PSI
Physical Superintelligence PBC

TECHNICAL REPORT PSI–TR–2026–0714 JULY 2026

WORKING DRAFT

Interstellar Precursor Mission to Alpha Centauri

Technical Feasibility Assessment

Minimum-Δv trajectories, solar-electric propulsion closure, vehicle mass budgets, program cost, and alternative targets

Prepared by	Physical Superintelligence PBC
Prepared for	The Fermi Explorer mission, a 501(c)(3) nonprofit organization
Date	July 16, 2026
Status	Working Draft
Verification	Staged independent audit, including a certified lower-bound proof with adversarial audit, a converged three-dimensional trajectory verification, and a seeded execution-dispersion analysis

PREPARED WITH PSI'S AUTONOMOUS PHYSICS-RESEARCH PLATFORM

The analyses, simulations, and proof-grade verification in this report were produced end-to-end by Physical Superintelligence's agentic research system under staged independent audit.

Abstract. This report assesses the feasibility of delivering a spacecraft with a payload of at least 1 kg to within 2,600 AU of Alpha Centauri within 500,000 years, launched by the end of 2029 on a rideshare to low Earth orbit, without gravity assists, for a program cost under \$10M. Four conclusions follow, in decreasing order of certainty. Trajectory-class feasibility is established: a verified, explicitly integrated solar-electric trajectory reaches the required cruise speed by multi-revolution perihelion pumping, without nuclear power. Quantitative closure of the 100 kg concept from low Earth orbit is not achieved by any trajectory found, a verdict stated conditionally because closure below the $\sim$22.3 $\mathrm{km\,s^{-1}}$ heliocentric crossover cannot be excluded within the 12-year operational horizon (no such trajectory found, referred to the design thrust level, comes within 0.6 $\mathrm{km\,s^{-1}}$ of it). Specification compliance is achieved at 100 kg from a geostationary-transfer-orbit drop-off—a delivery-orbit deviation from the ground rule—or at roughly 150 kg from low Earth orbit with fragile margins, and the end-2029 launch date is conditional on a propulsion qualification campaign within the demonstrated gridded-ion component class. The \$10M cost requirement is not met under conventional practice (median near \$16M); the gap is programmatic—bus, operations, and margin—not physical. Every headline result was independently recomputed from first principles, with worst-case discrepancy below 0.5%.

Contents

Nomenclature

a_0	thrust acceleration at 1 AU, full power $[\mathrm{m\,s^{-2}}]$	T	arrival (intercept) epoch from launch [yr]
Δv	velocity increment (propulsive budget) $[\mathrm{km\,s^{-1}}]$	v_e	exhaust velocity, $g_0 I_{sp}$ $[\mathrm{km\,s^{-1}}]$
v_∞	heliocentric hyperbolic excess speed $[\mathrm{km\,s^{-1}}]$	β	departure-asymptote tilt from the ecliptic [deg]
$v_{\infty,E}$	Earth-relative hyperbolic excess speed $[\mathrm{km\,s^{-1}}]$	μ	gravitational parameter $[\mathrm{km^3\,s^{-2}}]$
I_{sp}	specific impulse [s]	P	electrical power at 1 AU [kW]
r	heliocentric distance [AU]	m_0, m_f	initial (wet), final mass [kg]
C_3	characteristic energy, $v_{\infty,E}^2$ $[\mathrm{km^2\,s^{-2}}]$	E	specific orbital energy $[\mathrm{km^2\,s^{-2}}]$

AU = astronomical unit; ly = light-year; LEO/GTO = low Earth / geostationary transfer orbit;
SEP = solar-electric propulsion; PPU = power processing unit;
P10/P50/P90 = cost-distribution percentiles; RV = radial velocity.

Executive Summary

This report assesses the feasibility of a small, low-cost interstellar precursor mission. The spacecraft—approximately $100\,\mathrm{kg}$, propelled by solar-electric thrusters expelling xenon, delivered to low Earth orbit (LEO) as a rideshare payload, and launched before the end of 2029—must pass within $2600\,\mathrm{AU}$ of the Alpha Centauri AB barycenter within 500,000 years of departure, without gravity assists, for a total program cost below $10M. Every quantitative result in this assessment was produced and verified by the autonomous physics-research platform of Physical Superintelligence (PSI): agentic research systems developed the independent physics models, the trajectory simulation and optimization, and the probabilistic cost model, and independent verification agents replicated the reference analysis, established a certified lower bound for the heliocentric requirement under adversarial proof audit, carried a fully three-dimensional trajectory optimization to convergence, and ran a seeded, bit-reproducible targeting-dispersion Monte Carlo. Headline numbers were independently recomputed from first principles, and the staged audits were conducted under PSI's Get-Physics-Done verification framework; the computational campaigns and their acceptance gates are specified in Appendix E. Wherever a result's verification status matters, it is stated alongside the result. Eight findings follow.

1. The approach geometry and its optimum are settled. Using published astrometry for the Alpha Centauri AB barycenter (distance $4.365\,\mathrm{ly}$, space velocity $32.38\,\mathrm{km\,s^{-1}}$) [1], the departure-energy optimum lies at an arrival time of $T = 73{,}012$ years after launch. The required heliocentric hyperbolic excess speed (v_∞, the speed retained after escaping the Sun) is $23.64\,\mathrm{km\,s^{-1}}$, aimed $2.48°$ below the ecliptic plane; after borrowing Earth's in-plane orbital velocity of $29.78\,\mathrm{km\,s^{-1}}$, the Earth-relative excess is $18.59\,\mathrm{km\,s^{-1}}$, and a single perigee burn from a $400\,\mathrm{km}$ LEO requires an impulsive velocity increment (Δv) of $13.85\,\mathrm{km\,s^{-1}}$. The optimum is flat: any arrival between roughly 70,000 and 85,000 years costs less than $0.25\,\mathrm{km\,s^{-1}}$ extra (less than $0.1\,\mathrm{km\,s^{-1}}$ between roughly 67,000 and 80,000 years). Neither mission tolerance is binding—the $2600\,\mathrm{AU}$ miss allowance is worth only about $83\,\mathrm{m/s}$ of excess speed, and at the 500,000-year horizon the requirement *rises* to $19.93\,\mathrm{km\,s^{-1}}$ of impulsive Δv because the star is by then receding. Aiming tolerance is loose—$0.41°$, about $168\,\mathrm{m/s}$ of velocity dispersion—and can be met open-loop, with no terminal correction. These results agree with the reference analysis of the Fermi Explorer mission (a 501(c)(3) nonprofit) [12] to within 0.3% and were replicated independently from first principles to better than 0.5%.

2. The minimum Δv depends on the propulsion regime. A chemical vehicle exploiting the full Oberth benefit at perigee needs the $13.85\,\mathrm{km\,s^{-1}}$ floor. A solar-electric vehicle cannot exploit that benefit and needs 30.5–$31.6\,\mathrm{km\,s^{-1}}$ in total: a $7.6\,\mathrm{km\,s^{-1}}$ Earth-escape spiral (roughly 1.5 years through the radiation belts, with array degradation factored in) plus a heliocentric phase quoted across the reference interval $[22.9, 24.0]\,\mathrm{km\,s^{-1}}$. The upper end of that interval is a verified constructive bound—an explicitly integrated, timestep-refined pumping trajectory. The lower end, $22.9\,\mathrm{km\,s^{-1}}$, is a conservative reference rather than a certified bound: it derives from the best trajectory found by optimal-control analysis at a higher thrust acceleration ($a_0' = 1 \times 10^{-3}\,\mathrm{m/s^2}$: $22.66\,\mathrm{km\,s^{-1}}$ at the required cruise speed, $22.94\,\mathrm{km\,s^{-1}}$ at a slightly higher target variant), extended downward by a monotonicity argument valid for throttleable thrust. Because the indirect Pontryagin method converged there but a direct-collocation cross-check did not, that anchor rests on a single method. The requirement does, however, carry a certified lower bound of $16.56\,\mathrm{km\,s^{-1}}$, proved for throttleable-thrust trajectories with *unconstrained* flight time and perihelion no lower than $0.42\,\mathrm{AU}$. That freedom matters: with unlimited time an explicit trajectory reaches cruise speed for $18.87\,\mathrm{km\,s^{-1}}$ (after 33–58 years of powered flight), whereas the best 12-year operational trajectory found requires $23.98\,\mathrm{km\,s^{-1}}$, far above the free-time bound.

3. Solar power suffices for the trajectory class; no nuclear source is needed. Because array power falls as the inverse square of solar distance, a conventional *outward* escape spiral saturates: the maximum achievable v_∞ is approximately 0, 3.4, and $17.0\,\mathrm{km\,s^{-1}}$ at thrust accelerations $a_0 = 1.5$, 5, and $10 \times 10^{-4}\,\mathrm{m/s^2}$ respectively—all below the $23.4\,\mathrm{km\,s^{-1}}$ cruise floor (confirmed by two independent integrators). Perihelion pumping defeats this wall: retrograde thrust arcs near $1\,\mathrm{AU}$ shed angular momentum, dropping perihelion to $0.42\,\mathrm{AU}$, and prograde burn arcs then concentrate near perihelion where power is one to four times the 1-AU rating. The optimized profile at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m/s^2}$ delivers $23.98\,\mathrm{km\,s^{-1}}$ of Δv in 12.0 years of flight time (mostly coasting); the pumping mechanism was independently reconstructed during verification, and the result survives halving the assumed perihelion power cap at a cost of about $1.9\,\mathrm{km\,s^{-1}}$. For a 2029 launch the vehicle escapes the solar system around 2043; the arrival epoch is unaffected because the optimum is flat, and over 98% of the mission is ballistic coast. Fermi Explorer's reference analysis concluded that pure solar-electric propulsion does not close from LEO *at today's array specific power* (~60 W/kg), identifying a

pure-solar closing corner above roughly $100\,\mathrm{W/kg}$, and separately rejected a *single-pass* electric perihelion-burn architecture. Multi-revolution pumping amends both conclusions at once: by answering the single-pass objection, it closes at today's ~60 W/kg.

4. The 100 kg concept closes from GTO, not from LEO. With the required specific impulse of $2800\,\mathrm{s}$ and present-day component parameters, the minimum closing wet mass from LEO is 113.5–$146\,\mathrm{kg}$ across the 22.9–$24.0\,\mathrm{km\,s^{-1}}$ reference interval: the $100\,\mathrm{kg}$ concept misses by 1.0–2.8% of wet mass for every trajectory found, while $150\,\mathrm{kg}$ closes across the entire interval in the planar accounting (the measured out-of-plane cost overturns this at the conservative end; finding 5). The LEO verdict is conditional, not absolute: a $100\,\mathrm{kg}$ vehicle from LEO cannot be ruled out if some heliocentric trajectory below the ~22.3 $\mathrm{km\,s^{-1}}$ crossover exists within the 12-year operational horizon—but no trajectory found in this analysis, referred to the design thrust level, comes within $0.6\,\mathrm{km\,s^{-1}}$ of that crossover. Starting instead from a geostationary transfer orbit (GTO) rideshare cuts the departure requirement from 7.6 to $4.24\,\mathrm{km\,s^{-1}}$: the $100\,\mathrm{kg}$ concept then closes with margin (+4.1 points at the conservative end; even $80\,\mathrm{kg}$ closes). Under a 30% thrust-acceleration margin, the GTO start retains $100\,\mathrm{kg}$ closure at the favorable end of the interval (boundary $83.4\,\mathrm{kg}$) but needs approximately $110\,\mathrm{kg}$ to carry the margin across the whole interval (boundary $103.7\,\mathrm{kg}$ at the conservative end); the same margin drives the LEO boundary to 312–$910\,\mathrm{kg}$. The GTO start remains robust where the LEO start is fragile: under the worst-case out-of-plane increment, GTO closes from $80\,\mathrm{kg}$ while the $150\,\mathrm{kg}$ LEO vehicle fails (boundary near $198\,\mathrm{kg}$). The closure arithmetic was verified algebraically in an independent recomputation. The recommendation is to procure the GTO drop-off with a 100–$110\,\mathrm{kg}$ vehicle.

5. The out-of-plane cost is measured, not estimated, and it strengthens the GTO recommendation. The results above are planar, with the $2.48°$ aim below the ecliptic treated as an estimated increment between $22\,\mathrm{m/s}$ and $1.02\,\mathrm{km\,s^{-1}}$. A genuinely three-dimensional optimization measures that increment at $0.58\,\mathrm{km\,s^{-1}}$—inside the estimated band, and the achieved cost of the best trajectories found rather than a bound on the true optimum. The $100\,\mathrm{kg}$ GTO vehicle still closes at both ends of the 22.9–$24.0\,\mathrm{km\,s^{-1}}$ interval (+5.1/+3.1 points; +2.3 even at the $1.02\,\mathrm{km\,s^{-1}}$ worst case); the $150\,\mathrm{kg}$ LEO fallback closes at $22.9\,\mathrm{km\,s^{-1}}$ (+1.0 points) but fails at the conservative $24.0\,\mathrm{km\,s^{-1}}$ end (–0.8 points). The GTO drop-off recommendation is unchanged and strengthened.

6. Open-loop targeting confidence is 0.98 under the stated error budget. A seeded 20,000-trial Monte Carlo propagated execution and ephemeris errors to the encounter: the probability of passing within $2600\,\mathrm{AU}$ of the barycenter is 0.982 (95% confidence interval $[0.980, 0.983]$), independently reproduced by an analytic covariance calculation to 0.01%. Thrust-magnitude knowledge dominates the error budget, ahead of pointing, ephemeris, and departure residuals. One caveat deserves emphasis: this is a simulation conditioned on assumed engineering error levels, not on flight-measured performance. Doubling the assumed thrust-magnitude error alone drops the probability to 0.858; doubling the assumed pointing error, to 0.897—both below 0.9.

7. The cost gap is programmatic, not physical. A probabilistic parametric model (2026 USD) places the median program cost at \$15.7M–\$16.6M across the closing architectures, with essentially zero probability of meeting \$10M under conventional aerospace practice. The structure of the overrun is instructive: launch (\$1.0–1.3M) and propulsion (\$0.6–1.2M) each come in at roughly one third of their \$3M allocations, while the entire gap sits in the bus, operations, and program margin (approximately \$12M nominal against \$3M allocated). Because the mass-closing trajectory implies about 13 years of powered-flight custody, automation is structural to the cost estimate, not merely a communications convenience. The path to \$10M runs through lean automated operations and compressed integration and test—not through new physics.

8. A cheaper star exists. If the objective is a passage through *any* star's Oort-cloud environs rather than Alpha Centauri specifically, the M5V red dwarf LSPM J2146+3813 ($22.99\,\mathrm{ly}$) dominates: an impulsive floor of $9.30\,\mathrm{km\,s^{-1}}$ versus $13.85\,\mathrm{km\,s^{-1}}$, a cruise speed of $7.74\,\mathrm{km\,s^{-1}}$ versus $23.64\,\mathrm{km\,s^{-1}}$, and an arrival at 78,017 years—the same timeline. No perihelion pumping is needed at all, since the conventional-spiral ceiling exceeds the requirement. Its published future encounter with the Sun, $0.568\,\mathrm{pc}$ at +82.5 kyr, is among the tightest known [6]. A sweep of the same encounter catalog found no higher-confidence target inside the mission horizon (Section 8).

Table 1 condenses the findings and recommendations.

Table 1. Summary of findings and recommendations.

Topic	Finding	Recommendation
Arrival-time optimum	T = 73,012 yr; $v_\infty = 23.64\,\mathrm{km\,s^{-1}}$ aimed $2.48°$ below the ecliptic; flat to $< 0.25\,\mathrm{km\,s^{-1}}$ over 70–85 kyr	Choose the arrival epoch within the window for launch-phasing convenience
Minimum Δv from LEO	$13.85\,\mathrm{km\,s^{-1}}$ impulsive; 30.5–$31.6\,\mathrm{km\,s^{-1}}$ solar-electric ($24.0\,\mathrm{km\,s^{-1}}$ verified constructive bound; $22.9\,\mathrm{km\,s^{-1}}$ best-found single-method anchor; certified lower bound $16.56\,\mathrm{km\,s^{-1}}$, unconstrained flight time)	Size electric propulsion to the top of the reference interval
Solar-electric feasibility	Perihelion pumping delivers the full cruise speed ($23.98\,\mathrm{km\,s^{-1}}$ in 12.0 yr at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m/s^2}$); survives halving the perihelion power cap for $\sim 1.9\,\mathrm{km\,s^{-1}}$	No nuclear source required; qualify a gridded-ion propulsion string
Vehicle mass closure	$100\,\mathrm{kg}$ misses from LEO for every trajectory found (conditional: $\sim 22.3\,\mathrm{km\,s^{-1}}$ crossover caveat); closes from GTO, where $\sim 110\,\mathrm{kg}$ also carries a 30% thrust margin across the interval	Procure a GTO drop-off; fly a 100–110 kg vehicle
Out-of-plane cost	Measured at $0.58\,\mathrm{km\,s^{-1}}$ by three-dimensional optimization, inside the prior 22 m/s–$1.02\,\mathrm{km\,s^{-1}}$ planar estimate; GTO-100 closes at both bracket ends; LEO-150 fails at the conservative end	Keep the GTO drop-off; do not rely on the 150 kg LEO fallback
Targeting dispersion	$P(\mathrm{pass} \leq 2600\,\mathrm{AU}) = 0.982$ $[0.980, 0.983]$ under the stated error budget, dominated by thrust-magnitude knowledge; doubling assumed execution errors drops confidence below 0.9	Treat thrust-magnitude knowledge as a driving requirement; revalidate with measured hardware error levels
Program cost	Median $15.7–16.6M; probability of meeting $10M near zero under conventional practice; entire gap in bus and operations	Pursue lean automated operations and compressed integration and test
Alternative target	LSPM J2146+3813: $9.30\,\mathrm{km\,s^{-1}}$ impulsive floor, one-third the cruise speed, same arrival timeline	Prefer this target if the objective is any star's Oort cloud

Answers to the Mission's Questions

This section gathers in one place the report's verdicts on the four commissioned objectives and its answers to the Fermi Explorer team's direct questions. Each entry is a summary; the section it cites carries the full analysis.

Requirements Compliance

Table 2. Compliance of the four mission objectives (Table 3).

Objective	Verdict	Basis
Pass within 2600 AU of the Alpha Centauri AB barycenter within 500,000 yr	**Met** (by the recommended architecture)	Arrival at ~73,000 yr; the Fermi Explorer team's planned aim at the ~80,000-year ecliptic-plane crossing is equally acceptable, costing under $0.1\,\mathrm{km\,s^{-1}}$ over the optimum. With all execution and ephemeris error sources acting together under the stated engineering error budget, the probability of passing inside the sphere open-loop is 0.982 (95% confidence interval $[0.980, 0.983]$); the verdict degrades below 0.9 if either the assumed thrust-magnitude or pointing error doubles (Section 2.6, Section 6).
Payload of at least 1 kg (1U class)	**Met**	Carried explicitly as a fixed line item in every mass budget in this report; the recommended 100 kg architecture closes with the payload aboard at a margin of +3.1 to +5.1 points of wet mass (Section 5).
Launch before the end of 2029	**Schedule-conditional**	Rideshare procurement is routine, but the propulsion qualification campaign (Section 5) and MESSENGER-class thermal qualification, from a mid-2026 start, make an end-2029 launch aggressive. No full schedule analysis was performed; the verdict is conditional on one.
Program cost below \$10M (\$3M launch / \$3M propulsion / \$3M bus and operations)	**Not met** under conventional practice	Median estimate \$15.7M–\$16.6M. Launch and propulsion under-run their allocations; the entire gap sits in bus, operations, and margin—programmatic, not physical (Section 7). A lean-automation path is described in Section 7.5; it takes the Fermi Explorer team at their word: their stated tolerance for mostly-automated operations—a minimal antenna and eventual loss of contact included—is that path's load-bearing assumption, not an operational afterthought (Section 7.4).

Direct Answers

Q1 — What is the minimum spacecraft Δv from any low-Earth-orbit drop-off, with no gravity assist, arriving within 500,000 years? For impulsive (chemical) propulsion, $13.85\,\mathrm{km\,s^{-1}}$: a single perigee burn from a 400 km orbit capturing the full Oberth benefit (Section 3). For solar-electric propulsion, 30.5–$31.6\,\mathrm{km\,s^{-1}}$: a $7.6\,\mathrm{km\,s^{-1}}$ Earth-escape spiral plus a heliocentric phase whose cost spans the reference interval 22.9–$24.0\,\mathrm{km\,s^{-1}}$. The upper end of that interval is a verified constructive bound: an explicitly integrated, timestep-refined trajectory reaching cruise speed for $23.98\,\mathrm{km\,s^{-1}}$ in 12 years, the best 12-year operational figure found. The lower end rests on an optimal-control anchor at four times the design thrust acceleration, where an indirect method based on Pontryagin's minimum principle converges to $22.66\,\mathrm{km\,s^{-1}}$ at the required cruise speed, a result extended by a monotonicity argument valid for throttleable thrust; the anchor is single-method—direct collocation did not converge on this problem despite repeated documented attempts—and marks the best trajectory found, not a certified bound. The certified floor is $16.56\,\mathrm{km\,s^{-1}}$, proved for every admissible thrust program in the throttleable power-envelope class with the 0.42 AU perihelion bound and no constraint on flight time (the impulsive multi-burn limit is $16.50\,\mathrm{km\,s^{-1}}$, falling to $15.6\,\mathrm{km\,s^{-1}}$ if 0.25 AU is tolerable). Freed of the 12-year operational horizon, an admissible trajectory at $18.87\,\mathrm{km\,s^{-1}}$ exists at the price of 33–58 years of powered flight, so the operational figure is a schedule-shaped value, not the class optimum. These figures already exercise the offered freedom to choose the drop-off orbit within the low-Earth class: the injection plane is taken to contain the departure asymptote tilted $2.48°$ below the ecliptic, so no separate plane change is charged, and the altitude lever is weak—worth a few tenths of a $\mathrm{km\,s^{-1}}$ across low-Earth altitudes (Section 3). The drop-off choice that genuinely matters lies beyond low Earth orbit: a geostationary-transfer-orbit delivery cuts the electric-propulsion escape cost from 7.6 to $4.24\,\mathrm{km\,s^{-1}}$, and it is this single procurement change, rather than any selection among low orbits, that closes the 100 kg vehicle (Section 5).

Q2 — What mission profile achieves it? Rideshare to low Earth orbit; an escape spiral of roughly 1.5 years through the radiation belts; then a perihelion-pumped heliocentric phase—retrograde thrust arcs near 1 AU shed angular momentum to drop perihelion to 0.42 AU, and prograde burns concentrate there, where power peaks—lasting about 12 years, most of it coasting. The vehicle departs on an asymptote $2.48°$ below the ecliptic aimed at the 73,000-year intercept and coasts ballistically thereafter; over 98% of the mission is unpowered (Section 4).

Q3 — Does electric propulsion plus solar power work, given the inverse-square falloff of sunlight? Yes, for the trajectory class. Outward spirals saturate below the required cruise speed at the design thrust, but perihelion pumping does not, because its burns concentrate where power is highest; the mechanism survives halving the assumed perihelion power cap at a cost of about $1.9\,\mathrm{km\,s^{-1}}$ (Section 4). At the vehicle level, closure comes at 100 kg from a geostationary-transfer-orbit drop-off or at approximately 150 kg from low Earth orbit. The measured $0.58\,\mathrm{km\,s^{-1}}$ three-dimensional cost of the tilted departure asymptote leaves the first verdict intact at both ends of the reference interval but defeats the second at the interval's conservative end; both closures are conditional on qualifying a propulsion string within the demonstrated gridded-ion component class (Section 5). The solar-electric answer also holds at the targeting level: the open-loop aim delivered at the end of the powered phase passes inside the required sphere with probability 0.982 under the stated error budget (Section 6).

The Fermi Explorer Team's Five Arrival-Time Intuitions

All five were tested against the full optimization (Section 2.5):

1. The budget rises sharply beyond ~80,000 yr — **confirmed** (the impulsive requirement grows from $13.85\,\mathrm{km\,s^{-1}}$ at the optimum to $19.93\,\mathrm{km\,s^{-1}}$ at the 500,000-year horizon, once the star begins to recede);
2. the closest natural approach (~3 ly) is not the optimum — **confirmed**, with the epoch refined: the approach occurs at 27,955 yr with the star at 3.15 ly, against the anticipated ~20,000 yr at 3 ly, and intercepting there is the costliest option examined;
3. a tangential intercept near 60,000 yr minimizes cruise speed — **confirmed exactly** ($T^* = 58{,}422$ yr, at the star's transverse velocity of $23.38\,\mathrm{km\,s^{-1}}$, from a Sun-only analysis with no planetary departure assumed);
4. the ecliptic-plane crossing minimizes the departure Δv — **very nearly right** (the 79,786-yr crossing costs under $0.1\,\mathrm{km\,s^{-1}}$ more than the optimum; the basin is flat to better than $0.1\,\mathrm{km\,s^{-1}}$ for arrivals between roughly 67,000 and 80,000 yr);
5. the optimum lies near 73,000 yr, aimed a few degrees off-plane — **confirmed almost exactly** (73,012 yr, $2.48°$ below the ecliptic).

The Two Flexibility Branches

At $50M. The recommended use of the tier is an escape-injection rideshare: purchasing the departure energy from the launch provider closes the vehicle at approximately 60 kg and removes the belt-transit, array-degradation, and patience risks outright, at a median cost near $15.6M (Section 7.6). A dedicated multi-stage chemical kick stage sized to the $13.85\,\mathrm{km\,s^{-1}}$ impulsive departure floor buys the same physics less well: it is

Table 3. Mission requirements and ground rules.

Requirements	
Payload	At least $1\,\mathrm{kg}$ (1U CubeSat class)
Encounter	Pass within $2600\,\mathrm{AU}$ of the Alpha Centauri AB barycenter within 500,000 years of launch
Launch date	Before the end of 2029
Program cost	Below \$10M total: \$3M launch, \$3M propulsion, \$3M bus and mission operations
Ground rules	
Injection	Rideshare drop-off in any low Earth orbit
Trajectory	No gravity assists
Operations	Mostly automated operations acceptable; loss of contact during cruise is tolerable

more expensive and less reliable than purchasing the equivalent injection energy from the launch vehicle. A Jupiter gravity assist, legitimately back on the table at this tier, offers of order 10–$15\,\mathrm{km\,s^{-1}}$ of heliocentric excess but is window-contingent: its value depends on Jupiter phasing against the end-2029 launch date, and it adds transfer years and deep-space operations of exactly the class that dominates program cost (Section 7.6). Radioisotope generators cannot supply the kilowatt-class power the electric propulsion requires, and a fission reactor exceeds even this tier. Among methods beyond those the Fermi Explorer team named, the cheapest sits below the tier entirely: the geostationary-transfer-orbit drop-off of Section 5, which nearly halves the escape budget and closes the $100\,\mathrm{kg}$ vehicle without touching the additional funds. The remaining assist-class option—a solar-Oberth maneuver with a purchased kick stage—demands thermal protection of the Parker Solar Probe class and is assessed in Section 7.6.

Better stars. All three offered candidates were adjudicated. LSPM J2146+3813 is confirmed as the best target on offer: an impulsive floor of $9.30\,\mathrm{km\,s^{-1}}$ at a ~78,000-year arrival, reachable without perihelion pumping (Section 8.1). Lambda Serpentis is a sound solar-type second choice at $11.38\,\mathrm{km\,s^{-1}}$ and a ~152,000-year arrival, confirming the anticipated 150,000–200,000-year window (Section 8.2). Alpha-2 Librae should be removed from the list: its minimum-cruise-speed intercept lies at 1.03 million years, outside the 500,000-year horizon, and forcing an in-horizon arrival costs $19.74\,\mathrm{km\,s^{-1}}$—worse than Alpha Centauri itself—a verdict robust across the full $\pm 2\sigma$ range of its poorly determined systemic radial velocity of $-22.0\pm 5.8\,\mathrm{km\,s^{-1}}$ (Section 8.3). Beyond the offered set, the Bailer-Jones catalog of future stellar encounters [6] was swept for cheaper targets (Section 8.4). None displaces LSPM J2146+3813 as the best high-confidence target: the only entries that formally undercut it are two faint, unnamed Gaia M dwarfs whose catalog radial velocities carry 8–$11\,\mathrm{km\,s^{-1}}$ uncertainties—too weak to aim at. Every unflagged catalog candidate inside the 500,000-year horizon is an M dwarf, so among Sun-like stars Lambda Serpentis stands unchallenged.

1. Introduction

1.1 Mission Statement

The Fermi Explorer mission (a 501(c)(3) nonprofit) seeks the least expensive spacecraft mission that can reach the near environs of another star. The specific formulation under assessment is a probe that passes within $2600\,\mathrm{AU}$ of the Alpha Centauri AB barycenter—a distance comparable to the scale of a stellar Oort cloud—within 500,000 years of departure, launched before the end of 2029, on a total program budget below \$10M. The Fermi Explorer team has published its own reference analysis of the problem [12]. The present report is an independent feasibility assessment by Physical Superintelligence (PSI), built on separately developed models and data; it is intended both to answer the feasibility question on its own terms and to cross-check the reference analysis. Table 3 states the four mission requirements and the accompanying ground rules.

1.2 Concept Vehicle

The Fermi Explorer concept vehicle is a spacecraft of approximately $100\,\mathrm{kg}$ wet mass—most of it xenon propellant—that uses solar-electric propulsion to build up the required departure energy and carries the $1\,\mathrm{kg}$ payload. The operational posture is deliberately expendable: the mission is judged successful once the probe is placed on a verified intercept trajectory, and continued contact through the multi-millennial cruise is not required. Whether this concept closes—in trajectory physics, in vehicle mass, and in cost—is the central question of this report.

1.3 Scope

The assessment covers the interstellar approach geometry and the choice of arrival epoch; departure velocity budgets from low Earth orbit in both the impulsive (chemical) and low-thrust (solar-electric) regimes; the feasibility of solar-electric escape given the inverse-square decline of array power with solar distance; vehicle mass closure at the level of component parameters (thruster specific impulse and efficiency, array specific power, tankage, and power processing); a probabilistic program cost estimate against the $10M target and its internal allocations; and a survey of alternative stellar targets evaluated by identical methods. It does not cover detailed spacecraft bus design, launch-vehicle selection, or subsystem-level thermal and communications design. Where those disciplines constrain feasibility—for example, the thermal qualification implied by repeated close perihelion passes—the constraint is identified and its cost or mass consequence carried, but no design is attempted.

