing, and General Data Protection Regulation-like privacy rules, while implementing zero-trust cybersecurity architectures to safeguard assets.[55]

Energy advantage: Sustainability and efficiency in LEO

The energy advantage of LEO is rewriting the rules of sustainability and efficiency for digital infrastructure, offering organizations a chance to harness the relentless power of the sun and the unique thermodynamic benefits of space. Imagine data centers that never sleep, drawing up to 30 percent more solar energy than their terrestrial counterparts, and shedding heat effortlessly into the void, no longer bound by the constraints of land, water, or grid reliability. Now, with the emergence of beamed power technology, energy harvested in space can be transmitted wirelessly, using microwave or laser beams, from orbital platforms directly to receivers on Earth or other satellites, opening the door to a future where clean, continuous power flows wherever it's needed, untethered by traditional infrastructure.

But the real excitement lies in what's next: As advanced energy solutions like small nuclear reactors, high-capacity storage, and beamed power move from concept to reality, the resilience and sustainability of LEO data centers will only accelerate. Organizations bold enough to embrace this frontier will not only meet regulatory demands but also set the pace for innovation, shaping a digital ecosystem that is as sustainable as it is limitless.

Continuous solar power and thermodynamic benefits

LEO offers continuous access to solar energy, with orbital platforms receiving up to 30 percent more sunlight than terrestrial sites due to minimal atmospheric attenuation.[56] Radiative heat rejection via space-grade radiators enables efficient thermal management without water or land-intensive cooling systems.[57] Academic studies estimate that ODCs could achieve energy usage effectiveness below 1.1, compared to terrestrial averages of 1.5–2.0.[58] Land-neutral siting eliminates the environmental impact of large-scale data center construction, supporting sustainability goals.

Carbon offset and global impact

The potential for carbon offset is significant. If just 5 percent of global AI inference workloads were migrated to LEO data centers, annual carbon emissions from data center operations could be reduced by up to 20 million metric tons, equivalent to taking 4 million cars off the road.[59] Organizations seeking to meet aggressive sustainability targets can leverage orbital compute as a strategic tool, aligning with environmental, social, and governance (ESG) mandates and regulatory requirements.

Future energy innovations

Orbital platforms are exploring advanced energy solutions. Small nuclear reactors, such as the National Aeronautics and Space Administration's (NASA) Kilopower prototypes, promise continuous power for high-duty-cycle missions.[60] Advanced storage technologies and beamed power concepts are under development, with industry and government partnerships driving innovation.[38] These breakthroughs will further enhance the energy efficiency and resilience of LEO data centers.

Strategic opportunities for organizations

The opportunity to offset terrestrial AI compute to orbit is nothing short of transformative, offering organizations a chance to redefine their competitive edge in a rapidly evolving digital landscape. With launch costs plummeting and orbital platforms delivering unprecedented efficiency, resilience, and market access, the business case for this leap is stronger than ever, promising not only operational savings but also a seat at the table where new standards and governance are being set.

Yet the urgency to act cannot be overstated; those who hesitate risk falling behind as first movers capture new markets, shape industry norms, and establish themselves as leaders in the orbital cloud ecosystem. In this pivotal moment, bold decisions and strategic partnerships will determine which organizations thrive in the next era of global digital infrastructure.

Business case for terrestrial AI compute offset

Thinking strategically about offsetting terrestrial AI compute to orbit requires management teams to balance immediate operational gains with long-term transformational potential. In the near term, the dramatic reduction in launch costs and the emergence of commercial orbital platforms make it possible to achieve rapid return on investment (ROI) through bandwidth savings, energy efficiency, and enhanced resilience, positioning the organization ahead of competitors. Looking at the long term, first movers will not only benefit from lower total cost of ownership and improved risk profiles, but also gain influence over the standards, governance, and market access that will define the future of global digital infrastructure.

