Economic Fusion Moonshot Initiative

CLIMATE MOONSHOT PORTFOLIO

Billions to save trillions

Proposed $1 Billion / Five-Year R&D Program

STELLARATOR • LIQUID-METAL WALL • NEXT GENERATION MAGNETS

DESIGN FOR 100 GWe ANNUAL PRODUCTION

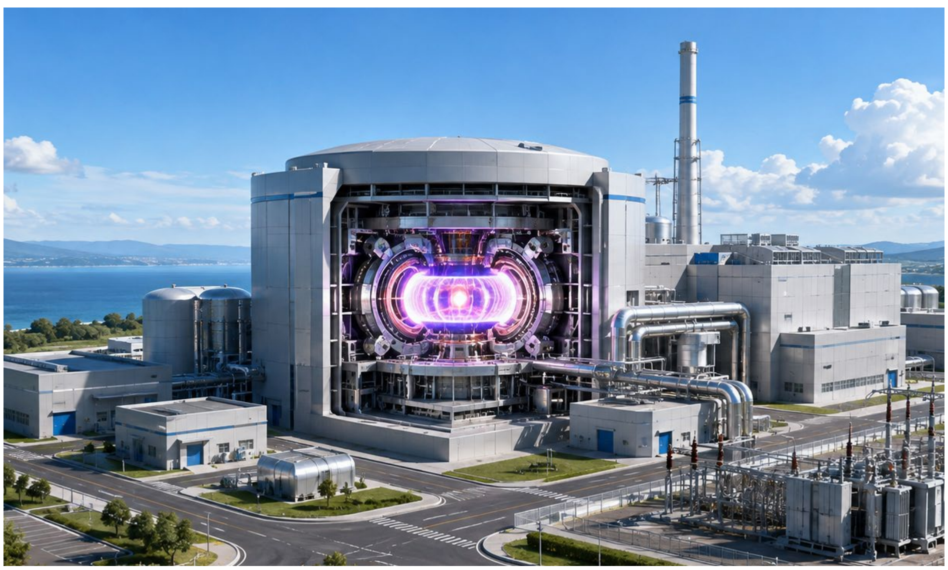

By Glenn Weinreb

Developed with assistance from ChatGPT

OPEN SOURCE

Executive Summary

Purpose. This proposed $1 billion, five-year Economic Fusion Moonshot Initiative is designed to determine whether a mature stellarator-based fusion power plant could produce net electricity in China for less than a new coal-fired power plant, and whether a fusion retrofit at an existing coal plant could beat the site-specific marginal cost of continued coal operation. The program is an engineering, experimental, manufacturing, and economic validation effort; it does not promise a commercial fusion plant within five years.

Economic hypothesis. A Volpe P1-like stellarator architecture combining a liquid-metal plasma-facing system, high-temperature-superconducting magnets, low recirculating power, high technical availability, China-localized manufacturing, and aggressive factory production may be able to reach the program targets: below $43/MWh-electric for greenfield power, preferably $40/MWh or less, and roughly $2,000/kW-net or less for mature greenfield CAPEX. Those values are targets to test, not assumed outcomes.

Five-year work. The portfolio allocates $900 million across ten planned technical areas plus a centrally controlled $100 million reserve. The ten areas cover systems engineering and a bottom-up China cost model; stellarator plasma validation; HTS magnets and low-cost manufacturing; liquid-metal wall, blanket, shielding, and heat extraction; tritium breeding and processing; net power, power conversion, and coal-retrofit integration; RAMI and long-life components; Chinese factory production and construction engineering; full-scale integrated sector prototypes and endurance testing; and independent verification, safety/licensing, cost audit, and challenge testing. Major hardware commitments follow explicit go/no-go gates, and reserve funds are released only for documented risk or additional evidence needs after independent review.

Success criterion. By Year 5, the program should deliver one or two controlled reference designs supported by complete power and tritium balances, audited CAPEX and LCOE models, a bottom-up RAMI model, construction and factory-production plans, full-scale or near-full-scale prototypes of the highest-risk systems, and an integrated engineering-sector demonstration. A credible YES would show how mature Chinese fusion could beat coal; a credible NO would identify which technical or economic assumptions fail and by how much. Either result would be decision-grade evidence for the next funding step.

Climate Moonshot Portfolio. We define “Climate Moonshot” as conducting R&D to the extent required to solve the climate problem. This would require a portfolio of R&D initiatives, an example of which is at www.APlanToSaveThePlanet.org. The document you are reading is one initiative within that portfolio. It uses a Common Framework that establishes shared rules for governance, procurement, data rights, independent review, and oversight. All documents within this portfolio are open-source and may be freely viewed, used, and modified.

Program at a Glance

The program is an economic product-development and validation effort, not a promise to build a complete commercial fusion plant in five years. The table below states the decision problem, the controlled reference cases, the main technical risks, and the evidence required before larger commitments.

Item	Economic Fusion Moonshot
Funding / duration	$1.0B over five years: $900M planned technical work + $100M sponsor-controlled reserve
Mission	Determine whether mature Chinese stellarator fusion can produce net electricity below the relevant coal alternative.
Baseline architecture	P1-like stellarator; liquid-metal plasma-facing system; HTS magnets; D-T baseline unless evidence supports a better economic alternative.
Controlled references	China greenfield fusion plant + China coal-plant retrofit.
Greenfield economic test	< $43/MWh-net; preferred <= $40/MWh; mature CAPEX target <= $2,000/kW-net.
Retrofit economic test	Below the site-specific marginal cost of continued coal operation.
Plant targets	Approximately 1 GWe net; >=95% technical availability target supported by bottom-up RAMI; <=36-month mature construction target.
Principal risks	Plasma operating point; magnet cost/manufacturing; liquid-wall MHD/materials; tritium closure; parasitic power; RAMI; integrated-sector performance.
Industrialization	Credible manufacturing and construction path first to tens, then toward approximately 100 GWe/year.
Evidence standard	Assumption -> simulation -> engineering calculation -> vendor quote -> component test -> integrated test -> operating data.
Year-5 outcome	Independent YES / NO / targeted redesign decision, supported by controlled designs, measured hardware evidence, audited CAPEX/LCOE, and quantified uncertainty.

Table of Contents

Economic Fusion Moonshot

The proposed $1B Economic Fusion Moonshot (EFM) program intends to design a fusion machine that produces electrical power at a cost less than fossil fuel-based power. This is commonly referred to as "economic fusion."

Our program builds on the work by Dr. Volpe who has published a design and cost model for a machine that meets our objectives. His design uses a stellarator architecture with liquid metal wall and next generation magnets.

China has recently added coal-fired generating capacity at very large scale. The EFM deployment case therefore uses new Chinese coal generation as the primary greenfield competitor and existing Chinese coal generation as the retrofit comparator. The objective is not to assume future coal additions, but to test whether fusion could become the lower-cost choice wherever new firm generation or replacement heat is required. If the economic target is achieved through mature high-volume manufacturing in China, the same manufacturing approach could then be adapted for deployment elsewhere.

For a list of relevant fusion papers, see the last few chapters of this document; and for a video summary of fusion moonshot, see:

```
Economic Fusion Moonshot [#14]
https://www.youtube.com/watch?v=CvZzGHSugy4
```

Primary Design Targets

Metric	Primary requirement / target
China greenfield LCOE	< $43/MWh-electric; preferred <= $40/MWh; stretch <= $35/MWh
China coal-retrofit LCOE	Below the site-specific marginal cost of continued coal operation; screening target <= $35/MWh-electric
China greenfield CAPEX	<= $2,000/kW-net (about $2.0B at 1 GWe net, ≤ $2/Watt)
China coal-retrofit CAPEX	No arbitrary hard ceiling initially; must be low enough to pass the retrofit LCOE test and must be reported in total dollars and $/kW-net
Commercial net output	Approximately 1,000 MWe net for the reference design
Technical plant availability	>= 95% design target, supported by bottom-up RAMI evidence
Mature construction duration	<= 36 months from major site-construction authorization to commercial operation; preferred approximately 24 months for major fusion-site work

Terminology

“Greenfield” means construction of an entirely new fusion power plant on a new or undeveloped site rather than reuse of an existing power plant. “Retrofit” refers to building a fusion machine at an existing fossil-fuel power plant and reusing suitable electrical and thermal equipment. It follows that “coal-retrofit” refers to building a fusion machine at an existing coal-fired power plant, and “China coal-retrofit” refers to doing this in China.

“Mature n^{th}-of-a-kind” refers to costs after many machines have been built and cost-reduction has occurred.

“MWe” refers to electrical power output capacity in units of mega-watts. “MWh-electric” refers to electrical energy in units of mega-watt-hours. And “MWh-thermal” refers to thermal energy (i.e. heat) in units of mega-watt-hours.

“LCOE” refers to the cost of electricity due to daily operations and original construction costs, and is often described in units of dollars-per-megawatt hour (i.e. $/MWh-electric levelized cost of electricity). “CAPEX” refers to the cost of construction and is often described in units of dollars-per-watt of generation capacity (i.e. $/Watt capital expenditure). “Technical availability” refers to the percentage of time a plant is available for operation over a year, and “Capacity Factor” refers to the percentage of actual output over a year, relative to running at maximum capacity 100% of the time.

Program Deliverables

The $1B program deliverable is:

> A validated, China-ready engineering and economic design package for an Economic Fusion Power Plant, supported by enough full-scale component and integrated-system testing to determine credibly whether it can meet the specified cost targets.

More concretely, at the end of five years there should be one or perhaps two down-selected designs with:

- a controlled 3-D plant design;
- complete mass, energy, heat, electrical, and tritium balances;
- full net-power calculation;
- bottom-up China BOM and supplier-based cost estimate;
- separate greenfield and coal-retrofit designs;
- audited CAPEX and LCOE models;
- RAMI/availability model;
- construction sequence and schedule;
- factory-production plan;
- maintenance and replacement plan;
- tritium inventory and breeding model;
- safety/licensing concept;
- experimentally validated liquid-wall, magnet, heat-transfer, and other critical subsystem data;
- full-scale or near-full-scale prototypes of the highest-risk components;
- an integrated sector or engineering mockup that demonstrates the interfaces among major systems;
- quantified uncertainty ranges for every major cost and performance assumption.

The cost model should explicitly show the three cases: **Volpe reference, China greenfield, and China coal retrofit**, using consistent accounting boundaries.

The program succeeds if, after spending $1B, it can say one of two things with strong evidence:

> **YES:** Here is a technically credible P1-like plant, here is the hardware evidence, here are Chinese supplier costs, and here is why mature greenfield electricity can plausibly reach, for example, $38–42/MWh and/or a retrofit can beat continued coal operation.

or:

> **NO:** The design cannot reach economic fusion because, for example, realistic parasitic load is too high, the liquid wall fails, tritium breeding requires too much blanket thickness, availability cannot exceed 85%, or the plant costs $3,500/kW even after Chinese localization.

That second result would still be an extremely valuable $1B outcome because it prevents tens of billions from being spent pursuing an architecture that cannot meet the economic objective.

The **ultimate deliverable** in one sentence is:

> A decision-grade, experimentally supported answer to whether a mature Chinese stellarator fusion plant can beat coal—and, if yes, a detailed engineering package showing exactly how to build it; if no, a quantified explanation of why not.

Critical Questions

Our $1B fusion moonshot hopes to answer the following question:

> Can a mature stellarator-based fusion plant deliver net electricity in China for less than new coal, and can a retrofit at an existing coal plant beat the marginal cost of continuing to burn coal?

More specifically, can a greenfield (i.e. new) fusion machine generate electrical power at a cost ≤ $43/MWh at the plant gate; and can a retrofit beat the site-specific marginal cost of continued coal operation?

Furthermore, the $1B program is intended to answer these questions:

1. **What is the true net electrical output?**

   Can a Volpe P1-class plant really export about 1 GWe, or do tritium systems, pumps, cryogenics, vacuum, magnets, heating, controls, and other parasitic loads reduce saleable output toward ~700 MWe? The current P1 estimates leave this unresolved.

2. **Can the complete plant actually be built cheaply enough in China?**

   What does every Volpe/ARIES cost account become when rebuilt from bottom-up Chinese vendor prices, fabrication costs, labor, installation, construction schedule, and financing rather than historical U.S.-derived assumptions?

3. **Can greenfield CAPEX reach roughly $2,000/kW-net or lower?**

   If not, which components prevent it—magnets, structures, blanket, heat-transfer system, buildings, turbine island, tritium systems, or indirect costs?

4. **How much cheaper can a coal-plant retrofit be?**

   Which existing assets can actually be reused—turbine, generator, condenser, cooling towers, transformers, switchyard, buildings, grid connection—and what modifications are required?

5. **Can technical availability approach 95%?**

   What fails, how often does it fail, how long does repair take, and which systems need redundancy, bypasses, isolation, modular replacement, or online maintenance? The result should come from a bottom-up RAMI model, not simply an assumed percentage.

6. **Does the liquid-metal wall actually produce the economic benefits we are counting on?**

   Can it handle the heat flux, reduce neutron damage, remain stable in magnetic fields, avoid excessive pumping power, limit corrosion and contamination, manage tritium, and allow high availability over long periods?

7. **Can the stellarator reach an economically useful plasma operating point?**

   Not necessarily infinite, but sufficiently good confinement and sufficiently low auxiliary power that the plant's total LCOE is minimized.

8. **Can the D-T fuel cycle close?**

   Is tritium breeding adequate after realistic penetrations and losses? How much startup inventory is needed? What do extraction, cleanup, accountability, permeation control, pumping, maintenance, and downtime cost?

9. **What is the least-cost power-conversion system?**

   For greenfield construction, does a high-efficiency advanced cycle win? For retrofit, is it cheaper to produce steam compatible with an existing Chinese coal turbine even if efficiency is somewhat lower?

10. **Can the machine be industrialized rather than built like a research experiment?**

    Which parts can be off-the-shelf? Which must be redesigned for automated Chinese factory production? How many factories, magnet-production lines, blanket modules, heat exchangers, and other components are needed to reach tens of GW/year and ultimately very large deployment rates?

11. **Can construction be fast enough?**

    Is approximately 24–36 months of mature site work realistic, particularly for a retrofit where much of the fusion island is installed while the coal plant is still operating?

12. **When every real cost is included, does the economics still work?**

    That means CAPEX, financing, O&M, replacements, parasitic power, outages, licensing, safety, waste, tritium systems, decommissioning, and refurbishment of reused coal assets—not just the fusion core.

What is Economic Fusion?

There are two primary types of nuclear power: fission and fusion.

Fission is the traditional form that generates electricity with uranium fuel. However, this is not popular due to meltdown risk, nuclear waste, nuclear bomb proliferation risk, and cost.

Fusion, on the other hand, does not have these issues; however, it is still in development. Typical fusion systems maintain a hot plasma within a donut-shaped reactor, as illustrated here.

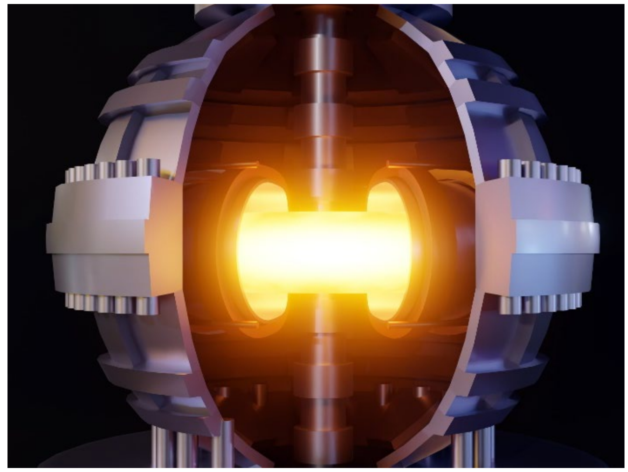

Economic Fusion

Commercial fusion timelines remain uncertain and vary widely across programs. EFM does not depend on a particular industry forecast. It asks a narrower question: can a concentrated five-year engineering and experimental program resolve the assumptions that determine whether a mature fusion plant could become economically preferable to fossil generation?

“Economic fusion” refers to generating electricity with fusion at a cost below that of fossil fuel-based electricity. This would require the fusion machine to operate reliably for extended periods, without failure, while keeping costs low.

If green energy were cheaper than fossil fuels, nations would have a powerful incentive to decarbonize, to save money, to remain competitive, to keep voters employed. In other words, cheaper green energy could potentially solve the carbon dioxide emissions problem.

Current fusion machines take years to construct and operate for seconds to minutes per experiment. Many scientists hope this will lead to economic fusion within a few decades. Alternatively, a surge of funding in key areas could potentially change this trajectory. For a video that summarizes this concept, see Economic Moonshot Video [#14].

Fusion Moonshot

Harm from climate change is likely to cost trillions of dollars; therefore, it is reasonable for governments and foundations to allocate billions of dollars to achieve economic fusion within several years. While this might seem unrealistic, consider the remarkable number of “gadgets” the U.S. designed and manufactured between 1939 and 1945.

In 1961, President Kennedy declared his goal to land a man on the moon by the end of the decade. In response, a program was established and funded. In theory, a similar approach could be applied to economic fusion.

But how might our society tackle fusion R&D differently? Let's explore.

Tokamak vs. Stellarator

The tokamak and the stellarator are two fusion designs that could potentially produce large amounts of electricity in the near future. The tokamak features a donut-like shape, while the stellarator is more bumpy.