1.4 Assessment Methodology

The analysis rests on four legs, each constructed independently of Fermi Explorer's models and code. The first is independent physics models: the kinematics of the encounter are computed from published astrometry of the Alpha Centauri system [1] and cross-checked against independent orbit and parallax determinations [2, 3], with closed-form and numerical treatments required to agree. The second is trajectory simulation and optimization: low-thrust escape trajectories are propagated with power-limited thrust models, and the key heliocentric velocity budget is quoted as a reference interval rather than any single point solution. The upper end of that interval is a verified constructive bound—an explicitly integrated, timestep-refined trajectory—while the lower end is the best trajectory found by optimal-control analysis at higher thrust, extended downward by a monotonicity argument; the interval carries the quantitative conclusions, and the certification status of each end is stated wherever it is used. The third is a probabilistic cost model: fourteen parametric line items with triangular uncertainty distributions, sampled by Monte Carlo, so that cost conclusions take the form of percentile ranges rather than point estimates.

The fourth leg is independent verification. The analysis was carried through PSI's Get-Physics-Done (GPD) verification framework, a staged, artifact-bound pipeline in which the manuscript's claims are indexed and audited by independent passes for mathematical soundness, physical support, literature positioning, and proof integrity, each pass writing machine-validated artifacts. Every headline quantity in this report was re-derived under that framework by an independent verification pass, working from the cited astrometry and the stated models rather than from the originating code. The worst kinematic discrepancy is below 0.5%, the low-thrust saturation result was confirmed by two independent integrators, and the mass-closure arithmetic was verified algebraically. Where a standard technique failed—direct collocation trajectory optimization does not converge in the low-thrust, multi-revolution regime—the failure is reported and no conclusion depends on it.

These four legs were executed end-to-end by PSI's autonomous physics-research platform. Autonomous research agents carried out the trajectory-optimization, mass-closure, cost, and dispersion campaigns, from model construction through converged results, and the independent re-derivation of every headline number was performed by separate verification agents working from the indexed claims rather than from the originating code. Claims advanced at theorem grade—principally the certified lower bound on the heliocentric velocity budget—were further subjected to an adversarial proof audit in which a dedicated agent attempted to defeat each step of the argument before the proof was accepted. Every Monte Carlo result in this report is seeded and bit-reproducible; the seed discipline and acceptance gates are documented in Appendix E. The platform's role is confined to the campaign this report documents; no conclusion rests on unaudited machine output.

Three elements of the verification campaign carry particular weight and are reported in full. The heliocentric velocity budget is anchored from below by a certified lower bound: an analytic proof, checked by a dedicated adversarial audit, that no admissible thrust program in the assessed class can depart the solar system for less than $16.56\,\mathrm{km\,s^{-1}}$ without a time-of-flight restriction. The floor of the feasibility argument therefore rests on a theorem rather than on the absence of better trajectories found. The out-of-plane cost of the tilted departure asymptote is computed directly on a fully three-dimensional integrated trajectory ($578\,\mathrm{m\,s^{-1}}$ at the flyable design point, quoted as the difference between the best trajectories found with and without the out-of-plane requirement rather than as a bound on the true optimum). And the probability of meeting the $2600\,\mathrm{AU}$ encounter criterion under execution and knowledge errors—departure residuals, thrust pointing and magnitude knowledge, and target ephemeris—is quantified by a seeded 20,000-sample Monte Carlo dispersion analysis, reported as a simulation conditioned on stated engineering error assumptions rather than on flight-measured performance. As a further anchor, the verification pass reproduces the geometric optimum of the Fermi Explorer reference analysis at the sub-percent level (Section 2.4 gives the per-quantity comparison), so the two analyses share a common baseline before their conclusions diverge.

1.5 Report Organization

The report proceeds from kinematics to program cost. It first establishes the approach geometry and the arrival-time optimum, then develops the departure velocity budgets from low Earth orbit in each propulsion regime. The solar-electric feasibility question follows, including the power-limited saturation of conventional escape spirals and the perihelion-pumping trajectory that overcomes it. Vehicle mass closure and its sensitivities are treated next, then an execution-dispersion analysis that propagates departure, pointing, thrust-knowledge, and ephemeris uncertainties through the multi-millennial cruise to the probability of meeting the encounter criterion, followed by the probabilistic cost assessment against the \$10M target. The report then evaluates alternative stellar targets under identical methods, compares the present results with the Fermi Explorer reference analysis, and closes with conclusions and recommendations. Appendices document the independent verification of each headline result.

2. Intercept Geometry and Arrival-Time Optimization

This section establishes the kinematic foundation of the mission: where the target will be, when it is cheapest to meet it, and how tightly the departure must be aimed. All results here follow exactly from the target's measured space motion; every headline quantity was independently recomputed to better than 0.5% (Section 1), and the departure optimum agrees with the Fermi Explorer reference analysis [12] at the sub-percent level (Section 2.4 gives the per-quantity comparison).

2.1 Kinematics of the Target

The mission aims at the barycenter of the Alpha Centauri AB binary. We adopt the orbit-corrected barycentric astrometry of Kervella et al. [1]: distance $d = 4.365$ ly and heliocentric space velocity $32.38\,\mathrm{km\,s^{-1}}$, resolved into a radial component $v_r = -22.40\,\mathrm{km\,s^{-1}}$ (approaching the Sun) and a transverse component $v_t = 23.38\,\mathrm{km\,s^{-1}}$. Over the 10^4–10^5-year horizons of interest the barycenter is modeled in uniform straight-line motion. This idealization costs almost nothing: a closed-form bound on the differential Galactic tidal acceleration between the Sun and the target limits the resulting perturbation over the 73,000-year flight to about $1.2\,\mathrm{AU}$ in position and under $0.2\,\mathrm{m\,s^{-1}}$ in velocity—negligible against the $2600\,\mathrm{AU}$ encounter sphere—and the linear model is validated externally against published orbit-integrated encounter solutions (Section A.4). Appendix A details the astrometric inputs, the coordinate transformation, and the closed-form results quoted below.

Figure 1 shows the resulting approach geometry. Today the barycenter lies roughly 3 ly below the ecliptic plane and is closing on the Sun; in the tens of millennia ahead it makes its closest natural approach, crosses the plane, and then recedes. Four epochs on this path organize the analysis and are collected, together with the 500,000-year requirement horizon, in Table 4.

2.2 The Lead-the-Target Relation

Once beyond the Sun's gravitational influence, the probe cruises in essentially a straight line at its hyperbolic excess speed v_∞. Let $\mathrm{r}_0 = d\,\hat{\mathrm{n}}$ be the target's present position ($\hat{\mathrm{n}}$ the unit line-of-sight vector) and $\mathrm{v}_\star$ its space

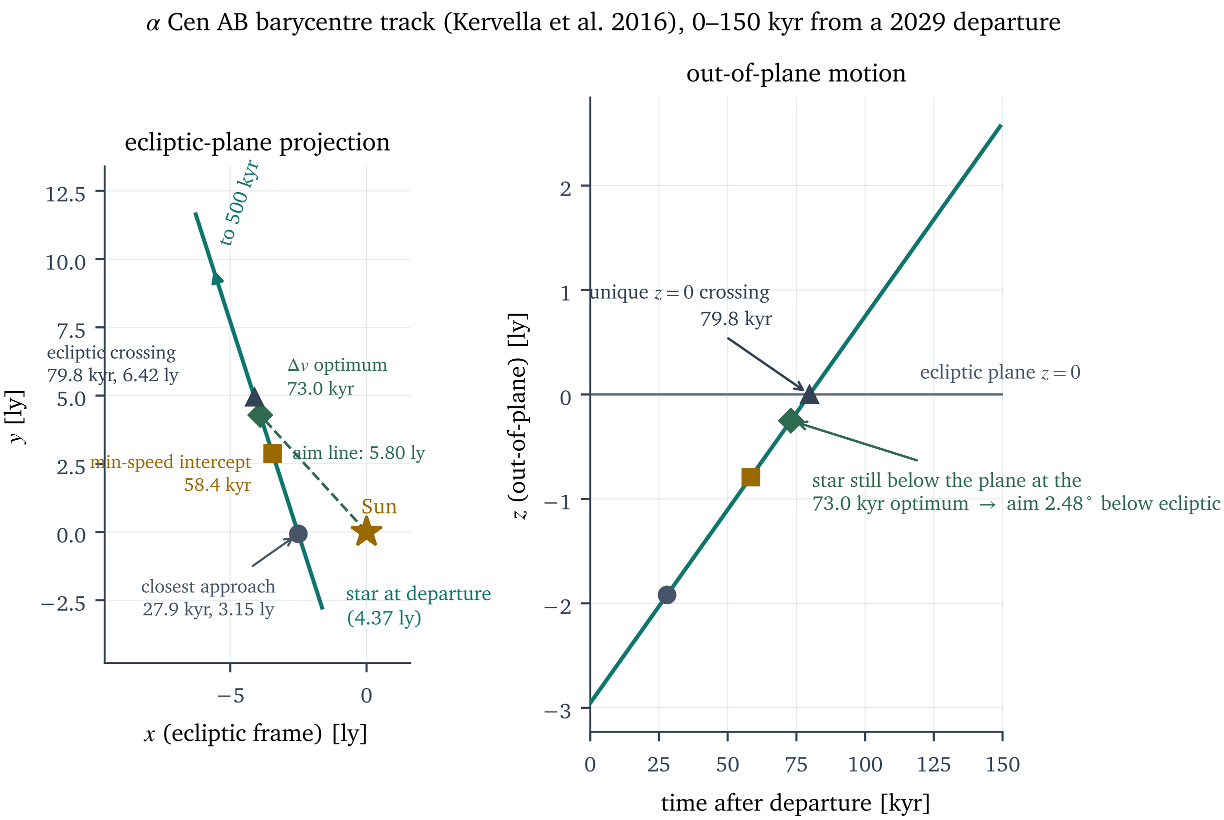

Figure 1. Approach geometry of the Alpha Centauri AB barycenter over the next 150,000 years, from the astrometry of [1]. (a) Projection of the star's path onto the ecliptic plane, with the Sun, the aim line to the 73.0 kyr intercept point, and the four characteristic epochs of Table 4 marked. (b) Out-of-plane coordinate z versus time: the star currently sits about 3 ly below the ecliptic and crosses the plane 79,786 years from launch.

velocity. Meeting the star at time T after departure requires the cruise velocity

$$\mathrm{V}(T) = \frac{\mathrm{r}_0}{T} + \mathrm{v}_\star, \tag{1}$$

i.e. the probe must fly not toward where the star is, but toward where it will be. The required velocity is the sum of a *closing* term r_0/T, which shrinks as the mission accepts a longer flight time, and the star's own velocity, which must be matched in full no matter how patient the mission is. Decomposing $\mathrm{v}_\star$ into its radial and transverse parts gives the scalar form

$$v_\infty^2(T) = \left(\frac{d}{T} + v_r\right)^2 + v_t^2. \tag{2}$$

Because the star is approaching ($v_r < 0$), the radial term vanishes identically at $T^* = d/|v_r| = 58{,}422$ years: at that arrival time the star's own approach closes the line-of-sight distance for free, and the required cruise speed reaches its absolute floor, $v_\infty = v_t = 23.38\ \mathrm{km\,s^{-1}}$ — exactly the star's transverse velocity. The identity is exact under linear motion for a *true* intercept (Appendix A), and the floor it sets is the single most important number in the study: *no trajectory, however patient, can truly intercept Alpha Centauri at a cruise speed below the star's transverse velocity.* The mission's 2600 AU miss allowance lowers the effective floor only slightly, to $23.17\ \mathrm{km\,s^{-1}}$ at the re-optimized epoch, and the headline optimum of Section 2.4 already includes that allowance. For earlier arrivals the closing term dominates and the requirement rises steeply; for later arrivals the star recedes and the requirement rises again, gently at first.

Table 4. Characteristic intercept epochs for the Alpha Centauri AB barycenter, from launch. The impulsive Δv is a single perigee burn from a 400 km low Earth orbit with full Oberth benefit (Section 2.4). Every entry was independently recomputed from the cited astrometry to better than 0.5%.

Event	Arrival T (yr)	Cruise v_∞ ($\mathrm{km\,s^{-1}}$)	Comment
Closest natural approach	27,955	—	star at 3.152 ly; costliest intercept examined
Minimum cruise speed	58,422	23.38	$v_\infty = v_t$ exactly (tangential intercept)
Departure-Δv optimum	73,012	23.64	aim $2.48°$ below ecliptic; $\Delta v = 13.85\ \mathrm{km\,s^{-1}}$
Ecliptic-plane crossing	79,786	24.14	star at 6.425 ly; within $0.1\ \mathrm{km\,s^{-1}}$ of optimum
Requirement horizon	500,000	30.61	star receding; $\Delta v = 19.93\ \mathrm{km\,s^{-1}}$ (not binding)

2.3 Characteristic Epochs

The first four epochs of Table 4 are distinct because they minimize different quantities. The closest *natural* approach (27,955 years, 3.152 ly) minimizes the distance flown but is by far the most expensive intercept examined — the probe must close 4.365 ly in under 28,000 years, driving the departure budget to more than twice the optimum (Figure 2). The tangential intercept (58,422 years) minimizes the cruise speed. The departure-Δv optimum (73,012 years) minimizes what actually matters — the propulsive cost of leaving Earth — and falls between the tangential intercept and the ecliptic-plane crossing (79,786 years) for the reason developed next.

2.4 Arrival-Time Optimization

The departure cost depends on more than the cruise speed. A probe departing the Earth–Moon system borrows Earth's orbital velocity of $29.78\ \mathrm{km\,s^{-1}}$, but the loan extends only to the component of the outbound asymptote that lies *in* the ecliptic plane: the departure epoch is chosen so that the in-plane projection of the asymptote aligns with Earth's motion, while any out-of-plane component must be bought with propellant. Because the intercept point sits below the ecliptic for arrivals earlier than the plane crossing, the optimizer faces a trade (Figure 3): wait until the star crosses the plane (no tilt, but a more distant, faster-receding target) or depart slightly out of plane earlier (a small tilt penalty, but a cheaper intercept).

We therefore minimized the impulsive departure Δv from a 400 km low Earth orbit jointly over the arrival time T and the aim direction, using patched-conic departure kinematics (Appendix A gives the closed-form chain: cruise $v_\infty \to$ Earth-relative excess $v_{\infty,E} \to$ perigee-burn Δv, with reference speeds $v_{\mathrm{circ}} = 7.67\ \mathrm{km\,s^{-1}}$ and $v_{\mathrm{esc}} = 10.85\ \mathrm{km\,s^{-1}}$ at 400 km, and solar escape speed $42.12\ \mathrm{km\,s^{-1}}$ at 1 AU). The result:

- optimum arrival $T = 73{,}012$ years after launch;
- required cruise speed $v_\infty = 23.64\ \mathrm{km\,s^{-1}}$, only $0.26\ \mathrm{km\,s^{-1}}$ above the absolute floor;
- departure asymptote tilted $\beta = 2.48°$ *below* the ecliptic;
- Earth-relative hyperbolic excess $v_{\infty,E} = 18.59\ \mathrm{km\,s^{-1}}$;
- impulsive departure $\Delta v = 13.85\ \mathrm{km\,s^{-1}}$ from a single perigee burn.

The same epoch and tilt minimize all three departure metrics examined (impulsive Δv, Earth-relative excess, and a low-thrust departure proxy), so the optimum is not an artifact of the impulsive assumption. Independent recomputation reproduces all six numbers characterizing the optimum to better than 0.5%, and Fermi Explorer's reference analysis [12] — computed with entirely separate code — places its optimum at $13.88\ \mathrm{km\,s^{-1}}$ at 72,800 years with a $2.4°$ tilt. The departure Δv, the arrival epoch, and the Earth-relative excess agree with ours to 0.3% or better; the optimum aim tilt agrees to $0.09°$ absolute, the tilt being ill-conditioned along the flat valley of Figure 2, where a small absolute spread is expected.

The flatness of the optimum has a practical consequence worth emphasizing: the Δv curve of Figure 2 varies by less than $0.25\ \mathrm{km\,s^{-1}}$ across arrival times from roughly 70,000 to 85,000 years, and by less than $0.1\ \mathrm{km\,s^{-1}}$ from roughly 67,000 to 80,000 years. Arrival-time selection inside this window costs essentially nothing and can be ceded entirely to launch-window phasing and operational convenience. In particular, the multi-year

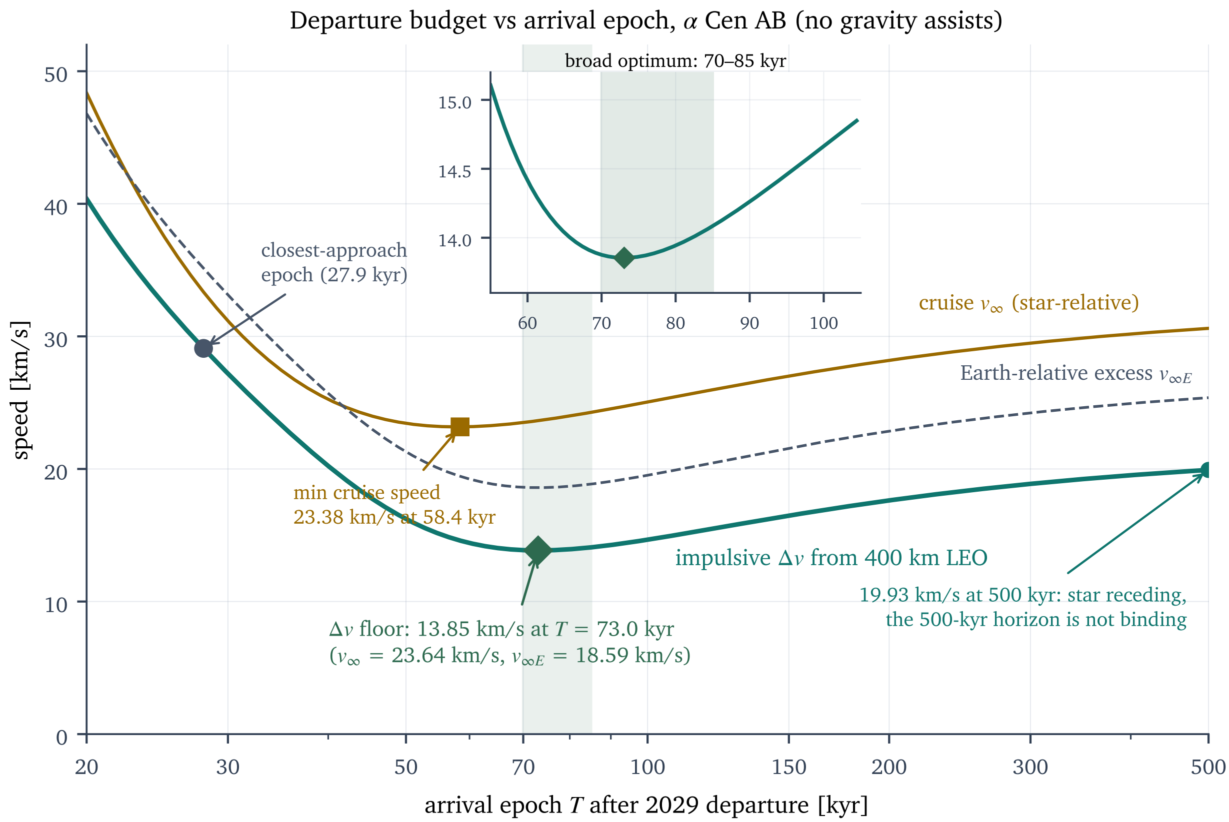

Figure 2. Departure cost versus arrival time. Impulsive Δv from a 400 km low Earth orbit and required cruise v_∞ as functions of arrival time on a logarithmic axis from 20 to 500 kyr, with the optimum, minimum-cruise-speed, and closest-approach epochs and the 500 kyr endpoint marked. The inset expands 55–105 kyr: the curve is flat to within $0.25\ \mathrm{km\,s^{-1}}$ across the shaded 70–85 kyr window, so the arrival-time choice within that window is free for other considerations. Past ~80 kyr the star recedes and the requirement rises monotonically, reaching $\Delta v = 19.93\ \mathrm{km\,s^{-1}}$ at the 500 kyr horizon.

low-thrust escape examined in later sections delays solar-system departure to roughly 2043 for a 2029 launch — negligible against a 73,000-year cruise and a flat optimum.

2.5 Design Intuitions Examined

The Fermi Explorer team advanced several geometric intuitions about where the optimum should lie. Each was tested against the full optimization, and the scorecard strongly endorses the team's physical reasoning.

1. *The departure budget rises sharply for arrivals much beyond ~80,000 years.* **Confirmed.** The required cruise speed grows from $24.14\ \mathrm{km\,s^{-1}}$ at the plane crossing to $30.61\ \mathrm{km\,s^{-1}}$ at 500,000 years, and the impulsive departure Δv from 13.9 to $19.93\ \mathrm{km\,s^{-1}}$: patience stops paying once the star begins to recede.
2. *The star's closest natural approach (~3 ly) is not the Δv optimum.* **Confirmed**, with the epoch refined: the closest approach occurs 27,955 years from launch at 3.152 ly, and intercepting there is the costliest option examined — the short flight time more than cancels the shorter distance.
3. *A tangential intercept near 60,000 years minimizes the cruise speed.* **Confirmed exactly:** $T^* = 58{,}422$ years, where the required v_∞ equals the star's transverse velocity of $23.38\ \mathrm{km\,s^{-1}}$ by the identity of Section 2.2.
4. *The ecliptic-plane crossing minimizes the departure Δv.* **Very nearly right, but not exact:** arriving at the

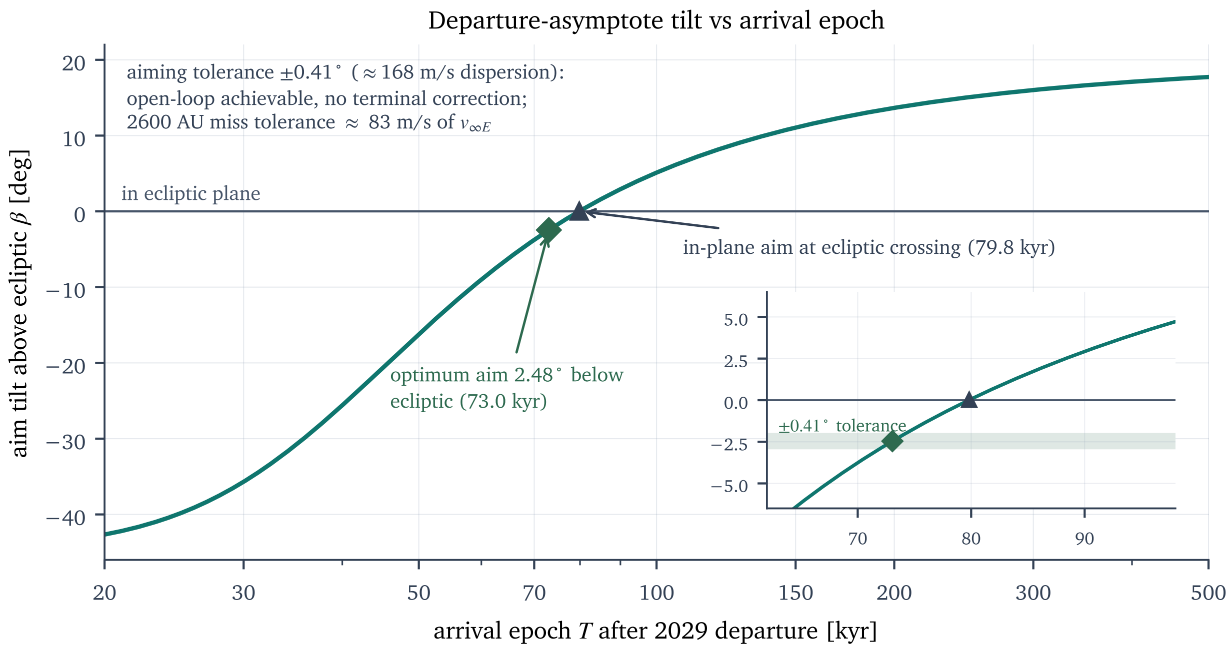

Figure 3. Departure aim tilt β versus arrival time; the inset shows the $\pm 0.41°$ open-loop aiming tolerance band around the optimum. Earth's orbital velocity of $29.78\ \mathrm{km\,s^{-1}}$ is borrowed in-plane only, so the tilt directly prices the out-of-plane component of the asymptote. The tilt passes through zero at the ecliptic-plane crossing (79,786 yr); the joint optimum accepts a $2.48°$ tilt to intercept the star several thousand years earlier, when it is closer and slower to chase.

crossing epoch of 79,786 years costs less than $0.1\ \mathrm{km\,s^{-1}}$ more than the true optimum. The intuition correctly identifies the dominant effect (avoiding the out-of-plane penalty) but misses the small gain from intercepting earlier.

5. *The true optimum lies near 73,000 years, aimed a few degrees off plane.* **Confirmed almost exactly:** 73,012 years, $2.48°$ below the ecliptic.

2.6 Miss Tolerance and Aiming Dispersion

The mission requires only a pass within $2600\,\mathrm{AU}$ of the barycenter, not a true intercept. In energy terms this tolerance is nearly worthless: it reduces the required Earth-relative excess by about $83\ \mathrm{m\,s^{-1}}$, which the escape-energy term dilutes to roughly $71\ \mathrm{m\,s^{-1}}$ of impulsive Δv. The miss tolerance relaxes the *targeting* problem, not the energy problem.

As a targeting relaxation, however, it is decisive. Hitting a $2600\,\mathrm{AU}$ sphere at interstellar range permits an asymptote pointing error of $0.41°$, equivalent to about $168\ \mathrm{m\,s^{-1}}$ of departure velocity dispersion. The sphere must absorb target-ephemeris error as well as execution dispersion, so the ephemeris term is budgeted first. Propagated to the 73,012-year encounter, the proper-motion uncertainty of the adopted astrometry [1] produces a cross-track dispersion of about $381\,\mathrm{AU}$ (1σ; some $147\,\mathrm{AU}$ using the cross-check orbital solution of [2]), consuming roughly 15% of the $2600\,\mathrm{AU}$ sphere. Section 6 carries this ephemeris covariance jointly with the execution error sources through a full Monte Carlo dispersion analysis and confirms that the combined budget closes. The radial-velocity uncertainty produces a much larger along-track term — about $4600\,\mathrm{AU}$ (1σ; some $1500\,\mathrm{AU}$ with the cross-check values) — but an along-track error does not widen the miss distance: it merely shifts the encounter epoch by roughly ±900 years, and the pass still occurs, at negligible cost given the flat optimum of Section 2.4. Netting the cross-track ephemeris term against the sphere thus leaves the great majority of the allowance for execution dispersion, and the combined dispersion budget closes with margin.

Within that residual allowance, the pointing tolerance is expected to be met open-loop — meaning no trajectory correction after escape, with navigation custody retained through the powered escape via star-tracker attitude knowledge and thrust telemetry — so no terminal guidance, no course-correction propellant, and no late-mission contact with the vehicle are required. The execution-dispersion Monte Carlo of Section 6, which carries the target-ephemeris covariance alongside the thrust and attitude errors, bears this expectation out; the requirement set is therefore consistent with a fully expendable, out-of-contact cruise lasting the entire 73,000-year coast.

2.7 The 500,000-Year Horizon Is Not Binding

Finally, the mission's 500,000-year arrival ceiling exerts no pressure on the design. The optimum lies at 73,012 years, deep inside the allowed window, and the requirement *rises* toward the horizon: by 500,000 years the star is receding and the budget has grown to $v_\infty = 30.61\ \mathrm{km\,s^{-1}}$ ($\Delta v = 19.93\ \mathrm{km\,s^{-1}}$). Relaxing the horizon would therefore buy nothing, and tightening it to as little as ~85,000 years would cost less than $0.1\ \mathrm{km\,s^{-1}}$. The binding constraints on this mission are propulsive and programmatic, not temporal; they are taken up in the sections that follow.

3. Departure Energetics from Low Earth Orbit

The intercept geometry of the preceding section fixes the departure requirement: the least expensive trajectory leaves Earth on a hyperbolic asymptote tilted $2.48°$ below the ecliptic plane, with a heliocentric hyperbolic excess speed $v_\infty = 23.64\ \mathrm{km\,s^{-1}}$ and arrival $73{,}012$ years after launch. By choosing the departure epoch so that the in-plane projection of the required velocity aligns with Earth's orbital motion, the spacecraft borrows Earth's 29.78 $\mathrm{km\,s^{-1}}$ of orbital speed in the ecliptic plane; only the small out-of-plane component, applied at an angle to that motion, must be bought propulsively. The resulting Earth-relative hyperbolic excess is $v_{\infty,E} = 18.59\ \mathrm{km\,s^{-1}}$. This section asks what that excess costs from a 400 km low Earth orbit, gravity assists being excluded by the mission ground rules. The answer depends strongly on the thrust regime, and that dependence drives the entire vehicle design.

3.1 The Impulsive Floor

For a burn that is short compared with the orbital period, the Oberth effect applies: propellant spent deep in the gravity well, where the vehicle moves fastest, buys more orbital energy per unit of Δv, because the energy gained from an increment $\mathrm{d}v$ scales with the instantaneous speed at which it is applied. A single impulsive burn at perigee of a 400 km orbit therefore achieves the departure asymptote at the minimum possible cost,

$$\Delta v_{\mathrm{imp}} \;=\; \sqrt{v_{\infty,E}^2 + v_{\mathrm{esc}}^2} \;-\; v_{\mathrm{circ}}, \tag{3}$$

where $v_{\mathrm{circ}} = 7.67\ \mathrm{km\,s^{-1}}$ is the circular orbital speed at 400 km altitude and $v_{\mathrm{esc}} = \sqrt{2}\,v_{\mathrm{circ}} = 10.85\ \mathrm{km\,s^{-1}}$ is the corresponding escape speed. Evaluating Equation 3 at $v_{\infty,E} = 18.59\ \mathrm{km\,s^{-1}}$ gives

$$\Delta v_{\mathrm{imp}} = 13.85\ \mathrm{km\,s^{-1}}. \tag{4}$$

This is the floor for any high-thrust (chemical) departure without gravity assist, and it captures the full Oberth benefit. Two checks support it. An independent replication from first principles recovered the figure — together with the optimal arrival epoch, aim tilt, and Earth-relative excess — with a worst-case discrepancy below 0.5%. And the Fermi Explorer reference analysis [12] agrees to within 0.3%, finding 13.88 $\mathrm{km\,s^{-1}}$ at an arrival epoch of $72{,}800$ years with a $2.4°$ tilt — a small offset consistent with the flatness of the optimum established in the geometry analysis.