The ROI for offsetting AI compute to orbit is compelling. Launch costs for orbital platforms have dropped by 95 percent over the past three decades, driven by reusable rockets and commercial launch providers.[61] Bandwidth savings, energy efficiency, and operational resilience translate into lower total cost of ownership and improved risk profiles. First-mover organizations gain competitive differentiation, access to new markets, and influence over emerging standards and governance frameworks.[62]

The Roadmap to Leadership

A practical guide to leadership

Understanding the orbital compute opportunity is necessary but insufficient. Capturing it requires deliberate action, staged investment, and organizational readiness. This section translates the strategic insights of the preceding analysis into a concrete roadmap for decision-makers. It outlines a phased approach, from foundational partnerships and pilot programs to scaled deployment and market leadership, tailored to different organizational starting points and risk tolerances. It also provides a set of board-level questions designed to pressure-test strategic assumptions and surface capability gaps. For organizations prepared to act, first-mover advantage remains available. For those who wait, the window is narrowing.

Our role in realizing the opportunity

KPMG provides wide-ranging advisory for organizations seeking to leverage orbital compute. Services include business case development, regulatory compliance, risk management, technical strategy, and partnership models. Our cross-industry collaboration and thought leadership help position clients to capture value and mitigate risk in the evolving orbital data center ecosystem.[31]

Management teams and board-level considerations

Management teams considerations

As organizations contemplate migrating AI workloads to ODCs, management teams must recognize that this shift involves far more than a one-time investment, it demands ongoing vigilance over operational costs and risk factors unique to space. The complexity of managing assets in orbit, from data transmission and maintenance to ensuring redundancy for mission-critical applications, requires a proactive and systematic approach to cost tracking and risk mitigation. By weighing these operational realities against the resilience and security benefits of orbital compute, management can make informed decisions that protect both the organization's investment and its long-term strategic interests.

Board-level questions

As boards confront the decision to offset AI compute to orbit, they stand at the threshold of a transformative shift in digital infrastructure, one that demands a rigorous, forward-looking assessment of both cost and risk. The stakes are high: orbital platforms promise unprecedented resilience and efficiency, but they also introduce new complexities and exposures that cannot be ignored. Only by benchmarking operational savings against the realities of space and embedding robust risk management from the outset can boards ensure their investments deliver sustainable strategic value.

Management teams and board-level considerations

How should we evaluate and manage the operational costs and risks of shifting AI workloads to ODCs?

Management teams need to conduct a detailed analysis of the operational costs involved in migrating AI workloads to orbit, including not only the initial investment in launch and hardware but also the ongoing expenses for maintenance, ground segment integration, and support. This means tracking the costs of data transmission, monitoring orbital assets, and ensuring redundancy for mission-critical applications. Teams must also actively manage risks such as exposure to space debris, radiation, and potential service interruptions, while balancing these against the benefits of increased resilience to terrestrial outages and cyber threats. Regular reviews of insurance coverage, vendor reliability, and compliance with evolving space regulations are essential to maintain operational continuity and protect the organization's investment.

What is the total cost of ownership and risk profile for offsetting AI compute to orbit?

When evaluating the total cost of ownership for offsetting AI compute to orbit, boards must look beyond initial launch and hardware expenses to include ongoing maintenance, integration with terrestrial ground segments, and end-of-life management for orbital assets. While the cost of deploying compute to orbit has dropped dramatically due to reusable launch vehicles and commercial innovation, organizations must also weigh the risks unique to space, such as exposure to radiation, space debris, and regulatory delays. These risks are counterbalanced by the resilience orbital platforms offer against terrestrial outages, cyber threats, and geopolitical disruptions. A comprehensive assessment should benchmark operational savings in energy and bandwidth against the capital intensity and insurance requirements of orbital infrastructure, ensuring that risk management and compliance strategies are embedded from the outset.

