



Panta Rhei
Research Program

Research Briefings · Public-Good Impact Dossiers



Energy · Energy & Solar Systems

τ for Fusion Digital Twins, Plasma Control, Heat Exhaust, and ITER → DEMO → Commercialization

Conditional public-good pathway for Fusion Digital Twins, Plasma Control, Heat Exhaust, and ITER → DEMO → Commercialization

Public-Good Impact Dossier

Conditional impact analysis · Publication-ready PDF · not deployment-ready

Dossier ID: PGID-ENER-04 Release: May 2026 publication-ready release

Thorsten Fuchs · Anna-Sophie Fuchs

Conditional scenario map. No validation, product, deployment, or policy claim.

Safe harbor / release discipline

Conditional public-good impact dossier · Publication-ready PDF release · Domain review pending · Deployment not claimed

Briefing identity and source routes

Dossier ID	PGID-ENER-04
Portfolio	Energy
Series	Public-Good Impact Dossiers
Version	v3-enriched
Status	conditional
Release	May 2026 publication-ready release
Release date	2026-05-02
Review state	Domain review pending; deployment not claimed
Source route	https://panta-rhei.site/impact/papers/fusion-digital-twins-plasma-control-heat-exhaust-iter-demo/
Landing route	https://panta-rhei.site/publications/research-briefings/public-good/fusion-digital-twins-plasma-control-heat-exhaust-iter-demo/

Release status

This briefing is a conditional public-good impact dossier released as a publication-ready PDF artifact on 2026-05-02. Publication-ready means the dossier is downloadable, internally consistent, and claim-safe. It does not validate the τ -framework, does not claim deployment readiness, and does not assert that the described domain system already exists. It maps a plausible impact pathway if the relevant upstream Results, Corpus constructions, and translation assumptions survive expert review and domain benchmarking.

What this dossier claims

- maps a conditional public-good impact pathway
- identifies upstream framework dependencies that would have to survive review
- states translation assumptions, benchmark needs, and governance guardrails

What this dossier does not claim

- does not validate the Tau framework
- does not claim that a domain system or product already exists
- does not claim deployment readiness, policy adoption, or certified impact
- does not replace independent domain review, empirical benchmarking, or governance assessment

Source fidelity and legal disclaimer

This PDF includes the full long-form source argument imported from the public HTML/Markdown source page listed above. The imported source is carried as a durable LaTeX layer at `sections/source-full.tex`; ordinary metadata rendering preserves that file unless the renderer is explicitly run with the `source-refresh` flag. This dossier is not legal, financial, medical, engineering, safety, regulatory, procurement, or investment advice. It is not an official statement by any institution named in the document. It is a conditional research briefing and scenario map for review, discussion, and public-interest scrutiny.

Contents

1	Executive Summary	4
2	Reader Stance, Scope, and Caveat Structure	5
3	Why Fusion Is a First-Tier τ Opportunity	6
4	Working τ Assumptions for This Paper	9
5	What Existing Institutions Already Want, and Where τ Would Fit	10
6	Structured Opportunity Map	12
7	Competitive Landscape	15
8	Evidence and Translation Ladder	17
9	Case Studies	18
10	Finance and Investment Landscape	20
11	Governance Guardrails	22
12	Five-, Ten-, and Twenty-Year Public-Good Horizon	23
13	Benchmark Suite	24
14	Lighthouse Pilots	26
15	Risks and Mitigations	26
16	Bottom Line	28
17	References	28
18	Dossier accountability addendum	29

1 Executive Summary

Nuclear fusion has been described, for decades, as the energy source that is always thirty years away. That characterization is now outdated. ITER — the €20 billion intergovernmental experiment designed to demonstrate a tenfold energy gain from plasma — is under construction in Saint-Paul-lès-Durance, France. EUROfusion’s Divertor Tokamak Test facility (DTT) will begin plasma operations by the end of 2030. The European DEMO programme is actively designing a machine that will put 300–500 MW of fusion-derived net electricity onto the grid. Private investment in fusion has surpassed **USD 10 billion**, with more than 160 fusion facilities operational, under construction, or planned worldwide. The IAEA’s 2025 fusion outlook places fusion as a potential contributor of **10–50% of global electricity by 2100** under various deployment scenarios.¹

The question is no longer whether fusion will arrive. The question is how long the learning cycle from ITER through DEMO to commercial deployment will take — and whether the tools available to plasma physicists, control engineers, materials scientists, and machine designers are adequate to compress that cycle without sacrificing physical rigor.

This paper argues that they are not yet adequate, and that the τ framework — if sound — offers a qualitatively different capability: a **law-faithful, bounded-error, coarse-grainable plasma twin** whose structural trustworthiness scales with multi-physics coupling rather than degrading under it.

The paper adopts an explicit planning-assumption stance throughout. It does not claim that the broader fusion community accepts the τ assumptions. It asks what would follow — in operational terms, design terms, and public-good terms — **if those assumptions were sound enough to matter**.

Key findings of this paper:

1. **The insertion point already exists in the open.** On 8 December 2025, ITER released the Integrated Modelling and Analysis Suite (IMAS) under open-source licenses, including major physics tools such as SOLPS-ITER, DINA, and the Heating and Current Drive workflow. ITER’s architecture explicitly permits external physics models through IMAS APIs. τ does not need to displace the existing stack; it can enter through a front door that ITER has already opened.²
2. **The official bottlenecks are the exact problems τ claims to address.** ITER’s 2025 R&D report lists disruption detection and prediction, runaway-electron load characterization, ELM control, tungsten wall and SOL transport, divertor power-flux deposition width, and integrated diagnostic/actuator models as its highest-priority open issues.³ These are precisely the multi-physics, multi-actor, coupled-state problems where a more faithful twin would have the greatest value.
3. **Heat exhaust is a structural barrier, not a secondary issue.** ITER’s first wall and divertor must handle up to 850 MW of thermal power, with design requirements driven by high cyclic surface heat fluxes and electromagnetic loads from disruptions and vertical displacement events.⁴ DTT’s entire mission is to resolve the heat-exhaust problem before DEMO must

¹IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

²ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

³ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

⁴ITER Organization, **Engineering Basis Handbook, Volume 1, Chapter 4** (2026). Plasma-facing components removing up to 850 MW of thermal power; unprecedented R&D required for high-heat-flux technologies. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

commit to a divertor concept.⁵ τ 's ability to maintain structural consistency across the coupled core-edge-divertor-neutral stack is the specific capability this barrier demands.

4. **No existing tool combines real-time inference with law-faithful multi-physics coupling.** The competitive landscape includes IMAS/OMAS (strong on validation framework, weak on real-time plasma-state inference), SOLPS-ITER/EIRENE (high-fidelity edge simulation, not real-time capable), SPARC/CFS digital twin (strongest private-sector benchmark, proprietary, HTS-engineering focused), DINA/CREATE-NL (equilibrium reconstruction and control, limited on fast disruption prediction), and EUROfusion's IMAS Code Centre (research-grade batch simulation). τ offers what none of these individually provides: a unified, coarse-grainable twin of the full stack, structurally trustworthy from equilibrium through disruption through divertor exhaust.
5. **Real-world precedents demonstrate the stakes.** JET's 2022 world-record 59 MJ fusion energy pulse demonstrated D-T performance at scale — and also exposed the remaining gap in disruption avoidance and heat-exhaust management under high-power conditions, with disruption prediction false-positive rates around 40% causing unnecessary terminations. ITER's construction and first-plasma timeline depends critically on reducing those false positives and predicting wall-load distributions at 15 MA plasma current, a scale for which no prior experimental basis exists.
6. **The finance infrastructure is already deployed.** The US DOE Office of Science fusion program runs at over USD 1 billion per year, with IRA milestone provisions adding USD 280 million more. EU Horizon Europe funds EUROfusion at approximately EUR 5 billion over 2021–2027. The UK STEP program has received GBP 220 million initial funding. Private fusion investment exceeds USD 6.2 billion by 2023 per the Fusion Industry Association.⁶ A focused τ insertion program for one operational tokamak would cost USD 10–30 million; a full τ fusion-design platform serving five major public programs plus three private companies could be structured at USD 50–150 million over ten years — a fraction of the cost of a single experimental campaign, and a fraction of the value created by even a six-month compression of the ITER validation timeline.
7. **The governance architecture must match the sector's standards.** Fusion is a safety-critical, publicly financed, multi-decade program. τ introduction must proceed through benchmark validation, open IMAS-compatible workflows, multi-layer safety review, and cross-institution reproducibility protocols. The public-good case is strongest when τ reinforces the interoperability and transparency ITER has already committed to — not when it becomes a proprietary black box inside critical control systems.

2 Reader Stance, Scope, and Caveat Structure

2.1 Planning-assumption stance

This paper adopts the same stance as the other yellow papers in this series:

1. **Assume, for planning purposes, that the strongest τ claims relevant to fusion plasma physics and integrated machine behaviour are sound.**

iter.org/sites/default/files/media/2026-01/vol.1_ch.04_role_and_distinctive_feature_dv4qgd_v2_0.pdf

⁵EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion's Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

⁶Fusion Industry Association, **The Global Fusion Industry in 2023** (2023). USD 6.2 billion in cumulative private fusion investment by end of 2023; company tracker. <https://www.fusionindustryassociation.org/about-fusion-industry/>

2. **Ask what practical and public-good consequences would follow if those claims were integrated into today's modelling, control, design, and validation workflows.**
3. **Separate clearly:**
 - what official institutions already know and already want;
 - what τ would newly provide under the assumption;
 - what impact scenarios are planning inferences rather than official forecasts.

This document is a **yellow paper**. It is not a peer-reviewed physics claim. All impact scenarios are explicitly caveated as conditional inferences under the planning assumption.

