Cislunar Crossroads: Why a Domain That Transports Its Own Disturbances Needs Polycentric Information Sharing

A Physics-Grounded Case for a Minimum Viable Information Floor in Cislunar Space Traffic Management

Dr. Nathaniel Dailey, FRAeS
2026-06-10

Abstract

A physics-grounded analysis shows that the Earth-Moon domain transports its own disturbances rather than containing them, which converts a debris question into a governance question. In the circular restricted three-body problem, the invariant manifolds of libration-point orbits spread perturbations across the whole domain, whereas an Earth radiation or debris belt confines them. A Phase 0 computation validated the mechanism to machine precision: an Earth-Moon L2 Lyapunov orbit amplifies a perturbation roughly 1363-fold per revolution [1], and of 150 fragments dispersed near L2, 45 escape the domain entirely while others reach lunar impact or Earth vicinity within a few orbital periods [1]. Because the domain transports rather than contains, and atmospheric drag clears debris only inside Earth's atmosphere but at no lunar altitude, the open question becomes what information federated participants share, and with whom. The answer is polycentric in Elinor Ostrom's sense [2], [3]: overlapping authoritative centers coordinating through agreed standards, anchored by a Minimum Viable Information floor of six items plus a graduated complete-set delta. A two-stage instrument, a Kobe-chart equilibrium metric paired with a toggle-and-regression model, ranks which items most improve space traffic management managerial capacity, and a full-factorial regression over fourteen items places all six floor items above the eight complete-set additions. At prototype stage these coefficients are conditional on a simulation and elicitation data-generating process, a designed instrument rather than measured outcomes from a standing federation. The paper maps a three-venue path toward AMOS 2026.

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY AND THE COLLABORATION ASK

Cislunar space, the region between geosynchronous orbit and the Moon, is filling faster than the governance rules written to keep it safe. Operational relay assets already sit there, alongside lunar landers, and a lengthening queue of planned national and commercial missions. Every tool that keeps Earth orbit tractable was built for Earth orbit and belongs to it: object catalogs, conjunction screening on a familiar cadence, disposal rules that lean on atmospheric drag. Carry those tools out to the Moon and they meet a domain that differs in kind, not degree. The difference starts with the physics.

An international team from the International Astronautical Federation's Technical Committee on Space Traffic Management, including the author, framed this question in *Moon to Mars: Challenges and Strategic Frameworks for Space Traffic Management in Cislunar and Cismartian Environments* [4], the direct predecessor to the present paper. The exposure is measurable, not hypothetical. In 2025 the Aerospace Corporation showed that a single high-energy breakup in low lunar orbit can persist for years and raise the daily collision risk for a nearby asset by one to two orders of magnitude [5]. Institutions are catching up rather than leading. The 2019 Orbital Debris Mitigation Standard Practices apply to Earth orbits (low/medium/geosynchronous/highly elliptical) only, and hold cislunar provisions for a later supplement [6]. The January 2025 revision of the international debris-mitigation guidelines kept its single sentence on cislunar space and added a full new section on constellations [7].

Four beats of physics make the stakes visible to a general reader. Earth orbit is a stable ellipse, predictable and self-cleaning. Move to the Moon and the motion reframes into a rotating frame, five libration points, a sharp stability divide. Holding the useful collinear-point orbits means a continuous fight against instability, which is why NASA put Gateway in the near-rectilinear halo orbit, the near-stable high ground. Once station-keeping goes, a breakup near L1 or L2 does not stay where it began; the invariant manifolds carry its fragments across the whole domain. That physics earns the governance argument.

That argument is polycentric, in Elinor Ostrom's sense [2], [3], and it stands as an explicit alternative to the tragedy-of-the-commons framing. A single central catalog for cislunar space-object data faces political and architectural headwinds, documented later in the section titled "Why Physical Sensitivity Drives a Governance Question." The workable answer is a federation of equals. Authoritative centers overlap, each holding its own data and coordinating through agreed standards, anchored by a Minimum Viable Information (MVI) participation floor with graduated sharing above it. It should be noted that the Department of War and Intelligence Community's MVI participation may not be possible or may be extremely limited due to sensitivities.

Cislunar space is a shared inheritance rather than a frontier to be claimed, and its stewardship is a charge no single nation or operator can carry alone.

This paper is the alignment artifact for a MITRE-led effort toward the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference in September 2026. Three venues carry the path forward: the International Astronautical Federation Astrodynamics Technical Committee (TC-26), to ratify the MVI as consensus; the IAA Space Traffic Management Committee, to carry the polycentric governance argument; and AMOS 2026, for the booth demonstrator and the conditional ranking of information-item value. The spine of that path is as follows. Orbits that carry the physics converge to machine precision, the equilibrium metric draws its clearing rate from those same orbits, one end-to-end regression pass ranks the information items, and the demonstrator runs them together in the engine. What this paper enables is one focused decision: align on the MVI floor and the two-stage measurement instrument as the spine that ranks the value of each information item, so the centers already cooperating in this domain have a common floor on which to build.

2. THE PHYSICAL STAKES: FOUR BEATS OF CISLUNAR DYNAMICS

Four beats carry the physics, each grounded in a Phase 0 result and the foundational dynamical-systems literature. These distinctions appear within the relevant beats: several cislunar orbits are stable, the collinear family is the hard case, and the point-mass model has a defined scope. The beats build to one conclusion: a domain that transports its own disturbances cannot be governed with instruments designed for a domain that contains them.

2.1 Framing: why two-body intuition diverges at the Moon

Most readers already hold a reliable mental model of Earth orbit. A satellite traces a closed ellipse around one dominant body, the orbit repeats, and the mathematics reduces to the two-body problem Kepler and Newton settled. At the Moon that intuition breaks, because two bodies of comparable gravitational influence pull on a spacecraft at once. The governing model becomes the circular restricted three-body problem (CR3BP): the spacecraft moves in a frame rotating with the Earth-Moon line [8]. There, a single conserved quantity, the Jacobi integral, replaces familiar Keplerian energy, and surfaces of zero velocity partition the regions a spacecraft can reach.

The local structure of that flow is well understood. Charles Conley showed that the flow near the collinear equilibrium points sorts into transit, asymptotic, and non-transit classes [9]. Koon, Lo, Marsden, and Ross globalized that local picture into the tube-like invariant-manifold framework that organizes the rest of the paper [10]. High-fidelity lunar and planetary ephemerides ground the idealized model against reality [11]. The Phase 0 computation fixes the Earth-Moon mass parameter at $\mu = 0.012150585$, conserves the Jacobi integral, and reproduces the zero-velocity surfaces [1]. At the Moon the closed ellipse gives way, and a different kind of mathematics sets where a spacecraft can reach.

2.2 Beat 1: stable and predictable, the closed ellipse

Earth orbit is the baseline the reader already trusts. There a satellite follows a closed two-body ellipse, its motion stays predictable far into the future, and the regime partly cleans itself: residual atmosphere drags low-altitude debris down until it reenters and burns. Governance instruments were built for exactly this regime. Catalogs hold because objects stay roughly where the catalog predicts. Conjunction screening works because the geometry repeats. Disposal rules lean on drag and on the natural dispersal an oblate Earth induces.

The two-body ellipse is the limiting case of the richer three-body dynamics [8], and its convenient features, drag-driven cleanup and predictable repetition, belong to Earth orbit alone. Drag needs an atmosphere to act in, and oblateness-driven dispersal is an Earth-orbit feature [5]. The comfort of Beat 1 is what the next three beats replace.

2.3 Beat 2: reframe at the Moon, the rotating frame and the five libration points

The frame that rotates with the Earth-Moon line reveals five libration (Lagrange) points, where gravitational and centrifugal forces balance. L1, L2, and L3 lie on the Earth-Moon line; L4 and L5 occupy the vertices of equilateral triangles with Earth and Moon. Phase 0 places the collinear points at normalized coordinates L1 = 0.83692, L2 = 1.15568, L3 = -1.00506, with L4 and L5 at $(0.5 - \mu,\ \pm\sqrt{3}/2)$ [1].

These points behave in two distinct ways. For a mass ratio as small as the Earth-Moon system's, the triangular points L4 and L5 are linearly stable, so a small nudge leaves a spacecraft oscillating in a bounded path. The collinear points L1, L2, and L3 are saddle-type and linearly unstable, so a small nudge grows without bound. Szebehely derives the analytic threshold: L4 and L5 hold their stability when the mass ratio sits below a Routh value of roughly 0.0385, well above the Earth-Moon value [8]. Phase 0 makes the contrast concrete. A perturbation near L4 stays bounded, with a maximum radial offset of 0.043; the same perturbation near L2 departs, reaching a maximum displacement of 1.47 in identical normalized units [1]. This dichotomy is the hinge the whole debris argument later turns on, because stability for a spacecraft is also stability for a fragment.

