

OpenFusion Frontier MANIFESTO

Fusion Reactor LinealFRC6T Gatling (CERN-DOI-REGISTERED)

© <https://openfusionfoundation.org>

2026

Contents

FOUNDATIONAL MANIFESTO: OpenFusionFrontier	1
1. The 4 Pillars: Our Vision	2
2. The 9 Practical Goals (Fewer promises, more pragmatism)	3
3. The Call: Get your hands dirty	4
4. Governance: How we organize ourselves without falling into chaos	4
5. From PDF to Steel (Money and the real world)	5
6. The AAA Method (Architect-Assistant-Auditors)	5
7. By way of closing	6
ANNEX E.1. Shield against skepticism (Technical and Financial Justification)	7
1. The Linear Trap and the Kinematic Limit (30 ms and 80 m)	7
2. The Gas Wall and Vacuum Saturation (The 7.5 second limit)	8
3. The Algorithmic Mirage and the Supply Scale (Zero <i>Unobtainium</i>)	8
4. The Open Source Defensive Moat (Financial Scale)	9
5. The Legal Vacuum (Liability and Governance)	9
ANNEX E.2. Map of Open Problems (A possible “guide” for “thesis advisors”)	10
PART I: THE BALLISTIC LINE (Phase by Phase)	10
PART II: PLANT SUBSYSTEMS (Macro-Engineering)	11
The Whitepaper: The true research quarry	12

FOUNDATIONAL MANIFESTO: OpenFusionFrontier

- **Project:** Fusion Plant “RIF-v101” (*Resilient Industrial Fusion*)
- **Architecture:** Linear Fusion Reactor FRC6T-Gatling(30+10)-Hybrid
- **Author:** (c) Leopoldo Cano Guardiola - leo@ulogicmind.ai - (DOI-Registered) <https://doi.org/10.5281/zenodo.22070202>
- **OpenFusion Frontier** <https://openfusionfoundation.org>
- **Document Version:** v101
- **Open License:** CERN OHL-S (Strongly Reciprocal)

Making nuclear fusion commercial and viable is probably the biggest technical headache we have ever faced. For decades, thousands of brilliant researchers have pushed the boundaries of physics to bring us closer to the goal. Their work is heroic. But we believe that a problem of this brutal scale is not going to be solved behind closed doors.

OpenFusionFrontier is born from a very simple idea: hyper-complex problems require collaborative intelligence. If open source was able to build the software infrastructure that sustains the world today (Linux, the internet, servers), it is time to apply that same logic to heavy metalwork and thermodynamics.

We are not building the design of this reactor to hide it in a safe. We bet on boring linear geometries, on catalog parts, and on using Artificial Intelligence as a tireless excavator. We release it under open source (CERN OHL-S) because we want humanity to own the blueprints.

A note before starting:

- *The design of this reactor (RIFv101) is not the “invention” of a solitary individual. It is, in reality, a massive puzzle-assembly exercise. The pieces—empirical plasma physics, cryogenics, materials science—already existed because they are pure knowledge generated over decades by countless researchers, laboratories, and organizations worldwide.*
- *My role in this project has been that of assembler (plus some contributions of my own to direct the design toward a specific geometry and operating mode). Artificial Intelligence has provided me with an unfathomable brute force of calculation, but AI completely lacks lateral thinking. It does not know how to invent new paradigms. My contribution has been to provide the “natural intelligence,” the intuition necessary to pull the design out of the traditional loop and find the way for those pieces to fit into a linear, viable architecture.*
- *Tracing this thermodynamic map is a milestone, but I am pragmatic. Coordinating the evolution of this system and scaling a collaborative network globally demands a bandwidth that exceeds any single individual. A hyper-complex project led by a single person is born dead. Therefore, this manifesto is also a call for management succession. OpenFusionFrontier seeks professionals who not only audit equations but who have the experience to coordinate teams, structure the technical committee (the “Kernel”), and lead the Foundation. The v101 design is the initial base camp. If you have the capacity to organize the technical expedition and guide it toward steel, the doors are open for you to take command. The person signing this document is only the initial architect; the true protagonist of this story will be the network that builds it.*

1. The 4 Pillars: Our Vision

- **(WHAT) An open-source fusion reactor** The heart of this project is the RIF-Plant and the design of the Gatling (30+10) linear reactor. We are not looking to protect a secret, but to test a design. It is a linear, modular, and crude machine, licensed under CERN OHL-S. The goal is to deliver 200 net MWe to the grid (scalable to 271 MWe at peaks) using engineering that can be built today, without

waiting for magic materials that do not exist.

