




India's Strategic Opportunity in **Space-Based Compute Infrastructure**

16 JUNE 2026

Table of Contents

Part I — The Compute Imperative and India’s Position	6
1.1 Where the World Is Headed	6
1.2 The Terrestrial Ceiling	7
1.2.1 Power Constraints	7
1.2.2 Water Constraints	8
1.2.3 Land Constraints	8
1.2.4 Regulatory and Permission Constraints	8
1.2.5 The Earth Observation Bottleneck	8
1.2.6 Compute as a Strategic Vulnerability	9
1.3 How does India see AI compute?	9
Part II — The Orbital Compute Paradigm: Race for Orbital Data Centres	11
2.1 The paradigm shift: From satellites that sense to infrastructure that computes	11
2.1.1 Entry wedge: The connecting link	11
2.2 The convergence thesis	13
2.3 The constraints	14
Part III — Legal and Governance Dimensions	15
3.1 International Legal Framework	15
3.1.1 The Outer Space Treaty and its successors	15
3.1.2 ITU framework and spectrum governance	16
3.1.3 Space sustainability and debris governance	16
3.1.4 Data governance in outer space: the emerging gap	17
3.2 India’s Domestic Legal and Regulatory Landscape	17
3.2.1 Indian Space Policy 2023 and IN-SPACe	18
3.2.2 Data protection and cybersecurity	18
3.2.3 Spectrum and telecommunications	19
Part IV — India’s Strategic Position and Policy Recommendations	20
4.2 Extend Existing National Frameworks	21
4.3 Develop a Longer-Term Legislative Roadmap for Orbital Infrastructure	2

Executive Summary

The rapid growth of artificial intelligence is driving unprecedented demand for compute infrastructure. AI is no longer constrained primarily by algorithms or semiconductor capability; increasingly, it is constrained by the physical systems required to support computation itself. While terrestrial data centres remain the foundation of today's digital economy, accelerating demand is intensifying pressure on power availability, cooling requirements, land acquisition, and infrastructure concentration—raising broader questions about the long-term scalability and geographic distribution of compute.

At the same time, another bottleneck is emerging at the edge of the digital economy. Modern Earth Observation (EO) constellations generate unprecedented volumes of imagery and sensor data across climate monitoring, disaster response, agriculture, logistics, and strategic applications. Yet transmission capacity has not scaled at the same pace. A substantial share of data generated in orbit is never transmitted, stored, or analysed because downlink bandwidth remains finite and economically selective. As sensing capabilities continue to improve, moving all data to Earth before processing becomes increasingly inefficient.

These two trends are beginning to converge. AI demand is progressively shifting from training-intensive architectures toward inference-led deployment, where value is created through rapid processing close to the point of data generation. In parallel, compute itself is becoming increasingly infrastructure constrained. Together, these forces are creating interest in Orbital Data Centres (ODCs) and orbital computing as a potential new layer of digital infrastructure—one that extends compute capabilities beyond terrestrial boundaries and shifts processing closer to where data originates.

Orbital compute does not replace terrestrial infrastructure; it complements it. By integrating onboard compute, storage, and AI inferencing, orbital systems enable selected workloads to be processed in orbit and return insights rather than raw datasets. This has implications not only for efficiency, but also for latency, resilience, continuity, and the economics of data-intensive applications.

This transition remains early-stage and orbital infrastructure today does not compete with terrestrial compute on cost alone. However, technological progress in launch economics, onboard accelerators, storage architectures, and inter-satellite connectivity is narrowing historical barriers. On TakeMe2Space's modelling, cost convergence begins to emerge around 2035–2036. The significance of that crossover is not immediate substitution—it is that the standards, governance models, operational experience, and strategic positioning that determine leadership will be established well before economic parity arrives.

India enters this transition with meaningful advantages: a mature space ecosystem, expanding private participation, growing AI ambition, and strong digital governance capability. The policy challenge is therefore not whether to replace existing frameworks, but how to progressively evolve them to support a new infrastructure category while it remains formative rather than after market structures consolidate.

This paper argues for a phased approach:

Recommendations:

1. Establish an IN-SPACe–MeitY Orbital Compute Regulatory Sandbox.

India should prioritise a supervised regulatory environment for testing orbital data processing, AI inferencing, cybersecurity controls, and compliance frameworks under real operational conditions.

2. Integrate Orbital Compute into Existing National Infrastructure Frameworks.

Initiatives such as the IndiaAI Mission and national cybersecurity programmes should progressively recognise orbital infrastructure as an emerging extension of India's digital and AI stack.