3.2 The Low-Thrust Picture

A solar-electric vehicle cannot claim the Oberth benefit at Earth. At thrust accelerations of order $10^{-4}\ \mathrm{m\,s^{-2}}$, escape from low Earth orbit proceeds as a slow spiral over many thousands of revolutions, with thrust applied at essentially all orbital phases. For such spirals the propulsive cost of escape from a circular orbit approaches the circular orbital speed of the starting orbit itself [8] — the well-known penalty of low-thrust escape. From 400 km this is 7.67 $\mathrm{km\,s^{-1}}$; the budget adopted here is 7.6 $\mathrm{km\,s^{-1}}$, expended over approximately 1.5 years.

Table 5. Departure budgets from a 400 km low Earth orbit at the optimal arrival epoch ($73{,}012$ yr), no gravity assist. The heliocentric interval is a reference interval: its upper end is a verified constructive bound, its lower end a best-found anchor (Appendix B).

Propulsion regime	Δv ($\mathrm{km\,s^{-1}}$)	Remarks
Impulsive (chemical)	13.85	single perigee burn; full Oberth
Solar-electric: Earth escape	7.6	spiral, ~1.5 yr; belt transit
Solar-electric: heliocentric	22.9–24.0	perihelion pumping (Section 4); reference interval
certified lower bound	16.56	throttleable class, unconstrained flight time, $r_p \geq 0.42$ AU
Solar-electric: total	30.5–31.6	reference interval

The spiral transits the radiation belts, and the vehicle sizing therefore includes a solar-array degradation factor of 1.25 for this phase.

The deeper consequence is that a slowly spiraling vehicle leaves Earth's sphere of influence with negligible Earth-relative excess speed. The entire heliocentric requirement must then be accumulated after escape, in interplanetary flight, where a second obstacle appears: solar-array power, and hence thrust, falls with distance from the Sun. Section 4 shows that a conventional outward spiral cannot reach the required cruise speed at all, and that a perihelion-pumping trajectory can. With that heliocentric strategy, the total solar-electric departure budget is

$$\Delta v_{\mathrm{SEP}} = \underbrace{7.6}_{\text{Earth escape}} + \underbrace{[22.9,\ 24.0]}_{\text{heliocentric}} = 30.5\text{–}31.6\ \mathrm{km\,s^{-1}}. \tag{5}$$

The heliocentric term is stated as an interval rather than a point value, and its two ends carry different evidentiary weight. The upper end is a verified constructive bound: an explicitly integrated, timestep-refined perihelion-pumping trajectory at an initial thrust acceleration of $2.5 \times 10^{-4}\ \mathrm{m\,s^{-2}}$ reaches the required cruise speed at 23.97 $\mathrm{km\,s^{-1}}$. The lower end, 22.9 $\mathrm{km\,s^{-1}}$, is the best trajectory found by optimal-control analysis at a higher thrust acceleration ($1{\times}10^{-3}\ \mathrm{m\,s^{-2}}$) — a single-method result, from an indirect Pontryagin solution whose direct-collocation cross-check did not converge — extended down to the design thrust level by a monotonicity argument valid for throttleable thrust. It is a best-found reference anchor, not a certified bound. The certified lower bound is 16.56 $\mathrm{km\,s^{-1}}$, proved for the throttleable-thrust class with unconstrained flight time under the design's 0.42 AU minimum solar distance (the impulsive multi-burn limit it supersedes is 16.50 $\mathrm{km\,s^{-1}}$; 15.6 $\mathrm{km\,s^{-1}}$ if the perihelion bound is relaxed to 0.25 AU). Quantitative closure results later in this report are quoted across the reference interval 22.9–24.0 $\mathrm{km\,s^{-1}}$ with an explicit caveat: closure of a 100 kg vehicle from low Earth orbit cannot be excluded if a heliocentric trajectory below the ≈22.3 $\mathrm{km\,s^{-1}}$ crossover exists within the 12-year operational horizon; no trajectory found in this analysis, when referred to the design thrust level, comes within 0.6 $\mathrm{km\,s^{-1}}$ of that threshold. The full argument is given in Appendix B. Table 5 summarizes the departure budgets.

The factor-of-two-plus spread between the impulsive floor and the solar-electric total is not an inefficiency to be engineered away; it is the structural price of low thrust, paid once through the loss of the Oberth effect at Earth and again through the heliocentric climb. Any credible solar-electric concept must therefore be sized against roughly 31 $\mathrm{km\,s^{-1}}$ of total Δv, not against the impulsive floor — a distinction that dominates the mass-closure analysis later in this report.

Both budgets already make full use of the freedom to select the delivery orbit within the low-Earth class. The quoted figures assume the injection plane is chosen to contain the tilted departure asymptote, so no separate plane change is charged, and the altitude lever is weak: raising the starting orbit reduces the spiral cost only in proportion to the circular speed — a few tenths of a $\mathrm{km\,s^{-1}}$ across low-Earth altitudes — while slightly eroding the Oberth benefit of an impulsive burn. The delivery-orbit choice that genuinely matters lies beyond low Earth orbit altogether: a geostationary-transfer-orbit drop-off cuts the electric-propulsion escape cost from 7.6 to 4.24 $\mathrm{km\,s^{-1}}$, and Section 5 shows that this single procurement change is what closes the 100 kg vehicle.

4. Solar-Electric Propulsion: the Power Wall and Perihelion Pumping

4.1 Power Model

Solar-array output falls with the inverse square of heliocentric distance r. The propulsion model used throughout this report takes

$$P(r) \;=\; P_1 \,\min\left[\left(\frac{r_1}{r}\right)^2,\, 4\right], \tag{6}$$

where P_1 is the array output at $r_1 = 1$ AU and the factor-of-four cap, which engages inside roughly 0.5 AU, represents a thermal limit on harvested power. The cap is an assumption rather than demonstrated practice: inner-heliosphere missions of the MESSENGER class survive 0.3–0.47 AU by off-pointing their arrays to *refuse* the elevated flux, not by harvesting it, so a cap of four leans optimistic. Its influence on every pumping result is therefore quantified directly in Section 4.6. At fixed specific impulse and thruster efficiency, thrust is proportional to power, so the available thrust acceleration inherits the same profile. The single sizing parameter is a_0, the initial thrust acceleration at 1 AU and full vehicle mass.

4.2 The Power Wall: Outward-Spiral Saturation

The conventional low-thrust escape strategy thrusts continuously outward from 1 AU. Such a spiral is self-defeating: the vehicle climbs away from its own power source, so thrust decays as it recedes, and the Δv that is still deliverable arrives at ever larger radii and ever lower orbital speeds, where each increment contributes least to orbital energy. The achievable hyperbolic excess therefore saturates as a function of a_0. The baseline analysis finds ceilings of approximately 0, 3.4, and 17.0 $\mathrm{km\,s^{-1}}$ at $a_0 = 1.5$, 5, and $10 \times 10^{-4}\,\mathrm{m\,s^{-2}}$ respectively — all below the 23.38 $\mathrm{km\,s^{-1}}$ cruise floor, the minimum speed any intercept requires, set by the target system's transverse velocity. The saturation curve, confirmed by two independent trajectory integrators that agree to 2.7%, is shown in Figure 4.

On outward spirals the present analysis and the Fermi Explorer reference analysis [12] reach the same qualitative conclusion. The agreement extends to the exit from the wall: the region above $a_0 \approx 1.2\times10^{-3}\,\mathrm{m\,s^{-2}}$ in which Figure 4 shows outward spirals clearing the cruise floor corresponds to the pure-solar closing corner the reference analysis itself identified, above a vehicle specific power of roughly 100 W/kg, and presented as its default high-technology route. The comparison between that route and the present result is taken up in Section 4.7. One caution is warranted, however: the reference analysis's published figure of roughly 20 $\mathrm{km\,s^{-1}}$ for a "realistic low-thrust" departure is a benchmark assumption rather than a model output — the reference trajectory engine itself yields 25.1 $\mathrm{km\,s^{-1}}$ for the outward spiral at the optimal epoch — and the associated vehicle sizing of roughly 500 kg inherits the optimistic figure.

4.3 Perihelion Pumping

Saturation is a property of the outward-spiral trajectory class, not of solar power. Escape requires the vehicle to leave the Sun; nothing requires it to thrust while leaving. A perihelion-pumping trajectory inverts the spiral's logic. Retrograde thrust arcs near 1 AU shed angular momentum, lowering perihelion to 0.42 AU — just inside the radius at which the thermal cap of Equation 6 engages, so that array power at perihelion stands at one to four times its 1-AU rating. Prograde burn arcs are then concentrated near perihelion, where power is highest and where, by the same Oberth logic that favors the perigee burn in the impulsive case, each increment of Δv buys the most orbital energy. Successive revolutions raise aphelion while the perihelion burns repeat, until the orbit reaches escape energy. In effect, the multi-revolution structure recovers for low thrust the perihelion-burn advantage that a single fast pass cannot provide at these thrust levels.

The ingredients of this maneuver are established practice, and it is worth stating which are old and which is new. The energetic advantage of burning deep in a gravity well was identified by Oberth [9], and the two-burn escape that exploits it — lower the periapsis, then burn at periapsis — is classical. Powered-flyby and solar-Oberth (perihelion-burn) escapes built on the same effect are established concepts in the interstellar-precursor literature [11], including in Fermi Explorer's own reference analysis [12]. Splitting a low-thrust maneuver into repeated apsis-centered arcs over many revolutions is likewise standard practice: SMART-1 escaped geostationary-transfer orbit on repeated perigee- and apogee-centered thrust arcs [10], and multi-revolution apsis-arc structures are routine in low-thrust trajectory optimization [7]. What the present analysis

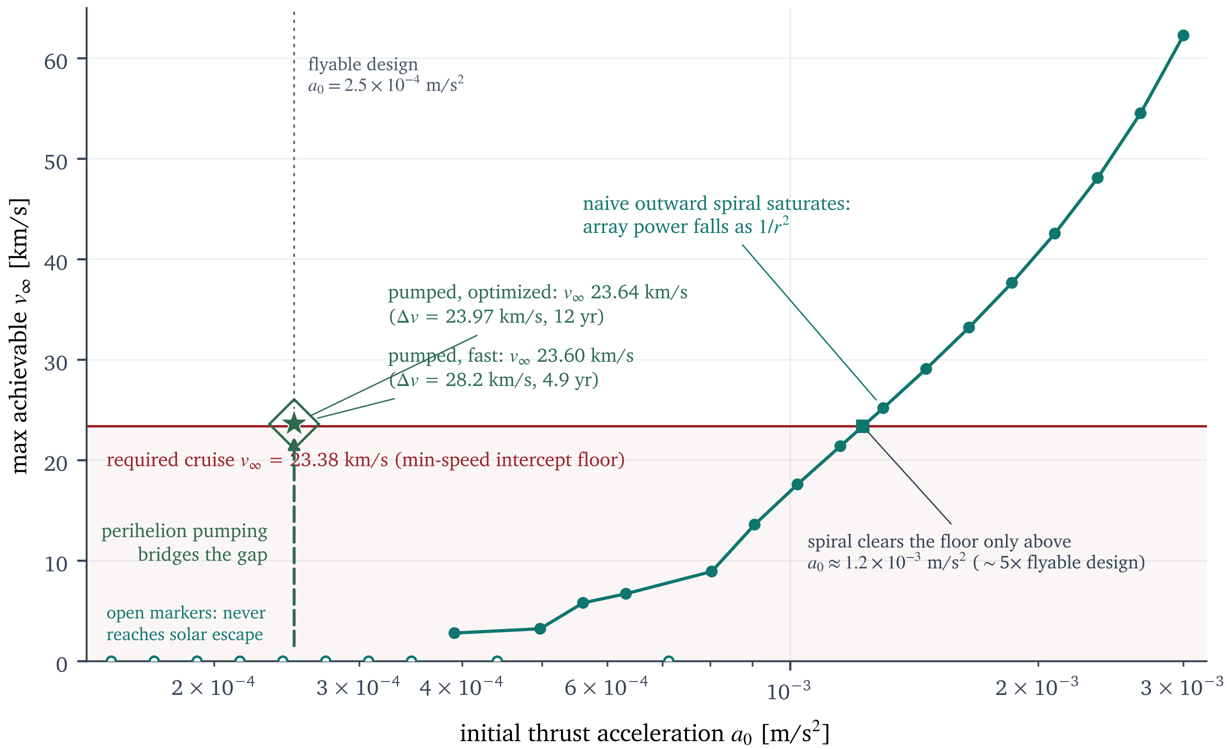

Figure 4. Achievable hyperbolic excess speed versus initial thrust acceleration a_0 (logarithmic axis) for a conventional outward low-thrust spiral. The curve saturates below the 23.38 $\mathrm{km\,s^{-1}}$ cruise floor (horizontal line) at every flyable thrust level, clearing it only above $a_0 \approx 1.2 \times 10^{-3}\,\mathrm{m\,s^{-2}}$ — roughly five times the design value. The two points plotted above the line are perihelion-pumped trajectories at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$: the saturation bounds outward spirals, not solar-electric propulsion as such.

contributes is the quantified closure of this mission class by that combination — cruise speed, Δv, flight time, and mass fraction — under solar power throttled by the $1/r^2$ law of Equation 6.

The optimized profile at $a_0 = 2.5{\times}10^{-4}\,\mathrm{m\,s^{-2}}$ (Figure 5) delivers the required cruise speed of 23.64 $\mathrm{km\,s^{-1}}$ for a total Δv of 23.97 $\mathrm{km\,s^{-1}}$, over a flight time of 12.0 years of which only about 1.3 years is cumulative thrusting; the remainder is coast between burn arcs. The final mass fraction is 0.418, consistent with the rocket equation at the 2800 s specific impulse assumed in the vehicle model, and the minimum solar distance is 0.42 AU. For a 2029 launch, escape from the solar system occurs around 2043. The arrival epoch is unaffected by the added departure time: as established in the geometry analysis, the departure-cost optimum is flat from roughly 70 to 85 thousand years, with a spread below 0.25 $\mathrm{km\,s^{-1}}$. Figure 6 places the phases on a common time axis.

4.4 The Patience Trade

A faster pumping profile exists: a 4.9-year variant reaches essentially the same cruise speed (23.60 $\mathrm{km\,s^{-1}}$) at a cost of 28.2 $\mathrm{km\,s^{-1}}$. Patience is therefore worth approximately 4 $\mathrm{km\,s^{-1}}$ of Δv, and the trade is decisive — the mass-closure analysis of the following section closes only with the patient 12-year profile. Against a cruise of 73,000 years, the additional seven years of powered flight are immaterial to the mission; their real cost is operational, in the duration of powered-flight custody, a point taken up in the cost analysis.

4.5 Bounding the Requirement

The 12-year cost is the best schedule found under an operational limit on powered-flight duration, not a physical floor, and the two claims can be separated quantitatively. Consider the admissible class against which the

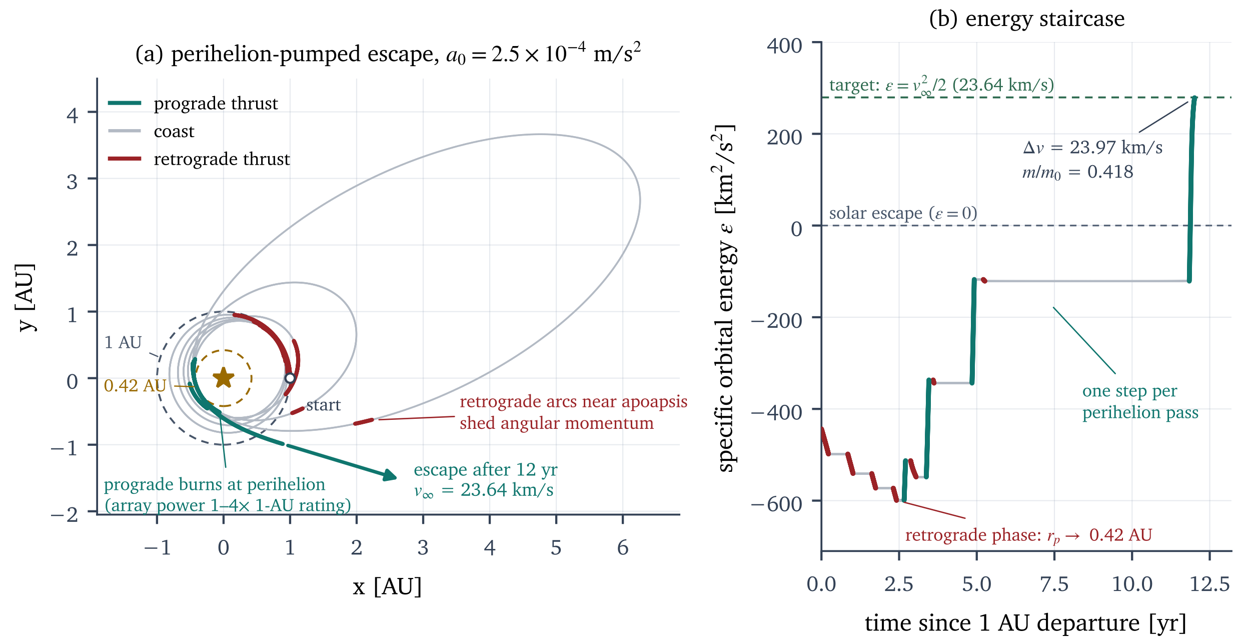

Figure 5. The optimized perihelion-pumping escape at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$. (a) Heliocentric trajectory colored by thrust state (retrograde, coast, prograde), with the 1 AU and 0.42 AU circles dashed. Retrograde arcs near 1 AU lower perihelion; prograde arcs concentrate near perihelion, where array power is at or above its 1-AU rating. (b) Specific orbital energy versus time, colored by thrust state: each prograde perihelion pass adds one step to the energy staircase; the vehicle reaches escape energy at 12.0 yr with Δv = 23.97 $\mathrm{km\,s^{-1}}$ and final mass fraction 0.418.

design is judged: any piecewise-continuous thrust program whose magnitude stays within the power-throttled envelope of Equation 6, whose heliocentric distance never falls below 0.42 AU, and which reaches the required hyperbolic excess of 23.64 $\mathrm{km\,s^{-1}}$, with no constraint on flight time. For every trajectory in this class, the total heliocentric Δv is at least 16.56 $\mathrm{km\,s^{-1}}$. The bound is proof-grade — a closed-form chain of analytic inequalities over the entire class, owing nothing to any optimizer's output — and it improves on the 16.50 $\mathrm{km\,s^{-1}}$ bi-parabolic floor, an impulsive-limit argument, by pricing the finite, power-throttled thrust directly.

The bound is not far from what is achievable. With flight time unconstrained, an explicitly integrated trajectory in the same admissible class reaches the required cruise speed for 18.87 $\mathrm{km\,s^{-1}}$, at the price of 33–58 years of powered flight. The unconstrained optimum is therefore pinned between 16.56 and 18.87 $\mathrm{km\,s^{-1}}$ — the lower end certified, the upper end demonstrated. The roughly 5 $\mathrm{km\,s^{-1}}$ separating that bracket from the 12-year cost is the continuation of the patience trade: it is bought back by program duration alone, so the 12-year figure is a schedule-driven quantity, not a limit set by the dynamics. One consequence deserves emphasis. The 18.87 $\mathrm{km\,s^{-1}}$ trajectory lies below the 22.26 $\mathrm{km\,s^{-1}}$ crossover at which a 100 kg vehicle departing from low Earth orbit would close (Section 5), so that concept is excluded by the 12-year schedule, not by the physics of the trajectory class.

On the 12-year horizon itself, the closure analysis carries the requirement as the reference interval 22.9–24.0 $\mathrm{km\,s^{-1}}$, and the two ends have different standing. The upper end rests on the verified constructive schedule: the integrated trajectory of Figure 5 costs 23.97 $\mathrm{km\,s^{-1}}$ at its production timestep, 23.98 $\mathrm{km\,s^{-1}}$ when quoted conservatively across successive timestep refinements, and the closure analysis rounds up to 24.0. The lower end, 22.9 $\mathrm{km\,s^{-1}}$, is a conservative reference derived from an optimal-control extremal computed at four times the design thrust acceleration — where the small revolution count keeps the problem numerically tractable; the extremal costs 22.66 $\mathrm{km\,s^{-1}}$ at the required cruise speed and 22.94 $\mathrm{km\,s^{-1}}$ at a slightly higher target variant — and carried down to the design level by monotonicity of the optimum in available thrust. That value must be characterized honestly: it comes from a single method, an indirect solution of the necessary con-

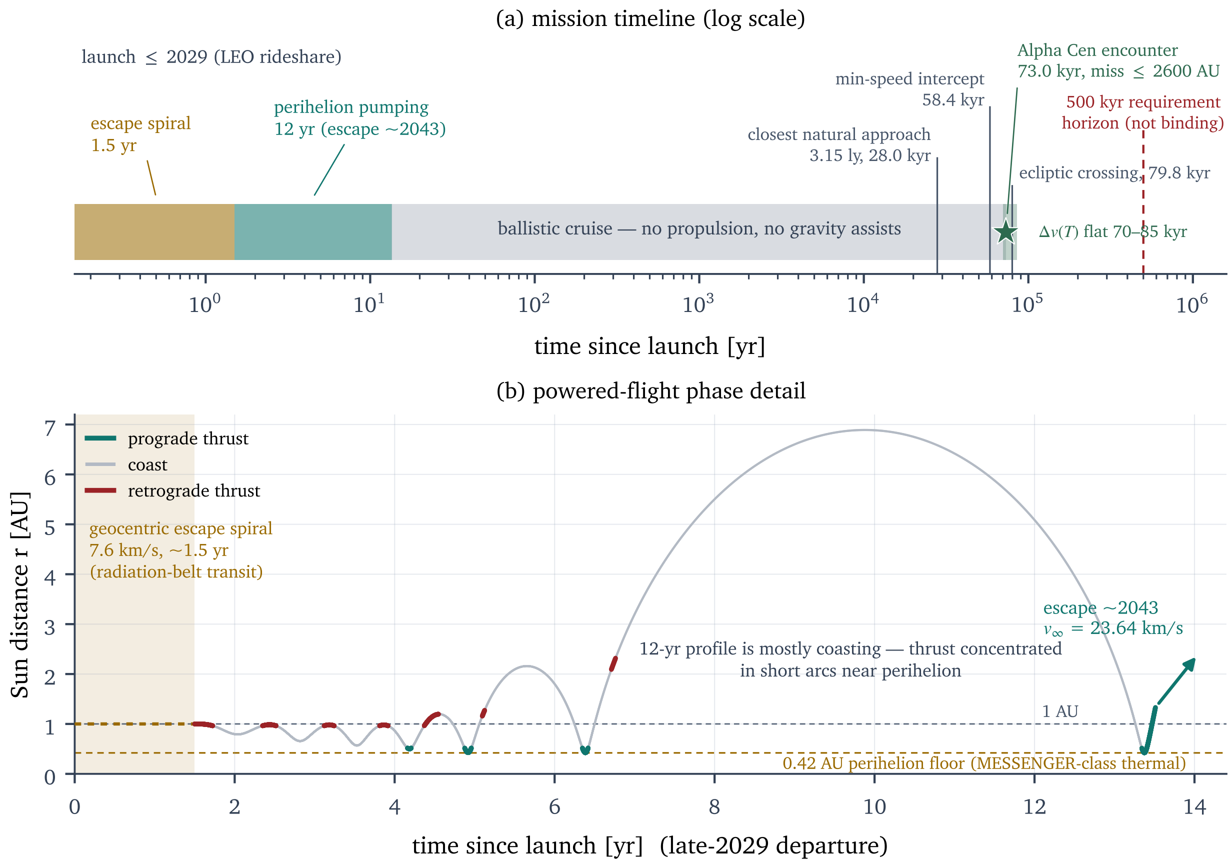

Figure 6. (a) Mission phases on a logarithmic time axis: Earth-escape spiral (~1.5 yr, 7.6 $\mathrm{km\,s^{-1}}$), heliocentric perihelion pumping (12 yr, 23.97 $\mathrm{km\,s^{-1}}$), and ballistic interstellar coast (73,000 yr) to the flyby; more than 98% of the mission is unpowered coast. (b) Powered-flight detail: heliocentric distance versus time, colored by thrust state, showing the escape spiral near 1 AU and pumping arcs down to the 0.42 AU perihelion floor.

ditions of optimal control (Pontryagin costate integration with shooting on the boundary conditions), and the intended cross-check by an independent direct collocation solver failed to converge despite repeated attempts at successively finer meshes. The anchor is therefore the best solution found by one converged method — not a value confirmed by two independent ones, and not a bound on the true optimum in either direction. (The extremal problem also carried no explicit perihelion floor; the computed trajectory nonetheless stays outside 0.6 AU, well clear of the 0.42 AU constraint.) Figure 7 lays out the resulting two-class structure.

4.6 Sensitivity to the Thermal Power Cap

Every quantitative pumping result in this section is computed under the factor-of-four cap of Equation 6. Because inner-heliosphere flight practice manages the 0.3–0.5 AU thermal environment by off-pointing arrays rather than by harvesting elevated flux, the cap was tested directly: the pumping schedule was re-optimized at the design $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$ with the cap reduced, all other model parameters held fixed. With power capped at three times the 1-AU rating, the optimized escape requires 24.10 $\mathrm{km\,s^{-1}}$ over 19.7 years; capped at two, 25.88 $\mathrm{km\,s^{-1}}$ over 15.1 years. Both re-optimized trajectories reach escape, and both are timestep-stable (half-timestep reruns change Δv by less than 0.3%).

The mechanism therefore survives a halving of the assumed cap at a cost of approximately 1.9 $\mathrm{km\,s^{-1}}$ of Δv and a few years of flight time; it does not fail. The two cases bound the physical situation: cap 4 is optimistic, since it assumes the arrays harvest at four times rating in an environment MESSENGER-class practice survives

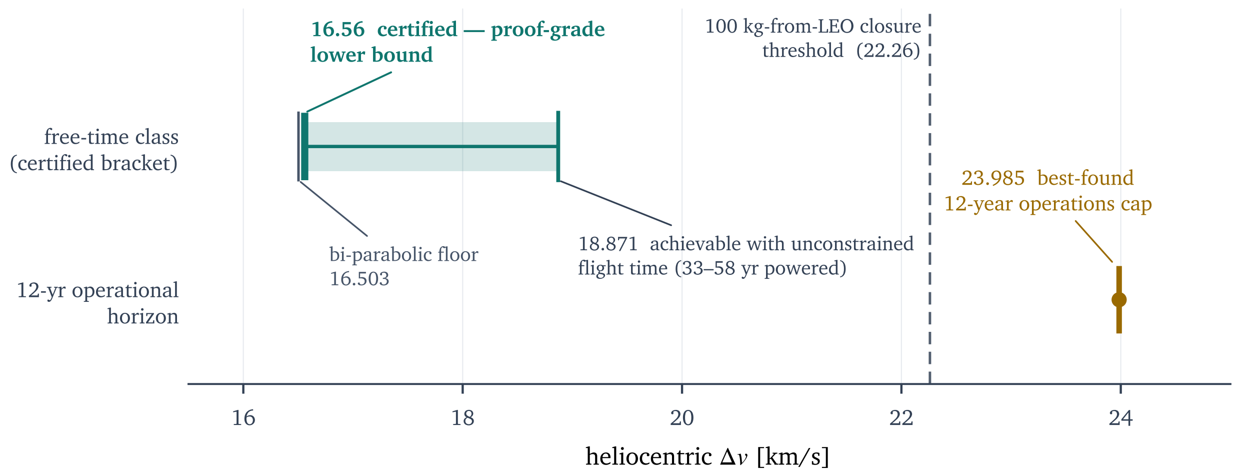

Figure 7. Two-class structure of the heliocentric Δv requirement. Upper row: with flight time unconstrained, the optimum is bracketed between a certified, proof-grade lower bound of 16.56 $\mathrm{km\,s^{-1}}$ and a demonstrated trajectory at 18.87 $\mathrm{km\,s^{-1}}$ (33–58 years of powered flight); the 16.50 $\mathrm{km\,s^{-1}}$ bi-parabolic floor that the certified bound supersedes is marked below the bracket. Lower row: on the 12-year operational horizon, the best schedule found costs 23.98 $\mathrm{km\,s^{-1}}$ — a best-found value from a converged search, not a bound on the true optimum. The dashed line marks the 22.26 $\mathrm{km\,s^{-1}}$ crossover below which a 100 kg vehicle from low Earth orbit would close: the unconstrained-flight-time trajectory lies below it, so that concept fails on schedule, not on physics.

by refusing that flux, while cap 2 is conservative. The mass-closure consequence falls to Section 5: the 100 kg vehicle departing from a geostationary-transfer-orbit drop-off retains closure at the cap-2 Δv requirement, with a reduced margin, and Section 5 traces the closing boundary.

4.7 Consequences

Perihelion pumping changes the feasibility verdict for this mission class, and the change is best stated against what the Fermi Explorer reference analysis actually concluded [12]. That analysis found, as this one does, that outward spirals saturate below the required cruise speed at every flyable thrust level, and it identified two ways past the wall. First, a pure-solar closing corner at high specific power: above a vehicle specific power of roughly 100 W/kg, outward-spiral solar-electric closure becomes available, and the reference analysis presented that high-specific-power route as its default. Second, a solar-Oberth architecture — lower the perihelion, then burn at perihelion — whose electric variant it rejected on the grounds that a single perihelion pass is too short for ion-class thrust to deliver meaningful Δv. On that basis it recommended nuclear-electric power for a near-term vehicle.

What the present result adds is the multi-revolution form of the perihelion burn, which removes the single-pass objection: the Δv that cannot be delivered in one pass is delivered over many, with the burn arcs re-concentrated at perihelion on every revolution. The required cruise speed is then reachable with solar power alone at today's ~60 W/kg system-level arrays — a trajectory-class result, not a component advance — and no nuclear reactor, no chemical kick stage, and no gravity assist is required. The amendment to the reference analysis is therefore precise and narrow: its power-wall finding stands; its high-specific-power closing corner remains valid but is no longer necessary; and its single-pass objection to the electric perihelion burn, correct for a single pass, does not apply to the multi-revolution schedule.

What solar-electric propulsion does demand, beyond the patience of the 12-year profile, is the operating point of the vehicle model: a specific impulse of at least 2800 s at total efficiency near 0.55 in a 0.5–1.0 kW string, with roughly 17,000–21,000 hours of single-string firing and 64–103 kg of xenon throughput. No flown Hall thruster covers that point — flown Hall specific impulse tops out near 1,650 s at this power class

— but it is bracketed by demonstrated gridded-ion hardware: NSTAR's flight-documented TH3–TH4 throttle points (0.91–1.02 kW, 2,843–2,942 s, efficiency 0.527–0.554) sit on the required corner, with 30,352 hours and 235 kg of single-string ground life and 16,265 hours of single-unit flight, and with dual-string redundancy and the standard 1.5× throughput qualification margin available as mitigations. The propulsion requirement is thus met by a demonstrated component class — gridded ion — at the price of a mission-specific qualification campaign; the full life and throughput budget is carried in Section 5.