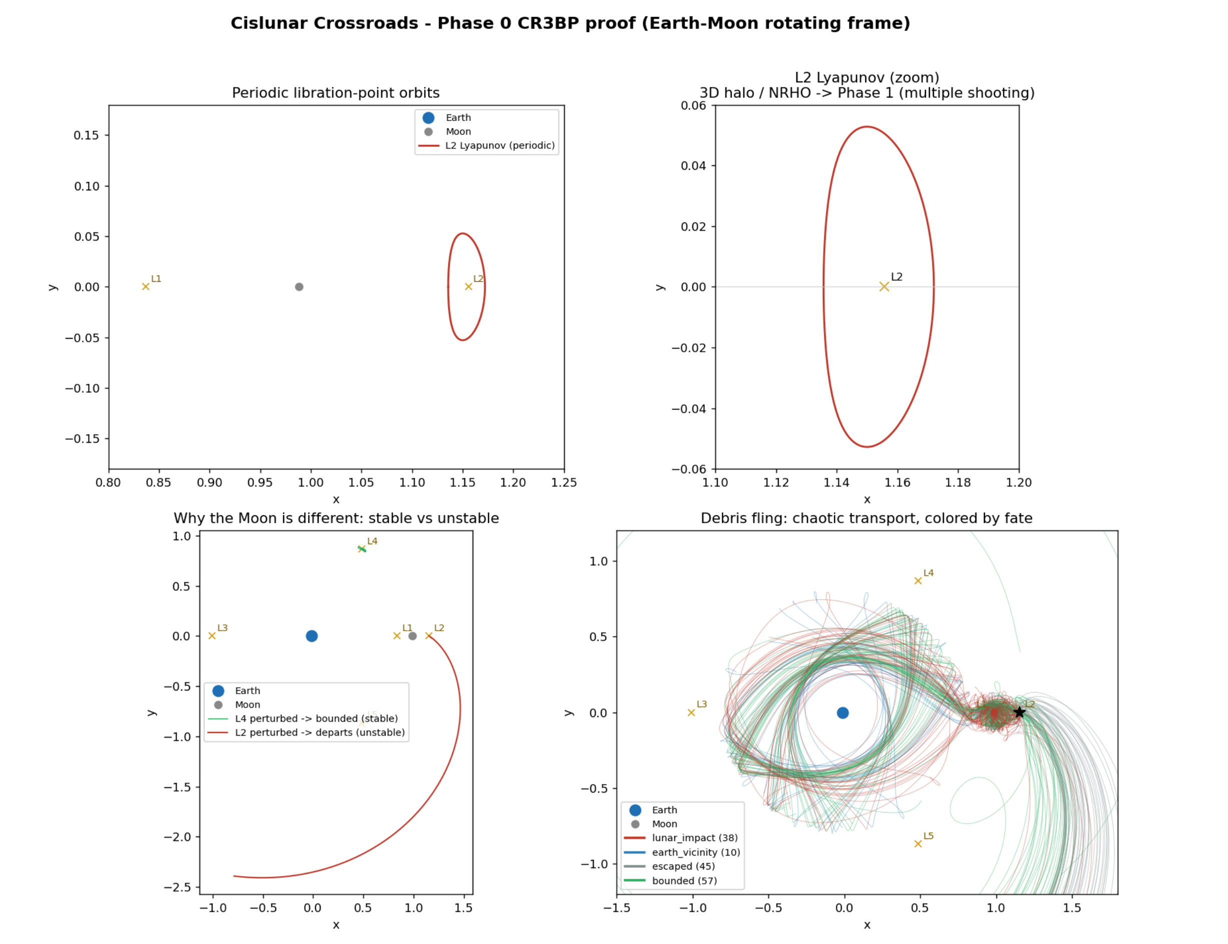

Figure 1: Cislunar dynamics overview. Libration-point locations, representative orbit families, the L4-versus-L2 stability contrast, and the fate distribution of fragments dispersed from the vicinity of L2. The stability contrast panel

maps to Beat 2; the orbit-family panel maps to Beat 3; the debris-fate panel maps to Beat 4. Computed in Phase 0 at $\mu = 0.012150585$ [1].

2.4 Beat 3: fight to stay, the unstable L2 family and the near-stable NRHO superhighway

Occupying the useful collinear-point orbits is a continuous fight against instability. The clearest quantification is the monodromy matrix, which captures how a small displacement grows over one full revolution of a periodic orbit. The Phase 0 L2 Lyapunov orbit uses initial conditions $x_0 = 1.135682$ and $v_{y0} = 0.102478$, with period $T = 3.38752$, and it closes to machine precision at a residual velocity of $|v_{x,f}| = 6.09 \times 10^{-13}$ [1]. The maximum eigenvalue of its monodromy matrix is approximately 1362.5, paired with its reciprocal of 7.34×10^{-4} [1]. A perturbation is therefore amplified roughly 1363-fold every revolution. Call this the fight-to-stay statistic. The theory of the monodromy matrix as the revolution map of the state-transition matrix is standard [10], and the numerical result rests on a high-accuracy DOP853 integrator at relative tolerance 10^{-12} and absolute tolerance 10^{-12} [12].

One scope note belongs where the 1362.5 figure is introduced: it is for the planar L2 Lyapunov family. The same instability that makes the fight-to-stay point vivid also pushes naive single-shooting continuation past its range once the correction targets a three-dimensional halo. That methodological limitation is now resolved. A multiple-shooting corrector converges the three-dimensional families directly, using eight patch points, $6N + 1$ unknowns, a state-transition-matrix-block Jacobian, and a least-squares Newton update. Seeds come from the JPL Three-Body Periodic Orbit Catalog, re-converged in the project's own dynamics, and the catalog Jacobi convention check closes to 5.3×10^{-15} [13]. Three converged orbits, given in Table 1, anchor the cislunar geometry the paper invokes.

Table 1: Three cislunar orbits converged with the multiple-shooting corrector, re-converged from JPL catalog seeds in the project dynamics. The stability index spans the dichotomy: the L2 halo is unstable, the 9:2 NRHO is near-stable by design, and the DRO is stable. All three close to better than 10^{-10} at a DOP853 tolerance of 10^{-12} [13].

Orbit	Period (TU / days)	Jacobi C	Stability index ν	Closure	Note
L2 halo	2.562699 / 11.13	3.016810	4.061 (unstable)	1.2×10^{-11}	The fight: the unstable collinear high-amplitude case
9:2 NRHO	1.513204 / 6.570	3.046343	1.327 (near-stable)	5.3×10^{-12}	NASA-selected high ground; perilune 3274 km, 0.124% from the 9:2 resonance target
DRO	5.893106 / 25.59	2.738211	1.000 (stable)	2.9×10^{-11}	Stable retrograde baseline

The unstable L2 halo carries a maximum monodromy eigenvalue near 8 and a stability index of 4.061, confirming the two-real-inverse-eigenvalue structure that makes it the three-dimensional counterpart of the planar fight-to-stay statistic. The 9:2 NRHO carries a stability index of 1.327, the near-stable high ground NASA selected for Gateway, with a perilune of 3274 km inside the published Gateway-compatible envelope. With three-dimensional family convergence established by multiple shooting, three scope boundaries remain: the ephemeris-fidelity transition out of the point-mass CR3BP, the mascon-driven low-lunar-orbit decay the point-mass model does not capture, and the conditionality of the regression on its data-generating process. Those three stand, carried in the Limitations section.

A perturbation that grows 1363-fold per revolution still yields to active control, and the Phase 0 station-keeping experiment shows the price. Left uncontrolled, the position error of a spacecraft near L2 climbs to 5.09 normalized units and flings it out of the vicinity within one to two revolutions. Under control, that error holds below 10^{-4} for a normalized delta-v expenditure of 0.0008 [1]. NASA's Gateway program once lived this trade in practice, publishing station-keeping budgets and continuous multi-year reference trajectories for its selected orbit [14], [15], [16].

That selected orbit is the near-rectilinear halo orbit (NRHO), the cislunar high ground. NASA chose it for Gateway because it occupies a dynamically privileged niche: stable enough for long-term operations, synchronized through a 9:2 resonance with the lunar synodic period to hold continuous sunlight, and open to near-continuous communications with Earth [14]. A four-body analysis confirms the engineering basis of that resonance [17], and the trade study that weighed the NRHO against the distant retrograde orbit (DRO) and other candidates selected it on operational grounds [18], [19].

The full stability picture belongs here: several cislunar orbits are stable. The DRO and the triangular points L4 and L5 are stable outright, and the NRHO is near-stable by design. The L1 and L2 collinear family is the hard case. Two facts matter equally here: the cost of the unstable regimes, and the operational value of the stable niches.

