

Antimatter and Spacetime Propulsion

An Integrated Research Architecture

Robert J. Hathaway III

First edition 10 October 2026 | Revised 18 October 2026

“The geometrical interpretation of the Eulerian energy thus opens new doors to generating realistic warp fields for laboratory study and potential future manned interstellar travel.”

Shaun D. B. Fell and Lavinia Heisenberg, 2021, Abstract [3]

Abstract

A useful prototype can be built as a coordinated research platform for antimatter handling, precision energy delivery and the physical sourcing of candidate spacetime geometries. Its first purpose is to turn an ambitious propulsion concept into measurable subsystem requirements. This paper connects the warp research of Erik Lentz, Alexey Bobrick and Gianni Martire, Shaun Fell and Lavinia Heisenberg, and subsequent investigators with NASA antimatter propulsion concepts, recent CERN results, large AI computing facilities, autonomous laboratories and quantum computing. It incorporates the design prompts of frozen antihydrogen pellets, an optional antihelium cushion and contact-free fuel regulation.

The proposed architecture has two development paths. One advances production, cooling, confinement, metering and energy recovery using known physics. The other establishes whether a controllable spacetime configuration can actually be sourced, formed, accelerated and extinguished. They share computing, materials, diagnostics and systems engineering. Antimatter power is connected to a future field actuator only after the second path establishes a physical mechanism. No published result reviewed here demonstrates an affordable, controllable faster-than-light vehicle.

The planning horizons are 2031, 2036, 2046 and beyond 2046. Engineering targets span laboratory pulses through eventual power and propulsion experiments. An aggressive scenario allows recursive AI improvement, selected quantum advantages and rapidly multiplying automated laboratories to compress research cycles. Faster-than-light dates remain conditional on a verified physics breakthrough. Quantitative estimates distinguish stored field energy, reactor power, fuel inventory, production electricity and heat rejection; confusing these quantities can produce errors of many orders of magnitude.

SpaceX-NASA cryogenic-transfer results provide an engineering precedent for storage infrastructure; the formation and confinement of solid antihydrogen remain separate research tasks. [19,36]

Keywords: antimatter propulsion; spacetime geometry; antihydrogen storage; cryogenics; recursive AI; interstellar transport

Principal architecture decision

Build a modular platform whose antimatter, computation and measurement systems have useful outcomes at every stage. Treat the spacetime actuator as a rigorously specified research interface. Preserve the possibility of dramatic discoveries without assigning an unsupported performance specification to missing physics.

Contents

Select a section title or page number to jump to that section.

1 Introduction and scope.....	3
2 The spacetime physics.....	4
3 Current theoretical requirements.....	5
4 Guiding system architecture.....	6
5 Fuel and power components.....	7
6 Field and support components.....	8
7 Antimatter production and preparation.....	9
8 Frozen fuel and the antihelium cushion.....	10
8.1 Starship cryogenics and antihydrogen storage.....	11
8.2 Cryogenic research and acceptance milestones.....	12
9 Metering antimatter flow.....	13
10 Power and inventory estimates.....	14
11 Production electricity and heat rejection.....	15
12 A physical field actuator.....	16
13 Verification and measurement.....	17
14 AI factories and recursive improvement.....	18
15 Quantum computation and automated laboratories.....	19
16 Accelerated research scenarios.....	20
17 Initial engineering envelope.....	21
18 Roadmap of linked innovations.....	22
19 Five year and ten year objectives.....	23
20 Twenty year and later objectives.....	24
21 Prototype stages and acceptance tests.....	25
22 Resources and research priorities.....	26
22.1 Prospective SpaceX research collaboration.....	27
22.2 Staged execution and continuing employment.....	27
23 Conclusions and quantitative limits.....	28
References for spacetime physics.....	29
References for antimatter engineering.....	30
References for enabling technologies.....	31
Concept inputs.....	31
References for cryogenics and research transition.....	32
Revision history.....	32

1 Introduction and scope

The intended prototype is a guiding architecture, not a fabrication drawing or a claim of demonstrated warp propulsion. The first hardware is a ground-based set of linked experiments at an established accelerator laboratory. Later systems may demonstrate energy conversion, an uncrewed conventional spacecraft, or a genuine spacetime effect if the required physical source becomes available. A passenger vehicle is a separate development program.

Four evidence labels are used throughout. Demonstrated means a measured laboratory result. Proposed means a published model or engineering concept. Target means an original planning objective in this paper. Conditional means that an unresolved physical mechanism or major production advance is required. The future ranges are scenario bands, not statistical confidence intervals or vendor forecasts.

Evidence category	Examples	Architecture consequence
Demonstrated	Trapped anti-atoms; antiproton transport; AI plasma control	Build repeatable modules around the measurements.
Published proposal	Warp geometries; solid antihydrogen storage; antimatter propulsion	Reproduce the calculations and identify experimental assumptions.
Engineering target	A production pilot; a metered pulse source; automated test capacity	Budget and measure progress against explicit requirements.
Physics dependent	A compact controllable FTL field; an accessible source for its tensor	Advance only after independent theoretical and experimental checks.

How optimism enters the estimates

The accelerated scenario assumes much better research agents, more capable numerical tools, increasingly autonomous fabrication and access to expanding cloud and quantum facilities. These developments can improve search, discovery and implementation. They do not, by themselves, establish that a prohibited configuration is possible or that an exotic source is accessible.

The central question is which constraints are technological and which are physical. Production yield, cooling efficiency, field tolerances and experiment throughput can plausibly improve substantially. Causality, energy conditions, quantum backreaction and the relationship between a desired metric and realizable matter require explicit solutions. A successful research program can discover a new mechanism, identify a rigorous limitation, or produce valuable antimatter and spacecraft technologies while the FTL branch remains open.

2 The spacetime physics

General relativity relates spacetime curvature to the complete stress-energy tensor. The tensor includes energy density, pressures, stresses and momentum flow. Supplying a large amount of energy does not automatically produce the desired geometry. A proposed actuator must create the appropriate distribution and obey the coupled evolution of the matter and gravitational fields.

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}$$

ALPHA gravity measurements support antihydrogen falling toward Earth. Antimatter is therefore not an experimentally established source of negative gravitational mass; its annihilation supplies ordinary positive-energy reaction products. [18]

Lentz proposed hyper-fast soliton geometries with positive Eulerian energy density and discussed possible plasma and electromagnetic sources. Bobrick and Martire developed a general moving-shell description and positive-energy subluminal configurations. Fell and Heisenberg studied geometric decompositions that yield positive energy density for a particular observer family and suggested optimizing the geometry to lower its requirements. These are distinct mathematical approaches; combining their names does not establish a single complete engine model. [1-3]

Research line	Result relevant to the prototype	Remaining requirement
Lentz	Candidate superluminal soliton geometry	Correct tensor, energy conditions and a realizable source.
Bobrick and Martire	Moving-shell framework and subluminal positive-energy examples	Propulsion, formation and a practical shell.
Fell and Heisenberg	Geometric interpretation and opportunities for optimization	Full observer analysis, sources, dynamics and energy reduction.
Fuchs and colleagues	2024 constant-velocity subluminal solution satisfying energy conditions	Efficient matter shell and a physical acceleration process.

The 2024 subluminal solution is especially useful as a verification benchmark because it includes nontrivial spatial geometry and lapse rather than relying only on a shift vector. Its illustrated shell still has a mass of 4.49×10^{27} kg. Its importance is the consistency of the construction within its stated regime, not a demonstrated engineering mass budget. [6]

In a proposed FTL spacetime, a passenger trajectory can remain locally timelike while travel is faster than the corresponding distant light journey. The prototype must compare actual emission and reception events, clocks and null trajectories. A coordinate velocity above c , a moving optical pattern or an effective optical metric alone does not demonstrate that transport.