The pumping result carries independent verification of a strong form. A separate implementation, developed from scratch without access to the baseline trajectory code, constructed its own pumping schedule and reached a cruise speed of 23.60 $\mathrm{km\,s^{-1}}$ at a Δv of 28.2 $\mathrm{km\,s^{-1}}$ — independently confirming that the power wall is broken at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$. The optimized 12-year schedule itself was timestep-refined, with the converged Δv changing by 0.06%. The outward-spiral ceiling that both trajectories defeat was, as noted above, reproduced by two independent integrators. Appendix B details the models, the control policy, and the bracketing argument.

4.8 Limitations

Three limitations of the trajectory analysis deserve plain statement. First, the baseline heliocentric optimization is planar; the $2.48°$ aim tilt enters through the Earth-relative excess. The three-dimensional cost of acquiring that tilt was measured directly: a three-dimensional trajectory optimization under the same perihelion floor, thrust envelope, and 12-year limit, constrained to place the departure asymptote within $0.04°$ of the target latitude $2.48°$ below the ecliptic, delivers the tilt for an out-of-plane increment of 0.58 $\mathrm{km\,s^{-1}}$ over the planar solution — 24.56 $\mathrm{km\,s^{-1}}$ against 23.98 $\mathrm{km\,s^{-1}}$, both quoted conservatively across successive timestep refinements (Figure 8). The measured increment falls where analytic reasoning places it: between the 22 m/s cost of steering the tilt within the escape burns at no geometric penalty and the 1.02 $\mathrm{km\,s^{-1}}$ cost of executing it as a separate plane change. Because the planar and three-dimensional totals are each the best solution found by a converged search rather than a certified optimum, their difference is a measured estimate of the out-of-plane cost, not a bound on it in either direction; its consequence for mass closure is carried in Section 5. Second, the repeated passes at 0.42 AU place the vehicle in a thermal environment requiring qualification comparable to inner-solar-system missions of the MESSENGER class; this is a cost and schedule item, not a physics risk, and it is priced in the cost analysis, while the companion assumption that the arrays continue to harvest at up to four times their 1-AU rating in that environment is tested directly in Section 4.6. Third, the mechanism has a genuine sensitivity: pumping fails to reach the required cruise speed below $a_0 \approx 2.25 \times 10^{-4}\,\mathrm{m\,s^{-2}}$, so an array degradation of 10–20% over the pumping phase would erase the margin of a 100 kg-class vehicle starting from low Earth orbit. This sensitivity, and the departure options that retire it, are treated quantitatively in the mass-closure analysis.

5. Vehicle Mass Closure

The trajectory analysis fixes what the propulsion system must deliver. For a solar-electric departure from a $400\,\mathrm{km}$ low Earth orbit (LEO), the total Δv is the sum of a $7.6\,\mathrm{km\,s^{-1}}$ Earth-escape spiral (roughly 1.5 years, transiting the radiation belts) and a heliocentric perihelion-pumping phase whose cost spans the reference interval of 22.9–$24.0\,\mathrm{km\,s^{-1}}$, for a total of 30.5–$31.6\,\mathrm{km\,s^{-1}}$. The two ends of that interval carry different pedigrees. The upper end is a verified constructive bound: an explicitly integrated, timestep-refined pumping schedule delivers the cruise speed for $24.0\,\mathrm{km\,s^{-1}}$ at the design thrust. The lower end, $22.9\,\mathrm{km\,s^{-1}}$, is the best trajectory found by optimal-control analysis at a higher thrust acceleration ($a_0' = 10^{-3}\,\mathrm{m/s^2}$; single-method — an indirect Pontryagin solution whose direct-collocation cross-check did not converge), extended down to the design thrust by a monotonicity argument valid for throttleable engines; it is a best-found reference anchor, not a certified bound. The certified lower bound lies well below both, at $16.56\,\mathrm{km\,s^{-1}}$, proved for the throttleable class with unconstrained flight time and the $0.42\,\mathrm{AU}$ perihelion bound (the impulsive multi-burn limit it supersedes is $16.50\,\mathrm{km\,s^{-1}}$; $15.6\,\mathrm{km\,s^{-1}}$ at $0.25\,\mathrm{AU}$). This section asks whether a vehicle of the concept class—approximately $100\,\mathrm{kg}$ at launch, predominantly xenon propellant, carrying a payload of at least $1\,\mathrm{kg}$—can absorb that budget within its own mass. It nearly can from LEO; it comfortably can from a geostationary transfer orbit (GTO).

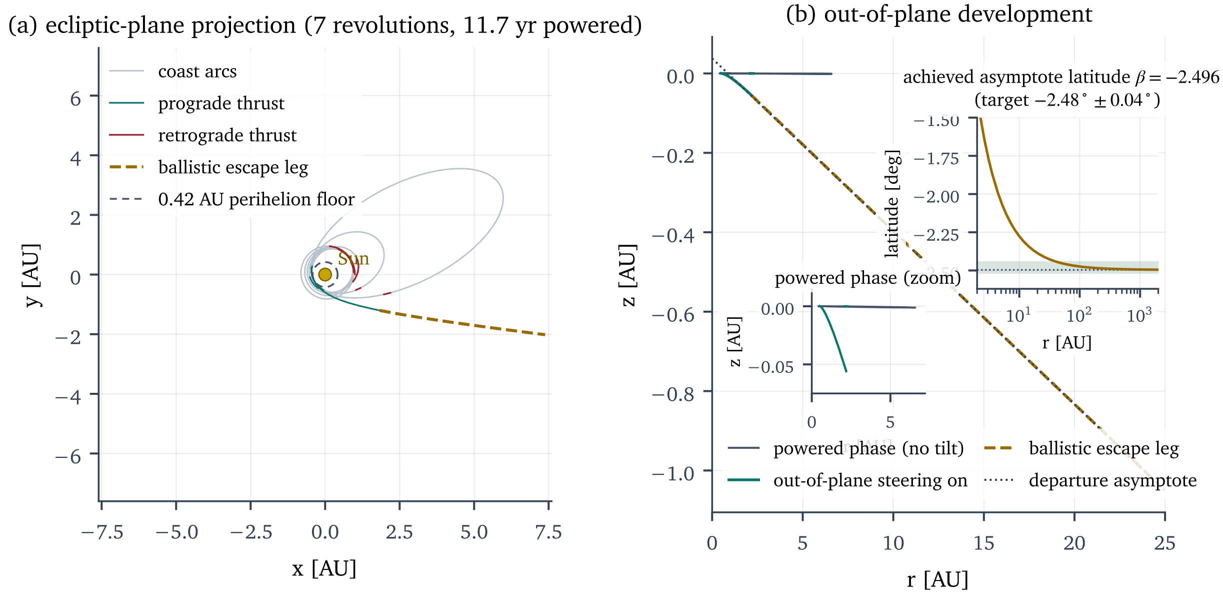

Figure 8. The three-dimensional perihelion-pumping escape at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$, with the departure asymptote directed $2.48°$ below the ecliptic (achieved within the $0.04°$ aiming gate). Out-of-plane excursions stay below 0.06 AU and are concentrated late in the escape, after the orbit is already hyperbolic. The measured cost of the tilt is 0.58 $\mathrm{km\,s^{-1}}$ over the planar schedule, for a total of 24.56 $\mathrm{km\,s^{-1}}$.

5.1 Mass Model

The mass-closure model combines the rocket equation with linear subsystem scalings; the full equations and the closure-boundary algebra are given in Appendix C. For a wet (launch) mass m_w and specific impulse I_{sp}, the propellant mass is $m_p = m_w\,[1 - \exp(-\Delta v_{tot}/v_e)]$, where $v_e = g_0 I_{sp}$ is the exhaust velocity and Δv_{tot} the total budget above. The thrust acceleration at 1 AU and wet mass, $a_0 \approx 2.5 \times 10^{-4}\,\mathrm{m/s^2}$, is sized so that the cumulative thrusting time fits within a four-year allowance—consistent with the optimized pumping profile, which accumulates roughly 1.3 years of thrust arcs over its twelve-year flight. The thrust $F = a_0 m_w$ then sets the 1 AU electrical power through the jet-power relation $P = F v_e/(2\eta)$ with thruster efficiency $\eta = 0.55$; closing designs require 0.5–1.0 kW at 1 AU. The subsystems scale linearly with that power and with propellant load: the solar array at 60 W/kg system-level specific power, penalized by a factor of 1.25 for radiation degradation during the belt transit of a LEO start; power processing at 6 kg/kW; xenon tankage at 12% of propellant mass; and a fixed 1 kg payload plus 8 kg bus; the expendable, out-of-contact posture permits a minimal antenna, whose mass is subsumed in the bus allowance. A design closes when the sum of these masses does not exceed m_w; margins below are quoted in percentage points of wet mass.

5.2 Closure Results

Figure 9 summarizes the closure results at both ends of the heliocentric reference interval. From LEO, the 100 kg concept misses closure by 1.0–2.8% of wet mass everywhere in the interval—a shortfall of one to three kilograms, never more. Because the interval's lower end is a reference anchor rather than a certified bound, this verdict is conditional: closure of a 100 kg vehicle from LEO cannot be excluded if some trajectory below the $\approx 22.3\,\mathrm{km\,s^{-1}}$ crossover exists within the 12-year operational horizon. None of the trajectories found in this analysis, referred to the design thrust level, comes within $0.6\,\mathrm{km\,s^{-1}}$ of that crossover. Equivalently, the minimum closing wet mass from LEO is 113.5–146 kg across the interval. Growing the vehicle to 150 kg closes the budget across the entire interval in this planar accounting, with margins of +0.2 to +2.0 points (revisited under the measured out-of-plane cost below). The closure arithmetic was verified algebraically by an independent analysis, which reproduced the boundary at 113.3 kg—agreement to better than 0.2% with the design-sweep value.

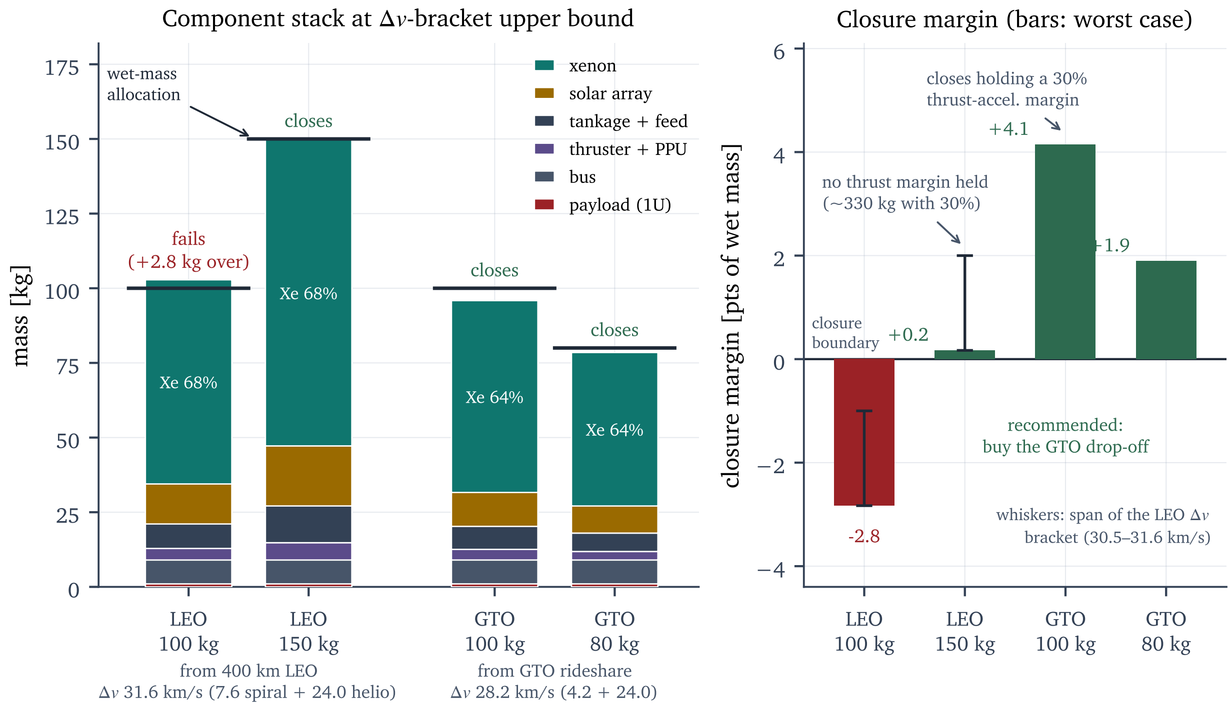

Figure 9. (a) Stacked component masses (payload, bus, thruster and power processing, tankage, solar array, xenon) at the conservative upper end of the heliocentric Δv reference interval; the horizontal cap on each bar marks the available wet mass. (b) Closure margin in percentage points of wet mass; whiskers span the reference Δv interval for the low-Earth-orbit cases. A 100 kg vehicle started from low Earth orbit exceeds its cap everywhere in the interval; 150 kg closes, and both 100 kg and 80 kg close from a geostationary transfer orbit.

Two features of the concept survive scrutiny unchanged. First, the low-Earth-orbit closing designs are 67–68% propellant and the transfer-orbit designs 63–64%—roughly two thirds of the vehicle, essentially all xenon: the concept's proportions are correct, and only its total is short from LEO. Second, closure is possible only with the patient pumping profile. The twelve-year optimized schedule delivers the heliocentric phase for $23.97\,\mathrm{km\,s^{-1}}$; a fast 4.9-year variant exists but costs $28.2\,\mathrm{km\,s^{-1}}$, and no examined vehicle closes at that price. The patience trade is worth roughly $4\,\mathrm{km\,s^{-1}}$ of Δv, and against a 73,000-year cruise the additional seven years of powered flight are immaterial to the mission.

Closure also constrains the thruster, through a mechanism worth stating precisely: the specific-impulse requirement arises from burn-time sizing against the pumping-feasibility thrust floor ($a_0 \gtrsim 2.25 \times 10^{-4}\,\mathrm{m/s^2}$), not from the rocket equation alone—lowering I_{sp} raises the propellant load while the burn-time allowance holds the thrust acceleration, and with it the power chain, up. For the LEO start, a specific impulse of at least 2800 s is required. From GTO the requirement relaxes: a 2400 s thruster operated at the pumping-feasibility thrust floor closes 100 kg, with the closure boundary at 93.5 kg at the conservative end of the reference interval. The 2800 s figure is therefore a LEO-start requirement rather than a mission-wide one—a distinction that matters for component selection, because the high specific impulse is what pushes the operating point beyond the flown Hall-thruster envelope.

The full operating point—$I_{sp} \geq 2800$ s at total efficiency $\eta \approx 0.55$ in a 0.5–1.0 kW string, with roughly 17,000–21,000 hours of cumulative single-string firing and 64–103 kg of xenon throughput—is not covered by any flown Hall thruster: flown Hall performance tops out near 1,650 s at this power class, more than 1,100 s short on specific impulse and roughly a factor of two short on demonstrated single-string hours. It is, however, bracketed by demonstrated gridded-ion hardware. The NSTAR engine's flight-documented TH3–TH4 throttle points (0.91–1.02 kW, 2,843–2,942 s, efficiency 0.527–0.554) sit essentially on the required corner, and the same device demonstrated 30,352 hours and 235 kg of xenon throughput in single-string ground life testing, with 16,265 hours of single-unit flight operation. Dual-string redundancy, standard practice for missions of this class, halves the per-string firing-time and throughput demands, and the demonstrated ground records already exceed the standard 1.5× throughput qualification margin applied to this mission's loads. One mass consequence is booked explicitly: a gridded-ion string is heavier than the 6 kg/kW assumed for power processing in the mass model; at 10–15 kg/kW, a string adds 3–9 kg at 0.5–1.0 kW and shifts the closure boundaries by a few kilograms. The propulsion system is therefore best characterized as a demonstrated component class (gridded ion), requiring a mission-specific qualification campaign.

Starting from a GTO rideshare changes the picture decisively. The departure Δv falls from 7.6 to $4.24\,\mathrm{km\,s^{-1}}$, and at the conservative (upper) end of the heliocentric reference interval an 80 kg vehicle closes with +1.9 points of margin, 100 kg with +4.1, and 150 kg with +7.1; margins are larger still at the lower end of the interval. The full closure grid appears in Appendix C.

5.3 Sensitivities

Thrust acceleration and array degradation. This is the sharpest risk in the design. The pumping strategy fails to reach the required cruise speed below $a_0 \approx 2.25\times10^{-4}\,\mathrm{m/s^2}$, so a 10–20% array degradation beyond the budgeted allowance erases the LEO-start closure margin entirely. Resizing the vehicle to carry a prudent 30% thrust-acceleration uplift (the burn-time-sized a_0 multiplied by 1.3, with thrust and the power chain rescaled accordingly) moves the closing boundaries sharply enough that the two ends of the reference interval must be quoted separately. From LEO, the minimum closing wet mass rises to 312 kg at the favorable end and 910 kg at the conservative end—three to nine times the concept. From GTO, the boundary rises to 83 kg at the favorable end and 104 kg at the conservative end: a 100 kg GTO vehicle carries the full margin at the favorable end but misses by approximately 4 kg at the conservative end, while a ∼110 kg GTO vehicle carries the margin across the entire interval. The two architectures are therefore not close substitutes: the GTO start remains within a few kilograms of closure under a stress that drives the LEO start to several times its size, though at 100 kg it does not carry the full 30% margin everywhere in the interval.

Out-of-plane departure. The optimal departure asymptote is tilted 2.48° below the ecliptic. A planar treatment can only bound the cost of acquiring this tilt, between 22 m/s (if steered within the escape burns) and $1.02\,\mathrm{km\,s^{-1}}$ (worst case, executed as a separate plane change). A genuinely three-dimensional re-optimization of the pumping profile measures it directly: the out-of-plane increment is $0.58\,\mathrm{km\,s^{-1}}$ (578 m/s), inside the planar band but far above its low end. Because both the planar and three-dimensional figures are the best trajectories found by a heuristic search rather than certified optima, their difference is not a bound on the true three-dimensional cost in either direction; it is the measured cost of the best profile found. Propagated through the closure algebra (by linear interpolation of the closure scan, not a re-run of the mass model), the verdicts sharpen. The 100 kg GTO vehicle closes at both ends of the heliocentric Δv reference interval, with +5.1 points of margin at the favorable end and +3.1 at the conservative end, and it retains +2.3 points even at the band's $1.02\,\mathrm{km\,s^{-1}}$ worst case. The 150 kg LEO fallback closes at the favorable end (+1.0 points) but fails at the conservative end (−0.8 points). The fragility the bounding analysis flagged—the conservative-end LEO closure boundary moving from approximately 146 to 198 kg—is now quantified: any out-of-plane increment above roughly $0.10\,\mathrm{km\,s^{-1}}$ eliminates the 150 kg LEO margin at the conservative end of the interval, and the measured value is more than five times that threshold. Steering the plane change within the escape burns—the maneuver the LEO architectures were counting on to stay near the 22 m/s low end—costs $0.58\,\mathrm{km\,s^{-1}}$ for the best trajectories found. On this axis the ranking is confirmed and sharpened: GTO is robust and LEO is not, and the measurement strengthens the recommendation of a GTO drop-off (Section 5.4). Figure 10 summarizes the re-evaluated margins.

Thermal environment. The pumping profile's perihelion of 0.42 AU places repeated passes inside the orbit of Mercury. The vehicle requires thermal qualification of the class demonstrated by the MESSENGER mission; this is established engineering, but it is priced as a qualification burden in the cost analysis rather than treated as free. The pumping numbers also assume solar power harvested up to four times the 1 AU rating near perihelion (Section 4.6); with the cap halved to two, the re-optimized heliocentric requirement rises to $25.88\,\mathrm{km\,s^{-1}}$ and the GTO-100 closure boundary moves to approximately 98 kg—still closing.

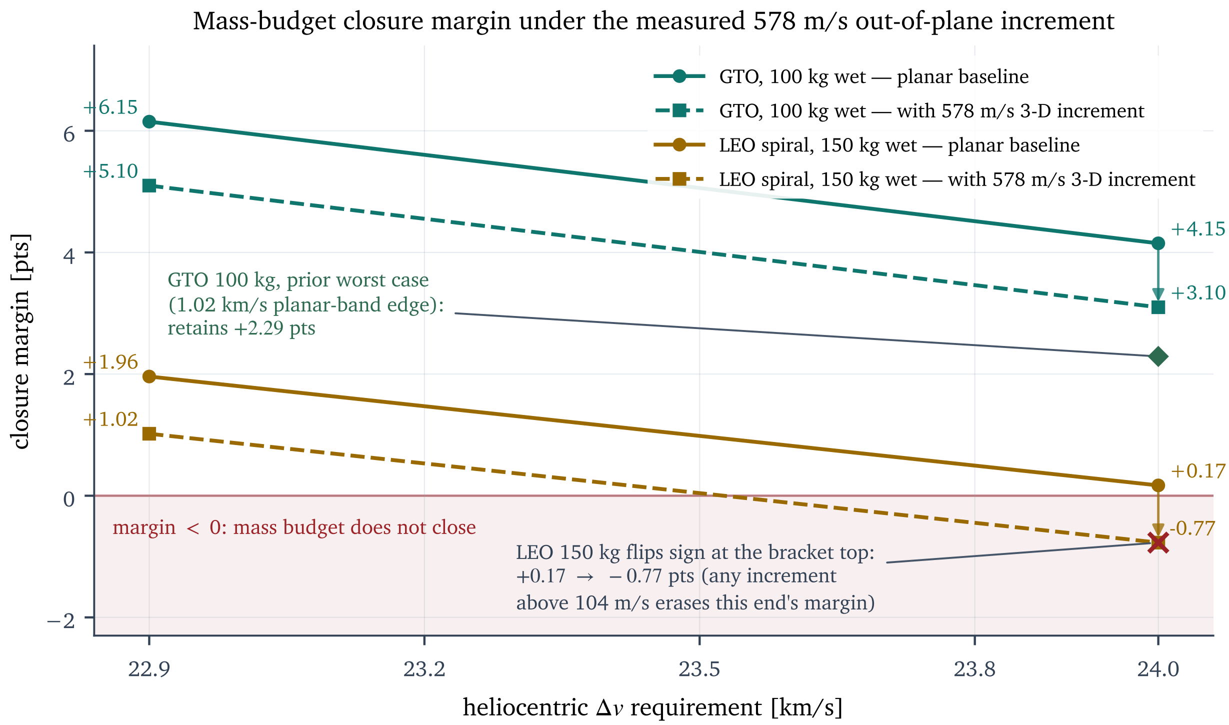

Figure 10. Closure margins re-evaluated with the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane departure cost, in percentage points of wet mass at both ends of the heliocentric Δv reference interval. The $100\,\mathrm{kg}$ vehicle from geostationary transfer orbit closes at both ends (+5.1 and +3.1 points, and +2.3 even at the planar band's $1.02\,\mathrm{km\,s^{-1}}$ worst case); the $150\,\mathrm{kg}$ low-Earth-orbit fallback closes at the favorable end (+1.0) but fails at the conservative end (−0.8). Any increment above roughly $0.10\,\mathrm{km\,s^{-1}}$ erases the conservative-end margin of the $150\,\mathrm{kg}$ low-Earth-orbit vehicle; the measured value is more than five times that threshold.

5.4 Recommendation

The recommended change is to procure a GTO drop-off in place of the specified low-Earth-orbit rideshare. The case is fourfold. First, it preserves the concept: $100\,\mathrm{kg}$ closes with +4.1 points of margin at the conservative end of the Δv reference interval, while the LEO start misses everywhere in the interval on the trajectories found. Second, it is robust where the LEO start is fragile: under a 30% thrust-acceleration uplift the GTO closure boundary stays within a few kilograms of the $100\,\mathrm{kg}$ concept (83–$104\,\mathrm{kg}$ across the interval, versus 312–$910\,\mathrm{kg}$ from LEO), and the GTO start absorbs the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane departure cost that defeats the $150\,\mathrm{kg}$ LEO fallback at the conservative end of the interval. Third, it nearly halves the departure budget (7.6 to $4.24\,\mathrm{km\,s^{-1}}$) and shortens the radiation-belt transit that drives array degradation. Fourth, GTO rideshares are a routine commercial service, so the deviation from the delivery specification is a matter of procurement rather than technology. If the drop-off cannot be changed, the fallback is a $150\,\mathrm{kg}$ vehicle from LEO, which closes across the entire reference interval before the out-of-plane cost is charged but is margin-fragile: it carries no allowance against array degradation and, once the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane departure cost is included, fails to close at the conservative end of the interval.

6. Targeting Confidence Under Execution and Ephemeris Dispersion

The aiming analysis of Section 2.6 established that the $2600\,\mathrm{AU}$ encounter sphere permits an open-loop asymptote pointing error of $0.41°$ — equivalent to approximately $168\,\mathrm{m\,s^{-1}}$ of departure velocity dispersion — and budgeted a one-sigma cross-track target-ephemeris term of approximately $381\,\mathrm{AU}$ against it. That budget was

Table 6. Error sources and one-sigma levels for the dispersion Monte Carlo. The ephemeris levels are published catalog uncertainties; the three execution-error levels are engineering assumptions — plausible for a tracked, telemetry-closed departure but not flight-measured — and Section 6.4 quantifies how the verdict depends on them.

Source	One-sigma level	Basis
Target ephemeris	parallax ± 0.61 mas; proper motion $\pm 3.9\ \mathrm{mas\,yr^{-1}}$ per component; radial velocity $\pm 0.30\ \mathrm{km\,s^{-1}}$	published uncertainties [1]; [2] as variant (± 0.38 mas, $\pm 1.5\ \mathrm{mas\,yr^{-1}}$, $\pm 0.10\ \mathrm{km\,s^{-1}}$)
Thrust-magnitude knowledge	$50\ \mathrm{m\,s^{-1}}$ on the achieved departure speed	engineering assumption (post-cutoff velocity-trim residual)
Asymptote pointing	$0.10°$ per axis, two transverse axes	engineering assumption (attitude/thrust-vector alignment residual)
Post-departure residual velocity	$25\ \mathrm{m\,s^{-1}}$ per axis, isotropic	engineering assumption (launch-window and departure-model residual)

assembled term by term. This section closes the question it raised: with all error sources acting together, what is the probability that the probe actually passes inside the sphere? The answer, from a 20,000-draw seeded Monte Carlo simulation reproduced independently by a first-principles analytic error-propagation model, is 0.982 under the stated error budget. That number is a simulation result conditioned on explicit engineering assumptions, not a flight-measured quantity; this section presents both the result and the conditions under which it degrades.

6.1 Method

Each Monte Carlo draw perturbs the departure state and the target's motion, then evaluates the encounter in closed form. The target barycenter is propagated in uniform straight-line motion — valid here to well under the encounter-sphere scale, since the bounded Galactic-tide perturbation over the flight is about 1.2 AU (Section 2.1) — and the probe cruises on a straight line at its perturbed asymptotic velocity, so the closest approach to the moving target has an exact analytic solution per draw. The aim point is the target barycenter at the 73,012-year encounter epoch, computed from the adopted astrometry [1]; a zero-error draw reproduces that aim to below 10^{-6} AU at the exact epoch, anchoring the pipeline before any error is injected. At this geometry, one $\mathrm{m\,s^{-1}}$ of transverse velocity error maps to 15.4 AU of miss distance, so the $168\ \mathrm{m\,s^{-1}}$ allowance of Section 2.6 corresponds almost exactly to the full $2600\ \mathrm{AU}$ sphere — the two analyses are consistent by construction.

Four error sources are drawn independently, at the one-sigma levels of Table 6. The target-ephemeris draws perturb the parallax, both proper-motion components, and the radial velocity within the published uncertainties of the Kervella et al. astrometry [1], with the Akeson et al. solution [2] run as a self-consistent sensitivity variant (its own central values, uncertainties, and aim point). The remaining three sources are execution errors, and their levels are engineering assumptions, stated as such: a $50\ \mathrm{m\,s^{-1}}$ knowledge error on the achieved departure speed, representing the residual velocity-trim error after a telemetry-closed burn termination rather than raw thruster bias (uncorrected thruster-level bias would be tens of times larger — the assumption presumes the burn is closed on tracking data); a $0.10°$-per-axis pointing error on the departure asymptote, about one quarter of the total $0.41°$ open-loop allowance; and a $25\ \mathrm{m\,s^{-1}}$-per-axis isotropic residual representing launch-window discretization and the finite fidelity of the patched-conic departure mapping. Every run is seeded and bit-reproducible, and before the combined production run each source was exercised alone and checked against its closed-form transverse dispersion, which it reproduces at the percent level.

6.2 Results

With all four sources active, 19,631 of 20,000 draws pass within $2600\ \mathrm{AU}$ of the barycenter:

$$P(\text{closest approach} \leq 2600\ \mathrm{AU}) = 0.982, \qquad \text{95\% confidence interval } [0.980,\ 0.983].$$

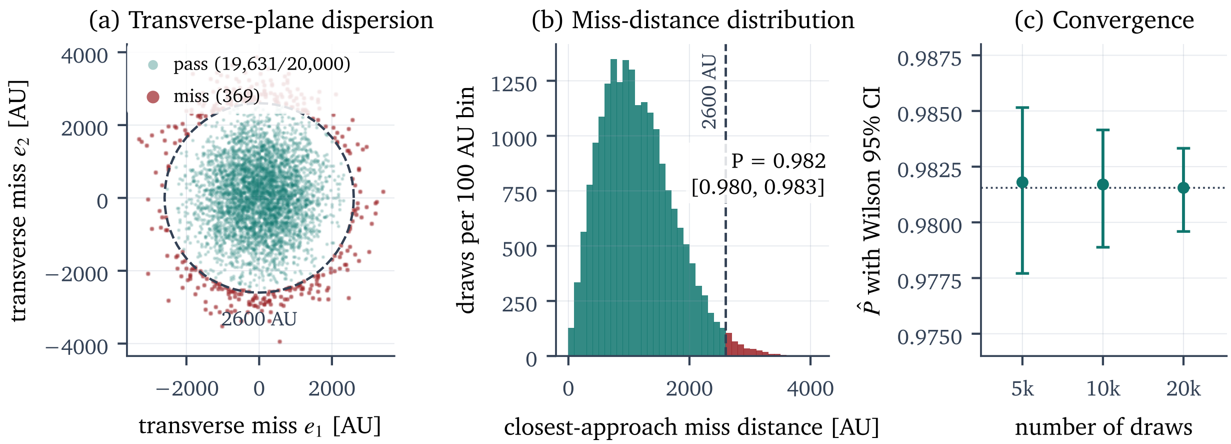

Figure 11. Encounter dispersion under the combined error budget, from the 20,000-draw seeded Monte Carlo. (a) Dispersion in the plane transverse to the approach velocity: 19,631 of 20,000 draws (teal) fall inside the $2600\,\mathrm{AU}$ requirement circle, 369 (red) outside. (b) Distribution of closest-approach miss distance against the requirement: the median miss is $1076\,\mathrm{AU}$, the 97.5th percentile $2499\,\mathrm{AU}$, and $P = 0.982$ of draws pass inside the sphere (95% confidence interval $[0.980, 0.983]$). (c) Convergence: the estimate is stable across the first 5,000, 10,000, and full 20,000 draws, with the confidence interval narrowing as $1/\sqrt{n}$. An independent analytic error-propagation model, sharing no code or random draws with the simulation, reproduces the success probability to 0.011%.