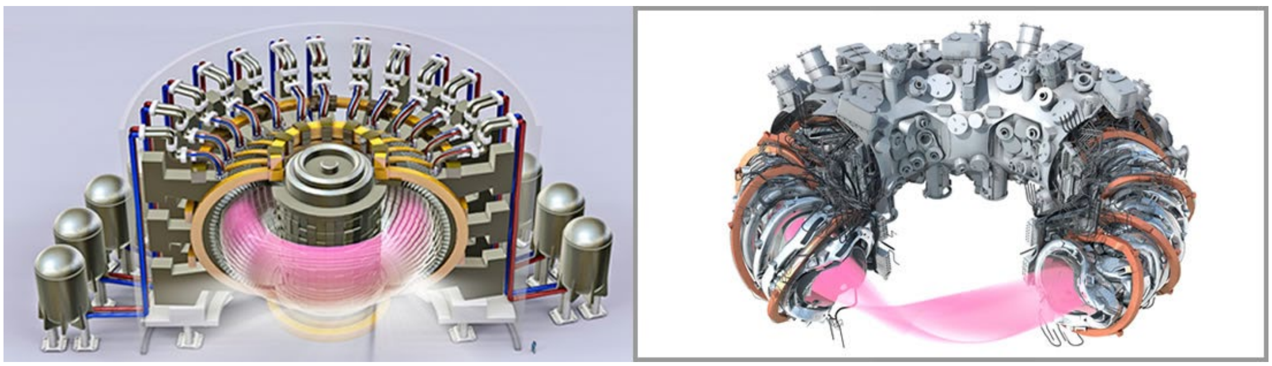

Figure 1: Tokamak (left) and stellarator (right).

These two approaches to fusion differ in more than geometry. Tokamaks rely on a very large toroidal plasma current, while stellarators generate the confining magnetic geometry primarily with external coils. Tokamaks must therefore manage current-driven disruptions and associated electromagnetic, thermal, and runaway-electron loads. Stellarators avoid that particular operating burden, although they introduce their own challenges in magnet geometry, optimization, construction, and maintenance.

Sustaining such currents is no small feat. For context, a typical household vacuum cleaner draws about 10 amps, making tokamak currents equivalent to a million vacuum cleaners operating simultaneously. If a "disruption" occurs and this current inadvertently contacts the containment vessel wall, severe damage can result. Tokamaks must contend with this risk; however, plasma in stellarators do not have these currents, and are therefore easier to maintain.

Damage is Bad

Damage within high-energy fusion devices is especially problematic because internal components typically become radioactive during operation. Radioactivity results from neutron activation and from the absorption of radioactive tritium gas into surfaces, which later outgas. As a result, human contact with these components must be avoided for many years after the reactor begins operation.

If ITER, a major tokamak project, experiences damage after starting high-energy operations (expected around 2034), repairs might be impossible due to radioactivity hazards. In other words, due to worker safety, the site might need to be abandoned—despite its roughly 50-year construction period.

Cost Advantages of Stellarators

Certain power-plant cost studies suggest that stellarators may offer economic advantages in some configurations. Potential advantages include steady-state operation without tokamak-style major current disruptions, different magnet and maintenance architectures, and the possibility of reducing recirculating power if an ignited or near-ignited operating point can be achieved. Whether those advantages survive complete engineering, maintenance, and cost accounting is a question for the EFM program to test rather than assume.

Making Electricity

A fusion power plant must capture heat from plasma-facing and blanket systems, transfer that heat through controlled coolant and heat-exchange loops, and convert it to electricity. Candidate working fluids and power cycles depend on the reactor architecture and may include liquid metals, molten salts, gases, water/steam, or other intermediate loops. The EFM program should choose the combination that minimizes plant-gate LCOE after parasitic power, reliability, safety, maintenance, and equipment cost are included.

To produce electricity cost-effectively, the fusion thermal output must be large enough that the electrical plant still exports substantial net power after plasma heating, magnets, cryogenics, pumps, vacuum, tritium processing, controls, and conventional station loads are included. Plasma confinement and magnetic-field performance have substantial experimental foundations, but the economically optimal integrated power-plant operating point remains uncertain and is one of the program questions to resolve.

First Wall

There are two primary methods for handling the plasma-facing "first wall." One involves flowing liquid directly in *front* of a metal plate, and the other involves flowing liquid *behind* a metal plate.

Liquid Metal Wall

In theory, a roughly 0.5m thick layer of flowing molten lithium metal could cover the internal surface of the plasma containment chamber (figure 2). This would be held in place by electric currents in the liquid metal itself.

Pushing Liquid Outward

As one might recall from physics class, a force is exerted on a current-carrying conductor in the presence of a magnetic field (figure 3). In other words, one can embed posts within a fluid, move current between the posts, and maintain an evenly distributed wall of liquid by pushing it toward the plate, even if the plate is above the fluid.

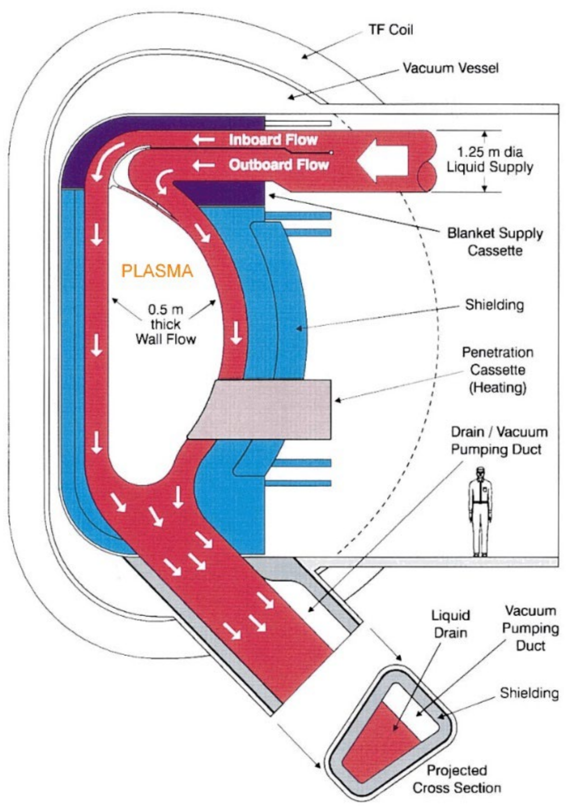

Figure 2: Plasma radiates against liquid wall (red), held in place against metal plate (blue) by electrical currents and by magnet fields. This image first appeared in a 2001 paper by Dr. Mohamed Abdou.

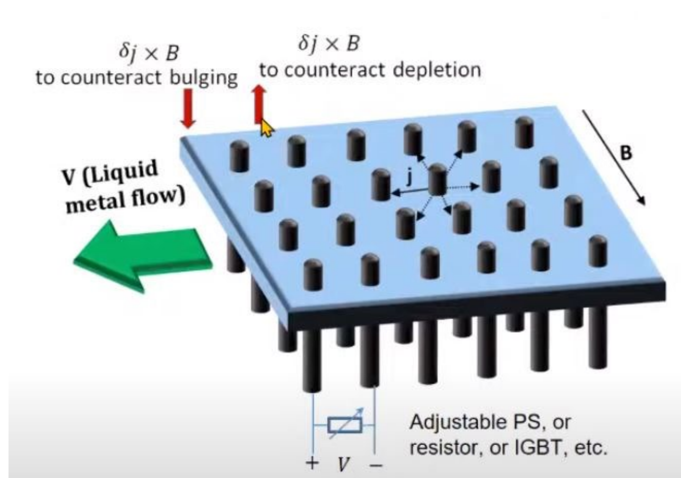

Figure 3: Flowing liquid is pushed toward plate by surrounding magnetic field and by electrical currents between posts. Courtesy Renaissance Fusion.

Internal Cooling System

The alternative cooling method is to expose a thin (~2mm) metal plate to neutron radiation and cool it by flowing molten liquid along its back surface. High neutron exposure can embrittle and damage solid plasma-facing structures, creating periodic replacement and remote-maintenance requirements. If a thick flowing liquid-metal wall can sufficiently reduce heat and neutron damage to underlying structures, it could reduce or eliminate some major scheduled replacement operations. Demonstrating whether that lifetime and availability advantage is real is a central objective of the EFM program.

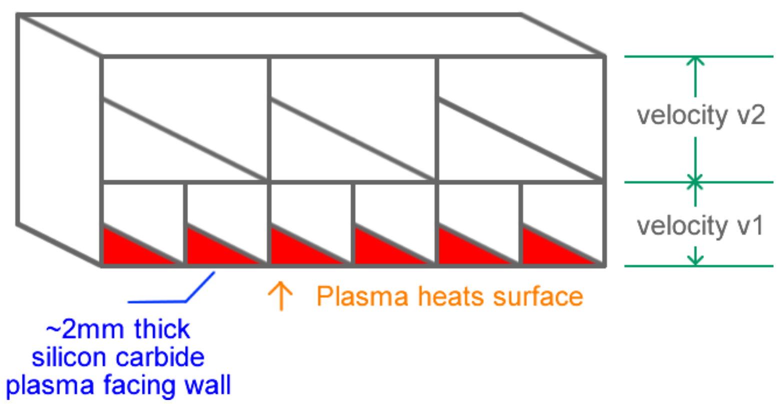

Figure 4: Liquid flows in channels behind metal plate (red) and removes heat.
Concept illustration by G. Weinreb.

Liquid Metal Wall vs. Internal Cooling System

There is a limit to how much power one can radiate against a metal plate. In theory, more power can be radiated against a wall of flowing liquid metal. Additionally, increasing the power per square meter of internal wall surface area reduces the reactor's physical size for a fixed amount of total power. This is good, since smaller physical size costs less to build.

In summary, a flowing liquid-metal wall could offer substantial advantages, including higher allowable heat loading, reduced damage to some underlying solid components, greater power density, and potentially fewer major in-vessel replacement operations. Those benefits are not assumed. Liquid-metal-wall technology remains insufficiently validated for a commercial power plant, and experiments are needed to establish surface stability, MHD behavior, corrosion, tritium handling, heat extraction, neutron shielding, component lifetime, pumping power, and maintainability.

For details, see liquid metal wall video, tutorial, design and theory.

Magnet Making Machines

Fusion devices contain large, expensive magnets whose cost could be reduced by developing specialized machines capable of: (a) producing superconducting REBCO-like tape more quickly and at a lower cost, (b) print large 2 to 200cm wide conductors onto a rotating tube, (c) print tiny 0.1cm to 2cm wide conductors directly onto a rotating tube, and (d) print conductors directly onto a magnet baseplate.

3D Magnet Printing

Stellarators require irregularly shaped magnets to confine plasma, and these magnets are expensive to fabricate using traditional methods. In theory, these magnets could be made by placing alternating layers of superconductor and “insulator” onto a rotating tube via an additive process such as ion deposition. For example, one might place five irregular shaped magnets onto each of 12 tubes, and place a large magnet at the end of each tube—as illustrated lower-right. This could potentially reduce stellarator cost. Therefore, it is reasonable to dramatically increase funding for the development of 3D magnet printing machines. For details, see this video.

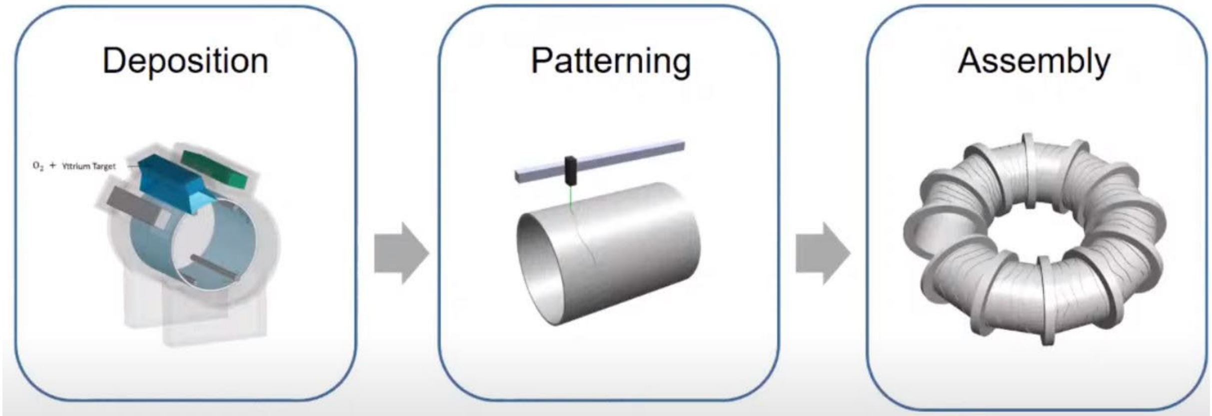

Figure 5: 3D magnet printing. Courtesy Renaissance Fusion.

Paper Only Designs

In theory, researchers can design an economic fusion machine entirely on paper, and publish their work. This has been done. For example, the Volpe/Frost P1 design is expected to generate $1GW_e$ of electrical power at a cost comparable to that of fossil fuel-based electricity.

Magnet costs account for ≤18% of the total, with the remaining expenses distributed across the facility. For a closer look at a fusion facility, check out this tour of the SPARC site.

For more details on the Prost/Volpe design, refer to their 2024 paper, design summary, energy diagram, cost breakdown, YouTube video, and cost model. This design incorporates new innovations, such as liquid metal wall and 3D magnet printing. However, as noted earlier, these technologies have yet to be validated in a working fusion machine.

The work by MIT and Commonwealth Fusion Systems (CFS) is also noteworthy. However, their papers refer to a more expensive approach that does not include liquid metal walls. As a result, their machine requires disassembly every one to five years, which is expensive. For a summary of parameters that describe a fusion machine, see fusion_reactor_design_parameters.pdf

Parallel, not Serial

The fusion research community is operating under the assumption continuously-operating fusion machines will not appear soon. As a result, relatively little funding has been directed toward developing the equipment required for continuous operation. For instance, specialized systems are needed to separate exhaust gases (such as helium, tritium, and deuterium) and extract tritium from a stream of liquid lithium. To accelerate economic fusion R&D, all subsystems must be developed concurrently, supported by equipment that tests these systems at maximum power for months.

High-Power Long-Duration Experiments

A commercial fusion plant would expose plasma-facing liquids, solids, pumps, heat exchangers, sensors, tritium systems, and structural materials to coupled thermal, magnetic, radiation, chemical, and maintenance stresses for years. No single experiment in this proposal is assumed to reproduce all of those conditions at once. The five-year EFM program therefore combines existing specialized facilities with new P1-specific MHD, heat, corrosion, magnet, tritium, maintenance, and integrated-sector tests.

A dedicated long-duration experimental fusion reactor could eventually provide uniquely valuable integrated neutron, plasma, materials, tritium, and heat-extraction data. Such a machine would be a subsequent Phase II or separate program, not a deliverable of this five-year $1 billion initiative.

The program should use CRAFT and other existing or planned fusion, irradiation, plasma-exposure, and materials facilities whenever they can answer the required question. New EFM capital should concentrate on P1-specific evidence gaps that cannot be resolved efficiently through those facilities.

The five-year EFM program therefore ends with controlled plant designs, validated subsystem and full-sector hardware, endurance data, manufacturing/construction evidence, and an independently audited economic decision. A complete power-producing fusion pilot or commercial demonstration would require a separate authorization and budget after the final G6 decision.

Fusion Companies are not Aiming to Fix the Climate Problem

Economic fusion could significantly impact the climate crisis. However, many fusion companies have a different focus. For example, Commonwealth Fusion Systems aims to develop a single fusion machine in the 2030's ("ARC") capable of producing electricity at a cost much higher than fossil fuel-based electricity. This single machine would be paid for by investors, rather than by the typical electricity customer.

In other words, it is easier to build a fusion machine that generates electricity at $250/MWh, funded by investors, than to build one that generates electricity at $40/MWh, funded by customers. While fusion companies are concentrating on the former, climate-focused funding is looking at the latter, at an accelerated timeline, and at a higher funding level.

Private fusion companies reasonably optimize around investor capital, technical milestones, commercialization schedules and first-of-a-kind plants. The Economic Fusion Moonshot asks a different question: what additional publicly or philanthropically funded R&D would be needed to enable climate-scale deployment (e.g. ≥ 100 GWe/year).

Government Support

The U.S. government currently allocates approximately $800 million annually to basic fusion research, supporting programs such as PPPL, ITER, and DIII. In theory, a *new* initiative could be launched to achieve economic fusion sooner through a top-down approach, led by the world's leading scientists. Such an initiative would integrate plasma physics, manufacturing engineering, and cost analysis. Priority would be given to de-risking key technologies, such as liquid metal walls and 3D-printed magnets.

Open Fusion

Companies and research laboratories often avoid sharing developed materials, fearing that open review could negatively impact their funding. However, achieving economic fusion sooner may require a shift toward greater transparency. Sharing experimental data and materials online for open review could accelerate progress by identifying and addressing problems more quickly. In theory, climate money could pay test laboratories, such as China's CRAFT, to share ("collaborate"). It could also fund scientists and engineers to conduct R&D and publish open-source materials online, maximizing the climate benefit per dollar spent.

Program Requirements

The requirements for the proposed $1B Economic Fusion Moonshot (EFM) program are specified in this chapter.

Primary Design Targets

The following table summarizes the primary goals of our initiative. All electricity values are net plant-gate electricity *after* all fusion-system and balance-of-plant parasitic loads.

For details on existing electrical power costs in China, see the "Cost Comparison: Electricity Generation Only" chapter within the Proposed $1B Nuclear Power Automated Construction document.

Metric	Primary requirement / target
China greenfield LCOE	< $43/MWh-electric; preferred <= $40/MWh; stretch <= $35/MWh
China coal-retrofit LCOE	Below the site-specific marginal cost of continued coal operation; screening target <= $35/MWh-electric
China greenfield CAPEX	<= $2,000/kW-net (about $2.0B at 1 GWe net, ≤ $2/Watt)
China coal-retrofit CAPEX	No arbitrary hard ceiling initially; must be low enough to pass the retrofit LCOE test and must be reported in total dollars and $/kW-net
Commercial net output	Approximately 1,000 MWe net for the reference design
Technical plant availability	>= 95% design target, supported by bottom-up RAMI evidence
Mature construction duration	<= 36 months from major site-construction authorization to commercial operation; preferred approximately 24 months for major fusion-site work

Economic Fusion Moonshot Program Objectives

Our mission is to develop fusion electricity that is economically preferable to fossil-fuel electricity in the Chinese deployment cases that matter most for rapid decarbonization. The program has two distinct economic tests (“greenfield” and “retrofit”).