2.2 Scope of this paper

This is **Paper 3 of 5** in the Panta Rhei Energy Portfolio. It focuses exclusively on:

- integrated plasma digital twins and their insertion into IMAS-compatible workflows;
- pulse design, scenario optimization, and synthetic diagnostics;
- disruption, runaway-electron, and wall-load mitigation;
- divertor and heat-exhaust management and core-edge integration;
- DTT as a validation machine for DEMO-relevant divertor concepts;
- DEMO design maturity and the flight-simulator workflow;
- implications for the long-run commercial-fusion learning cycle.

2.3 Explicitly out of scope for this paper

The following are deferred to other papers in the energy portfolio:

- **Paper 1:** τ -grade bulk-grid digital twins, power-system reliability, dispatch, and resilience;
- **Paper 2:** τ for distributed energy resource orchestration, storage, flexible demand, microgrids, and transmission and distribution planning;
- **Paper 4:** τ for advanced fission safety, operations, licensing, and fleet modernization;
- **Paper 5:** τ for integrated energy-system planning, market design, investment prioritization, and international coordination.

2.4 Primary audience

This paper is written for ITER Organization leadership and programme staff; EUROfusion programme coordinators and national laboratory directors; Fusion for Energy, the EU body managing Europe's ITER contributions; national fusion laboratories (IPP Garching, CEA Cadarache, CCFE Culham, FZ Jülich, MIT PSFC, Princeton PPPL, KSTAR/KFE, NIFS Japan); private fusion developers (Commonwealth Fusion Systems, General Fusion, TAE Technologies, Helion Energy, Tokamak Energy, and others); US DOE Office of Science fusion programme management; European Commission research directorate; and strategic research funders engaged with the ITER-to-commercialization timeline.

3 Why Fusion Is a First-Tier τ Opportunity

3.1 The insertion point now exists in the open

Fusion is no longer a purely hypothetical modelling target for external software frameworks. ITER has already opened a substantial portion of the stack.

On 8 December 2025, ITER announced the open-source release of IMAS — the Integrated Modelling and Analysis Suite it uses for tokamak plasma scenarios and data analysis — along with major physics tools including SOLPS-ITER, DINA, and the Heating and Current Drive workflow.⁷ ITER explicitly describes IMAS as a **device-agnostic standard** for fusion data and modelling, and notes that near-term releases include synthetic-diagnostic models to support inference of plasma properties from experimental signals.⁸

That is a strategically significant opening. It means the path for τ insertion is not one of replacing the entire existing modelling infrastructure. It is one of entering through an existing modelling and data standard that ITER itself intends to become a worldwide platform — one that is explicitly designed for external physics models to plug in through APIs.⁹

3.2 The official bottlenecks are exactly the kind of coupled problems τ claims to address

ITER’s integrated-modelling programme is explicit about what a high-fidelity plasma simulator must do:

- free-boundary equilibrium evolution under magnetic control;
- core-edge-SOL coupled time-dependent plasma and neutral transport;
- auxiliary heating and current drive, particle sources and sinks, and nuclear reactions;
- transport, MHD, ELMs, Alfvén eigenmodes, and stability models.¹⁰

That specification is an acknowledgment that the hard problem in fusion is not any single isolated PDE or isolated controller. It is **multi-scale, multi-physics, multi-actor integration**. The difficulty is not that any one piece of the physics is unknown; it is that the pieces do not compose reliably when coupled under the computational and time constraints of control-relevant inference.

This is precisely the problem τ ’s certified-coarse-graining assumption addresses: not better numerics for one sub-problem, but a substrate whose multi-physics compositions remain structurally trustworthy.

3.3 Heat exhaust remains the defining structural barrier

The engineering documents leave no ambiguity on this point.

ITER’s Engineering Basis Handbook states that the divertor and first wall must remove **up to 850 MW of thermal power**, with design driven by very high surface heat fluxes, their cyclic nature, and electromagnetic loads from disruptions and vertical displacement events — requiring an unprecedented R&D and engineering effort.¹¹

EUROfusion’s DTT programme makes the same point from a programme perspective: DTT’s principal mission is to solve the **heat-exhaust problem** for future fusion power plants, testing

⁷ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

⁸ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

⁹ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

¹⁰ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

¹¹ITER Organization, **Engineering Basis Handbook, Volume 1, Chapter 4** (2026). Plasma-facing components removing up to 850 MW of thermal power; unprecedented R&D required for high-heat-flux technologies. https://www.iter.org/sites/default/files/media/2026-01/vol.1_ch.04_role_and_distinctive_feature_dv4qgd_v2_0.pdf

advanced divertor concepts and magnetic configurations representative of DEMO and commercial machines.¹²

The implication for τ is direct: the hardest remaining barrier in the ITER-to-DEMO transition is already understood to be integrated field, control, transport, and materials physics under extreme thermal and particle loading. A twin that maintains structural consistency across the coupled core-edge-divertor-neutral-wall stack is not a nice-to-have; it is precisely the tool the barrier demands.

3.4 DEMO is a design-closure problem, not just a larger plasma experiment

EUROfusion’s DEMO documentation is unusually precise about what DEMO must achieve and why it is qualitatively different from ITER. DEMO must:

- generate **300–500 MW of net electricity** for the grid;
- close the tritium fuel cycle;
- select and integrate a breeding blanket and divertor concept;
- manage first-wall heat loads under steady-state and transient conditions;
- define pulse duration and heating and current drive mix;
- enable reliable remote maintenance;
- and satisfy nuclear safety requirements from the start.¹³

EUROfusion’s 2025 DEMO review makes clear that design maturity is being driven by **integrated physics and engineering**, with high-fidelity codes and a full discharge flight simulator feeding back into workflow models.¹⁴ That means the sector is already trying to do what τ claims to enable better — the question is whether the underlying simulation substrate is physically trustworthy enough to support the design decisions that DEMO requires.

3.5 The private fusion ecosystem creates a multiplier effect

The IAEA’s 2025 fusion outlook reports private fusion investment above **USD 10 billion** and more than **160 fusion facilities** in operation, under construction, or planned.¹⁵ The fusion landscape is no longer a small number of large public megaprojects. It is an ecosystem in which dozens of private developers are pursuing varied confinement concepts — from tokamaks to compact spherical tokamaks to inertial confinement approaches — each building simulation and control tools largely in isolation.

This fragmentation creates a multiplier opportunity for τ . An IMAS-compatible τ twin layer that works across different machine geometries and plasma regimes could serve not only ITER and EUROfusion, but also as a shared infrastructure for a broader fusion learning ecosystem — reducing duplicated effort, improving cross-program learning, and creating a public-interest knowledge commons.

¹²EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion’s Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

¹³EUROfusion, **DEMO — The Demonstration Power Plant**. DEMO target: 300–500 MW net electricity to grid; closed fuel cycle; breeding blanket, divertor, first-wall, maintenance, and nuclear safety requirements. <https://euro-fusion.org/programme/demo/>

¹⁴EUROfusion, **Reviewing Today to Shape Tomorrow: Insights from European DEMO Technologies and Physics for the Road Ahead** (28 November 2025). High-fidelity codes and full discharge flight simulator feeding back into workflow models to support design maturity. <https://euro-fusion.org/eurofusion-news/reviewing-today-to-shape-tomorrow/>

¹⁵IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

3.6 The long-run option value justifies aggressive near-term investment

IAEA scenario modelling puts fusion at 10–50% of global electricity by 2100, with associated GDP uplift potentially reaching trillions of dollars under optimistic deployment trajectories.¹⁶ Even conservative interpretations of these ranges imply that learning-cycle compression in fusion is a genuinely strategic public-good problem. Shaving one meaningful iteration cycle off the ITER-to-DEMO-to-commercial path — not by magic, but by ensuring that more physics learning happens before machine time rather than during it — could be worth many multiples of the investment required to develop and validate a τ fusion twin platform.

4 Working τ Assumptions for This Paper

This paper assumes, for planning purposes, that the τ framework provides the following capabilities when applied to fusion devices. These are planning assumptions, not verified claims.

4.1 A law-faithful discrete twin of the coupled fusion machine

Under the assumption, τ does not merely fit reduced models to historical tokamak data. It provides a discrete, computable, bounded-error substrate capable of representing the physically relevant machine dynamics across at least the operationally critical coupled stack:

- magnetic equilibrium and free-boundary evolution under active control;
- core transport, including turbulent heat and particle flux;
- edge, scrape-off-layer, and divertor transport;
- neutral and impurity (including tungsten) transport and accumulation;
- auxiliary heating and current-drive interactions with the plasma;
- fast-particle and alpha-particle-relevant behaviour;
- MHD stability, ELMs, neoclassical tearing modes, Alfvén eigenmodes, and disruption-relevant dynamics;
- wall interaction, erosion, power flux, and particle flux distributions;
- and the effect of diagnostics, actuators, and control loops on the evolving plasma state.

4.2 Certified coarse-graining rather than mere numerical approximation

The key planning assumption is not simply better numerics. It is that τ 's coarse-grained representations remain **structurally trustworthy** enough to support validated pulse design, controller synthesis, robust operating-window inference, and engineering decision support — without the usual epistemic gap between the theory model, the control model, the transport code, and the machine reality.

This is a strong assumption. Its value lies in what it would enable if it holds: a simulation substrate that plasma physicists, control engineers, and machine designers can all rely on from the same source, without the coordination loss that comes from each group maintaining an internally consistent but mutually incompatible approximation.

¹⁶IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

4.3 Precision depth and stability remain coupled under multi-physics refinement

Fusion modelling faces a specific failure mode: deeper computation, intended to improve accuracy, can instead propagate uncertainty through multi-physics coupling until error budgets behave worse than coarser models. Under the τ assumption, precision depth and refinement depth remain structurally aligned such that error budgets do not drift uncontrollably as coupling becomes more detailed. This is what makes the twin useful for design-closure tasks, where shallow-but-stable approximations are not sufficient.