3. Develop a Long-Term Legislative Roadmap.

India should begin clarifying the applicability of domestic data governance, cybersecurity, and critical infrastructure protections to space-based compute systems, ensuring regulatory certainty while preserving flexibility for innovation.

Part I — The Compute Imperative and India's Position


The defining infrastructure contest of this decade is not over chips. It is over the physical regime required to run them. Afterall what is AI without infrastructure?

1.1 Where the World Is Headed

The world is headed toward mass compute and inference, and the logic underneath it is simple: the one with faster access to information wins. The game is no longer only about massive compute, but about massive compute delivered the fastest way — and, just as critically, the most accurate and secure way, for storage as much as for processing. That reframing is what turns compute from a commercial line item into national infrastructure.

The demand for AI compute is growing so rapidly that infrastructure itself is becoming a strategic asset. xAI announced Colossus, a supercluster combining 200,000 GPUs, and that capacity is now being resold as AI-compute-as-a-service “at significant scale,” to support the growing demand for AI training and inference. That compute capacity is now being commercialised at significant scale, reports indicate that Anthropic has committed approximately USD 1.25 billion per month through 2029 for compute services, while Google is expected to pay around USD 920 million per month between 2026 and 2029 under similar arrangements.

Beyond these developments, major industry players such as NVIDIA continue to invest heavily in AI infrastructure, committing tens of billions of dollars across computing ecosystems, cloud capacity, and related investments, including a recent USD 2 billion cloud expansion initiative.



The AI industry is booming, with little sign of slowing. Data-centre investment reached an all-time high of **USD 61 billion** in 2025, the ecosystem already supports **5.5 million jobs** and contributes approximately **USD 927 billion to global GDP**, and the two workloads driving it are still early in steep growth curves. Earth Observation data and AI both linked to their time-sensitive needs. The Earth Observation market is projected to expand from **USD 7.68 billion to USD 14.55 billion** over 2026–2034 (an **8.31% CAGR**), while the AI inference market grows from **USD 117.80 billion to USD 312.64 billion** (a **12.98% CAGR**).

India is not left behind either: investments of around **USD 70 billion** in data centres¹ — particularly AI data centres, are already underway, with a further **~USD 90 billion** announced, and Budget 2026–27 has backed the sector with a **tax holiday to 2047** to position India among the leading global destinations for AI and cloud infrastructure.

But these data centres are increasingly contested, heavily protested over their power², water, and land demands, and, as concentrated critical infrastructure, targeted in conflict³. And the load is only rising: the International Energy Agency's Energy and AI analysis projects worldwide data-centre electricity demand to **more than double, from 415 TWh in 2024 to roughly 945 TWh by 2030**, slightly more than Japan's entire consumption today, and to continue toward **roughly 1,200 TWh by 2035**, with AI as the primary driver and demand from AI-optimised data centres alone set to **more than triple**. In the United States, data centres are on course to account for nearly **half of all electricity-demand growth to 2030**.

The answer is not to constrain growth in compute demand, AI adoption, digital transformation, and data generation will continue to accelerate. Rather, the challenge is to expand beyond the physical constraints that increasingly shape terrestrial infrastructure. To understand why alternative architectures are beginning to attract attention alongside the existing ones, it is first necessary to examine the limitations emerging within the terrestrial model itself.

1.2 The Terrestrial Ceiling

Every gigawatt of terrestrial AI compute now competes for four scarce and increasingly contested resources. None is an absolute limit today. But each rises as compute demand compounds, and each is fundamentally a function of the same thing: being built on contested ground inside a populated jurisdiction. For a country like India, the four are already visible.

1.2.1 Power Constraints

Hyperscale infrastructure in India is now measured in gigawatts, not floor space. Large-scale data-centre investments are already being proposed in Andhra Pradesh, including a planned 1.5 GW cluster in Visakhapatnam⁴, as part of the state's broader ambition to build approximately 6.5 GW of compute capacity over the coming years⁵. Even a single 100 MW hyperscale facility can consume electricity equivalent to a large village, and AI workloads create highly dynamic

¹ <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2221894®=3&lang=2>

² <https://www.cbc.ca/news/canada/ai-data-centre-opposition-explained-9.7225796>

³ <https://responsiblestatecraft.org/saudi-data-center-iran-war/>

⁴ https://www.fortuneindia.com/technology/google-to-build-15-billion-1-gw-ai-data-centre-in-andhra-pradesh-state-targets-65-gw-digital-capacity/134657?utm_source=chatgpt.com

⁵ https://www.fortuneindia.com/technology/google-to-build-15-billion-1-gw-ai-data-centre-in-andhra-pradesh-state-targets-65-gw-digital-capacity/134657?utm_source=chatgpt.com#google_vignette

demand patterns that stress grid stability and require dedicated substations and transmission infrastructure. Power is not merely an input cost here; it is the rate limiter on how fast capacity can be built at all.