The median miss distance is $1076\,\mathrm{AU}$ — comfortably inside the sphere — and the 97.5th percentile is $2499\,\mathrm{AU}$, just under the requirement. The result is stable: recomputing the probability on the first 5,000 and first 10,000 draws gives 0.982 and 0.982 respectively, with the confidence-interval width halving from 5,000 to 20,000 draws exactly as the expected $1/\sqrt{n}$ scaling predicts. As an independent check, a first-principles analytic model propagates the same one-sigma levels through the closed-form encounter geometry to a two-dimensional transverse error covariance and integrates the resulting Gaussian over the $2600\,\mathrm{AU}$ disc by quadrature; sharing no code path or random draws with the simulation, it too gives 0.982, agreeing with the Monte Carlo result to 0.011%. Under the Akeson et al. variant astrometry, whose formal uncertainties are tighter, the probability rises to 0.992. The simulation also confirms the along-track argument of Section 2.6: radial errors shift the encounter epoch — by about ± 1250 years, one sigma, with all sources combined — but the pass still occurs, at negligible cost on the flat optimum of Section 2.4.

6.3 Which Error Source Dominates

Rerunning the simulation with one source active at a time ranks the contributors by RMS transverse miss: thrust-magnitude knowledge ($764\,\mathrm{AU}$) leads, followed by asymptote pointing ($647\,\mathrm{AU}$), target ephemeris ($617\,\mathrm{AU}$), and the post-departure residual ($547\,\mathrm{AU}$) (Figure 12a). The four single-source mean-square misses sum to the full-simulation value to within 0.2%, confirming that the sources act nearly independently, and a total-effect variance decomposition preserves the same leader. No single source dominates outright — all four sit within a factor of 1.4 — but the ordering carries a design lesson: the largest contributor is neither the star's ephemeris nor the pointing system but knowledge of the achieved departure speed. Improving thrust calibration and burn-termination velocity knowledge is therefore the highest-leverage mitigation available to the program, and it is a ground-and- telemetry investment rather than a flight-hardware one.

6.4 What the Verdict Is Conditioned On

Because three of the four error levels are engineering assumptions, the honest form of the result is the sensitivity table, not the headline number (Figure 12b). Doubling the assumed pointing error (to $0.20°$ per axis) drops the success probability from 0.982 to 0.897; doubling the thrust-magnitude knowledge error (to $100\ \mathrm{m\,s^{-1}}$)

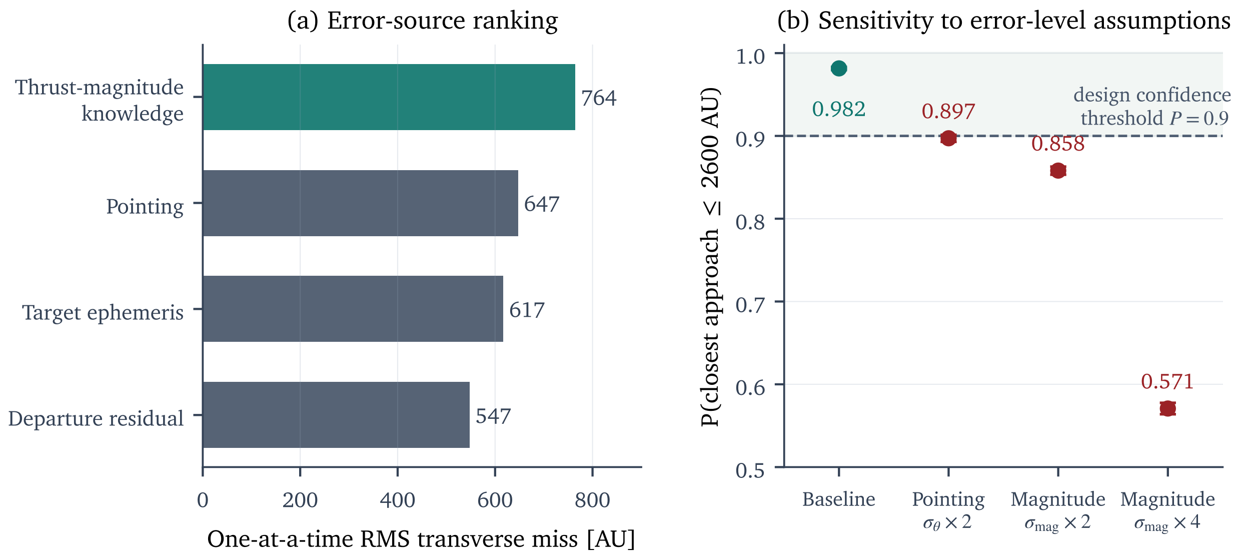

Figure 12. (a) Error-source ranking by one-at-a-time RMS transverse miss: thrust-magnitude knowledge (764 AU) leads pointing (647 AU), target ephemeris (617 AU), and the departure residual (547 AU), so thrust calibration is the highest-leverage mitigation. (b) Sensitivity of the success probability to the assumed error levels: the baseline 0.982 falls to 0.897 when the pointing assumption is doubled, to 0.858 when the thrust-magnitude assumption is doubled, and to 0.571 when it is quadrupled — each excursion below a 0.9 design confidence threshold. The verdict holds under the stated budget but is not robust to factor-of-two error growth.

drops it to 0.858; quadrupling that error (to $200\ \mathrm{m\,s^{-1}}$) leaves 0.571 — a coin flip. Each of these excursions falls below a 0.9 design confidence threshold. The verdict is therefore firm under the stated budget but not robust to factor-of-two growth in the execution errors: the mission earns its 0.982 through pre-flight thrust calibration and a telemetry-closed departure, not by geometric margin alone.

6.5 A Catalog Systematic, Held Separate

One effect is deliberately excluded from the probability above and reported separately. The two published astrometric solutions for the system disagree on the barycenter's proper motion by $26.1\ \mathrm{mas\,yr^{-1}}$ — formally about 6.7 times the quoted uncertainty of the adopted catalog — which propagates to a cross-track aim offset of roughly 860 AU at the encounter epoch (a first-order hand calculation gives 863 AU, in agreement). This is a systematic error, not a statistical one: averaging the two catalogs would manufacture false confidence, so the analysis instead adopts one solution [1], runs the other [2] as a variant, and reports the offset explicitly. If the full offset were real and the adopted aim wrong, the success probability would fall to 0.962 — degraded, not broken, because the offset consumes only a third of the sphere. The clean mitigation is observational: roughly seventeen years elapse between launch commitment and the final departure asymptote, and a modern re-reduction of the system's astrometry over that interval — anchored to current-generation space-based catalogs — would resolve the discrepancy before the aim is committed. That mitigation is stated, not modeled; no credit for it is taken in any number above.

In summary, the targeting question raised in Section 2.6 is closed in the affirmative: under the stated ephemeris covariance and execution error budget, the probe passes within 2600 AU of the barycenter with probability 0.982, open-loop, with no post-escape correction. The result is conditioned on execution-error levels that must be earned through thrust calibration and a tracked departure, and on an astrometric discrepancy that a pre-departure re-reduction can retire.

7. Program Cost Assessment

The mission requirements set a total program cost target of \$10M, allocated as \$3M for launch, \$3M for propulsion, and \$3M for the spacecraft bus and mission operations. Because the three categories sum to \$9M, the residual \$1M is treated throughout this section as unallocated program reserve. This section estimates what the architectures that close the mass budget actually cost, compares that cost structure against the allocation, and identifies where—and on what terms—the \$10M target can be met.

7.1 Cost Model

Program cost is estimated with a probabilistic parametric model comprising fourteen cost line items. Each line item carries a triangular uncertainty band (low, nominal, high) expressed in 2026 dollars per unit of its cost driver, anchored to published price bases and vendor-class quotations for small-spacecraft hardware, rideshare launch services, and deep-space operations. The drivers are physical quantities taken directly from the mass-closure model: launch cost scales with wet mass; the thruster and power-processing string scales with input power; the solar arrays scale with array wattage at 1 AU; propellant and tankage scale with the xenon load; and operations cost scales with the duration of powered flight. Architectures that employ perihelion pumping carry an additional thermal-qualification increment on the integration-and-test line, reflecting the high-temperature array and structure qualification required for repeated passes at 0.42 AU (Appendix D).

The total cost distribution is obtained by Monte Carlo sampling with 20,000 draws. One modeling choice deserves an explicit caveat: the line-item draws are treated as statistically independent. Real programs exhibit positive correlation between line items—a schedule slip inflates operations and integration together—and positive correlation widens the total-cost distribution without changing its mean. The percentile ranges reported below are therefore mildly optimistic in the tails, and the quoted probability of remaining under \$10M is, if anything, an overestimate. The program-margin line partially compensates by carrying program-level risk as its own band.

The physical inputs to the model rest on independently checked results. The mass-closure arithmetic that supplies the wet mass, propellant load, and power level was replicated algebraically by an independent check, which located the closure boundary at 113.3 kg for a low-Earth-orbit start; the underlying heliocentric Δv budget is carried across the reference interval of 22.9–$24.0\,\mathrm{km\,s^{-1}}$, whose upper end is a verified constructive bound and whose lower end is the best-found reference anchor from optimal-control analysis (Section 3).

7.2 Results

Figure 13(b) summarizes the total-cost distributions for the three architectures that close the mass budget. The low-Earth-orbit start at 150 kg wet mass carries a median (P50) cost of \$16.6M, with a 10th-to-90th-percentile range of \$14.2M to \$19.4M. The geostationary-transfer-orbit start at 100 kg—the recommended architecture—carries a P50 of \$16.2M (\$13.8M to \$19.0M), and its 80 kg variant a P50 of \$15.7M (\$13.3M to \$18.4M). Under conventional 2026 aerospace practice, the probability of completing any of these architectures under \$10M is effectively zero.

Two features of these results deserve emphasis. First, all three architectures land within about \$1M of one another: total cost is nearly insensitive to the choice of architecture, because the dominant cost lines are fixed or duration-driven rather than mass-driven. Second, the overrun does not fall where the allocation anticipated it, as the following comparison shows.

7.3 Comparison with the Mission's Allocation

Mapping the line items onto the mission's three budget categories (Figure 13(a)) shows that the \$3M/\$3M/\$3M split is misallocated in a favorable direction. Launch comes in at \$1.0M to \$1.3M and propulsion at \$0.6M to \$1.2M across the closing architectures—each roughly one third of its \$3M allocation. The entire overrun is concentrated in the bus, operations, and program-margin category: approximately \$12M nominal against the \$3M allocated. The top cost drivers, in order, are powered-flight operations, program margin and management, and the spacecraft bus itself.

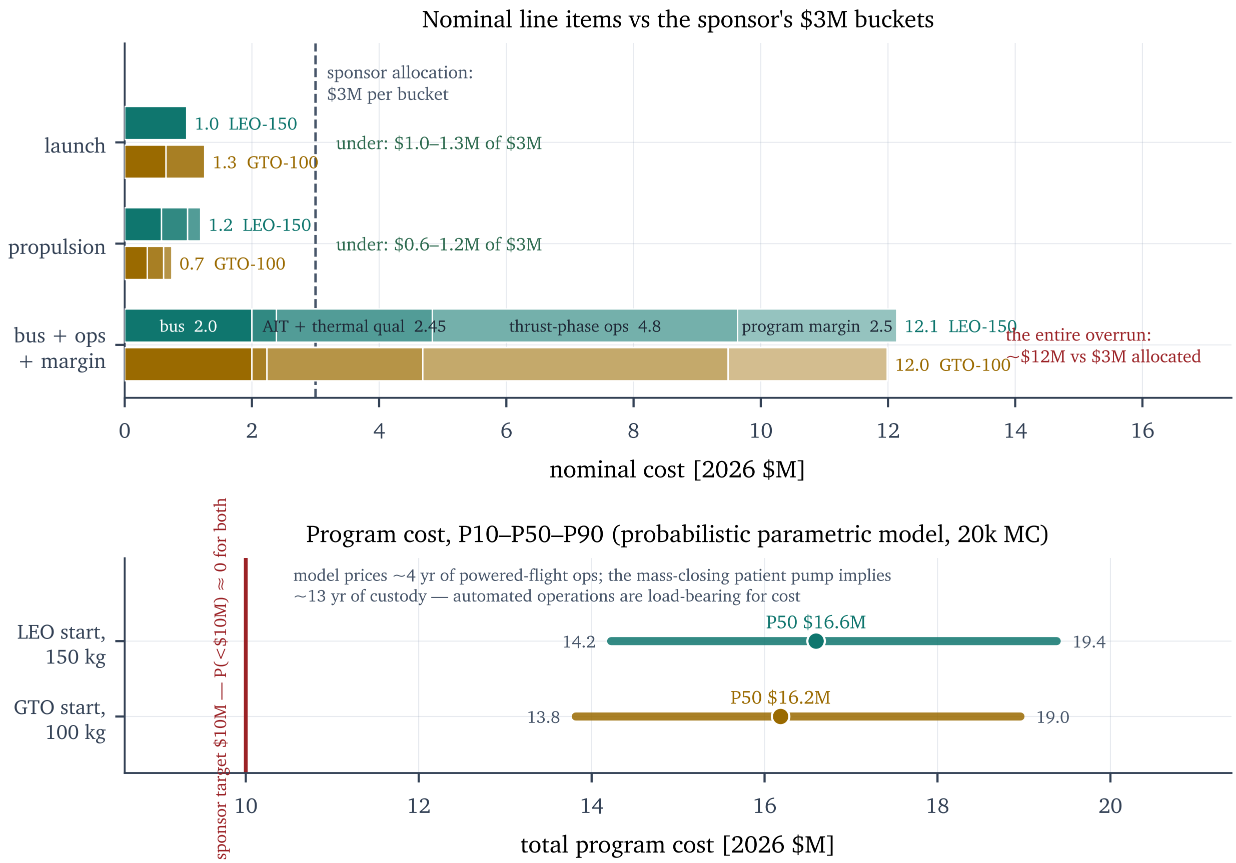

Figure 13. (a) Nominal cost by mission budget category for the two primary architectures, compared with the $3M/$3M/$3M allocation: launch and propulsion come in at roughly one third of their allocations, while the bus-and-operations category absorbs the entire overrun. (b) 10th–50th–90th-percentile total-cost ranges from the 20,000-draw Monte Carlo, compared with the $10M program target.

7.4 Operations Duration

In one respect the reported percentiles understate the true commitment. The model prices approximately four years of powered-flight operations, whereas the patient perihelion-pumping profile that closes the mass budget implies roughly thirteen years of powered-flight custody—the Earth-escape spiral followed by the twelve-year pumping phase, most of it coasting. Under traditional staffed operations, the additional custody years would push the median cost higher still. Automation of the thrust phase is therefore load-bearing for program cost, not merely a communications-link convenience: the mission concept's stated tolerance for mostly-automated operations, including possible loss of contact, is a cost-structural feature of the design, not an operational afterthought.

7.5 The Path to $10M

The gap between the $16M-class estimates and the $10M target is programmatic, not physical. The launch and propulsion categories are already comfortably within budget; the roughly $12M bus-and-operations figure prices 2026 aerospace norms—a staffed operations floor, a conventional smallsat bus, and a full environmental-qualification campaign augmented for the perihelion thermal environment. Meeting the $10M total requires compressing the conventional bus, integration-and-test, and operations practice by roughly 40 percent. That is aggressive but not implausible for a team with small-satellite heritage operating under an expendable, largely autonomous mission posture—which is precisely the posture the mission requirements describe. The discriminating cost issue for this program is therefore the credibility of a lean, automated operations concept, not any question of physics.

7.6 The $50M Tier

A higher program tier of $50M buys certainty rather than new physics. The most effective use of the additional budget is an escape-injection rideshare: purchasing a near-zero-hyperbolic-excess-energy departure from the launch provider eliminates the Earth-escape spiral entirely. The vehicle then closes at approximately 60 kg wet mass, and the three principal risk carriers of the baseline architecture—the radiation-belt transit, the array-degradation margin on the pumping phase, and the twelve-year patience requirement—are removed outright. The median cost of that architecture is approximately $15.6M; the $50M tier is barely touched.

The additional budget does not usefully purchase alternative power or propulsion. Radioisotope generators cannot supply the kilowatt-class power the electric-propulsion system requires; a fission reactor exceeds even the $50M tier. A multi-stage chemical kick vehicle capable of the $13.85\,\mathrm{km\,s^{-1}}$ impulsive departure floor is both more expensive and less reliable than buying the equivalent injection energy from the launch vehicle.

Gravity assists warrant direct assessment at this tier, because the $50M budget suspends the no-assist ground rule under which the rest of this report operates. At scoping fidelity, a Jupiter gravity assist can supply of order 10–15 $\mathrm{km\,s^{-1}}$ of heliocentric hyperbolic excess—a flyby of the most massive planet, operating near its $\sim$5.9 $\mathrm{km\,s^{-1}}$-class maximum turning benefit for slow arrivals—which would substitute for most of the perihelion-pumping phase. Three costs stand against that benefit. First, the maneuver requires phase alignment of Jupiter with the departure asymptote; the alignment recurs on synodic windows, and a fixed end-2029 launch date makes it a matter of luck rather than design. Second, the routing adds roughly two to six years of inner- and outer-system transfer before the heliocentric cruise begins. Third, it adds deep-space navigation and operations effort of exactly the class the cost model above identifies as dominating program cost. The other assist-class option is the solar-Oberth route with a purchased kick stage examined in Fermi Explorer's own reference analysis [12]: a heat-shielded vehicle executing a chemical burn at a perihelion near ten solar radii—a thermal-protection development of the Parker Solar Probe class.

The honest conclusion is that at $50M the escape-injection rideshare remains the recommended use of funds: it removes the identified risk carriers directly, with no launch-window dependence and no new development. A Jupiter assist is a launch-window-contingent opportunity worth assessing against the actual 2029 departure geometry—that window assessment is a follow-up item, not performed here—rather than a baseline. The recommendation at the $50M tier is therefore the same in kind as at the $10M tier—solar-electric propulsion with the escape spiral bought down or bought out—executed with margin instead of on the margin.

8. Alternative Targets

The mission requirements admit flexibility in the choice of stellar target, and the machinery developed for Alpha Centauri applies to any star with well-determined astrometry. This section evaluates the three alternative targets offered by the Fermi Explorer team with the identical method used throughout this report—exact linear kinematics for the stellar motion, the same departure geometry from a $400\,\mathrm{km}$ low Earth orbit, and a miss tolerance set to 1% of each target's current distance—and then takes up the team's invitation to suggest better targets with a sweep of the full close-encounter catalog (Section 8.4). Astrometry for the alternative targets is taken from Gaia Data Release 3 [3], except for Alpha-2 Librae, whose solution derives from the revised Hipparcos reduction [4]; the Alpha Centauri solution is that of Kervella et al. [1]. As with the primary target, every alternative-target result was replicated independently from first principles; the worst-case kinematic discrepancy is below 0.5%.

Figure 14 presents the impulsive departure Δv as a function of arrival time for all four targets, and Table 7 summarizes the optima. Two of the three alternatives are viable, and one of them is energetically superior to the primary target.

8.1 LSPM J2146+3813: The Energetically Preferred Target

LSPM J2146+3813 is an M5V red dwarf at a present distance of $22.99\,\mathrm{ly}$. Although it lies more than five times farther away than Alpha Centauri today, its space motion carries it toward the Sun: the published en-

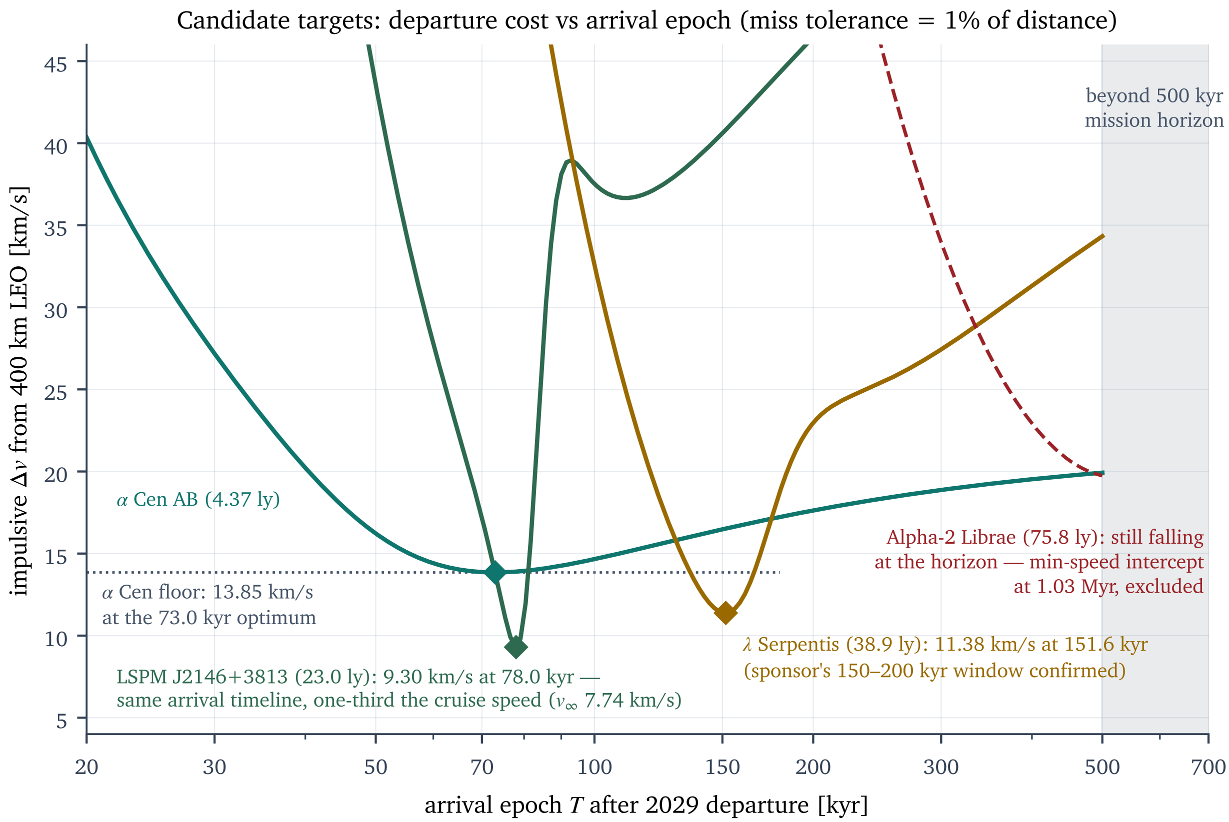

Figure 14. Minimum impulsive departure Δv from a $400\,\mathrm{km}$ low Earth orbit versus arrival time for the four candidate targets, with the minima of the three in-horizon targets marked. LSPM J2146+3813 (highlighted) achieves the lowest floor of any target examined, at essentially the same arrival epoch as Alpha Centauri; Alpha-2 Librae never approaches competitiveness inside the $500{,}000$-year horizon.

counter analysis of Bailer-Jones [6], based on Gaia astrometry, gives a perihelion of $0.568\,\mathrm{pc}$ at $+82.5$ thousand years—one of the tightest known future stellar encounters. Our independent linear-kinematics computation is consistent with that published result.

The consequences for mission design are striking. The impulsive departure floor is $9.30\,\mathrm{km\,s^{-1}}$ at a 78,017-year arrival, with the departure asymptote tilted only $0.84°$ below the ecliptic. The required heliocentric cruise speed is $7.74\,\mathrm{km\,s^{-1}}$—one third of the $23.64\,\mathrm{km\,s^{-1}}$ demanded by Alpha Centauri—on essentially the same arrival timeline ($78{,}000$ versus $73{,}000$ years).

The reduction in cruise speed transforms the propulsion problem. The conventional outward solar-electric spiral, which saturates below the Alpha Centauri requirement at every thrust level examined (Section 4), comfortably exceeds the LSPM requirement: the spiral ceiling of $17.0\,\mathrm{km\,s^{-1}}$ at a thrust acceleration of $10{\times}10^{-4}\,\mathrm{m/s^2}$ is more than double the $7.74\,\mathrm{km\,s^{-1}}$ needed. No perihelion pumping is required at all—and with it disappear the $0.42\,\mathrm{AU}$ thermal environment, the twelve-year pumping phase, and the thrust-acceleration margin sensitivity that dominate the risk posture of the Alpha Centauri design. A vehicle that closes its mass budget only narrowly against Alpha Centauri closes comfortably against this target.

The choice before the Fermi Explorer mission is therefore straightforward. If the objective is to reach another star's Oort cloud, LSPM J2146+3813 is the energetically preferred target: a lower departure floor, a far lower cruise speed, a benign trajectory with no perihelion passes, and the same arrival epoch. If the objective is Alpha Centauri specifically, the primary analysis of this report stands. The catalog sweep of Section 8.4 extends this conclusion from the offered candidates to the full encounter catalog.

Table 7. Alternative-target comparison at each target's departure-Δv optimum. All departures from a 400 km low Earth orbit, no gravity assists; miss tolerance 1% of current distance.

Target	Distance (ly)	Arrival epoch (kyr)	Impulsive Δv floor ($\mathrm{km\,s^{-1}}$)	Published closest encounter
Alpha Centauri AB	4.37	73.0	13.85	3.15 ly at +28.0 kyr[a]
LSPM J2146+3813	22.99	78.0	9.30	0.568 pc at +82.5 kyr [6]
Lambda Serpentis	38.86	151.6	11.38	2.30 pc at +167 kyr [5]
Alpha-2 Librae	75.80	excluded[b]	19.74[b]	—

[a]Closest natural approach computed in this analysis (Section 2); the star itself passes no closer. [b]The minimum-cruise-speed intercept falls at 1.03 million years, outside the 500,000-year horizon; the quoted Δv is the best achievable within the horizon.

8.2 Lambda Serpentis

Lambda Serpentis is a G0V star—a close solar analog—at 38.86 ly. Its impulsive departure floor is 11.38 $\mathrm{km\,s^{-1}}$ at a 151,607-year arrival, which confirms the 150,000–200,000-year arrival window anticipated by the Fermi Explorer team. The published encounter analysis gives a perihelion of 2.30 pc at +167 thousand years [5], consistent with our computation. Lambda Serpentis is a sound second alternative: cheaper to reach than Alpha Centauri, though at roughly double the flight time and without the exceptional encounter geometry of LSPM J2146+3813. Its solar spectral type may carry independent scientific weight in target selection.

8.3 Alpha-2 Librae: Excluded

Alpha-2 Librae, at 75.80 ly, should be removed from the candidate list. Its minimum-cruise-speed intercept falls at 1.03 million years, outside the 500,000-year mission horizon; constrained to arrive within the horizon, it requires an impulsive departure of 19.74 $\mathrm{km\,s^{-1}}$—substantially worse than Alpha Centauri itself.

The exclusion is robust to the weakest input in the target set. Alpha-2 Librae is a spectroscopic binary, and its systemic radial velocity is correspondingly uncertain: $-22.0 \pm 5.8\ \mathrm{km\,s^{-1}}$. Sweeping that velocity across its full $\pm 2\sigma$ range moves the minimum-speed intercept epoch, but at no point within the range does the epoch enter the 500,000-year horizon. The conclusion therefore does not depend on the precise value of the poorly determined radial velocity: under no plausible kinematic solution is Alpha-2 Librae competitive within the mission constraints.

8.4 Catalog Sweep: No Better Target Found

The Fermi Explorer team invited suggestions of better targets, and the candidates examined above are only the offered set. To answer the invitation directly, all 61 entries of the Bailer-Jones close-encounter catalog [6] were swept with the same departure-cost metric used throughout this section: minimum impulsive Δv from a 400 km low Earth orbit, arrival within the 500,000-year horizon, and the same miss rule (with the 2,600 AU allowance substituted where an encounter's own perihelion already satisfies the percentage rule). Of the 61 catalog rows, 29 are future encounters, 9 fall within the 500,000-year horizon, and 5 of those carry radial velocities free of the catalog's quality flags.

The sweep confirms LSPM J2146+3813 as the best high-confidence target. Only two catalog entries formally undercut its 9.30 $\mathrm{km\,s^{-1}}$ floor: the anonymous distant M dwarfs Gaia DR3 398496965625177216 (238 ly, 8.78 $\mathrm{km\,s^{-1}}$) and Gaia DR3 4155835025908320640 (311 ly, 8.82 $\mathrm{km\,s^{-1}}$). Their apparent advantage does not survive inspection. It rests on weak radial velocities (standard errors of 8–11 $\mathrm{km\,s^{-1}}$, against which the quoted floors are not robust) and on a near-degenerate encounter geometry: both stars pass close to the Sun themselves, so the departure cost collapses toward the metric's bare solar-escape floor while the mission value collapses with it—a flyby at 150–280 $\mathrm{km\,s^{-1}}$ relative speed arriving at +330,000 to +470,000 years, against LSPM J2146+3813's 83 $\mathrm{km\,s^{-1}}$ encounter at +78,000 years. That is strictly worse mission value at nominally lower departure cost.

Two further results bound the search. No non-M-dwarf in the catalog approaches Lambda Serpentis's 11.38 $\mathrm{km\,s^{-1}}$: every unflagged future encounter inside the horizon is an M dwarf, so Lambda Serpentis stands as the best solar-type option available. And Gliese 710, whose celebrated 0.06 pc perihelion is the tightest known future stellar approach, lies at +1.29 million years—well outside the 500,000-year horizon—and is therefore not a candidate under the mission constraints.

The answer to Fermi Explorer's invitation is thus explicit: the sweep found no better target. LSPM J2146+3813 is the energetically preferred target catalog-wide, not merely within the offered set.

9. Conclusions and Recommendations

The analysis condenses to eight findings and six recommendations, followed by an overall assessment layered by certainty.