Management teams and board-level considerations

What steps can we take to optimize energy use and sustainability as we integrate orbital compute into our infrastructure?	How do energy and sustainability metrics compare to terrestrial options?
On a day-to-day basis, management teams should monitor and optimize the energy profile of both terrestrial and ODCs, leveraging the continuous solar power and efficient cooling available in orbit. This involves selecting workloads that are best suited for orbital processing, tracking energy usage effectiveness, and reporting on sustainability metrics to meet internal and external ESG requirements. Teams should also stay informed about new developments in space-based energy solutions, such as advanced solar arrays and nuclear power options, and be prepared to adapt operational strategies as these technologies mature. By integrating orbital compute into sustainability reporting and collaborating with facilities and procurement, management can ensure that the organization's digital infrastructure supports long-term environmental goals.	Boards should consider that ODCs benefit from continuous solar power and radiative cooling, which can yield energy usage effectiveness metrics superior to those of terrestrial facilities. Unlike ground-based data centers, which are constrained by land, water, and grid limitations, orbital platforms eliminate the environmental impact of large-scale construction and water-intensive cooling systems. The potential for carbon offset is substantial, as even a modest migration of AI workloads to orbit could result in significant reductions in global emissions. Looking ahead, advancements in space-based energy, such as small nuclear reactors and beamed power, promise to further enhance sustainability and operational resilience, positioning orbital compute as a strategic lever for meeting ESG and regulatory targets.

Management teams and board-level considerations

How do we ensure compliance with legal, regulatory, and cybersecurity requirements in our orbital compute operations?

Management teams are responsible for implementing robust processes to comply with export controls, spectrum licensing, and international data protection laws as they relate to orbital compute. This includes maintaining up-to-date documentation, coordinating with legal and compliance departments, and ensuring that contracts with vendors and partners clearly define responsibilities for data sovereignty and jurisdiction. Cybersecurity must be embedded in daily operations, with regular audits, incident response drills, and supply chain assessments to protect against threats unique to space-based platforms. By fostering a culture of compliance and security, management can safeguard the organization's reputation and avoid costly disruptions or penalties.

What legal, regulatory, and cybersecurity frameworks must be addressed?

The legal and regulatory landscape for orbital compute is complex and evolving. Boards must ensure compliance with export controls, spectrum licensing, and international data protection laws, recognizing that data processed in orbit may be subject to multiple jurisdictions. The Outer Space Treaty and Registration Convention assign control to the state of registry, but customer and operator domiciles may also impose requirements, especially for sensitive or personal data. Cybersecurity is paramount; orbital platforms must be designed with zero-trust architectures and robust supply chain assurance to mitigate autonomous threats and escalation risks. A unified compliance strategy, harmonizing contractual, technical, and operational safeguards, is essential to navigate the multijurisdictional demands and protect both assets and reputation.

Management teams and board-level considerations

What practical actions can we take to position our organization as a leader in the orbital compute market?

To establish operational leadership, management teams should prioritize the rapid piloting of orbital AI workloads, develop partnerships with Space-as-a-Service (SPaaS) providers, and integrate hybrid cloud architectures that allow seamless migration between terrestrial and orbital environments. This requires cross-functional coordination to upskill staff, update operational playbooks, and ensure readiness for new workflows and technologies. Engaging with industry groups, sharing leading practices, and contributing to the development of standards will help the organization stay ahead of competitors and respond quickly to market opportunities. By focusing on execution, continuous improvement, and proactive collaboration, management can drive the organization’s success in the emerging orbital compute ecosystem.

How can organizations capture first-mover advantage in the orbital compute market?

To capture first-mover advantage, boards should prioritize early engagement with SPaaS providers and pilot orbital AI workloads that align with strategic business objectives. By participating in industry forums and regulatory consultations, organizations can help shape emerging standards and governance frameworks, positioning themselves as leaders in the evolving orbital cloud ecosystem. Building internal capabilities, through upskilling teams for hybrid operations and investing in technical, legal, and risk management expertise, will be critical for operational success. Ultimately, those who move early will gain access to new markets, differentiated resilience, and influence over the rules and leading practices that will define the future of orbital compute.