1.2.2 Water Constraints

Cooling turns compute into a major water consumer. Google reported approximately 31 billion litres of water consumption across its global data-centre fleet in 2024. The problem is acute in India specifically, where a recent assessment found that more than 50% of the country's data centres sit in water-stressed regions. In Visakhapatnam — where major hyperscale facilities are planned — groundwater availability is already among the lowest in Andhra Pradesh, placing digital expansion in direct competition with local water security.

1.2.3 Land Constraints

The physical footprint of compute is expanding rapidly. Google's proposed AI Hub in Andhra Pradesh has been allocated 480 acres of land — before accounting for the supporting substations, transmission corridors, renewable-energy assets, cooling systems, and connectivity networks that surround any large cluster. As compute scales, it competes for scarce land and places growing pressure on local infrastructure and communities.

1.2.4 Regulatory and Permission Constraints

Even where power, water, and land exist, the right to use them is becoming its own ceiling. Data centres currently lack a dedicated environmental-clearance category in India despite their substantial resource requirements, and large projects are attracting rising scrutiny over energy sourcing, water consumption, transmission infrastructure, and community impact. Social and regulatory acceptance is now a binding constraint in its own right — and, unlike the other three, it tends to tighten precisely as projects grow large enough to matter.

1.2.5 The Earth Observation Bottleneck

While the first four constraints describe a ceiling that hardens over the next decade, Earth Observation shows the same physics already binding today. Modern constellations generate petabytes of imaging data daily across climate monitoring, disaster response, agriculture, and defence intelligence. But, less than 10% is ever downlinked to the ground⁶. The constraint is no longer data acquisition; it is the bandwidth, latency, and terrestrial processing capacity needed to downlink, store, and compute that stream within operationally meaningful timeframes. This is the structural inefficiency of the current architecture in its sharpest form: data is generated in orbit

⁶ https://www.researchgate.net/publication/222009691_Selecting_Pixels_for_Kepler_Downlink

but computed on Earth. It is also the clearest near-term case for shifting computation to the point of capture — performing inference in orbit.

1.2.6 Compute as a Strategic Vulnerability

On top of these physical ceilings sits a qualitative risk. Recent conflicts have underscored that data centres are now critical information infrastructure as well as commercial infrastructure — and that centralised compute is vulnerable to disruption. As compute concentrates in a small number of large terrestrial sites, scale itself becomes a systemic risk, creating a growing requirement for redundancy and geographic distribution in sovereign digital systems.

These constraints do not represent an immediate ceiling, but they do make future compute growth increasingly dependent on scarce terrestrial resources. For nations pursuing AI sovereignty, the challenge is no longer whether compute demand will grow, but how it can be sustained securely, economically, and at scale. It is against this backdrop that India's position becomes relevant.

1.3 How does India see AI compute?

Across emerging technologies that remain capital-intensive, infrastructure-heavy, or still evolving in commercial viability, early-stage momentum is typically shaped by significant government participation and policy support.

India has not been a bystander to this shift. Through the IndiaAI Mission, the Government of India has committed ₹10,371.92 crore to building domestic AI capability and technological sovereignty, and a central pillar of that effort — the IndiaAI Compute Pillar — has already onboarded more than 58,000 GPUs⁷, providing affordable, subsidised access for startups, researchers, academia, and industry⁸. The strategic logic is explicit in national policy: compute is treated as sovereign infrastructure, on par with energy and connectivity, because the capacity to train and run AI domestically determines whether a nation's data, models, and economic value remain under its own control.

India has already signalled that compute is a strategic national priority and has backed it with public capital. The open question is whether that is enough to compete with demand that is compounding globally.

The answer is increasingly emerging within India's own strategic neighbourhood that the frontier of the contest has already moved off the ground. **The next, must be Orbital Infrastructure.**

Orbital Chenguang, a Beijing-based startup, has secured USD 8.45 billion in state-backed financing to build storage and compute in space⁹. The race to extend sovereign compute into orbit is no longer hypothetical — it is funded and underway elsewhere.

⁷ <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2229397®=48&lang=2>

⁸ <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2225781®=3&lang=2>

⁹ <https://spacenews.com/china-backs-orbital-data-center-startup-with-8-4-billion-in-credit-lines/>

Part II — The Orbital Compute Paradigm: Race for Orbital Data Centres

BIS Research projects the in-orbit data centre market at USD 1.77 billion by 2029, growing at a compound annual growth rate of 67.4 percent to USD 39.09 billion by 2035. Against a global space industry projected to reach USD 1.8 trillion by 2035¹⁰, that positions orbital compute as one of the fastest-growing subsectors in the economy.