9.1 Findings

F1. Departure geometry and optimum. The minimum-Δv departure for Alpha Centauri falls at $T = 73{,}012$ years, requires a heliocentric cruise speed $v_\infty = 23.64\,\mathrm{km\,s^{-1}}$ aimed $2.48°$ below the ecliptic, and corresponds to an Earth-relative hyperbolic excess of $18.59\,\mathrm{km\,s^{-1}}$ and an impulsive Δv of $13.85\,\mathrm{km\,s^{-1}}$ from a $400\,\mathrm{km}$ low Earth orbit. The Δv curve is flat from roughly 70,000 to 85,000 years (spread below $0.25\,\mathrm{km\,s^{-1}}$; below $0.1\,\mathrm{km\,s^{-1}}$ from roughly 67,000 to 80,000 years), leaving the arrival epoch free for launch-window phasing. Every number in this finding was independently recomputed from first principles, with worst-case discrepancy below 0.5%.

F2. The formal constraints are not binding. The 2,600 AU miss tolerance is worth only about 83 m/s of Earth-relative excess speed (about 71 m/s of impulsive Δv); it relaxes targeting precision, not the energy problem. The 500,000-year horizon is likewise not binding: the star is receding by then, and the requirement at the horizon *rises* to $v_\infty = 30.61\,\mathrm{km\,s^{-1}}$ ($\Delta v = 19.93\,\mathrm{km\,s^{-1}}$). The aiming tolerance is $0.41°$, equivalent to roughly 168 m/s of velocity dispersion. A 20,000-draw dispersion simulation with all execution and ephemeris error sources acting together confirms that this budget is achievable open-loop, with no terminal correction maneuver: the probability of passing inside the sphere is 0.982 (95% confidence interval $[0.980, 0.983]$), reproduced to 0.011% by an independent analytic error-propagation model (Section 6). The result is conditioned on the stated engineering error budget, with thrust-magnitude knowledge at burn cutoff the dominant contributor: doubling that assumption (50 m/s) drops the probability to 0.858, and doubling the assumed pointing error ($0.1°$ per axis) drops it to 0.897—both below 0.9.

F3. A conventional solar-electric spiral cannot reach cruise speed. Because array power falls as the inverse square of solar distance (capped at four times the 1 AU rating inside roughly 0.5 AU by the array thermal limit), an outward low-thrust spiral saturates: the maximum achievable v_∞ is approximately 0, 3.4, and $17.0\,\mathrm{km\,s^{-1}}$ at thrust accelerations of 1.5, 5, and $10 \times 10^{-4}\,\mathrm{m/s^2}$ respectively—all below the $23.38\,\mathrm{km\,s^{-1}}$ cruise floor. This saturation was confirmed by two independent trajectory integrators, agreeing to within 2.7%.

F4. Perihelion pumping defeats the power wall. Retrograde thrust arcs near 1 AU shed angular momentum and lower perihelion to 0.42 AU; prograde burn arcs then concentrate near perihelion, where available power is one to four times the 1 AU rating. The optimized profile at $a_0 = 2.5\times10^{-4}\,\mathrm{m/s^2}$ delivers the heliocentric phase for $\Delta v = 23.98\,\mathrm{km\,s^{-1}}$ in 12.0 years of flight time (mostly coasting, with about 1.3 years of cumulative thrust) at a final mass fraction of 0.418. The total solar-electric budget from low Earth orbit is 30.5–$31.6\,\mathrm{km\,s^{-1}}$: a $7.6\,\mathrm{km\,s^{-1}}$ Earth-escape spiral plus a heliocentric phase quoted across the reference interval 22.9–$24.0\,\mathrm{km\,s^{-1}}$. The upper end of the interval is a verified constructive bound—the explicitly integrated, timestep-refined pumping schedule. The lower end is anchored by optimal-control analysis at a higher thrust acceleration ($a_0' = 1\times10^{-3}\,\mathrm{m/s^2}$, four times the design value), where an indirect method based on Pontryagin's minimum principle converges to $22.66\,\mathrm{km\,s^{-1}}$ at the required cruise speed, extended downward by a monotonicity argument valid for throttleable thrust. The anchor is single-method—direct collocation did not converge on this problem in any documented attempt—so it is the best trajectory found by one method, not a certified bound. The certified floor is $16.56\,\mathrm{km\,s^{-1}}$, proved over every admissible thrust program in the throttleable power-envelope class with the 0.42 AU perihelion bound and no constraint on flight time (the impulsive multi-burn limit is

$16.50\,\mathrm{km\,s^{-1}}$; $15.6\,\mathrm{km\,s^{-1}}$ at $0.25\,\mathrm{AU}$). The gap between that floor and the operational figure is real but schedule-shaped: with flight time unconstrained, an admissible trajectory reaching cruise speed for $18.87\,\mathrm{km\,s^{-1}}$ exists, at the price of 33–58 years of powered flight; within the 12-year operational horizon the best trajectory found is $23.98\,\mathrm{km\,s^{-1}}$. A fast 4.9-year variant costs $28.2\,\mathrm{km\,s^{-1}}$; the patience trade is worth about $4\,\mathrm{km\,s^{-1}}$, and only the patient profile closes the mass budget. For a 2029 launch, escape occurs around 2043, with no effect on the arrival epoch.

F5. Mass closure is achievable within a demonstrated component class, but not at 100 kg from low Earth orbit by any trajectory found. The required operating point—a specific impulse of at least $2800\,\mathrm{s}$ at total efficiency near 0.55 in a 0.5–$1.0\,\mathrm{kW}$ string, with roughly 17,000–21,000 hours of single-string firing and 64–$103\,\mathrm{kg}$ of xenon throughput—is not covered by any flown Hall thruster (flown Hall specific impulse tops out near $1{,}650\,\mathrm{s}$ at this power class). It is bracketed by demonstrated gridded-ion hardware: NSTAR's flight-documented TH3–TH4 throttle points (0.91–$1.02\,\mathrm{kW}$, $2{,}843$–$2{,}942\,\mathrm{s}$, efficiency 0.527–0.554), with $30{,}352$ hours and $235\,\mathrm{kg}$ of single-string ground life and $16{,}265$ hours of single-unit flight; dual-string redundancy and the standard $1.5\times$ throughput qualification margin are the available mitigations. Closure therefore rests on a demonstrated component class (gridded ion), requiring a mission-specific qualification campaign. The $2800\,\mathrm{s}$ requirement is specific to the low-Earth-orbit start; from a geostationary transfer orbit a $2400\,\mathrm{s}$ thruster closes at the pumping-feasibility thrust floor (boundary near $93.5\,\mathrm{kg}$; Section 5). From low Earth orbit, the minimum closing wet mass is 113.5–$146\,\mathrm{kg}$ across the 22.9–$24.0\,\mathrm{km\,s^{-1}}$ reference interval; the $100\,\mathrm{kg}$ concept misses closure by 1.0–2.8% of wet mass for every trajectory found—closure cannot be excluded outright, since it would follow from a heliocentric trajectory below the ~22.3 $\mathrm{km\,s^{-1}}$ crossover within the 12-year operational horizon, but no trajectory found at the design thrust level comes within $0.6\,\mathrm{km\,s^{-1}}$ of that crossover—while $150\,\mathrm{kg}$ closes throughout in the planar accounting. From the transfer orbit, where the departure budget falls from 7.6 to $4.24\,\mathrm{km\,s^{-1}}$, an $80\,\mathrm{kg}$ vehicle closes with +1.9 points of margin, $100\,\mathrm{kg}$ with +4.1, and $150\,\mathrm{kg}$ with +7.1. Under a 30% thrust-acceleration margin, the transfer-orbit start retains $100\,\mathrm{kg}$ closure only at the favorable end of the interval (boundary $83.4\,\mathrm{kg}$ at $22.9\,\mathrm{km\,s^{-1}}$ versus $103.7\,\mathrm{kg}$ at $24.0\,\mathrm{km\,s^{-1}}$; approximately $110\,\mathrm{kg}$ carries the margin throughout), while the same margin drives the low-Earth-orbit boundary to 312–$910\,\mathrm{kg}$. Under the measured out-of-plane departure cost of $0.58\,\mathrm{km\,s^{-1}}$—obtained from a genuinely three-dimensional re-optimization of the pumping profile, inside the planar analysis's $22\,\mathrm{m/s}$-to-$1.02\,\mathrm{km\,s^{-1}}$ band but far above its low end—the transfer-orbit start closes at both ends of the reference interval (+5.1 and +3.1 points of margin, and +2.3 points even at the band's worst case) while the $150\,\mathrm{kg}$ low-Earth-orbit fallback fails at the conservative end (-0.8 points). Steering the plane change during powered flight—the maneuver the low-Earth-orbit architectures were counting on to stay near the band's low end—costs $0.58\,\mathrm{km\,s^{-1}}$ for the best trajectories found, against the roughly $0.10\,\mathrm{km\,s^{-1}}$ that margin can absorb. The closing designs are 63–68% propellant, essentially all xenon: the concept's proportions are correct. The closure arithmetic was verified algebraically in an independent recomputation, which reproduced the closure boundary at $113.3\,\mathrm{kg}$.

F6. The Fermi Explorer reference analysis is confirmed on geometry and amended on propulsion. Our independent computation agrees with the reference analysis [12] on the departure geometry to within 0.3% (their optimum: $13.88\,\mathrm{km\,s^{-1}}$ at 72,800 years, tilt 2.4°). Two material differences remain. First, their published figure of roughly $20\,\mathrm{km\,s^{-1}}$ for a realistic low-thrust departure is a benchmark assumption rather than a model output—their own trajectory engine yields $25.1\,\mathrm{km\,s^{-1}}$ for the outward spiral—and their vehicle sizing of roughly $500\,\mathrm{kg}$ inherits the optimistic figure. Second, the amendment on propulsion must be stated precisely: their conclusion was that pure solar-electric propulsion does not close from low Earth orbit *at today's array specific power* (about $60\,\mathrm{W/kg}$)—their own analysis identifies a pure-solar closing corner above roughly $100\,\mathrm{W/kg}$—and they separately rejected a *single-pass* electric perihelion-burn architecture. Multi-revolution perihelion pumping amends both conclusions at once: it closes the mission with solar power at today's ~60 W/kg by spreading the perihelion burn over many revolutions, directly answering the single-pass objection (Section 4). Closure is at 113.5–$146\,\mathrm{kg}$ from low Earth orbit or approximately $100\,\mathrm{kg}$ from a geostationary transfer orbit, within the demonstrated gridded-ion component class.

F7. The \$10M cost target fails under conventional practice, for programmatic rather than physical reasons. The median program cost is \$15.7M–\$16.6M across the mass-closing architectures, with a probability of remaining under $10M that is effectively zero under 2026 aerospace norms. The allocation, however, is misjudged in a favorable direction: launch ($1.0M–$1.3M) and propulsion ($0.6M–$1.2M) each come in at roughly one third of their $3M allocations, and the entire overrun sits in the bus, operations, and program-margin category (roughly $12M nominal against $3M allocated). The cost model prices about four years of powered-flight operations, whereas the mass-closing profile implies roughly thirteen years of custody: automation is essential to the cost estimate, not merely to communications. At a $50M tier, an escape-injection rideshare closes the vehicle at approximately $60\,\mathrm{kg}$ and removes the belt-transit, array-degradation, and patience risks outright, at a median cost near $15.6M; radioisotope generators cannot supply the required kilowatt-class power, and a fission reactor exceeds even that tier.

F8. An energetically superior target exists. LSPM J2146+3813, an M5V dwarf with a published closest approach of $0.568\,\mathrm{pc}$ at +82.5 thousand years [6], requires an impulsive departure floor of $9.30\,\mathrm{km\,s^{-1}}$ (versus 13.85) and a cruise speed of $7.74\,\mathrm{km\,s^{-1}}$ (versus 23.64) on the same arrival timeline as Alpha Centauri. No perihelion pumping is needed: the conventional spiral ceiling exceeds the requirement more than twofold. Lambda Serpentis is a sound second alternative ($11.38\,\mathrm{km\,s^{-1}}$ at $151{,}607$ years, confirming the Fermi Explorer team's anticipated window); Alpha-2 Librae is excluded, robustly across the $\pm 2\sigma$ range of its uncertain systemic radial velocity. A sweep of the Bailer-Jones encounter catalog found no higher-confidence target inside the $500{,}000$-year horizon and no Sun-like challenger to Lambda Serpentis (Section 8).

9.2 Recommendations

R1. Procure a geostationary-transfer-orbit drop-off with a 100–110 kg vehicle. This preserves the concept vehicle with +4.1 points of closure margin at the conservative end of the reference interval and, decisively, remains robust where the low-Earth-orbit start is fragile: a 30% thrust-acceleration margin is carried at $100\,\mathrm{kg}$ at the favorable end of the interval and at approximately $110\,\mathrm{kg}$ throughout (against 312–$910\,\mathrm{kg}$ from low Earth orbit); under the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane departure cost, the transfer-orbit start closes at both ends of the interval while the $150\,\mathrm{kg}$ low-Earth-orbit fallback fails at the conservative end. If the delivery orbit cannot be changed, grow the vehicle to $150\,\mathrm{kg}$, which closes from low Earth orbit across the reference interval before the out-of-plane cost is charged but is margin-fragile: it carries no allowance against array degradation and, once the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane cost is included, fails at the conservative end of the interval even with the plane change steered during powered flight.

R2. Select a gridded-ion propulsion string and start its qualification campaign at once. For the low-Earth-orbit start, closure requires a specific impulse of at least $2800\,\mathrm{s}$; from the recommended transfer-orbit start, a $2400\,\mathrm{s}$ thruster closes at the pumping-feasibility thrust floor (boundary near $93.5\,\mathrm{kg}$), so the $2800\,\mathrm{s}$ requirement is specific to the low-Earth-orbit contingency. The operating point is not covered by any flown Hall thruster (flown Hall tops out near $1{,}650\,\mathrm{s}$ at this power class) but is bracketed by demonstrated gridded-ion hardware—NSTAR's TH3–TH4 throttle points, with $30{,}352$ hours and $235\,\mathrm{kg}$ of single-string ground life. Thruster selection should therefore target the demonstrated gridded-ion class, with dual-string redundancy and the standard $1.5\times$ throughput qualification margin, and the mission-specific qualification campaign should be treated as the pacing schedule item for an end-2029 launch.

R3. Adopt an automation-first operations architecture. The mass-closing trajectory implies roughly thirteen years of powered-flight custody, while the cost estimate prices four; staffed operations over the full duration would push the median cost higher still. Autonomy during the escape spiral and pumping phase is essential to both the mass budget and the cost estimate, and the mission's stated tolerance for mostly-automated operations—including possible loss of contact—should be exploited as a structural design feature, not held in reserve.

R4. Scope thermal qualification to the perihelion environment early. The pumping profile executes repeated passes at $0.42\,\mathrm{AU}$, inside the orbit of Mercury, requiring thermal qualification of the class demonstrated by the MESSENGER mission. This is established engineering but a real cost and schedule item, and it should enter the qualification program from the start rather than as a late addition; together with the propulsion qualification campaign (R2), it is one of the two items on which the end-2029 launch date is conditional.

R5. Weigh LSPM J2146+3813 as the energetically dominant alternative. If the mission objective is to reach another star's Oort cloud rather than Alpha Centauri specifically, this target offers a lower departure floor, one third the cruise speed, no perihelion passes, and the same arrival epoch—and a vehicle that is marginal against Alpha Centauri closes comfortably against it.

R6. Carry four items into the next analysis phase. First, thrust-magnitude calibration: the dispersion analysis identifies velocity knowledge at burn cutoff as the dominant targeting error—at the assumed $50\,\mathrm{m/s}$ residual the probability of passing inside the encounter sphere is 0.982, but doubling that assumption drops it to 0.858 and quadrupling it to 0.571—so telemetry-closed thrust calibration and cutoff accuracy must be carried as explicit propulsion requirements and folded into the qualification campaign (R2). Second, an ephemeris re-reduction: the two published Alpha Centauri astrometric solutions disagree by a systematic offset worth roughly 860 AU of cross-track aim at the encounter epoch—held separate from the headline probability, not averaged into it—and roughly seventeen years of pre-departure astrometry are available to resolve it before the asymptote is committed. Third, a detailed bus design, since the bus and its operations concept—not launch or propulsion—carry the entire cost overrun. Fourth, a full schedule analysis: the end-2029 launch verdict is conditional on the propulsion and thermal qualification campaigns, and no integrated schedule has been constructed.

9.3 Overall Assessment

The assessment resolves into four layers, in decreasing order of certainty.

Trajectory-class feasibility: established. A solar-electric vehicle can reach the required cruise speed with solar power alone via multi-revolution perihelion pumping. This rests on a verified constructive trajectory—explicitly integrated and timestep-refined, delivering $24.0\,\mathrm{km\,s^{-1}}$ at the design thrust acceleration—and the mechanism was independently reconstructed during verification. The physics that appeared to forbid the outcome, the inverse-square decay of solar power, constrains only outward spirals; it does not constrain trajectories that do their burning at or inside 1 AU.

Quantitative closure of the 100 kg concept from low Earth orbit: not achieved within the operational horizon. No trajectory found in this analysis closes 100 kg from the specified low-Earth-orbit start within the twelve-year powered-flight horizon that the component-life and burn-sizing assumptions price; the shortfall is 1.0–2.8% of wet mass across the 22.9–$24.0\,\mathrm{km\,s^{-1}}$ reference interval. Closure would require a heliocentric trajectory below the ~$22.3\,\mathrm{km\,s^{-1}}$ crossover, and the verdict is inherently horizon-scoped: with flight time unconstrained, an admissible trajectory below the crossover exists—$18.87\,\mathrm{km\,s^{-1}}$, at the price of 33–58 years of powered flight, far beyond demonstrated thruster life—so no certified lower bound above the crossover is possible once flight time is freed, while the certified floor over the full class stands at $16.56\,\mathrm{km\,s^{-1}}$. Within the twelve-year horizon, no flyable-thrust trajectory found comes within $1\,\mathrm{km\,s^{-1}}$ of the crossover (best found: $23.98\,\mathrm{km\,s^{-1}}$), and even the best-found trajectory at four times the design thrust, $22.66\,\mathrm{km\,s^{-1}}$, remains $0.4\,\mathrm{km\,s^{-1}}$ above it.

Specification compliance: conditional. The mission closes at 100 kg from a geostationary-transfer-orbit drop-off—a delivery-orbit deviation from the low-Earth-orbit ground rule—or at approximately 150 kg from low Earth orbit, a margin-fragile variant that the measured $0.58\,\mathrm{km\,s^{-1}}$ out-of-plane departure cost defeats at the conservative end of the reference interval. The launch-before-end-2029 objective is schedule-conditional on the propulsion qualification campaign within the demonstrated gridded-ion component class and on MESSENGER-class thermal qualification; no full schedule analysis was performed.

Cost: the \$10M requirement fails under conventional practice. The median estimate is \$15.7M–\$16.6M against the \$10M target, with launch and propulsion each well under their allocations and the entire gap concentrated in the bus, operations, and margin category. The gap is programmatic, not physical; the path to the target runs through lean automated operations and compressed integration and test—a posture the mission's own ground rules invite.

The discriminating question for this program is therefore not whether the trajectory class can be flown—it can—but whether the propulsion string can be qualified in time, and whether the operations concept can be built lean enough to afford the years it takes to fly it.

References

[1] Kervella, P., Mignard, F., Mérand, A., and Thévenin, F., "Close stellar conjunctions of α Centauri A and B until 2050," *Astronomy & Astrophysics*, Vol. 594, A107, 2016.

[2] Akeson, R., et al., "Precision Millimeter Astrometry of the α Centauri AB System," *Astronomical Journal*, Vol. 162, 14, 2021.

[3] Gaia Collaboration, "Gaia Data Release 3," *Astronomy & Astrophysics*, Vol. 674, A1, 2023.

[4] van Leeuwen, F., "Validation of the new Hipparcos reduction," *Astronomy & Astrophysics*, Vol. 474, p. 653, 2007.

[5] Bailer-Jones, C. A. L., et al., "New stellar encounters discovered in the second Gaia data release," *Astronomy & Astrophysics*, Vol. 616, A37, 2018.

[6] Bailer-Jones, C. A. L., "Stars That Approach within One Parsec of the Sun: New and More Accurate Encounters from Gaia DR3," *Astrophysical Journal Letters*, Vol. 935, L9, 2022.

[7] Betts, J. T., *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd ed., SIAM, Philadelphia, 2010.

[8] Edelbaum, T. N., "Propulsion Requirements for Controllable Satellites," *ARS Journal*, Vol. 31, p. 1079, 1961.

[9] Oberth, H., *Wege zur Raumschiffahrt*, R. Oldenbourg Verlag, Munich, 1929.

[10] Racca, G. D., et al., "SMART-1 mission description and development status," *Planetary and Space Science*, Vol. 50, No. 14–15, pp. 1323–1337, 2002.

[11] McNutt, R. L., Jr., et al., "Near-term interstellar probe: First step," *Acta Astronautica*, Vol. 162, pp. 284–299, 2019.

[12] Fermi Explorer mission team, reference analysis, `https://fermiexplorer.github.io`, 2026.

A. Astrometric Inputs and Intercept Kinematics

This appendix documents the astrometric inputs for all four candidate targets, the coordinate transformation applied to them, and the closed-form kinematic results used throughout Section 2. The closed forms serve two purposes: they make the geometry transparent, and they provide exact internal checks on the numerical sweeps. Every quantity in this appendix was rederived from first principles in independently written code; the worst disagreement across all kinematic quantities was below 0.5%.

A.1 Astrometric Inputs and Provenance

For Alpha Centauri the orbital motion of the A and B components contaminates single-star astrometric solutions, so we adopt the orbit-corrected barycentric solution of Kervella et al. [1], cross-checked against the independent orbital study of [2]: distance 4.365 ly and space velocity $32.38\ \mathrm{km\,s^{-1}}$, resolved into a radial component $v_r = -22.40\ \mathrm{km\,s^{-1}}$ (approaching) and a transverse component $v_t = 23.38\ \mathrm{km\,s^{-1}}$. The two single stars, LSPM J2146+3813 and λ Serpentis, use Gaia DR3 astrometry [3] directly. α^2 Librae is too bright and too disturbed (it is itself a spectroscopic binary) for a clean Gaia solution; we use the Hipparcos re-reduction of van Leeuwen [4]. Its systemic radial velocity, $-22.0 \pm 5.8\ \mathrm{km\,s^{-1}}$, is by far the weakest input in Table 8; Section A.4 examines whether that uncertainty could alter the star's exclusion.

Table 8. Adopted astrometry for the four candidate targets. Distances and space motions are transformed to the heliocentric ecliptic frame as described in Section A.2. The published-encounter column lists the closest-approach solutions from the stellar-encounter literature, used as an external cross-check in Section A.4.

Target	Type	d (ly)	Astrometric source	Published encounter
Alpha Centauri AB	binary barycenter	4.365	Kervella et al. [1]	—
LSPM J2146+3813	M5V	22.99	Gaia DR3 [3]	0.568 pc at +82.5 kyr [6]
λ Serpentis	G0V	38.86	Gaia DR3 [3]	2.30 pc at +167 kyr [5]
α^2 Librae	spectroscopic binary	75.80	Hipparcos [4]	—

A.2 Transformation to the Heliocentric Ecliptic Frame

All intercept computations are performed in heliocentric ecliptic Cartesian coordinates, with x toward the vernal equinox and z toward the north ecliptic pole. Each catalog provides, in the ICRS equatorial frame, the right ascension α, declination δ, parallax ϖ, proper-motion components $\mu_{\alpha *} \equiv \mu_\alpha \cos\delta$ and μ_δ, and radial velocity v_r. The transformation proceeds in four steps.

1. *Distance.* $d = 1/\varpi$, with ϖ in arcseconds and d in parsecs. At the parallax precision of the sources in Table 8 the simple inversion is appropriate; no distance prior is needed.
2. *Position.* The unit line-of-sight vector in equatorial coordinates is $\hat{\mathrm{n}} = (\cos\delta\cos\alpha,\ \cos\delta\sin\alpha,\ \sin\delta)$, and the position vector is $\mathrm{r}_0 = d\,\hat{\mathrm{n}}$.
3. *Velocity.* With $\hat{\mathrm{e}}_\alpha$ and $\hat{\mathrm{e}}_\delta$ the local east and north unit vectors on the sky,
$$\mathrm{v}_\star \;=\; v_r\,\hat{\mathrm{n}} \;+\; \kappa\, d\,(\mu_{\alpha *}\,\hat{\mathrm{e}}_\alpha + \mu_\delta\,\hat{\mathrm{e}}_\delta)\,, \qquad \kappa = 4.740\ \mathrm{km\,s^{-1}}\ \mathrm{per}\ (\mathrm{arcsec\,yr^{-1}\,pc}), \tag{7}$$
where κ is the standard conversion between proper motion at distance d and transverse velocity.
4. *Rotation.* Both vectors are rotated from the equatorial to the ecliptic frame by a rotation about the common x-axis through the obliquity of the ecliptic, $\varepsilon \simeq 23.44°$.

The star is then propagated as $\mathrm{r}(t) = \mathrm{r}_0 + \mathrm{v}_\star t$. Perspective acceleration and Galactic tidal acceleration are neglected; both omissions are validated against the published encounter literature in Section A.4, and an internal check that integrates the target's motion through the Galactic potential leaves the intercept epoch essentially unchanged.

A.3 Closed-Form Intercept Kinematics

Lead-the-target relation and the minimum-speed identity. A probe cruising in a straight line at velocity V meets the star at time T if $\mathrm{V}T = \mathrm{r}_0 + \mathrm{v}_\star T$, i.e. $\mathrm{V}(T) = \mathrm{r}_0/T + \mathrm{v}_\star$. Writing $\mathrm{v}_\star = v_r\,\hat{\mathrm{n}} + v_t\,\hat{\mathrm{t}}$, with $\hat{\mathrm{t}}$ the unit vector transverse to the line of sight in the plane of the star's motion,

$$v_\infty^2(T) \;=\; \left(\frac{d}{T} + v_r\right)^2 + v_t^2. \tag{8}$$

For an approaching star ($v_r < 0$) the radial term vanishes at

$$T^* \;=\; \frac{d}{|v_r|}, \qquad v_\infty(T^*) \;=\; v_t, \tag{9}$$

so the minimum-speed intercept occurs when the star's own approach exactly closes the line-of-sight distance, and the minimum cruise speed equals the star's transverse velocity — an exact identity under linear motion, not an approximation. For Alpha Centauri, Equation 9 with the adopted astrometry gives $T^* = 58{,}422$ years and $v_\infty = 23.38\ \mathrm{km\,s^{-1}}$, reproducing the numerical sweep exactly; the identity thus serves as a built-in consistency check on the sweep machinery. The independent re-derivation reproduced these epochs to better than 0.5%.

Ecliptic-plane crossing. In ecliptic coordinates the star's out-of-plane position is $z(t) = z_0 + v_z t$, so the plane crossing (when it exists, i.e. when z_0 and v_z have opposite signs) occurs at

$$T_\times \;=\; -\frac{z_0}{v_z}. \tag{10}$$

For Alpha Centauri, which currently lies below the ecliptic and is moving north, Equation 10 gives $T_\times = 79{,}786$ years, at which epoch the star is 6.425 ly from the Sun.

Closest natural approach. Minimizing $s(t) = |\mathrm{r}_0 + \mathrm{v}_\star t|$ gives the epoch and distance of the star's closest approach to the Sun:

$$t_{\mathrm{ca}} \;=\; -\frac{\mathrm{r}_0\cdot\mathrm{v}_\star}{|\mathrm{v}_\star|^2} \;=\; \frac{d\,|v_r|}{|\mathrm{v}_\star|^2}, \qquad s_{\mathrm{ca}} \;=\; \frac{|\mathrm{r}_0\times\mathrm{v}_\star|}{|\mathrm{v}_\star|} \;=\; \frac{d\,v_t}{|\mathrm{v}_\star|}. \tag{11}$$

For Alpha Centauri, $s_{\mathrm{ca}} = 4.365 \times 23.38/32.38 = 3.152$ ly at $t_{\mathrm{ca}} = 27{,}955$ years.

Earth-relative hyperbolic excess. The departure is treated as a patched conic. To leave the solar system with asymptotic speed v_∞, the heliocentric speed required at Earth's distance is

$$v_{\mathrm{dep}} \;=\; \sqrt{v_\infty^2 + v_{\mathrm{esc},\odot}^2}, \qquad v_{\mathrm{esc},\odot}(1\,\mathrm{AU}) = 42.12\ \mathrm{km\,s^{-1}}. \tag{12}$$

When the departure asymptote is tilted an angle β out of the ecliptic with its in-plane projection aligned with Earth's orbital velocity $v_\oplus = 29.78\ \mathrm{km\,s^{-1}}$ (the alignment is set by choosing the departure epoch within the year), the law of cosines gives the required excess speed relative to Earth:

$$v_{\infty,E}^2 \;=\; v_{\mathrm{dep}}^2 + v_\oplus^2 - 2\,v_{\mathrm{dep}}\,v_\oplus\cos\beta. \tag{13}$$

At the departure optimum ($v_\infty = 23.64\ \mathrm{km\,s^{-1}}$, $\beta = 2.48°$), Equation 12 gives $v_{\mathrm{dep}} = 48.30\ \mathrm{km\,s^{-1}}$ and Equation 13 gives $v_{\infty,E} = 18.59\ \mathrm{km\,s^{-1}}$, matching the values quoted in Section 2.4.

Impulsive departure from low Earth orbit. A single tangential burn at perigee of a circular 400 km orbit, with circular speed $v_{\mathrm{circ}} = 7.67\ \mathrm{km\,s^{-1}}$ and local escape speed $v_{\mathrm{esc}} = 10.85\ \mathrm{km\,s^{-1}}$, captures the full Oberth benefit and requires

$$\Delta v \;=\; \sqrt{v_{\infty,E}^2 + v_{\mathrm{esc}}^2} \;-\; v_{\mathrm{circ}} \;=\; \sqrt{18.59^2 + 10.85^2} - 7.67 \;=\; 13.85\ \mathrm{km\,s^{-1}} \tag{14}$$

at the optimum. This is the impulsive floor of Table 4. Each quantity defining the optimum (T, v_∞, β, $v_{\infty,E}$, and Δv) was confirmed independently to better than 0.5%, and the optimum agrees with the Fermi Explorer reference analysis [12] ($13.88\ \mathrm{km\,s^{-1}}$ at 72,800 years, $2.4°$ tilt) to within 0.3%.

Miss tolerance and aiming dispersion. Requiring only a pass within $2600\,\mathrm{AU}$ of the barycenter, rather than a true intercept, reduces the required Earth-relative excess by approximately $83\ \mathrm{m\,s^{-1}}$; after dilution by the escape-energy term in Equation 14 this is worth only about $71\ \mathrm{m\,s^{-1}}$ of impulsive Δv. The same sphere admits an asymptote pointing error of $0.41°$, equivalent to roughly $168\ \mathrm{m\,s^{-1}}$ of departure velocity dispersion — the basis for the open-loop guidance conclusion of Section 2.6.