NOW/NEAR/NEXT

A hypothesis and practical roadmap for AI compute offset to orbit

NOW (1-2 years): Technical validation and early adoption

To get started, organizations can run small-scale experiments by sending AI tasks to space using rented hardware or services, allowing them to test how well orbital computing works for their needs. At the same time, they should set up strong compliance and cybersecurity programs, drawing on expert guidance from firms like KPMG, to make sure all regulations are met and risks are managed from day one. Select targeted orbital workloads by prioritizing missions with high downlink costs or latency sensitivity.[53] Organizations should use low-power language processing units (LPUs) and specialized processors to protect against space radiation and power fluctuations, taking advantage of their energy efficiency and fault tolerance.

By combining these with smart error correction and robust thermal management, ODCs can maintain reliable performance even during challenging conditions like eclipses. Architect for radiation and power by choosing LPU/system on a chip field-programmable gate arrays (SoCFPGA) accelerators, implementing error-correcting code and triple modular redundancy, and designing radiators/battery margins for eclipses.[63] To adopt SPaaS contracts, use hosted payloads (LoftOrbital) and cloud-integrated ground services (Azure Orbital) to reduce cycle time and capital burden[5]:

Pilot projects

Run select AI inference tasks directly on satellites for missions where cost and latency are most critical.[52]

SPaaS engagement

Leverage hosted payloads and cloud-integrated ground services to accelerate deployment and reduce capital intensity.[4]

Skill development

Upskill operational teams to manage both orbital and ground segments as hybrid architectures emerge.[49]

Hardware design

To make the battery last a very long time and handle tasks efficiently, we use special low-power chips (LPUs) along with flexible, reprogrammable "all-in-one" chips called SoCFPGAs. Engineer for radiation and power constraints using LPUs, SoCFPGAs, and robust thermal management.[28]

Compliance & cybersecurity

Map International Traffic in Arms Regulations, Export Administration Regulations, and classifications; plan for controlled space access exceptions; and institute zero-trust cyber controls across space, ground, and link segments.[32]

NEAR (3-5 years): Scaling and integration

To get started, organizations can run small-scale experiments by sending AI tasks to space using rented Teams should identify opportunities to scale mesh nodes within their orbital compute architecture, ensuring that distributed processing and redundancy are built into the system for greater resilience and performance. At the same time, they should actively pursue integration of energy innovations such as advanced solar arrays, energy storage solutions, or even small nuclear reactors, to support continuous and efficient operations in orbit. Deploy multitenant orbital compute with optical cross-links, interoperate with Space Development Agency (SDA) standards, and formalize failover playbooks.[12] Integrate high-efficiency solar, advanced storage, and small nuclear pilots where mission profiles demand continuous power.[38]

Codify data governance by implementing cryptographic separation of customer data, jurisdictional mapping, and contractual mechanisms addressing registry/sovereignty.[64] As organizations prepare for the next phase of orbital compute, management teams must also anticipate the need for postquantum cryptography to safeguard sensitive data and communications against future quantum-enabled threats. Integrating quantum-resistant encryption protocols into orbital data center architectures will be essential for maintaining long-term security and regulatory compliance as quantum computing capabilities mature:

Mesh network deployment

Expand orbital compute nodes with optical cross-links and distributed tasking for resilience and scalability.[12]

Energy innovation

Integrate advanced solar, storage, and nuclear pilots to support continuous, high-duty-cycle workloads.[65]

Data governance

Implement cryptographic separation, jurisdictional mapping, and contractual mechanisms for sovereignty and compliance.[64]

Market expansion

Formalize failover playbooks and interoperate with SDA standards to support defense, industrial, and commercial anchors.[38]

Regulatory harmonization

Collaborate with government and industry partners to establish standardized security protocols and harmonized data regulations for space.[39]