Requirement	China greenfield fusion	China coal-plant retrofit
Competitor	Construction of a new Chinese coal-fired power plant	Continued operation of an existing Chinese coal-fired power plant
Reference cost	$43 to $58/MWh-electric LCOE	$35 to $45/MWh-electric marginal generation cost
Hard economic test	Fusion LCOE < $43/MWh-electric at the	Fusion retrofit LCOE < the site-specific

	plant gate	marginal cost of continued coal operation
Preferred target	<= $40/MWh-electric	<= $35/MWh-electric screening target
Capital-cost boundary	Entire new fusion power plant	Incremental fusion island, integration, refurbishment/modification, and all required new spending; reused assets disclosed separately
Primary decision	Build fusion instead of new coal	Retire the coal heat source and retrofit the existing electrical island with fusion

The reference benchmarks should be refreshed with current Chinese data during the program. The low end of each relevant cost range is the primary screening hurdle so that economic success is not created by choosing only an expensive comparator.

Controlled Reference Designs

The program should freeze two controlled reference designs early enough that all engineering teams use the same geometry, power balance, site assumptions, accounting boundaries, and economic definitions.

- **Reference A** - China Greenfield P1.

  A mature nth-of-a-kind P1-class stellarator of approximately 1 GWe net, using a liquid-metal plasma-facing system and D-T fuel unless a better economic alternative is demonstrated. The reference includes the complete new power plant and site infrastructure.

- **Reference B** - China Coal-Retrofit P1.

  The same general fusion architecture integrated with a representative approximately 1-GWe Chinese ultra-supercritical coal power plant. In Year 1, the program should nominate one actual or fully specified reference unit, including age, steam conditions, turbine-generator rating, cooling system, electrical island, switchyard, remaining useful life, and site layout. Young, midlife, and old coal-plant cases should also be evaluated as sensitivities.

Changes to either reference design must be configuration-controlled. Economic comparisons should not allow one team to assume a nearly new coal turbine while another assumes a 30-year-old turbine, or otherwise change the reference boundary to obtain a favorable result.

Five-Year Program Boundary and Phase II

The $1 billion EFM initiative does not construct or operate a complete 1-GWe fusion power plant and does not assume that a dedicated multi-billion-dollar fusion neutron source can be built inside the five-year budget. It develops the controlled reference designs, manufacturing methods, component prototypes, P1-specific experimental facilities, integrated engineering sector, endurance data, and independent economic evidence needed to decide whether a larger pilot should be authorized.

A later Phase II may include a power-producing pilot, a dedicated long-duration fusion test machine, or a first commercial-scale demonstration. Such a program should be authorized only after the G6 review identifies a credible plant architecture, quantified remaining risks, a validated path to the economic target, and a financing and licensing plan appropriate to the next scale.

Public R&D and Manufacturing-Capital Assumptions

We assume billions of dollars may be available as climate-focused public or philanthropic R&D capital to reduce the cost of the mature fusion machine. For example, a program could develop automated equipment that manufactures REBCO-like conductor, then replicate those manufacturing machines at contract manufacturers so conductor is sold near production cost rather than at a price intended to recover the original R&D investment. Put differently, “climate money” (which is not looking for a return on investment) could potentially: (a) support the design of an automated REBCO tape production machine, (b) give the design away for free, (c) have dozens of these machines built at a contract manufacturing facility, and then (d) sell these machines at cost to interested parties.

This kind of "climate money" may legitimately reduce future REBCO tape production cost, but the accounting must remain transparent. Every major economic review should show two views: (1) project-owner REBCO plant economics, which may exclude previously spent public R&D and public-good tooling that has become sunk cost; and (2) fully loaded societal economics, which separately identifies public R&D, grants, factory/tooling subsidies, and other externally funded capital. The headline fusion-plant LCOE may use project-owner economics, but public contributions must not disappear from the record.

Ultimately, the EFM program asks the following question:

> How can public and philanthropic “climate money” — funding that does not require a financial return — be used to make mature fusion electricity cheaper than fossil-fuel electricity? Put differently, what can be developed and given away for free to make green energy cost less than fossil fuel? For example, in theory, climate money could develop an open-source fusion machine, or an automated REBCO production machine.

The objective is not to make fusion appear cheap by indefinitely subsidizing each power plant. It is to use one-time R&D and public-good donations to reduce the underlying mature cost of fusion.

Open Research and Data Requirement

Portfolio-wide open-source, data, foreground/background-IP, release, audit-access, and transition rules are specified in the Moonshot Common Framework. The EFM additions below apply because the program relies on unusually consequential cost models, fusion-specific test data, and manufacturing equipment intended to lower industry-wide costs.

Fusion-specific open-research additions

- Release the controlled plant power balance, China cost-model methods, RAMI model, tritium inventory model, test protocols, calibration records, source code, interface specifications, failures, and negative results in reusable form, subject only to narrowly justified restrictions.
- For magnet, liquid-metal, tritium, power-conversion, and integrated-sector work, retain enough source-level evidence that an independent team can reproduce the claimed effect on net power, availability, CAPEX, and LCOE.
- When climate-focused funding develops manufacturing equipment primarily to reduce industry-wide cost - for example automated REBCO-like conductor or patterned-magnet production equipment - publish sufficient design/process information for competitive replication when lawful and technically appropriate.
- Any supplier confidentiality, export-control, security, or background-IP restriction that prevents full release must be identified before award and carried explicitly into the final evidence and transition package.

CAPEX and LCOE Requirements

For China greenfield construction, the mature CAPEX target is <= $2,000/kW-net. At 1 GWe net, this corresponds to total capital cost of <= $2.0B. If net output differs materially from 1 GWe, the CAPEX target scales with net output and must still be reported in both total dollars and $/kW-net.

For the China coal-retrofit case, the governing requirement is LCOE rather than an arbitrary fixed CAPEX ceiling. Retrofit CAPEX must include the fusion island, integration with the existing coal plant, owner and engineering costs, contingency, financing during construction, and any refurbishment or replacement required to make reused assets reliable for the fusion plant economic life. All reused coal-plant assets excluded from incremental CAPEX must be listed explicitly.

Coal Retrofit Strategy

Retrofits should reuse existing assets whenever that lowers LCOE, including the turbine-generator, condenser, cooling system, feedwater equipment, water treatment, switchyard, transformers, grid connection, transmission lines, control facilities, buildings, roads, rail, workshops, cranes, and site utilities.

For every major existing coal-plant asset, the program must determine whether the least-cost option is to reuse it as-is, refurbish it, modify it, or replace it. The reference site must have sufficient remaining useful life in reused assets; otherwise, refurbishment or replacement cost must be included in fusion CAPEX and LCOE.

Because a retrofit can use sunk assets, the economic analysis must show two asset-accounting views: (1) an incremental retrofit case, charging the fusion project for new spending, refurbishment, modification, and integration actually required; and (2) a transparency case that also assigns reused assets their remaining economic value or opportunity cost where material. The headline retrofit test is the incremental case, while the transparency case prevents the apparent advantage from being hidden in unpriced asset reuse.

The commercial design target is approximately 1,000 MWe net output so that a P1-class fusion retrofit can make full use of a roughly 1-GW coal-plant electrical island. Net output, rather than fusion power or gross generator output, is the economic metric.

The fusion heat system should preferentially produce steam conditions compatible with the existing Chinese coal turbine-generator when that minimizes LCOE. The program must compare reuse of the existing steam cycle with modified steam cycles, supercritical-CO2 cycles, or other power-conversion systems and select the lowest net $/MWh option, not automatically the option with the highest thermodynamic efficiency.

Net Power and Plant Power Balance

The program shall use consistent plant-power terminology so that "net" electricity cannot exclude inconvenient internal loads:

P-fusion = fusion thermal power produced by the plasma

P-thermal = useful thermal power delivered to the electrical power-conversion system

P-gross = gross generator electrical output

P-recirc = fusion-system electrical consumption, including plasma heating/current drive where applicable, magnets, cryogenics, tritium systems, vacuum, and fusion-specific pumping

P-aux = conventional balance-of-plant auxiliary electrical consumption, including secondary pumps, heat rejection, controls, buildings, and other station service

P-net = P-gross - P-recirc - P-aux

All LCOE, $/kW, and commercial-capacity claims must use P-net exported at the plant gate. Volpe's optimistic approximately 1.01-GWe result and the approximately 700-MWe case that includes up to about 320 MWe of omitted loads should be treated as scenarios to resolve, not as predetermined answers.

Availability and Capacity Factor

Technical availability and capacity factor are related but not identical. Technical availability is the fraction of time the plant is physically capable of operating. Capacity factor is actual net electricity generated divided by nameplate net output times 8,760 hours (one year). A technically available plant can have a lower capacity factor because of grid curtailment or economic dispatch.

The design target is >= 95% technical plant availability. This is a target to achieve and demonstrate, not an assumption that may be used without support. Economic results must separately report capacity-factor cases at 80%, 85%, 90%, and 95%, together with the best demonstrated or bottom-up RAMI-supported technical availability. Grid curtailment and economic dispatch must be reported separately from equipment unavailability.

The program must develop a bottom-up reliability, availability, maintainability, and inspectability (RAMI) model that identifies every major source of scheduled and unscheduled downtime. Where economically justified, pumps, heat exchangers, vacuum equipment, tritium-processing trains, controls, power electronics, and other balance-of-plant systems should use N+1 redundancy, isolation valves, bypasses, modular replacement, and online maintenance so a single failure does not unnecessarily shut down the fusion plant.

Stellarator Baseline

The baseline confinement concept is a stellarator rather than a tokamak. The economic rationale is steady-state magnetic confinement without the large externally driven toroidal plasma current required by a tokamak, thereby avoiding tokamak-style major current disruptions and their associated electromagnetic, thermal, and runaway-electron loads. The design must still address all remaining steady and transient plasma-facing loads.

Baseline and Challenger Architectures

P1 remains the controlled baseline so that the program can accumulate comparable engineering and economic evidence. During roughly the first 12–18 months, however, limited challenger studies should be maintained for alternatives that could materially change LCOE: magnet/current-surface manufacturing, liquid-wall material and flow architecture, power-conversion choices, maintenance geometry, and one or more credible stellarator design variants. Challengers must use the same China accounting boundary, net-power definitions, RAMI requirements, and deployment assumptions as the baseline.

The purpose is not to reopen broad fusion research. At G1 and G2, alternatives that do not show a credible economic advantage or risk-reduction benefit should be stopped; the remaining hardware budget should converge rapidly on one or at most two controlled reference architectures.

Minimize LCOE Electrical Power Cost Instead of Maximizing Q-plasma

The program must not assume infinite Q-plasma or ignition merely because the machine is a stellarator. It should seek an ignited or near-ignited operating point that minimizes auxiliary heating and recirculating electrical power. If a design with finite auxiliary heating produces lower LCOE, that design is preferable. Economic fusion, not maximum Q-plasma, is the objective.

Liquid Metal Wall

The baseline P1 architecture uses a liquid-metal plasma-facing wall because it may permit higher thermal loading, continually renew the plasma-facing surface, reduce neutron damage to solid structures, and support higher availability. These benefits must be demonstrated rather than assumed. The liquid-wall system must show acceptable heat-flux capability, MHD pressure drop and pumping power, surface stability, evaporation and plasma contamination, corrosion, tritium behavior, neutron shielding, leak management, maintainability, and component lifetime consistent with the LCOE and 95% availability targets.

Tritium Breeding and Fleet Startup Inventory

For a D-T baseline, the plant must demonstrate a credible tritium self-sufficiency strategy with adequate breeding margin, startup and operating inventory, extraction and processing, accountability, permeation control, and reserve requirements. The capital cost, parasitic electrical load, maintenance burden, and downtime associated with the complete tritium fuel cycle must be included in the economic model.

A single plant being tritium-self-sufficient is not enough for a rapidly expanding fleet. The program must model startup tritium inventory and breeding growth for deployment rates of tens of GWe per year and ultimately approximately 100 GWe per year. If startup tritium availability constrains fleet growth, the design must reduce inventory, improve breeding/extraction, change commissioning strategy, or evaluate an alternative fuel-cycle architecture.

Lowest-Cost Design and Manufacturing

The program should buy off-the-shelf Chinese industrial equipment whenever practical and when doing so reduces lifecycle cost without compromising safety, reliability, maintainability, licensing, or required performance. Fusion-specific custom equipment should be limited to functions that cannot economically be met with standardized industrial equipment.

Components that cannot be purchased off the shelf must be designed for high-volume Chinese factory production using automated fabrication, standardized mechanical/electrical/thermal/data interfaces, interchangeable modules, automated inspection, and multiple qualified suppliers where practical. The mature design should be optimized for serial production rather than one-off scientific-machine construction.

Buildings and site civil works must be aggressively minimized. Volpe's model estimates roughly $430M for buildings, but the China designs should use compact standardized industrial buildings, outdoor equipment, factory skids, modular structures, and existing coal-site structures where they lower LCOE. Only the minimum necessary volume should require specialized nuclear-grade construction.

The mature deployment target is no more than approximately 36 months from major site-construction authorization to commercial operation, with a preferred mature target of approximately 24 months for major fusion-site work. In a retrofit, much of the fusion island should be constructed while the coal plant remains in operation, with final steam, electrical, and control tie-ins completed during the shortest practical planned outage.

Fleet-Scale Manufacturing and Deployment

A reactor that meets the LCOE target but cannot be manufactured at climate-relevant scale does not meet the mission. The mature architecture must show a credible industrial path first to tens of GWe per year and ultimately to approximately 100 GWe per year of deployment.

The program must quantify annual requirements and bottlenecks for HTS conductor, lithium and other blanket materials, structural materials, tritium startup inventory, specialized manufacturing machines, factory floor area, high-temperature pumps and heat exchangers, power-conversion equipment, skilled labor, ports/rail/heavy transport, and other long-lead items. The analysis should identify the number, size, capital cost, and lead time of factories required to support each deployment-rate scenario.

Cost Calculations and Accounting Rules

All economic comparisons must distinguish first-of-a-kind costs from mature China nth-of-a-kind costs. Learning, serial manufacturing, construction repetition, financing, supplier scale, and reused coal-plant assets must be stated explicitly rather than embedded in unexplained cost reductions.

The retrofit case shall use a formal incremental LCOE definition: annualized incremental fusion-retrofit CAPEX plus fusion O&M, refurbishment, replacement, consumables, and other incremental costs, divided by net MWh sold. This must be compared with the site-specific marginal cost of continued coal operation using a consistent operating period and currency basis. The coal marginal-cost model shall identify fuel, variable O&M, pollution-control consumables and energy, and other costs that actually change with an additional MWh of coal generation.

All LCOE calculations must state the financing basis explicitly. At a minimum, report sensitivity to 3%, 5%, and 7% real cost of capital, the assumed construction draw schedule, construction duration, inflation convention, tax treatment where applicable, and plant economic life. A financing advantage must not be hidden inside an unexplained LCOE reduction.

Every design review must report a complete audited cost model in $/kW-net and $/MWh-net, including CAPEX, fixed and variable O&M, replacement costs, parasitic electricity, scheduled and unscheduled downtime, financing, construction duration, plant life, decommissioning where material, and any refurbishment of reused coal assets. Costs and performance values must be identified as measured, vendor-quoted, detailed engineering estimates, simulations, or assumptions.

The economic target may not be achieved by excluding required safety, licensing, quality assurance, security, radioactive-material management, waste handling, tritium-control, or decommissioning obligations. The program should minimize these costs through design, standardization, passive features, and reuse where justified, but all required functions and their costs must remain inside the system boundary.

For each of the three scenarios - Volpe P1 Reference, China greenfield, and China coal retrofit - the program must estimate total and direct CAPEX, $/kW-net, net electrical output, technical availability, capacity factor, annual O&M, replacement cost, construction duration, and LCOE ($/MWh-net). The same accounting boundary and financing assumptions should be used wherever possible so differences reflect real engineering and localization effects rather than inconsistent bookkeeping.

China Cost Localization

Volpe's plant-cost model largely inherits ARIES cost accounts that were developed from U.S. power-plant and construction cost data and then escalated to later dollars, while a number of fusion-specific accounts use P1-specific calculations or other technical sources. The program must therefore trace the basis of every major cost account before localizing it. For China planning, costs should be rebuilt in constant 2026 U.S. dollars using mature China nth-of-a-kind assumptions rather than applying one blanket "China discount."

For every cost account in Volpe's model, the program should build a table showing: (a) Volpe/ARIES reference cost and cost-basis year, (b) China greenfield cost, and (c) China coal-retrofit incremental cost. Each row should show the physical quantity or sizing basis, Chinese vendor or material price, factory fabrication, installation labor, site equipment, engineering, construction schedule, refurbishment or reuse credit, evidence source, uncertainty range, and evidence class. Reused assets must not be double-counted as both a savings and a separately omitted cost.

Requirements Verification and Evidence Maturity

Every major requirement shall have a defined verification method and evidence maturity. A technical improvement is valuable only to the extent that it credibly lowers net LCOE, raises availability, increases net electrical output, reduces CAPEX, shortens construction or maintenance outages, improves manufacturability, or reduces another economically important risk.

