Magnetic Matter Movers Program Gate Plan — MTS

Magnetic Transportation Systems (internal program designation) · Phase 0 → Mars: entry conditions, validation targets, kill criteria, funding sources

Same discipline as the Phase 0 checklist, applied to the whole launch architecture. Early gates are contractual (numbers in ink); later gates are directional and get re-baselined as earlier phases return data. No step skips. (Not to be confused with the Matter Stack display project, which has its own checklist.)

How to Read This

Each phase lists: **Entry** (what must already be true) · **Do** (the work) · **Validate** (predicted numbers the phase must confirm — the actual-vs-predicted discipline) · **Gate evidence** (what passing looks like) · **Kill / pivot** (the result that honestly ends or redirects the phase — deciding these *before* the data arrives is what keeps the program honest) · **Cost & source**.

Phase 0 — Benchtop (ACTIVE)

Entry: none. **Do:** Build Sheets 001 + 002 per the Phase 0 Checklist. **Validate:** Inductrack flat-drag signature; droplet thrust-vs-power; trap-calibrated μN metrology. **Gate evidence:** the filled Phase 0 Checklist package. **Kill/pivot:** none — Phase 0 cannot fail, only teach. **Cost:** ~$270, self-funded.

Phase 1 — Paper & Photons (~$0–10k, evenings)

Entry: Phase 0 gate passed. **Do:** (1) Split-step wave-optics propagation sim — turbulence screens + thermal blooming feedback, slant vs. relay-vertical; (2) Lab Note 001 from Phase 0 data; (3) HX coupon test spec; (4) SIS Draft A + demand-registry instrument; (5) high-power droplet campaign (3B safety gate first) + beam-riding offset map; (6) torsion-pendulum thrust stand. **Validate:** sim reproduces scaling-law results within ~2× (N_D 28 relay-vertical / 270 slant; delivered-flux fractions); offset map shows self-centering regime exists at some power. **Gate evidence:** propagation white paper draft; SIS Draft A circulated (CONFERS/SmallSat submission); thrust stand calibrated against trap results. **Kill/pivot:** wave-optics shows blooming uncorrectable even relay-vertical with 300 m array → architecture pivots (larger site dispersal, higher-altitude site, partial-microwave hybrid) BEFORE any hardware money. **Cost & source:** self + first Etsy catalog revenue.

Phase 2 — Flight Flux on a Bench (~$10k–500k)

Entry: Phase 1 gate; propagation paper survives external review. **Do:** 1–10 kW fiber laser on HX coupons — 10 kW on 2 cm² = full 48 MW/m² flight flux; hydrogen-side and steam-side coupon campaigns (C-C/TaC and oxide/Ir coating sets); cycle-life testing (300↔3,000 K); steam microthruster on the torsion stand (measured Isp); subscale coilgun bench (multi-stage capacitor coilgun + null-flux track v2); **cadence endurance rig** — automated subscale coilgun fired at shots-per-minute pace to **5,000 cycles** with per-shot condition monitoring and a cost-per-shot ledger (cadence risk is cycle-count-invariant: switch life, coil thermal cycling, and flywheel fatigue retire at benchtop scale). **Validate:** coupon survives flight flux × 462 s × ≥100 cycles (extended protocol: ≥1,000 automated cycles); measured microthruster Isp within 20% of temperature prediction; coilgun stage efficiency ≥ published EM-launcher benchmarks; endurance rig completes 5,000 cycles with replacements logged and cost-per-shot flat. **Gate evidence:** coupon life dataset; Isp measurement; two more lab notes. **Kill/pivot:** coating life < ~10 cycles at flight flux after iteration → de-rate to Isp 850 baseline (payload 21%→19%, architecture survives) or accelerate laser-sustained-plasma track; if BOTH fail → beamed stage dies, gun+tug-only architecture (freight-first, no delicate-cargo line). **Cost & source:** SBIR/STTR (AFWERX/SpaceWERX), angel; this is the first external-money phase — Phase 1's papers

are the application.

Phase 3 — 1 MW, First Flights (~$10–50M)

Entry: Phase 2 gate; SIS Draft B out (feedback incorporated). **Do:** 1 MW array (≈50 modules) + AO; 10 kg vehicles, suborbital hops; full-panel HX (tiled coupons); beam-riding closed-loop control in flight; tracking-loss budget measurement. **Validate:** end-to-end efficiency vs. the 0.43%-collapse benchmark — target ≥100× better via AO + cooperative beacon (i.e., ≥40% moving-target transport); panel hot-spot survival with commanded flux shaping (work item 11); vehicle Δv within 15% of prediction. **Gate evidence:** repeated flights of one reusable airframe at sustained cadence (≥2 flights/week for a month); efficiency ledger; investor-grade data room. **Kill/pivot:** moving-target efficiency stuck near fixed-hardware limits → relay-mirror dependence increases (pull relay coating work forward); airframe non-reusable after N flights → cost model re-run, proceed only if $/kg trajectory holds. **Cost & source:** venture + strategic (this is where SpinLaunch-class raises happened).