- **(WHO) People who want to build real things** We want to pull talent out of the drawers. If you are a student, researcher, or engineer, you have the opportunity to step out of pure theory and get your hands dirty in a project with real impact. Here your talent materializes in solving alloy fatigue, optimizing a vacuum pump, or calculating the angle of a magnetic field that will end up on a manufacturing blueprint.
- **(HOW) AAA Methodology: Humans and AI working together** Taming fusion demands a different way of working. We use the AAA framework: the **Architect** (the person who provides common sense and intuition), the **Assistant** (the AI that processes massive data and balances equations without getting tired), and the **Auditors** (the community of peers that reviews to ensure the machine isn't hallucinating impossible things). Together, we avoid analysis paralysis.
- **(WHY) Sovereignty and common sense** Energy dictates what a society can or cannot do. By opening these blueprints, atomic knowledge ceases to be the monopoly of a few countries. We want any consortium or nation to have the necessary map to assemble its own plant, lower the price of energy, and stop depending on burning things to keep the lights on.

2. The 9 Practical Goals (Fewer promises, more pragmatism)

1. **Pragmatic Viability:** We flee from the static donut (Tokamak). We designed a straight tube where the plasma flies. It is much easier to deal with heat and pressures if the thermonuclear fire is in motion.
2. **"Gatling" System (Pure Resilience):** We cannot afford blackouts. That's why we don't build one giant tube, but 40 independent tubes. 30 fire sequentially and 10 rest regenerating their vacuum. It is the principle of a multi-cylinder engine block applied to nuclear fusion.
3. **Use everything:** We extract energy twofold. We brake the plasma magnetically to generate direct electricity (DEC) and we use the molten salt walls (FLiBe) to catch the immense heat from neutrons and drive commercial CO2 turbines.
4. **Perfect the AAA tool:** Building the reactor is important; fine-tuning the working method between humans and AI models to solve deadlocked problems is vital. It is our new engine of innovation.
5. **Executable Knowledge ("Knowledge-as-Code"):** Enough with dead PDFs. We document every advance with equations and parameters that AI can read, understand, and integrate directly into its next simulation.
6. **Shipyard modularity:** The plant is divided into cartridges. If a module breaks, a crane pulls it out and puts another one in hot. This allows different groups to work simultaneously on isolated parts of the reactor without stepping on each other's toes.
7. **Only catalog technology (COTS):** If a component cannot be bought, forged, or adapted using current petrochemical, naval, or aerospace industry, the design is discarded. Zero science fiction.

8. **Documenting failures:** Sometimes, knowing where *not* to go is more useful than getting it right. We rigorously record every failed calculation and every “dead end” so no one in the community wastes time tripping over the same stone again.
9. **The Empirical Digital Twin:** Before pouring a single gram of concrete, the design has to add up in simulators fed with real, boring industrial data. That is the only way a regulatory agency will look us in the face.

3. The Call: Get your hands dirty

This is a call to arms. To the veteran engineer bored in his office, to the researcher suffocated by bureaucracy, and to the professional who wants to make a real impact. History is not watched from the sidelines.

But, very specifically, this is an open invitation to the **academic world**.

Every year, thousands of Bachelor’s Theses, Master’s theses, and Doctorates are written that end up gathering dust in a digital archive that almost no one reads. Let’s change that. We invite students and their thesis advisors to use this architecture as their test bed. If you are going to spend months calculating a material’s thermal fatigue or simulating the fluid dynamics of a molten salt, do it on a live problem. Let your thesis serve to optimize a nozzle’s magnetic field or the laminar flow of FLiBe. Give that academic effort a purpose that ends up turned into steel (and add your name to that of the engineers and physicists who ultimately end up solving the real construction problems).

Entering this project is not a walk in the park. It will force you to review physics, to tame Artificial Intelligence so it doesn’t invent data, and to defend your numbers before a community that doesn’t forgive a calculation error. But it is the best mental gym in the world.