2.1 The paradigm shift: From satellites that sense to infrastructure that computes

Understanding this shift requires moving from broad strategic signals to the first practical implementations. The most important development at this stage alongside patient investments in large-scale orbital data centres, is the emergence of entry-level systems that can demonstrate real workloads in space, validate unit economics, and establish operational credibility. It is this “entry wedge” — small, functional compute systems deployed in orbit — that defines the starting point of the orbital compute paradigm.

2.1.1 Entry wedge: The connecting link

What makes this transition materially different from earlier space infrastructure cycles is not simply the ability to place compute in orbit, but the fact that the compute substrate itself is no longer changing.

What if, unlike previous generations of space systems that required bespoke avionics stacks and aerospace-specific software environments, the emerging orbital compute layer is increasingly built on the same underlying hardware architecture that already dominates terrestrial AI — GPU-accelerated, CUDA-compatible systems deployed in containerised workflows.

Because, this continuity matters more than the deployment location itself. It implies that orbital infrastructure is not introducing a new compute paradigm but extending to an existing one. In practical terms, models trained in terrestrial environments, that do not need to be re-engineered for space; can be ported with minimal modification, executed on-orbit, and reintegrated into standard AI pipelines on Earth. The implication is a reduction in both cognitive and technical friction, a prerequisite for any infrastructure layer that intends to scale beyond niche scientific use.

It is in this context that early “entry wedge” systems become important — not as finished architectures, but as validation points for whether orbital compute can behave like a standard extension of cloud-native infrastructure.

One such implementation is being developed in India by **TakeMe2Space (TM2S)**, with its first MOI-1A satellite scheduled for launch in October 2026. The programme advances from single-satellite technology demonstration towards operational orbital compute infrastructure.

¹⁰ <https://www.weforum.org/publications/space-the-1-8-trillion-opportunity-for-global-economic-growth/>

Further, through subsequent expansions into a constellation, interconnected through optical inter-satellite links, the company aims of reaching a capacity of 50kW orbital datacentre. The system is not positioned as a standalone data product, but as a deployable compute infrastructure in orbit, designed to physically twin terrestrial AI workflows without modification in space conditions.

Rather than introducing a new software paradigm, the platform is structured around four functional primitives that mirror terrestrial cloud infrastructure:

- **On-orbit inference and Earth observation analytics**, where workloads are executed directly on satellite hardware and only structured outputs are transmitted to Earth (OrbitLab)
- **Disaster-resilient archival for high-value datasets**, designed for long-horizon storage use cases where orbital persistence provides an alternative to terrestrial data centre concentration risk (OrbitVault)
- **Deployer Infrastructure**, by simply enabling operators to run their own AI models directly in orbit, processing data at the moment it is captured and on their own schedule
- **Regulated-sector sandbox deployments**, where controlled pilots — for instance in banking or insurance — can be executed within a governed orbital environment to evaluate data security, latency, and compliance implications

The significance of this class of systems lies less in their scale and more in their compatibility. If orbital compute uses the same GPU architectures and developer toolchains as terrestrial AI systems, then the transition from ground to orbit becomes incremental rather than discontinuous. The bottleneck shifts from **“can this be built?”** to **“how quickly can workloads migrate?”**

This is further reinforced by parallel activity within India’s broader space-tech ecosystem, where multiple private actors across different layers of the stack have begun exploring orbital data centre and in-space compute architectures that implicitly assume computation will increasingly move closer to orbit rather than remain entirely Earth-bound.

Alongside startups like Agnikul Cosmos, including its collaboration with NeevCloud on orbital AI infrastructure and firms such as Pixxel in partnership with Sarvam AI have also begun examining compute-enabled orbital architectures¹¹, collectively extending the conversation beyond sensing and connectivity into storage, AI execution, and distributed orbital compute.

The direction of travel is therefore not isolated to a single organisation, but indicative of an


¹¹ <https://economictimes.indiatimes.com/tech/startups/space-company-pixxel-and-ai-firm-sarvam-join-hands-to-build-orbital-data-centre-satellite/articleshow/130795476.cms?from=mdr>

emerging design space being tested across launch, compute, storage, and orbital operations. This trajectory may be further strengthened as launch economics continue to evolve, given its direct impact on costing. Emerging Indian launch companies including Skyroot Aerospace, through composite-intensive vehicle design and manufacturing innovation, Agnikul Cosmos through highly configurable launch systems enabled by extensive additive manufacturing, and EthereumX through ambitions for reusable heavy-lift architectures, are collectively pursuing pathways that could reduce the long-term marginal cost of access to orbit. While launch remains a meaningful constraint today, sustained reductions in per-kg launch costs would further improve the economic case for persistent orbital compute and data infrastructure.