3 Current theoretical requirements

Positive energy for one observer family is weaker than the weak energy condition, which requires nonnegative measured energy for every timelike observer. Santiago, Schuster and Visser exposed this distinction in analyses of warp geometries. A subsequent preprint by Celmaster and Rubin directly challenges the Lentz construction and its derivation. The scope of each result matters: the criticisms constrain specific constructions and classes; they should not be expanded into an unsupported theorem covering every modified-gravity theory. [4,5]

Fell and Heisenberg calculated about 9.25×10^{43} J for their numerical example. They also proposed further geometric optimization. That example is an energy integral for the specified configuration, not a measured reactor power or a universal lower bound on future designs. It neither establishes kilowatt operation nor proves that every possible design must retain planetary mass-equivalent requirements. [3]

Research continues. A 2026 revision by Garattini and Zatrımaylov examines a bubble in a de Sitter background under a special relation between its motion and cosmic expansion. Its local and averaged energy statements and background assumptions require careful treatment before any application to a local vehicle. Santos-Pereira's 2025 thesis explores fluid and charged sources in Alcubierre-related models. Neither establishes a complete controllable compact FTL machine. [8,9]

Constraint	Required evidence before a transport claim
Energy conditions	Analyze the complete tensor in all relevant observer directions, with controlled numerical error.
Physical source	Give a matter or quantum-field model that produces the tensor and satisfies its own equations.
Causal formation	Evolve the device from accessible initial data; identify controls available after horizons form.
Stability	Bound perturbations, radiation and quantum backreaction over the intended operating period.
Usable transport	Show formation, travel, steering, braking and termination with a consistent energy and momentum budget.

Olum's superluminal-travel result, quantum-inequality restrictions, and semiclassical warp-instability studies belong in the automated test suite. Each must be applied with its actual assumptions. Removing one difficulty leaves the other requirements to be solved. [10-12]

4 Guiding system architecture

The architecture supports parallel progress without requiring the reactor and the spacetime model to mature together. Fuel preparation and field research exchange measured parameters through a common simulation and control platform. Conventional power can run early field and sensor experiments.




Figure 1 Prototype architecture. Solid links are engineering or analysis interfaces. The dashed link requires a verified spacetime source and actuator. Antimatter power alone does not close that link.

The power interface reports pulse energy, spectrum, efficiency, waste heat, available electrical power and timing. The field interface must report the source model, required stresses, geometry, confinement energy, loss rate, control bandwidth and observable curvature. Either interface may fail its feasibility test independently.

A common data platform preserves source provenance, coordinate conventions, solver versions, calibration histories and all failed trials. Independent implementations validate candidate results before they affect hardware. Rapid local controllers operate the instruments; remote AI factories supply design, simulation and supervisory planning.

5 Fuel and power components

The following requirements form a system decomposition. Existing technologies are starting points; their integration at greater antimatter inventory remains a research task. Component ranges elsewhere in the paper are targets rather than purchased specifications.

Module	Function and principal interface	Evidence or development need
Production source	Generate antiprotons and positrons; export yield, energy spectrum, electricity and target loads.	Dedicated source optimization rather than routine harvesting of LHC collision products.
Collection and deceleration	Accept a broad phase-space distribution and deliver a cold beam to traps.	Acceptance, energy spread and losses dominate useful yield.
Charged particle cooling	Reduce emittance and thermal energy; transfer controlled batches.	Electron cooling, plasma manipulation and modular traps.
Neutral synthesis	Combine antiprotons with positrons; characterize state and kinetic energy.	Antihydrogen is demonstrated; scalable molecule formation remains a milestone.
Fuel storage	Hold inventory without material contact; measure temperature, loss and position.	Atoms first; molecules and microcrystals only after verification.
Meter and transfer	Deliver a known amount to the reaction region; stop delivery on a fault.	Charged-particle extraction first; noncontact pellet delivery later.
Annihilation and conversion	Measure product spectra and convert a defined fraction to usable power or thrust.	Charged products, photons, secondary particles and neutrinos need separate budgets.

The antihelium buffer is an optional storage research branch. It is excluded from the minimum working prototype so that failure to produce a bulk antihelium fluid cannot stop atom trapping or reaction control. Power conversion is likewise replaceable: the field laboratory can use ordinary electrical power while antimatter work matures.

6 Field and support components

Module	Function and principal interface	Required development
Metric and source model	Convert a transport objective into a full stress-energy distribution and causal history.	Verified equations and realizable matter dynamics.
Field actuator	Produce the source distribution in a bounded test volume.	No validated FTL actuator exists; keep this interface conditional.
Curvature metrology	Compare clocks, light paths, matter motion and tidal response with predictions.	Independent instruments and complete electromagnetic, thermal and vibration budgets.
Magnets and structures	Provide confinement and transfer fields, alignment and mechanical support.	Field uniformity, stress, radiation tolerance and quench behavior.
Vacuum and cryogenics	Suppress contact losses and remove thermal loads at multiple temperatures.	Cryopumping, clean assembly and separated hot and cold systems.
Local control and diagnostics	Track inventory, timing, fields, radiation and faults.	Deterministic response; AI supervisory proposals bounded by validated operating limits.
Heat rejection and shielding	Handle deposited energy, radiation and secondary particles.	Mass and thermal design can outweigh the fuel mass.
Vehicle systems	Provide conventional motion, navigation, communications and remote operation.	Uncrewed flight before any crew integration; normal propulsion remains necessary.
Research computing	Run AI search, high-precision solvers, quantum subproblems and experiment planning.	Performance measured on verified outcomes, not advertised FLOPS alone.

Integration must include failure propagation. A magnet quench can change particle trajectories; a trap loss can create a thermal or radiation load; that load can disturb a neighboring trap. Segmenting inventory and isolating interfaces limits cascading faults. Small early inventories let the team study these effects without treating an untested bulk storage system as a prerequisite.

7 Antimatter production and preparation

A useful fuel factory is an end-to-end chain. A proton-target source produces antiprotons, collection optics capture an accepted fraction, deceleration removes kinetic energy, and cooling prepares particles for confinement. Positrons require their own source and accumulation system. The first objective is increased usable inventory per unit electricity and facility time, not a larger number of nominal collisions.

Recent CERN work demonstrates meaningful improvements. ALPHA's beryllium-assisted positron cooling increased anti-atom accumulation rates by approximately eightfold and enabled more than 15,000 simultaneously confined antihydrogen atoms. CERN reported that the sample could be accumulated in under seven hours. This advance improves precision experiments; the corresponding mass remains about 2.5×10^{-20} g. [14,15]

A March 2026 GBAR preprint reports 56 percent trapping of the ELENA beam, about 6.4 million antiprotons per shot and more than 64 million accumulated in under 35 minutes. The demonstrated inventory corresponds to about 1.1×10^{-16} g. It is a charged-particle record in that apparatus, not a neutral fuel pellet. [17]

The September 2026 BASE-STEP paper reports transportation of 92 antiprotons over a 7.5 km route without particle loss or measurable degradation of vacuum. This validates autonomous trap operation and logistics at tiny inventory. It supports the architecture's modular handling approach while leaving dense fuel transport to future work. [16]

Where AI can improve the factory

The factory model should jointly optimize target survival, particle production spectrum, collection acceptance, deceleration, cooling time and trap loading. Improving one variable while worsening another can reduce final yield. Accelerator simulations, surrogate models, active learning and instrument feedback can explore this trade space faster. The acceptance criterion is independently measured delivered inventory and energy use.

For scheduling estimates, a line delivering 6.4 million antiprotons every two minutes continuously would collect approximately 3 picograms per year. This arithmetic assumes no downtime and uses the per-shot result as a hypothetical sustained rate. It is not CERN's annual output. It makes the scale of the later nanogram, microgram and gram targets explicit.

Parallel lines and dedicated facilities may outperform this benchmark. Better energy sources lower operating cost, but manufacturing antimatter remains energy storage: the source must supply the energy needed to create the fuel. Recovering less energy than was used to manufacture it is compatible with a useful high-energy-density propellant.

8 Frozen fuel and the antihelium cushion

The pellet proposal should be formulated as solid molecular antihydrogen, anti-H₂. Ordinary bulk hydrogen solidifies as a molecular material. Creating that antimatter analogue would require sufficiently cold anti-atoms, controlled molecule formation, removal of association energy and nucleation without contact with ordinary surfaces. Cooling isolated trapped atoms to a low temperature does not automatically produce a solid.