Join the network. A small adjustment in a valve’s fluid dynamics or the thermodynamic correction of a coolant are real victories. Your name will remain in the machine’s repository. Here you have an open blueprint, a pile of technical problems waiting for someone with common sense, and a purpose worthy of your talent.

4. Governance: How we organize ourselves without falling into chaos

So that thousands of people can design a nuclear power plant without the project becoming a chaos of endless meetings, we copy the open-source software structure:

- **The Foundation (OpenFusionFrontier):** It is the legal and institutional umbrella (openfusion-foundation.org). It pays for servers, protects the open license, and talks to the adults (governments, regulators, and investors).
- **The “Kernel” (Technical Committee):** A core group of maintainers who assemble the pieces contributed by the community and ensure the physics keeps adding up.

We do not vote on physics by majority; we use mathematical meritocracy. No one gets to the Kernel on charisma. You get there if your equations balance, if your knowledge nodes work, and if your solutions survive the Auditors' attack. If you are right, your code gets in. If not, it is discarded.

But let's be clear: this model is not carved in stone. It is a germinal idea, a moving organism. Coordinating the intelligence of thousands of people on a global scale is, in itself, a hyper-complex system. Just as we iterate the reactor's thermodynamics, the Foundation itself will constantly study and refine these self-organization rules. The goal is to discover how to channel collective talent without bureaucracy suffocating it. Building the machine is the technical challenge; building the human network capable of designing it is our true experiment.

5. From PDF to Steel (Money and the real world)

It's one thing to be an engineer, and another to be naive: We know that writing equations is free, but pouring thousands of tons of concrete and buying superconducting magnets costs billions.

We have nothing against private capital or energy companies; in fact, we need them to enter. The deal is simple: OpenFusionFrontier eliminates the immense risk of the "blank page". We hand them the architecture, the baseline physics, and the initial blueprints. In exchange for using this design for free to make money selling energy, the license requires that empirical data from the real plant flow back openly to the network. They make a monumental business; humanity inherits the technology.

6. The AAA Method (Architect-Assistant-Auditors)

If you want to survive hyper-complexity while working with an AI, the workflow rules must be the following (**AAA Innovation Method (Architect-Assistant-Auditors)**):

- **The architecture is imposed by the human:** AI is programmed to propose exclusively what already exists and what is statistically popular. The human architect has the inescapable obligation to provide intuition and make the lateral leap.
- **AI is the heavy excavator:** When you face specific problems, the machine will dive into the infinity of human knowledge. It will tell you exactly which *doped alloy solves your structural friction*. It invents nothing new, but locates what already exists at superhuman speed.
- **The Red-Team algorithm:** If you ask the AI to validate your design, it will complacently agree with you and tell you that you are a genius. That is the direct path to failure. You must order it to look for vulnerabilities and attempt to destroy your architecture without mercy. Only after that adversarial attack should you demand that it propose solutions within a strict limit of three iterations.
- **Drift versus Convergence:** It is natural to think that iterating infinite times with the AI will progressively refine the design. This is a dangerous lie. Without anchors, the language model suffers a progressive "drift" toward absurdities of all kinds. To force the machine to converge, you have to nail **Inviolable Axioms** to the wall and document a strict list of **Dead Ends**. These are failed

solutions that the machine is absolutely forbidden to propose again.

- **Understand the music, ignore the lyrics:** Throughout development, the AI will spit out hundreds of equations at you. Trying to memorize them or fighting hand-to-hand with every mathematical detail is useless. What really matters is grasping the physical intuition and the structural mechanics pulsing behind the problem. You are the conductor of the orchestra and your only job is to find the perfect melody (“do not try to beat a calculator at multiplication”).