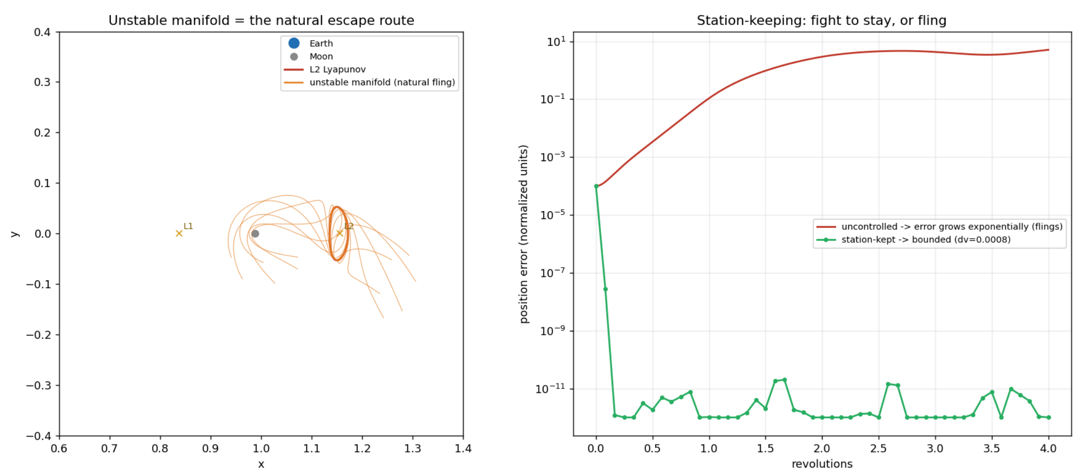

Figure 2: **Manifold transport and station-keeping**. Left, the unstable invariant manifold of an L2 orbit traces the natural escape route a departing object follows. Right, the fling-versus-hold contrast: an uncontrolled spacecraft departs within one to two revolutions, while active station-keeping holds position for a small delta-v cost. Computed in Phase 0 [1].

Table **2** lays out the stability picture across all four beats, regime by regime.

Table 2: **Cislunar regime stability summary**. The stable niches and the unstable collinear family are the two ends of the dichotomy that drives the debris argument.

Regime	Stability character	Governing basis
L4 / L5 triangular points	Linearly stable	Routh threshold; Phase 0 bounded perturbation (max offset 0.043) [1], [8]
Distant retrograde orbit	Stable	Trade-study baseline [18]
Near-rectilinear halo orbit	Near-stable, NASA-selected	9:2 synodic resonance; Earth-comms continuity [14], [17]

L1 / L2 collinear family	Linearly unstable, saddle-type	Monodromy max eigenvalue ~1362.5; uncontrolled error grows to 5.09 [1]

2.5 Beat 4: lose, and it flings, a breakup near L1/L2 rides the manifolds across the whole domain

When station-keeping ends, or when a conjunction or breakup happens near L1 or L2, the fragments depart from where they formed. The same invariant manifolds that route low-energy spacecraft transfers also route debris, carrying fragments across the entire cislunar domain: past the L1 and L2 transit regions, to lunar impact, to Earth vicinity, or out to heliocentric space within a few orbital periods.

The Phase 0 fling experiment is the quantitative payoff. A Gaussian velocity dispersion applied to a common position near L2 seeded a cloud of 150 fragments, each propagated for 12 normalized time units in the rotating-frame CR3BP. Of those 150 fragments, 45 escaped the domain, 38 impacted the Moon, 10 reached Earth vicinity, and 57 remained temporarily bounded [1]. Across the run, 93 of the 150 fragments, 62.0 percent, departed the bounded region within 12 normalized time units by manifold transport. That clearing rate is a computed transport input, not a rhetorical figure, and it carries forward into the governance instrument. It is the natural-regeneration term behind the equilibrium metric presented later: in the measurement-instrument section, the domain's own dispersion sets the baseline against which fragmentation pressure is gauged. The transit timing is consistent with the manifold-tube transit times in the foundational literature [10]. The domain transports its disturbances across the whole region, where an Earth belt confines them. Table 3 gives the fate breakdown and mechanism.

Table 3: Debris fate breakdown for 150 fragments dispersed in velocity from the vicinity of L2, propagated to 12 normalized time units in the Phase 0 rotating-frame CR3BP [1].

Fate	Count (of 150)	Manifold mechanism
Escaped the domain	45	Unstable manifold tube to the exterior realm
Lunar impact	38	Transit orbit captured into the Moon realm
Earth vicinity	10	Transit through the L1 neck to the Earth realm
Temporarily bounded	57	Trapped between zero-velocity surfaces over the simulation horizon

The transport network is concrete and flight-demonstrated. Three flown missions rode these structures: Japan's Hiten reached the Moon through a weak-stability-boundary transfer in 1991 [20]; NASA's GRAIL pair followed low-energy lunar trajectories in 2011 [21]; and ARTEMIS transferred from Earth-Moon L2 halo orbits to lunar orbit across 2010 and 2011 [22], [23]. The underlying mathematics, that the manifold tubes and the weak stability boundary describe the same geometric structure, is rigorously established [10], [24], [25], [26]. What ferries a spacecraft cheaply is exactly what ferries debris widely.

The debris paradox follows. The general cislunar volume partly clears itself: the competing pull of Earth and Moon disperses fragments on dynamical rather than slow secular timescales, sending many to escape or to lunar impact rather than into a long-lived belt. The exception is decisive. The stable niches, the NRHO, the DRO, and the triangular points, concentrate and amplify exposure, because stability for a spacecraft is stability for a fragment. The Aerospace Corporation's counterfactual makes that niche-persistence concrete: debris from a lunar-orbit breakup can persist beyond five years in a long-lived low orbit, clearing in roughly eight months from a shorter-lived one. The same conjunction can raise a nearby asset's daily collision risk by one to two orders of magnitude, with cumulative annual risk on the order of twenty times a comparable Earth-orbit breakup [5].

Direct modeling of NRHO breakup debris confirms the niche finding [27]; halo-class breakup studies establish the regime baseline [28]; and manifold-based admissible-region work shows how sharply the relevant phase space concentrates once the geometry is exploited [29]. Forensic expansion metrics for breakup clouds are advancing alongside [19], [30].

One more scope note belongs in this beat. The Phase 0 fling experiment runs on a point-mass CR3BP and supplies qualitative transport-and-fling intuition, with the persistence numbers sourced separately at higher fidelity. The multi-year low-lunar-orbit persistence figures come from the Aerospace study and its high-fidelity lunar gravity field. Low-lunar-orbit decay traces to mass concentrations in the non-spherical lunar gravity field, an effect high-fidelity gravity models capture, so this paper sources that beat from those models in place of the Phase 0 point-mass model [5]. The qualitative claim, that the domain flings across the whole region where an Earth belt confines, is what the Phase 0 numbers establish, to machine precision.

3. WHY PHYSICAL SENSITIVITY DRIVES A GOVERNANCE QUESTION

The physics of the previous section carries a direct governance consequence. Because debris disperses domain-wide along the manifolds while atmospheric drag operates at Earth, the passive mechanism that makes low-Earth-orbit debris management tractable belongs to Earth orbit alone. Three things set Earth apart: drag acts inside the atmosphere, a given delta-v produces a far larger orbit change at the Moon than at Earth, and oblateness-driven dispersal is an Earth-orbit feature [5]. Outside the stable niches the domain governs itself through physical dispersion, and in the general volume that dispersion is precisely what carries a local breakup into everyone else's operating regions.

Physics locates where the exposure sits and how long it lasts. The authority to use that knowledge, and the shared information needed to act in time, comes from governance rather than from physics [31]. Physical sensitivity and domain-wide transport together make a collective-action problem, and coordinating many independent operators is a governance task. The open question moves past whether exposure exists. It becomes what information federated participants share, and with whom, to manage a domain that transports its own disturbances.

Institutions hold this question open, and the record is current, as Table 4 sets out. The 2019 Orbital Debris Mitigation Standard Practices organize four objectives for Earth orbits, including low, medium, geosynchronous, and highly elliptical regimes, and leave cislunar space and "other orbits" for a later supplement [6]. The international debris-mitigation guidelines, revised in January 2025, tightened low-Earth-orbit disposal and added a full new section on constellations, and they carried the single-sentence cislunar provision in Section 5.3.3 forward unchanged from the prior revision [7]. That new constellation section shows the institution can act, adding regime coverage once it decides to. The cislunar provision, held to a single line, marks a decision for this region that still lies ahead. The choice is deliberate, and the capacity to close it is already on record.

Table 4: **The open cislunar provisions at a glance**. The institution demonstrably can add substantive regime coverage; the cislunar provisions stand open by choice.

Instrument	Cislunar coverage	What changed recently
ODMSP (2019 update)	Four objectives govern Earth orbits; cislunar provisions open for a later supplement [6]	Cislunar supplement pending
IADC-02-01 Rev 4 (Jan 2025)	Section 5.3.3 "Other Orbits" stands at one sentence [7]	New Section 5.5 on constellations added; LEO disposal tightened to 90% minimum / 99% goal

4. THE POLYCENTRIC ANSWER: FEDERATION, THE MVI FLOOR, AND THE COMPLETE-SET DELTA

The framing here is Elinor Ostrom's polycentric governance of common-pool resources, not a tragedy-of-the-commons framing. From there the argument runs through four moves: the federation model, the MVI floor as a schema, the complete-set delta, and the trust and provenance dimension.