NASA has discussed neutral antihydrogen microcrystals, diamagnetic or electrostatic suspension and extraction from hypothetical frozen inventory. Its concepts are a useful starting point for storage engineering, but formation and reliable high-density operation remain development tasks. [19]

The first experiment should use ordinary hydrogen particles as analogues. It should measure nucleation, sublimation, radiation absorption, electrostatic charging, force margins and position control during disturbances. A later anti-atom experiment should separately establish molecular formation and survival. Antimatter microcrystals become an architecture option only after those results.

Role of a hypothetical antihelium jacket

A liquid antihelium layer could, in principle, surround a solid anti-H₂ core without matter-antimatter annihilation at their interface. Under the usual symmetry assumptions, its thermodynamic behavior is expected to mirror helium. A helium-4 analogue could have a superfluid phase. Its possible value is temperature equilibration and management of internal motion, not a permanent source of refrigeration.

NASA experiments have formed ordinary solid hydrogen particles in ordinary liquid helium. That validates a limited material analogy, not an antimatter container. An antihelium jacket would still require its own external suspension, an evaporation and radiation model, and heat removal without ordinary-material contact. Superfluidity does not remove the need to confine the outer surface or guarantee mechanical stability. [23]

Anti-helium nuclei have been studied in collider experiments, but the required neutral atomic fluid is an entirely different production and storage problem. An incoming ordinary gas particle would annihilate in the outer layer and inject energetic products; the layer would not provide a benign shield. Its usefulness must be tested against the simpler bare-pellet approach. [24]

The decision criterion is whether a jacket reduces total fuel loss, heat load and control difficulty enough to justify the extra rare antimatter and confinement system. Early research should compare bare pellets, segmented charged microcrystals and helium-coated ordinary analogues. A bulk antihelium bath is a post-molecular-storage extension, with no defensible completion date today.

8.1 Starship cryogenics and antihydrogen storage

Starship provides an engineering precedent for cryogenic operations and infrastructure that could support future advanced propulsion. SpaceX describes a system using subcooled liquid methane and oxygen for Earth orbit, the Moon, Mars and beyond. A 2025 NASA-SpaceX study reports transfer of multiple metric tons of liquid oxygen between tanks aboard one Starship in flight. Its thermal-fluid models, low-gravity liquid positioning and vapor-management results inform spacecraft cryogenic analysis. They do not demonstrate antimatter storage or transfer between separate vehicles. [36,37]

Material or system	Reference regime	Relevance to the architecture
Liquid methane	About 111 K / -162 °C	Normal boiling-point reference for Starship fuel; actual subcooled conditions differ.
Liquid oxygen	About 90 K / -183 °C	Normal boiling-point reference for oxidizer; informs heat-leak and transfer analysis.
Solid molecular hydrogen	Near 14 K / -259 °C	Low-pressure phase analogue for solid molecular antihydrogen; lower temperatures may be needed to limit sublimation.
Webb MIRI cryocooler	Below 7 K	Demonstrated spacecraft refrigeration at a closer temperature regime; cooling capacity and loads must be matched.

These are reference regimes rather than interchangeable operating specifications. Antihydrogen phase behavior is expected to mirror hydrogen under matter-antimatter symmetry. Methane and oxygen infrastructure can inform thermal shielding and spacecraft operations, but a separate lower-temperature stage and heat balance are required for condensed fuel. [38-40]

Temperature density and molecular formation

CERN laser cooling has reduced some trapped antihydrogen atoms to kinetic energies corresponding to approximately 0.012 K. This is an atomic-motion measurement, not a demonstration of a dense molecular or solid phase. Formation of anti-H₂, removal of association and condensation heat, nucleation without material contact and long-term survival remain distinct milestones. NASA's solid-antihydrogen concepts provide a starting point for their investigation. [13,19,41]

Ordinary helium can cool the surrounding cryostat and superconducting equipment without contacting the antimatter. An antihelium cushion would require its own production, confinement and thermal model. Its role remains an optional comparison against bare-pellet storage, as described in Section 8.

8.2 Cryogenic research and acceptance milestones

The proposed research transfers methods and system interfaces where justified. Ordinary-propellant pumps, valves and transfer lines are not suitable for direct antimatter contact. A condensed antimatter module needs a dedicated field-based confinement and extraction path inside a separately engineered thermal environment. SpaceX-NASA cryogenic work and NASA's broader fluid-transfer program provide relevant spacecraft precedents. [19,36,42]

Stage	Research task	Evidence for advancement
C1 Thermal model	Evaluate staged insulation, shields, cryocoolers, residual gas and parasitic loads using public data.	Closed heat and electrical-power balances, cooling capacity at each temperature, mass and uncertainty.
C2 Ordinary analogue	Study hydrogen particle formation, noncontact suspension, thermal stability and metering through a qualified host.	Measured nucleation, sublimation, position stability, transfer survival and recovery after disturbances.
C3 Antimatter synthesis	Investigate molecular association, cluster formation and condensed phases with an agreed antimatter facility.	Species identification, formation yield, inventory, lifetime and independent verification of confinement.
C4 Flight qualification	Assess a dedicated storage module after usable inventory and supply are demonstrated.	Thermal-vacuum, vibration and acceleration performance; full mission and refrigeration budgets.

NASA has formed ordinary solid hydrogen particles in liquid helium. These experiments support analogue design but do not establish wall-free antimatter nucleation. A bare-pellet reference and an optional cushion configuration should be compared under the same heat input, disturbance and inventory assumptions. A cushion advances only if its measured or independently validated benefit exceeds its added production and confinement burden. [19,23]

Useful transfer from Starship research includes thermal modeling, instrumentation, low-gravity operations and eventual depot logistics. Colder spacecraft technology, such as Webb's cryocooler, informs the inner thermal stage. Neither establishes the cooling power or confinement performance of a proposed antimatter device without a matched engineering analysis. [36,40,42]

C1 and the design of C2 are appropriate early study tasks. Host access and quotations determine which analogue measurements can be purchased. C3 and C4 require separately justified resources and successful earlier stages. This sequence complements the prototype gates in Section 21 and retains the October 2026 planning origin.

9 Metering antimatter flow

The regulator controls how much antimatter reaches ordinary matter. The first practical analogue to a fictional dilithium regulator is a measured extraction and transfer system, physically separated from storage and the reaction region. It does not require a material that lets antimatter pass through its atoms without interacting.

The earliest prototype should use counted charged-particle batches. A later neutral system could transfer small calibrated packets or charge selected microcrystals before steering them. Laser-induced charging and extraction are prospective approaches discussed in NASA's concepts. A neutral pellet cannot be steered as if it were a cloud of charged antiprotons; its magnetic response and electrical charging need separate characterization. [19]

Control variable	Measured quantity	Development objective
Packet inventory	Particle count or calibrated pellet mass	Bound pulse energy before transfer.
Repetition and duty cycle	Arrival time and spacing at the reaction region	Set average power without overloading conversion hardware.
Transfer efficiency	Delivered fraction versus unintended loss	Prevent wall interception and unmeasured reactions.
Reaction products	Energy, angular distribution and deposited heat	Calibrate usable power and shielding requirements.
Shutdown response	Fuel delivered after a stop command	Limit the last packet and prevent subsequent deliveries.

Useful average power is determined by antimatter mass flow and conversion efficiency when equal ordinary mass is supplied. The regulator therefore closes its loop on measured power, loss and timing rather than on an assumed relation between pellet rate and warp speed.

$$P_{\text{use}} = 2\eta_{\text{conv}} \frac{d\bar{m}}{dt} c^2$$

Annihilation of antiprotons with protons produces pions and subsequent decay products. Some energy appears in charged particles that magnetic fields can guide; some appears in gamma rays or escaping neutrinos. Ordinary magnetic nozzles do not guide gamma rays. Secondary nuclear interactions complicate heavier targets. NASA's propulsion studies analyze different ways to couple these products to a working fluid or direct thrust. [20,21]

Compare three conversion branches: deposited heat driving electricity or hot propellant; magnetic coupling of charged products to thrust; and antimatter-assisted fission or fusion, where antimatter initiates additional nuclear energy release. The last branch could reduce antimatter demand but adds ignition, confinement, fuel and reaction-gain requirements; NASA's published studies are conceptual estimates. The tables here describe direct annihilation. A nuclear-assisted system needs a separate energy balance and cannot inherit their assumed efficiency. [20,21]

The small-inventory prototype is primarily a diagnostic experiment, not a net-power plant. It must first measure losses and conversion. High-power metering becomes relevant only when production, stored inventory and heat handling can support it.