Requirement	Target	Verification / evidence expected
Greenfield LCOE	< $43/MWh-electric; preferred <= $40/MWh	Audited bottom-up China NOAK model, consistent financing assumptions, Chinese vendor quotations, and progressively more measured component/integration data
Retrofit LCOE	< site-specific marginal coal cost; screening <= $35/MWh	Reference-site incremental lifecycle model using actual/reconstructed coal marginal costs and audited fusion retrofit costs
Greenfield CAPEX	<= $2,000/kW-net	Bill of materials, Chinese supplier quotes, factory/site labor estimates, construction schedule, indirects, contingency, and financing
Net electrical output	Approximately 1,000 MWe net	Complete integrated plant power balance followed by subsystem and integrated-test measurements
Technical availability	>= 95%	Bottom-up RAMI model, redundancy/maintenance design, component endurance data, integrated operating evidence, and explicit outage budget
Liquid-metal wall	Performance sufficient for LCOE and availability targets	Long-duration heat/MHD/corrosion/plasma-compatibility tests, leak/failure testing, and integrated sector prototype
Tritium self-sufficiency and fleet growth	Adequate plant breeding margin and scalable startup inventory	Neutronics, extraction/process tests, inventory model, measured retention/permeation data, and fleet growth analysis
Construction duration	<= 36 months mature; preferred approximately 24 months major site work	Resource-loaded construction model, factory-production plan, site logistics simulation, and representative installation demonstrations

Manufacturing scale	Credible path to tens of GWe/year and ultimately approximately 100 GWe/year	Factory/supply-chain model with material quantities, machine counts, floor area, labor, long-lead equipment, capital cost, and ramp schedule

Evidence maturity should progress from assumption -> simulation -> engineering calculation -> vendor quote -> component test -> integrated test -> operating data.

The standalone Stage Gates and Go/No-Go Criteria chapter converts this evidence ladder into program-level funding decisions. Area-specific hold points may add stricter technical readiness criteria, but they do not replace the common G0–G6 decision architecture.

Program Budget

Prompt to AI: Given $1B over 5 years, how would you (supreme intelligence) achieve the objectives specified in the previous chapter?

Response From AI:

Given the requirements specified in the previous chapter, I would run the $1 billion as an industrial product-development program rather than a conventional fusion-science program. The planned baseline is $900 million across ten technical areas, with a separate $100 million sponsor-controlled reserve. The objective remains specific: determine whether a mature Chinese stellarator can meet the greenfield and retrofit economic tests, reach roughly $2,000/kW-net, export about 1 GWe net, approach 95% technical availability, and be manufacturable at climate-relevant scale.

Budget Summary

My proposed allocation would be:

Work package	5-year budget	Main question answered
Area 1. Systems engineering, China cost model & controlled reference designs	$65M	What does the complete plant actually cost in China?
Area 2. Stellarator plasma / confinement / operating-point validation	$70M	Can P1 reach an economically useful plasma state?
Area 3. HTS magnets, structures & low-cost manufacturing	$110M	Can the complete magnet system be built cheaply and repeatedly enough?
Area 4. Liquid-metal wall, blanket, shielding & heat extraction	$155M	Does the liquid-wall architecture work well enough to create an economic advantage?
Area 5. Tritium breeding, processing & fleet-startup inventory	$80M	Can the D-T fuel cycle close and scale to a rapidly growing fleet?
Area 6. Net power, BOP, power conversion & coal-retrofit integration	$90M	What net power is actually exported, and what coal-site assets can be reused economically?
Area 7. RAMI, maintenance, redundancy & long-life components	$80M	Can technical availability credibly approach 95%?
Area 8. China factory production, automation & construction engineering	$65M	Can the plant be manufactured and installed at tens to approximately 100 GWe/year?
Area 9. Full-scale integrated sector prototypes & endurance testing	$140M	Do the critical systems work together at realistic scale and for long durations?
Area 10. Independent verification, safety/licensing,	$45M	Would independent experts accept the technical

Work package	5-year budget	Main question answered
cost audit & challenge testing		and economic case?
Central reserve	$100M	What additional tests, redesigns, supplier substitutions, or integration work become necessary as evidence emerges?
Total	$1.000B	$900M planned technical program + $100M sponsor-controlled reserve.

The key is that most planned technical money should eventually go into hardware that resolves economic uncertainties, not into papers or generic plasma research. The $100M reserve remains outside contractor baselines until a documented technical, supplier, safety, integration, or evidence need justifies release under the Common Framework.

How the programs fit together

The programs should not operate as separate research silos. They form a continuous loop.

Systems engineering (Area 1) defines the economically important uncertainties. **Plasma** (Area 2), **magnets** (Area 3), **liquid walls** (Area 4), **tritium** (Area 5), **power systems** (Area 6) and **RAMI** (Area 7) experimentally reduce those uncertainties. Manufacturing and construction determine whether the resulting machine can actually be mass-produced in China. The **integrated sector** (Area 9) tests whether the pieces work together. The **independent team** (Area 10) challenges the resulting conclusions.

Then the economic model is updated and the cycle repeats.

That is exactly the program logic implied by Program Requirements final deliverable: not merely a better stellarator concept, but a decision-grade, experimentally supported answer about whether mature Chinese stellarator fusion can beat coal.

At the end of five years, the $1B should therefore buy something much more specific than “fusion research.” It should buy enough engineering, experiments, Chinese cost information and integrated hardware to establish whether the following chain is credible:

Physics → Engineering → Net Power → Availability → Manufacturing → CAPEX → LCOE

If the answer is yes, the next program would logically move toward a **full pilot / first commercial demonstration**. If the answer is no, the program should identify exactly which requirements fail and how much improvement would be necessary to make the architecture economic.

Staged Spending and Reserve Release

The program should not obligate the full $1 billion at the start.

Technical areas receive staged authority through G0–G6. Reserve draws are separate decisions: they require a documented problem or additional evidence need, alternatives analysis, independent review, revised cost/schedule baseline, and an accountable owner.

An illustrative annual authorization profile is shown below. The reserve column is an upper envelope, not an entitlement; unused reserve may remain unspent.

Year	Planned technical	Reserve available up to	Maximum authorization	Primary emphasis
Year 1	$110M	$10M	$120M	Freeze references/economics; answer first kill questions; start discriminating tests.
Year 2	$160M	$20M	$180M	Subsystem prototypes, supplier quotes, challenger down-selection, early endurance work.
Year 3	$210M	$30M	$240M	Full-scale subsystem testing; architecture down-selection; integrated-sector fabrication begins.
Year 4	$250M	$25M	$275M	Integrated sector, long-duration testing, manufacturing and retrofit demonstrations.
Year 5	$170M	$15M	$185M	Endurance closure, independent validation, final China engineering/economic package.
Total	$900M	$100M	$1.000B	Reserve remains sponsor-controlled and may remain partly unspent.

The largest planned commitments occur only after the early program has identified the dominant technical and economic uncertainties.

First-year kill questions

I would spend roughly the first $110M of planned technical work trying to answer five questions quickly, before committing the large prototype budget or making substantial reserve draws:

1. After all parasitic loads, is there a credible path to roughly 1 GWe net?
2. After replacing ARIES/U.S. costs with Chinese NOAK costs, is there a credible path to ≤$2,000/kW-net?

3. Does the liquid-wall/blanket configuration survive serious MHD, corrosion, shielding and tritium analysis?
4. Does bottom-up RAMI show any plausible path toward 95% availability?
5. Can the D-T startup inventory support rapid fleet growth?

If the answers are bad, redesign P1 before committing hundreds of millions to hardware.

What I would require at the end of Year 5

Your requirements document already defines the overall deliverable very well: a decision-grade, experimentally supported answer about whether mature Chinese stellarator fusion can beat coal.

I would make the final package contain four major items:

A. China Greenfield P1	A complete design with independently audited CAPEX, net electrical output, technical availability, O&M, replacement cost, construction schedule, and LCOE, together with uncertainty ranges.
B. China Coal-Retrofit P1	A design tied to an actual/reference Chinese coal plant, with real interface engineering and incremental LCOE.
C. Experimental evidence package	Results from the magnets, liquid wall, blanket, tritium system, heat-transfer equipment, reliability testing and integrated sector.
D. Industrialization package	Factories, supply chain, construction plan, workforce, materials and capital required to build: 10 → 30 → 100 GWe/year.

And I would require one final table:

Metric	Requirement	Best demonstrated/model-supported result	Confidence
Greenfield CAPEX	≤$2,000/kW-net	?	?
Greenfield LCOE	<$43/MWh	?	?
Retrofit LCOE	< marginal coal	?	?
Net output	~1 GWe	?	?
Availability	≥95%	?	?
Construction	≤36 months	?	?

Tritium fleet scale	~100 GWe/yr path	?	?
Manufacturing scale	~100 GWe/yr path	?	?

That table should probably be the **single most important deliverable of the entire $1B program**.

The philosophy I would use is:

> Spend relatively little proving fusion is scientifically interesting; spend most of the $1B proving or disproving every assumption required for fusion electricity to cost less than coal.

That follows Program Requirements much more closely than a conventional fusion R&D portfolio.

Program Crosswalk

Prompt to AI: Show how the ten technical areas and central reserve map to the program questions, principal evidence, economic outputs, stage gates, and major contract packages.

Response From AI:

This crosswalk connects the controlled $1 billion budget to the decision architecture. Area budgets, contract-package references, experiments, and gates are mappings of the same work rather than additional funding. Every major result should flow into the controlled plant model so subsystem optimization cannot become detached from plant-gate economics.

Area / budget	Principal uncertainty	Main evidence	Economic output	Gate / contract packages
Area 1. $65M	Complete plant and China cost basis	Controlled designs; power/cost model; supplier evidence	CAPEX, LCOE, uncertainty	All gates; CP-1A/B
Area 2. $70M	Economically useful plasma operating point	Validated simulation + targeted stellarator experiments	Fusion power, auxiliary power, machine size	G1-G3; CP-2A/B
Area 3. $110M	Magnet installed cost and manufacturability	Conductor cells; patterned/current-surface tests; full magnet sector	CAPEX, cryogenic load, availability	G2-G4; CP-3A-C
Area 4. $155M	Liquid-wall feasibility, lifetime and parasitic power	Facilities A/B/C; heat/ MHD/ corrosion/ endurance/ fault tests	P-net, lifetime, outage hours, CAPEX/O&M	G2-G5; CP-4A-C
Area 5. $80M	Tritium closure and fleet growth	Neutronics; extraction/process tests; dynamic inventory model	P-net, availability, startup inventory, deployment rate	G2-G5; CP-5A/B
Area 6. $90M	Net saleable power and least-cost power cycle/retrofit	Integrated load model; BOP tests; reference coal-site engineering	P-net, retrofit CAPEX, LCOE	G1-G5; CP-6A/B
Area 7. $80M	95% availability feasibility	RAMI model; endurance; remote-maintenance and replacement trials	Availability, O&M, replacement cost	G2-G5; CP-7A/B
Area 8. $65M	Factory scale and construction duration	Production-cell data; supplier model; installation demonstrations	NOAK CAPEX, schedule, GWe/year	G3-G6; CP-8A/B
Area 9. $140M	Cross-system integration	Full-scale/near-full-scale integrated sector + endurance/fault campaigns	Validates coupled cost, power, RAMI assumptions	G4-G6; CP-9A/B
Area 10. $45M	Independent credibility	Replication, cost audit,	Confidence in final	All gates; CP-10A/B

Area / budget	Principal uncertainty	Main evidence	Economic output	Gate / contract packages
		safety/licensing review	YES/NO	
Reserve / $100M	Unknown failures and extra evidence needs	Only sponsor-approved incremental work	Protects mission without preallocating contingency	Draws only after documented review

Funding Source Requirements

We assume the Funding Source aims to maximize climate benefit per dollar spent, as opposed to maximize financial return on an investment. This is sometimes referred to as “climate money”, and it differs from the traditional investment portfolio. Sources of climate money include high-net-worth individuals, foundations, and governments.

The portfolio-wide Funding Source requirements are specified in the Moonshot Common Framework.

Stage Gates and Go/No-Go Criteria

Prompt to AI: Define the program-level gates that control progression from reference-design economics through critical-technology tests, integrated-sector work, industrialization evidence, and any later Phase II.

Response From AI:

Passing a gate authorizes only the next R&D increment. It does not authorize a commercial fusion plant. Gate criteria should be fixed before the associated major commitment, use the same controlled economic model, and receive independent review. Area-specific technical hold points may be stricter, but they are subordinate to the common G0-G6 architecture.

Gate / timing	Required evidence	Decision enabled	If not met
G0 - Mission baseline / Month 6	Coal comparators; accounting rules; greenfield/retrofit references; preliminary plant power balance; evidence taxonomy.	Freeze mission and common comparison boundary.	Correct baselines before technology teams optimize against inconsistent targets.
G1 - Reference-design feasibility / Month 12	Controlled designs; bottom-up China cost reconstruction; first five kill questions; initial challenger comparison.	Continue baseline and limited challengers into larger risk-reduction work.	Redesign reference architecture, change challenger mix, or terminate assumptions that cannot preserve an economic path.
G2 - Critical-technology feasibility / Months 18-24	Plasma, magnet, liquid-wall, tritium, net-power and RAMI evidence with quantified uncertainty.	Authorize large subsystem prototypes and down-select major architecture choices.	Repeat discriminating tests, substitute approach, reduce scope, or stop failed technology branch.
G3 - Architecture down-selection / ~Month 30	Integrated cost/power/RAMI/tritium model updated with measured component evidence and supplier data.	Select one or at most two controlled reference architectures for integrated hardware.	Do not commit major integrated-sector fabrication until a credible economic path exists.
G4 - Integrated-sector authorization / ~Month 36	Subsystem qualification; frozen interfaces; facility/test readiness; safety and independent review.	Fabricate and operate the integrated engineering sector.	Delay, redesign interfaces, substitute components, or scale the test to close readiness gaps.
G5 - Endurance & industrialization readiness / ~Month 48	Long-duration reliability, maintenance, manufacturing, cost, tritium and net-power evidence; multi-system fault recovery.	Freeze mature-China design basis and complete final validation package.	Run targeted endurance/reliability work or reduce claims; no automatic progression.
G6 - Final Economic Fusion	Audited CAPEX/LCOE, P-net, availability, tritium closure, construction plan,	YES: authorize Phase-II planning; NO: stop	No commercial/pilot commitment without a

Gate / timing	Required evidence	Decision enabled	If not met
Decision / Month 60	manufacturing scale, integrated-sector evidence and independent challenge.	architecture; TARGETED: fund only quantified remaining gaps.	decision-grade technical and economic case.

Program-level stop rule

Stop or redirect a workstream when progressively stronger evidence no longer supports a credible path to the applicable China economic target; when safety, tritium, maintainability, or manufacturing constraints cannot be resolved economically; when a challenger clearly dominates the baseline; or when the expected value of another test is lower than competing uses of the remaining funds.

Economic Traceability and Acceptance Criteria

Prompt to AI: Define common acceptance evidence so every major experiment shows how its result changes net power, CAPEX, availability, construction time, deployment scale, or LCOE.

Response From AI:

Every major workstream should deliver a short plant-effect record with the controlled baseline assumption, measured or independently supported result, uncertainty range, model variable changed, resulting delta in $/kW-net and/or $/MWh-net, and the associated gate recommendation. A successful component demonstration is not sufficient if the integrated economic effect is negative or unquantified.

Technical result	Required measured / controlled output	Economic translation	Advancement criterion
Plasma / confinement	Fusion power, auxiliary heating, impurity sensitivity, uncertainty	P-fusion -> P-net; machine size; heating CAPEX/O&M -> LCOE	Measured/validated range still supports the economic target after uncertainty.
Magnets	Conductor $/kA-m, structural mass, field accuracy, yield, cryogenic load, assembly time	Magnet CAPEX; P-recirc; outage/replacement -> $/kW-net and LCOE	Complete installed magnet system, not tape alone, remains economically credible.
Liquid wall / blanket	Heat flux, surface stability, MHD pressure drop, pumping/current power, corrosion, lifetime.	P-recirc; component size; replacement interval; outage hours; O&M	Measured lifetime/power-density benefit exceeds added complexity and cost.
Tritium	TBR margin, extraction rate, working inventory, permeation, process power, train reliability	CAPEX; P-recirc; availability; startup inventory; fleet doubling time	Plant closure and fleet-growth path remain feasible with realistic penetrations/losses.
Power conversion / retrofit	Gross efficiency, all auxiliary loads, steam/thermal interfaces, reuse/refurbishment scope	P-net; incremental retrofit CAPEX; LCOE	Net plant-gate electricity and retrofit reuse are supported by controlled engineering.
RAMI / maintenance	Failure rates, MTTR, planned outage hours, replacement time, redundancy benefit	Technical availability; replacement/O&M; lost MWh -> LCOE	Bottom-up outage budget supports the claimed availability range.
Manufacturing / construction	Yield, labor hours, cycle time, factory CAPEX, installation time, logistics	China NOAK CAPEX; construction financing; GWe/year ramp	Serial production and construction assumptions are supported by evidence.
Integrated sector	Interface fit, thermal/mechanical behavior, fault detection/recovery, maintenance access, actual prototype cost	Cross-checks coupled CAPEX, P-net, RAMI and schedule assumptions	Integrated evidence does not reveal unmodeled penalties that break the economic case.

Major Purchase Orders / Contract Packages

Prompt to AI:

Translate the $900 million planned technical program into a manageable set of major competitively awarded contract packages with explicit technical acceptance evidence.

Response From AI:

The packages below are a procurement map, not additional budget. Each amount is already inside its Area allocation. Detailed solicitations may split or combine these packages, but the program should preserve the same scope boundaries, controlled reference-design interfaces, milestone evidence, and independent acceptance logic.