4.1 Not a central catalog: the international community is converging on federation

The strongest objection to a federated design deserves its full force first. A single authoritative catalog would be simpler to query and audit, and simpler to hold accountable, than a federation of nodes whose data definitions demand continuous reconciliation. The objection holds on simplicity. The feasibility evidence points the other way.

A single United Nations-custodied central repository for cislunar space-object data faces two empirical headwinds, one political and one architectural. Politically, data sovereignty, commercial sensitivity, and geopolitics leave states and operators inclined to keep authoritative data in their own custody. A mandated central catalog would reach partial population at best, and partial population yields a thinner resource than a well-specified federation. Architecturally, the multi-body structure of the domain distributes the relevant data by nature, across many sensors, operators, and vantages, and that favors a design that meets the data where it is generated.

The international community is, of its own accord, converging on a federated answer: overlapping authoritative nodes that interoperate through agreed standards. The convergence is documented, not asserted, across independent vantages. Space situational awareness discussions in the United Nations Committee on the Peaceful Uses of Outer Space (COPUOS) Scientific and Technical Subcommittee show it directly. The United States described interoperating national and regional nodes as the preferred design over a central repository, and Peru proposed a national-node model. An industry participant offered the civil-aviation analogy, where safety emerged from interoperating national and regional systems operating alongside one another in place of a single trust authority [32]. A practitioner with sixteen years of operating experience confirmed that the Space Data Association has run a federated, machine-to-machine data exchange among commercial operators for that entire span [32].

Expert-group recommendations reflect that federated architecture [33], and thirty years of voluntary consensus across thirteen agencies in the inter-agency debris coordination process stands as an existence proof that durable, standards-bound federation holds at multi-decade horizons [34]. Federated cooperation already operates in cislunar space: the 19th Space Defense Squadron formally acknowledges NASA's Multi-Agency Collaborative Deconfliction and Conjunction Assessment Project, hosted at the Jet Propulsion Laboratory, as the service provider for cooperative cislunar conjunction assessment [35]. The architecture this paper characterizes is the one the community is already building.

4.2 Ostrom's polycentric governance as the theoretical anchor

Ostrom's polycentric governance of common-pool resources is the theoretical anchor for this federated answer. Her central finding, drawn from decades of empirical study of fisheries, irrigation systems, and forests, is that resource users frequently do build durable, self-organizing institutions that avoid both privatization and central state control. Those institutions take the form of overlapping governance units coordinating through agreed rules [2]. Her later generalization holds that polycentric systems, with many decision centers operating at different scales, can outperform both markets alone and states alone on collective-action problems, global environmental change among them [3].

The orbital commons is frequently cast as a Hardin tragedy-of-the-commons scenario, unchecked use consuming the shared resource. Polycentric governance is the documented alternative to that framing, and the framing this paper adopts. Two peer-reviewed papers make the Ostrom-to-space-debris connection directly. Morin and Richard offered the first peer-reviewed application of polycentric governance to space debris, reading the orbital commons through a polycentric frame in place of the Hardin frame [36]. Morin and Couette then assessed the viability of a polycentric system for orbital debris and found the structural characteristics of polycentricity already present in the global network of space organizations, with the norm-promoting arrangements still to be established [37].

That finding is the exact opportunity this paper's MVI floor is designed to fill. A polycentric structure becomes a governing one once paired with shared commitments, and the MVI floor specifies the minimum behavioral commitments that give the structure its operating teeth. One caveat holds throughout: Ostrom's design principles were developed for terrestrial common-pool resources, so their application to orbital and cislunar space is analogical reasoning rather than direct transfer. The Morin-Richard and Morin-Couette papers justify the analogy, which serves here as theoretical motivation in place of a direct empirical claim. Peer-reviewed work applying Ostrom's framework to cislunar space specifically remains open territory, so this paper stakes an original governance contribution rather than extending an established synthesis. Table 5 maps Ostrom's eight design principles to their cislunar SSA correlates and previews how the MVI items realize them.

Table 5: **Ostrom's eight design principles mapped to cislunar SSA correlates**. The mapping is the conceptual bridge from polycentric theory to the MVI floor.

Ostrom design principle	**Cislunar SSA correlate**
Clearly defined boundaries	Orbit-family and regime declaration (MVI item 4)
Rules suited to local conditions	Three-body-aware ephemeris and covariance standards (MVI item 2)
Collective-choice arrangements	Federated node participation via TC-26 and STM consensus
Monitoring	Independent tracking observations (complete-set delta)
Graduated sanctions	Graduated sharing tiers above the MVI floor
Conflict-resolution mechanisms	Cislunar-adapted conjunction screening and CDMs (complete-set delta)
Recognition of rights to organize	National and regional nodes retaining their own data
Nested enterprises for larger systems	Overlapping federated nodes interoperating through standards

4.3 The MVI floor: the six items every participant shares

This paper proposes a Minimum Viable Information floor: six items that set the participation requirement for any cislunar operator in the polycentric architecture. The six are chosen to be individually necessary and collectively sufficient to enable basic cooperative conjunction screening while an operator keeps full data sovereignty. The set is principled, derived from the information elements NASA's cislunar conjunction-assessment activity uses today, which runs catalog-free, screens on self-reported ephemerides, and works to a roughly two-week horizon [35]. This is the paper's answer to the Morin-Couette opportunity: the floor names the behavioral commitments that carry a polycentric structure over into a governing one. Table 6 presents the schema. It should be noted that the Department of War and Intelligence Community may not be able to share all of these fields due to mission sensitivities.

Table 6: **The Minimum Viable Information floor**. Six items constitute the participation requirement for any cislunar operator in the polycentric architecture.

#	MVI item	Why it is necessary	Format standard	Source anchor
1	Responsible-party identity and 24/7 point of contact	A named counterpart enables coordination of a maneuver or response to a warning	UN registration fields	[7], [38], [39]
2	Predicted ephemeris plus realistic covariance, three-body-aware, OEM-class, replacing TLE/SGP4	Three-body-aware sets represent cislunar motion; screening relies on realistic uncertainty	CCSDS OEM (502.0-B-3)	[19], [32], [40], [41], [42]
3	Maneuverability and maneuver intent	Declared intent defines a maneuverable asset's future state	CCSDS CDM (508.0-P-1.1)	[35], [43]
4	Orbit family and regime	Orbit family determines which manifold structures and which exposures are operationally relevant	Orbital classification framework	[27], [44], [45]
5	Physical and RF signature	Signature enables custody maintenance and characterization of uncorrelated tracks	Characterization data products	[28], [29]
6	Mission phase and disposal plan	Disposal intent is the forward-looking control on debris creation	IADC Section 5.3.3; ODMSP escape clause	[6], [7], [46]

A second caveat attaches to the schema: the six items are author-proposed. Their grounding is operational practice in cooperative cislunar screening, plus the per-item anchors in Table 6, and adoption as an international standard remains the next step. Ratifying the six-item floor as consensus is the deliverable this paper maps to TC-26.

4.4 The complete-set delta and the trust dimension

Above the floor, more information is modeled to raise managerial capacity for space traffic management. The complete set adds eight items, each independently togglable, whose marginal contribution is exactly what the Stage 2 instrument is built to measure. Table 7 lists them.

Table 7: **The complete-set delta**. The Stage 2 instrument ranks how much each of eight information items above the MVI floor contributes to managerial capacity.

#	Complete-set delta item	What it adds above the floor	Source anchor
1	Independent tracking observations	External verification of self-reported state	[41]
2	Orbit-determination metadata and provenance	Trust in the pedigree of a state estimate	[47], [48]
3	Maneuver history and propulsion model	Predictive power for future maneuvers	[49]
4	Cislunar-adapted CDMs and screening results	Shared, comparable conjunction assessment	[43]
5	Immediate breakup and anomaly reports	Time-critical hazard alerting	[5], [28]
6	UCT characterization and light curves	Resolution of uncorrelated tracks into known objects	[29]
7	Common frame, time, and ephemeris standards	Interoperability across federated nodes	[11], [40]
8	Spectrum and comms plans and disposal verification	Closure of the regulatory and end-of-life loop	[46], [50]

Trust, alongside data, is the load-bearing dimension. An information-sharing architecture that establishes the pedigree of a state estimate lets an operator commit to a maneuver under adversarial or commercial-sensitivity conditions. The architecture proposed here builds on a documented lineage adapted to current provenance primitives. Its entry point is the Shared Space Information Sharing Ecosystem, the canonical layer for operator coordination [51], extended to great-power competition contexts [52], to the requirement for trusted stakeholder coordination [53], and to the broader sustainability framework [54].

The current trust layer carries forward through two primary sources: a federated evidence architecture with a trustworthy-data engine and a trusted-supplier program [47], and a transparency-and-traceability framework that roots provenance in a recognized cybersecurity meta-framework for AI-human space missions [48]. Zero-trust treatment of dual-use sensor networks [55] and market-based compliance scoring through tradable debris credits [56] show how the trust layer composes with incentive design. The standards themselves already exist. For every information element in the relevant survey there is a corresponding established data standard, so the remaining work is adoption and assignment of responsibility; the formats are already in hand [32], [40], [43].