10 Power and inventory estimates

All values in this section are original arithmetic using $c = 299,792,458$ m/s. Complete annihilation of an antimatter mass m with equal ordinary mass releases total energy $2mc^2$ across all products. A nominal useful fraction of 30 percent is an illustrative target, not a demonstrated full-engine efficiency.

$$E_{\text{total}} = 2mc^2$$

Antimatter per packet	Total energy	Useful energy at 30 percent
1 femtogram	0.180 J	0.054 J
1 picogram	180 J	54 J
1 nanogram	180 kJ	54 kJ
1 microgram	180 MJ	54 MJ
1 milligram	180 GJ	54 GJ

Dense fuel reduces tank volume, but it also makes a single lost packet significant. A first microcrystal program should work toward the lowest mass that can be characterized reliably. Even a one-nanogram packet contains much more energy than a normal laboratory particle-trapping event.

Useful average power	Antimatter flow at 30 percent	Fuel for one continuous year
100 W	1.85 pg/s	58.5 micrograms
10 kW	0.185 ng/s	5.85 milligrams
1 MW	18.5 ng/s	0.585 grams
100 MW	1.85 micrograms/s	58.5 grams
1 GW	18.5 micrograms/s	585 grams

A 1 MW burst for one hour would consume about 66.8 micrograms at the same efficiency. Continuous operation requires 8,760 times the one-hour inventory each year. Nominal burst power, stored fuel and sustainable factory throughput must therefore be specified independently.

For design sensitivity, evaluate useful fractions of 1, 10, 30 and 50 percent. Fuel scales inversely with that fraction. The upper cases assume substantial advances in product capture and conversion; they should not be used to certify a vehicle mass budget.

11 Production electricity and heat rejection

Define production efficiency as the fraction of facility electricity embodied in the rest energy of a newly created captured antimatter amount and its equal-mass matter counterpart. This definition includes collection, cooling and preparation losses. It is an accounting convention for this estimate; facilities may report differently defined efficiencies.

$$\frac{d\bar{m}}{dt} = \frac{\eta_{\text{prod}} P_{\text{input}}}{2c^2}$$

Assumed overall efficiency	Yield per year at 1 MW input	Yield per year at 1 GW input
0.00001 percent	17.6 nanograms	17.6 micrograms
0.001 percent	1.76 micrograms	1.76 milligrams
0.1 percent	0.176 milligrams	0.176 grams

These are sensitivity cases, not measurements of an existing source. Producing 10 grams per year would require about 57 GW at the optimistic 0.1 percent efficiency, or a new source achieving about 5.7 percent at 1 GW. This gives the extreme roadmap target a concrete infrastructure requirement. A 0.1 percent production efficiency followed by 30 percent conversion returns only 0.03 percent of the original factory electricity as useful onboard energy.

Heat rejection is a separate vehicle constraint. If all energy not converted to useful output is deposited onboard, a 1 MW useful reactor at 30 percent efficiency rejects about 2.33 MW. In reality some energy may escape directly as radiation, changing both the thermal and radiation design. Treat that partition explicitly.

$$A = \frac{Q}{\epsilon \sigma T^4}$$

Radiator temperature	Effective emitting area for 2.33 MW	Design implication
300 K	About 5,645 m ²	Low temperature demands a large structure.
600 K	About 353 m ²	Requires compatible hot conversion hardware.
900 K	About 70 m ²	Smaller area with tougher materials and interfaces.

The area calculation assumes emissivity 0.9, an unobstructed view of cold space and no incoming load. It is total radiating surface area, not necessarily projected panel area. Plumbing, redundancy and unfavorable view factors increase it. Cryogenic traps need separate low-temperature heat removal; a hot radiator cannot directly maintain an anti-atom trap.

12 A physical field actuator

The field module begins as an interface specification. For each candidate metric, the team computes the required stress-energy tensor, then asks whether an actual matter, electromagnetic or quantum-field model can produce it. Its constitutive equations, conservation laws and stable dynamics must accompany the metric. Treating a tensor derived from a desired geometry as a buildable plasma distribution skips this essential step.

The geometry team should search beyond restrictive metric assumptions where justified, including lapse, spatial curvature and physically appropriate boundary conditions. It should minimize peak energy density, integrated source energy, stress, formation work, radiation and sensitivity to fabrication error together. Lowering a volume integral while increasing unachievable local pressures is not a feasible optimization.

For orientation, ordinary energy creates extraordinarily weak gravity at laboratory scale. The dimensionless scale of a weak gravitational perturbation from energy E distributed over length L is approximately GE divided by $c^4 L$, up to geometry-dependent factors.

$$h_{\text{scale}} \sim \frac{GE}{c^4 L}$$

For $E = 1$ MJ and $L = 1$ m, this scale is about 8.3×10^{-39} . The entire annihilation energy of one gram of antimatter with one gram of matter gives about 1.5×10^{-30} at the same length. These are dimensional estimates, not detector transfer functions, but they show why a conventional high-power magnetic apparatus does not become a useful warp device just by concentrating its power.

Three field development branches

The first branch uses ordinary gravity and well-characterized electromagnetic matter sources to establish calculation and measurement consistency. The second studies quantum stress-energy, including negative-energy effects where permitted and their quantum-inequality bounds. The third tests explicitly defined modifications of gravity or new matter couplings against independent experimental constraints. [10-12]

An optical or condensed-matter analogue can reproduce aspects of an effective metric for waves. It is valuable for studying propagation and control, but it does not demonstrate curvature that transports an arbitrary material payload. Report analogue and gravitational results separately.

The eventual actuator must specify stored source energy, creation time, maintenance losses, available control channels and the method of stopping. Reactor power is the rate at which energy is supplied or losses replaced; it is not the same quantity as the energy integral of the spacetime. Recoverable field energy should be tracked separately from dissipated energy.

13 Verification and measurement

A candidate geometry must survive more than one codebase. Use symbolic tensor calculations, a high-precision numerical implementation and an independent solver. Start with known metrics and the published Warp Factory toolkit as benchmarks. Compare coordinate systems, tetrads, mesh convergence and boundary treatments. A search agent should propose candidates and tests, while independent numerical and analytic checks determine acceptance. [7]

Test	Observable or calculation	Acceptance principle
Tensor reconstruction	Einstein tensor and the proposed source tensor	Agreement within controlled numerical and model error.
Energy conditions	Minimum contractions over timelike and null directions	Analytic or certified bounds where possible; sampled passes are insufficient.
Dynamics	Evolution from initial data through shutdown	Causal controls and consistent constraints at every step.
Stability	Perturbations and quantum backreaction	Growth bounded over the full operating interval.
Physical measurement	Clock comparison, optical paths, matter trajectories and tidal response	Independent sensors reproduce the predicted effect after systematic errors.
Transport	Arrival events compared with reference null trajectories	A geometric and operational travel result rather than a coordinate artifact.

The first instrument platform should combine radiation calorimetry, spectroscopy, field mapping, vacuum diagnostics and trap image-current or annihilation measurements. Any gravitational experiment adds clock or interferometric observables appropriate to a stated predicted signal. Its sensitivity target must follow from that model, rather than being selected to imply that laboratory energy has become detectable curvature.

Blind changes, ordinary-matter control runs, sensor substitution, polarity and geometry variations, and independent replication help distinguish a physical signal from thermal expansion, magnetic forces, charging, refractive-index shifts and vibration. A null result with a calibrated sensitivity bound is useful evidence for refining the model.