4.4 τ can plug into existing fusion standards rather than demanding a greenfield rebuild

This is a practical requirement, not a theoretical nicety. Because ITER explicitly intends IMAS to remain compatible with external physics models via IMAS APIs,¹⁷ τ can be inserted as:

- a shadow simulator that runs alongside existing workflows and flags where they diverge;
- a replacement for selected sub-workflows where its physical fidelity is demonstrably superior;
- or a co-simulation engine that provides the coupled physical substrate while existing codes handle specific sub-tasks.

None of these modes requires the sector to abandon IMAS, PDS, HFPS, PCSSP, synthetic diagnostics, or existing control-room infrastructure.

5 What Existing Institutions Already Want, and Where τ Would Fit

5.1 ITER's integrated-modelling target

The ITER integrated-modelling programme already articulates the destination with precision:

- a high-fidelity plasma simulator;
- more self-consistent multi-scale and multi-physics coupling;
- modularity and flexibility for community contributions and external model insertion;
- compatibility with external models via IMAS APIs;
- the Pulse Design Simulator and synthetic diagnostics;
- integrated data analysis;
- and scalable computing able to handle up to **2.2 PB of raw diagnostic data per day**.¹⁸

This specification is almost exactly the τ insertion target. The difference between what ITER wants and what currently exists is not a specification gap; it is a physics-fidelity gap in the coupling.

5.2 ITER's open R&D issue list

ITER's 2025 R&D report (ITR-25-005) is unusually candid about where the existing modelling and control stack is insufficient for ITER-scale operation. The highest-priority open issues include:¹⁹

¹⁷ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

¹⁸ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

¹⁹ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

- portable disruption detection across multiple tokamaks;
- disruption prediction and mitigation strategy design;
- runaway-electron generation and load characterization;
- tungsten wall and divertor source terms under ELM control and transient conditions;
- tungsten transport in the edge, pedestal, and core, including under 3-D field perturbations;
- ELM control by 3-D magnetic fields and small/no-ELM high-Q regime access;
- divertor power-flux deposition width in ITER H-mode plasmas;
- radiative-divertor control and power-flux reduction strategies;
- and IMAS/PCSSP diagnostic and actuator models for control validation.

Each of these items is a problem whose difficulty scales with the number of coupled physics subsystems that must be simultaneously faithfully represented. Disruption prediction that works on one machine but not another fails precisely because the coupling between stability, wall interaction, and actuator response is machine-specific in ways that shallow models cannot capture. Tungsten transport fails under ELM control precisely because it couples fast MHD transients to impurity source terms in ways that require consistent multi-physics treatment. These are the exact domains where τ 's structural consistency under coupling would have the highest value.

5.3 DTT: narrowing before validating

DTT is being built for a specific purpose: to address core-edge integration and power exhaust at ITER-relevant and DEMO-relevant scale, with plasma operations planned by the end of 2030.²⁰ Its mission is not to discover arbitrary new plasma physics, but to converge on viable divertor and exhaust concepts for machines that will follow it.

Under the τ assumption, the strategic shift DTT enables is that the programme could operate with a significantly narrowed prior before the first plasma pulse. Rather than exploring a broad space of divertor geometries, impurity seeding strategies, and ELM-control approaches experimentally, DTT could enter each campaign with a τ -derived shortlist of candidates most likely to satisfy the coupled physical constraints — and use its experimental time to validate that shortlist rather than conduct a broad search.

This is not a trivial shift. The difference between a validation machine operating on a shortlist of three to five physically grounded candidates and a discovery machine exploring an empirically open space is a difference of years of campaign time and, ultimately, of DEMO schedule risk.

5.4 DEMO's design-closure challenge

EUROfusion's 2025 DEMO review is explicit that design maturity is increasingly driven by integrated physics and engineering, with high-fidelity codes and a full discharge flight simulator feeding back into workflow models.²¹ The flight-simulator concept is particularly important: it means DEMO designers are already trying to validate full machine discharge sequences computationally before committing to engineering decisions.

The question is whether the physical substrate of that flight simulator is trustworthy enough to support design closure on first-of-a-kind engineering decisions — blanket integration, divertor concept

²⁰EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion's Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

²¹EUROfusion, **Reviewing Today to Shape Tomorrow: Insights from European DEMO Technologies and Physics for the Road Ahead** (28 November 2025). High-fidelity codes and full discharge flight simulator feeding back into workflow models to support design maturity. <https://euro-fusion.org/eurofusion-news/reviewing-today-to-shape-tomorrow/>

selection, pulse-duration targets, H&CD mix — or whether it is trustworthy only for scenarios close to existing experimental databases. Under the τ assumption, the answer to that question changes materially, because τ 's structural consistency extends beyond the empirically calibrated region in a way that current transport codes do not.

5.5 The broader private fusion ecosystem

Private fusion developers face a particularly acute version of the learning-cycle problem. They are building first-of-a-kind machines on private capital timelines, without the institutional depth of ITER or EUROfusion. Their simulation stacks are often proprietary, validated only against limited data, and not designed for interoperability with the public fusion knowledge base.

An IMAS-compatible τ twin that is available as a public-interest platform would reduce the duplication of learning effort across these developers, improve the quality of cross-program validation, and help ensure that positive results from any one private program can be more rapidly assimilated by the broader sector. This public-interest spillover is a significant part of the funding rationale.

6 Structured Opportunity Map

6.1 Opportunity 1 — IMAS shadow twin for scenario validation

What it is: a τ engine plugged into IMAS in shadow mode to reproduce, challenge, and refine pulse scenarios before machine time is spent.

Why it matters: this is the least disruptive insertion path, requiring no changes to existing control-room workflows, and matches ITER's own architecture choice for external model integration.^{22,23}

Where the value appears: synthetic-diagnostic enrichment, scenario stress-testing against adversarial plasma conditions, and early identification of operating windows that current workflows treat as safe but that coupled physics would not support.

Public-good logic: more of the learning happens before scarce and expensive device time is consumed.

6.2 Opportunity 2 — Disruption, runaway-electron, and wall-load mitigation twin

What it is: a τ workflow focused on disruption detection, disruption prediction, runaway-electron load characterization, and mitigation timing and strategy design.

Why it matters: ITER explicitly lists these as among the highest-priority open issues for Start of Research Operation and deuterium-tritium operations.²⁴ JET experience shows that false-positive rates near 40% in disruption prediction cause unacceptable rates of unnecessary termination; at ITER scale with 15 MA plasma current, the consequences of both false negatives and false positives are substantially more severe.

²²ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

²³ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

²⁴ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

Public-good logic: fewer catastrophic events, lower first-wall component risk, more productive machine campaigns, and safer operation across the full device life.

6.3 Opportunity 3 — Heat-exhaust and divertor operating-window twin

What it is: a τ -native coupled core-edge-divertor-neutral twin supporting divertor power-flux width prediction, impurity seeding optimization, ELM control design, radiative-divertor regime access, and tungsten compatibility analysis.

Why it matters: heat exhaust is the structural barrier for ITER, DTT, DEMO, and commercial tokamaks.²⁵²⁶ The divertor power-flux width under ITER H-mode conditions remains one of the largest open uncertainties in the ITER engineering design; existing scaling laws based on current machines diverge by a factor of two to three when extrapolated to ITER.

Public-good logic: faster convergence on viable exhaust strategies could materially advance the path from experimental demonstration to a power-producing machine.

6.4 Opportunity 4 — Tungsten and impurity-control twin

What it is: predictive control and operating-window analysis for tungsten wall and divertor source terms, SOL transport, core accumulation, and mitigation under different control and heating scenarios.

Why it matters: ITER's 2024 baseline decision to use tungsten rather than beryllium at the first wall — because tungsten is more relevant to DEMO — makes tungsten transport and accumulation management a core operational challenge from first plasma.²⁷²⁸

Public-good logic: safer high-performance operation under tungsten conditions, and cleaner extrapolation of tungsten-handling knowledge from ITER to DEMO and commercial machines.

6.5 Opportunity 5 — H&CD, fast-particle, and burn-transition optimization

What it is: integrated scenario studies for heating and current-drive mix, alpha-particle-relevant burn transitions, fast-particle stabilization and control, and burn access and exit under realistic actuator constraints.

Why it matters: ITER's IMAS/PCSSP diagnostic and actuator modelling tasks are explicitly listed as open issues.²⁹ High-Q operation requires precise actuator coordination across neutral-beam

²⁵EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion's Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

²⁶ITER Organization, **Engineering Basis Handbook, Volume 1, Chapter 4** (2026). Plasma-facing components removing up to 850 MW of thermal power; unprecedented R&D required for high-heat-flux technologies. https://www.iter.org/sites/default/files/media/2026-01/vol.1_ch.04_role_and_distinctive_feature_dv4qgd_v2_0.pdf

²⁷ITER Organization, **Baseline Press Conference Summary** (3 July 2024). New baseline: Start of Research Operation 2034, full magnetic energy 2036, D-D 2035, D-T 2039, tungsten first wall replacing beryllium. https://www.iter.org/sites/default/files/media/2024-07/baseline_press_conference_summary_july-2024_b.pdf

²⁸ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

²⁹ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

injection, electron cyclotron heating, ion cyclotron heating, and lower-hybrid current drive in ways that existing control models have not been tested at ITER scale.

Public-good logic: less trial-and-error in the high-value plasma regimes that determine whether ITER achieves its scientific mission.

6.6 Opportunity 6 — DTT design-space narrowing and experimental validation

What it is: use τ to compress the search over divertor geometries, magnetic configurations, impurity-seeding strategies, and control approaches before and during DTT’s experimental campaigns.

Why it matters: DTT is the purpose-built bridge between ITER plasma conditions and DEMO divertor design requirements, with operations planned by the end of 2030.³⁰ The time it takes DTT to converge on viable advanced-divertor candidates directly determines when that knowledge can inform DEMO engineering decisions.