5. THE MEASUREMENT INSTRUMENT: EQUILIBRIUM METRIC AND THE TOGGLE-AND-REGRESSION MODEL

The measurement instrument anchors the paper. Written for reproducibility, its description applies a representation-fidelity rule: every modeled result must remain faithful to the available evidence, state uncertainty and scope limits explicitly, and avoid implying more than the underlying knowledge supports. The Abstract and the Limitations apply the same rule.

5.1 The open policy provision this instrument measures against

The instrument measures against the documented open cislunar provision. The governance-question section and Table 4 showed that the 2019 standard practices govern Earth orbits directly and leave cislunar provisions open for a later supplement [6], and the January 2025 revision of the international guidelines tightened low-Earth-orbit disposal and added a full constellations section while the cislunar provision stands at a single sentence [7]. That new constellations section shows the institution can and will add substantive regime coverage. A cislunar supplement built on the same process is therefore feasible, and the open provision reflects a choice within the institution's demonstrated capacity [35], [57]. The instrument then quantifies how much closing each piece of that information opportunity would raise managerial capacity.

5.2 Stage 1: the Kobe-chart equilibrium metric

Stage 1 carries the Kobe-chart methodology into the cislunar NRHO and halo corridor. The chart comes from fisheries stock assessment and was previously adapted to low Earth orbit in the team's prior research-design work [58]. Corridor stock status runs along the horizontal axis, $x = N/N_{cap}$, the count of resident objects relative to an estimated carrying capacity. Fragmentation pressure runs up the vertical axis, $y = F/F_{eq}$, the rate of debris-generating events relative to the corridor's natural clearing rate. The four-quadrant green, yellow, orange, and red space carries over from the SD04 construction. Figure 3 retains the Kobe chart's fisheries labels to preserve the two-threshold analogy. In the source method, "overfished" marks a stock below its biomass target, while "overfishing" marks removals above the sustainable rate. In this adaptation, the stock variable is resident-object load, so a value above one marks corridor saturation; a fragmentation-pressure value above one marks debris generation above natural clearing. The quadrants therefore distinguish stock and pressure threshold states, not literal fishing conditions.

The cislunar contribution grounds the clearing term. F_{eq} rests on the 62.0 percent invariant-manifold transport rate computed in Beat 4, the natural-regeneration mechanism the domain supplies for itself. The metric therefore weighs fragmentation pressure against a clearing rate that the Phase 0 computation provides, not against an assumed constant [1], [59]. Table 8 places the corridor on the chart across three traffic scenarios.

Table 8: **The three Kobe scenario points**. $N_{cap} = 20$ reflects MADCAP 14-day screening geometry and the NRHO 9:2 stability niche; F_{eq} is grounded in the 62.0 percent Phase 0 transport rate. Scenario stock and pressure inputs are demo assumptions with documented provenance [59].

Scenario	Year	$x = N/N_{cap}$	$y = F/F_{eq}$	Quadrant	Verdict
Sparse	2026	0.150	0.200	Green	Sustainable
Moderate	2030	0.400	0.700	Green	Sustainable
Unmanaged	2035	1.100	1.700	Red	Critical

One finding earns its own sentence. On the chart the corridor skips the yellow and orange quadrants outright, running from green through green straight to red: unmanaged growth trips the stock threshold and the pressure threshold at nearly the same instant, the nonlinear conjunction-risk accumulation Highsmith and colleagues document for the cislunar regime [5]. There is no gradual slide through an intermediate caution band. Both limits give way almost at once as unmanaged traffic reaches corridor saturation.

The equilibrium metric places one candidate equilibrium and three demo-assumption scenarios against a theory-predicted clearing rate. It is a diagnostic instrument, not a longitudinal measurement of realized outcomes from a standing federation. Its 62.0 percent transport rate is a Phase 0 physics output; the carrying capacity and the scenario inputs remain documented demo assumptions.

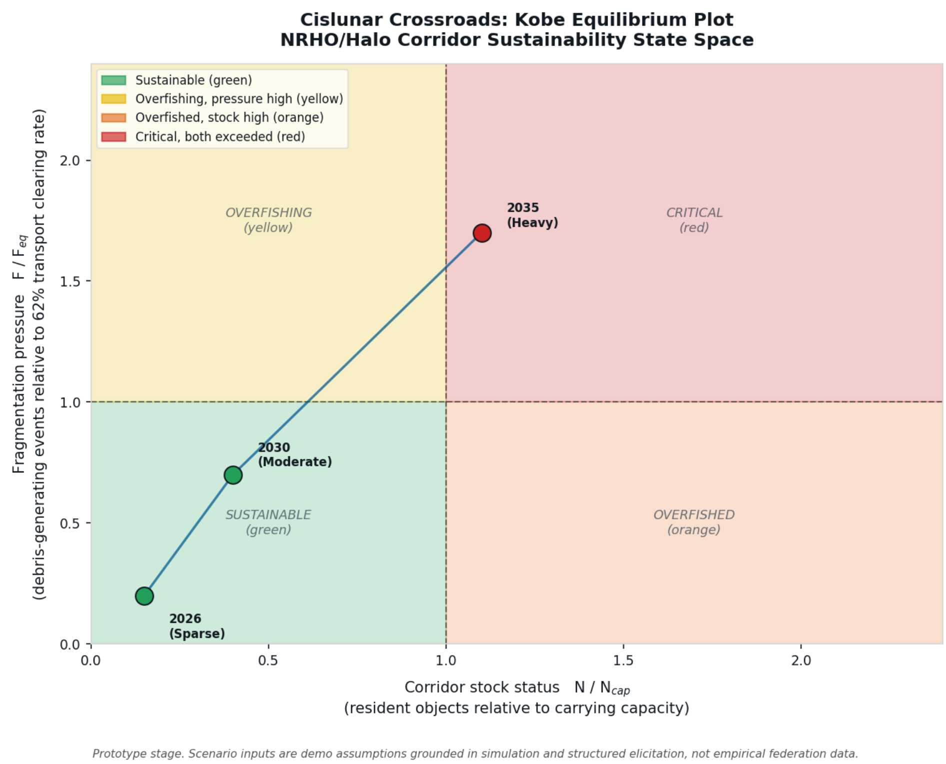

Figure 3: **Cislunar Crossroads Kobe equilibrium chart**. Corridor stock status (N/N_{cap}, horizontal) against fragmentation pressure (F/F_{eq}, vertical) for the three traffic scenarios. The clearing rate F_{eq} is grounded in the 62.0 percent invariant-manifold transport rate from Phase 0 (93 of 150 fragments departed the bounded region within 12 TU). The trajectory skips the yellow and orange quadrants, crossing both thresholds nearly simultaneously. Scenario inputs are demo assumptions with documented provenance; this is a diagnostic instrument, not a longitudinal measurement from a standing federation [59].

5.3 Stage 2: the toggle-and-regression model on STM managerial capacity

Stage 2 is a multivariate linear regression. Every MVI item and every complete-set delta item enters as an independently togglable binary predictor, with space traffic management managerial capacity as the response variable. Sweeping the toggle configurations builds the design matrix, and the data-generating process behind it is a simulation over MVI and complete-set combinations that structured elicitation from domain experts can augment.

The model reports a full statistical battery: standardized regression coefficients ranked by magnitude, 95% confidence intervals, p-values, the coefficient of determination R^2, and variance inflation factors that diagnose multicollinearity among correlated information items. Those standardized coefficients answer the AMOS-facing empirical question, namely which information items contribute most to managerial capacity per unit of disclosure. Fixing the full reporting plan in advance lets a reviewer weigh the method before any numbers exist.

The instrument now runs end to end. Ordinary least squares fits a full-factorial sweep over the fourteen toggles, $2^{14} = 16,384$ configurations drawn three times for 49{,}152 observations, and returns a ranked, separable coefficient set. That fit explains the response with $R^2 = 0.8384$, an F-statistic of 18{,}211 at $p < 10^{-300}$, and a variance inflation factor of 1.000 across all fourteen items [60].

The orthogonal full-factorial design is what drives the VIF to near unity. Under a balanced sweep the predictors are uncorrelated by construction, so the coefficients separate cleanly and the diagnostic carries no signal yet. It turns load-bearing once the synthetic sweep gives way to real participation data: operators who share ephemeris also tend to share covariance and identity, those correlations inflate VIF, and the confidence intervals on correlated items widen. Reporting it now pins the diagnostic in place for that transition. The R^2 remains below unity because the documented noise floor retains modeled variance around the deterministic signal instead of imposing a deterministic mapping.

Table 9 reports both the top of the ranking and the structural result: all six floor items outrank all eight complete-set delta items. Ephemeris plus realistic covariance leads at standardized $\beta = 0.373$, the identity-and-point-of-contact item follows at 0.352, maneuverability and intent at 0.349, orbit family at 0.292, and signature at 0.274; the delta tail runs from 0.206 down to 0.124 [60].