2.2 The convergence thesis

Orbital datacentres are currently constrained by launch mass economics, radiative heat rejection limits, radiation-induced hardware degradation, and orbit-dependent connectivity; however, as early demonstrations emerge, the field is entering a narrowing window in which experimental deployments can still shape foundational architectural standards before orbital compute systems transition into standardised, large-scale infrastructure buildouts.

The economic heart of the case is a crossover. Terrestrial data-centre costs are rising, by an estimated 7%–15% per year, as power, water, land, and permitting all tighten under the demand growth¹². Orbital compute costs, by contrast, will fall, driven principally by the collapse in launch cost per kilogram (rideshare launch has reduced the dominant historical cost line by more than an order of magnitude over the past fifteen years) and by improving power and compute density per satellite.



Two cost curves moving in opposite directions intersect. Based on TakeMe2Space’s modelling, a potential economic “sweet spot” begins to emerge around 2035–2036, the period at which the cost of a unit of orbital compute approaches parity with an equivalent unit of terrestrial compute.

¹² <https://www.jll.com/en-us/insights/market-outlook/data-center-outlook>

Under the model's assumptions, this convergence occurs at approximately \$7,000 per kW per month, reflecting not a sudden decline in orbital costs alone but the interaction of two opposing trends: increasing terrestrial compute costs and gradual reductions in the cost of operating compute infrastructure in orbit.

Before that point, orbital compute is a premium product justified by its non-cost properties (sovereignty, continuity, jurisdictional neutrality, latency to space-origin data). After it, orbital compute competes on price for an expanding share of workloads.

The convergence thesis reframes the entire proposition. The question is not "is orbital compute cheaper than the ground today?" — it is not — but "when does it become cheaper, and what is the value of having built the capability, the standards, and the operational heritage before that date arrives?"

2.3 The constraints

Orbital compute faces four principal challenges: thermal management, radiation exposure, bandwidth limitations, and launch economics. In orbit, waste heat cannot be dissipated through air or water cooling and must instead be radiated into space, making thermal rejection the primary engineering constraint on compute density. Electronics must also operate reliably in a radiation environment that can degrade components and disrupt operations. At the same time, data transmission from orbit remains spectrum-constrained, while deployment and replacement of hardware depend on launch availability and cost. Together, these constraints historically prevented compute infrastructure from scaling beyond terrestrial environments.

The key difference today is that most of these barriers have become increasingly manageable. Launch costs show a decreasing curve through reusable rockets and rideshare missions, with a growing number of Indian providers entering the market. Advances in fault-tolerant architectures, shielding, and error-correcting systems enable commercial-class processors to operate reliably in orbit. On-orbit inference changes the economics of data transmission by processing information where it is generated and downlinking insights rather than raw data. Thermal management remains the defining challenge, but it is now a bounded engineering problem with established solution pathways rather than a fundamental limitation. As a result, orbital compute has shifted from a theoretical concept to a technically credible infrastructure proposition.

Part III — Legal and Governance Dimensions

3.1 International Legal Framework

The governance of outer space activities has, since the late 1960s, rested on a body of international treaty law that was designed for a fundamentally different technological context: state-operated launch vehicles, scientific missions, and communications satellites operated by national agencies or intergovernmental organisations. That framework has subsequently supported through complementary conventions, ITU regulatory instruments, and soft-law guidelines developed through the Committee on the Peaceful Uses of Outer Space (COPUOS). The emergence of orbital compute infrastructure — privately operated, commercially deployed, and functionally equivalent to a terrestrial data centre in its processing activities — raises legal questions that no existing instrument addresses directly. The analysis that follows examines the applicable international framework first, then India's domestic regulatory position.

3.1.1 The Outer Space Treaty and its successors

The Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, Including the Moon and Other Celestial Bodies (Outer Space Treaty, OST) of 1967 provides the constitutional framework for space law¹³. Its core provisions are well established. Article II prohibits national appropriation of outer space by claim of sovereignty. Article VI makes states internationally responsible for national activities in outer space, whether conducted by governmental or non-governmental entities, and requires that national activities be carried out in conformity with the Treaty¹⁴. Article VIII provides that a state retains jurisdiction and control over space objects registered to it, and ownership is unaffected by the object's presence in outer space¹⁵.

