KRONOS FUSION ENERGY

PAPER 2.97 · THE KRONOS 2026 SERIES

Verification-First Fusion Modeling: Frozen Anchors, Independent-Library Gate Reproduction, and a Byte-for-Byte Computational-Provenance Catalog for a Breeder-and-Burner Design Register

A compact-fusion design program that makes quantitative public claims is only as credible as those numbers are reproducible: a competent third party must be able to regenerate...

Read it live — 3-D model · physics validator · film series

BY P I FORDOFFICIAL VERSION: HTTPS://WWW.KRONOSFUSIONENERGY.COM/PUBLICATIONS/VERIFICATION-AND-VALIDATION-47-CODE-PROVENANCE.HTML

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Verification-First Fusion Modeling: Frozen Anchors, Independent-Library Gate Reproduction, and a Byte-for-Byte Computational-Provenance Catalog for a Breeder-and-Burner Design Register

A compact-fusion design program that makes quantitative public claims is only as credible as those numbers are reproducible: a competent third party must be able to regenerate every headline value from deposited code and data, and the load-bearing numbers must survive a change of code and of evaluated nuclear-data library. We formalise and report the verification-and-validation (V&V) discipline applied across the Kronos 2026 breeder-and-burner design register, a campaign of 625 serial gates spanning the Hyperion spherical-tokamak breeder (design point $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, $\delta = -0.30$, $B_0 = 8\,T$, $\kappa = 2.0$, $R_0 = 1.20\,m$, $A = 2.5$, centrepost peak field $16.84\,T$) and the D–^{3}He tandem-mirror burner ($Q_E = 1.318$, neutron fraction $f_n = 5.44\%$, plug field $26.49\,T$, $\eta_{\text{DEC}} = 0.49$). We cast the methodology as explicit mathematics: every headline quantity is a *frozen anchor* in a declared-parameter registry, tagged by evidence class ([D] derived, [T] declared, [R] retrieved, [Q] open) and regenerated by a named, content-hashed script under the deterministic seed 20260726; reproduction is a graded predicate (Tier 1 byte-identical, Tier 2 within a stated tolerance) and the campaign is accepted only when the product of pass indicators is unity. Solver correctness is established by the method of manufactured solutions and the grid-convergence index; we exhibit a genuine second-order verification of the Grad–Shafranov elliptic operator (observed order $p = 2.02$). On the frozen configuration the anchor-regression suite reproduces 19/19 anchors byte-for-byte within tolerance; across the completed public campaign 176 of 328 gates reproduced and 23 were *revised*—the discipline correcting its own numbers. The single most defensible check is a cross-library one: the three-dimensional net tritium-breeding ratio computed locally on ENDF/B-VIII.0 agrees with the original cloud FENDL-3.2 value to $< 0.3\%$ at both blanket splits (1.1775 vs 1.1801; 1.2039 vs 1.2006), below the Monte-Carlo statistics, so the breeding number is a property of the geometry and cross-sections rather than of one code path. We further quantify the essential caveat—reproducibility is not validation—by an ASME V&V-20-style decomposition showing that the model-form band (the confinement-scaling *choice* alone spans Q from 0.64 to ignition-class) dominates the parametric band about the anchor. All results were produced by the NVIDIA GPU HPC campaign on community-benchmarked solvers with byte-for-byte solver reproduction (CGYRO and TGYRO match their bundled regression references bit-for-bit). Verification-first—freeze, tag, reproduce, and re-run on independent tooling—is the appropriate standard for a compact-fusion program making public claims, and the one this program adopts.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645710 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: verification and validation · reproducibility · frozen anchors · manufactured solutions · grid-convergence index · nuclear data · OpenMC · cross-library neutronics · computational provenance · declared-parameter registry