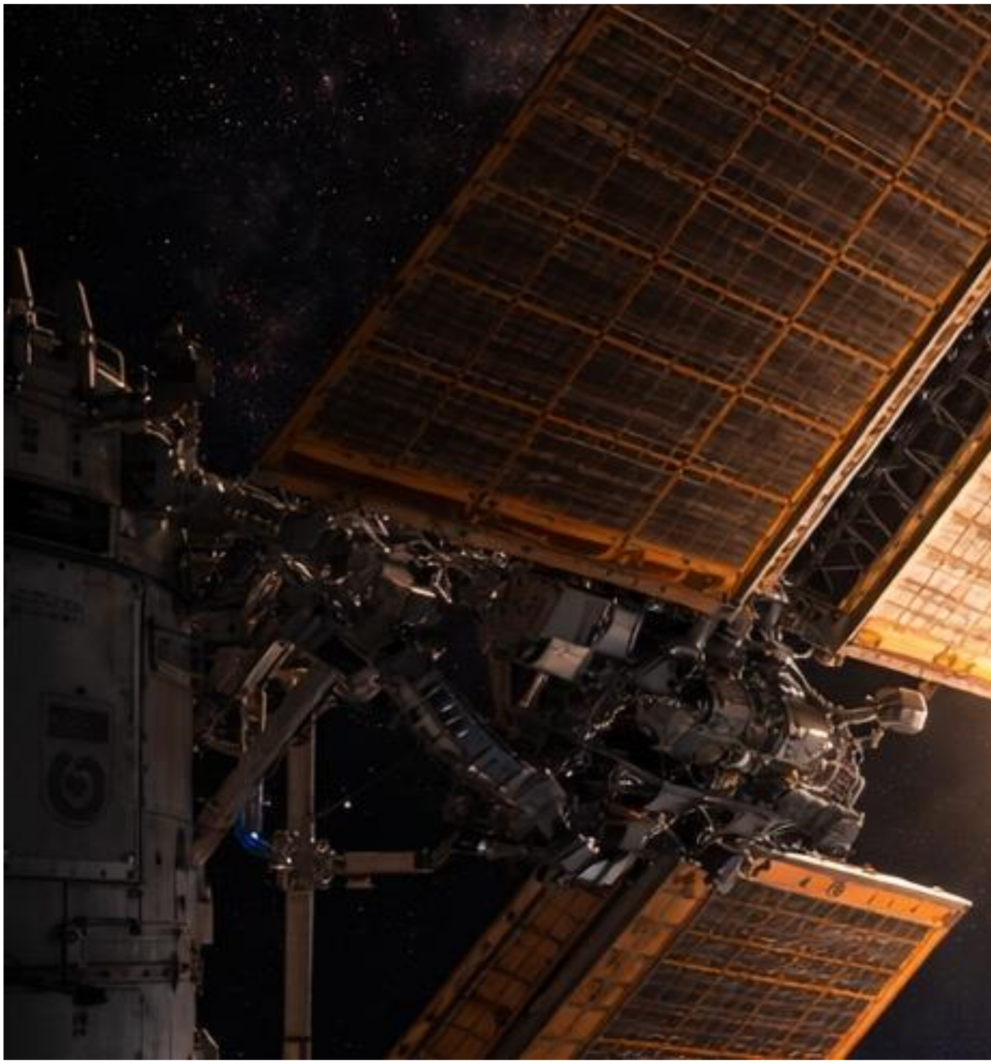

NEAR NEXT (5-10 years): Market integration and ecosystem maturity

Operationalize orbital compute-as-a-service with SLA-backed performance guarantees. Converge terrestrial and orbital clouds, enabling policy-aware workload migration. Look for ways to offer Compute-as-a-Service in orbit with SLA-backed processing / exploitation / dissemination, sovereign partitions, and threat-resilient operations, aligned to defense and industrial anchors.[4]

Actively engage and collaborate on global norms for cybersecurity, export controls, and data protection[34]:

Compute-as-a-Service

Offer on-demand orbital computing with SLA-backed guarantees, enabling seamless workload management between ground and space data centers.[34]

Cloud convergence

Abstract workload placement across terrestrial and orbital clouds using policy-aware schedulers that optimize latency, energy, and regulatory constraints.[5]

Global norms

Participate in cross-industry forums to advance zero-trust in space and harmonize export/data regimes for orbital compute.[54]

Ecosystem leadership

Shape the global market for orbital cloud, influencing standards, governance, and leading practices.[34]

The world is moving from cloud on Earth to cloud in orbit, and it's happening faster than most leaders realize.