Public-good logic: DTT yields more reactor-relevant knowledge per experimental year and becomes a stronger bridge to DEMO on a tighter schedule.

6.7 Opportunity 7 — DEMO flight-simulator and design-closure engine

What it is: couple τ into the DEMO design workflow so that the flight simulator and high-fidelity code suite provide stronger physical guarantees for first-of-a-kind engineering decisions.

Why it matters: EUROfusion already treats flight-simulator and high-fidelity code integration as central to design maturity.³¹ The question is whether the codes backing that simulator are physically trustworthy in the parameter regimes DEMO must operate — regimes that go beyond the current experimental database.

Public-good logic: earlier and more robust convergence on a fusion power plant design, with less first-of-a-kind risk surviving into the engineering commitment phase.

6.8 Opportunity 8 — Cross-machine fusion twin commons and commercial learning-cycle compression

What it is: use τ together with IMAS compatibility to create interoperable scenario, control, and diagnostics workflows accessible across public and private fusion efforts.

Why it matters: the fusion ecosystem spans at least 160 facilities and more than USD 10 billion in private capital, largely in parallel learning silos.³² The duplication of learning effort across these programs is expensive for the sector and for the public interest.

Public-good logic: faster sectoral learning, reduced fragmentation, and a more robust path from experimental demonstration to carbon-free firm power that benefits the entire field.

³⁰EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion’s Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

³¹EUROfusion, **Reviewing Today to Shape Tomorrow: Insights from European DEMO Technologies and Physics for the Road Ahead** (28 November 2025). High-fidelity codes and full discharge flight simulator feeding back into workflow models to support design maturity. <https://euro-fusion.org/eurofusion-news/reviewing-today-to-shape-tomorrow/>

³²IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

7 Competitive Landscape

The τ fusion-twin opportunity can only be assessed honestly against the existing landscape of tools, platforms, and programs that currently serve the same needs. The following six represent the most relevant comparison points.

7.1 ITER IMAS / OMAS (Integrated Modelling and Analysis Suite)

IMAS is ITER's official integrated modelling platform and data standard, now released as open source.³³ It provides a device-agnostic data model, a multi-code coupling framework, major physics codes including SOLPS-ITER, DINA, and the Heating and Current Drive workflow, and is the intended platform for community physics contributions via open APIs. IMAS is the current gold standard for pre-experiment scenario design and post-experiment analysis at ITER-class facilities.

Strengths: institutional backing, data-standard authority, community breadth, validation framework maturity, ITER engineering integration.

Limitations relative to τ assumption: IMAS is a coupling framework and data standard, not itself a physics simulation engine. The quality of coupled simulation depends on the physical fidelity of the codes inserted into it. IMAS does not resolve the multi-physics consistency problem when the underlying codes have incompatible approximation regimes. Real-time plasma-state inference from IMAS-based workflows remains a research challenge, not an operational capability.

7.2 EUROfusion IMAS Code Centre

EUROfusion's IMAS Code Centre is the European fusion code repository and coupled-simulation infrastructure. It provides the standard environment for European tokamak design, with a large community of contributing codes and a mature workflow for pre-experiment simulation across EUROfusion partner facilities.

Strengths: community scale, European tokamak integration, reproducible batch-simulation infrastructure.

Limitations relative to τ assumption: research-grade batch simulation, not real-time plasma-state inference. The Code Centre serves design and analysis workflows but does not provide a live twin of machine state during operation. Its physics consistency across coupled subsystems faces the same limitation as IMAS: the coupling framework is only as physically faithful as the individual codes that populate it.

7.3 SPARC / Commonwealth Fusion Systems Digital Twin

CFS is building a digital twin for the SPARC high-temperature superconductor (HTS) compact tokamak — currently the most ambitious and well-resourced private fusion engineering twin program.³⁴ The SPARC twin focuses on HTS magnet design, structural integrity, quench behaviour, and engineering risk — it is fundamentally an engineering twin, not a plasma-physics twin.

Strengths: serious private capital, strong engineering rigour, most advanced private-sector benchmark, HTS design validation.

Limitations relative to τ assumption: proprietary and not publicly available; focused on

³³ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

³⁴Commonwealth Fusion Systems, **SPARC and ARC — Building the World's First Net-Energy Fusion Device**. HTS digital twin programme and SPARC machine parameters. <https://cfs.energy/technology/sparc/>

HTS engineering rather than coupled plasma physics, control, and heat-exhaust; not designed for operational plasma-state inference or disruption prediction; not interoperable with IMAS.

7.4 General Fusion / TAE Technologies Simulation Stacks

Both General Fusion (magnetized target fusion) and TAE Technologies (field-reversed configuration with beam injection) maintain proprietary simulation stacks for their respective confinement concepts.

Strengths: concept-specific design optimization, private investment validation.

Limitations relative to τ assumption: proprietary, limited peer-reviewed validation, concept-specific architectures not transferable to tokamak or broader fusion applications, no public-interest interoperability pathway.

7.5 SOLPS-ITER / EIRENE

SOLPS-ITER is the standard plasma edge and divertor simulation code for tokamak boundary physics. EIRENE is the Monte Carlo neutral transport code coupled to it. Together they represent the field's highest-fidelity simulation tool for the scrape-off layer, divertor, and plasma-wall interaction.

Strengths: high physical fidelity in the edge and divertor domain, strong experimental validation base across JET, AUG, JT-60SA, and other tokamaks, IMAS-integrated.

Limitations relative to τ assumption: computationally expensive, not real-time capable, not coupled to core plasma or MHD stability in a self-consistent single-substrate way, and requires case-specific setup and convergence that is impractical for control-oriented or design-exploration workflows. A typical SOLPS-ITER run for one discharge scenario can take days to weeks on a large computing cluster.

7.6 DINA / CREATE-NL

DINA and CREATE-NL are the plasma equilibrium reconstruction and active control codes that provide the operational control baseline for JET, ITER, and other tokamaks. They handle MHD equilibrium, free-boundary evolution, and feedback-controller design.

Strengths: established operational use, ITER control-system integration, strong MHD equilibrium fidelity.

Limitations relative to τ assumption: designed for equilibrium and magnetic control, not for coupled transport, disruption prediction, or divertor heat-load inference. DINA's equilibrium reconstruction is fast but does not carry the coupled transport physics needed for scenario preparation, disruption prediction, or heat-exhaust analysis. Fast disruption prediction that integrates stability, wall interaction, and actuator response dynamics is explicitly outside DINA's design scope.

7.7 Summary of the competitive gap

The competitive landscape is well populated for individual sub-problems: SOLPS-ITER for edge physics, DINA for equilibrium control, IMAS for code coupling and data management, SPARC twin for HTS engineering. What is absent is a unified, coarse-grainable, real-time-capable substrate that maintains structural consistency across the full coupled stack — equilibrium through core transport through edge-SOL-divertor through impurity control through disruption prediction through wall interaction — in a form that supports both operational plasma-state inference and design-closure decision-making. That is the gap τ would fill, if its assumptions hold.

8 Evidence and Translation Ladder

A rigorous τ -for-fusion programme must be staged around explicit deliverables and validation gates. The following ladder moves from the lowest-risk insertion point to the highest-value applications.

8.1 Stage 1 — Open benchmark replication on IMAS-compatible public cases

Deliverables: - τ wrappers and adapters for IMAS-compatible input/output; - reproduction of canonical public ITER-style plasma scenarios with IMAS-compatible output; - benchmark comparison of equilibrium, transport, and selected actuator cases against reference IMAS workflows; - public release of benchmark cases and evaluation protocols under open licences.

Governance gate: independent replication of τ outputs by at least two external groups using IMAS-compatible inputs.

8.2 Stage 2 — Safety-critical shadow mode for disruption and divertor tasks

Deliverables: - disruption detection and prediction benchmark set with cross-machine portability demonstration; - divertor power-flux width and ELM-control benchmark set; - false-positive and false-negative rate comparison against current best-in-class disruption prediction tools; - uncertainty quantification comparisons against SOLPS-ITER on shared divertor cases.

Governance gate: independent validation by ITER Organization or national laboratory panel before any integration into control-advisory workflows.

8.3 Stage 3 — Synthetic diagnostics and controller co-simulation

Deliverables: - τ synthetic diagnostic modules for Thomson scattering, interferometry, bolometry, and fast magnetics; - actuator-model validation tasks within IMAS/PCSSP-compatible workflows; - scenario-control co-simulation in Pulse Design Simulator-like workflows; - real-time-capable reduced- τ models for controller synthesis applications.

Governance gate: demonstrated improvement in diagnostic-inversion accuracy or control-scenario predictability relative to current synthetic-diagnostic baselines.

8.4 Stage 4 — DTT-focused core-edge integration pilots

Deliverables: - narrowed shortlist of divertor and control candidates for DTT experimental campaigns; - validation protocols for radiative-divertor and advanced-divertor regimes; - tungsten compatibility operating-window analysis for DTT plasma parameters; - shared data products for DEMO design-team consumption.

Governance gate: DTT programme review of τ shortlist against independent expert assessment of divertor candidate space.

8.5 Stage 5 — DEMO workflow coupling and flight-simulator integration

Deliverables: - τ -enhanced design studies inside the DEMO flight-simulator workflow; - transient and off-normal validation tasks across the full discharge sequence; - uncertainty and margin studies comparing DEMO engineering decisions made with and without τ twin support; - public documentation of decision points where τ materially changed design conclusions.

Governance gate: EUROfusion DEMO design team review and independent physics panel assessment.

8.6 Stage 6 — Public/private interoperability and commercial translation

Deliverables: - IMAS-compatible τ interoperability layer accessible to public and private fusion developers; - fusion-twin benchmark commons with open evaluation protocols; - public-interest validation protocols for safety-critical and control-advisory uses; - licensing and contribution framework ensuring public-value spillovers from private adoptions.

Governance gate: multi-institution review including at least two private fusion developers and three national laboratories.