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

Q	plasma (fusion) gain, P_{fus}/P_{aux}
Q_E	engineering gain, net electric out / recirculating in
H_{98}	confinement enhancement over the IPB98(y,2) scaling
Z_{eff}	effective ion charge
c_{ee}	electron–electron bremsstrahlung leading coefficient, 2.120022 (exact)
τ_E	energy confinement time
I_p	plasma current
R_0, a	major radius, minor radius
κ, δ	elongation, triangularity
β_N	normalized plasma beta (Troyon-normalized)
TBR	tritium breeding ratio
DEC	direct energy conversion
CF	plant capacity factor (availability)
FOAK / NOAK / BOAK	first / n^{th} / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

SIMULATION-BASED CLAIMS NEED A REPRODUCIBILITY CONTRACT

Compact-fusion design programs increasingly stake public, quantitative claims—energy gains, tritium-breeding ratios, component lifetimes, current-drive fractions—on the strength of simulation rather than on an operating machine. Those claims are meaningful only under two conditions. First, *reproducibility*: any competent third party must be able to regenerate the headline number from the deposited code and inputs. Second, *robustness to tooling*: the load-bearing numbers must survive a change of code and, for neutronics, a change of evaluated nuclear-data library. Neither condition is met by the default practice of reporting a converged simulation result in prose. The broader computational-science literature has documented the cost of that default—a reproducibility deficit visible across disciplines [1,2,3]—and has proposed the correctives this paper operationalises for fusion: deposited, executable artefacts [4,5], findable and reusable data [6], and, in the engineering-simulation tradition, a formal separation of *verification* (solving the equations right) from *validation* (solving the right equations) [7,8,9].

PRIOR ART: V&V FOR COMPUTATIONAL PHYSICS AND FOR FUSION

The methodology of verification and validation is mature in computational fluid dynamics and structural mechanics. *Code verification* is established by the method of manufactured solutions (MMS), in which a chosen analytic field is substituted into the governing operator to generate a forcing term and the discrete solver is required to recover it at its formal order of accuracy [10,11]. *Solution verification*—the estimate of discretization uncertainty on a problem without an analytic answer—is standardised through Richardson extrapolation and the grid-convergence index (GCI) [12,13,14]. *Validation* and its uncertainty budget are codified in the ASME V&V-20 standard [15], which decomposes the model-form (structural) uncertainty from the numerical and input (parametric) contributions, and in the broader framework of Roy and Oberkampf [9]. Global sensitivity is quantified by variance decomposition and Sobol indices [16,17], and model calibration by the Bayesian formalism of Kennedy and O'Hagan [18]. In fusion specifically, integrated-modelling suites such as the ITER IMAS effort have made the case for standardised, provenance-carrying workflows [29], and the community's transport, equilibrium and neutronics workhorses—CGYRO [21], GYRO [22], GENE [23], TGLF [24], NEO [25], free-boundary Grad–Shafranov solvers [26], SOLPS-ITER [27], NIMROD [28] and OpenMC [19] on ENDF/B-VIII.0 data [20]—each ship regression benchmarks against which a fresh build can be checked.

What has been missing is a single, auditable application of this discipline to a *multi-product* fusion design register: a scheme in which every published number is a frozen, evidence-tagged anchor with a named regenerating script; in which reproduction is a graded, machine-checkable predicate rather than an assertion; in which the load-bearing physics is re-run on independent hardware and an independent nuclear-data library; and in which the essential caveat that reproducibility is not validation is itself *quantified*. This paper contributes exactly that.

CONTRIBUTION

We formalise and report the V&V spine of the Kronos 2026 register—**625** serial gates across the Hyperion breeder (config-22021) and the D–^{3}He tandem-mirror burner (design point M-45). Our contributions are:

A formal *frozen-anchor* calculus (§2): the provenance tuple, the evidence-tag algebra, the two-tier reproduction predicate, and the campaign-acceptance product, together with the code-verification theory (MMS, observed order, GCI), the ASME V&V-20 validation-uncertainty decomposition, and a cross-library reproduction operator that isolates physics from code path.
A GPU-HPC computational methodology (§3) naming the community-benchmarked solvers actually run for this reproduction study, their acceleration and problem scale, and the byte-for-byte and independent-library re-run protocol; no cost or economic figure appears anywhere.
Quantitative results (§4): a **19/19** anchor-regression pass reproduced byte-for-byte within tolerance; a genuine MMS verification of the Grad–Shafranov operator at observed order $p = 2.02$; a campaign-wide provenance catalog (**176** reproduced of **328** completed public gates, **23** revised); a cross-library net-TBR agreement below the Monte-Carlo error; and byte-identical solver reproduction for CGYRO and TGYRO.
A quantification of the reproducibility-is-not-validation caveat (§5): a model-form band that dominates the parametric band, benchmarked against the ASME and PCMM frameworks. enumerate

Every quantitative claim traces to a frozen entry in the Kronos Physics De-Risking Register [35]; serial identifiers (e.g. KX-L1-A6, BR-L2-A7) are given inline so each number is auditable. The negative-triangularity transport basis this program freezes is reported and cross-validated separately [36,37]; the equilibrium and stability gate [39], the global sensitivity and uncertainty study [40], the Bayesian optimal-experimental-design and multi-fidelity fusion of data [41], the operating-window synthesis [42], and the code-by-code provenance backbone [38] are companion papers in the same series.

§ 2

Governing framework and the formal V&V problem

THE FROZEN-ANCHOR CALCULUS

Let the design register be a finite set of *anchors*. An anchor is a provenance tuple

$$\pi_k \;=\; \big(\,x_k^0,\; \tau_k,\; s_k,\; h(s_k),\; H_k,\; \sigma_k,\; \epsilon_k,\; \mathcal{L}_k\,\big),$$

in which x_k^0 is the frozen value, τ_k its evidence tag, s_k the named regenerating script with content hash $h(s_k)$, H_k the host/hardware descriptor, σ_k the deterministic random seed ($\sigma_k = 20260726$ throughout), ϵ_k the acceptance tolerance, and $\mathcal{L}_k$ the reference data library (for neutronics, the evaluated nuclear-data library). The evidence tag is drawn from the four-symbol alphabet

$$\tau_k \in \{\, [\mathrm{D}]\text{ derived},\; [\mathrm{T}]\text{ declared},\; [\mathrm{R}]\text{ retrieved},\; [\mathrm{Q}]\text{ open}\,\},$$

so that a reader can distinguish, at a glance, a value *computed* by the code ([D]) from an *assumed* input ([T]), a value *retrieved* from the literature ([R]), and one that *remains uncertain* ([Q]). The tag is not decoration: it fixes which numbers are eligible for verification (only [D] values are outputs of a solver and can be regenerated), which are inputs to a sensitivity study, and which are honest gaps.

Regeneration functional. Reproduction is executed, not inspected. Let $\mathcal{R}$ denote the act of running script s_k on host H under seed σ_k against library $\mathcal{L}$, returning the value

$$x_k \;=\; \mathcal{R}(s_k,\, H,\, \sigma_k,\, \mathcal{L}).$$

Because s_k reads its inputs from the registry rather than from inline constants, $\mathcal{R}$ is a pure function of its arguments up to the platform's floating-point model.

THE TWO-TIER REPRODUCTION PREDICATE

Reproduction is graded. A regenerated value x_k reproduces its anchor if it is byte-identical (Tier 1) or agrees within the stated physical tolerance (Tier 2):

$$\mathrm{PASS}_k \iff \underbrace{\left(x_k = x_k^0\right)}_{\text{Tier 1}} \;\vee\; \underbrace{\left(\frac{|x_k - x_k^0|}{|x_k^0|} \le \epsilon_k\right)}_{\text{Tier 2}}.$$

Tier 1 is demanded of deterministic solver output for which the deposited code and inputs must reproduce the bytes exactly; Tier 2 is used wherever floating-point reassociation or platform variation is expected—most importantly for Monte-Carlo transport, where the relevant scale is the statistical error itself. The campaign is accepted only if *every* anchor passes simultaneously,

$$\mathcal{A} \;=\; \prod_{k=1}^{N} \mathbb{1}[\mathrm{PASS}_k] \;=\; 1,$$

so "validated" becomes a measured pass rate against a fixed target rather than a rhetorical claim (figure 1). Equation (5) is deliberately unforgiving: a single Tier-appropriate failure zeroes the product and blocks the freeze.