A.4 Consistency with Published Encounter Catalogs

The linear-motion model can be checked against the stellar-encounter literature, which derives closest-approach solutions by integrating stellar orbits through a Galactic potential with full propagation of astrometric uncertainty. Two of our targets have published solutions. For LSPM J2146+3813, Bailer-Jones [6] finds a perihelion of $0.568\,\mathrm{pc}$ at $+82.5\,\mathrm{kyr}$ — one of the tightest known future stellar encounters; our linear-motion perihelion agrees to 0.2%. For λ Serpentis, the published solution [5] gives $2.30\,\mathrm{pc}$ at $+167\,\mathrm{kyr}$; ours agrees to 1.4%. Agreement at the percent level with orbit-integrated solutions confirms that straight-line kinematics is adequate at these ranges over 10^5-year horizons, and validates the neglect of perspective and Galactic accelerations in Section A.2.

For α^2 Librae no comparably tight published solution exists, and its systemic radial velocity is uncertain at the $\pm 5.8\ \mathrm{km\,s^{-1}}$ level. Its exclusion from the target set is nonetheless robust: across the $\pm 2\sigma$ range of that radial velocity, the minimum-speed intercept epoch of Equation 9 remains outside the 500,000-year mission horizon (the nominal value is 1.03 Myr), and forcing an arrival inside the horizon requires an impulsive departure of $19.74\ \mathrm{km\,s^{-1}}$ — worse than Alpha Centauri itself. The exclusion therefore does not depend on the quality of the weakest astrometric input.

The two published-encounter cross-checks above concern the targets carried forward into the analysis; the systematic sweep of the full 61-row encounter catalog that produced the candidate list — its filters, exclusions, and ranking metric — is documented in Section E.7.

B. Trajectory Models and Verification

This appendix documents the trajectory models behind the departure and heliocentric results of Section 3 and Section 4. It covers the control policy and optimization of the perihelion-pumping escape; the bounding analysis for the heliocentric Δv requirement, comprising a certified lower bound, constructive upper bounds at two flight-time horizons, and a best-found reference anchor; and the independent verification applied to each headline number.

B.1 Equations of Motion and Thrust Law

The heliocentric phase is modeled as planar two-body motion about a point-mass Sun with a power-throttled, steerable thrust:

$$\ddot{\mathrm{r}} = -\mu_\odot \frac{\mathrm{r}}{r^3} + \frac{F(r)}{m}\,\hat{\mathrm{u}}, \tag{15}$$

$$F(r) = F_1 \ \min\left[\left(\frac{r_1}{r}\right)^2, 4\right], \tag{16}$$

$$\dot{m} = -\frac{F(r)}{g_0\, I_{\mathrm{sp}}}, \tag{17}$$

where r is the heliocentric position with magnitude r, $\mu_\odot$ is the solar gravitational parameter, m is the vehicle mass, F_1 is the thrust at $r_1 = 1$ AU, g_0 is standard gravity, and I_{sp} is the specific impulse. The thrust direction $\hat{\mathrm{u}}$ is restricted to prograde, retrograde, or off (coast) along the velocity vector, as commanded by the control policy described below. This three-state law at full envelope thrust is the policy class of the constructive schedule; the bounding analysis later in this appendix embeds it in a broader throttleable class, over which the heliocentric minimum is defined. The factor-of-four cap in Equation 16 represents the array thermal limit, engaging inside roughly 0.5 AU. The characteristic parameter is the initial 1-AU thrust acceleration $a_0 \equiv F_1/m_0$, with m_0 the initial mass; because thrust at fixed power is independent of mass, the instantaneous acceleration grows as m_0/m while propellant depletes, and the integrated results include this effect.

B.2 Earth-Escape Spiral

The geocentric escape phase is not integrated revolution-by-revolution; it is budgeted by the standard low-thrust result that continuous-thrust escape from a circular orbit costs approximately the circular orbital speed of the starting orbit [8]. From a 400 km orbit ($v_{\mathrm{circ}} = 7.67\ \mathrm{km\,s^{-1}}$) the budget carried is 7.6 $\mathrm{km\,s^{-1}}$, expended over approximately 1.5 years. The spiral transits the radiation belts, and the array is sized with a degradation factor of 1.25 for this phase. A vehicle escaping by slow spiral leaves Earth's sphere of influence with negligible Earth-relative excess speed, so the full heliocentric requirement is built subsequently under Equation 15–Equation 17.

B.3 Perihelion-Pumping Control Policy and Optimization

The pumping escape is generated by a four-parameter feedback policy operating on the osculating orbit:

(i) the target perihelion radius, to which retrograde thrusting drives the orbit (the optimizer selects approximately 0.42 AU, just inside the thermal-cap radius);

(ii) the fraction of each revolution, centered on aphelion, over which retrograde thrust is applied while the perihelion is being lowered;

(iii) the heliocentric radius bounding the prograde burn window, so that accelerating arcs fire only near perihelion where power is at or above the 1-AU rating;

(iv) the specific-orbital-energy threshold at which pumping terminates and the vehicle thrusts continuously to escape.

The four parameters were optimized numerically — a coarse parameter sweep followed by local refinement — with total Δv to reach the required cruise speed of 23.64 $\mathrm{km\,s^{-1}}$ as the objective, at $a_0 = 2.5 \times 10^{-4}\ \mathrm{m\,s^{-2}}$. Concretely, the search drew 400 seeded uniform-random candidates over the parameter box and polished the five best by Nelder–Mead simplex descent; an independently seeded re-optimization with a larger budget reproduced the same schedule and carried 23.98 $\mathrm{km\,s^{-1}}$ from its most conservative timestep rung. The trajectories themselves are integrated by a fixed-step kick-drift-kick scheme with the thrust policy frozen across each step, a step law of $\max(600\ \mathrm{s},\ 0.002 \times$ local circular period) scaled by the refinement factor, and a rocket-equation mass update on thrusting steps; on every accepted schedule the integrated Δv matches the rocket-equation value implied by the final mass to $10^{-13}\ \mathrm{km\,s^{-1}}$. Seeds, acceptance gates, and both optimization records are detailed in Section E.2. The best schedule (Figure 5) delivers the cruise speed at a Δv of 23.97 $\mathrm{km\,s^{-1}}$ over 12.0 years of flight time, of which only about 1.3 years is cumulative thrusting (0.9 years retrograde, 0.4 prograde), with a minimum solar distance of 0.42 AU. The final mass fraction of 0.418 agrees with the rocket equation evaluated at $I_{\mathrm{sp}} = 2800$ s, an internal consistency check on the mass integration. Convergence was established by timestep refinement, with one caveat to disclose: at one-half the production step the integration initially terminated at the simulation-horizon cap short of cruise speed; with the horizon extended, the half-step run reaches cruise speed at a Δv of 23.969 $\mathrm{km\,s^{-1}}$, and the quarter-step run at 23.985 $\mathrm{km\,s^{-1}}$ — agreement within 0.07%, with the residual refinement drift covered by carrying the upper bound at 24.0 $\mathrm{km\,s^{-1}}$.

B.4 Bounding the Heliocentric Requirement

Let $\Delta v^*(a_0)$ denote the minimum heliocentric Δv that reaches the required cruise speed (heliocentric hyperbolic excess of 23.64 $\mathrm{km\,s^{-1}}$) from the post-escape state at 1 AU under Equation 15–Equation 17 with initial 1-AU thrust acceleration a_0. The minimum is taken over the throttleable admissible class: piecewise-continuous thrust programs with free in-plane direction, magnitude commandable anywhere in the power envelope, $0 \le F \le F_1 \min[(r_1/r)^2,\ 4]$, and heliocentric distance $r(t) \ge 0.42$ AU — the minimum solar distance carried by the thermal design, which the constructive schedule respects ($r_{\min} = 0.423$ AU). Throttleability is physically justified — a flight electric-propulsion string throttles continuously through its power processing unit — and the three-state bang-off policy class of the preceding subsection is a subset of this class, so every constructive result above is admissible a fortiori. The class definition places no constraint on flight time; where a result below instead refers to the 12-year operational horizon carried by the mission design, that restriction is stated explicitly. The report carries four quantities for $\Delta v^*(2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}})$, of different evidentiary standing: a certified lower bound, an impulsive floor retained as a consistency reference, constructive upper bounds at two flight-time horizons, and a best-found reference anchor.

Monotonicity on the throttleable class. Δv^* is non-increasing in a_0. Under a larger acceleration bound $a_0' > a_0$, any admissible throttleable program at bound a_0 remains admissible: its commanded thrust magnitude at every instant lies inside the enlarged envelope and is realized by scaling the throttle setting by a_0/a_0' relative to full thrust. The trajectory, the mass history, and the expended Δv are unchanged, so the feasible set only grows with a_0 and the minimum cannot increase. The argument requires throttleability — a bang-off class restricted to full envelope thrust does not contain the rescaled program — which is why the admissible class is stated explicitly above.

Certified lower bound. Every trajectory in the admissible class expends a heliocentric Δv of at least 16.56 $\mathrm{km\,s^{-1}}$:

$$\Delta v^*(2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}) \ \ge\ 16.56\ \mathrm{km\,s^{-1}}. \tag{18}$$

This is a proof-grade bound on the true class optimum — not the cost of any particular trajectory — established by a six-step bounding argument in which every step is a closed-form inequality: (i) the trajectory is truncated at first attainment of the target energy, and propellant accounting converts the mass history into an explicit ceiling on achievable thrust acceleration as a function of the total Δv expended; (ii) the final escape leg is shown to make exactly one transit of the region near the perihelion floor — the available thrust falls an order of magnitude short of what lingering there would require; (iii) the duration of that single transit is bounded by a comparison argument, capping both the Δv deliverable and the orbital energy gained inside the transit window; (iv) every remaining increment of orbital energy must then be purchased at a price no better than the reciprocal of the fastest admissible speed near the perihelion floor, and integrating that price over the required energy gain bounds the cost of the escape leg; (v) the initial segment, ending when the orbit first reaches escape energy, is bounded below by $(\sqrt{2}-1)$ times Earth's orbital speed, about 12.3 $\mathrm{km\,s^{-1}}$, by a separate self-contained argument; (vi) the pieces close through a self-consistency (fixed-point) check on the assumed total: if the total Δv were below 16.56 $\mathrm{km\,s^{-1}}$, the thrust ceiling of step (i) would be tight enough that steps (ii)–(v) force a total of at least 16.5659 $\mathrm{km\,s^{-1}}$ — a contradiction. Because the admissible class carries no flight-time constraint, the bound holds a fortiori for every time-restricted subclass, including the 12-year operational horizon; like the class itself, it is a statement about the planar model of Equation 15. The argument was audited adversarially step by step; every inequality in the chain was additionally tested against three independently integrated trajectories (21 falsification checks, zero violations), and the assembled constants were recomputed from first principles, matching the derivation to better than 10^{-5} relative. The complete chain and the numerical harness applied to it are documented in Section E.5.

Impulsive floor (consistency reference). Removing the thrust-magnitude bound altogether ($a_0 \to \infty$) only enlarges the feasible set, so Δv^* is also bounded below by the impulsive multi-burn minimum over the same class, whose $r \geq 0.42$ AU constraint is what makes the floor nontrivial (without any perihelion bound the impulsive floor decays to the 1-AU escape cost, about 12.3 $\mathrm{km\,s^{-1}}$). That minimum is the bi-parabolic transfer: an escape burn at 1 AU costing $v_{\mathrm{esc}} - v_{\mathrm{circ}} \approx 12.3\ \mathrm{km\,s^{-1}}$, a vanishing redirection burn at great distance, and a periapsis burn at 0.42 AU costing $\sqrt{v_\infty^2 + v_{\mathrm{esc}}^2(r_p)} - v_{\mathrm{esc}}(r_p) \approx 4.2\ \mathrm{km\,s^{-1}}$ at $v_\infty = 23.64\ \mathrm{km\,s^{-1}}$ — a floor of 16.5 $\mathrm{km\,s^{-1}}$ (15.6 $\mathrm{km\,s^{-1}}$ if the perihelion bound is relaxed to 0.25 AU). The certified finite-thrust bound of Equation 18 exceeds this floor by 0.06 $\mathrm{km\,s^{-1}}$, as any sound bound must.

Constructive upper bounds at two horizons. Two explicitly integrated admissible trajectories bracket the optimum from above, at different flight-time horizons. At the 12-year operational horizon — the component-life and burn-sizing basis of the mission design — the optimized pumping schedule of the preceding subsection reaches the cruise speed at 23.97 $\mathrm{km\,s^{-1}}$, carried as 24.0 $\mathrm{km\,s^{-1}}$; an independent re-optimization at the same horizon reproduced 23.98 $\mathrm{km\,s^{-1}}$. With the flight-time restriction removed, a substantially cheaper admissible schedule exists: a long-horizon trajectory of the same pumping family, integrated at four timestep resolutions and quoted at the most conservative value across them, reaches the same cruise speed at 18.87 $\mathrm{km\,s^{-1}}$ with about 58 years of flight (the family of such schedules found spans 33–58-year flight times). Both figures are the best schedules found, not optima. Combined with Equation 18, the unconstrained-flight-time optimum is bracketed as

$$16.56\ \mathrm{km\,s^{-1}} \ \leq\ \Delta v^*(2.5 \times 10^{-4}) \ \leq\ 18.87\ \mathrm{km\,s^{-1}}, \tag{19}$$

while the 12-year-constrained optimum is known only to lie between 16.56 and 24.0 $\mathrm{km\,s^{-1}}$ — the lower bound does not exploit the time limit, and how much of the 18.87-to-24.0 gap is recoverable within 12 years is open.

Best-found reference anchor. At a higher thrust acceleration the optimal-control problem is numerically tractable, and the following problem was solved as a reference point: planar two-body dynamics per Equation 15 with the power-throttled thrust envelope per Equation 16, initial 1-AU thrust acceleration $a_0' = 1 \times 10^{-3}\ \mathrm{m\,s^{-2}}$ (four times the design value), departure from the post-escape heliocentric state at 1 AU with negligible hyperbolic excess, terminal condition equal to the required cruise speed, flight time fixed at 0.82 years (matching the elapsed time of a pumping heuristic at the same target; with flight time left free this problem degenerates toward the impulsive limit), and expended Δv as the objective. An indirect method based on Pontryagin's minimum principle (costate integration with shooting on the boundary conditions) converges cleanly to an extremal expending 22.66 $\mathrm{km\,s^{-1}}$ at the 23.64 $\mathrm{km\,s^{-1}}$ cruise speed; the same solver bit-for-bit reproduces the previously computed 22.94 $\mathrm{km\,s^{-1}}$ extremal at a 24.14 $\mathrm{km\,s^{-1}}$ target variant. The solve imposes no perihelion-floor path constraint, but the extremal never descends below 0.62 AU, so the 0.42 AU floor is inactive. A direct-collocation cross-check of this extremal was attempted and did not converge: attempts at 24, 48, and 60 mesh segments either stalled at their iteration limits near the initializing heuristic — some 18% above the indirect value — or failed re-integration checks. The 22.66 $\mathrm{km\,s^{-1}}$ value therefore carries single-method provenance and is labeled best-found: an extremal satisfying first-order optimality conditions whose global optimality is unproven — a best-found upper reference on the optimum at the higher thrust level, not a bound on anything at the design thrust level. No conclusion in this report rests on it; it is carried only as a plausibility reference for what a stronger thruster buys.

What the report's conclusions rest on. The certified statement is the lower bound of Equation 18, paired with the demonstrated upper bounds of 18.87 $\mathrm{km\,s^{-1}}$ (unconstrained flight time) and 24.0 $\mathrm{km\,s^{-1}}$ (12-year horizon). By the closure algebra of Appendix C, a 100 kg vehicle from low Earth orbit closes if a heliocentric trajectory below the ≈22.3 $\mathrm{km\,s^{-1}}$ crossover (22.26 $\mathrm{km\,s^{-1}}$) exists. That threshold is crossed by unconstrained-flight-time solutions — the demonstrated 18.87 $\mathrm{km\,s^{-1}}$ schedule sits 3.4 $\mathrm{km\,s^{-1}}$ below it, at the price of 33–58 years of flight — but by no known 12-year solution: the best 12-year schedule sits 1.7 $\mathrm{km\,s^{-1}}$ above the crossover, and nothing between 16.56 and 24.0 $\mathrm{km\,s^{-1}}$ is excluded for the 12-year class. The closure statements in the body of the report are therefore scoped to the 12-year operational horizon, with the unconstrained-flight-time result stated alongside them as the demonstrated long-duration alternative.

B.5 Independent Verification

Each headline trajectory result was independently replicated from first principles, without access to the baseline implementation and with the explicit objective of refuting the claimed value; the figures quoted below are drawn from that replication.

- *Departure kinematics.* The optimal arrival epoch, aim tilt, Earth-relative hyperbolic excess, and the 13.85 $\mathrm{km\,s^{-1}}$ impulsive floor were all reproduced, with a worst-case discrepancy below 0.5%. The same quantities agree with the Fermi Explorer reference analysis [12] to within 0.3%.
- *Outward-spiral saturation.* The power-wall ceiling of Figure 4 was reproduced by two independently written trajectory integrators, agreeing to 2.7%.
- *Perihelion pumping.* The mechanism was reconstructed independently: a from-scratch implementation designed its own pumping schedule and reached a cruise speed of 23.60 $\mathrm{km\,s^{-1}}$ at a Δv of 28.2 $\mathrm{km\,s^{-1}}$ over 4.9 years — the fast variant discussed in Section 4 — confirming both the breaking of the power wall at $a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$ and the cost of the impatient profile. The optimized 12-year schedule reaches the same cruise speed at 23.97 $\mathrm{km\,s^{-1}}$.
- *Heliocentric lower bound.* Beyond the adversarial step-by-step audit of the derivation, every inequality in the six-step chain was tested numerically against three independently integrated trajectories — 21 falsification checks in all, with zero violations — and the constants entering the final fixed-point evaluation were recomputed from first principles in a separate harness, agreeing with the derivation to better than 10^{-5} relative.

B.6 A Note on Direct Collocation

Direct transcription of the low-thrust escape at the design thrust acceleration — Hermite–Simpson collocation over the full heliocentric phase — was attempted and did not converge, a failure reported plainly here because its mode is structural rather than incidental. The optimal solution alternates coast arcs lasting months with perihelion passes whose dynamical timescale is days: a mesh fine enough to resolve the fast passes, repeated across many revolutions, grows to an intractable problem size, while a coarser mesh aliases exactly the burn arcs on which the solution depends. This difficulty of direct methods in many-revolution low-thrust problems is well documented [7]. Direct collocation also failed at the four-times-higher anchor thrust level of the bounding analysis above, where the revolution count is small: repeated attempts across mesh resolutions stalled near their initializing heuristic rather than descending to the extremal found by the indirect method, which is why the anchor value carries single-method provenance. No conclusion in this report depends on direct collocation: the heliocentric budget rests on the certified lower bound and the explicitly integrated constructive schedules, and the anchor is carried as a best-found reference with its caveat attached.

C. Mass-Closure Model

This appendix documents the vehicle mass-closure model summarized in Section 5: the governing equations and their coefficients, the algebra behind the sharp threshold in closing wet mass, the full closure grid, and the independent verification of the arithmetic.

C.1 Model Equations

Let m_w denote the wet (launch) mass, $I_{\rm sp}$ the thruster specific impulse, and $v_e = g_0 I_{\rm sp}$ the exhaust velocity, with g_0 the standard gravitational acceleration. For $I_{\rm sp} = 2800\,\mathrm{s}$, $v_e = 27.46\,\mathrm{km\,s^{-1}}$.

Δv budget. The total budget is the sum of a departure leg and a heliocentric leg,

$$\Delta v_{\rm tot} = \Delta v_{\rm dep} + \Delta v_{\rm hel}, \qquad \Delta v_{\rm hel} \in [22.9,\ 24.0]\,\mathrm{km\,s^{-1}}, \tag{20}$$

where the heliocentric interval is the perihelion-pumping reference interval of the trajectory analysis (Section 3): its upper end a verified constructive bound, its lower end a best-found reference anchor. For a LEO start, $\Delta v_{\rm dep} = 7.6\,\mathrm{km\,s^{-1}}$: a slow low-thrust escape spiral from a circular orbit costs approximately the initial circular speed [8], here evaluated with eclipse and steering effects included. For a GTO start, the departure leg is computed from a standard $250 \times 35{,}786\,\mathrm{km}$ geostationary transfer orbit, whose semi-major axis is $a = 24{,}396\,\mathrm{km}$. The energy-based low-thrust escape cost is the local circular speed at that semi-major axis, with a 5% multiplicative allowance covering eclipse outages and steering losses:

$$\Delta v_{\rm dep} \ = \ 1.05\sqrt{\mu_\oplus/a} \ = \ 1.05 \times 4.04\,\mathrm{km\,s^{-1}} \ = \ 4.24\,\mathrm{km\,s^{-1}}, \tag{21}$$

with $\mu_\oplus$ the Earth gravitational parameter. Because this leg is modeled rather than explicitly integrated, a $\pm10\%$ sensitivity is carried through the closure algebra: raising the leg to $4.66\,\mathrm{km\,s^{-1}}$ lowers the GTO-100 margins of Table 9 by 0.6–0.7 percentage points (about 1.5 points per $\mathrm{km\,s^{-1}}$ of departure-leg growth), so the GTO-100 case still closes at both ends of the reference interval. The interval ends correspond to total budgets of 30.5 and $31.6\,\mathrm{km\,s^{-1}}$ from LEO, and 27.14 and $28.24\,\mathrm{km\,s^{-1}}$ from GTO.

Propellant. The rocket equation gives the propellant fraction and mass,

$$\zeta \ = \ \frac{m_p}{m_w} \ = \ 1 - \exp\!\left(-\frac{\Delta v_{\rm tot}}{v_e}\right), \qquad m_p = \zeta\, m_w. \tag{22}$$

Phase-split burn-time sizing. The thrust acceleration a_0—defined at wet mass and at the 1 AU power rating—is a design variable bounded below by an operations constraint on cumulative thrusting time. Splitting the budget into its two phases, the propellant fractions of the initial wet mass expended in each phase are

$$\zeta_{\rm dep} = 1 - e^{-\Delta v_{\rm dep}/v_e}, \qquad \zeta_{\rm hel} = \left(1 - \zeta_{\rm dep}\right)\left(1 - e^{-\Delta v_{\rm hel}/v_e}\right), \tag{23}$$

and the cumulative burn time is modeled as

$$t_b \ = \ \left(K_{\rm dep}\,\zeta_{\rm dep} + K_{\rm hel}\,\zeta_{\rm hel}\right)\frac{v_e}{a_0}, \tag{24}$$

where the ideal constant-thrust burn time for a phase is $\zeta\, v_e/a_0$ and the dimensionless factors K correct it for the realities of each phase: $K_{\rm dep} = 1.54$ for the LEO spiral (the reciprocal of a 0.65 eclipse duty cycle) and 1.2 for the GTO spiral; $K_{\rm hel} = 1.84$ for the pumped heliocentric phase, calibrated against the integrated pumping trajectories, whose burn arcs are localized near perihelion where available power meets or exceeds the 1 AU rating. The model sets a_0 to the smallest value for which $t_b \le 4\,\mathrm{yr}$; the resulting $a_0 \approx 2.5{\times}10^{-4}\,\mathrm{m/s^2}$ is consistent with the optimized pumping profile, which accumulates roughly 1.3 years of thrust within its twelve-year flight duration. Because Eq. (24) involves only mass *fractions*, the sized a_0 is independent of m_w—a property the closure algebra below exploits.

Power and subsystem masses. Thrust and power at the 1 AU rating follow from

$$F = a_0\, m_w, \qquad P = \frac{F\, v_e}{2\eta}, \tag{25}$$

with thruster total efficiency $\eta = 0.55$. The subsystem masses are

$$m_{\rm arr} = f_{\rm rad}\,\frac{P}{\alpha}, \qquad m_{\rm pp} = \kappa\, P, \qquad m_{\rm tank} = f_t\, m_p, \tag{26}$$

where $\alpha = 60$ W/kg is the system-level array specific power, $f_{\rm rad}$ the radiation-degradation factor on the array (1.25 for the LEO start with its long belt transit; 1.15 for the shorter GTO transit), $\kappa = 6$ kg/kW the power-processing specific mass, and $f_t = 0.12$ the xenon tankage fraction. The fixed mass is $m_{\rm fix} = 9$ kg (1 kg payload plus 8 kg bus). The design closes when the required mass fits under the wet mass; the margin fraction is

$$f_m \;=\; \frac{m_w - m_{\rm req}}{m_w}, \qquad m_{\rm req} = m_{\rm fix} + m_p + m_{\rm tank} + m_{\rm arr} + m_{\rm pp}, \tag{27}$$

with closure defined by $f_m \geq 0$. Thrust and power are sized at the 1 AU rating; the variation of available power along the trajectory is the province of the trajectory analysis, not the mass model.

C.2 Closure-Boundary Algebra

Every term in $m_{\rm req}$ except the fixed 9 kg scales linearly with wet mass: m_p and $m_{\rm tank}$ through the propellant fraction ζ, and $m_{\rm arr}$ and $m_{\rm pp}$ through $P = a_0 m_w v_e/(2\eta)$ with a_0 independent of m_w. The requirement therefore collapses to

$$m_{\rm req} = m_{\rm fix} + \varphi\, m_w, \qquad \varphi \;=\; \zeta\,(1+f_t) \;+\; \frac{a_0 v_e}{2\eta}\left(\frac{f_{\rm rad}}{\alpha} + \kappa\right), \tag{28}$$

where φ is the *variable mass fraction*: the kilograms of propellant, tankage, array, and power processing required per kilogram of wet mass. The margin becomes $f_m = 1 - \varphi - m_{\rm fix}/m_w$, so closure requires

$$\varphi < 1 \qquad \text{and} \qquad m_w \;\geq\; \frac{m_{\rm fix}}{1-\varphi}. \tag{29}$$

The boundary is sharp because each additional kilogram of wet mass buys only $(1-\varphi)$ kilograms of net margin. From LEO, $1-\varphi$ is 0.06–0.08, so the 9 kg of fixed mass forces the wet mass above roughly 110–150 kg before the margin turns positive. And if $\varphi \geq 1$—which occurs when the thrust acceleration is roughly doubled, since the power chain in Eq. (28) is proportional to a_0—no wet mass closes at any size.

Worked examples at the interval ends make the boundary concrete. For the LEO start ($I_{\rm sp} = 2800$ s, $v_e = 27.46\ {\rm km\,s^{-1}}$): at $\Delta v_{\rm tot} = 30.5\ {\rm km\,s^{-1}}$, Eq. (22) gives $\zeta = 0.671$, so the propellant-plus-tankage term is $\zeta(1+f_t) = 0.751$; the burn-time-sized a_0 makes the power chain $a_0 v_e/(2\eta) \approx 6.3$ W per kilogram of wet mass, and the combined array-plus-processing specific mass is $f_{\rm rad}/\alpha + \kappa = 0.0268$ kg/W, contributing 0.169; hence $\varphi = 0.921$ and $m_w \geq 9/(1-\varphi) = 113.3$ kg. At $\Delta v_{\rm tot} = 31.6\ {\rm km\,s^{-1}}$, $\zeta = 0.684$, the power chain contributes 0.173, $\varphi = 0.938$, and the boundary rises to 145.6 kg. These reproduce the 113.5–146 kg range quoted in Section 5, the small differences reflecting the finite step of the design sweep. The same algebra for the GTO start gives $\varphi \approx 0.85$–0.87 and a closure boundary near 59–68 kg, comfortably below the smallest (80 kg) case examined—which is why every GTO entry in the grid closes.

C.3 Closure Grid

Table 9 gives the full closure grid: both departure modes, wet masses of 80–150 kg, and both ends of the heliocentric Δv reference interval. Margins are in percentage points of wet mass (Eq. (27)); a design closes when its margin is non-negative. The grid exhibits the structure derived above: from LEO the 100 kg vehicle misses everywhere in the interval (by 1.0–2.8 points) and 150 kg closes everywhere; from GTO every case closes, with margins that survive the conservative end of the interval.

C.4 Independent Verification

The closure arithmetic was replicated independently, from the model statement alone, by a separate verification analysis. The replication reproduced the closure boundary algebraically at 113.3 kg—in agreement with

Table 9. Mass-closure grid. m_p is propellant mass and P the required electrical power at 1 AU; the margin is the closure margin f_m in percentage points of wet mass, with negative values indicating failure to close. Common parameters: $I_{\rm sp} = 2800\,$s, $\eta = 0.55$, $\alpha = 60\,$W/kg, $\kappa = 6\,$kg/kW, $f_t = 0.12$, 1 kg payload plus 8 kg bus. Departure legs: LEO $7.6\,\mathrm{km\,s^{-1}}$, GTO $4.24\,\mathrm{km\,s^{-1}}$.

		$\Delta v_{\rm hel} = 22.9\,\mathrm{km\,s^{-1}}$			$\Delta v_{\rm hel} = 24.0\,\mathrm{km\,s^{-1}}$		
Departure	m_w (kg)	m_p (kg)	P (kW)	margin	m_p (kg)	P (kW)	margin
LEO	80	53.7	0.50	−3.29	54.7	0.52	−5.08
LEO	100	67.1	0.63	−1.04	68.4	0.64	−2.83
LEO	120	80.5	0.76	+0.46	82.0	0.77	−1.33
LEO	150	100.6	0.95	+1.96	102.5	0.97	+0.17
GTO	80	50.2	0.46	+3.90	51.4	0.47	+1.90
GTO	100	62.8	0.58	+6.15	64.2	0.59	+4.15
GTO	120	75.3	0.69	+7.65	77.1	0.71	+5.65
GTO	150	94.2	0.87	+9.15	96.4	0.89	+7.15

Eq. (29) above and within 0.2% of the design-sweep value—and confirmed the two pivotal grid entries at the lower end of the interval: a 100 kg vehicle from LEO misses closure by approximately one kilogram, while 120 kg closes. The same check also exercised the no-close condition of Eq. (29): sizing the thrust acceleration at roughly twice the nominal value drives the variable mass fraction φ above unity, so that no wet mass of any size closes.