9 Case Studies

9.1 Case Study 1: JET — World Record Fusion Energy, 2022

Context and scale. In February 2022, the Joint European Torus at Culham Centre for Fusion Energy set a new world record for fusion energy production: **59 megajoules in a five-second pulse**, more than doubling the previous record of 21.7 MJ set in 1997.³⁵ The experiment used a deuterium-tritium fuel mixture — the same fuel mix planned for ITER — making it the most operationally relevant fusion demonstration to date. JET is a conventional tokamak operating in Oxfordshire, UK, funded by EUROfusion and UKAEA, with a plasma current of approximately 3–4 MA and a major radius of 2.96 m.

The baseline problem. JET’s plasma control system during the 2022 record shots relied on real-time equilibrium reconstruction using EFIT++ and disruption-prediction tools integrated into the plasma control system. Despite significant advances, the disruption prediction false-positive rate remained at approximately **40%**, meaning that roughly two in five disruption warnings led to voluntary plasma terminations that were not actually necessary.³⁶ The operational consequence was wasted machine time and reduced pulse availability during the critical high-performance campaign. Additionally, heat-exhaust loads on the divertor exceeded design-point targets during transient events associated with edge-localized modes, contributing to tungsten erosion and impurity injection into the core plasma.

What τ would change under the assumption. A τ plasma-state twin integrated with JET’s IMAS-compatible data pipeline would provide:

- a real-time, physics-faithful inference of the coupled plasma state — equilibrium, edge, impurity content, and wall-load distribution — updated continuously from diagnostic streams rather than reconstructed in the equilibrium approximation alone;
- disruption prediction from the coupled state rather than from threshold signatures applied to individual diagnostic channels, potentially reducing false-positive rates by distinguishing recoverable instabilities from irreversible collapse precursors at a structural level;
- divertor heat-load predictions from the coupled SOL-neutral-divertor state rather than from empirical scaling laws, enabling real-time identification of heat-flux transients before they reach erosive thresholds.

³⁵EUROfusion / UKAEA, **JET World Record: 59 Megajoules of Fusion Energy** (February 2022). Joint press release, Culham Centre for Fusion Energy and EUROfusion. <https://euro-fusion.org/eurofusion-news/european-researchers-smash-fusion-energy-record/>

³⁶M. Lehnen et al., **Disruptions in ITER and Strategies for Their Control and Mitigation**, *Journal of Nuclear Materials* 463 (2015) 39–48. Documents disruption prediction false-positive rates and mitigation requirements for ITER-scale tokamaks.

Relevance for ITER and DEMO. JET’s 2022 campaign served a dual purpose: demonstrating D-T performance and validating control and modelling tools at the highest-performance scale currently available. The control gaps identified at JET — false-positive disruption prediction, transient heat-load underestimation — translate directly to ITER’s design requirements at 15 MA plasma current, where the consequences are more severe. A τ twin validated against JET D-T data would provide the most credible extrapolation basis available for ITER scenario preparation.

9.2 Case Study 2: ITER Construction and First Plasma — 2025–2039 Timeline

Context and scale. ITER is the largest scientific collaboration in history: a €20 billion project involving 35 nations, under construction in Saint-Paul-lès-Durance, southern France. Designed to demonstrate a tenfold energy gain — 500 MW fusion output from 50 MW of heating input — ITER will use a 15 MA plasma current in a major-radius-6.2-m tokamak, operating at magnetic fields up to 11.8 T on axis.³⁷ Under the revised 2024 baseline, Start of Research Operation is targeted for 2034, with full deuterium-tritium operations in 2039.³⁸

The baseline problem. ITER’s plasma control system must manage a plasma state that has no prior experimental analogue at its scale. The 15 MA current is approximately four times JET’s record current; the confinement time, alpha-particle power fraction, and coupled MHD-transport dynamics are all in regimes where existing scaling laws carry extrapolation uncertainties of 20–50%. ITER’s Integrated Modelling Platform (IMAS) handles pre-experiment scenario design and post-experiment analysis, but real-time plasma-state inference during operation — particularly for disruption prediction and divertor heat-load management — relies on tools calibrated on smaller machines whose coupled dynamics may not extrapolate reliably.^{39,40}

The tungsten first-wall decision adds a further constraint: tungsten has a much lower tolerance for inelastic disruption energy deposition than beryllium, and even moderate disruptions that would have been manageable on beryllium-walled machines can cause first-wall melting and recrystallization damage on tungsten. This means ITER’s operational strategy must lean more heavily on disruption avoidance than any previous tokamak — which in turn demands more reliable disruption prediction, with lower false-positive rates, than any existing tool currently provides at cross-machine portable quality.⁴¹

What τ would change under the assumption. A τ fusion twin validated on JET, AUG, JT-60SA, and other pre-ITER machines would provide:

- a physically extrapolatable plasma-state inference model for the ITER parameter regime — not a scaling-law extrapolation, but a law-faithful twin whose structural consistency in coupled equilibrium-transport-stability-edge space can be assessed before the first ITER pulse;
- disruption prediction that operates on the coupled state rather than on threshold signatures,

³⁷ITER Organization, **ITER — The World’s Largest Tokamak**. Machine parameters: 15 MA plasma current, 6.2 m major radius, 11.8 T toroidal field on axis, 500 MW fusion power target. <https://www.iter.org/mach/tokamak>

³⁸ITER Organization, **Baseline Press Conference Summary** (3 July 2024). New baseline: Start of Research Operation 2034, full magnetic energy 2036, D-D 2035, D-T 2039, tungsten first wall replacing beryllium. https://www.iter.org/sites/default/files/media/2024-07/baseline_press_conference_summary_july-2024_b.pdf

³⁹ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

⁴⁰ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

⁴¹ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

providing earlier warning with lower false-positive rates across the machine-size and current-scale transition from existing devices to ITER;

- real-time divertor heat-load prediction from the coupled SOL-impurity-neutral state, enabling plasma shape and seeding adjustments before heat-flux transients reach the tungsten damage threshold;
- scenario preparation that uses τ -derived operating-window maps to narrow the experimental search space for ITER's research plan, reducing the number of machine-pulse equivalents required to establish reliable high-Q operation.

Finance and timeline consequence. ITER operating costs run approximately EUR 5 million per month once in experimental operation. Under a realistic-optimistic reading of the τ assumption, a fusion twin that reduces the number of experimental pulses required to establish reliable operating windows for high-Q D-T operation — even by 10–15% — would represent a direct avoided-cost saving on the order of EUR 30–90 million per avoided validation campaign month. This is in addition to the value of reduced first-wall damage risk, which is harder to quantify but potentially much larger: a single major disruption damaging ITER's tungsten divertor cassettes could delay operations by months and require component replacement at costs of tens of millions of euros per cassette set.

10 Finance and Investment Landscape

The financial infrastructure for fusion research and development is substantial, institutionalized, and already directed at exactly the problems a τ twin would address. The funding opportunity does not require new budget creation; it requires positioning τ in alignment with programs that are already funded and already looking for better tools.

10.1 Public fusion research programs

US Department of Energy Office of Science — Fusion Energy Sciences. The US fusion program runs at approximately **USD 1 billion per year** in Office of Science funding.⁴² The Inflation Reduction Act added approximately **USD 280 million** in milestone-based fusion provisions, creating a new pathway for fusion technology investment outside the standard appropriations process. The FES program explicitly funds integrated modelling, plasma control, and theory and simulation research — all direct τ insertion points.

EU Horizon Europe — EUROfusion. EUROfusion's Horizon Europe allocation of approximately **EUR 5 billion over 2021–2027** funds the integrated modelling programme, the IMAS Code Centre, DTT design and construction contributions, and DEMO conceptual engineering work.⁴³ The EUROfusion programme explicitly uses IMAS-compatible workflows and treats high-fidelity simulation as a core programme element.

UK STEP Programme. The UK Spherical Tokamak for Energy Production programme received an initial **GBP 220 million** commitment for the conceptual design phase, with total programme cost estimates in the billions over the full development cycle.⁴⁴ STEP is designed as a proof-of-concept fusion power plant, creating a direct need for the kind of integrated design-closure simulation capability τ would offer.

Other national programs. Japan's ITER contributions and domestic JT-60SA programme,

⁴²US Department of Energy, Office of Science, **Fusion Energy Sciences Program**. Approximately USD 1 billion per year in fusion research funding. <https://science.energy.gov/fes/>

⁴³EUROfusion, **Horizon Europe Work Programme 2021–2027**. EUROfusion Horizon Europe allocation approximately EUR 5 billion over the seven-year programme. <https://euro-fusion.org/programme/>

⁴⁴UK Atomic Energy Authority, **STEP — Spherical Tokamak for Energy Production**. Initial GBP 220 million commitment for conceptual design phase. <https://step.ukaea.uk/>

Korea’s KSTAR programme, China’s EAST and CFETR programmes, and India’s SST-2 programme all represent substantial national investments in fusion physics and modelling — each creating a potential τ integration partner within an IMAS-compatible framework.

10.2 Private fusion investment

The Fusion Industry Association’s 2023 report documents **USD 6.2 billion in private fusion investment** by that date, growing from near zero in 2018.⁴⁵ By 2025, total private investment has exceeded USD 10 billion per IAEA estimates.⁴⁶ The major private investors include Commonwealth Fusion Systems (USD 1.8 billion raised through 2023), Helion Energy (USD 570 million from OpenAI’s Sam Altman and others, with a USD 2.7 billion conditional commitment from Microsoft for power purchase), TAE Technologies (USD 1.2 billion), and General Fusion (USD 340 million).

Private fusion companies face a compressed validation timeline relative to public programs. They are building first-of-a-kind machines without the institutional modelling depth of ITER or EUROfusion. A τ twin platform that is public, IMAS-compatible, and validated against major public datasets would reduce their modelling investment burden and improve the quality of their physics predictions at each design decision point.