ID / area	Major contract package	Budget	Principal deliverable	Technical acceptance evidence
CP-1A / A1	Controlled greenfield + retrofit reference designs	$30M	Configuration-controlled plant models, interfaces, power/tritium/RAMI accounting boundary.	Independent design review; traceable quantities; common reference assumptions frozen.
CP-1B / A1	China cost, supplier and economic-model package	$35M	Bottom-up China BOM, supplier evidence, lifecycle cost/LCOE model, uncertainty database.	Auditable source evidence and reproducible cost model; no blanket China discount.
CP-2A / A2	Stellarator physics / operating-point validation	$35M	P1-relevant confinement, stability, impurity and auxiliary-heating model.	Validated simulation and quantified P-fusion/P-aux uncertainty.
CP-2B / A2	Existing-machine experiments + plasma/wall compatibility	$35M	Targeted experiments resolving economically important plasma assumptions.	Reproducible data tied to controlled reference operating point.
CP-3A / A3	REBCO-like conductor production and cost cell	$35M	Automated representative production line/process cell with yield and cost data.	Throughput, Ic, yield, labor/material use and $/kA-m independently verified.
CP-3B / A3	Patterned current-surface / magnet manufacturing	$35M	Curved-substrate current paths, metrology, repair and repeatability demonstrations.	Field/placement tolerances, defect handling and manufacturing cycle time demonstrated.
CP-3C / A3	Magnet structure, joints, cryogenics and integrated sector	$40M	Representative high-field structural/current-surface magnet sector.	Load/field/fault/cryogenic tests plus actual prototype and installation cost.
CP-4A /	Facility A - high-field	$45M	P1-relevant thick flowing wall	Surface coverage/stability, MHD

ID / area	Major contract package	Budget	Principal deliverable	Technical acceptance evidence
A4	thick-wall MHD		and electromagnetic-control test.	pressure drop, pumping/current power and recovery quantified.
CP-4B / A4	Facility B - hot-wall endurance and heat extraction	$45M	Thousands-of-hours materials, pumps, valves, sensors and heat-extraction testing.	Corrosion, deposits, leakage, sensor drift, component degradation and maintenance data.
CP-4C / A4	Facility C - integrated wall/blanket sector + external campaigns	$65M	Largest practical wall/blanket/heat-extraction sector plus selected exposure/irradiation work.	Integrated fault/recovery, shielding, tritium-relevant and RAMI evidence.
CP-5A / A5	Breeding/neutronics and fleet-startup inventory	$40M	3-D TBR analysis, dynamic inventory and deployment-rate model.	Real penetrations/losses; quantified breeding margin and startup constraint.
CP-5B / A5	Tritium extraction, processing, permeation and reliability rigs	$40M	Representative process trains and retention/permeation experiments.	Extraction rate, inventory, process power, cleanup and redundancy performance.
CP-6A / A6	Net-power/BOP model + representative load tests	$45M	Complete P-gross/P-recirc/P-aux/P-net model anchored by equipment tests.	All significant loads traced; net output independently reviewed.
CP-6B / A6	China coal-retrofit + power-conversion engineering	$45M	Reference coal-site interface, reuse/refurbish/replace decisions and least-cost power cycle.	Incremental CAPEX/LCOE based on controlled site data and constructible interfaces.
CP-7A / A7	RAMI, outage budget and redundancy architecture	$40M	Bottom-up component RAMI model and economic redundancy design.	Failure/repair assumptions traceable; outage budget closes at claimed availability.
CP-7B / A7	Endurance, remote maintenance and rapid-replacement trials	$40M	Long-duration component stands and physical maintenance mockups.	Measured replacement/recovery time and degradation data support RAMI model.
CP-8A / A8	China factory model, supplier mapping and production cells	$30M	Factory counts, tooling, BOM flow, bottleneck and deployment model.	Supplier/process evidence supports 10/30/~100 GWe/year scenarios.
CP-8B / A8	Modular construction and installation demonstrations	$35M	Large-module logistics, alignment, connection, inspection and retrofit tie-in trials.	Measured installation hours, tolerances and construction critical path.
CP-9A / A9	Integrated-sector design, procurement and fabrication	$60M	Full-scale/near-full-scale integrated P1 engineering sector.	Readiness review; controlled interfaces; actual cost and assembly evidence.
CP-9B /	Integrated-sector	$80M	Long-duration operation,	Measured coupled performance,

ID / area	Major contract package	Budget	Principal deliverable	Technical acceptance evidence
A9	endurance, fault and recovery campaigns		repeated cycling, induced faults and maintenance.	recovery time, degradation and reliability growth.
CP-10A / A10	Independent technical/economic and China-cost audit	$20M	Independent challenge of cost, power, RAMI, manufacturing, construction and LCOE.	Reproducible findings with direct access to source evidence.
CP-10B / A10	Safety/licensing, replication and challenge tests	$25M	Safety/licensing concept plus independent replication of critical claims.	Required safety functions remain inside cost boundary; key claims independently reproduced.

Area 1. Systems Engineering, China Cost Model, and Controlled Reference Designs — $65M

Prompt to AI: Explain Project Area 1.

Response From AI:

This work package is the economic and engineering backbone of the Economic Fusion Moonshot. Its purpose is to prevent the individual technical teams from optimizing their own subsystems without understanding whether their work actually reduces the cost of electricity. The program should continuously maintain a complete, configuration-controlled model of the entire power plant, including physics, geometry, thermal systems, electrical systems, parasitic power, maintenance, construction, manufacturing, CAPEX, O&M, availability, and LCOE.

The program would initially freeze the two controlled reference designs defined in the requirements. Reference A would be the mature China greenfield P1-class stellarator. Reference B would use substantially the same fusion architecture but would be integrated with a representative approximately 1-GWe Chinese ultra-supercritical coal plant. The retrofit reference plant should be specific enough that engineers know steam temperatures and pressures, turbine-generator ratings, condenser characteristics, cooling equipment, feedwater system, transformers, switchyard arrangement, site geometry, remaining equipment life, and other important characteristics.

A major Year-1 task would be reconstructing the existing Volpe/ARIES cost model. Each major cost account should be decomposed into physical quantities rather than simply applying a percentage discount for China. If the magnet support system requires a certain number of tonnes of specialized steel, for example, the program would estimate Chinese material cost, forming, machining, welding, inspection, transportation and installation. Pumps and heat exchangers would be priced through Chinese suppliers. Civil works would be estimated using Chinese construction methods and labor. The result should distinguish the original Volpe/ARIES reference cost, a mature China greenfield cost, and a China coal-retrofit incremental cost, exactly as noted in Program Requirements.

The team should obtain budgetary quotations from actual suppliers wherever practical. Estimates would carry an evidence classification such as assumption, engineering calculation, supplier quotation, prototype measurement or operating data. Uncertainty ranges would be maintained for every economically significant parameter.

The plant model would also contain a complete power balance:

$$P_{net} = P_{gross} - P_{recirc} - P_{aux}$$

This calculation is central because a nominally 1-GWe power plant costing $2B is a $2/W plant only if approximately 1 GWe is actually exported. If realistic parasitic loads reduce output to 700 MWe, the same $2B machine becomes roughly $2.86/W-net and the LCOE changes substantially.

The systems team should therefore serve as an independent arbiter between technical groups. A liquid-wall group might propose a stronger pump that improves wall stability but consumes another 20 MWe. The systems model would determine whether the reliability improvement is worth the additional recirculating power. The magnet group might propose another 10 cm of shielding that dramatically extends magnet lifetime but increases machine radius and magnet cost. Again, the program should calculate the net LCOE effect rather than letting either subsystem optimize itself independently.

Every quarter or six months, this group should publish updated numbers for both reference designs: fusion thermal power, thermal power delivered to the turbine, gross electric power, recirculating power, net power, CAPEX, O&M, replacement cost, technical availability, construction time and LCOE.

The Year-5 deliverable would be the authoritative economic model underlying the final YES/NO decision. It should be independently auditable and detailed enough that outside engineering organizations can reproduce the conclusions.

The fundamental question this $65M answers is:

> When all components, parasitic loads, construction costs, financing, maintenance and Chinese manufacturing conditions are accounted for consistently, does the complete plant actually meet the economic-fusion requirements?

Area 2. Stellarator Plasma, Confinement, and Operating-Point Validation — $70M

Prompt to AI: Explain Project Area 2.

Response From AI:

The plasma program should be substantially different from a conventional fusion-physics program. The purpose is not to obtain the most scientifically impressive confinement result or the highest possible plasma gain. Its purpose is to determine whether the stellarator can operate at the plasma conditions required by the lowest-cost power plant.

The baseline architecture is a stellarator because it can operate in steady state without relying on the large externally driven toroidal plasma current associated with tokamak disruptions. The requirements nevertheless explicitly state that stellarator operation does not automatically imply ignition or infinite Q-plasma. The correct objective is the operating point that minimizes LCOE after auxiliary heating, recirculating power, magnet requirements, machine dimensions and plant costs have all been included.

The program would validate the P1 assumptions concerning confinement, density, temperature, beta, alpha heating, impurity behavior and plasma stability. It would use high-fidelity simulations together with experiments on existing stellarators where those experiments can answer P1-relevant questions. The work should be deliberately linked to the engineering design. A plasma configuration that produces slightly better confinement but requires substantially more complicated magnets may not be preferable.

Particular attention should be given to the interaction between plasma performance and the liquid-metal wall. The liquid surface may introduce impurities, alter recycling and affect edge conditions. If the plasma requires an unrealistically pristine boundary condition, the apparent economic advantage may disappear once an actual liquid wall is present.

Auxiliary heating is particularly important. The correct output from this program is not simply a predicted Q-plasma. Engineers need the electrical power required to establish and sustain the desired operating point, including conversion efficiency between grid electricity and plasma heating. That number feeds directly into P-recirc.

The program should also explore economic sensitivity to machine parameters. Instead of assuming P1 is exactly optimum, engineers should examine somewhat different magnetic fields, aspect ratios, radii and plasma conditions. For each variation, the systems team would calculate the resulting costs of magnets, blanket, structures, heating, buildings and balance of plant.

The program should use staged evidence. Early years can rely more heavily on validated simulations and existing machines. Later years should concentrate on experiments specifically chosen because they resolve economically important uncertainties.

A key milestone might state:

> At Year 2, the integrated physics model shall predict the reference plant's fusion power and required auxiliary heating with quantified uncertainty.

A later milestone might require that the relevant confinement and impurity assumptions have experimental support sufficient to constrain net-electric output within an agreed range.

The economic success criterion should be expressed in terms of plant impact. For example, rather than saying "confinement improved by 10%," a milestone should say that the new confinement evidence changes predicted net electrical output from 830 to 930 MWe, or reduces required machine radius, or lowers auxiliary power by 25 MWe.

The key question is:

> Can a P1-class stellarator sustain a plasma operating state that produces enough fusion power, with sufficiently low auxiliary power and sufficiently practical engineering constraints, to support <$43/MWh electricity?

Area 3. HTS Magnets, Magnet Structures, and Low-Cost Manufacturing — $110M

Prompt to AI: Explain Project Area 3.

Response From AI:

The magnet program should develop the complete installed magnetic-field system rather than treating REBCO tape cost as the sole problem.

A high-field stellarator requires superconducting conductor, structural support capable of reacting enormous magnetic forces, electrical joints, cryogenic systems, quench protection, power equipment, insulation, precision positioning and extensive quality assurance. Consequently, a cheap conductor does not necessarily produce a cheap magnet.

The program should establish a target cost for the complete magnet system derived from the $2,000/kW-net plant target. That target should then be decomposed into conductor, structural material, fabrication, cryostat, electrical equipment, cryogenics, factory labor, installation and inspection.

One major research track would develop much lower-cost manufacturing of REBCO-like conductor. The requirements specifically allow public or philanthropic "climate money" to finance production technology (e.g. automated machine that makes REBCO tape) that would later be made widely available, provided the subsidy is disclosed separately from plant economics. Automated deposition, patterning, inspection and joining equipment could therefore be legitimate moonshot investments.

A second major track would investigate simplified magnet manufacturing. The Volpe/Renaissance approach is interesting partly because patterned conductors or current surfaces may allow a different manufacturing architecture than traditional individually wound complex stellarator coils. The program should compare candidate approaches based on complete installed cost, field accuracy, reliability and scalability.

Structural cost deserves nearly as much attention as superconducting material. Magnetic forces may require hundreds or thousands of tonnes of high-strength structure. Optimization should therefore minimize total conductor-plus-structure cost rather than magnetic field alone.

The team should manufacture increasingly realistic components: conductor samples, joints, patterned panels or coil sections, structural assemblies and ultimately one or more representative full-scale magnet sectors. Tests should measure field accuracy, mechanical deformation, joint resistance, thermal performance, quench behavior and manufacturing reproducibility.

The manufacturing program must also ask whether the magnets can eventually be produced at climate scale. A solution requiring several years of hand assembly per plant is unacceptable even if technically excellent. Engineers should develop production cells, automated metrology, standardized interfaces and inspection procedures designed for repeated factory manufacture.

The mature manufacturing model should answer questions such as how many conductor lines and magnet factories are required for 10, 30 and ultimately ~100 GWe/year of plant production.

Important outputs include an installed magnet-system $/kW figure, manufacturing hours per plant, REBCO quantities, structural material quantities, cryogenic demand and the effect of magnet design on reactor size.

Specific Hardware Demonstrations

Area 3 should convert the magnet concept from a paper geometry into a measured manufacturing and structural system. Existing work on patterned wide conductors and magnetic-force minimization should be treated as starting evidence; the EFM program should determine the current paths, fields, forces, materials, manufacturing sequence, and installed cost required by the controlled P1 reference design.

- **Conductor-production demonstration:** develop and operate an automated REBCO-like conductor production line or representative process cell, measuring throughput, yield, critical current, material use, labor, equipment cost, and resulting $/kA-m and $/m.
- **Patterned-current-surface demonstration:** manufacture representative wide-conductor or directly patterned current paths on P1-relevant curved substrates; measure field accuracy, conductor placement tolerance, defects, repairability, and repeatability.
- **Magnet-force and structure demonstration:** establish the P1 field and force map, including on-axis field, peak field, current, local electromagnetic loads, stored energy, and fault loads; then build and mechanically test a representative structural sector using candidate steels or other structural materials.
- **Joints, cryogenics, and fault protection:** test electrical joints, current leads, cryogenic interfaces, insulation, quench/fault detection, and safe energy extraction at representative current and field conditions.
- **Full-scale or near-full-scale magnet-sector demonstration:** integrate conductor/current surface, supporting structure, cryogenic interfaces, metrology, and installation features so manufacturing time, alignment time, field error, structural deformation, and actual prototype cost can be measured.

Every magnet demonstration should report the quantities that drive the plant cost model: conductor length and current capacity, peak magnetic field, force per sector, structural mass, cryogenic load, manufacturing hours, yield/scrap rate, installation hours, maintenance strategy, and projected mature-China installed cost.

Area 3 Technical Hold Points

These area-specific hold points add technical readiness requirements within the program-level G0-G6 stage-gate architecture; they do not create a separate funding system.

- **Year 1:** freeze the reference magnet geometry and current paths; produce a field/force map, conductor specification, structural concept, manufacturing process flow, and bottom-up China cost target.
- **Year 2:** demonstrate representative conductor manufacture, patterned current paths, joints, metrology, and material/structural coupons with measured cost and performance.
- **Year 3:** build a large structural/current-surface segment and test mechanical loads, field quality, cryogenic interfaces, faults, and repeatable automated manufacture.
- **Year 4:** build and operate a full-scale or near-full-scale integrated magnet sector, including manufacturing and installation trials.
- **Year 5:** deliver a supplier-quoted, factory-scaled magnet production package for 10, 30, and approximately 100 GWe/year, with installed $/kW-net contribution and independently reviewed uncertainty.

The core question is:

> Can the complete high-field stellarator magnetic system—not merely the superconducting tape—be manufactured, installed and maintained cheaply enough for a ~$2,000/kW-net power plant?

Area 4. Liquid-Metal Wall, Blanket, Shielding, and Heat Extraction — $155M

Prompt to AI: Explain Project Area 4.

Response From AI:

This should be the largest individual technical program because the liquid-metal system carries several of the most important economic assumptions in the entire P1 concept.

The proposed liquid wall potentially solves a longstanding fusion problem by continually renewing the plasma-facing surface. If successful, it may support much greater surface heat loading than conventional solid walls, protect underlying structures from neutron damage, reduce scheduled replacement and contribute to the ≥95% availability objective. But Program Requirements correctly requires these advantages to be demonstrated rather than assumed.

Recent research provides a growing technical basis for this work. Puente et al. (2026) examine thick flowing liquid-metal walls as an integrated multifunctional plasma-facing system; Hu et al. (2024) review recent liquid-metal plasma-facing-component experiments; Wynne et al. (2025) validate the open-source FreeMHD solver against free-surface liquid-metal measurements; and the U.S. DOE SciDAC-5 program is developing coupled predictive models of plasma-liquid-metal interactions in plasma-facing components and breeding blankets. These studies, together with the liquid-wall papers listed later in this document, provide starting points for Area 4; the EFM program would extend the work toward the heat flux, magnetic field, neutron exposure, lifetime, reliability, tritium, and cost conditions required by a commercial fusion plant.

What Has Already Been Demonstrated

Experiment	What was demonstrated	What remains unresolved
Renaissance Skyfall 0 / 1a	Liquid-metal circulation loops and permanent-magnet pumps were built and operated.	Established practical loop and pumping experience; did not test a fusion-relevant thick wall.
Renaissance Skyfall 1b	A thick flowing GaInSn free surface was operated in a circular chamber at about 0.3 T. Injected current (about 2 kA in a representative run) produced Lorentz forces that changed & supported the flow.	Demonstrated electromagnetic control of a thick flowing layer. The attached surface was visibly nonuniform, and diagnostics were not yet sufficient to identify all instability modes.
PPPL LMX-U	Free-surface liquid-metal channel flow was tested in transverse magnetic fields with different electrically conducting wall conditions.	Showed that MHD drag/pileup and the electrical conductivity of the backing structure materially affect flow; the data have also been used to

		validate free-surface MHD models.
EAST FLiLi and LTX-beta	Flowing or deposited lithium was exposed to actual tokamak plasmas. EAST demonstrated improved flowing-lithium coverage and substantial heat removal; LTX-beta measured unusually low global recycling with lithium-coated walls.	Supports plasma compatibility and possible plasma-performance benefits, but these were not P1-type thick-wall demonstrations.
OLMAT liquid-tin tests	Tin-filled capillary-porous targets were exposed to intense particle/heat pulses, including tests near the facility's maximum heat flux.	Demonstrated strong heat-load tolerance and evidence of vapor shielding, while also exposing risks such as leakage, droplet/particle ejection, dryout and replenishment.
MIFILM experiment	A 3-D microstructured substrate greatly improved liquid-metal wetting and maintained stable spreading under a transverse magnetic field up to 1.6 T.	Shows that engineered backing surfaces can improve liquid-film stability; it remains a thin-film/substrate result rather than a thick wall.
Renaissance 850°C corrosion tests	Candidate steels were statically exposed to liquid tin at 850°C for 24 hours; all suffered significant corrosion, with 316L performing best of the tested steels.	Shows that corrosion/material compatibility is a first-order design issue and that long-duration dynamic testing is essential.