The closure results also inherit the pedigree of their Δv inputs, and that pedigree is asymmetric. The upper end of the heliocentric reference interval, $24.0\,\mathrm{km\,s^{-1}}$, is a verified constructive bound: an explicitly constructed, optimized pumping schedule whose integration was timestep-refined. The lower end, $22.9\,\mathrm{km\,s^{-1}}$, is the best trajectory found by optimal-control analysis at a higher thrust acceleration ($a_0' = 10^{-3}\,\mathrm{m/s^2}$), obtained by a single method—an indirect Pontryagin solution whose direct-collocation cross-check did not converge—and extended down in thrust by a monotonicity argument valid for throttleable thrust. It is a best-found reference anchor, not a certified bound; the certified lower bound is $16.56\,\mathrm{km\,s^{-1}}$, proved for the throttleable class with unconstrained flight time at the 0.42 AU perihelion bound (superseding the $16.50\,\mathrm{km\,s^{-1}}$ impulsive multi-burn limit). Closure margins quoted across the interval are therefore conditional at its lower edge: a 100 kg LEO vehicle would close if a heliocentric trajectory below the $\approx 22.3\,\mathrm{km\,s^{-1}}$ crossover exists within the 12-year operational horizon, and no trajectory found in this analysis, once referred to the design thrust level, comes within $0.6\,\mathrm{km\,s^{-1}}$ of that crossover. The pumping mechanism itself was reconstructed independently by a separate analysis; the closure conclusions of Section 5 rest on no single computation.

D. Cost Model Detail

This appendix documents the program cost model summarized in Section 7: the line-item values for the two primary architectures, the definitions of the cost drivers, the thermal-qualification increment for perihelion passes, the statistical-independence assumption and its directional effect, and what the estimate deliberately excludes.

D.1 Line-Item Values

Table 10 lists the nominal line-item values, in millions of 2026 dollars, for the two primary mass-closing architectures: the low-Earth-orbit start at 150 kg wet mass and the recommended geostationary-transfer-orbit start at 100 kg. The model carries fourteen triangular line items; closely related sub-items are aggregated into the ten groups shown, and the rows are organized under the mission's three budget categories.

Table 10. Nominal cost line items (millions of 2026 dollars) for the two primary architectures, grouped by mission budget category. Columns may not sum exactly to the subtotals and totals because of rounding.

Line item	Cost driver	LEO start, 150 kg	GTO start, 100 kg
Launch			
Rideshare launch	wet mass	0.98	0.65
Escape-injection premium	wet mass	—	0.61
Subtotal (allocation $3.0M)		*0.98*	*1.26*
Propulsion			
Thruster and power-processing string	input power	0.58	0.36
Xenon propellant, loaded	propellant mass	0.41	0.26
Tankage and feed system	propellant mass	0.21	0.13
Subtotal (allocation $3.0M)		*1.19*	*0.74*
Bus and mission operations			
Spacecraft bus	fixed	2.00	2.00
Solar arrays	array wattage	0.39	0.24
Integration, test, and qualification	fixed	2.45	2.45
Powered-flight operations	burn duration	4.80	4.80
Program margin and management	fixed	2.50	2.50
Subtotal (allocation $3.0M)		*12.14*	*11.99*
Total, nominal		14.31	13.99
Total, median (P50), 20,000-draw Monte Carlo		16.60	16.19

The Monte Carlo median exceeds the sum of nominal values because the triangular uncertainty bands are right-skewed: for most hardware and services line items the credible upside in price is larger than the credible downside, so the median of the summed distribution lies above the sum of the modes.

D.2 Cost-Driver Definitions

Each line item is priced per unit of a physical driver taken from the mass-closure model output for the architecture being costed. The drivers are:

- **Fixed.** Independent of vehicle scale: the spacecraft bus, the integration-test-and-qualification campaign, and the program margin and management line.
- **Wet mass.** Sets the rideshare launch line and the escape-injection premium. The premium depends on the departure orbit: zero for a low-Earth-orbit drop-off, a factor of 0.45 for a geostationary-transfer-orbit drop-off (roughly the energetic halfway point to escape), and the full amount when the launch vehicle provides escape injection directly.

- **Input power.** The thruster and power-processing string is priced per kilowatt of electric-propulsion input power at 1 AU.
- **Array wattage.** The solar arrays are priced per watt of array output at 1 AU.
- **Propellant mass.** The loaded xenon and the tankage and feed system are each priced per kilogram of xenon carried.
- **Powered-flight duration.** Operations cost is priced per year of powered flight, with a floor of one year. As noted in Section 7.4, the model prices approximately four years of powered-flight operations, whereas the patient pumping profile that closes the mass budget implies roughly thirteen years of custody, so under staffed-operations assumptions the operations line is a lower bound.

Because these drivers come directly from the mass-closure model, the cost estimate inherits that model's verification status: the closure arithmetic supplying the wet mass, propellant load, and power level was replicated algebraically by an independent check.

D.3 Perihelion Thermal-Qualification Increment

Architectures that employ perihelion pumping make repeated solar passes down to 0.42 AU. Their integration-test-and-qualification line therefore carries a multiplier of 1.4 relative to a benign-cruise mission, representing a high-temperature qualification campaign for the arrays and structure of the kind flown on previous inner-heliosphere missions. The values tabulated in Table 10 include this multiplier. The increment follows the heliocentric trajectory, not the departure mode: an escape-injection architecture that still pumps at perihelion retains it in full.

D.4 Independence Assumption

The Monte Carlo total is the sum of statistically independent triangular draws, one per line item; the (low, nominal, high) parameters of each triangle, the sampling seed, and the closure scan supplying the cost drivers are recorded in Section E.7 (Table 12). Independence has two consequences. First, it makes the variance decomposition exact—the variance of the total is exactly the sum of the line-item variances—which is what makes the ranking of cost drivers reported in Section 7.3 well defined. Second, it is mildly optimistic in the tails: real programs exhibit positive correlation between line items (a schedule slip inflates operations and integration together; a propellant-market shock co-moves with launch demand), and positive correlation widens the total-cost distribution—raising the 90th percentile and lowering the probability of remaining under $10M—while leaving the mean unchanged. The reported percentile ranges should therefore be read with this bias in mind; the program-margin line partially compensates by carrying program-level, correlated risk as a band of its own.

D.5 Scope Exclusions

The estimate prices the flight segment, its launch, and its powered-flight operations. It excludes development of the 1 kg payload itself, which is priced as furnished; monitoring or custody of the vehicle after the end of powered flight, including any attempt to maintain contact during the interstellar coast; ground-segment capital beyond the tracking and operations services carried in the operations rate; launch insurance; and technology development beyond the first-unit qualification increments already carried in the propulsion and integration lines. None of these exclusions alters the structural conclusion of Section 7: the cost problem is concentrated in the bus, operations, and margin category, and its resolution is programmatic rather than physical.

E. Computational Campaign and Verification Record

The preceding appendices document the physics: the kinematic closed forms, the trajectory models and their bounding analysis, the mass-closure algebra, and the cost structure. This appendix documents the computations themselves — the search algorithms, sample counts, random seeds, timestep-refinement ladders, acceptance gates, and runtimes behind every headline number, together with the independent checks applied to each and the negative results retained alongside the positive ones. It records enough of each computation — inputs, seeds, gates, and independent checks — that every figure in this report can be reproduced from what is documented here, and shows exactly which computations were verified by which independent means.

E.1 Campaign Structure and Reproducibility Conventions

The analysis was executed as a series of self-contained computational campaigns, each of which records its inputs, seeds, and results in a single machine-readable artifact as part of the campaign's audit discipline. Four conventions are enforced uniformly across the campaigns.

- *Pinned environment.* Every campaign ran in one pinned Python environment, fixed by a locked dependency manifest — Python 3.14.3, NumPy 2.5.0, SciPy 1.18.0, macOS on arm64 — and each artifact embeds the environment block under which it was produced.
- *Named seeds.* Every stochastic computation uses an explicitly seeded generator (NumPy PCG64), with the seed recorded in the artifact next to the result it produced; deterministic computations record that no random numbers were used. Where reruns were performed, artifacts reproduce byte-for-byte.
- *Code digests.* Artifacts record SHA-256 digests of the scripts that produced them, binding each result to the exact code version.
- *Guarded execution.* Long-running solves ran under wall-clock guards with the guard limits recorded; runtimes are logged as provenance only — no result depends on them. Several producers refuse to write output at all if any of their internal self-checks fails.

Table 11 maps each headline number to its producing campaign and the independent check applied to it. The verification is layered: every load-bearing quantity was tested by at least one method that shares no code with its producer.

Two further verification layers sit above the per-campaign gates. First, every headline quantity was re-derived in a separate verification pass, using independently written implementations that work from the cited astrometry and the stated models rather than from the originating code; agreement was gated at 0.5% for the kinematic quantities — the worst observed discrepancy falls below that gate — with case-by-case tolerances elsewhere as reported in the sections that follow. Second, the assembled results are pinned by an automated regression suite of more than 220 checks, which recomputes the headline numbers and the intermediate quantities feeding them and fails on any value drift; the suite ran green against every result quoted in this report.

E.2 Perihelion-Pumping Schedule Search

The four-parameter pumping policy of Appendix B was optimized twice, by independently seeded searches. Both use the same trajectory integrator: a fixed-step kick-drift-kick scheme with the thrust policy frozen across each step, an adaptive step law $\Delta t = \max(600\ \mathrm{s},\ 0.002\,P_{\mathrm{circ}}(r)) \times s$, where $P_{\mathrm{circ}}(r)$ is the local circular-orbit period and s the refinement scale, and a rocket-equation mass update on every thrusting step.

The original search drew 400 uniform-random candidates over the parameter box from a generator seeded at 11, then polished the five best by Nelder–Mead simplex descent; the best polished schedule reaches the cruise speed at 23.966 $\mathrm{km\,s^{-1}}$ in 11.90 years. Timestep refinement exposed the seam disclosed in Appendix B: at half step the run initially terminated at the simulation-horizon cap short of cruise speed, and completing it at an extended horizon gives 23.969 $\mathrm{km\,s^{-1}}$, with the quarter-step run at 23.985 $\mathrm{km\,s^{-1}}$.

The re-optimization repeated the search from an independent seed (20260716) with the perihelion floor raised from the search-phase 0.25 AU to the 0.42 AU class floor: 802 coarse evaluations over the box $r_p \in [0.42, 0.6]$ AU, aphelion fraction $\in [0.55, 0.98]$, burn radius $\in [0.45, 1.0]$ AU, termination energy

Table 11. Headline results, producing campaigns, and independent checks. Each campaign recorded its seeds and acceptance record in its machine-readable artifact.

Result	Producing campaign	Independent check
Departure optimum (73,012 yr; $v_\infty = 23.64\ \mathrm{km\,s^{-1}}$; 2.48° tilt; $13.85\ \mathrm{km\,s^{-1}}$ floor)	Deterministic departure-geometry scan	Closed-form identities (Appendix A); from-scratch replication, $< 0.5\%$
12-yr pumping schedule, $24.0\ \mathrm{km\,s^{-1}}$	Seeded schedule search and re-optimization (Section E.2)	Timestep ladder; rocket-equation fingerprint; independent re-optimization
Free-flight-time schedule, $18.87\ \mathrm{km\,s^{-1}}$	Six-cell long-horizon sweep (Section E.3)	Four-resolution timestep ladder
Certified lower bound, $16.56\ \mathrm{km\,s^{-1}}$	Six-lemma bounding chain (Section E.5)	Adversarial proof audit; 21-check numerical harness (Section E.5)
Reference anchor, $22.66\ \mathrm{km\,s^{-1}}$	Indirect optimal-control solve (Section E.4)	Direct-collocation cross-check attempted; unconverged (single-method, so labeled)
Out-of-plane increment, 578 m/s	Three-dimensional trajectory sweep (Section E.3)	Bit-level planar-embedding regression; independent re-integration
Intercept probability, 0.982	Seeded dispersion Monte Carlo (Section E.6)	Closed-form analytic propagation, 0.01% agreement
Mass-closure boundaries and grid	Sixteen-point closure scan (Section E.7)	Algebraic replication (Appendix C)
Catalog sweep and target ranking	Catalog encounter sweep (Section E.7)	Published-encounter solutions (Section A.4)

$\in [-220, -20]\ \mathrm{km^2\,s^{-2}}$, followed by Nelder–Mead polish of the eight best candidates (400 function evaluations each). Each surviving candidate was then integrated on a three-rung timestep ladder ($s = 1, 0.5, 0.25$) and accepted only if it passed five gates: (i) the integrated Δv matches the rocket-equation value implied by the final mass to within $10^{-13}\ \mathrm{km\,s^{-1}}$ (a mass-bookkeeping fingerprint; worst case 6×10^{-14}); (ii) the ladder spread is within 1% (measured 0.077%); (iii) the minimum solar distance respects the 0.42 AU floor on every rung; (iv) the commanded thrust never exceeds the power envelope, checked pointwise after the fact; and (v) the seeds and integrator metadata are recorded in full. The winner is the schedule from the original search — the re-optimization found nothing better from a fresh seed — and the carried value is the most conservative rung of its ladder, 23.9845 $\mathrm{km\,s^{-1}}$, quoted as 24.0. The full search took approximately 4 seconds of wall time; the schedule search is cheap, and the 12-year value is limited by the policy class, not by search effort.

E.3 Long-Horizon and Three-Dimensional Trajectory Sweeps

Long-horizon probe. The unconstrained-flight-time upper bound of Section B.4 comes from a six-cell sweep over horizons of 60 and 150 years crossed with three values of an aphelion-switch radius (5, 10, 20 AU) governing a fifth policy parameter that widens the prograde-burn window once the orbit grows large. Each cell drew 400 seeded coarse candidates (seed 20260715) and polished the best by Nelder–Mead. The sweep ran under a pre-registered stop rule: halt when a schedule passes all acceptance gates — cruise speed reached, perihelion floor respected, timestep ladder within 1% with an additional eighth-step corroboration rung, envelope compliance, and the rocket-equation fingerprint — at a Δv below the 22.26 $\mathrm{km\,s^{-1}}$ closure crossover. The rule fired at 18.871 $\mathrm{km\,s^{-1}}$ (ladder maximum) in the 60-year cell with a 10 AU switch radius, on a trajectory using about 58 years of flight; gate-passing schedules across the cells span 33–58-year flight times at 18.9–20.5 $\mathrm{km\,s^{-1}}$. One cell (20 AU switch, 150-year horizon) produced a nominally cheaper schedule (18.68 $\mathrm{km\,s^{-1}}$) that failed the extended ladder gate — timestep-fragile at its horizon cap — and was excluded rather than quoted; the exclusion is recorded in the artifact.

Three-dimensional sweep. The out-of-plane verification of the $2.48°$ departure tilt was performed with a three-dimensional extension of the same propagator, in which the steering direction acquires an out-of-plane component commanded by up to four additional parameters. Five sweep cells were run, each with its own seed (20260715 through 20260719), 600 coarse candidates, and eight polished finalists: tilt applied within the terminal escape burn only; tilt applied during the retrograde pumping arcs beyond an activation radius; both combined; a variant with the pump parameters frozen at the planar optimum and only the steering free; and a control cell targeting zero tilt, which recovered the planar 23.966 $\mathrm{km\,s^{-1}}$ result exactly. The direction requirement is a hard gate of $0.04°$ on the achieved asymptote latitude against the $-2.48°$ target — one tenth of the $0.41°$ open-loop aiming budget, so the trajectory gate does not consume the dispersion allowance of Section 6; during the coarse search the gate is relaxed to a penalty of 0.413 $\mathrm{km\,s^{-1}}$ per degree of direction error to keep the search from starving.

Before the sweep, the 3-D propagator itself was certified by a four-part self-test battery, recorded alongside the sweep results: (i) a planar-embedding regression — with the steering zeroed, the 3-D code must reproduce the planar ladder (23.9661, 23.9835, 23.9845 $\mathrm{km\,s^{-1}}$) bit-identically, with the out-of-plane coordinate exactly zero at every step (bitwise, not merely small); (ii) a cross-method asymptote check, comparing the departure direction constructed analytically from the osculating orbit at thrust cutoff against an independent high-order integration (SciPy DOP853, relative tolerance 10^{-10}, absolute 10^{-12}-scaled) coasting to 100 AU — agreement to $7{\times}10^{-12}$ radians; (iii) validation of the closed-form out-of-plane gain factor against finite-difference velocity-kick experiments in three orbital regimes; and (iv) measurement of the radial distribution of retrograde thrust in the planar optimum, which showed all retrograde Δv inside 2.3 AU and none beyond 4 AU, fixing the activation-radius search box.

The winning cell is the combined parameterization: carried 3-D value 24.5627 $\mathrm{km\,s^{-1}}$ (ladder maximum over the same three-rung refinement and gate set as the planar search), against the planar carried value of 23.9845 $\mathrm{km\,s^{-1}}$ — an increment of 578.19 m/s, inside the $[22\ \mathrm{m/s},\ 1.02\ \mathrm{km\,s^{-1}}]$ band predicted by the closed-form bracket of Section 4. Because both endpoints are best-found constructive schedules, the increment is a difference of upper bounds, not a bound on the true optimal increment; it is quoted with that caveat. The retrograde-arc-only cell produced no candidate passing the full gate set, a negative recorded in the artifact rather than discarded. The winner was then re-integrated by a method sharing nothing with the fixed-step propagator — SciPy DOP853 at relative tolerance 10^{-10} with the control law evaluated continuously at every derivative call rather than frozen per step — reproducing the Δv to within 0.3% of every ladder rung (0.02% of the finest) and the asymptote latitude to $0.013°$.

Finally, the increment was propagated through the mass-closure grid by the margin slopes of the Appendix C scan: the recommended 100 kg GTO architecture keeps margins of $+5.1$ and $+3.1$ points at the two ends of the heliocentric reference interval, while the 150 kg LEO architecture moves to $+1.0$ and -0.8 points — the conservative end flips sign, since that case tolerates only about 104 m/s of growth. This is the basis for the statement in Section 5 that the three-dimensional cost leaves the GTO recommendation intact while removing the conservative-end margin of the LEO alternative. The full 3-D campaign, all five cells and ladders, took 103 seconds of wall time.

E.4 Optimal-Control Solves and the Direct-Collocation Record

Indirect solves. The reference-anchor extremals of Section B.4 were produced by an indirect method: the costate equations of Pontryagin's minimum principle are integrated with SciPy's DOP853 at relative tolerance 10^{-8} and absolute tolerance 10^{-10}, with the bang-bang thrust switching smoothed by a continuation schedule that sharpens the switching function in three stages ($\varepsilon = 0.3, 0.1, 0.03$), and the boundary conditions closed by shooting from five random multistarts (seed 20260702). Convergence requires the shooting residual below 10^{-6}; the accepted solutions close it to 4.4×10^{-13}. The 23.64 $\mathrm{km\,s^{-1}}$-target solve converges in 0.4 s to 22.660 $\mathrm{km\,s^{-1}}$ (minimum solar distance 0.62 AU, so the perihelion floor is inactive), and the same solver reproduces the previously computed 22.944 $\mathrm{km\,s^{-1}}$ extremal at the 24.14 $\mathrm{km\,s^{-1}}$ target variant bit-for-bit — an environment-drift check, not an independent verification.

Direct-collocation attempts. Four attempts were made to confirm the anchor extremal by direct transcription, and all four failed; the record is retained because the failures are informative. The first three — uniform-mesh Hermite–Simpson at 24, 48, and 60 segments — stalled at their iteration limits near the initializing heuristic or failed re-integration, as summarized in Appendix B. The fourth was the strongest attempt the problem structure permits: a 100-segment Hermite–Simpson transcription on a nonuniform mesh with nodes concentrated where the dynamics are fast (node density proportional to a blend of $r^{-3/2}$-warped and uniform time, giving a five-to-one ratio of largest to smallest segment), warm-started from the indirect extremal itself, with the steering and throttle histories seeded from the costate solution. The nonlinear program converged by every internal measure — successful termination in 148 iterations, collocation defects of 4.4×10^{-10} — claiming 21.62 $\mathrm{km\,s^{-1}}$, below the indirect value. Re-integrating the claimed control history through the dynamics, however, returns 21.18 $\mathrm{km\,s^{-1}}$ of actual expenditure, a 2.1% discrepancy exceeding the 2% acceptance gate: the transcription had exploited discretization slack rather than found a better trajectory. The attempt is recorded as a rejected solution, and the anchor retains its single-method label.

Abandoned design-thrust solves. Optimal control at the design thrust acceleration itself ($a_0 = 2.5 \times 10^{-4}\,\mathrm{m\,s^{-2}}$, where a solution would sharpen the closure verdict directly) was attempted and formally closed out as a negative result under a documented disposition procedure. Two solver jobs — a segment-grid direct solve and a free-horizon indirect solve — ran to approximately 20.4 and 20.1 CPU-hours respectively (essentially full CPU over roughly 26 hours of wall time each) without completing a single case. Rather than silently discarding them, the close-out procedure froze the environment, verified the process identities, captured stack samples (both processes were live inside the integrator's derivative evaluations, not hung), hashed every log before and after interruption, and adjudicated all recovered output against the certified interval: no converged admissible trajectory existed in any log. One recovered line — a 120-segment direct attempt reporting 12.9 $\mathrm{km\,s^{-1}}$ with a failed convergence flag — is explicitly recorded as inadmissible, since it lies below the certified 16.56 $\mathrm{km\,s^{-1}}$ bound and therefore cannot correspond to any real trajectory. The roughly 40 CPU-hours of abandoned solves are the strongest available evidence that the many-revolution problem at design thrust is out of reach for these methods, which is why the heliocentric budget rests on the certified bound and the constructive schedules instead.

E.5 Lower-Bound Certificate and Numerical Falsification Harness

The certified 16.56 $\mathrm{km\,s^{-1}}$ bound of Section B.4 is a chain of six closed-form lemmas, machine-checked end to end. In summary: the first lemma truncates the trajectory at first attainment of the target energy and converts the propellant budget into an explicit ceiling on thrust acceleration as a function of assumed total Δv; the second proves the escape leg makes a single transit of the inner region — every radial-velocity zero is a strict upward crossing, and hovering near the perihelion floor would require some thirteen times the available thrust; the third bounds the duration of that transit by a comparison argument on the radial dynamics, evaluated at a shell radius of 0.46 AU (slightly above the 0.42 AU floor, where the bound is sharpest), capping both the Δv deliverable and the energy gained inside the window; the fourth assembles the escape-leg cost as an integral of the reciprocal shell speed over the required energy gain; the fifth is a self-contained proof that the initial segment, up to first escape energy, costs at least $(\sqrt{2}-1)$ times Earth's orbital speed, via a verification function whose gradient inequality reduces to a two-variable polynomial bound; and the sixth closes the chain by self-consistency, evaluating the assembled bound at an assumed total of 16.56 $\mathrm{km\,s^{-1}}$ and obtaining 16.5659 $\mathrm{km\,s^{-1}}$ — a contradiction for any cheaper trajectory. The closure is insensitive to the assumed total: the assembled bound varies only between 16.566 and 16.564 $\mathrm{km\,s^{-1}}$ as the assumption ranges from 16.503 (the impulsive floor) to 17 $\mathrm{km\,s^{-1}}$. Two simpler relaxation strategies were tried first and are recorded as negative results: a scalar energy relaxation certifies only 11.39 $\mathrm{km\,s^{-1}}$ and a two-variable box relaxation at most 13.31 $\mathrm{km\,s^{-1}}$, both below the impulsive floor and therefore useless; their pointwise inequalities were salvaged as components of the main chain.

Because the bound is a theorem-grade claim, its audit was adversarial: a dedicated review pass, conducted separately from the derivation, attempted to defeat each lemma in turn — probing its case splits, boundary values, and sign conventions for counterexamples — and the chain was accepted only after every such attempt failed. The chain was then attacked numerically by a dedicated falsification harness: seven inequality targets drawn from the lemmas — turning-point sign structure, the radial-velocity lower bound, the transit-window duration cap, the in-window energy cap, the pointwise shell-speed bounds, and the polynomial gradient bound with its boundary values — were each evaluated on three trajectories of different character: the optimized pumping schedule, the anchor extremal, and a freshly integrated ballistic orbit grazing the shell (DOP853 at relative tolerance 10^{-12}, energy conserved to 6×10^{-15} over the arc). Of the 21 checks, none was violated; checks whose case-split premises do not apply to a given trajectory are recorded as vacuous rather than passed.

The tightest margin, -9×10^{-15} on the radial-velocity bound along the ballistic trajectory, is machine-epsilon-scale — the inequality is achieved with equality there, exactly as the comparison argument predicts, which is itself a nontrivial confirmation that the bound is sharp rather than slack. The polynomial gradient bound was additionally checked in a cancellation-free factored form on a 400×400 grid and against a 1,000-point random polynomial-identity cloud (seed 20260715), and the eight assembled constants of the fixed-point evaluation were recomputed from first principles in the harness, matching the derivation to 5.4×10^{-6} relative.

E.6 Aiming-Dispersion Monte Carlo

The intercept-probability analysis of Section 6 is a 20,000-draw Monte Carlo — a dedicated closed-form dispersion-propagation module driven by a thin campaign script — in which each draw perturbs the four error sources — target ephemeris (catalog parallax, proper-motion, and radial-velocity covariances), pointing (two transverse rotation angles), velocity magnitude, and the injection residual — and propagates the miss at closest approach in closed form, with no numerical integration per draw. Every run in the campaign carries its own named seed, fifteen in all (42–55): the headline run (seed 42), the alternate-catalog variant (43), four one-at-a-time source isolations (44–47), four leave-one-out complements (48–51), three assumption-sensitivity doublings (52–54), and the cross-catalog systematic diagnostic (55). The producer is deterministic — reruns are byte-identical — and refuses to write its artifact if any of its 29 internal self-checks fails.

The headline result is $P = 0.98155$ (19,631 of 20,000 draws inside the 2,600 AU sphere), with a Wilson 95% interval of $[0.97959,\ 0.98332]$ and an exact Clopper–Pearson interval agreeing to 1.7×10^{-4}. Convergence was established on prefix subsets of the single seed-42 stream: the 5,000-, 10,000-, and 20,000-draw intervals are mutually contained, and the interval width contracts by a factor of 0.502 from 5,000 to 20,000 draws against the 0.5 expected from $1/\sqrt{n}$ scaling. Ranked by one-at-a-time root-mean-square miss, the sources contribute 764 AU (velocity magnitude), 647 AU (pointing), 617 AU (ephemeris), and 547 AU (injection residual); a leave-one-out total-effect ranking promotes ephemeris above pointing once interactions with the geometry are included, and the variance budget closes to 0.2% — the interaction residual between the summed one-at-a-time second moments and the full-run second moment — confirming the sources act nearly independently. The headline probability was verified by a genuinely independent method: a first-principles linear-Gaussian propagation of the transverse covariance through the closed-form gains, with the disc probability evaluated by 400-node Gauss–Legendre quadrature, sharing no code or random numbers with the Monte Carlo. It gives $P = 0.98166$, a relative difference of 1.1×10^{-4} against an acceptance gate of 2×10^{-2}. Substituting the alternate astrometric catalog self-consistently (centers, uncertainties, and aim point together; seed 43) gives $P = 0.9919$, and the deliberate cross-catalog experiment — aiming with one catalog while drawing truth from the other — produces the ≈ 859 AU cross-track systematic offset discussed in Section 6.5 ($P = 0.9623$; a first-order analytic prediction of the offset agrees to 1.1%). The assumption-sensitivity rows are part of the record: doubling the pointing dispersion lowers P to 0.897, doubling the magnitude dispersion to 0.858, and quadrupling it to 0.571, which is why the velocity-trim requirement is flagged as the binding engineering assumption. The artifact embeds the SHA-256 digests of the producing code and the pinned-environment block.

E.7 Catalog Sweep and Cost Monte Carlo

Catalog sweep. The alternative-target scan of Section 8 rests on a full sweep of the Bailer-Jones stellar-encounter catalog [6]: all 61 rows of the published close-encounter table (retrieved from the VizieR service on 2026-07-14, with Gaia DR3 astrometry for the survivors fetched the same day) were screened for future encounters with perihelion inside 1 pc arriving within the 500,000-year horizon and free of the catalog's dubious-astrometry and binarity flags. The filter removes 32 past encounters, 20 whose perihelia lie beyond the horizon, and four flagged rows (each exclusion recorded with its reason — all four flagged stars carry spurious radial velocities of hundreds of $\mathrm{km\,s^{-1}}$); the eight survivors were then ranked by the minimum impulsive departure from low Earth orbit, alongside reference solutions for Alpha Centauri, Proxima, and λ Serpentis computed identically. The metric has a floor of 8.757 $\mathrm{km\,s^{-1}}$ — bare solar escape from a 400 km orbit with full Earth-velocity alignment — which no target can undercut. Two anonymous distant M dwarfs formally beat LSPM J2146+3813 (8.78 and 8.82 against 9.30 $\mathrm{km\,s^{-1}}$) but only marginally above the floor and only with weak radial velocities (±8–11 $\mathrm{km\,s^{-1}}$) and near-degenerate miss geometry, which is why Section 8 carries LSPM J2146+3813 as the practical minimum-energy target and λ Serpentis (11.38 $\mathrm{km\,s^{-1}}$) as the best Sun-like alternative; the sweep confirms every unflagged surviving candidate is an M dwarf.

Table 12. Triangular (low, nominal, high) parameters of the cost Monte Carlo, in 2026 dollars per driver unit. The escape-injection premium is scaled by departure mode (0 from LEO, 0.45 from GTO, 1.0 for direct escape injection); the integration-and-qualification line carries the 1.4 perihelion thermal-qualification multiplier for pumping trajectories (Appendix D). Operations are priced per year of powered flight with a one-year floor.

Line item	Unit	Low	Nominal	High
Spacecraft bus	fixed	\$1.2M	\$2.0M	\$4.0M
Thruster and power-processing string	per kW	\$0.35M	\$0.60M	\$1.2M
Solar arrays	per W	\$250	\$400	\$800
Xenon propellant, loaded	per kg	\$2,500	\$4,000	\$12,000
Tankage and feed system	per kg Xe	\$1,000	\$2,000	\$4,000
Rideshare launch	per kg	\$5,000	\$6,500	\$9,000
Escape-injection premium	per kg	\$8,000	\$13,500	\$25,000
Integration, test, and qualification	fixed	\$1.0M	\$1.75M	\$3.0M
Powered-flight operations	per year	\$0.6M	\$1.2M	\$2.5M
Program margin and management	fixed	\$1.0M	\$2.5M	\$5.0M

Cost Monte Carlo. The cost distributions of Section 7 are sums of ten independent triangular draws, one per line item, sampled 20,000 times from a generator seeded at 7 in a single vectorized call of the cost model. Table 12 gives the per-unit (low, nominal, high) parameters behind the nominal values of Table 10; each row carries a sourcing note in the model itself. Percentiles at this sample count are stable to well under the \$0.1M quotation grain. The costed architectures come from the sixteen-point closure scan: both departure modes, four wet masses, and both ends of the heliocentric reference interval — the same scan that populates the closure grid of Appendix C — with the full cost model attached to each closing case of interest. The variance decomposition, exact under independence, attributes about 64% of the cost variance to the powered-flight operations line, which is what makes the structural conclusion of Section 7 insensitive to the hardware bands.