These are **the three legs of the AAA method (Architect-Assistant-Auditors)**:

1. **Intuition (The Architect):** AI looks for the average. The human looks for the shortcut or the geometric madness (e.g.: “if neutrons go in a straight line, let’s bend the tube”). The human sets the course.
2. **Brute Force (The Assistant):** The AI receives the idea and begins calculating thousands of iterations, tolerances, and metal catalogs. It must do this under a strict rule: it is forbidden to propose solutions that are in our “Dead End Cartography” (the log of things we already know explode). We do not ask the AI to discover new physics. We use it for massive parametric sweeps, crossing materials catalogs, and geometric optimization. The human dictates the architecture and the machine does the computational dirty work (in physics, engineering, and materials). As the research advances we must set “immutable axioms” that cannot be bypassed under any circumstances, and also document “design guidelines” and “design prohibitions” that guide the AI in successive iterations (which it sometimes respects and sometimes does not).
3. **Common Sense (The Auditors):** The community looks at the AI’s result. Can this be manufactured or is it an impossible shape to machine? If the part survives human critique, it is added to the blueprints. If it fails, it goes to the Dead End.

This is how we get the machine to empower us without letting ourselves be dragged down by its hallucinations.

7. By way of closing

Designing a 200 MW reactor is fascinating, but what we are really building here is a human proof of concept. We are proving that, faced with hyper-complexity, locking fifty geniuses in a secret room no longer works. The only way out is to open the board, pool global talent, and use AI as leverage.

If we can design a Sun on Earth by organizing ourselves this way, imagine what other great historical problems we could get out of the way. The future is not patented. It is built.

Welcome to the frontier.

ANNEX E.1. Shield against skepticism (Technical and Financial Justification)

Subjecting a project of this scale to extreme stress is our main rule. Any physicist, engineer, or investor reading our proposal will attack us (rightfully so) citing the historical bottlenecks of nuclear fusion. This annex functions as our thermal shield. This is how the RIF v101 architecture deactivates traditional vulnerabilities, uniting industrial pragmatism, heavy metalwork, and unshakeable mathematical justification (this is just a summary: full details are in the Whitepaper).

§ 1. The Linear Trap and the Kinematic Limit (30 ms and 80 m)

The Objection: *“The plasma is going to escape out the ends. Furthermore, in linear geometries, the gas destabilizes in fractions of a second, preventing thermal viability from being reached and causing the reaction to collapse before generating net energy.”*

Our Solution: That is true. The safe flight time before the FRC (*Field-Reversed Configuration*) plasma becomes irretrievably destabilized is established at an operating regime of merely **30 milliseconds**. This is not a hypothesis; it is an empirical threshold exhaustively validated by current industry (TAE Technologies’ experimental reactors). We do not fight against that limit; we accept it and build the machine around it.

To guarantee the absolute cohesion of the plasma during its flight, the RIF architecture discards active electronic sensors (which would be destroyed by radiation in seconds) and integrates an ecosystem of dynamic and passive stabilization:

- **Neutral Beam Injection (NBI) at rest:** Tangential injection of fast ions in the initial chamber (strictly *before* the plasma starts flying) equips the gas with a kinetic armor. It acts as a pre-flight gyroscope, suppressing macroscopic tilt instability without generating subsequent aerodynamic friction.
- **Active Edge Biasing:** We use static electrodes to control the rotation of the outer boundary layer. This electrical shear effect brakes centrifugal inertia, preventing the “smoke ring” from tearing itself apart as it spins.
- **Passive Magnetic Bumper (W-Re Tiles):** If the projectile deviates a single millimeter from its central axis during flight, it would touch the wall and melt the reactor. To prevent this, we embed flat Tungsten-Rhenium (W-Re) tiles in the bottom of the steel channels, bathed by the liquid salt. If the plasma tries to drift to one side, its own magnetic field naturally induces eddy currents in these ceramics. This unleashes an automatic repulsive force (Lenz’s Law). The plasma repels itself and returns to the center. The course self-corrects with no moving parts or electronics.

On the other hand, our combustion chamber imposes a non-negotiable length of **80 meters**. This immense extension dilutes the neutronic radiation flux, avoiding localized ablation of the wall. To cross that steel chassis before the plasma’s useful life expires (the 30 ms), the projectile demands an unalterable

ballistic cruising velocity (v):

$$v = \frac{d}{t} = \frac{80 \text{ m}}{0.030 \text{ s}} \approx 2666 \text{ m/s}$$

The plasma crosses hell at Mach 7.8, delivering its heat to the salt mantle (FLiBe) and leaving the combustion chamber exactly before instability destroys it. Simple math, pure kinematics.