For an orbital compute operator, Article VIII is the most immediately operative provision. The launching state — for an India entity, through ISRO — retains jurisdiction over the space object and all activities performed aboard it. This provides a clear jurisdictional anchor for Indian regulatory authority over its orbital platforms. **The gap, as discussed, is that Indian domestic law has not yet translated this treaty-based jurisdictional entitlement into workable regulatory instruments for compute and data processing activities.**

¹³ Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space (Outer Space Treaty) (adopted 27 January 1967, entered into force 10 October 1967) 610 UNTS 205, art II (non-appropriation principle).

¹⁴ Outer Space Treaty (n 40001), art VI (state responsibility for national activities in outer space, including those of non-governmental entities).

¹⁵ Outer Space Treaty (n 40001), art VIII (jurisdiction and control of launching state over space objects and personnel).

The Liability Convention of 1972 establishes absolute liability for damage caused by space objects on the surface of the Earth or to aircraft in flight, and fault-based liability for damage caused in outer space¹⁶. In the context of orbital compute, liability questions arise most acutely in relation to debris-generating events. If an ODC satellite fragments – whether through collision, propulsion failure, or end-of-life disposal – and the resulting debris causes damage to another state’s space assets, the launching state bears liability under the Convention. This places **an additional argument for debris mitigation standards as a condition of orbital compute authorisation**.

The Registration Convention of 1976 requires states to register space objects in a national registry and submit relevant information to the Secretary-General of the United Nations¹⁷. India maintains a national registry and has complied with registration obligations for ISRO-launched satellites. **The Convention does not create a registry for data processed aboard registered objects** – a gap that is directly relevant to the jurisdictional analysis of orbital data processing.

3.1.2 ITU framework and spectrum governance

The International Telecommunication Union’s Radio Regulations govern the use of radio-frequency spectrum and geostationary and non-geostationary orbit slots¹⁸. Orbital compute operations require uplink capacity for workload delivery, downlink capacity for results transmission, and, in multi-satellite array configurations, inter-satellite links. Each of these requires spectrum coordination through the ITU process, which operates on timelines of years and confers priority rights on early filers. India’s Department of Telecommunications and ISRO coordinate ITU filings for domestic operators – creating a planning constraint that new entrants must account for¹⁹.

3.1.3 Space sustainability and debris governance

The COPUOS Long-Term Sustainability Guidelines, adopted in 2019, represent the most recent multilateral consensus on responsible space operations²⁰. Guideline 2.2 addresses spacecraft design and the mitigation of space debris, incorporating the inter-agency Space Debris Mitigation Guidelines and recommending post-mission disposal within 25 years of mission end for objects in low Earth orbit.

¹⁶ *Convention on International Liability for Damage Caused by Space Objects (Liability Convention)* (adopted 29 March 1972, entered into force 1 September 1972) 961 UNTS 187, arts II, III.

¹⁷ *Convention on Registration of Objects Launched into Outer Space (Registration Convention)* (adopted 12 November 1974, entered into force 15 September 1976) 1023 UNTS 15.

¹⁸ *ITU Radio Regulations (Edition of 2024)* (International Telecommunication Union, Geneva), arts 9, 22 (coordination and notification procedures for satellite networks; inter-satellite service definitions).

¹⁹ *ITU Radio Regulations (n 40013)*, art 22.2 (protection of radio astronomy and space research services); Appendix 4 (inter-satellite service coordination procedures).

²⁰ *Committee on the Peaceful Uses of Outer Space, ‘Guidelines for the Long-term Sustainability of Outer Space Activities’* (UNOOSA, 2019) UN Doc A/74/20, annex II <https://its.unoosa.org/unoosa_its/en/resources/guidelines/guidelinesabout.html> accessed 6 June 2026.

The Guidelines are voluntary and non-binding, and their implementation depends on national regulatory frameworks. The practical implication for ODC operators is that debris mitigation planning is both an ethical obligation and, increasingly, a commercial risk management requirement: a debris-generating event by an Indian-licensed operator would trigger liability under the Liability Convention and potential reputational consequences for India's broader commercial space programme.

No binding international instrument currently imposes specific obligations on the operators of orbital data processing facilities as distinct from satellite operators generally. The existing framework treats orbital compute infrastructure as a subset of satellite operations, with all attendant obligations applying without modification. This treatment is legally coherent for the moment but will require elaboration as orbital compute scales from demonstration to operational infrastructure²¹.