10.3 Cost scenarios for τ fusion-twin development

Scenario A — Integrated real-time control accelerator for one operational tokamak. Scoping τ as a shadow-mode twin for one major operating device (JET archive data, AUG, or JT-60SA), focused on disruption prediction and divertor heat-load management. Estimated cost: **USD 10–30 million over three to five years**, producing validated benchmarks, IMAS-compatible adapters, and demonstrated performance improvements on at least two benchmark tasks from ITER’s open R&D issue list.

Scenario B — Full τ fusion-design platform serving five major public programs and three private developers. A comprehensive IMAS-compatible τ twin platform covering the full coupled stack from equilibrium through disruption through divertor, with validated performance across multiple machine configurations, integrated with DTT design workflows and DEMO flight-simulator infrastructure. Estimated cost: **USD 50–150 million over ten years** — analogous in scale to a mid-size national fusion laboratory’s computational infrastructure investment.

Benefit-cost reference calculation. ITER operating costs run approximately EUR 5 million per month in experimental operation. If τ reduces the ITER validation timeline by six to eighteen months — by narrowing operating windows before expensive experimental scans, reducing the machine-pulse equivalents required to establish reliable high-Q operation, and enabling more of the scenario learning to happen in pre-pulse simulation — the avoided operational cost alone is **EUR 30–90 million**. This does not include the value of reduced first-wall damage risk, reduced disruption-driven schedule delays, or the design-closure acceleration for DEMO that better ITER physics knowledge enables. On this accounting, even the USD 50–150 million full-platform scenario has a cost-benefit ratio of two to five before counting the long-run option value of accelerating commercial fusion deployment.

⁴⁵Fusion Industry Association, **The Global Fusion Industry in 2023** (2023). USD 6.2 billion in cumulative private fusion investment by end of 2023; company tracker. <https://www.fusionindustryassociation.org/about-fusion-industry/>

⁴⁶IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

11 Governance Guardrails

Fusion is a safety-critical, publicly financed programme operating on multi-decade timescales. τ introduction must be designed with governance guardrails that match those characteristics. The following principles are non-negotiable from a public-good standpoint.

11.1 Simulation does not replace hardware completion

Even under the strongest τ assumption, simulation cannot substitute for: - materials qualification under neutron fluence and thermal cycling; - component manufacturing and quality assurance; - tritium breeding and fuel-cycle system validation; - blanket and divertor module testing under fusion-relevant loads; - remote-maintenance system validation; - or licensing-grade hardware evidence in national and international safety regimes.

τ compresses the learning cycle; it does not eliminate it. Any programme document that presents simulation results as substituting for hardware validation is misrepresenting the technology.

11.2 Safety-critical uses require multi-layer independent validation

Any τ application that informs disruption mitigation timing, wall-protection control decisions, or plasma-state inference in real-time control must be subject to: - independent physics validation by a panel drawn from ITER Organization, a national laboratory, and an independent academic group; - shadow operation with documented comparison against existing tools before any advisory or actuation role; - regular re-validation as machine operating conditions evolve; - clear failure-mode documentation and fallback to existing conservative control strategies.

11.3 Open standards must be reinforced, not undermined

Because ITER has committed to making IMAS an open, device-agnostic worldwide standard,⁴⁷ the public-good case for τ requires that its fusion applications reinforce that standard rather than creating a proprietary dependency. τ fusion-twin outputs should be IMAS-compatible, reproducible, and available under licences that prevent lock-in. A τ that becomes a black-box dependency inside safety-critical fusion control is not consistent with the public-interest mandate of the ITER programme.

11.4 Claims must be benchmarked, not narrated

Fusion has historically attracted over-optimistic claims about simulation capabilities. τ should therefore be introduced exclusively through: - explicit benchmark tasks with published evaluation criteria; - reproducible workflows using publicly available reference cases; - clear error envelopes and uncertainty quantification; - and cross-institution validation involving at least two independent groups.

No operational or design-advisory role for τ should be proposed until at least Stage 2 of the deployment ladder has been completed and independently assessed.

⁴⁷ITER Organization, **Release of IMAS Infrastructure and Physics Models as Open Source** (8 December 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>

11.5 Public and private benefits should be aligned, not in tension

The fusion ecosystem now spans public megaprojects and privately funded developers in ways that create potential conflicts between proprietary competitive advantage and public-interest knowledge sharing. A τ fusion programme should be structured so that: - benchmark data and validation protocols are publicly available; - private developers who adopt τ contribute to rather than draw down the public validation commons; - intellectual property arrangements do not prevent ITER, EUROfusion, or national laboratories from using τ -derived insights freely; - and the governance framework explicitly monitors for private capture of publicly funded development.

12 Five-, Ten-, and Twenty-Year Public-Good Horizon

12.1 In five years (2026–2031)

The realistic near-term contribution of τ to fusion is not fusion electricity on the grid, but measurably better use of the experimental time ITER and DTT will begin generating. The best-case five-year picture includes:

- IMAS-compatible τ shadow twins demonstrating improved scenario stress-testing and synthetic-diagnostic enrichment on at least two benchmark tasks from ITER's open R&D list;
- disruption-prediction benchmarks showing reduced false-positive rates relative to current tools on cross-machine data sets;
- divertor heat-load prediction benchmarks showing improved accuracy relative to empirical scaling laws on DTT-relevant configurations;
- an open fusion-twin benchmark commons with IMAS-compatible protocols accessible to both public and private fusion developers.

The public-good effect at five years is primarily in **research productivity** and **risk reduction** — more useful physics learned per experimental year, fewer machine-days lost to avoidable disruptions and unnecessary terminations, and a cleaner prior for DEMO design-team decision-making.

12.2 In ten years (2031–2036)

The hopeful medium-term picture under the τ assumption is a structural shift in how ITER and DTT are operated:

- ITER's research plan for D-T operation is informed by τ -derived operating-window maps that narrow the experimental search to physically grounded candidate scenarios, reducing the number of pulse-equivalents required to establish reliable high-Q operation;
- DTT operates as a validation machine for a τ -derived shortlist of advanced-divertor candidates rather than a broad exploratory search programme;
- DEMO's flight-simulator workflow carries smaller physics uncertainty into engineering design decisions, with documented improvements in design-margin confidence;
- private fusion developers have access to a public τ twin benchmark commons that reduces their individual modelling investment burden.

The quantitative effect at ten years is plausibly on the order of one to two experimental campaign cycles avoided or shortened for ITER and DTT combined — representing avoided operational costs in the range of EUR 50–200 million and schedule benefits worth multiples of that in option value.

12.3 In twenty years (2036–2046)

The large public-good possibility at the twenty-year horizon is that τ contributes materially to pulling the DEMO design-maturity timeline forward. Under a realistic-optimistic reading:

- DEMO or equivalent programmes reach production-grade design confidence with less empirical drift — fewer first-of-a-kind risks surviving into the engineering commitment phase;
- commercial fusion developers inherit a more thoroughly validated knowledge base, reducing the cost and risk of the transition from demonstration to commercial deployment;
- the fusion sector’s documented learning curve is steeper — more physics reliably discovered per year — creating a positive feedback loop between modelling confidence and experimental productivity.

If the IAEA’s long-horizon fusion-share range proves directionally correct,⁴⁸ the system-level value of even a modest learning-cycle gain at the DEMO and early commercial phase is potentially very large — comparable in economic terms to the value of a few years of earlier access to a clean-firm low-carbon electricity source at scale.

13 Benchmark Suite

A serious τ -for-fusion programme must stand or fall on explicit, publicly evaluable benchmarks. The following suite covers the eight most important validation dimensions.

13.1 Benchmark 1 — ITER baseline scenario reproduction

Reproduce a canonical ITER 15 MA D-T baseline scenario with IMAS-compatible outputs and compare against reference IMAS/HFPS workflows on equilibrium evolution, core transport, H&CD power balance, and plasma current profile.

13.2 Benchmark 2 — Cross-machine disruption detection and mitigation timing

Demonstrate whether τ improves portable disruption detection and prediction quality on cross-machine cases (JET, AUG, JT-60SA, KSTAR) relevant to ITER scale. Primary metric: false-positive rate reduction at fixed false-negative rate for disruptions occurring within the last 50 ms before thermal quench.⁴⁹

13.3 Benchmark 3 — Divertor power-flux width and 3-D field response

Benchmark divertor power-flux deposition width and the effect of 3-D resonant magnetic perturbation fields on divertor wetted area and peak heat flux, against SOLPS-ITER reference cases and available experimental data from JET, AUG, and EAST.⁵⁰

⁴⁸IAEA, **Fusion Energy in 2025: Six Global Trends to Watch** (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD 10 billion. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>

⁴⁹ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

⁵⁰ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

13.4 Benchmark 4 — Tungsten source, transport, and accumulation control

Demonstrate coupled prediction and feedback control of tungsten source terms (sputtering and ELM-driven erosion), SOL transport, and core accumulation under H-mode and various ELM-control strategies, compared against coupled EDGE2D-EIRENE and STRAHL reference cases.⁵¹

13.5 Benchmark 5 — H&CD, fast-particle, and burn-transition scenario

Benchmark actuator choices, scenario trajectory optimization, and control requirements for burn access and stable high-Q operation, including fast-particle-driven Alfvén mode activity, against reference IMAS/PCSSP workflow outputs.⁵²⁵³

13.6 Benchmark 6 — Synthetic diagnostics and control-model validation

Demonstrate τ -derived synthetic diagnostics for Thomson scattering, ECE, interferometry, and fast-magnetics channels, and validate inferred plasma-state accuracy against IMAS/PCSSP-compatible reference cases and experimental data.⁵⁴

13.7 Benchmark 7 — DTT heat-exhaust design-space narrowing

Demonstrate that τ can produce a ranked shortlist of advanced-divertor and impurity-seeding candidates for DTT experimental campaigns, with documented physical rationale, that a DTT expert panel assesses as covering the viable space at higher specificity than current empirical shortlisting methods.⁵⁵

13.8 Benchmark 8 — DEMO transient validation with the full discharge flight simulator

Demonstrate that τ improves confidence in transient and off-normal discharge sequences inside EUROfusion’s full discharge flight-simulator workflow, with quantified improvements in margin and uncertainty estimates for at least three first-of-a-kind engineering decisions in the DEMO conceptual design.⁵⁶

⁵¹ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

⁵²ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

⁵³ITER Organization, **Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan** (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten transport, ELM control, divertor power flux, IMAS/PCSSP models. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>

⁵⁴ITER Organization, **Integrated Modelling Programme** (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self-consistency goals, 2.2 PB/day data scale. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf

⁵⁵EUROfusion, **DTT Steps Up Progress Towards Tackling Fusion’s Power Exhaust Challenge** (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing components, operations by end of 2030. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>

⁵⁶EUROfusion, **Reviewing Today to Shape Tomorrow: Insights from European DEMO Technologies and Physics for the Road Ahead** (28 November 2025). High-fidelity codes and full discharge flight simulator feeding back into workflow models to support design maturity. <https://euro-fusion.org/eurofusion-news/reviewing-today-to-shape-tomorrow/>

14 Lighthouse Pilots

The following five pilots represent the highest-value, most tractable entry points for demonstrating τ capability to fusion stakeholders.