Table 9: **Stage 2 ranked standardized coefficients**. All six MVI floor items rank above all eight complete-set delta items. Diagnostics: $R^2 = 0.8384$, F $= 18{,}211$ ($p < 10^{-300}$), VIF $= 1.000$ (orthogonal full-factorial design), $n = 49{,}152$. Coefficients are conditional on the documented data-generating process [60].

Rank	Information item	Tier	Standardized $\boldsymbol{\beta}$	95% CI
1	F2 predicted ephemeris + realistic covariance	Floor	0.373	[0.369, 0.376]
2	F1 responsible-party identity + 24/7 POC	Floor	0.352	[0.348, 0.355]
3	F3 maneuverability + maneuver intent	Floor	0.349	[0.346, 0.353]
4	F4 orbit family / regime	Floor	0.292	[0.289, 0.296]
5	F5 physical / RF signature	Floor	0.274	[0.270, 0.278]
6	F6 mission phase + disposal plan	Floor	0.246	[0.243, 0.250]
7 to 14	All eight complete-set delta items	Delta	0.206 down to 0.124	(per item, all $p \approx 0$)

The framing rule for this result stays fixed. The ranking recovers the elicited weights, confirming that the pipeline reproduces its specified inputs. The current result establishes instrument behavior against the elicited prior. Its empirical value begins when federation data produce a ranking that differs from that prior; the difference becomes the finding. This qualification applies wherever coefficients or capacity scores appear and defines the Astrodynamics Technical Committee contribution as an instrument ready for testing against longitudinal federation data.

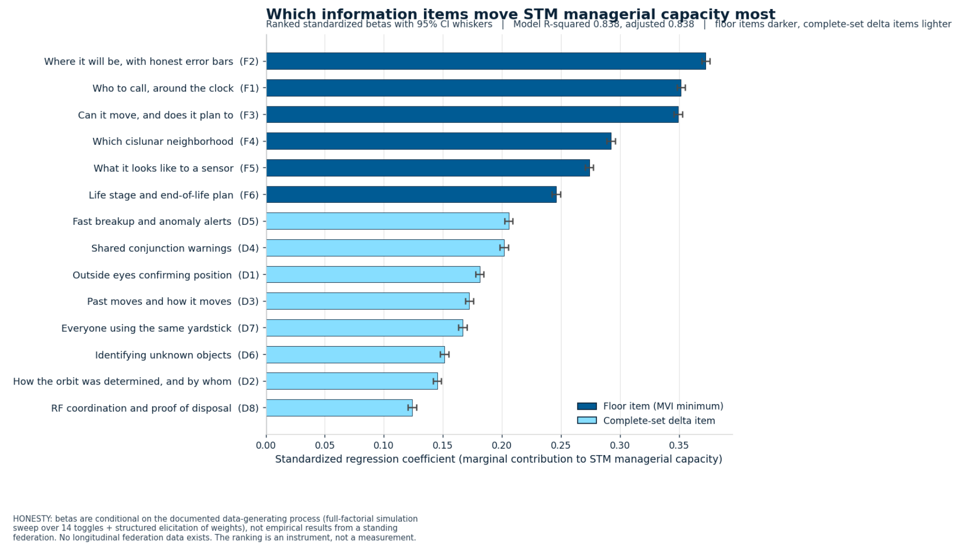

Figure 4: Stage 2 ranked standardized-coefficient (tornado) plot. Standardized regression coefficients with 95% confidence intervals for the fourteen MVI information items, estimated by OLS over a full-factorial 2^{14} design (49,152 observations, $R^2 = 0.8384$, VIF 1.000). The six floor items (darker) rank uniformly above the eight complete-set delta items (lighter), consistent with the elicitation prior. Coefficients are conditional on the documented data-generating process; the ranking is an instrument to be refined against real federation participation data, not an empirical result from a standing federation [60].

5.4 The demonstrator and the venue mapping

The instrument links to the AMOS 2026 booth demonstrator and the three-venue strategy. The booth makes the four-beat physics experience tangible and the live toggle study of Stage 2 interactive, so a visitor can switch information items on and off and watch the modeled managerial-capacity response move. Each venue carries one piece: TC-26 to ratify the MVI floor as consensus, the IAA Space Traffic Management Committee for the polycentric governance argument, and AMOS 2026 for the demonstrator and the conditional information-item ranking. Across all three venues, the trust and provenance layer connects to the MVI's trust dimension.

The demonstrator is a working engineering prototype. Unreal Engine 5.7 holds the assembled and saved stakes level, which carries Earth and the Moon, the five libration-point region markers and their labels, six hero-orbit splines including the converged 9:2 NRHO and L2 halo, an NRHO hero craft, a station-keeping craft on the halo with delta-v accounting, and the 150-fragment debris field, all driven by a double-precision C++ CR3BP module. In-engine physics validates against the Python reference. Jacobi drift holds at 1.8×10^{-15} over a full period at a substep of $1/4096$ TU, and endpoint position closes to 2.6×10^{-4} DU under display-grade fixed-step RK4 gates, a tolerance band kept deliberately distinct from the 10^{-12} science gates, with the distinction recorded in the validation log [13]. The four-beat level therefore stands physics-correct in-engine.

The booth-fidelity visual pass, the Niagara particle showpiece and the Quest 3 packaging, remains in progress and human-supervised by design. A baseline pass already stands: the bodies carry NASA-sourced surface imagery, the six hero-orbit families render in the project palette over a star field, and the libration-point regions and craft sit at their computed coordinates, so the assembled architecture reads at a glance while the particle showpiece is finished.

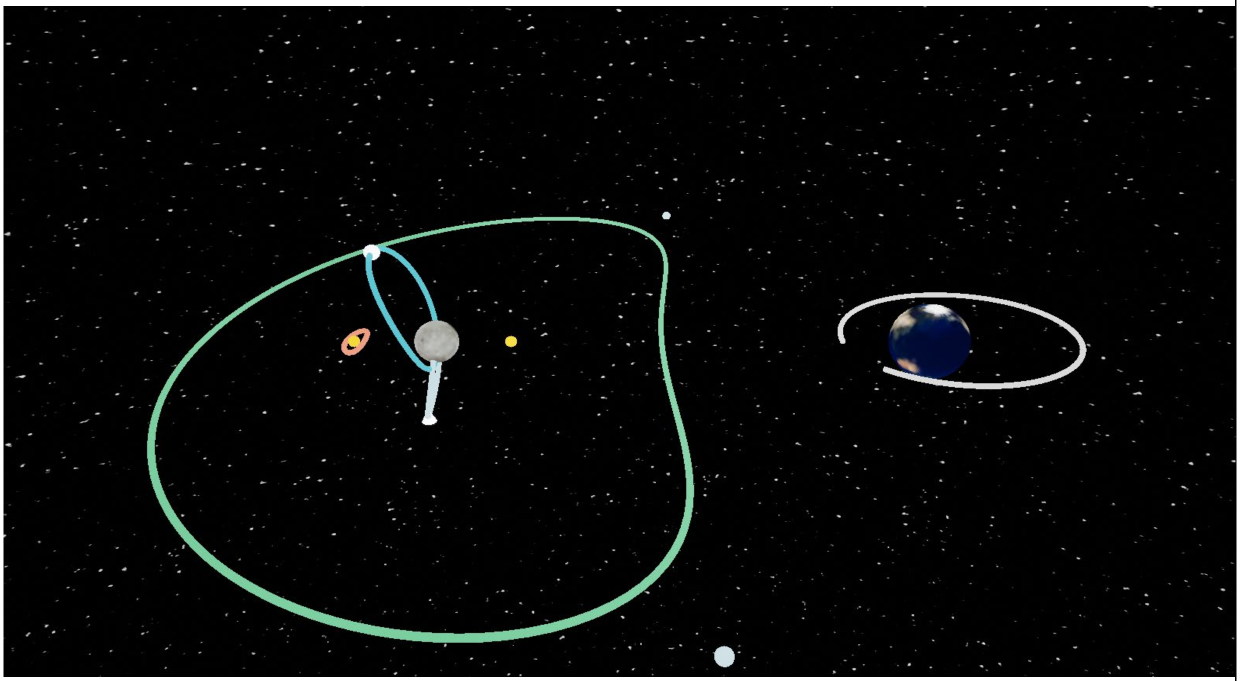

Figure 5: **The Cislunar Crossroads demonstrator running in Unreal Engine 5.7**. The converged 9:2 NRHO and the L2 halo stand at the Moon, the distant retrograde orbit encircles the system as the stable refuge, and the 150-fragment field is seeded at the L2 neck, over the Earth-Moon system with NASA surface imagery. The orbit families are drawn in the MITRE palette, and the physics is validated in the engine against the Python reference at a Jacobi drift of 1.8×10^{-15}. The baseline architecture and lighting pass are in place; the Niagara particle showpiece for the fling and the Quest 3 headset build are the named next step [13].