Selected Images and Drawings

The following images show the progression from demonstrated liquid-metal hardware to a next-generation lithium facility now under construction. They are included because they make clear which parts of the liquid-wall problem have already been tested and which parts remain to be demonstrated at P1 scale.

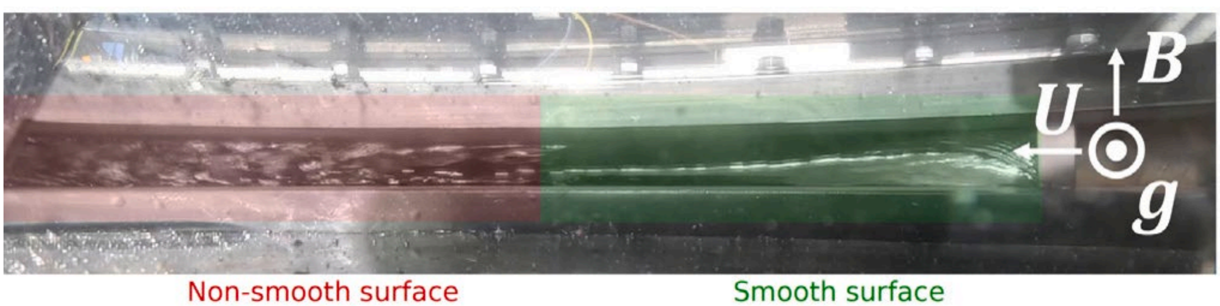

Figure 6: Renaissance Fusion Skyfall 1b experimental run. The GaInSn free surface shows a visibly non-smooth region and a smoother attached region under magnetic-field/current conditions. Source: Puente et al. (2026), Fig. 26. Open-source paper (page 15).

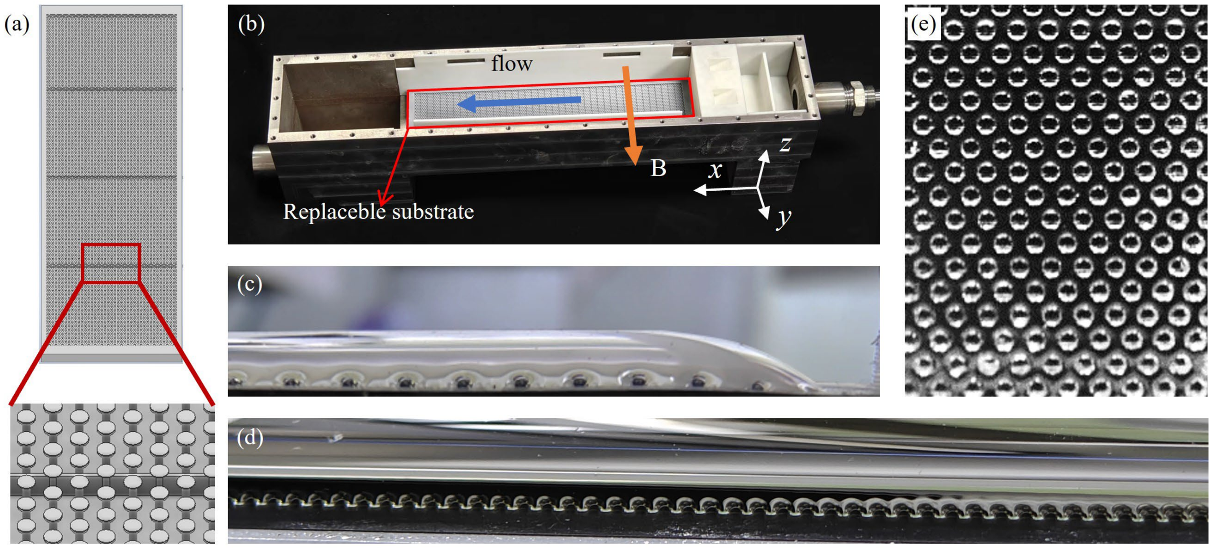

Figure 7: MIFILM experimental test section and microstructured substrate used to improve free-surface liquid-metal spreading in transverse magnetic fields up to 1.6 T. Source: Wang et al. (2025), Fig. 6. Open-source figure.

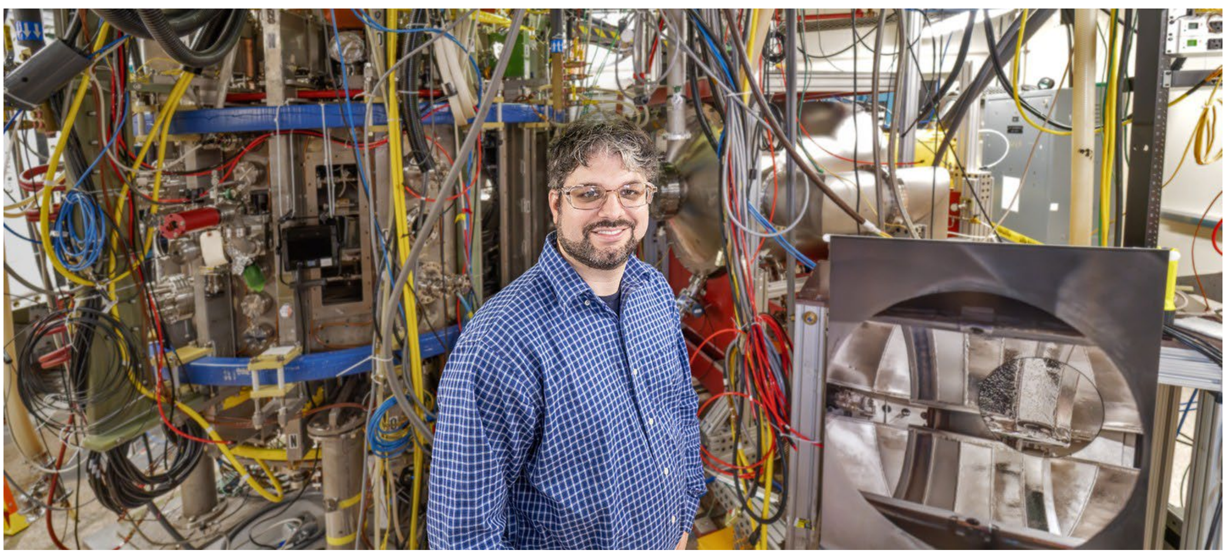

Figure 8: PPPL LTX-beta experimental fusion device; the inset photograph shows the device interior and a small pool of lithium. This is important evidence that liquid lithium can be used in an operating plasma-confinement experiment. Photo credit: Elle Starkman/PPPL Office of Communications. Open PPPL source page.

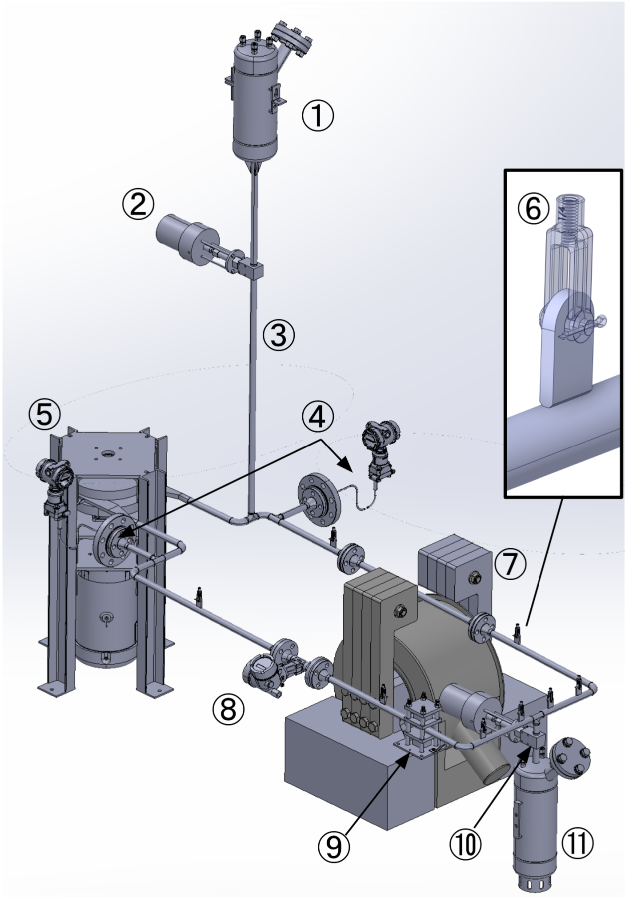

①	Expansion Tank
②	Throttle Valve (locked open)
③	3/4”-S40 SS pipe
④	Diaphragm Pressure Sensors
⑤	Moving Magnet Pump (MMP)
⑥	Welded-Tab Hanger Support
⑦	Low-Temp Superconducting Coil with Yoke ~ 0.5T
⑧	Vortex Flowmeter
⑨	MHD Flowmeter
⑩	Pneumatic Bellow-Sealed Valve
⑪	Dump Tank

Figure 9: CAD drawing of PPPL's LEAP Phase-I lithium loop, which is under construction. The design includes an expansion tank, moving-magnet pump, approximately 0.5-T superconducting coil, pressure sensors, flowmeters, valves, and dump tank. Source: Xu et al. (2026), Fig. 6. Open-source drawing.

Experiments Now Being Developed / Commissioned

Skyfall Hot is no longer best described as a purely not-yet-run facility. Renaissance Fusion reports that in October 2025 it successfully integrated and tested its Nicola100 HTS magnets with a 700°C flowing-liquid-metal loop. The August 2026 Puente paper describes the broader Skyfall Hot campaign as commissioning toward liquid tin at about 850°C in an approximately 1-T magnetic field, with additional results expected from the campaign. EFM should therefore use the latest operating data from Skyfall Hot rather than duplicate capabilities that are already being exercised.

PPPL's Lithium Experimental Application Platform (LEAP) is under construction as a modular, room-scale argon secondary-containment facility for a staged flowing-lithium program. Its Phase-I design includes lithium circulation, heating, diagnostics, magnetic-field exposure and a dump/containment architecture; later stages are intended to support more realistic open-surface and device-interface experiments. Figure 9 above shows the Phase-I CAD design.

Renaissance Fusion also describes later planned stages: Skyfall Helical Field would use pure lithium and seek full internal-wall coverage in a helical magnetic field; Skyfall Stellarator would progress toward Li-LiH, an approximately 10-T stellarator-like field, pebble management and a power-plant-module prototype. These later stages should be treated as planned experiments rather than completed evidence.

The implication for EFM is to build on, coordinate with and where practical use these existing programs rather than duplicate them. EFM capital should concentrate on the P1-specific experiments that are still missing.

The first research area is plasma-facing heat handling. Experiments need to establish whether the liquid surface remains stable under representative heat flux, magnetic fields and geometry. Measurements should include surface waves, evaporation, splashing, local dryout and impurity generation.

The second issue is magnetohydrodynamics. Conducting liquid metals moving through strong magnetic fields can experience substantial MHD forces and pressure drops. These forces may increase pumping power, complicate flow distribution and introduce new failure modes. Since pumping electricity comes directly out of P-net, MHD performance is an economic variable as well as a fluid-mechanics problem.

A third area is materials compatibility. Long-duration loops should expose candidate structural materials, piping, pumps and heat exchangers to realistic liquid chemistry and temperatures. Tests must run for sufficiently long periods to observe corrosion, mass transport, deposition and degradation that would not appear during short experiments.

A fourth task is neutron shielding and blanket performance. The liquid/solid blanket combination has to protect magnets and structures while remaining compact. Every additional centimeter of blanket and shield can increase machine radius, magnet dimensions and total plant cost. The optimization therefore needs to balance radiation lifetime against machine size.

A fifth task is heat extraction. The system must move gigawatts of thermal power from the fusion environment into a useful power-conversion system. Pumps, manifolds, piping, tanks and heat exchangers must be designed around low installed cost, low pumping electricity, high reliability and easy isolation.

The program should progress through increasingly integrated facilities. Early work can use small liquid-metal loops and high-heat-flux stands. Later facilities should combine strong magnetic fields with representative flow and temperature. Ultimately the program should construct a large sector including liquid-wall geometry, blanket material, structural backing, manifolds and heat extraction.

Long-duration operation is essential. Where practical, prototypes should operate continuously for thousands of hours to identify corrosion, pump degradation, deposits, leaks and sensor failures.

Fault testing should deliberately induce problems such as pump failure, flow blockage, leakage, loss of cooling and sensor faults. Engineers need to understand both whether these events are survivable and how long recovery takes.

The final deliverable should translate all of this into economically useful numbers: installed cost, required pumping power, component life, replacement interval, expected downtime, heat-transfer efficiency and uncertainty.

Specific Experimental Facilities and Demonstrations

EFM Area 4 should focus its major new capital on three P1-specific experimental systems. Smaller plasma-exposure, high-heat-flux, irradiation and tritium tests should use existing specialized facilities wherever practical.

- **Facility A - P1 high-field thick-wall MHD test:** reproduce the controlled P1 wall geometry and the most important P1-relevant magnetic-field directions and dimensionless MHD parameters. Begin with a safer surrogate metal and progress toward lithium where justified. Measure full surface coverage, layer thickness, waves, dry spots, splashing, electrical current paths, wall-conductivity effects, pressure drop, pumping/current power, and recovery after deliberately imposed disturbances. Diagnostics should be sufficient to quantify surface and subsurface instabilities rather than relying primarily on visual inspection. For an AI generated concept image, click here.

- **Facility B - P1 hot-wall endurance and heat-extraction test:** operate the down-selected liquid, structural materials, coatings, pumps, valves, sensors, manifolds and heat-extraction equipment at the reference temperature and representative heat loads for thousands of hours where practical. Measure corrosion and mass transfer, deposition, chemistry control, evaporation, leaks, pumping power, sensor drift, component degradation, shutdown/restart behavior, maintenance time and projected replacement interval. For an AI generated concept image, click here.

- **Facility C - P1 integrated wall/blanket sector:** combine liquid-wall geometry, backing structure, blanket/shield, manifolds, pumps, heat extraction, sensors, isolation, leak management and maintenance access in the largest practical P1 sector. Perform repeated startup/shutdown cycles and deliberately induce pump loss, partial blockage, loss of current/control, loss of cooling, leaks and sensor failures. Measure detection, isolation, safe response and recovery time as direct inputs to the RAMI and LCOE models. For an AI generated concept image, click here.

- **Supporting external experiments:** use OLMAT-, EAST-, LTX-, linear-plasma- or equivalent facilities for plasma/high-heat-flux exposure; use the best available neutron/irradiation facilities to anchor damage and shielding models; and coordinate tritium extraction, retention and permeation work with Area 5. Area 4 should not build a dedicated fusion-neutron source unless a later stage-gate review shows that no existing or planned facility can provide the required evidence.

The economic outputs must be explicit: MW/m² heat capability, MHD pressure drop, pump/current MWe, wall/blanket installed cost, backing-structure dose/lifetime, tritium inventory contribution, heat-transfer efficiency, expected replacement interval, outage hours, and the resulting change in plant $/kW-net and $/MWh.

Area 4 Technical Hold Points

These area-specific hold points add technical readiness requirements within the program-level G0-G6 stage-gate architecture; they do not create a separate funding system.

- **Year 1:** freeze the P1 liquid-wall/blanket geometry, material candidates, heat and neutron loads, MHD similarity targets, failure cases and economic performance requirements; complete Facility A design and use existing facilities for inexpensive discriminating tests.
- **Year 2:** build and operate Facility A; demonstrate controlled thick-wall flow at P1-relevant MHD conditions and quantify surface stability, wall-conductivity effects, pressure drop, pumping/current power and disturbance recovery.
- **Year 3:** build and begin long-duration operation of Facility B; down-select liquid, structural materials, coatings, pumps, heat-extraction equipment and flow/control architecture using measured corrosion, thermal and reliability data.
- **Year 4:** construct and begin endurance operation of Facility C, integrating the liquid wall, blanket/shield, heat extraction, instrumentation, isolation and maintenance systems; perform deliberate fault and recovery tests.
- **Year 5:** complete endurance testing and deliver experimentally anchored lifetime, parasitic-power, maintenance, shielding, tritium and installed-cost models sufficient for the final greenfield and retrofit LCOE decision.

The main question is:

> Does a real liquid-metal wall/blanket/heat-extraction system provide enough lifetime, power density and availability benefit to justify its complexity and enable coal-competitive electricity?

Area 5. Tritium Breeding, Processing, and Fleet Startup Inventory — $80M

Prompt to AI: Explain Project Area 5.

Response From AI:

The tritium program must establish both **single-plant fuel-cycle closure** and **rapid fleet growth**.

For a D-T plant, tritium cannot simply be treated as an externally available commercial fuel. The reactor has to breed sufficient tritium, extract it rapidly, process it, control inventories and return it to the plasma. Every loss or delay increases the amount of tritium tied up in the plant.

The first task is detailed neutronics. The actual 3-D geometry should include realistic penetrations, ports, piping, diagnostics and structural materials. Breeding-margin calculations that assume an uninterrupted ideal blanket are not adequate.

