Compliance · Technical Paper

The Constraints We Put in Orbit, Part II: Sovereignty Doesn't Subtract in Orbit, It Stacks

Part II of III. Going to orbit doesn't subtract a regulator, it stacks several incompatible sovereigns onto one moving object, and the seams between their legal theories are where accountability dies.

Jennifer McKinney · August 15, 2026 · 14 min read

This is Part II of a three-part series on what orbital AI infrastructure actually escapes, and what it doesn't. Part I examined why heat, not bandwidth, is the real physical ceiling on orbital data centers. Part III turns to governance and trust.

Movement II: Sovereignty Doesn't Subtract in Orbit, It Stacks

The escape fantasy meets the treaty

If the physics movement was about a constraint that hides inside a law of nature, this one is about a constraint that hides inside a tangle of laws made by people. The intuition driving the orbital escape fantasy is that space is a gap in the map, lawless, so regulated or shady activity can slip through it. That intuition fails on a single 1967 treaty provision, and the way it fails is the whole point.

Article VIII of the Outer Space Treaty says the state that registers a space object retains jurisdiction and control over that object and its personnel (Outer Space Treaty). There is no territorial gap. A satellite is not a lawless flag-of-nobody; it is, legally, a floating piece of its registry state, the way a ship or an embassy is. Scholars describe the mechanism as the state's pre-existing jurisdiction being *extended* into orbit by the act of registration (Outer Space Treaty). So a US-registered orbital data center is, for regulatory purposes, sitting in the United States: HIPAA, the SEC, the FDA, the export-control regime, the IRS, all of it reaches the box. Going to orbit does not exit a state's law. It carries that law along.

And the treaty closes the obvious corporate dodge. Article VI makes states internationally responsible for *all* national activities in space, expressly including those of non-governmental entities, and requires governments to authorize and continuously supervise private operators, so that states cannot escape liability by outsourcing their space programs (Outer Space Treaty). You cannot launder accountability through a shell company in orbit any more than you can on the ground. So the headline answer to "would the same laws apply" is yes: the regulatory regime of the registry state goes with the hardware, and the responsibility of that state goes with the activity. The naive "space is an escape hatch" theory is foreclosed by treaty text that has been in force since before the Moon landing.

Why the satellite is the wrong unit of analysis

Here is where I have to turn the criticism on the legal commentary itself, including the sources that establish the clean answer I just gave, because they share a flaw worth naming plainly.

Every law-firm analysis of this question reaches for the same example: a single registered satellite, one flag state, and asks which law governs it. They reach for it because it is the *easiest* case in the entire problem space. It maps cleanly onto maritime precedent, it resolves tidily, and a tidy resolution makes a publishable client-facing post. But a real orbital data center is not a satellite. It is a constellation of physically distinct objects, possibly launched by different providers, registered to different states, flying in formation, with compute and storage spread across nodes, inter-satellite links carrying data between objects under different jurisdictions, and workloads that migrate between nodes for thermal and power load-balancing. Tenant data may be sharded across multiple physical objects in multiple registries *simultaneously*.

At that point the question "which law governs the satellite" is close to meaningless, because there is no *the* satellite, and the data is not *on* a satellite, it is smeared across a moving mesh. The maritime flag-state analogy quietly dies here. A ship is one hull with one flag. This is more like a single cargo container whose contents are continuously redistributed across a fleet sailing under a dozen flags while underway. Nobody writing these opinions wants to touch that, because it does not resolve, and an unresolvable problem does not make a tidy post. So the neatness in the existing commentary is achieved by motivated example-selection: choosing the toy that permits the clean answer and declining to mention that the real object dissolves it. That is worth flagging in this piece not as a gap in the literature but as a methodological criticism *of* it.

Where the law is genuinely breaking, in two live theories

Strip away the toy example and the real fault line appears. Article VIII attaches cleanly to a *physical object* because the treaty was written for things you can point at and register. Data in transit, beamed between nodes or down to a ground station, is a different animal, and the law

has not caught up. Two competing theories are fighting over this right now, in early 2026, and they serve opposite interests.