Why does this matter? Because the way we process data is about to change forever. ODCs can cut delays, keep operations running even when Earth systems fail, and unlock new business models like SPaaS. The burning platform is clear: if your organization waits, then you risk falling behind in a market that could redefine global connectivity and resilience. The opportunity? Start small pilot on-orbit analytics, partner with hosted payload providers, and build cybersecurity and compliance into your design now. Those who act first will shape the rules and capture the value.

The shift from **cloud on Earth** to **cloud in orbit** is underway. ODCs and SPaaS will not replace terrestrial hyperscalers; they will **augment** them, serving space-native workloads where **latency, bandwidth, autonomy, and resilience** are paramount.

As we push our computational infrastructure beyond the atmosphere, we are not merely launching servers, we are launching a new set of rules for commerce, security, and law. The ultimate success of this endeavor will not be measured by processing speeds or latency reduction, but by our ability to answer the profound challenges that arise when our data inhabits the space above nations.

This new frontier demands a new brand of leadership, one that recognizes that a commitment to cybersecurity, responsible stewardship of the orbital environment, and equitable legal frameworks is not only a matter of compliance, but also the very foundation of a lasting digital legacy in the cosmos.

Enterprises that move now, via hosted payloads, on-orbit inference pilots, and zero-trust architectures, will be positioned to capture first-mover advantages in a market projected to expand rapidly with **space-enabled digital infrastructure.**[12, 44]

"We are moving from cloud on Earth to cloud in orbit, and this shift will redefine how data is processed, protected, and governed. On-orbit data centers won't replace terrestrial clouds; they will extend them where resilience, autonomy, and latency matter most. Enterprises that act now will shape the future of global connectivity."

Jonathon Gill, Global Head of Industrial Markets & Aerospace, KPMG UK

Leadership at the edge, seizing the orbital compute advantage

In moments of profound technological transition, true leadership is defined not by passive observation, but by the courage to act at the frontier of possibility. The migration from terrestrial to orbital compute is no longer a distant vision, it is unfolding now, reshaping the very foundations of AI, energy, and global digital infrastructure. Those who seize this moment, piloting orbital AI workloads, forging strategic partnerships, and embedding compliance and cybersecurity from the outset, will not merely adapt to change, they will set the rules, capture the value, and lead at the edge of everywhere.

"Earth observation data will present organizations with new challenges in sifting through the sheer volume of available data. Orbital edge computing presents the opportunity to provide immediate and actionable intelligence to personnel in the field, accelerating the ability to respond to events as they unfold."

Neil Rae, Head of Space, KPMG UK

Authors and Contributors

Brian Miske

Americas Space Lead
National Ignition Leader, KPMG
bmiske@kpmg.com

Jonathon Gill

Global Sector Lead IM&D, KPMG
jonathon.gill@kpmg.co.uk

Joe Cassidy

EMA Head of TMT, KPMG
joe.cassidy@kpmg.co.uk

Chad Seiler

Head of TMT, KPMG
cseiler@kpmg.com

Chris Brown

Ireland Head of Strategy, KPMG
christopher.brown@kpmg.ie

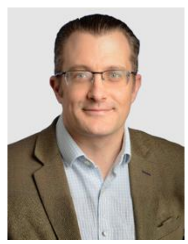

Neil Rae

UK Head of Space, KPMG
neil.rae@kpmg.co.uk

Angie Gildea

Partner, KPMG
agildea@kpmg.com

David Pessah

Discovery Ignition Leader, KPMG
dpessah@kpmg.com

Jozy Assa-Jaconson

Director, Ignition, KPMG
jassajacobson@kpmg.com

Ready to explore the *space economy?*

Contact KPMG today to discover how your organization can thrive in the Space Economy.

+Space Ignition Workshops

Explore the growing space economy! This immersive, design-thinking-fueled workshop takes you into the heart of the space revolution, examining the wealth of opportunities orbiting our planet. Through highly engaging, interactive exercises, you'll uncover the near-term realities and future potential of space-based ventures.