The second task is extraction and processing. Experiments should determine extraction rates from the proposed liquid-metal system, tritium retention, permeation and cleanup performance. Processing equipment must be sized and its capital cost and electrical consumption incorporated into the plant model.

The third task is inventory reduction. Tritium residing in blanket liquids, piping, processing systems, storage and fueling systems is unavailable for starting another reactor. The program should therefore optimize the entire system for both safety and low working inventory.

The fourth task is reliability. A tritium-processing train that forces the entire plant offline whenever one pump or separator fails will conflict with the 95% availability target. Parallel processing trains, isolation, spare capacity and online maintenance should therefore be examined.

The fifth issue is fleet startup. This is unusual but essential to the Economic Fusion Moonshot. Even a reactor with a tritium breeding ratio above unity may require years to generate enough surplus tritium to start additional machines. Program Requirements therefore requires explicit modeling of fleet growth into tens and ultimately about 100 GWe/year.

The program should construct a dynamic tritium model covering perhaps 20–30 years. Inputs would include startup inventory per plant, doubling time, breeding ratio, extraction efficiency, decay, reserves and commissioning schedules. The model should determine whether tritium availability constrains the planned manufacturing rate.

If it does, the program should investigate lower-inventory startup procedures, more aggressive breeding, temporary external supplies, staged commissioning or alternative fuel-cycle concepts.

By Year 5, the design should have an auditable tritium mass balance just as rigorous as its electrical power balance.

The key question is:

> Can the plant breed and process enough tritium to operate reliably—and can a rapidly expanding Chinese fleet obtain enough startup inventory to grow toward climate-relevant deployment rates?

Area 6. Net Power, Balance of Plant, Power Conversion, and Coal-Retrofit Integration — $90M

Prompt to AI: Explain Project Area 6.

Response From AI:

This program attacks one of the largest uncertainties in the existing P1 economic picture: how much electricity is actually available for sale.

The requirements explicitly distinguish P-fusion, P-thermal, P-gross, P-recirc, P-aux, and P-net. They also identify the optimistic approximately 1.01-GWe case and the roughly 700-MWe scenario that includes up to about 320 MWe of additional loads as an uncertainty the program must resolve.

Approximately half of the work package should therefore build a detailed power-consumption model backed progressively by component measurements.

Every significant electrical load must be included: heating, magnet systems, cryogenics, tritium processing, primary liquid-metal pumping, secondary pumping, vacuum, cooling, controls and other station service.

Large pumps and heat-transfer equipment should be tested at representative operating points or scaled using validated engineering models. Cryogenic and tritium-system power should similarly be determined from actual equipment designs rather than generic percentages.

The second half of this program should study power conversion and the coal-retrofit interface.

For the greenfield plant, engineers should compare conventional steam cycles and advanced cycles based on complete LCOE. Higher thermodynamic efficiency is desirable only when the added equipment cost and complexity are economically justified.

For retrofit, the situation is more interesting. A fusion heat source may be able to supply an existing coal turbine-generator, avoiding the purchase of a new turbine, generator, condenser, cooling tower, transformers and much of the electrical island. But this requires matching temperature, pressure, flow and reheat conditions.

The program should therefore select a reference Chinese ultra-supercritical coal plant and design the actual thermal interface.

Each existing system would be classified:

> reuse unchanged / refurbish / modify / replace.

The engineering team should evaluate whether preserving the original steam conditions is worthwhile or whether turbine modification results in lower total cost.

The work should include dynamic operation, startup, shutdown and load changes, not merely steady-state thermodynamics.

An important deliverable would be a retrofit interface package detailed enough that a Chinese engineering organization could estimate construction quantities and cost.

The decisive question is:

> After every internal electrical load is included and the least-cost power cycle is selected, how many net megawatts actually cross the plant gate—and what does the existing coal infrastructure allow us to avoid buying?

Area 7. RAMI, Maintenance, Redundancy, and Long-Life Components — $80M

Prompt to AI: Explain Project Area 7.

Response From AI:

The 95% technical-availability requirement makes reliability engineering a first-order economic activity rather than an afterthought.

A 95% availability plant can tolerate only about 438 unavailable hours per year (8760 × 5%). That allowance must cover scheduled maintenance as well as component failures. A fusion machine containing thousands of pumps, valves, sensors, power electronics, vacuum components, tritium components, cryogenic equipment and specialized reactor hardware therefore needs maintenance philosophy designed into it from the beginning.

The program should create a bottom-up RAMI model (i.e. Reliability, Availability, Maintainability and Inspectability) listing every significant component, failure mode, failure rate, repair time, inspection requirement and replacement interval.

For each important system, engineers should ask whether redundancy is cheaper than downtime. Examples include N+1 pumps, parallel tritium-processing trains, redundant power converters, redundant sensors and bypassable heat exchangers.

Isolation is equally important. A failed secondary pump should not require draining the entire fusion blanket. Valves, segmented loops and bypasses should permit maintenance while most of the plant remains operating.

Component modularity should be designed around replacement time. If a pump normally takes three days to disconnect because of inaccessible piping, redesigning the interface to allow a four-hour module replacement may be economically worthwhile even if the pump assembly costs more.

Long-duration endurance tests are therefore essential. Pumps, valves, seals, electrical equipment, vacuum systems and sensors should run for thousands of hours under representative temperature, magnetic-field, radiation or fluid conditions where appropriate.

Remote maintenance should be evaluated wherever radiation prevents human access. Maintenance concepts should be tested physically using mockups rather than relying only on computer animations.

The program should create an explicit annual outage budget. Each subsystem receives an allowed number of hours of scheduled and unscheduled downtime, and the sum must remain compatible with the ≥ 95% plant target.

If a subsystem cannot meet its outage budget, engineers can increase redundancy, improve durability, redesign interfaces or accept a lower plant availability and calculate the resulting LCOE.

This work package should also examine the existing coal-plant equipment in the retrofit case. Reusing an old turbine might save hundreds of millions of dollars initially but become uneconomic if it materially reduces future availability.

The main question is:

> Can all planned and unplanned maintenance actually fit inside the outage budget required for ~95% technical availability, and what equipment architecture is necessary to achieve it economically?

Area 8. China Factory Production, Automated Manufacturing, and Construction Engineering — $65M

Prompt to AI: Explain Project Area 8.

Response From AI:

The Economic Fusion Moonshot is not successful merely because one reactor can be built cheaply. Program Requirements explicitly requires a credible route first to tens and eventually approximately 100 GWe per year.

This program therefore designs the production system for the reactor.

Engineers should take the plant bill of materials and determine which components should come from existing Chinese industries, which require new dedicated factories, and which require specially developed manufacturing machines.

Major categories include superconducting conductor, magnet structures, blanket modules, liquid-metal equipment, heat exchangers, pumps, vacuum equipment, tritium skids, structural modules, piping systems and instrumentation.

For each component, the team should define an industrial manufacturing process including raw materials, machine tools, cycle time, labor, inspection, transportation and installation.

The program should then model several deployment levels: perhaps 10 GWe/year, 30 GWe/year and 100 GWe/year.

For each rate it should calculate:

- number of factories;
- factory floor space;
- tooling;
- factory CAPEX;
- annual raw materials;
- labor;
- port and rail requirements;
- manufacturing bottlenecks;
- lead times.

Automated manufacturing deserves particular attention. If the $1B program can develop a machine that reduces a critical component's manufacturing cost dramatically, the machine itself can be replicated and deployed throughout the supply chain.

Construction engineering is equally important. The plant should be divided into the largest practical factory-finished modules and skids, transported to site and assembled with standardized interfaces.

For greenfield construction, the mature target is ≤36 months from major construction authorization to commercial operation, with approximately 24 months preferred for major site work.

For retrofit, the objective is even more specific: install as much of the fusion plant as possible while the coal unit continues operating, then perform final steam, electrical and control tie-ins during a short-planned outage.

The program should create detailed construction simulations, crane plans, module logistics, workforce schedules and critical-path analyses.

Representative installation demonstrations should be performed using actual large modules or mockups. This provides real data on alignment time, welding, bolting, pipe connections and inspection.

The fundamental question is:

> If the design works economically on paper, can China actually manufacture and install enough identical plants fast enough to matter for climate change?

Area 9. Full-Scale Integrated Sector Prototypes and Endurance Testing — $140M

Prompt to AI: Explain Project Area 9.

Response From AI:

This is the principal hardware-integration program and probably the most visible physical deliverable of the moonshot.

The $1B fund should not attempt to construct a complete 1-GWe fusion plant. Instead, it should build enough full-scale or near-full-scale integrated hardware to determine whether the important engineering assumptions survive contact with reality.

The integrated sector should combine as many representative systems as practical: magnet and support structure, vacuum boundary, liquid-metal wall, blanket/shielding, heat-extraction manifolds, piping, pumps, sensors, tritium-relevant systems, maintenance access and remote handling.

Different effects may need to be reproduced in separate facilities. High heat flux could be supplied electrically or with beams; strong magnetic fields could be provided by representative magnets; neutron damage may require irradiation facilities rather than a fusion neutron source.

What matters is that interfaces are tested.

Fusion concepts often look attractive when individual subsystems are analyzed independently. Integration reveals the hidden problems: a pipe cannot be routed because a magnet is in the way; a maintenance tool cannot reach a valve; shielding added to protect the magnet prevents adequate heat extraction; a sensor fails in the magnetic environment; a pump requires more electricity than expected.

The sector should therefore undergo repeated thermal cycling, startup and shutdown, pump transients, simulated faults and maintenance operations.

Long-duration testing is particularly important. The objective is not a record-setting few minutes of operation but data on thousands of hours of service, because the economic targets depend heavily on corrosion, component life and maintenance frequency.

Engineers should deliberately introduce failures: pump loss, valve malfunction, partial blockage, leakage, sensor failure, loss of cooling and other realistic conditions. They should measure detection time, safe response and recovery time.

The prototype should also support construction and maintenance studies. Teams can measure how many worker-hours are required to assemble a sector, make connections, inspect it and replace damaged modules.

Actual prototype costs are valuable because they allow the cost team to compare engineering-model estimates with real manufacturing experience.

The integrated sector does not need to prove the entire fusion power plant. Its purpose is to eliminate the largest uncertainty between subsystem research and a future pilot plant.

By Year 5 it should provide measured values for pump power, thermal performance, leakage, corrosion, assembly time, maintenance time, structural behavior and other inputs to the economic model.

The main question is:

> When the major P1 engineering systems are physically combined at realistic scale, do they operate, survive, maintain and cost approximately what the economic model assumes?

Area 10. Independent Verification, Safety/Licensing, Cost Audit, and Challenge Testing — $45M

Prompt to AI: Explain Project Area 10.

Response From AI:

The planned $45M should create an independent challenge function rather than another technical development group.

The people who design an ambitious engineering system naturally become invested in making it succeed. That makes independent scrutiny particularly important when the headline result may be a claim that fusion can produce electricity for $35–43/MWh.

A separate organization should audit the China cost model, supplier quotations, manufacturing assumptions, power balance, reliability model, tritium analysis, construction schedule and LCOE.

It should have access to the underlying data rather than merely reviewing presentation slides.

Independent engineers might find, for example, that the main design estimates a heat exchanger at $80M but comparable Chinese equipment and installation indicate $150M. Or they might conclude that a purported 95% availability number overlooks a mandatory annual inspection outage.

The purpose is not to make the project pessimistic. It is to ensure the final YES or NO result is credible to policymakers, utilities, investors and other fusion researchers.

Part of this budget should address safety and licensing concepts. The economic model cannot reach its target by simply omitting functions required for tritium control, radioactive-material management, security, quality assurance, waste handling or decommissioning. The requirements explicitly place those items inside the economic boundary.

Safety-by-design could nevertheless reduce cost. Lower inventories, passive decay-heat removal, smaller radioactive zones and modular confinement systems may simplify both licensing and construction.

An independent team should also assess whether Chinese localization assumptions are defensible. Cost reductions should arise from identified materials, factories, processes, labor and financing rather than an unexplained "China factor."

Unexpected failures and discoveries should be funded from the separate $100M central reserve rather than from an embedded Area 10 contingency. Area 10 should document the technical issue, alternatives, and independent evidence needed for any proposed reserve draw, but the sponsor retains release authority under the Common Framework.

A reasonable internal split would be roughly $15M for independent engineering/economic audit, $10M for Chinese cost and supplier validation, $10M for safety/licensing work, and $10M for replication or challenge tests. These planned amounts total $45M; contingency is held separately in the central reserve.

The central question is:

> Would independent experts looking at the same engineering data conclude that the claimed CAPEX, net power, availability and LCOE are credible?

Central Reserve — $100M

Prompt to AI:

Define how the sponsor-controlled $100 million reserve should be used without turning it into preallocated contractor contingency.

Response From AI:

The reserve is not an eleventh technical work package and is not automatically spent. It remains under Funding Source control and may be released only when evidence from the staged program identifies a mission-critical uncertainty, failure, supplier problem, or additional validation need that is more valuable than competing uses of the remaining funds.

Reserve use	Example trigger	Required evidence before release
Critical technology redesign	Liquid-wall instability, magnet fault, tritium-processing failure, unanticipated parasitic load.	Failure data, alternatives, revised experiment, independent technical review, cost/schedule delta.
Supplier / fabrication recovery	Supplier failure, low yield, long-lead replacement, prototype rework.	Procurement record, replacement options, updated delivery risk and accountable owner.
Additional endurance / replication	Degradation appears late, independent result conflicts, uncertainty remains decision-critical.	Why another campaign changes the G5/G6 decision; frozen acceptance criterion.
Integration / safety closure	Integrated-sector interface, fault-recovery, safety or licensing gap emerges.	Deficiency report, corrective design, independent readiness/acceptance plan.
Final decision closure	Limited extra work could resolve a narrow uncertainty before G6.	Expected value of information exceeds competing uses of reserve; no automatic draw.

Every reserve draw should identify the triggering evidence, alternatives considered, independent review, amount requested, revised baseline, accountable owner, technical acceptance evidence, and effect on the next stage-gate decision. Unused reserve may remain unspent.

Program Governance, Contracting, and Open Research

Portfolio-wide Managing Organization selection, fiduciary governance, procurement, milestone-payment logic, open-source/data rights, financial controls, conflicts, independent review, management oversight, audit, reserve/change control, remedies, transition, and master agreement terms are specified in the Moonshot Common Framework. This chapter states only the additional management and contracting requirements unique to the Economic Fusion Moonshot.

Fusion-specific management requirements

- Maintain one accountable Program Director and systems-integration function controlling the greenfield and coal-retrofit reference designs, requirements database, configuration, integrated power balance, RAMI model, tritium balance, China cost model, interfaces, and final YES/NO recommendation.
- Retain an initiative-specific Scientific and Engineering Review Board spanning stellarator physics, HTS magnets and structures, liquid-metal MHD/materials, blankets and shielding, tritium, heat extraction, power conversion, RAMI, remote maintenance, manufacturing, and construction.
- Use independent economic/cost and safety/licensing reviewers to challenge supplier quotations, Chinese localization, net-power assumptions, availability, tritium closure, radioactive-material controls, quality assurance, construction schedule, and LCOE before major claims or scale-up decisions.
- Tie technical stage gates to the economic mission: large prototype, facility, or integrated-sector commitments proceed only when the evidence still supports a credible path to the applicable greenfield or retrofit economic target.
- The $100M central reserve remains under Funding Source control and follows the Common Framework reserve/change-control discipline; a draw must document the failure or uncertainty, alternatives, independent review, revised baseline, accountable owner, amount, and acceptance evidence.

Program-specific contracting and handoff additions

- Every major fusion award should identify the controlled reference-design interface it supports and report the evidence class of important inputs: assumption, simulation, engineering calculation, supplier quotation, component test, integrated test, or operating data.
- Contracts for magnet, liquid-metal, tritium, heat-transfer, power-conversion, manufacturing, and integrated-sector work should deliver editable source files and the quantities that drive plant eco-

nomics, including installed cost, parasitic power, lifetime, outage hours, manufacturing time, yield, replacement interval, and uncertainty.

- Background IP or controlled information that could prevent independent validation, competitive manufacture, or transfer to a successor team must be disclosed before award and handled under the common framework rather than discovered at closeout.

Scientific and Engineering Advisory Panel

Prompt to AI:

Define the fusion-specific scientific and engineering expertise needed to review the EFM technical program while relying on the Common Framework for generic governance and management oversight.

Response From AI:

The advisory panel should be small enough to make decisions but broad enough to challenge the coupled physics, engineering, manufacturing, safety, and economic assumptions. Members may recommend experiments and redesigns, but final spending authority remains with the program governance structure defined in the Common Framework.

Panel seat / expertise	Principal review responsibility	Independence focus
Stellarator physics	Confinement, operating point, plasma stability, auxiliary heating and uncertainty.	Challenge physics assumptions that drive P-fusion and P-recirc.
Fusion systems engineering	Reference architecture, interfaces, integrated power/tritium/RAMI models.	Prevent subsystem optimization from bypassing plant-level requirements.
HTS magnets / structures	Conductor, current surfaces/coils, magnetic loads, joints, cryogenics, quench/fault protection.	Review complete installed magnet cost and manufacturability.
Liquid-metal MHD / plasma-facing systems	Free-surface stability, electromagnetic control, pumping power, heat flux and failure modes.	Independent assessment of Area 4 feasibility and scaling.
Materials / blanket / shielding	Corrosion, activation, neutron damage, shielding, heat extraction, component lifetime.	Challenge lifetime and compactness assumptions.
Tritium / fuel cycle	Breeding, extraction, permeation, accountability, inventory and fleet startup.	Verify single-plant closure and fleet-growth constraints.
Power conversion / thermal systems	Heat transfer, turbines/cycles, pumps, parasitic loads and coal-retrofit interfaces.	Verify P-net and least-cost power-cycle choices.
RAMI / remote maintenance	Failure rates, redundancy, maintenance architecture, remote handling and outage budget.	Challenge the >=95% availability claim.