14.1 Pilot 1 — τ -IMAS Shadow Twin Challenge

A public benchmark pilot in which τ reproduces and stress-tests canonical ITER-style plasma scenarios using IMAS-compatible input/output. The pilot is run as a public challenge with ITER Organization participation and evaluation. Deliverable: a documented comparison of τ versus reference IMAS workflows on at least three scenarios from the ITER research plan, with public release of benchmark data and evaluation protocols.

14.2 Pilot 2 — Cross-Machine Disruption and Runaway Benchmark

A cross-machine benchmark pilot focused on disruption detection, runaway-electron generation threshold prediction, and mitigation-timing quality, conducted on JET D-T archive data, AUG disruption database, and KSTAR operational data. Conducted as a collaboration with CCFE, IPP Garching, and KFE, with independent evaluation by ITER Organization. Primary deliverable: a false-positive/false-negative characterization of τ disruption prediction relative to current best-in-class tools, with statistical significance assessment.

14.3 Pilot 3 — DTT Heat-Exhaust Narrowing Pilot

A focused programme developed in collaboration with the DTT programme team and EUROfusion, aimed at using τ to narrow the advanced-divertor and control design space before and alongside early DTT campaigns. Deliverable: a τ -derived shortlist of three to five divertor configuration and seeding strategy candidates, with documented physical rationale, submitted to the DTT programme review before DTT's first experimental campaign.

14.4 Pilot 4 — DEMO Workflow Co-Design Pilot

A collaboration with EUROfusion's DEMO team in which τ is used as a high-fidelity co-design and uncertainty-reduction engine inside the DEMO flight-simulator workflow. Deliverable: documented comparison of DEMO engineering decision confidence with and without τ twin support on at least three design decision points in the current DEMO conceptual design phase.

14.5 Pilot 5 — IMAS-Compatible Open Fusion Twin Commons

A public-interest repository of benchmark cases, IMAS-compatible adapters, synthetic-diagnostic modules, and evaluation protocols, accessible to both public fusion programs and private fusion developers. Governed by an open-source licence and a multi-institutional steering committee including ITER Organization representation, at least two national laboratories, and at least two private fusion developers. Deliverable: an operational public repository with documented use by at least five independent groups within three years of launch.

15 Risks and Mitigations

15.1 Risk 1 — τ claims do not hold at fusion-relevant coupling depth

Description: The τ assumptions may not survive contact with the specific multi-physics coupling complexity of a tokamak plasma — the combination of free-boundary MHD, turbulent transport, fast-particle dynamics, atomic and molecular physics of neutrals, and plasma-wall interaction may be sufficiently intractable that the certified-coarse-graining property does not deliver the claimed structural consistency.

Mitigation: Stage-gated deployment with explicit benchmark criteria at each stage. No operational or design-advisory role is proposed until Stages 1 and 2 are independently validated. If early benchmarks show τ is not materially superior to existing tools, later stages are not funded.

15.2 Risk 2 — Adoption barriers in a conservative institutional environment

Description: ITER and EUROfusion operate under strict engineering discipline and have invested decades in existing tools. Introducing a new physics substrate requires institutional credibility, long validation timelines, and champions inside the programme.

Mitigation: enter through IMAS compatibility rather than replacement; focus initial pilots on tasks where existing tools are acknowledged as insufficient (disruption cross-machine portability, real-time heat-load prediction); build institutional partnerships with ITER Integrated Modelling Programme management before proposing operational roles.

15.3 Risk 3 — Private fusion developers capture public development investment

Description: private fusion companies may use publicly funded τ development to create proprietary competitive advantages without contributing to the public knowledge commons.

Mitigation: explicit licence terms requiring publication of benchmark results from τ applications; governance framework ensuring private adoption contributes to rather than draws down public benchmark commons; contractual requirements for cross-validation data sharing as a condition of IMAS-compatibility licence.

15.4 Risk 4 — Safety-critical over-reliance before sufficient validation

Description: enthusiasm for τ 's theoretical properties leads to premature deployment in disruption mitigation or wall-protection roles before the validation evidence warrants it.

Mitigation: mandatory multi-layer independent validation requirement at Stage 2 before any control-advisory role; explicit prohibition on autonomous control actions based on τ inference before Stage 4 completion; continuous shadow operation period with human-supervised comparison before any real-time advisory integration.

15.5 Risk 5 — Benchmark tasks are designed to favour τ rather than to test it

Description: benchmark criteria defined by τ developers may be chosen to highlight strengths and avoid weaknesses, reducing the public-good informational value of the validation exercise.

Mitigation: benchmark criteria defined jointly with ITER Organization, an independent national laboratory panel, and an academic review group; all benchmark data published publicly; evaluation conducted by a panel with no financial interest in τ adoption outcomes.

16 Bottom Line

Fusion is a particularly strong τ opportunity — not because it is technically easy, but because the official institutions already describe the problem in almost the same terms τ claims to solve.

ITER's integrated-modelling programme specifies a high-fidelity coupled plasma simulator. EURO-fusion's DEMO programme needs a flight simulator with physically trustworthy coarse-graining. DTT's mission requires the ability to narrow a design space rather than search it exhaustively. The 2025 ITER R&D report enumerates, concretely and specifically, the multi-physics coupled problems where existing tools are insufficient. The IMAS open-source release creates the insertion architecture. The financial infrastructure is deployed.

The core public-good argument is correspondingly simple:

If τ can help transform fusion from a discipline that discovers too much of its operating truth only on expensive, rare, and irreplaceable machines into one that validates more of that truth before the machine is built or operated, then it shortens the learning cycle from ITER through DEMO and improves the probability that fusion becomes a practical clean-energy contributor sooner and more safely.

Learning-cycle compression by even one meaningful iteration — converting one empirical search campaign into a validation campaign, or pulling DEMO engineering design closure forward by two to four years — would be worth many multiples of the investment required to develop and publicly validate a τ fusion-twin platform.

The question for planners and funders is not whether that is a valuable public good. It self-evidently is. The question is whether the τ framework can deliver the physics-fidelity promise it makes — and the only way to answer that question is through the staged benchmark programme this paper has described.

17 References

Source: Full manuscript text integrated from Public-Good Briefing draft.

18 Dossier accountability addendum

The following addendum records the release-facing accountability layer for this dossier: claim boundaries, baseline evidence, upstream dependencies, translation assumptions, scenario bands, scorecard rationales, benchmark requirements, governance guardrails, and related Panta Rhei surfaces. It is intentionally downstream of the full source argument above.

Impact thesis

A Public-Good Briefing on how τ could advance fusion energy through plasma-physics digital twins, improved plasma control, heat-exhaust management, and the ITER-to-DEMO-to-commercialization pathway. The v3 impact thesis is conditional: a Tau-grade law-faithful plasma, control, heat-exhaust, and machine-design twin would become valuable if it improves benchmarked public decisions while preserving transparent uncertainty, reviewability, and governance control.

18.1 Public-good burden and baseline evidence

A Public-Good Briefing on how τ could advance fusion energy through plasma-physics digital twins, improved plasma control, heat-exhaust management, and the ITER-to-DEMO-to-commercialization pathway. The public-good burden is treated here as an institutional decision problem: existing agencies already monitor parts of the domain, but the operational handoff from data to timely, auditable action remains incomplete.

18.1.1 External evidence baseline

- **ITER Organization**, Release of IMAS Infrastructure and Physics Models as Open Source (8 December 2025) [10]: source-page evidence item.
- **ITER Organization**, Baseline Press Conference Summary (3 July 2024). New baseline: Start of Research Operation 2034, full magnetic energy 2036, D-D 2035, D-T 2039, tungsten first wall replacing beryll [7]: source-page evidence item.
- **ITER Organization**, Integrated Modelling Programme (14th ITER International School on Integrated Modelling, 30 June 2025). Includes HFPS, PDS, IMAS API compatibility for external physics models, self- [9]: source-page evidence item.
- **ITER Organization**, Required R&D in Existing Fusion Facilities to Support the New ITER Research Plan (ITR-25-005, 30 May 2025). Disruption detection and prediction, runaway-electron loads, tungsten tr [11]: source-page evidence item.
- **EUROfusion**, DTT Steps Up Progress Towards Tackling Fusion’s Power Exhaust Challenge (26 March 2025). DTT mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing compone [2]: source-page evidence item.
- **EUROfusion**, DEMO [1]: The Demonstration Power Plant. DEMO target: 300–500 MW net electricity to grid; closed fuel cycle; breeding blanket, divertor, first-wall, maintenance, and nuclear safety requirements.
- **EUROfusion**, Reviewing Today to Shape Tomorrow: Insights from European DEMO Technologies and Physics for the Road Ahead (28 November 2025). High-fidelity codes and full discharge flight simulat [3]: source-page evidence item.
- **IAEA**, Fusion Energy in 2025: Six Global Trends to Watch (28 October 2025). Fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above USD [6]:

source-page evidence item.