Local controllers should stop extraction before the next packet, isolate compartments and record the final delivered inventory. The cloud can optimize the next experiment, but networking delay and unavailable communications cannot be part of a critical confinement loop. A future superluminal vehicle would require an additional causal-control analysis; cloud connectivity cannot be assumed during transport.

14 AI factories and recursive improvement

There is a concrete basis for planning around much larger computing systems. RIKEN's FugakuNEXT program targets operation around 2030 and zettascale AI performance at FP8 precision. Its May 2026 basic-design publication gives the program an engineering basis. A zettaFLOP is 10^{21} floating-point operations per second. This is a plan for an AI-oriented performance regime, not a guaranteed 10^{21} sustained double-precision operations for arbitrary numerical relativity. [25,26]

NVIDIA's June 2026 scientific-computing announcement reports a rack-scale configuration with more than 7 exaFLOPS of AI-for-science performance and 5 petaFLOPS of native FP64 capability. Those figures differ by roughly 1,400 in nominal rate and refer to different numerical operations. Mixed precision, surrogates and machine learning can exploit the faster hardware, while decisive physical checks still require quantified errors and suitable precision. [27]

Recursive improvement could accelerate the research system itself. AI agents may improve solver kernels, compilers, experiment selection, models and chip designs, then use those improvements in their own development. Anthropic reports increasing AI contributions to its internal development while explicitly stating that fully autonomous recursive self-improvement has not been reached and is not inevitable. Its software-output measurements support acceleration within that setting; they are not a measured rate for discovering new spacetime physics. [28]

Google DeepMind reports AlphaEvolve applications across algorithms, computing infrastructure and scientific problems. Such systems suggest a useful prototype workflow: propose a change, run a reproducible evaluator, keep verified gains and repeat. The program should optimize independent replication rate, usable captured fuel and material lifetime rather than code volume or plausible explanations. [29]

A physically grounded improvement loop

Each model update should be tested on held-out benchmarks and new measurements. Mathematical claims need independent tensor calculations; material claims need synthesis and characterization; control gains need hardware tests outside the training conditions. A stronger model can increase the rate of proposed ideas, but shared evaluators can amplify a common mistake. Preserve independent measurements and diverse verification methods.

Even sustained exponential growth in a selected capability eventually meets bottlenecks in experiments, fabrication, energy, data or verification. If half of a sequential project cycle is irreducible hardware time, making its computation a million times faster reduces that cycle by only about twofold. Parallel laboratories and automated manufacturing must expand alongside compute to obtain much larger calendar gains.

15 Quantum computation and automated laboratories

Quantum computation is best treated as a specialized accelerator within the larger research platform. Plausible uses include selected electronic-structure and defect calculations, small quantum-field models, and optimization problems with a demonstrated advantage. A useful algorithm must include error correction, state preparation, input loading, measurement cost and the precision of its returned answer.

IBM's roadmap targets Starling in 2029 with 200 logical qubits and circuits containing 100 million gates. This is an announced target, not present capability. The Willow error-correction experiments demonstrate suppression of logical errors below a code threshold; they do not yet establish universal practical quantum acceleration. A July 2026 experiment reports reinforcement-learning stabilization of quantum error correction against injected drift, illustrating a real interaction between AI and quantum control. [32-34]

For a materials or field subproblem, the program should estimate logical qubits, gate count, error budget and elapsed time, then compare the complete quantum workflow with the best classical method at the same accuracy. Quantum optimization has opportunities but no generic guarantee of exponential advantage for engineering search. Its value must be established for the actual workload. [35]

Computing layer	Best initial use	Validation requirement
AI agents and surrogates	Generate and screen candidates; write and optimize software	Held-out prediction and independent numerical checks.
Classical HPC	High-precision relativity, particle transport and multiphysics	Convergence, conservation and measured benchmark agreement.
Quantum accelerators	Selected quantum matter or optimization kernels	End-to-end benefit at equal accuracy, including overhead.
Automated laboratories	Fabricate, characterize and repeat experiments	Calibrated measurements and replicated physical improvements.

AI plasma control has already been demonstrated in a tokamak, and autonomous materials laboratories have linked computation, synthesis and characterization. These are relevant precedents for adaptive traps and materials development, although transfer to antimatter must be tested. [30,31]

The proposed cloud system federates laboratories through common data formats, queues and provenance. It dispatches simulations to high-precision systems, screening to AI hardware and suitable subproblems to quantum processors. It measures completed verified experiments per month. Distributed experimental capacity, reliable fabrication and instrument uptime determine whether extreme compute growth produces extreme physical progress.

16 Accelerated research scenarios

The figures below are original planning bands for effective verified research throughput relative to a comparable 2026 program. They are not forecasts of intelligence, raw FLOPS, the probability of discovery or the speed of every physical experiment. The selected portfolio includes numerical screening, analysis, software verification and automated laboratory trials. Reaching the high bands requires many parallel facilities as well as better algorithms.

Scenario	Five years	Ten years	Twenty years
Steady improvement	3 to 20 times	10 to 100 times	30 to 1,000 times
Accelerated AI and automation	10 to 100 times	100 to 10,000 times	1,000 to 100,000 times
Recursive AI breakthrough	100 to 1,000 times	1,000 to 100,000 times	10,000 to 1,000,000 times

The extreme scenario assumes capable AI research agents improve their own tools and models, while robotic laboratories and fabrication capacity expand fast enough to supply new grounding data. Its upper bounds describe the possible scale of an organized research enterprise under those assumptions. They are intentionally aggressive stress tests for architecture and resource allocation.

A capability doubling every six months would grow about 1,000 times in five years and about one million times in ten years if sustained. This arithmetic illustrates why recursive improvement could change planning. It does not establish that general research ability follows that trend. Extending the same exponential for twenty years would disregard resource, evaluation and experimental bottlenecks.

Where additional capacity would go

Expand source simulations and end-to-end capture searches first, then high-precision metric and matter-source verification, radiation transport and materials models. Build enough ordinary-material test stations to challenge surrogate predictions. Use fault-tolerant quantum machines for subproblems with proven end-to-end benefit. By 2036, useful quantum workloads might be substantially larger than the announced 200-logical-qubit 2029 target, but the paper does not assume that a specific larger machine or useful FTL calculation will exist.

Beyond 2046, the decisive variable is the accumulated verified knowledge and physical infrastructure. An ASI-enabled program could discover a new source interaction, a more useful geometry or a much better fuel factory. The architecture is designed to absorb those discontinuities through replaceable interfaces. It also accommodates rigorous exclusions of candidate mechanisms.

17 Initial engineering envelope

The first prototype is a ground-based research assembly attached to an established antimatter source. The ranges below select a modest experiment class for integration studies. They are targets to refine with a host laboratory and measured loads, rather than a claim that every value is achieved simultaneously.

Parameter	Initial planning range	Purpose
Active trap or test volume	0.1 to 1 m characteristic dimension	Integrate diagnostics and manipulate small inventories.
Superconducting confinement field	1 to 5 T as a starting class	Build around established trap methods; larger fields require a separate design.
Cold-region pressure	10^{-12} to 10^{-16} mbar target band	Study loss as a measured function of residual gas and cryogenic conditions.
Neutral atomic energy scale	1 to 100 mK research objective	Improve trapped populations; this does not imply molecular condensation.
Counted antiproton packet	10^5 to 10^8 particles	About 30 microjoules to 30 millijoules total annihilation energy.
External conventional pulsed energy	1 kJ to 1 MJ test-system class	Study power and actuator interfaces without using bulk antimatter.

Cryogenic autonomous traps have demonstrated extremely low residual-gas conditions, and antihydrogen laser cooling is established. Different trap regions, particle species and operating cycles need different temperatures and field profiles. A 1 to 100 mK kinetic-energy scale for atoms is distinct from the thermal environment of the cryostat and the temperature of any future condensed pellet. [13,16]

The packet range is an inventory target for isolated pulses, not a claim of continuous supply. A repetition rate must be derived from measured delivered particles per second, trap replenishment time and reactor diagnostics. Historical NASA high-density trap studies also distinguish designed storage capacity from an actually commissioned antimatter inventory. [22]

Inverse energy budget for a later field demonstrator

A future test whose physical source fits in 1 GJ to 1 TJ would be a substantially different engineering project. At 30 percent conversion, depositing that energy requires about 18.5 micrograms to 18.5 milligrams of antimatter, plus equal ordinary matter. The values are hypothetical energy envelopes chosen to investigate system implications. No reviewed warp theory establishes that such energy is sufficient.