The first is **data follows the data subject**. This is the current default, and it's why the physics-style escape fails twice over. Data-protection law mostly does not care where the server sits. GDPR Article 3(2) is built so that where the processing physically happens is irrelevant; what triggers it is whether you target or monitor people located in the EU, and the person's nationality or residence is irrelevant so long as they are in the EU (European Data Protection Board). The European regulators explicitly extend "monitoring" to health-status tracking and to wearables and smart devices, think a connected pacemaker manufactured in the United States but implanted in a patient in France (European Data Protection Board). The published clinical-trial analogy is exact: a French sponsor receiving data processed in Bangladesh is still fully under GDPR because the processing serves the EU establishment of the controller (European Data Protection Board). Swap Bangladesh for low Earth orbit and nothing changes. Under this theory, moving the compute to orbit moves the hardware out of a territory but does not move the data out of the law, because the law attaches to the human, not the coordinates.

The second is the **digital flag state**, and I want to be careful about how I use it, because the source matters more than the substance. A February 2026 Bloomberg Law opinion argues that a US-licensed satellite should be treated as US soil for data-privacy and export-control purposes regardless of its orbital position, with the uplink and downlink as "customs gates," and that once data enters the "high seas" of orbit it is subject only to the flag state until it returns to Earth (Bloomberg Law). This analysis does not cite that argument as an authority, because as legal analysis it is not a novel one; it is a retelling of the standard maritime analogy already circulating across a dozen law-firm blogs. It appears here as a *specimen*. Notice what it's *for*: the same author is explicit that forcing compliance with a patchwork of national data laws "every 90 minutes" would "cripple the industry." (Bloomberg Law) That is an industry-favoring conclusion wearing the costume of legal analysis, and its purpose is to *collapse* the stack of sovereigns back down to one permissive flag. Hold that example; in Part III it becomes Exhibit A for how the interpretive layer produces the operators' preferred answer. For now the point is narrow: this is a deregulation proposal dressed as a clarity proposal, and it should not be mistaken for the former.

There is a third voice worth surfacing precisely because it sits outside both camps and outside industry money entirely. Payal Arora, Professor of Inclusive AI Cultures at Utrecht University and a digital anthropologist who advises UN e-government, UNESCO, and UNICEF rather than any operator, reframes the whole question along an axis the law-firm commentary never touches: equity. Her observation is that data localization policies have been a lever for domestic bargaining power on Earth, but orbital compute could render these mechanisms moot, and that if citizen-generated data is processed in orbit, sovereignty becomes ambiguous between the

country of origin, the launching state, and the operator. The consequence she names is the one nobody selling the architecture will say aloud, that moving infrastructure beyond national reach risks deepening digital dependence for much of the developing world. This is the same stacking problem seen from below: not "which sovereign wins the conflict" but "which populations lose their leverage entirely when the data leaves the jurisdiction that was protecting them." That framing deserves more weight than the law-firm version for a simple structural reason: her funding does not evaporate if her conclusion is unfavorable to operators.

So the instinct that the beam differs from the object is correct, and the punchline is that nobody has authoritatively decided which way it resolves, and the two proposed answers are aligned with opposite interests.

The multi-tenant trap, where the sovereigns stack

Now the strongest version of the problem, and the one that proves this movement's thesis. Picture the realistic data center: a country owns the physical hardware, and it hosts other countries' data. Does Article VI bind the same way? The mechanics are not symmetrical, and the asymmetry is the finding.

Article VI binds the state of the *operator*, not the state of the *data*. It makes the registry state responsible for the *activity*, the launch, the operation, the supervision, of running the facility lawfully. It does not transform foreign data into data belonging to the operator or make the state under which the operator is registered responsible for the content of the data. So if a US-registered data center hosts German citizens' medical records, the US is responsible under Article VI for operating the facility lawfully, while the EU's data-protection law independently reaches the *data itself* via the targeting rule. Two separate legal hooks land on the same physical object from different directions. The hull answers to its flag state; the data answers to its subjects' states; and those are not the same state.