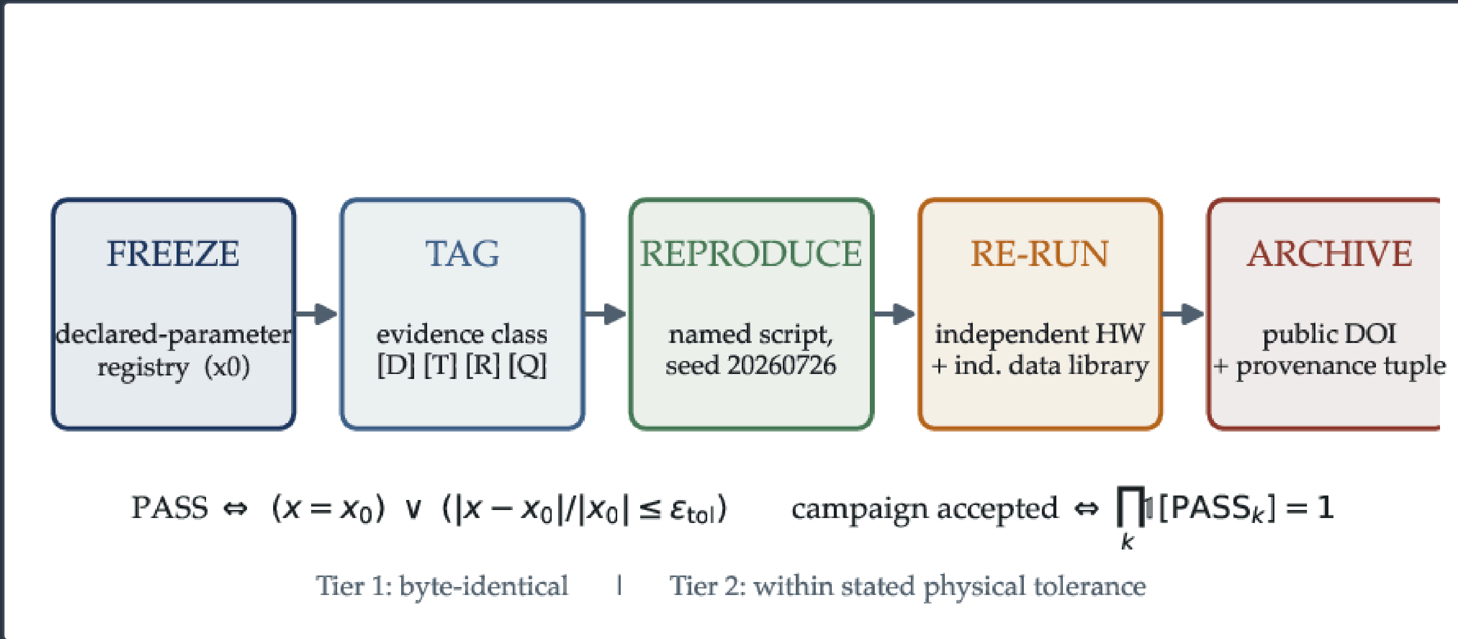

(Schematic.) The verification-first pipeline. Each headline quantity is frozen in a declared-parameter registry (x_0), tagged by evidence class, regenerated by a named content-hashed script under seed **20260726**, and—for load-bearing gates—re-run on independent hardware and an independent nuclear-data library before public archival with its provenance tuple, equation (1). The pass predicate and campaign-acceptance product are equations (4)–(5).