Relative to the Fell-Heisenberg example, fitting that envelope would require about 32 to 35 orders of magnitude less energy. A claimed route must explain this physically and supply the complete tensor and dynamics. An ASI-assisted search may find a different mechanism or a restriction, but a large computing multiplier alone cannot be substituted for those missing reductions.

18 Roadmap of linked innovations

The map distinguishes engineering progress from physical discovery. Its dates are planning horizons counted from October 2026. Later field demonstrations depend on successful tests, even in the strongest AI scenario. Earlier verified breakthroughs could move milestones forward; absent breakthroughs, hardware work continues on the engineering branches.

	2031 Five years	2036 Ten years	2046 Twenty years	2046 plus Later horizon
Research platform	Verified solvers and automated trials	Larger AI research and quantum kernels	Expanded autonomous laboratories	Continuously validated research
Antimatter fuel	Better capture and modular traps	Molecule and cluster experiments	Microcrystals if formation succeeds	Mission-scale supply and containment
Power and flight	Counted reaction pulses	Measured energy conversion	Burst power or uncrewed propulsion	Qualified continuous operation
Spacetime physics	Source-model search and critical tests	Physical field test if a source is established	Replicated and scalable effect if available	FTL transport only after all physics gates
Antihelium branch	Ordinary helium analogues	Quantify benefit against bare storage	Neutral fluid research after molecule success	Bulk jacket only if it helps containment

Dates do not replace experimental acceptance tests.

Figure 2 Dependency roadmap. Blue cells show engineering objectives. Amber cells require new physical results. The antihelium branch follows successful neutral molecular storage. Horizon dates are conditional planning targets.

Three scenarios

Steady progress expands existing research tools, source utilization and conventional spacecraft engineering. Accelerated discovery combines strong AI research agents, larger facilities, automated laboratories and selected useful quantum algorithms. The extreme scenario additionally assumes sustained recursive improvements and a verified new spacetime source mechanism. ASI is an explicit assumption in this third scenario, not a dated prediction.

The practical ordering is source yield before bulk fuel; controlled molecules before crystals; measured power conversion before a flight reactor; validated source physics before a spacetime actuator; and verified formation, control and braking before a transport demonstration. Each branch has usable results before the full sequence succeeds.

19 Five year and ten year objectives

By 2031

The most credible objective is an integrated ground laboratory using an existing accelerator, increased trapping throughput, ordinary-material pellet analogues, precision pulse metering and a verified tensor-analysis platform. AI agents should help operate simulation and analysis pipelines; quantum resources should be used only after workload benchmarks show value. The spacetime goal is a stronger physical model or a sharper limitation, with no assumed FTL demonstration.

By 2036

The accelerated objective is a dedicated production pilot, modular traps, experiments on antihydrogen molecules or small clusters, and measured reaction-energy coupling. A low-inventory uncrewed flight experiment could be investigated if qualification and supply permit it. Genuine field hardware remains dependent on a source mechanism and a detectable predicted signal.

Measure	2031 accelerated target	2036 breakthrough target
Captured antimatter per year	0.1 to 10 ng	0.1 to 100 micrograms
Modular charged inventory	10^9 to 10^{12} antiprotons	10^{12} to 10^{15} antiprotons
Neutral atomic inventory	10^5 to 10^7 anti-atoms	10^8 to 10^{12} anti-atoms
Pellet development	Ordinary analogues; molecule pathways	Verified anti-molecules or clusters if formation succeeds
Reaction demonstration	Counted pulses and energy accounting	Pulsed conversion with measured deposited heat
Spacetime result	Independent tests of candidate theories	A source-based model or a calibrated physical test if available

The numerical bands are ambitious original targets. They refer to an integrated program with expanding source capacity, not an extrapolated promise from a single CERN trap. Neutral and charged inventories have different limits and must not be added as if they were interchangeable storage technologies.

At the illustrative three-picogram annual benchmark, the 2031 production band requires about 30 to 3,000 times more delivered inventory. The 2036 band requires roughly 30,000 to 30 million times more. The upper ten-year case therefore needs major source, collection and parallel-facility breakthroughs, not just improved software.

20 Twenty year and later objectives

By 2046

In an aggressive scenario, a neutral storage breakthrough and industrial source infrastructure could support microcrystal fuel handling, short power bursts or an uncrewed antimatter-assisted propulsion demonstrator. A production objective of 0.01 to 10 grams per year is a stretch band requiring very large advances in efficiency and facility capacity. The upper end has the explicit tens-of-gigawatts requirement shown in Section 11 at 0.1 percent efficiency.

For neutral stored inventory, a picogram to a microgram is a useful experimental target band if molecules and condensed phases become feasible. It is not enough for a long-duration megawatt vehicle. A gram-class power mission needs independently demonstrated inventory, containment reliability and annual fuel supply. The small fuel volume of a solid does not establish the mass of its supporting system.

Branch	Optimistic 2046 outcome	Beyond 2046 requirement
Fuel factory	Industrial pilot or multiple dedicated sources	Sustainable grams to kilograms for the selected mission.
Storage and regulation	Verified microcrystals if molecular milestones succeed	Long-lived inventory and reliable transfer under flight loads.
Power and flight	Burst power or antimatter-assisted uncrewed propulsion	Continuous power, full radiation qualification and mission supply.
Spacetime source	Replicated field effect if a physical mechanism is found	Stable scalable transport, causal control and braking.
Antihelium cushion	Ordinary analogues and a quantified benefit model	Neutral antihelium fluid production and independent confinement.

The conditional superluminal route

A genuinely new physical result within the next five to ten years could justify a small field prototype during the following decade. If that prototype proves a controllable, scalable effect, the 2046-plus period could include an uncrewed transport experiment. The statement is a sequence of necessary conditions, not a forecast that such a result will occur. There is no evidence-based year for a crewed FTL starship.

ASI could change both discovery speed and the range of mechanisms considered. A complete source model might arrive earlier than expected, or stronger reasoning could show that a promising family fails. Engineering success in the antimatter branch remains valuable in either case. Any eventual FTL route must satisfy or explicitly revise the relevant physical constraints and agree with existing experimental evidence.

21 Prototype stages and acceptance tests

Stage	Prototype content	Acceptance test
P0 2026 to 2031	Ordinary analogues, source access, counted antimatter batches, simulation and diagnostics	Reproduce benchmarks; account for packet delivery and loss; validate each measurement chain.
P1 2031 to 2036	Improved capture, modular confinement, neutral synthesis research and reaction coupling	Increase delivered yield; measure conversion and heat; demonstrate reliable stop and recovery.
P2 2036 to 2046	Conditional microcrystals, larger inventory, an uncrewed energy or propulsion experiment	Demonstrate storage lifetime, transfer survival and full power budget under the selected mission loads.
P3 after physics validation	A source-based field apparatus using any appropriate power source	Measure the predicted genuine geometric effect and independently reproduce it.
P4 after P3 scaling	Uncrewed transport with controllable formation and termination	Demonstrate causal operation, navigation, braking and comparison with reference light propagation.

Requirements that cannot be delegated to a calendar

The physics gate requires a complete tensor and source model, controlled mathematical error, a causal creation history, stability, and an observable prediction. The actuator gate requires a measured physical effect with independent replication. The transport gate additionally requires an operational trajectory and safe termination. Passing an energy-density check or increasing reactor power does not pass these gates.

The fuel gate requires a measured increase in delivered inventory rather than particles created elsewhere in the facility. The storage gate requires survival and loss bounds at the intended inventory, temperature and disturbance level. The power gate requires an energy balance across useful output, escaping products and deposited heat. The flight gate requires qualification of the entire system, including conventional spacecraft functions.