The multi-tenant reality compounds this. The ISS is the closest existing analog and it already contains modules owned by five different nations, in which the laws of each nation apply, held together by a lattice of bilateral agreements (ESA: ISS Legal Framework). Digital Regulation Platform Now collapse that into a *single* facility where tenant A is EU health data, tenant B is US financial data, tenant C is a third country's government records, all on shared infrastructure owned by a fourth state. There is no clean doctrine for that. The honest current answer is that private contracts fill the void, commentators concede that in the absence of universal rules, private contracts and industry standards do the defining (European Data Protection Board). Translation: the law doesn't resolve it, so the service-level agreement does, which is a fragile place to keep social-security-grade PII.

A firmware-and-breach scenario sharpens this to a point. Imagine an exploit in the *firmware*, the responsibility of the hardware operator, hence of the flag state, that lets an attacker exfiltrate *tenant data*, which falls under the jurisdiction of the states where those data subjects reside. Now ask the questions with no settled answer. Whose breach-notification clock starts, and when, given that breach-disclosure rules vary widely even among US states (European Data Protection Board)? Which state's computer-crime law was violated, the flag state whose hardware was penetrated, or the states whose nationals' data was taken? Can a foreign regulator with an unambiguous legal right to investigate forensically image a satellite it cannot physically reach? The telemetry, command, uplink, and downlink segments are themselves exposed attack surfaces (Digital Regulation Platform), and they sit under different legal control. The applicable law is arguably *clear*, everyone's law applies to their piece. What collapses is attribution and enforcement across the seams.

Where the jurisdiction lands

A practitioner analysis from 2026 lays out the live scenario without flinching: a US-operated orbital data center processing EU data, transiting US and EU airspace every 90 minutes, with China and Russia asserting their laws apply during transit and the launch state claiming Article VIII jurisdiction, and asks flatly which privacy law governs (Jones Walker). Its own conclusion is that no existing doctrine was designed for processing that is simultaneously mobile, continuous, and non-resident, and that orbital AI shatters the assumption that decisions happen in identifiable jurisdictions under authority tied to physical location (Jones Walker). That same source estimates a roughly 18-to-24-month window before geopolitical competition makes coordination meaningfully harder.

So the second constraint did exactly what the first did. It did not vanish on the way up. Going to orbit does not subtract a regulator; it *adds* them, stacking several incompatible sovereigns onto one moving object through different legal theories, flag state for the hull, data-subject state for the data, transit states asserting opportunistically, and the seams between those theories are where breaches hide and enforcement dies. The constraint relocated from "one clear jurisdiction" to "a tangle no one can enforce across," and the fix the industry prefers is to make the tangle disappear by shrinking everyone's protection down to the most permissive flag.

I want to anchor that thesis in voices that are not selling anything, because it is the kind of claim that is easy to dismiss as rhetoric when it comes from the person making the argument. A legal analysis published by the International Bar Association in July 2026 reaches the same conclusion in practitioner terms: orbital data centers "operate in the physical vacuum of outer space, but not in a regulatory vacuum," and operators face cumulative obligations under GDPR, NIS2, the Outer Space Treaty, and the Registration Convention simultaneously, meaning the regulatory stack does not simplify in orbit, it grows (International Bar Association). And an

international-space-law treatment in the energy-sector press, again not a vendor, reaches the same antiquation conclusion the physics movement reached, that current filings and investments place real pressure on a space-law framework that was built for exploration rather than commerce, and that the non-appropriation principle prohibits sovereignty without preventing de facto exclusion through sheer volume and regulatory recognition (POWER Magazine). Three independent vantages, the anthropologist, the international legal practitioners, and the space-law commentator, converging on relocation-not-escape is worth more than any single authority, because convergence from unaligned sources is the closest a question this new can get to corroboration.

One caution belongs on the record here. There is essentially no settled case law yet; this movement synthesizes treaty text and legal commentary, and that absence of precedent is itself the finding. Any firm legal *claim* in this piece should be read as contested and, ideally, checked against someone who practices space or data-protection law. And the "skirting regulations" framing should stay structural, the architecture permits regime-shopping, rather than attaching intent to any named operator, both because it's more defensible and because naming names turns analysis into a defamation problem.

Continue to Part III: The Governance Gap Is Filled by Discretion. Or go back to Part I: The Heat Problem.

Sources

- Outer Space Treaty
- European Data Protection Board
- Bloomberg Law
- ESA: ISS Legal Framework
- Digital Regulation Platform
- Jones Walker
- International Bar Association
- POWER Magazine