Phase 4 — 10 MW Pilot, First Revenue (~$50–200M)

Entry: Phase 3 gate. **Do:** 10 MW array; 100 kg-class vehicles to 100+ km, near-orbital; first paying microsat customers designed to SIS; relay-mirror flight unit (coating validated in orbit); pulsed-power flywheel subscale for the gun line; **cadence campaign** — sustained shots/day on dummy/water payloads at marginal (energy-only) cost: low bookings never idle the asset; the cycle ledger itself is the sales instrument for the next booking. **Validate:** $/kg on the published price sheet met within 2×; SIS-designed third-party payload flies successfully; relay coating absorptivity in orbit ≤ spec. **Gate evidence:** revenue; manifested backlog; demand-registry conversions. **Kill/pivot:** no SIS payload demand materializes at pilot prices → freight-gun-first re-sequencing (bulk water market to depots doesn't need third-party payload designers). **Cost & source:** venture + first government anchor-tenant contracts (CLPS-pattern).

Phase 5 — Full Cargo Stack ($4–12B)

Entry: Phase 4 gate + anchor tenant signed (governments are the only rational ahead-of-demand buyer — this gate is political, not technical, and should be named as such). **Do:** 475 MW array + 6.8 km maglev track (Beam Line); 4.9 km coilgun freight line (Gun Line) in parallel; 2–3 relay mirrors; LEO water dock + 8 steam tugs (Depot). **Validate:** 21% payload fraction; $11/kg marginal energy; ramp to 100+ flights/yr (Falcon-parity gate) then 500+ ($870/kg). **Kill/pivot:** flight-rate demand stalls below ~50/yr for 3+ years → system pivots to gun-dominant bulk operations (water/propellant), beamed line held at pilot scale. **Cost & source:** project finance + anchor tenancy; laser $/W cost-curve timing is the single largest cost lever — re-bid annually.

Phase 6 — Crew (Crew Beam, $23–84B)

Entry: Phase 5 operating history (flight heritage is the human-rating currency) + demonstrated seat demand ≥60/yr at chemical prices. **Do:** 38 km 3-g track; 3.8 GW array; 6-seat vehicles; solid abort motor (the one chemical component, by design). **Kill/pivot:** seat demand never materializes → chemical crew taxis remain the permanent answer (1.1% of expedition mass — an acceptable end state, per the study).

Phase 7 — Mars Pipeline

Entry: Depot operating; lunar route debugged (Moon first: daily windows, 3-day feedback). **Do:** aerocapture slug variant (M-class heatshield); Mars-orbit port emplacement (one-time tug campaign); window-cadence operations. Phase 7b: Phobos momentum bank, only after route proven. **Validate:** aerocapture corridor vs. gun dispersion (work item 13 — this number is measured at the Moon and Earth first).

The Standing Rules

1. **No step skips.** Each phase's credibility is the next phase's collateral. *(Scope: binds spending and public claims — never thinking, prototyping, or exploring ahead of phase.)*
2. **Kill criteria are written before data arrives** — and rewritten only at gate reviews, never mid-phase.
3. **Every phase ships something:** a product, a paper, a dataset, or revenue. A phase that only consumes is off-pattern and gets re-scoped.
4. **Actual-vs-predicted tables at every scale.** The droplet got one; the 475 MW array gets one. Same columns, more zeros.
5. **Re-baseline forward, never backward:** later-phase numbers update as early phases return data; early-phase gate records are immutable.

Amendments 6–9 adopted at rule review, Jul 19, 2026 — closing two gaps (cadence, operator safety), adding one cap, loosening one over-tightening:

6. **The heartbeat.** Minimum cadence: one commit or one bench hour per week, and a 30-minute monthly self-review against the open work items. Deliberately tiny — small enough to survive any bad week, regular enough that the program can never silently flatline. Quality rules control *what* gets done; this is the only rule that controls *whether*. Missing two consecutive monthly reviews is itself a gate-review trigger.
7. **Operator interlocks — the zero-judgment category.** The machines get watchdogs and hardware cutoffs; so does the operator. Non-negotiable, no "just this once": (a) no laser energized with the enclosure open, ever, including alignment; (b) no high-power or stored-energy test while exhausted or after a full work shift; (c) once tests carry real energy (Class 3B, spinning mass at speed, charged coil banks), someone knows a test is running — a message before and after. Safety text in build sheets is information; this rule is enforcement.
8. **The burn cap.** A monthly all-projects spending ceiling set against the bridge-income budget, revised only at monthly review; no debt for hardware, ever. Gates control what may be bought; the cap controls how much, in total, regardless of how many gates are open.
9. **The sandbox exemption.** Untracked, undocumented play at the bench is encouraged, not merely allowed. Checklists govern *gate evidence*, not curiosity — spin the drum, chase the flicker, breadboard off-roadmap, no ink required. The rules bind the irreversibles (money, claims, eyes and fingers) and the cadence; they never bind wonder, which is what the program runs on.