Collaborate with your team and subject matter professionals to identify strategies that can benefit your organization and team, ensuring you are well positioned in the expanding space market now, near, and next.

Take a deep dive into the specifics of the space economy and discover how your organization can benefit today. This focused workshop offers practical insights into the applications of space technology with tailored use cases for your team. Collaborate with other participants to identify potential partnerships within the space sector, fostering ventures that amplify your impact. You will leave with clear strategies for leveraging immediate opportunities and aligning with space-based assets and technologies for future growth.

KPMG Ignition is ready to transform your organization's capabilities and contribute to the future of space exploration and commerce.

Brian Miske
Americas Space Lead
National US
Ignition Leader

David Pessah
Executive Director
Ignition Leader

Jozy Assa-Jacobson
Director
Ignition

References

1. NASA. "Cutting-Edge Computing Goes Spaceborne." NASA Spinoff, 2025
2. OrbitsEdge Inc. "Edge in Space." OrbitsEdge, Date unavailable
3. Data Center Frontier. "Data Centers in Orbit? Space-Based Edge Computing Gets a Boost." October 10, 2022
4. Axiom Space. "Axiom Space to Launch Orbital Data Center Nodes to Support National Security, Commercial, International Customers." Axiom Space, 2025
5. Loft Orbital. "About Us." Loft Orbital, Date unavailable
6. TechFundingNews. "Loft Orbital's €170M Funding Boosts 'Space-as-a-Service' Model, Advancing European Space Autonomy." TechFundingNews, 2025
7. NASA. "2.0 Complete Spacecraft Platforms." NASA SmallSat Institute, Date unavailable
8. Microsoft Azure. "New Azure Space Products Enable Digital Resiliency and Empower the Industry." Microsoft Azure Blog, 2022
9. Bradley Denby and Brandon Lucia. "Orbital Edge Computing: Machine Inference in Space." Clemson University, 2019
10. Zeqi Lai, Qian Wu, Hewu Li, Mingyang Lv, and Jianping Wu. "OrbitCast." ICNP, 2021
11. Yahoo Finance. "Orbital Data Center Launching to ISS to Advance Space Computing." Yahoo Finance, 2025
12. Cutter Consortium. "On-Orbit Data Centers: Mapping the Leaders in Space-Based AI Computing." Cutter Consortium, 2025
13. Data Center Dynamics. "On-Orbit Computing Startup Sophia Space Raises $3.5m, Promises Orbital Data Centers." Data Center Dynamics, 2025
14. Hackaday. "Space-Based Datacenters Take the Cloud Into Orbit." June 19, 2025
15. Deepak Vasisht and Ranveer Chandra. "A Distributed and Hybrid Ground Station Network for Low Earth Orbit Satellites." University of Illinois, 2021
16. World Economic Forum. "Cyber Resilience in Space Is Essential for Economic Security." World Economic Forum, 2025
17. NASA. "11.0 Ground Data Systems and Mission Operations." NASA SmallSat Institute, 2025
18. Tech Xplore. "Scientists Propose Carbon-Neutral Data Centers in Space." Tech Xplore, 2025
19. Wikipedia Contributors. "Electrical System of the International Space Station." Date unavailable
20. University of Luxembourg. "Space as a Service." University of Luxembourg, 2022
21. Kratos. "Space-as-a-Service: Building an Ecosystem of Commercial Infrastructure." Kratos, 2025
22. Viva Technology. "How Space-as-a-Service Is Transforming Access to Satellites." VivaTech, 2025
23. Justin Goodwill, Christopher Wilson, and James MacKinnon. "Current Technology in Space v4 Briefing." NASA Technical Reports Server, 2024
24. NASA. "Open Source Radiation Hardened by Design Technology." NASA Technical Reports Server, 2016
25. Jonathan A. Pellish. "ASEC 2017 Tutorial on Radiation Effects." NASA, 2017
26. Christian Poivey. "Total Ionizing and Non-Ionizing Dose Radiation Hardness Assurance." CERN, 2019
27. NASA NEPP. "Single Event Effect Criticality Analysis." NASA, 1996
28. Microchip Technology. "Radiation-Tolerant PolarFire® SoC FPGAs Offer Low Power, Zero Configuration Upsets, RISC-V Architecture for Space Applications." Microchip Technology, 2024
29. Renesas. "Enable High-Performance, Low-Power Inference for Your Edge AI Applications." February 13, 2024
30. The Space Review. "Prospects for Orbital Data Centers." The Space Review, 2024
31. Holland & Knight. "New Export Control Rules Present Key Regulatory Changes for Space-Related Items, Activities." Holland & Knight, 2024