- **ITER Organization**, Engineering Basis Handbook, Volume 1, Chapter 4 (2026). Plasma-facing components removing up to 850 MW of thermal power; unprecedented R&D required for high-heat-flux technologies [8]: source-page evidence item.
- **EUROfusion / UKAEA**, JET World Record: 59 Megajoules of Fusion Energy (February 2022). Joint press release, Culham Centre for Fusion Energy and EUROfusion [4]: source-page evidence item.

18.2 Current institutional landscape

The relevant landscape includes public agencies, research infrastructures, standards bodies, development-finance channels, and domain review communities represented in the evidence base, including EUROfusion, ITER Organization. These references are evidence and adoption surfaces, not endorsements or deployment partners.

18.3 Capability gap

The practical gap is a benchmarkable translation gap: current systems expose useful data or partial models, but they do not yet provide a single law-faithful, bounded-error decision layer for law-faithful plasma, control, heat-exhaust, and machine-design twin.

18.4 Tau framework dependency map

Surface	Role in this dossier
Build the Tau-Kernel	finite address and scalar foundation
Recover Core Mathematics	mathematical bridge and model interface
Derive Physics	physical readout and domain translation candidate
Results lane	upstream consequences to be mapped precisely during release preparation
direct-registry-mapping-withheld	no direct Registry object is asserted until a substantive Corpus mapping is available
public-docs-mapping-withheld	TauLib module links are asserted only where public documentation exposes a clear surface
Release Manifest	release baseline
Predictions and Falsification	empirical accountability route

18.5 Translation assumptions and missing engineering

Required domain model: **law-faithful plasma, control, heat-exhaust, and machine-design twin.**

First benchmarkable test: IMAS-compatible scenario, disruption, divertor, and synthetic-diagnostic tasks against ITER/EUROfusion benchmark workflows.

- domain-specific model construction
- data ingestion and validation
- benchmark harness
- pilot protocol

- independent review workflow







18.6 Impact mechanism chain

Public-good burden → external evidence baseline → τ capability hypothesis → upstream Results / Corpus / Verify dependency → translation assumptions → benchmarked pilot → governed adoption pathway.

18.7 Scenario bands

Band	Scenario summary	Confidence
Conservative	A narrow shadow-mode pilot improves one bounded decision task for Fusion Digital Twins, Plasma Control, Heat Exhaust, and ITER → DEMO → Commercialization without operational authority.	medium
Realistic	A reviewed prototype strengthens several public-sector workflows for Fusion Digital Twins, Plasma Control, Heat Exhaust, and ITER → DEMO → Commercialization after benchmark comparison with incumbent systems.	medium-low
Optimistic	A reusable public-good intelligence layer becomes plausible for Fusion Digital Twins, Plasma Control, Heat Exhaust, and ITER → DEMO → Commercialization after external validation and transparent governance review.	low

18.8 Impact scorecard

Public-good scale	 5/5	The affected public-good burden is large or institutionally significant within the portfolio.
Tau fit	 4/5	The proposed pathway depends on coupled state, bounded uncertainty, and compositional modelling rather than isolated prediction alone.
Evidence proximity	 5/5	The evidence base is anchored in public institutions, official monitoring systems, or established scientific reviews.
Measurability	 3/5	A first benchmark can be framed against incumbent public datasets, institutional records, or operational decision metrics.
Adoption readiness	 2/5	Adoption remains conditional on domain review, governance fit, data access, and institutional integration.
Equity leverage	 3/5	The pathway can prioritize underserved or vulnerable populations where public access and safeguards are built in.

18.9 Candidate pilot pathways

shadow-mode fusion-twin benchmark with public machine data and independent plasma-physics review

18.10 Benchmark suite and success metrics

Type	Incumbent line	base- Required benchmark	Tau Success metric	Validator
translation benchmark	current public or institutional systems in the domain	IMAS-compatible scenario, and diagnostic against ITER/EU-ROfusion benchmark workflows	pre-registered accuracy, latency, uncertainty, or decision-quality metric	independent domain reviewers
governance benchmark	existing audit, disclosure, and reporting practice	transparent assumption, data, model, and failure-mode disclosure	reviewable evidence pack and adverse-outcome protocol	public-sector or expert governance panel
equity benchmark	current service-quality, or exposure disparities	documented path-way for underserved or vulnerable without exclusion	distributional benefit and risk review before pilot expansion	equity, community, or public-interest review process

18.11 Governance and risk guardrails

- Human oversight for any operational use.
- Public benchmark disclosure before institutional adoption.
- Equity access review for underserved or vulnerable communities.
- Data-rights and privacy controls for operational datasets.
- Misuse-prevention and adverse-outcome monitoring.
- Adverse-outcome monitoring with a documented escalation path.
- External domain review before pilot expansion.

18.12 Related Results / Corpus / Verify / Publications

This dossier is downstream of Results, Corpus, Verify, and Publications surfaces. It is not a Registry object. Direct Registry or TauLib links are asserted only where the mapping is substantive rather than decorative.

18.13 Bibliography and external evidence

References

- [1] EUROfusion. Demo. <https://euro-fusion.org/programme/demo/>, 2026. The Demonstration Power Plant. DEMO target: 300–500 MW net electricity to grid; closed fuel cycle; breeding blanket, divertor, first-wall, maintenance, and nuclear safety requirements.
- [2] EUROfusion. Dtt steps up progress towards tackling fusion’s power exhaust challenge (26 march 2025). dtt mission, heat-exhaust focus, advanced divertor concepts, tungsten plasma-facing compone. <https://euro-fusion.org/partner-news/dtt-steps-up-progress-towards-tackling-fusions-power-exhaust-challenge/>, 2026. source-page evidence item.
- [3] EUROfusion. Reviewing today to shape tomorrow: Insights from european demo technologies and physics for the road ahead (28 november 2025). high-fidelity codes and full discharge flight simulat. <https://euro-fusion.org/eurofusion-news/reviewing-today-to-shape-tomorrow/>, 2026. source-page evidence item.
- [4] EUROfusion / UKAEA. Jet world record: 59 megajoules of fusion energy (february 2022). joint press release, culham centre for fusion energy and eurofusion. <https://euro-fusion.org/eurofusion-news/european-researchers-smash-fusion-energy-record/>, 2026. source-page evidence item.
- [5] Thorsten Fuchs and Anna-Sophie Fuchs. τ for fusion digital twins, plasma control, heat exhaust, and iter \rightarrow demo \rightarrow commercialization. <https://panta-rhei.site/impact/papers/fusion-digital-twins-plasma-control-heat-exhaust-iter-demo/>, 2026. Current public full-text source for dossier fusion-digital-twins-plasma-control-heat-exhaust-iter-demo.
- [6] IAEA. Fusion energy in 2025: Six global trends to watch (28 october 2025). fusion at 10–50% of global electricity by 2100; more than 160 facilities globally; private investment above usd. <https://www.iaea.org/newscenter/news/fusion-energy-in-2025-six-global-trends-to-watch>, 2026. source-page evidence item.
- [7] ITER Organization. Baseline press conference summary (3 july 2024). new baseline: Start of research operation 2034, full magnetic energy 2036, d-d 2035, d-t 2039, tungsten first wall replacing beryll. https://www.iter.org/sites/default/files/media/2024-07/baseline_press_conference_summary_july-2024_b.pdf, 2026. source-page evidence item.
- [8] ITER Organization. Engineering basis handbook, volume 1, chapter 4 (2026). plasma-facing components removing up to 850 mw of thermal power; unprecedented r&d required for high-heat-flux technologies. https://www.iter.org/sites/default/files/media/2026-01/vol1.1_ch.04_role_and_distinctive_feature_dv4qgd_v2_0.pdf, 2026. source-page evidence item.
- [9] ITER Organization. Integrated modelling programme (14th iter international school on integrated modelling, 30 june 2025). includes hfps, pds, imas api compatibility for external physics models, self-. https://www.iter.org/sites/default/files/media/2025-07/i-2_pinches.pdf, 2026. source-page evidence item.
- [10] ITER Organization. Release of imas infrastructure and physics models as open source (8 december 2025). <https://www.iter.org/node/20687/release-imas-infrastructure-and-physics-models-open-source>, 2026. source-page evidence item.
- [11] ITER Organization. Required r&d in existing fusion facilities to support the new iter research plan (itr-25-005, 30 may 2025). disruption detection and prediction, runaway-electron loads, tungsten tr. <https://www.iter.org/sites/default/files/media/2025-06/itr-25-005-final.pdf>, 2026. source-page evidence item.

- [12] Panta Rhei Research Program. Public-good briefing landing page. <https://panta-rhei.site/publications/research-briefings/public-good/fusion-digital-twins-plasma-control-heat-exhaust-iter-demo/>, 2026.



Panta Rhei Research Program

Public-Good Impact Dossier

**τ for Fusion Digital Twins, Plasma Control, Heat Exhaust,
and ITER → DEMO → Commercialization**

Dossier ID: PGID-ENER-04 Portfolio: Energy Release: May 2026
publication-ready release

Conditional scenario map. Domain review pending. Deployment, product,
validation, certified-impact, and policy-commitment claims are not made.

Public contact and review routes

Website: panta-rhei.site

Contact: panta-rhei.site/engage/contact/

Public discussion: github.com/orgs/Panta-Rhei-Research/discussions

General: hello@panta-rhei.site

Corrections: errata@panta-rhei.site

Media: press@panta-rhei.site