§ 2. The Gas Wall and Vacuum Saturation (The 7.5 second limit)

The Objection: *“You cannot empty the tube of gases in time. Evacuating the residual gas between supersonic pulses requires an impossible suction volume. Your reactor will asphyxiate.”*

Our Solution: Aerodynamically, emptying a tube of this scale to achieve a perfect ultra-high vacuum cannot be done in milliseconds. The physics of Free Molecular Flow imposes an incompressible barrier: the residual gas takes a minimum of **7.5 seconds** to physically travel to the sinks. In nuclear physics, 7.5 seconds between shots is an eternity that ruins commercial viability.

To manage to clean the chamber in that time, we avoid brute force suction and use distributed chemical pumping. We line the 427 m² of the interior wall with an NEG (Zirconium-Vanadium-Iron) film that “swallows” hydrogen isotopes instantly on pure contact, supported by massive perimeter turbopumps for the inert helium ashes.

To validate the 7.5 second limit, we assume the full true volume of the pipe (V) is **602,320 Liters**. The initial exhaust pressure is $P_1 = 10^{-3}$ Torr and the target pressure required to fire again is $P_2 = 10^{-8}$ Torr. We calculate the required Effective Pumping Speed (S_{eff}):

$$S_{eff} = \frac{V}{t} \ln \left(\frac{P_1}{P_2} \right)$$

$$S_{eff} = \frac{602320}{7.5} \ln(10^5) \approx \mathbf{924,595 \text{ L/s}}$$

The immense **50 m²** surface area of NEG panels mounted in the terminal plenums delivers a passive rate of **1,000,000 L/s**, comfortably validating the equation and covering the safety margin.

Additionally, to prevent the electrical grid from suffering a blackout every 7.5 seconds, the architectural solution is **the Gatling System**. We have placed 40 independent tubes in a matrix. While one tube takes its 7.5 seconds to swallow ashes, dozens of parallel tubes are firing sequentially and rotationally. This is how we mask the thermodynamic limit and deliver continuous base load.

§ 3. The Algorithmic Mirage and the Supply Scale (Zero Unobtainium)

The Objection: *“Your reactor will require materials that do not yet exist in laboratories (Unobtainium). And even if they did exist, there is no industrial capacity to manufacture them at that scale.”*

Our Solution: On the contrary. We have de facto forbidden the use of experimental metals and paper physics. We impose the strict principle of COTS (*Commercial Off-The-Shelf*) catalog engineering. We use nuclear-grade steel (Eurofer), FLiBe salts for the thermal mantle, and standard Tungsten alloys for the protection tiles.

For the immense 6.0 Tesla magnets wrapping the barrel, we flee from expensive and scarce experimental high-temperature superconductors (REBCO). We demand exclusively conventional Niobium-Titanium (NbTi) braided cables, bathed in liquid helium at 4 K. Will it be a logistical hell to get so much commercial cable for 40 80-meter tubes? Yes. Building this reactor will force a drastic scaling of global manufacturing in the petrochemical, naval, and MRI industries. But it is a logistical and assembly challenge, not having to invent new materials physics.

Giving up science fiction doesn't mean it's easy. Integrating extreme vacuum, 4 K cryogenics, and 600 °C liquid metals in the same industrial warehouse will be one of the biggest engineering nightmares of our century. But it is a nightmare of manufacturing, tolerances, and industrial logistics, not of unattainable theoretical physics.

§ 4. The Open Source Defensive Moat (Financial Scale)

The Objection: *"No one is going to put billions into an open source infrastructure. Releasing intellectual property scares off venture capital because any competitor can use the blueprints for free the next day."*

Our Solution: Releasing the design is not financial naivety; it is pragmatism. The defensive moat of this project does not lie in retaining a theoretical patent, but in the brutal capacity for industrial execution. The real business model (the multi-billion dollar monopoly) lies in building the plant, controlling the steel and magnet supply chain, electromechanical integration, and the final sale of energy.

Money in large infrastructure is made by operating the facility and collecting the light bill. Freeing the core radically lowers the initial R&D risk for consortiums. We hand them the thermodynamic design of the core; they assume the operational risk (set at a strict limit of **\$4.5 billion USD CAPEX**) and compete to build the best chassis. They do monumental business; humanity keeps the base technology.