If an experiment contradicts a candidate field model, preserve the result and redirect that branch. If the physical source remains absent, keep P3 and P4 as research interfaces while advancing P0 through P2. This prevents a large hardware commitment from resting on a missing coupling between power and geometry.

22 Resources and research priorities

Program spending should follow measured improvements and verified physics. The following order-of-magnitude envelopes are original planning allowances in 2026 US dollars. They are not supplier quotations. They exclude construction of a global AI factory, an existing host accelerator, fuel mass production for a full mission and any unverified spacetime hardware.

Program element	Illustrative envelope	What it funds
First five-year research platform	\$25 million to \$150 million	Teams, source access, computing allocations, traps, analogue benches and metrology.
Expanded integration program	\$0.2 billion to \$2 billion	Additional source capacity, automated labs, materials and thermal engineering.
Dedicated production pilot	\$1 billion to \$10 billion	Accelerator and target infrastructure, capture, cooling and industrial throughput research.
Qualified uncrewed flight program	\$1 billion to \$10 billion	Mission-specific power, containment, spacecraft integration and demonstration.
FTL hardware	Not presently estimable	Requires a validated source mechanism before cost or mass estimates are meaningful.

These are separate work-package ranges with potential overlap; they must not be summed without a defined schedule and system boundary. The wide bands reflect different facility partnerships and scope. Review them after the first measured yield, thermal and reliability data become available.

Highest value innovations

First, improve capture-to-storage efficiency, because every later experiment benefits from more usable particles. Second, improve simulation verification and source-based metric search, because a correct impossibility result or a viable mechanism both prevent misdirected hardware spending. Third, expand automated fabrication and characterization so computation can produce physical evidence faster. Fourth, develop contact-free neutral handling and measured conversion. The antihelium jacket follows a demonstrated benefit over simpler storage.

The program should report useful particles per joule, verified experiments per month, trap lifetime at stated inventory, energy converted per delivered packet, heat and shielding mass per useful watt, and replicated field predictions. These measures remain meaningful across AI capability changes. They also reveal whether a claimed breakthrough merely transfers the bottleneck to another module.

22.1 Prospective SpaceX research collaboration

SpaceX's published Moon, Mars and beyond objectives create a strategic overlap with the long-term exploration goals of this architecture. The company is a prospective spacecraft and cryogenic engineering partner; Elon Musk and SpaceX leadership are potential strategic contacts. Any research role, resource contribution or funding arrangement is subject to technical review and agreement. [37]

Stage of engagement	Potential contribution	Decision evidence
Technical review	Assess public cryogenic results, storage requirements and candidate mission interfaces.	A defined partner problem, technical counterpart and scoped benefit calculation.
Component research	Negotiate a modeling, instrumentation or ordinary-material test package, potentially with in-kind resources.	Measured performance, costed scope, background intellectual property and publication terms.
Future flight assessment	Evaluate launch or module support with SpaceX or another qualified provider after laboratory milestones.	Validated storage, sufficient fuel supply, agreed spacecraft interfaces and a funded mission case.

The immediate partnership case is a specific engineering result with value before an antimatter-powered starship exists. Thermal models, diagnostic methods and validated research-control tools could address conventional cryogenic needs while enabling later fuel work. Starship progress supplies relevant experience and possible future infrastructure; FTL transport still requires the independent source-physics milestones in Sections 2, 3 and 12.

22.2 Staged execution and continuing employment

The author can begin with part-time architecture, semantic AI software, source review and reproducible models while continuing professional employment. The initial purpose is a technically reviewed study and a costed next experiment. Laboratory work requires a qualified host, specialist investigators and confirmed access. A change to full-time research depends on funding sufficient to support the agreed effort and program responsibilities.

Hathaway's proposed role is system architecture, semantic AGI/ASI research strategy and integration. Cryogenic, antimatter and spacetime specialists lead their experimental and theoretical work, with independent verification. The initial program can proceed without a SpaceX commitment, and each prospective partner is approached with a bounded deliverable and resource request.

The companion funding proposal translates this research architecture into a nine-month planning study and a phased investment roadmap. Its \$500,000 initial planning request comprises separate \$225,000 NASA-facing and \$275,000 private packages. Those figures are proposed budgets pending host rates, quotations, eligibility and award terms; the technical paper's larger research envelopes remain separate. [43]

23 Conclusions and quantitative limits

A hypothetical prototype can serve as a useful architecture now. Its immediately buildable core combines accelerator access, precision traps, ordinary-material storage analogues, counted delivery, radiation measurement, classical multiphysics and AI-assisted research. Better source engineering, autonomous laboratories and larger computing facilities can make this platform substantially more capable over five and ten years.

Twenty-year expectations become much more sensitive to assumed breakthroughs. Molecule formation and microcrystal storage could change containment density. Efficient industrial production could make antimatter a useful specialized propellant or energy store. A verified spacetime source could change the whole program. Each needs its own demonstration; none follows automatically from the existence of powerful AI, quantum processors or an energy-rich fuel.

The guiding architecture retains the pellet, cushion and regulated-flow ideas as researchable modules. It couples their engineering to a complete physical source model rather than an assumed warp coil. The strongest defensible optimism is that radically better research and fabrication could explore far more possibilities, replicate results faster and discover mechanisms unavailable to present engineering. Assigning a dated FTL capability requires evidence that has not yet been established.

Quantitative assumptions

Quantity	Value or definition
Antiproton or approximate anti-atom mass	1.6726×10^{-27} kg; positron mass neglected in small estimates.
Energy per antimatter mass	$2c^2 = 1.7975 \times 10^{17}$ J/kg when equal matter is supplied.
Year and duty cycle	365 days; annual power table assumes uninterrupted operation.
Conversion fraction	30 percent for examples; sensitivity range 1 to 50 percent.
Production fraction	Pair-equivalent captured rest energy divided by factory electricity.
Radiator model	Emissivity 0.9; no incoming flux; area counts emitting surfaces.
Future dates and ranges	Conditional planning targets; no probabilities or guaranteed completion dates.

Reproduction checks: 15,000 anti-atoms have mass about 2.51×10^{-20} g. One nanogram contains about 5.98×10^{14} anti-atoms. A 1 MW useful system at 30 percent requires 18.5 ng/s. Field energy and maintenance power must be calculated separately for every candidate model.

References for spacetime physics

Primary research, institutional technical reports and identified provider plans. Sources were checked on 10 October 2026. Preprints and vendor targets are labeled. Numerical scenarios, prototype envelopes and program budgets are original estimates in this paper.

[1] Lentz, E. W. Breaking the Warp Barrier Hyper-Fast Solitons in Einstein-Maxwell-Plasma Theory. 2021 journal paper; arXiv preprint 2006.07125.

<https://arxiv.org/abs/2006.07125>

[2] Bobrick, A. and Martire, G. Introducing Physical Warp Drives. 2021.

<https://arxiv.org/abs/2102.06824>

[3] Fell, S. D. B. and Heisenberg, L. Positive Energy Warp Drive from Hidden Geometric Structures. 2021 journal version.

<https://arxiv.org/abs/2104.06488>

[4] Santiago, J., Schuster, S. and Visser, M. Generic warp drives violate the null energy condition. 2021 preprint and subsequent journal publication.

<https://arxiv.org/abs/2105.03079>

[5] Celmaster, B. and Rubin, S. Violations of the Weak Energy Condition for Lentz Warp Drives. 2025 preprint.

<https://arxiv.org/abs/2511.18251>

[6] Fuchs, J. and colleagues. Constant Velocity Physical Warp Drive Solution. 2024.

<https://arxiv.org/abs/2405.02709>

[7] Helmerich, C. and colleagues. Warp Factory A Numerical Toolkit for the Analysis and Optimization of Warp Drive Geometries. 2024 preprint.

<https://arxiv.org/abs/2404.10855>

[8] Garattini, R. and Zatrimeylov, K. Warp Drive in a De Sitter Universe. 2025 preprint revised 17 June 2026.

<https://arxiv.org/abs/2502.13153v4>

[9] Santos-Pereira, O. L. The Warp Drive Superluminal Travel within General Relativity. 2025 doctoral thesis.