Panel seat / expertise	Principal review responsibility	Independence focus
China manufacturing / construction	Supplier base, factories, automation, modular construction, logistics and schedule.	Test NOAK cost and 10/30/~100 GWe/year assumptions.
Fusion economics / cost engineering	CAPEX/LCOE accounting, financing, uncertainty, comparison with coal.	Ensure common boundaries and evidence classes.
Safety / licensing / radioactive materials	Tritium confinement, activation, waste, QA, security and licensing concept.	Ensure required safety functions remain inside the economic boundary.
Independent chair / integration	Set review agenda, reconcile cross-disciplinary findings, issue gate recommendations.	No line responsibility for delivery teams whose work is being judged.

Top Fusion Cost and Economic Modeling Papers

The following papers explore methodologies for estimating fusion electricity production costs.

Fusion Costing Methodology

Whyte (2026), "Criteria for the economic viability of fusion power plants"

Provides explicit economic criteria for judging whether a fusion power plant can compete with alternatives; useful for defining the program's CAPEX, availability, and LCOE tests.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/2026_Whyte_Criteria_for_the_economic_viability_of_fusion_power_plants.pdf

Woodruff (2026), "A costing framework for fusion power plants"

Proposes a structured costing framework for fusion plants and helps make accounting boundaries and cost assumptions more comparable across designs.

https://arxiv.org/pdf/2601.21724

Woodruff (2026), "Extension of the fusion power plant costing standard"

Extends fusion costing methods and is useful for building a consistent, auditable plant-level cost model rather than comparing isolated subsystem prices.

https://arxiv.org/pdf/2602.19389

Chapman (2025), "Developing Integrated Cost Models for Fusion Power Plants"

Discusses integrated fusion cost modeling and supports the program requirement to connect physics, engineering, construction, maintenance, and financing to LCOE.

https://doi.org/10.1007/s10894-025-00515-1

Electricity-Market Value / Maintenance

Schwartz (2022), “The value of fusion energy to a decarbonized United States electric grid”

Examines the system value of fusion in a decarbonized grid and helps distinguish plant-gate LCOE from broader electricity-system value.

https://arxiv.org/pdf/2209.09373

Schwartz (2024), “Valuing maintenance strategies for fusion plants as part of a future electricity grid”

Quantifies how fusion maintenance strategy and outage timing affect economic value, directly supporting the RAMI and availability work.

https://arxiv.org/pdf/2405.01514

Maintenance Engineering / Remote Replacement

Bachmann (2022), “Conceptual study of the remote maintenance of the DEMO breeding blanket”

Provides a detailed remote-maintenance concept for activated fusion components and is relevant to replacement time, access, tooling, and outage-duration assumptions.

https://doi.org/10.1016/j.fusengdes.2022.113077

Commercialization Roadmap / R&D Gap Coverage

U.S. Department of Energy (2026), “Fusion Science and Technology Roadmap”

Summarizes U.S. fusion science and technology priorities and provides a useful external roadmap against which the EFM program can check major R&D gaps.

https://www.energy.gov/fusion/articles/fusion-st-roadmap

Top Tokamak Papers

"ARC" is a Tokamak reactor that is currently being built by Commonwealth Fusion Systems. For details, see https://en.wikipedia.org/wiki/ARC_fusion_reactor.

For a summary of parameters that describe a fusion machine, see fusion_reactor_design_parameters.pdf

Original ARC Power-Plant Design (2015)

Sorbom (2015), "ARC: A compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets"

Defines the original ARC high-field tokamak power-plant concept and provides a major technical and economic comparator for the stellarator baseline.

https://arxiv.org/pdf/1409.3540

ARC Physics Basis Papers (2026)

Hillesheim (2026), "Overview of the physics basis for the ARC fusion power plant"

Summarizes the updated ARC physics basis and provides a current tokamak benchmark for confinement, operating point, and plant-level physics assumptions.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/overview-of-the-physics-basis-for-the-arc-fusion-power-plant.pdf

Sweeney (2026), "ARC disruption physics and strategy"

Addresses disruption risks and mitigation in ARC, illustrating a major tokamak-specific engineering burden that the stellarator baseline seeks to avoid.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/arc-disruption-physics-and-strategy.pdf

Leuthold (2026), "ARC physics basis–magnetohydrodynamics"

Describes ARC MHD stability and operating constraints and provides a tokamak comparison for the EFM plasma-validation work.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/arc-physics-basis-magnetohydrodynamics.pdf

Eich (2026), “Power and particle exhaust for the ARC fusion power plant”

Analyzes ARC power and particle exhaust, offering a useful comparison with the EFM liquid-metal wall and heat-exhaust strategy.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/power-and-particle-exhaust-for-the-arc-fusion-power-plant.pdf

Howard (2026), “Performance and transport in the ARC tokamak”

Evaluates ARC plasma performance and transport and provides a contemporary benchmark for power-producing magnetic-confinement assumptions.

https://www.ma2life.org/g/fusion/fusion_cfs_arc_2026/performance-and-transport-in-the-arc-tokamak.pdf

SPARC & HTS Magnet Papers (2020-2024)

“SPARC” is a Tokamak reactor that is currently being built by Commonwealth Fusion Systems. For details, see: https://en.wikipedia.org/wiki/SPARC_(tokamak).

Creely (2020), “Overview of the SPARC tokamak”

Defines the SPARC machine and physics mission, providing context for high-field compact tokamak development and its relationship to ARC.

https://www.ma2life.org/g/fusion/fusion_cfs_sparc_2024/Sparc_Overview-of-the-sparc-tokamak.pdf

Vieira (2024), “Design, Fabrication, and Assembly of the SPARC Toroidal Field Model Coil”

Documents fabrication and assembly of a large high-field HTS model coil and is directly relevant to manufacturability, joints, structure, metrology, and quality control.

https://www.ma2life.org/g/fusion/fusion_cfs_sparc_2024/Sparc_Design_Fabrication_and_Assembly_of_the_SPARC_Toroidal_Field_Model_Coil.pdf

Whyte (2023), “Experimental Assessment and Model Validation of the SPARC Toroidal Field Model Coil”

Reports experimental validation of the SPARC model coil and provides evidence on high-current HTS magnet performance under representative mechanical and electromagnetic loads.

https://www.ma2life.org/g/fusion/fusion_cfs_sparc_2024/Sparc_Magnet_2024_Experimental_assessment_and_model_validation_of_th.pdf

Chen (2023), “Modeling of HTS high-current stacked conductors with defective tapes in different locations”

Examines defect effects in stacked HTS conductors and is relevant to conductor yield, quality assurance, fault tolerance, and manufacturing cost.

https://www.ma2life.org/g/fusion/fusion_cfs_sparc_2024/Sparc_Modeling_of_HTS_high-current_stacked_conductors_with_defective_tapes_in_different_locations.pdf

Hartwig (2023), “The SPARC Toroidal Field Model Coil Program”

Summarizes the model-coil development program and provides a useful precedent for staged design, fabrication, testing, and validation of fusion-scale HTS magnets.

https://www.ma2life.org/g/fusion/fusion_cfs_sparc_2024/Sparc_Toroidal_Field_Model_Coil_Program.pdf

Tritium Fuel Cycle

Meschini (2023), “Modeling and analysis of the tritium fuel cycle for ARC- and STEP-class D-T fusion power plants”

Models tritium inventories and fuel-cycle behavior for D-T plants and directly informs EFM requirements for plant closure, working inventory, and fleet startup.

https://doi.org/10.1088/1741-4326/acf3fc

Top Stellarator Papers

Reactor Economics / Plant Design

For a summary of parameters that describe a fusion machine, see fusion_reactor_design_parameters.pdf

Prost (2024), "Economically optimized design point of high-field stellarator power-plant"

Provides the central high-field stellarator economic design point on which the EFM P1-like reference architecture and cost hypothesis are based.

https://iopscience.iop.org/article/10.1088/1741-4326/ad142e

Famà (2024), "Thermodynamic and economic analyses of the retrofit of existing electric power plants with fusion reactors"

Evaluates fusion retrofits of existing power plants and is the key reference for EFM's coal-retrofit strategy and reuse of the electrical/thermal island.

https://www.ma2life.org/g/fusion/fusion_volpe_2024/Volpe_2024_Thermodynamic_and_economic_analyses.pdf

Famà (2023), "An optimized power conversion system for a stellarator-based nuclear fusion power plant"

Optimizes power conversion for a stellarator plant and informs the comparison between greenfield efficiency and retrofit compatibility with existing turbines.

https://www.sciencedirect.com/science/article/abs/pii/S0196890422013504

Magnets / Stellarator Design

Botelho (2024), "Simplified magnet design and manufacture based on patterning of wide conductors"

Proposes simplified stellarator magnet manufacturing using patterned wide conductors, directly motivating EFM work on lower-cost, automatable magnet production.

https://arxiv.org/pdf/2409.20143

Robin (2021), “Minimization of magnetic forces on stellarator coils”

Explores reduction of magnetic forces in stellarator coils and is relevant to reducing support-structure mass, magnet complexity, and installed cost.

https://arxiv.org/pdf/2103.13195

Plasma Compatibility of Lithium

Avril (2026), “Particle-in-Cell simulations of lithium redeposition in plasma sheaths under oblique magnetic fields”

Models lithium redeposition near plasma-facing surfaces and supports assessment of impurity control and plasma compatibility for lithium-wall concepts.

https://www.ma2life.org/g/fusion/fusion_volpe_2024/Avril_2025_Particle_in_Cell_simulations_of_lithium_redeposition_in_plasma_sheaths.pdf

Avril (2026), “Global gyrokinetic study of lithium core turbulent transport with and without transport barrier”

Studies lithium transport in the plasma core and helps evaluate whether lithium-facing systems can be compatible with the confinement required by the reference plant.

https://www.ma2life.org/g/fusion/fusion_volpe_2024/Avril_2025_Gyrokinetic_modelling_of_lithium_turbulent_transport_in_core_plasma-13.pdf

Background / Reference

Waganer (2013), “ARIES Cost Account Documentation”

Documents ARIES cost accounts that underlie many historical fusion cost models and is essential for tracing which P1 costs require bottom-up China localization.

https://aries.pppl.gov/LIB/REPORT/ARIES-ACT/UCSD-CER-13-01.pdf

Abdou (2001), “On the exploration of innovative concepts for fusion chamber technology”

Presents APEX liquid-wall and advanced chamber concepts that form an important historical foundation for thick liquid-metal wall architectures.

https://doi.org/10.1016/S0920-3796(00)00433-6 | University repository: https://fti.neep.wisc.edu/fti.neep.wisc.edu/pubs/moreb85a.html?pub=UWFDM-1153

Presentation

Acheson (2025), “Skyfall 2: a 1.2m diameter HTS pseudo-Helmholtz magnet for liquid metal MHD studies”

Describes an HTS magnet intended for liquid-metal MHD experiments and is relevant to scaling magnetic-field test capability for P1 wall studies.

https://indico.global/event/12526/contributions/134341/

Thea Energy / Helios

Swanson (2025), “Overview of the Helios Design: A Practical Planar Coil Stellarator Fusion Power Plant”

Presents a modern planar-coil stellarator power-plant concept and provides an alternative stellarator architecture for comparison with the P1 baseline.

https://arxiv.org/pdf/2512.08027

Type One Energy / Infinity Two

Clark (2025), “Breeder blanket and tritium fuel cycle feasibility of the Infinity Two fusion pilot plant”

Evaluates blanket and tritium feasibility for a stellarator pilot plant and offers a useful independent comparison for EFM breeding and inventory assumptions.

https://doi.org/10.1017/S002237782500039X

Sheffield (2016), “Catalyzed D-D Stellarator Reactor”

Explores a D-D stellarator fuel-cycle concept and provides a reference for considering alternatives if D-T tritium availability becomes a fleet-scale constraint.

https://www.osti.gov/pages/biblio/1261560

Top Liquid Metal Wall Papers

Prost (2024), "Compact fusion blanket using plasma facing liquid Li-LiH walls and Pb pebbles"

Defines a compact liquid Li-LiH/Pb blanket concept closely related to the EFM wall/blanket architecture and its shielding and tritium functions.

https://www.ma2life.org/g/fusion/fusion_volpe_2024/Volpe_Compact_fusion_blanket_using_plasma_facing_liquid_Li-LiH_walls_and_Pb_pebbles.pdf

Giovacchini (2024), "Modeling the free-surface magnetohydrodynamics of thick liquid metal walls for fusion"

Models thick free-surface liquid-metal MHD behavior and directly supports Facility A similarity targets, stability analysis, and pumping-power estimates.

https://arxiv.org/pdf/2412.13598

Castro (2021), "Lithium, a path to make fusion energy affordable"

Argues that lithium plasma-facing systems may reduce major fusion engineering and cost burdens and provides broad motivation for the EFM liquid-wall path.

https://doi.org/10.1063/5.0042437 | Open repository: https://www.osti.gov/pages/biblio/1850218

Sun (2023), "Magnetohydrodynamics in free surface liquid metal flow relevant to plasma-facing components"

Reports free-surface liquid-metal MHD behavior relevant to plasma-facing components and helps define flow stability, drag, and backing-wall effects.

https://www.ma2life.org/g/fusion/fusion_volpe_2024/Volpe_Sun_Liquid_Metal_Wall_Sun_2023_Nucl_Fusion_63_076022.pdf

Kessel (2019), "Critical Exploration of Liquid Metal Plasma-Facing Components in a Fusion Nuclear Science Facility"

Reviews the engineering challenges of liquid-metal plasma-facing components and identifies the performance, materials, safety, and integration questions that must be tested.

https://doi.org/10.1080/15361055.2019.1610685

Puente (2026), "Flowing liquid metal walls as plasma facing components: An integrated approach for a multifunctional system"

Reviews thick flowing liquid-metal wall development as an integrated multifunctional system and summarizes recent experimental progress and remaining scale-up gaps.

https://doi.org/10.1016/j.nme.2026.102203 | Open-access PDF: https://www.ma2life.org/g/fusion/fusion_cc/Puente_Flowing_liquid_metal_walls_2026.pdf

Smolentsev (2026), "Progress and status of the SciDAC-5 project on simulation of plasma-liquid metal interactions in plasma-facing components and breeding blankets"

Describes coupled modeling of plasma-liquid-metal interactions and breeding blankets, providing computational tools that can be benchmarked against EFM experiments.

https://doi.org/10.1016/j.fusengdes.2026.115885

Hu (2024), "Recent progress in the development of liquid metal plasma facing components for magnetic fusion devices"

Reviews recent liquid-metal plasma-facing-component experiments and provides a broad evidence base for selecting EFM materials, geometries, and test priorities.

https://doi.org/10.1016/j.nme.2024.101776 | Open-access PDF: https://www.ma2life.org/g/fusion/fusion_cc/Hu_Recent_progress_liquid_metal_PFCs_2024.pdf

Wynne (2025), "FreeMHD: Validation and verification of the open-source, multi-domain, multi-phase solver for electrically conductive flows"

Validates the open-source FreeMHD solver against electrically conductive flow cases and supports model-to-experiment comparison for free-surface MHD tests.

https://doi.org/10.1063/5.0230242 | Open-access PDF: https://www.ma2life.org/g/fusion/fusion_cc/FreeMHD_Validation_and_verification_2025.pdf

Khodak (2024), "Coupled model for liquid lithium plasma facing components"

Develops coupled modeling for liquid-lithium plasma-facing components and is relevant to linking plasma conditions, heat transfer, and liquid-metal behavior.

https://doi.org/10.1016/j.fusengdes.2024.114651

Altıntaş (2026), "Magnetohydrodynamic behavior of free-surface liquid metal plasma-facing components in fusion reactor environments"

Analyzes free-surface liquid-metal MHD behavior under fusion-relevant conditions and supports EFM stability and pressure-drop design calculations.

https://doi.org/10.1063/5.0315892 | Open-access PDF: https://www.ma2life.org/g/fusion/fusion_cc/Altintas_MHD_free_surface_liquid_metal_PFCs_2026.pdf

Xu (2026), "Lithium Experimental Application Platform (LEAP): Secondary-Containment Architecture for Flowing Liquid Lithium in Fusion Systems"

Describes the LEAP flowing-lithium facility and its secondary-containment architecture, helping EFM identify existing test capabilities that should be used rather than duplicated.

https://arxiv.org/abs/2605.16329 | Open-access PDF: https://www.ma2life.org/g/fusion/fusion_cc/LEAP_Secondary_Containment_Architecture_2026.pdf

Zuo (2019), "Results from an improved flowing liquid lithium limiter with increased flow uniformity in high power plasmas in EAST"

Reports flowing-lithium operation in high-power plasma conditions and provides operating evidence on coverage, heat handling, and plasma interaction.

https://doi.org/10.1088/1741-4326/aaedcb

Maan (2024), "Estimates of global recycling coefficients for LTX-beta discharges"

Quantifies low recycling in lithium-conditioned LTX-beta plasmas and is relevant to understanding how lithium surfaces may affect plasma fueling and edge behavior.

https://doi.org/10.1063/5.0177604

Castro (2023), "Physics and Technology Research for Liquid-Metal Divertor Development, Focused on a Tin-Capillary Porous System Solution, at the OLMAT High Heat-Flux Facility"

Reports high-heat-flux testing of liquid-tin plasma-facing concepts and provides evidence on heat tolerance, replenishment, leakage, and failure modes.

https://doi.org/10.1007/s10894-023-00373-9

Wang (2025), “Experimental investigation of a novel liquid metal plasma facing component with pre-filled microstructures”

Tests a microstructured liquid-metal plasma-facing component and shows how engineered substrates can improve wetting and flow stability in magnetic fields.

https://arxiv.org/abs/2505.08512