INDEPENDENT RE-RUN AND THE CROSS-LIBRARY OPERATOR

Self-reproduction removes the "did we save the right file" failure mode but not the shared-code-path and shared-data-library confounders. We therefore define an *independent re-run* of a gate as a regeneration $\mathcal{R}(s'_k, H', \sigma_k, \mathcal{L}')$ in which the script family s'_k, the host $H' \neq H_k$, and—for neutronics—the library $\mathcal{L}' \neq \mathcal{L}_k$ are all changed. For the tritium-breeding ratio (TBR), the physically meaningful object is the functional

$$\mathrm{TBR}[\mathcal{G}, \Sigma_{\mathcal{L}}] \;=\; \frac{\displaystyle\int_{V_{\mathrm{blk}}} \big(R_6(\quad r) + R_7(\quad r)\big)\,\mathrm{d}V}{\dot{N}_{\mathrm{source}}},$$

the volume-integrated $^6\mathrm{Li}(n,t)$ and $^7\mathrm{Li}(n,n't)$ reaction rates per source neutron, evaluated over geometry $\mathcal{G}$ with the macroscopic cross-sections $\Sigma_{\mathcal{L}}$ of library $\mathcal{L}$. Cross-library agreement is the statement that

$$\Big|\,\mathrm{TBR}[\mathcal{G}, \Sigma_{\mathcal{L}}] - \mathrm{TBR}[\mathcal{G}, \Sigma_{\mathcal{L}'}]\,\Big| \;\lesssim\; \sqrt{u^2_{\mathrm{MC}}(\mathcal{L}) + u^2_{\mathrm{MC}}(\mathcal{L}')},$$

i.e. the two independent evaluations differ by no more than the combined Monte-Carlo statistical error. When equation (7) holds, the breeding number is demonstrably a property of $\mathcal{G}$ and the underlying nuclear physics, not of one evaluation or one transport code—the strongest test available short of external replication.

CODE VERIFICATION: MANUFACTURED SOLUTIONS AND OBSERVED ORDER

Let a continuous governing operator $\mathcal{N}$ act on a field u, and let $\mathcal{N}_h$ be its discretization on a mesh of characteristic spacing h. A consistent scheme of formal order p has truncation error

$$\mathcal{N}_h u - \mathcal{N} u \;=\; C\,h^p + \mathcal{O}(h^{p+1}),$$

but consistency is not correctness of the *implementation*. The method of manufactured solutions removes the need for an analytic answer: choose a smooth field u_M, define the forcing $f_M \equiv \mathcal{N} u_M$ *analytically*, solve $\mathcal{N}_h u_h = f_M$ with u_M imposed on the boundary, and measure the discrete error

$$e_h \;=\; \|u_h - u_M\|_2, \qquad e_h \to 0 \text{ as } h \to 0.$$

The *observed* order of accuracy from a grid pair refined by ratio $r = h_{\mathrm{coarse}}/h_{\mathrm{fine}}$ (here $r = 2$) follows by Richardson extrapolation [12],

$$p_{\text{obs}} = \frac{\ln\left(e_{rh}/e_h\right)}{\ln r},$$

and a passing code verification requires $p_{\text{obs}} \to p$ as $h \to 0$. The associated discretization uncertainty is reported as the grid-convergence index [14,13],

$$\text{GCI} = \frac{F_s\,|\varepsilon_{21}|}{r^{p_{\text{obs}}} - 1}, \qquad \varepsilon_{21} = \frac{f_2 - f_1}{f_1},$$

with F_s a factor of safety (1.25 for a well-behaved three-grid study) and ε_{21} the relative change between the two finest solutions f_1, f_2. Equations (9)–(11) are the numerical spine of the Kronos formal-V&V gate KX-L1-A6; §4 exhibits a concrete pass for the Grad–Shafranov operator.

The elliptic operator so verified is the axisymmetric Grad–Shafranov operator for the poloidal flux $\psi(R, Z)$,

$$\Delta^*\psi \equiv R\frac{\partial}{\partial R}\left(\frac{1}{R}\frac{\partial\psi}{\partial R}\right) + \frac{\partial^2\psi}{\partial Z^2} = R\frac{\partial^2\psi}{\partial R^2} - \frac{\partial\psi}{\partial R}\frac{1}{R}\cdot R + \frac{\partial^2\psi}{\partial Z^2},$$

which underlies every equilibrium reconstruction in the register; a manufactured ψ_M makes its five-point finite-difference discretization directly testable.