<https://arxiv.org/abs/2508.20348>

[10] Olum, K. D. Superluminal travel requires negative energies. 1998.

<https://arxiv.org/abs/gr-qc/9805003>

[11] Finazzi, S., Liberati, S. and Barcelo, C. Semiclassical instability of dynamical warp drives. 2009.

<https://arxiv.org/abs/0904.0141>

[12] Pfenning, M. J. and Ford, L. H. The unphysical nature of Warp Drive. 1997 paper; corrected preprint version 2001.

<https://arxiv.org/abs/gr-qc/9702026>

References for antimatter engineering

[13] ALPHA Collaboration. Laser cooling of antihydrogen atoms. Nature 2021.

<https://doi.org/10.1038/s41586-021-03289-6>

[14] ALPHA Collaboration. Be assisted simultaneous confinement of more than 15000 antihydrogen atoms. Nature Communications 2025.

<https://doi.org/10.1038/s41467-025-65085-4>

[15] Harris, R. and CERN. Breakthrough in antimatter production. 18 November 2025.

<https://home.web.cern.ch/breakthrough-antimatter-production/>

[16] Leonhardt, M. and BASE Collaboration. Road transport of trapped antiprotons. Nature 16 September 2026.

<https://doi.org/10.1038/s41586-026-11019-z>

[17] Lee, B. and colleagues. Record accumulation of antiprotons in a Penning-Malmberg Trap and their preparation for improved production of antihydrogen beams. 22 March 2026 preprint.

<https://arxiv.org/abs/2603.21168>

[18] ALPHA Collaboration. Observation of the effect of gravity on the motion of antimatter. Nature 2023.

<https://doi.org/10.1038/s41586-023-06527-1>

[19] LaPointe, M. and NASA. Antimatter Propulsion. 2020 presentation NTRS 20200001904; includes solid antihydrogen concepts.

<https://ntrs.nasa.gov/citations/20200001904>

[20] Frisbee, R. H. How to build an antimatter rocket for interstellar missions systems level considerations in designing advanced propulsion technology vehicles. 2003 conference paper NTRS 20060043411.

<https://ntrs.nasa.gov/citations/20060043411>

[21] Schmidt, G. R., Gerrish, H. P., Martin, J. J., Smith, G. A. and Meyer, K. J. Antimatter Requirements and Energy Costs for Near-Term Propulsion Applications. 1999 NTRS 19990110316.

<https://ntrs.nasa.gov/citations/19990110316>

[22] Smith, G. A. and Kramer, K. J. Enabling Exploration of Deep Space High Density Storage of Antimatter. 1999 NASA record NTRS 20020043303.

<https://ntrs.nasa.gov/citations/20020043303>

[23] Palaszewski, B. Solid Hydrogen Experiments for Atomic Propellants. 2001 NASA technical memorandum NTRS 20020036213.

<https://ntrs.nasa.gov/citations/20020036213>

[24] CERN and ALICE Collaboration. ALICE estimates how transparent the Milky Way is to antimatter. 2022 research announcement and linked Nature Physics study.

<https://home.cern/news/news/physics/alice-estimates-how-transparent-milky-way-antimatter>

References for enabling technologies

- [25] RIKEN Center for Computational Science. FugakuNEXT. Program overview accessed 10 October 2026.
<https://www.r-ccs.riken.jp/en/fugaku-next/>
- [26] RIKEN Center for Computational Science. Release of the FugakuNEXT Basic Design Technical Report. 29 May 2026.
<https://www.r-ccs.riken.jp/en/outreach/topics/20260529-1/index.html>
- [27] NVIDIA. NVIDIA Vera Rubin Delivers World-Class Supercomputers for Science. 22 June 2026 manufacturer announcement.
<https://nvidianews.nvidia.com/news/nvidia-vera-rubin-delivers-world-class-supercomputers-for-science>
- [28] Anthropic Institute. When AI builds itself. 2026 institutional account accessed 10 October 2026.
<https://www.anthropic.com/institute/recursive-self-improvement>
- [29] Google DeepMind AlphaEvolve team. AlphaEvolve How our Gemini-powered coding agent is scaling impact across fields. 7 May 2026 institutional account.
<https://deepmind.google/blog/alphaevolve-impact/>
- [30] Degraeve, J. and colleagues. Magnetic control of tokamak plasmas through deep reinforcement learning. Nature 2022.
<https://doi.org/10.1038/s41586-021-04301-9>
- [31] Szymanski, N. J. and colleagues. An autonomous laboratory for the accelerated synthesis of inorganic materials. Nature 2023; see January 2026 author correction.
<https://doi.org/10.1038/s41586-023-06734-w>
- [32] IBM. Quantum 2030 and the Starling roadmap. Roadmap accessed 10 October 2026; 2029 target.
<https://www.ibm.com/roadmaps/quantum/2030/>
- [33] Google Quantum AI and collaborators. Quantum error correction below the surface code threshold. Nature 2025; see published correction.
<https://doi.org/10.1038/s41586-024-08449-y>
- [34] Sivak, V. and colleagues. Reinforcement learning control of quantum error correction. Nature 8 July 2026.
<https://doi.org/10.1038/s41586-026-10759-2>
- [35] Abbas, A., Ambainis, A., Augustino, B. and colleagues. Challenges and opportunities in quantum optimization. Nature Reviews Physics 2024.
<https://doi.org/10.1038/s42254-024-00770-9>

Concept inputs

Robert J. Hathaway III supplied the design prompts concerning frozen antihydrogen pellets, an optional antihelium cushion, regulated antimatter flow and accelerated AI-enabled development. He also selected and integrated references on spacetime geometry, warp-drive concepts and faster-than-light research, including Lentz, Bobrick and Martire, and Fell and Heisenberg. His primary research is philanthropic semantic AGI/ASI and recursively self-improving intelligence, pursued since 1980 according to his research chronology. These statements identify his conceptual and architectural inputs; experimental validation, publication priority and patent novelty require separate evidence.

References for cryogenics and research transition

References 1-35 retain the original technical baseline and numbering. The additional sources below were reviewed on 18 October 2026.

[36] Hauser D and colleagues, including NASA and SpaceX coauthors. In-Space Cryogenic Propellant Transfer: Modeling and Validation of Two-Phase Flow Dynamics in Low-Gravity. 2025, NTRS 20250003325.

<https://ntrs.nasa.gov/citations/20250003325>

[37] SpaceX. Starship. Public vehicle propellants and Earth-orbit, Moon, Mars and beyond mission objectives.

<https://www.spacex.com/vehicles/starship/mars>

[38] NASA. Development of New Technologies for Modified Collins Cycle Expanders. 16 December 2020.

<https://www.nasa.gov/directorates/stmd/space-tech-research-grants/development-of-new-technologies-for-modified-collins-cycle-expanders/>

[39] NIST. Kelvin ITS-90. Hydrogen triple-point reference and low-temperature metrology.

<https://www.nist.gov/si-redefinition/kelvin/kelvin-its-90>

[40] NASA Science. Webb Cryocooler. Spacecraft refrigeration below 7 K.

<https://science.nasa.gov/mission/webb/cryocooler/>

[41] CERN. ALPHA cools antimatter using laser light for the first time. 31 March 2021.

<https://home.cern/alpha-cools-antimatter-using-laser-light-for-the-first-time/>

[42] NASA. Fluid Transfer. Public overview of in-space cryogenic transfer development.

<https://www.nasa.gov/fluid-transfer/>

[43] Hathaway RJ III. Antimatter and Spacetime Propulsion Research Proposal. Revision 1.1, 18 October 2026. Author-held companion funding and implementation proposal, 44 pages.

Revision history

10 October 2026. Original 27-page technical architecture and conditional roadmap.

18 October 2026. Added Sections 8.1 and 8.2 on Starship cryogenics and staged storage research, and Sections 22.1 and 22.2 on prospective collaboration and execution alongside continuing employment. Added references 36-43, expanded contribution attribution, restored the professional title and opening quotation, and produced an editable Word edition with native equations and a matching PDF. Original quantitative targets and technical section numbering are retained.

18 October 2026 navigation update. Added a linked contents page with page references and return links; pagination updated.