3.1.4 Data governance in outer space: the emerging gap

No international treaty governs the processing, storage, or transfer of data aboard orbital infrastructure. The Outer Space Treaty provides little clarity on data. The Registration Convention registers objects, not data. The ITU framework governs the electromagnetic transmission of data but not its content or processing. GDPR and comparable national data protection regimes apply extraterritorially based on the location of data subjects, not the location of processing, but their enforcement mechanisms presuppose accessible operators with identifiable assets in the regulating jurisdiction.

The structural gap is as follows. An operator incorporated in a low-regulation jurisdiction, launching from a permissive launch state, and serving customers globally from an orbital platform, can credibly argue that no single national data governance regime applies to its processing activities. The Outer Space Treaty provides jurisdictional authority to the state of registry but does not specify what that jurisdiction must require in terms of data governance. This creates the conditions for a flag-of-convenience dynamic analogous to maritime law, where registration in permissive jurisdictions is used to minimise regulatory exposure.

Resolving this gap requires either multilateral agreement – which is unlikely in the near term given the divergent interests of major spacefaring states – or unilateral extraterritorial regulation anchored to operator identity and data subject location, on the model of GDPR Article 3. India has both the constitutional authority and the policy interest to pursue the latter approach. The recommendations in Part IV address – how?

3.2 India's Domestic Legal and Regulatory Landscape

India's regulatory framework for space activities has historically been organised around ISRO, with private participation permitted on terms defined by executive policy. The Indian Space Policy 2023 and the associated Norms, Guidelines and Procedures issued by IN-SPACe in 2024 represent the most significant liberalisation of that framework to date, creating a single-window authorisation mechanism for private space activities. That framework, however, was designed with launch vehicles, satellite manufacturing, ground stations, and EO data commercialisation in

²¹ United Nations Office for Outer Space Affairs, 'Space Debris Mitigation Guidelines of the Committee on the Peaceful Uses of Outer Space' (UNOOSA, 2010) UN Doc A/62/20, annex.

mind. Orbital compute infrastructure — in which the satellite is primarily a processing node rather than a sensing or communications platform and will require new operational characteristics which need evolution alongside this existing regulatory foundation.

3.2.1 Indian Space Policy 2023 and IN-SPACe

The Indian Space Policy 2023 authorises IN-SPACe to serve as the single-window regulatory and promotional body for private space activities in India²². The Norms, Guidelines and Procedures (NGP) issued in 2024 specify the authorisation process for satellite manufacturing, launch vehicle development, ground segment operations, and the commercialisation of space-derived data²³. ODC operations — the provision of compute services from an orbital platform — do not appear in any listed category of permissible activity. The practical result is that operators seeking to deploy ODC infrastructure must characterise their platforms within existing categories (most commonly, Earth observation satellite or communications satellite).

This regulatory mismatch carries three consequences. First, operators access a licensing framework designed for a different class of activity, creating uncertainty about the scope and limits of their authorisation. Second, the data governance obligations attached to EO or communications satellite authorisations do not address the specific risks of commercial data processing in orbit. Third, orbital data centre (ODC) architectures are not yet explicitly defined as a distinct category within existing regulatory frameworks. Current IN-SPACe oversight is primarily structured around established satellite functions, while compute-centric systems introduce additional dimensions that will likely require further definitional evolution over time.

3.2.2 Data protection and cybersecurity

The Information Technology Act 2000 attaches jurisdiction to computer resources within Indian territory²⁴. Section 43A, which imposes liability for failure to protect sensitive personal data, and Section 70B, which gives CERT-In incident reporting powers²⁵, both presuppose an accessible operator with locatable assets in India. An ODC operator with no Indian ground infrastructure falls in gap. The SPDI Rules require reasonable security practices but provide no guidance on what those practices look like for hardware subject to radiation events, thermal cycling, and debris impact.

²² *Indian Space Policy 2023 (India, Department of Space)* para 3.2 (IN-SPACe as single-window regulatory body); para 6 (private sector participation categories).

²³ *Norms, Guidelines and Procedures for Implementation of Indian Space Policy 2023 (IN-SPACe, 2024)*, para 2 (authorisation framework); para 5 (permitted private space activities).

²⁴ *Information Technology Act 2000 (India)*, ss 43A (liability for failure to protect sensitive personal data), 70B (CERT-In incident reporting powers); *Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules 2011 (India)*, r 8.

²⁵ *Information Technology Act 2000 (India)*, s 70B(6), (7) (CERT-In powers to issue directions and impose criminal liability for non-compliance; respondent must be identifiable and locatable).