§ 5. The Legal Vacuum (Liability and Governance)

The Objection: *"An internet project operated by volunteer contributors cannot obtain nuclear licenses or assume billion-dollar civil liabilities in the event of a radioactive accident."*

Our Solution: OpenFusionFrontier is never going to operate a physical plant. A reactor is not a GitHub repository; regulatory agencies require legal entities, concrete, and massive insurance guarantees. Therefore, we establish an absolute firewall between the design of the machine and its actual operation.

The core team (Kernel) designs, evaluates, audits equations, and certifies parameters under strict mathematical meritocracy on the global network. We deliver the initial map. But commercial construction and exploitation will always fall upon legal and regulated energy operators (governments or corporations). They assume land ownership, capital risk, and the criminal and civil liability of piloting the ship under

international laws. We only fine-tune the compass so they don't crash.

PS: This is the state of the art of the RIF v101 architecture. It is not a perfect map, it is the first base camp. If you see a thermal flaw, a magnetic miscalculation, or a better geometric solution, the code is yours. Improve it."

ANNEX E.2. Map of Open Problems (A possible “guide” for “thesis advisors”)

Saying that the v101 thermodynamic design is mapped out does not mean the machine is finished. It means we have managed to corner the paradoxes of theoretical physics until they become pure engineering, fluid thermodynamics, and materials strength problems.

This is our transparent log of “open paths.” It is a catalog of technical nightmares and industrial bottlenecks. We invite universities, thesis advisors, engineers, and developers to choose one of these cracks, simulate it, resolve it, and contribute their knowledge node to the OpenFusionFrontier *Kernel*.

If you are looking for a topic with historical impact for your next research, choose your battle:

§ PART I: THE BALLISTIC LINE (Phase by Phase)

The plasma's behavior along the stretched “S” tube. Each phase of the projectile presents its own thermodynamic demons.

- **Phase 1 (Birth and Injection): Aerodynamic drag.** We inject Neutral Beams (NBI) at rest to equip the FRC plasma with a gyroscopic shell. We need to mathematically model whether that sub-population of high-energy ions generates cross friction or Kelvin-Helmholtz instabilities once the ballistic projectile begins its massive acceleration at an angle of **21.28°**.
- **Phase 2 (The First Turn): The stress of the blind wall.** The insertion curve creates a sacrificial Tungsten wall so that backward-traveling neutrons do not fry the injector. We need to calculate the exact neutronic erosion of that elbow and the magnetic fluid dynamics of the plasma as it takes the curve without crashing into the outer wall.
- **Phases 3 and 4 (The Central Barrel): DPA, shear, and passive bumper.** In this immense 80-meter straight, the plasma travels at Mach 7.8 for 30 milliseconds. We have three open fronts:
 1. *Parametric shear:* We require simulations of Edge Biasing (electrodes) to calculate the exact electron leakage rate crossing field lines while we force the rotation of the plasma's outer layer.
 2. *DPA fatigue:* Probabilistic models of Displacements Per Atom damage. How many shots can Eurofer steel withstand under a 14.1 MeV fast flux before becoming brittle as glass?
 3. *Eddy current heating:* If the plasma deviates, it induces braking currents in the Tungsten-Rhenium

tiles. We need the exact thermal dissipation equation: how many megawatts of local heat will the tile absorb while magnetically repelling the projectile?

- **Phase 5 (The Nozzle): Electromagnetic induction (DEC).** We drop the magnetic field from 6.0 to 0.75 Teslas. The plasma expands and grazes a copper cage to brake and generate gigawatts of electricity. We seek the optimum geometric topology and braiding pattern of those conductive coils to maximize harvesting without entering destructive electromagnetic resonance.
 - **Phases 6 and 7 (Second Turn and Gas Target): The inelastic collision.** The surplus gas turns at **30.6°** and we shoot a dense curtain of cold neutral Helium at it. We need massive ion-atom interaction simulations to ensure the gas steals electrons quickly, neutralizes the plasma, and decouples it from the magnetic field without reflecting a sonic shockwave backward (choking the barrel).
 - **Phases 8 and 9 (Sinks and UHV): The 7.5-second limit.** The Zirconium-Vanadium-Iron film absorbs cold hydrogen instantly. But during the rest phase, the NEG wall must be baked at **550 °C** to extract the gas. We need to calculate how many thousands of thermal expansion and contraction cycles this nanometric coating can withstand before flaking and falling into the chamber.
-

§ PART II: PLANT SUBSYSTEMS (Macro-Engineering)

The peripheral infrastructure that keeps the barrel alive. This is where heavy industry must prove its capability.