32. Federal Register. “Export Administration Regulations: Revisions to Space-Related Export Controls.” October 23, 2024
33. Akin Gump Strauss Hauer & Feld LLP. “First Significant Changes in Over a Decade to US Export Controls on Space-Related Items and Activities.” 2024
34. CISA. “Recommendations to Space System Operators for Improving Cybersecurity.” CISA, 2024
35. The White House. “Space System Cybersecurity Industry Perspectives Report.” January 2025
36. KPMG. “Cybersecurity in the Space Sector.” KPMG, 2023
37. Karen Cunninghum, John Carr, and Brandon Lewis. “Spacecraft Electrical Power Systems.” NASA Technical Reports Server, 2018
38. World Nuclear Association. “Nuclear Reactors and Radioisotopes for Space.” World Nuclear Association, 2025
39. Military & Aerospace Electronics. “Nuclear Power Satellites Space Stations.” Military & Aerospace Electronics, 2023
40. NASA. “Space-Based Solar Power.” NASA, 2024
41. NASA Office of International and Interagency Relations. “Export Control Briefing.” Date unavailable
42. McKinsey. “Building Data Centers Bigger, Faster.” McKinsey, 2025
43. KPMG. “From Crisis to Opportunity: Building a Bridge Between AI and Sustainable Energy.” KPMG, 2025
44. World Economic Forum. “Space Is Booming. Here's How to Embrace the $1.8 Trillion Opportunity.” World Economic Forum, 2024
45. International Energy Agency. “Data Centres and Energy – Global Trends.” International Energy Agency, 2024
46. World Economic Forum. “AI and Sustainability: The Energy Challenge.” World Economic Forum, 2024
47. U.S. Department of Energy. “Electricity Consumption in U.S. Data Centers.” U.S. Department of Energy, 2024
48. CBRE Research. “Global Data Center Trends.” CBRE, 2024
49. Uptime Institute. “Annual Data Center Survey.” Uptime Institute, 2023
50. McKinsey & Company. “Scaling Data Centers: Grid and Resilience Challenges.” McKinsey, 2024
51. NASA Goddard Space Flight Center. “AI Onboard: Feasibility and Use Cases.” NASA Technical Reports Server, 2024
52. HPE and NASA. “Spaceborne Computer Missions.” NASA Spinoff, 2024
53. Microsoft. “Azure Orbital Ground Station: Bridging Satellite and Cloud.” Microsoft Azure Blog, 2024
54. European Space Agency. “Space Data Governance and Compliance.” European Space Agency, 2024
55. OECD. “Space Economy: Energy and Sustainability.” OECD, 2024
56. Springer Nature. “Thermal Management in Space Platforms.” Springer, 2023
57. MDPI. “Energy Usage Effectiveness in Orbital Data Centers.” MDPI, 2024
58. Novaspace. “Space Economy Forecasts.” Novaspace, 2024
59. NASA. “Kilopower: Small Nuclear Reactors for Space.” NASA, 2023
60. Space Foundation. “Global Space Economy Report.” July 18, 2024
61. KPMG. “AstroEconomy: The Next Frontier.” KPMG, 2025
62. Microchip. “Radiation-Tolerant PolarFire SoC FPGAs.” Microchip, 2024
63. NASA. “Legal Overview: Registration and Jurisdiction in Space.” NASA, 2024
64. World Nuclear Association. “Nuclear Power for Space Applications.” World Nuclear Association, 2024

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**