A physical floor sets the demonstrator's decision-velocity beat, and the federated architecture works within it. Round-trip light time to lunar distance is roughly 2.6 seconds, and human-performance data show that a control latency of that magnitude imposes on the order of a 150% task-time penalty, so the fastest cislunar hazard events resolve on board rather than from an Earth console in real time [61]. NASA has already run multi-year autonomous operations at lunar distance: the Apollo Lunar Surface Experiments Package operated autonomously for roughly eight years, the direct precedent for pre-authorized autonomous decision-making in this regime [62]. The architecture therefore sorts transactions into time-critical and non-time-critical classes, each with its own latency and trust requirements [32], and it rides a delay-and-disruption-tolerant transport layer whose foundational standard carries a MITRE engineer as co-author of record [63].

6. LIMITATIONS AND FUTURE WORK

Table 10 consolidates the scope qualifications introduced throughout the paper and pairs each with a forward action.

Three scope notes fall on the physics side. First, single shooting is no longer a limitation: a multiple-shooting corrector with eight patch points and a state-transition-matrix-block Jacobian converges the three-dimensional L2 halo, the 9:2 near-rectilinear halo, and the distant retrograde orbit directly from JPL catalog seeds, each closing to better than 10^{-10} with the catalog Jacobi convention check at 5.3×10^{-15} [10], [13]. Second, the dynamics remain a point-mass CR3BP, so the ephemeris-fidelity transition to a DE440-class force model with solar gravity and lunar oblateness remains future work, and mass-concentration-driven low-lunar-orbit decay falls outside the point-mass model; quantitative low-lunar-orbit persistence comes from high-fidelity studies that capture that effect [5], [11]. Third, a high-fidelity lunar gravity field enters the next model iteration for operational fidelity.

Two scope notes fall on the governance and instrument side. Fourth, the six-item MVI floor is author-proposed, grounded in operational cooperative-screening practice, with ratification as consensus proposed as the next step. Fifth, and most consequential, the Stage 2 regression coefficients are conditional on a simulated and elicited data-generating process, characterizing a designed instrument that real longitudinal federation data will convert into a measured result; that data is the decisive future input.

This roadmap follows from the scope notes. With the three-dimensional halo and NRHO families now converged by multiple shooting, the next steps carry those seeds into an ephemeris-fidelity force model, add a high-fidelity lunar gravity field for the low-lunar-orbit regime, and move the regression from a simulated and elicited data-generating process to real federation data as the federation stands up.

Table 10: **Limitations and future-work matrix**. Scope qualifications and corresponding forward actions consolidated from the in-section analyses. Multiple shooting resolves three-dimensional family convergence; ephemeris fidelity, mascon effects, and the data-generating process define the remaining work.

#	Scope note	Current status	Forward action
1	3D halo / NRHO / DRO generation	Resolved: multiple-shooting converged all three to better than 10^{-10} closure	Carry seeds into ephemeris-fidelity correction
2	Point-mass CR3BP; ephemeris fidelity	Qualitative transport reproduced in the point-mass baseline	DE440-class force model with solar gravity and J2
3	Mascon LLO decay out of scope of point-mass model	LLO persistence sourced from high-fidelity studies	Introduce high-fidelity lunar gravity field
4	MVI six-item floor author-proposed	Grounded in operational practice	TC-26 ratification as consensus
5	Regression conditional on simulated/elicited DGP	Instrument runs end to end; coefficients recover the prior	Transition to real federation data

7. CONCLUSION

Four physical facts carry the argument. Earth orbit is a stable, predictable, self-cleaning ellipse. At the Moon, the natural frame becomes a rotating one, with five libration points and a sharp divide between the stable niches and the unstable collinear family. Keeping a useful orbit means fighting an instability that amplifies a perturbation by roughly 1363-fold each revolution, and the near-rectilinear halo orbit is the near-stable high ground NASA chose for Gateway. Once that hold breaks, the invariant manifolds scatter fragments across the entire domain, where an Earth belt would have kept them contained. The physics makes the stakes plain. That is what earns the governance argument that follows.

The answer is a polycentric federation, in Ostrom's sense and as the documented alternative to Hardin's. A Minimum Viable Information floor anchors it, and a measurement instrument ranks which information items most improve managerial capacity. The evidentiary boundary is explicit: the instrument's current coefficients are conditional on a simulated and elicited data-generating process; longitudinal data from a functioning federation will produce the measured result.

Two objections deserve direct answers. A central catalog would be simpler. A system that states populate voluntarily is more likely to achieve broad coverage as a well-specified federation than as a mandated central repository. Shared commitments turn polycentric structure into a governing one, exactly as Morin and Couette establish: polycentricity's structural characteristics are already in place, and the norm-promoting arrangements are the next step [37]. The MVI floor meets that opportunity: it specifies the minimum behavioral commitments that convert an existing structure into a governing one.

Cislunar space will be shared across generations, and the question before this community is not whether to act together but how soon, and on what common ground. The centers that operate there already cooperate; what they lack is an agreed floor to build upon. The ask here is a measured one, in keeping with a shared endeavour: align the MITRE team and its partners on the Minimum Viable Information floor and the two-stage measurement instrument, and carry them forward along the three-venue path through TC-26, the IAA Space Traffic Management Committee, and the AMOS 2026 demonstrator. The authors note that the MVI floor may not be possible for some national assets due to sensitivities an recommend as much participation as possible in these cases while protecting national security. A floor agreed early is a generosity to every operator who follows, and the surest foundation for the cooperation, stewardship, and shared discovery this new frontier will ask of us all.

DATA AND CODE AVAILABILITY

The computational results are reproducible from the project artifacts. The validated CR3BP dynamics, the multiple-shooting orbit corrector, the converged orbit parameters, and the 150-fragment fling simulation reside in the **physics_proof/** directory, and an independent verification pass confirms all 47 physics gates, together with the in-engine Jacobi-drift validation [1], [13]. The Kobe equilibrium metric, its scenario inputs, and the toggle-and-regression instrument, including the full-factorial design matrix and the ranked-coefficient results, are stored in the **governance_study/** directory [59], [60]. The Unreal Engine 5.7 demonstrator project, with its assembled stakes level and its double-precision C++ CR3BP module, is maintained alongside as the AMOS 2026 booth artifact. Every numerical input carries a documented provenance tag distinguishing validated physics outputs from demo assumptions, and the regression coefficients are labeled throughout as conditional on the documented data-generating process.

REFERENCES

[1] N. Dailey, "Cislunar crossroads phase 0 CR3BP physics proof: Validated data and figures." Local dataset., Jun. 2026.

[2] E. Ostrom, *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press, 1990. doi: **10.1017/cbo9781316423936**.

[3] E. Ostrom, "Polycentric systems for coping with collective action and global environmental change," *Global Environmental Change*, vol. 20, no. 4, pp. 550–557, 2010, doi: **10.1016/j.gloenvcha.2010.07.004**.

[4] K. Anilkumar *et al.*, "Moon to mars: Challenges and strategic frameworks for space traffic management in cislunar and cismartian environments," *Acta Astronautica*, vol. 229, pp. 211–217, 2025, doi: **10.1016/j.actaastro.2024.12.056**.

[5] D. E. Highsmith and L. K. Newman, "Risk to lunar surface and orbital assets from a high energy breakup in low lunar orbit." Presented at the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 2025. NASA Conjunction Assessment / MADCAP public briefing., 2025.

[6] Executive Office of the President, National Space Council, "Orbital debris mitigation standard practices (november 2019 update)." White House National Space Council, 2019.

[7] Inter-Agency Space Debris Coordination Committee, "IADC space debris mitigation guidelines, revision 4," IADC, IADC-02-01 Rev 4, Jan. 2025. Available: **https://www.iadc-home.org/documents_public/view/id/318**

[8] V. Szebehely, *Theory of orbits: The restricted problem of three bodies*. New York: Academic Press, 1967.

[9] C. C. Conley, "Low energy transit orbits in the restricted three-body problem," *SIAM Journal on Applied Mathematics*, vol. 16, no. 4, pp. 732–746, 1968, doi: 10.1137/0116060.

[10] W. S. Koon, M. W. Lo, J. E. Marsden, and S. D. Ross, *Dynamical systems, the three-body problem and space mission design*. Marsden Books, 2011.

[11] R. S. Park, W. M. Folkner, J. G. Williams, and D. H. Boggs, "DE440 and DE441: Lunar and planetary ephemerides," *The Astronomical Journal*, vol. 161, no. 3, p. 105, 2021, doi: 10.3847/1538-3881/abd414.

[12] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving ordinary differential equations I: Nonstiff problems*, 2nd ed. Berlin: Springer, 1993. doi: https://doi.org/10.1007/978-3-540-78862-1.

[13] N. Dailey, "Cislunar crossroads multiple-shooting orbit convergence and in-engine validation." Local dataset. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/physics_proof/**, Jun. 2026.