Subsystem A: Liquid Mantle Thermohydraulics (FLiBe)

- *Hartmann's blockage:* Liquid salt flows parallel to the 6.0 Tesla field with a grazing tilt of 8 degrees. We need massive 3D CFD (Computational Fluid Dynamics) simulations validating the use of macro-fins and W-Re tiles to comfortably assimilate the master flow of **3920.59 kg/s** in the Peak Phase.
- *Surface tension vs. Bubbling:* To extract freshly created Tritium, we inject a counter-current hurricane of Helium microbubbles in tall vacuum towers. Heavy fluid dynamics characterization is required: what is the exact bubble diameter and pressure to break the salt's stochastic surface tension?
- *Beryllium kinetics:* To neutralize the remaining hydrofluoric acid, we divert salt to Beryllium pebble beds. We need the exact stoichiometric equation of passive dissolution as a function of flow rate and temperature.

Subsystem B: Superconductivity, Cryogenics, and Neutronics

- *The hyper-compact shield:* Starfire is less than a meter from Niobium-Titanium coils at **4 K** (-269 °C). We require advanced neutronic modeling (Monte Carlo / MCNP) to optimize the millimeter thickness of the Boron/Water shield. A single escaped neutron reaching the magnet will cause a *quench* (thermal collapse).

- *Lorentz structural stress:* The 6.0 Tesla magnets attract and repel each other with forces of thousands of tons. We seek the optimal civil design of the titanium tie-rods holding the magnetic backbone to prevent chassis buckling.
- *Civil parasitic consumption:* We need to optimize the thermodynamic cascade of the immense Helium liquefaction compressors. The cryogenic electrical toll must be minimized so as not to ruin the profit margin to the electrical grid.

Subsystem C: Civil Architecture and Kinematic Chain

- *Friction in dry gravitational bearings:* The central steel tube grows linearly up to **72 cm** through expansion. It rests freely on ceramic rollers seated in “V” rails. We require extreme tribology studies at **600 °C** in a secondary radiation environment to ensure the bimodal maneuver slides the cartridge cleanly into the **8 m** clear-width diagonal technical corridor.
- *Hydrodynamics of the underground manifold (The Trench):* Forty tubes pour liquid salt by gravity into the building’s backbone. We need open-channel flow models to prevent cavitation, acoustic resonances, and thermal waves when unifying the flow.

Subsystem D: Balance of Plant (BoP) and Commercial Cycle

- *Isochoric buffering in Chlorides (TES):* The Gatling reactor emits violent beats at 4 Hz. We dump that asymmetric heat into gigantic passive pools of cheap chloride salts. We seek the exact volumetric topology of these tanks so the heat flattens out and flows completely continuously and docilely toward the turbines.
- *Palladium-Silver membrane filtration:* Separating Tritium from residual Helium demands pressurized atomic sieves at **400 °C**. We need to optimize the alloy thickness to maximize flow rate without the compressor pressure bursting the filter.

§ The Whitepaper: The true research quarry

This list is just the surface map. If you are a researcher, student, or thesis advisor, you should know that these bullet points are barely an executive summary of the reactor’s technical gaps. The **technical Whitepaper** of the RIF v101 architecture contains the complete mathematical, geometric, and thermodynamic topography.

There you will find the exact kinematic regimes, required step-by-step vacuum pressures, magnetic field tolerances, and the physical behavior of molten salts. We invite you to dissect that main document. Every engineering assumption, every parametric limit, and every equation we have documented is, in itself, a seed for dozens of new research lines, Master’s theses, or Doctoral dissertations. Use the Whitepaper not as a static manual, but as your primary quarry to extract, analyze, and extend this guide of open problems to their ultimate consequences.