3.2.3 Spectrum and telecommunications

Orbital compute operations require licensed spectrum for uplink, downlink, and inter-satellite communications. In India, spectrum for satellite communications is coordinated through the Department of Telecommunications (DoT) and subject to TRAI oversight, in alignment with ITU filings made by ISRO on behalf of Indian operators. At present, there is no dedicated spectrum allocation category specific to orbital data centre (ODC) architectures. Plus, the planning constraint is lead time: ITU coordination for new spectrum assignments operates on multi-year timelines, and operators must factor this into constellation deployment planning.

Part IV — India's Strategic Position and Policy Recommendations

The question before Indian policymakers is not whether orbital compute infrastructure will become strategically significant. The trajectory of investment, the pattern of competitive activity among major spacefaring economies, and the structural logic of compute demand all point in the same direction. The question is whether India engages with that transition as an active participant — contributing technology, setting norms, and building industrial capacity — or as a late adopter that inherits the standards and dependencies others have defined.

Orbital data centres (ODCs) represent an emerging extension of digital and space infrastructure that may benefit from progressively tailored policy and statutory support as deployment scales and operational models mature. A phased governance approach may therefore be appropriate: in the near term, adapting and extending existing legal frameworks and institutional mechanisms to provide regulatory certainty for early deployment; and over the longer term, developing dedicated guidance and, where appropriate, sector-specific regulatory instruments informed by operational experience and technological evolution.

Here are few near-term recommendations:

4.1 Immediate Priority: Establish an Orbital AI Regulatory Sandbox

The immediate policy priority should be the establishment of a joint MeitY-IN-SPACe regulatory sandbox for orbital computing, with participation from CERT-In and relevant sectoral regulators where appropriate.

India already has a tested institutional precedent for this approach. RBI's Regulatory Sandbox and the subsequent Inter-Operable Regulatory Sandbox (IoRS) demonstrated that emerging technologies need not wait for comprehensive legislation; instead, regulators can supervise controlled deployments, understand real compliance questions empirically, and develop rules grounded in operational evidence rather than theoretical assumptions.

Orbital compute presents a comparable regulatory challenge. Questions relating to orbital data storage, in-orbit processing, AI inferencing, cybersecurity controls, auditability, and compliance with sector-specific obligations cannot be fully resolved ex ante.

A supervised orbital sandbox would enable controlled testing under real operational conditions and generate empirical evidence to directly inform future guidance, implementation standards, and eventual statutory development. The objective should not be regulatory relaxation—it should be regulatory learning.

4.2 Extend Existing National Frameworks

In parallel, existing national digital, AI, and cybersecurity frameworks should progressively evolve to recognise orbital infrastructure as an emerging extension of India's digital stack.

This includes consideration of space-based compute and AI systems within initiatives such as the IndiaAI Mission, and evaluation of the applicability of existing Critical Information Infrastructure protections and associated cybersecurity standards to orbital platforms supporting strategically significant sectors.

Where possible, implementation should prioritise extension of existing institutions and frameworks before creating new administrative structures.

4.3 Develop a Longer-Term Legislative Roadmap for Orbital Infrastructure

As deployment scales and operational practices mature, India should consider a longer-term legislative pathway for orbital data and compute infrastructure.

This may include clarifying the applicability of Indian data governance and cybersecurity frameworks to Indian-operated space-based infrastructure, establishing legal certainty around data processed, stored, or analysed in orbit, and developing fit-for-purpose regulatory treatment for orbital data centres and in-orbit compute architectures.

The objective is not to replicate terrestrial regulation in orbit, but to ensure continuity of governance as compute expands beyond terrestrial boundaries.

Conclusion

A phased approach—combining targeted executive action in the near term with longer-term legislative evolution—would allow India to build a secure, trusted, and globally competitive orbital infrastructure ecosystem while governance frameworks remain formative, rather than adapting after strategic and commercial standards have already consolidated elsewhere.

SIA-India

SIA-India is a non-profit organization representing the interests of the space industry in India. Our membership includes satellite operators, manufacturers, suppliers, startups, academic institutions, and law firms. The association is committed to serve as a voice for the industry and represent the industry's interests at policy making and regulatory bodies. SIA-India also offers a platform to forge international collaborations and provides an interface with the government.

TAKEME2SPACE

TakeMe2Space is building Orbital Data Centres. The company specialises in designing and developing subsystems and satellites with high performance compute. They have completed two missions on ISRO POEM successfully and will be launching MOI-1a, a 6U CubeSat with multispectral imaging and on-orbit inferencing capabilities in October, 2026. Following MOI-1a, the company will launch a 5kW orbital compute constellation and lay the foundation for a 1 GW data centre in space by 2036

<https://tm2.space>