[14] J. Williams, D. E. Lee, R. J. Whitley, K. A. Bokelmann, D. C. Davis, and C. F. Berry, "Targeting cislunar near rectilinear halo orbits for human space exploration," NASA Johnson Space Center, AAS 17-267; NTRS 20170001352, 2017.

[15] D. C. Davis *et al.*, "Orbit maintenance and navigation of human spacecraft at cislunar near rectilinear halo orbits," NASA Johnson Space Center, AAS 17-269; NTRS 20170001354, 2017.

[16] D. E. Lee, "White paper: Gateway destination orbit model: A continuous 15 year NRHO reference trajectory," NASA Johnson Space Center, NTRS 20190030294, 2019.

[17] K. K. Boudad, K. C. Howell, and D. C. Davis, "Dynamics of synodic resonant near rectilinear halo orbits in the bicircular four-body problem," *Advances in Space Research*, vol. 66, no. 9, pp. 2194–2214, 2020, doi: 10.1016/j.asr.2020.07.044.

[18] R. Whitley and R. Martinez, "Options for staging orbits in cislunar space," in *Proceedings of the IEEE aerospace conference*, 2016. doi: 10.1109/AERO.2016.7500635.

[19] A. De Riz, M. Grattagliano, and P. Di Lizia, "Initial orbit determination and refinement of cislunar targets with a priori dynamical constraints," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 3*, 2026.

[20] E. A. Belbruno and J. K. Miller, "Sun-perturbed earth-to-moon transfers with ballistic capture," 1990.

[21] R. B. Roncoli and K. K. Fujii, "Mission design overview for the gravity recovery and interior laboratory (GRAIL) mission," in *AIAA/AAS astrodynamics specialist conference*, 2010.

[22] D. C. Folta, M. Woodard, K. Howell, C. Patterson, and W. Schlei, "Applications of multi-body dynamical environments: The ARTEMIS transfer trajectory design," in *Acta astronautica*, 2012, pp. 237–249. doi: 10.1016/j.actaastro.2011.11.007.

[23] M.-K. Chung *et al.*, "Trans-lunar cruise trajectory design of ARTEMIS: From the earth–moon L2 halo orbit to the lunar orbit," in *AIAA/AAS astrodynamics specialist conference*, 2010.

[24] J. S. Parker and R. L. Anderson, *Low-energy lunar trajectory design*. in JPL deep space communications and navigation series, monograph 12. Pasadena, CA: Jet Propulsion Laboratory / California Institute of Technology, 2014.

[25] G. Gómez, J. Llibre, R. Martínez, and C. Simó, *Dynamics and mission design near libration points, Vols. I–IV*. Singapore: World Scientific, 2001.

[26] E. Belbruno, M. Gidea, and F. Topputo, "Weak stability boundary and invariant manifolds," *SIAM Journal on Applied Dynamical Systems*, vol. 9, no. 3, pp. 1061–1089, 2010, doi: 10.1137/100794882.

[27] N. Boone, R. A. Bettinger, and B. D. Little, "Risks from spacecraft breakup events in near rectilinear halo orbits," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2023. doi: 10.64861/JIHC4550.

[28] C. J. Wetterer *et al.*, "Cislunar debris from halo object breakups," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2023.

[29] C. Petion and B. Jones, "Manifold-based admissible regions for cislunar initial orbit determination," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2025.

[30] D. Ottoboni, A. Muciaccia, and C. Colombo, "Investigation of debris cloud expansion using metric-based analysis," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 3*, 2026.

[31] N. Dailey, "Cislunar crossroads PRISMA synthesis, track A: physics." Local document. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/paper/synthesis/track_A_physics.md**, May 2026.

[32] UN COPUOS Scientific and Technical Subcommittee SSA Expert Group, "SSA expert group discussion group transcript (otter.ai, 232 lines)." Internal working document, UN COPUOS STSC, 63rd session, Vienna, February 2026, 2026.

[33] UN COPUOS Scientific and Technical Subcommittee SSA Expert Group, "EG-SSA recommendations document." UN COPUOS STSC 63rd session, Vienna, February 2026, 2026.

[34] Inter-Agency Space Debris Coordination Committee, "IADC terms of reference, revision 12.1," IADC, IADC-93-01 ToR Rev 12.1, 2025.

[35] N. Dailey, "Cislunar crossroads PRISMA synthesis, track B: governance." Local document. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/paper/synthesis/track_B_governance.md**, May 2026.

[36] J.-F. Morin and B. Richard, "Astro-environmentalism: Towards a polycentric governance of space debris," *Global Policy*, vol. 12, no. S3, 2021, doi: 10.1111/1758-5899.12950.

[37] J.-F. Morin and C. Couette, "The missing ingredients for a polycentric governance system of orbital debris," *Global Environmental Politics*, vol. 25, 2025, doi: 10.1162/glep_a_00775.

[38] United Nations General Assembly, "Convention on registration of objects launched into outer space." UN Treaty Series, Vol. 1023, No. 15020, 1975.

[39] J. Ferrier and J. Roberts, "Measuring compliance to the UN registration convention," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 1*, 2026.

[40] Consultative Committee for Space Data Systems, "Orbit data messages, blue book," CCSDS, 502.0-B-3, 2023.

[41] J. Evans *et al.*, "Visibility-informed covariance for periodic orbits in the cislunar regime," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2025.

[42] C. Heidrich and M. Holzinger, "Cislunar initial orbit determination for maneuvering objects via ephemeris-grounded methods," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2025.

[43] Consultative Committee for Space Data Systems, "Conjunction data message, pink book (draft recommendation)," CCSDS, 508.0-P-1.1.

[44] N. Dailey, "Orbital classification framework for cislunar space domain awareness." International Astronautical Federation (IAF) TC-26 presentation, 2025.

[45] J. White, J. Bagchi, and J. Latif, "OrbX: A framework for orbital capacity characterization," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 4*, 2026.

[46] J. Wysack, "Cislunar missions end-of-life disposal strategies," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2024. doi: 10.64861/GDKE6454.

[47] J. McGillis and H. Reed, "Space trustworthy data: Imperative for a sustainable space economy," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 4*, 2026.

[48] R. Stilwell, H. Reed, and N. Dailey, "Enhancing trust in AI-human enabled space missions using transparency and traceability," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 4*, 2026.

[49] A. K. Mashiku, L. K. Newman, and D. E. Highsmith, "NASA conjunction assessment, risk analysis, and artificial intelligence/machine learning," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2025.

[50] J. Shinners and J. Roberts, "Compliance to ITU milestone approach to non-GEO assignments," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 1*, 2026.

[51] H. Reed, N. Dailey, R. Stilwell, and B. Weeden, "A shared space information sharing ecosystem (SISE) for improved SSA and space traffic management," in *AIAA ASCEND*, 2021.

[52] N. Dailey, K. Toner, and D. Murphy, "Employing a shared space information sharing ecosystem as a mechanism for promoting constructive U.S.–China space relations," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2022.

[53] N. Dailey, H. Reed, and N. Tsamis, "Space sustainability and traffic management requires trusted space stakeholder coordination," in *Advanced maui optical and space surveillance technologies (AMOS) conference*, 2023. doi: 10.64861/MVHC2799.

[54] H. Reed, N. Dailey, and R. Stilwell, "Space information sharing ecosystem: Expanding the framework for space sustainability," in *International astronautical congress (IAC)*, 2022.

[55] J. Zhang, "Securing dual-use space sensor networks for federated space traffic management," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 5*, 2026.

[56] J. Kadri and J. Shishkina, "Space debris credits: A market-based framework for governance and environmental stewardship in orbit," in *Space traffic management conference (STC26), austin, TX, 18–19 february 2026, theme 5*, 2026.

[57] N. Dailey, "Cislunar crossroads cross-cluster confirmation matrix: Opus 4.8 adversarial re-qualification." Local document. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/paper/synthesis/cross_cluster_confirmations_48.md**, May 2026.

[58] N. Dailey, R. Stilwell, *et al.*, "Kobe plot methodology adapted to orbital debris carrying capacity." IRD Project Research Designs V2.0, SD04. Internal research design document, Sep. 2024.

[59] N. Dailey, "Cislunar crossroads kobe equilibrium metric: Computed inputs and scenario points." Local dataset. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/governance_study/**, Jun. 2026.

[60] N. Dailey, "Cislunar crossroads MVI toggle-and-regression instrument: Ranked coefficients and diagnostics." Local dataset. **D:/Documents/MITRE/NASA/JPL/AMOS Booth Demo/CISLUNAR_CROSSROADS/governance_study/**, Jun. 2026.

[61] D. F. Lester, "The case for human spaceflight," Southwest Research Institute / University of Texas, 2017.

[62] NASA, "Apollo lunar surface experiments package (ALSEP): Mission description and operations." NASA Manned Spacecraft Center, 1969.

[63] V. G. Cerf *et al.*, "Delay-tolerant networking architecture," IETF, RFC 4838, Apr. 2007. Available: **https://www.rfc-editor.org/rfc/rfc4838**