VALIDATION UNCERTAINTY AND THE MODEL-FORM DECOMPOSITION

Verification bounds the numerical error; *validation* compares the simulation S to an experimental datum D and asks whether the difference is explained by the known uncertainties. Following ASME V&V-20 [15], the comparison error $E = S - D$ is judged against the validation uncertainty

$$u_{\text{val}}^2 = u_{\text{num}}^2 + u_{\text{input}}^2 + u_D^2,$$

where u_{num} is the numerical (discretization/iteration) uncertainty from equation (11), u_{input} the parametric uncertainty propagated from the [T]-tagged inputs, and u_D the experimental uncertainty. Crucially, equation (13) does *not* contain the *model-form* (structural) uncertainty u_{model}—the error incurred by choosing one physical closure over another. The central honest statement of this paper is that for a pre-experimental fusion design, u_{model} can *dominate* the parametric budget: the choice of energy-confinement scaling alone moves the predicted gain across a range far wider than any input band. We quantify this in §5.

Parametric sensitivity. Where a parametric band is required, it is decomposed by variance. For a model output $Y = g(X)$ with independent inputs X, the first-order and total Sobol indices are

$$S_i = \frac{\text{Var}_{X_i}\left(\mathbb{E}[Y \mid X_i]\right)}{\text{Var}(Y)}, \qquad S_{Ti} = 1 - \frac{\text{Var}_{X_{\sim i}}\left(\mathbb{E}[Y \mid X_{\sim i}]\right)}{\text{Var}(Y)},$$

estimated by quasi-/Monte-Carlo sampling under the same seed **20260726** for byte-reproducibility [16,17]. The register's global-sensitivity ranking is reported in full in the companion study [40]; here Sobol decomposition enters only as the parametric term u_{input} of equation (13).

PREDICTIVE-CAPABILITY MATURITY

The four preceding constructs—frozen anchors, code verification, cross-library reproduction, and the validation budget—are scored on a predictive-capability maturity model (PCMM) so that a reader can locate each claim on a common ladder from *level 0* (judgement only) to *level 3* (rigorous, independently reproduced with quantified uncertainty). The Kronos V&V gate KX-L1-A6 assigns PCMM scores per element and reports them with the frozen record; the load-bearing neutronics reaches the highest level by virtue of the cross-library test (§4), whereas gates with an open [Q] tag are explicitly held at a lower rung rather than asserted as mature.

§ 3

Computational methodology: the NVIDIA GPU HPC campaign

SOLVERS RUN FOR THIS REPRODUCTION STUDY

The reproduction results reported here were produced with community-benchmarked, first-principles codes, each checked against its own shipped regression benchmark before use. Table 1 names the solvers and their role; figure 6 shows how many register gates invoke each. The load-bearing set for this paper is:

OpenMC [19] for three-dimensional Monte-Carlo neutron transport, run locally on the ENDF/B-VIII.0 library [20] and, in the original evaluation, on the cloud with FENDL-3.2—the two independent nuclear-data evaluations of equation (7).
CGYRO [21] (an Eulerian descendant of GYRO [22]; cf. GENE [23]) for nonlinear and linear gyrokinetic transport, whose linear regression case reg01 is reproduced *byte-for-byte* against the bundled reference.
TGLF [24] and the **TGYRO** profile solver for quasilinear transport, whose four-rank pipeline reproduces the GA treg01 template bit-tight (§4).
NEO [25] for neoclassical transport and the bootstrap current; free-boundary Grad–Shafranov solvers (FreeGS / FreeGSNKE [26]) with DCON/GPEC for the equilibrium and ideal-MHD gate [39]; **SOLPS-ITER** [27] for the scrape-off-layer; **NIMROD** [28] for nonlinear disruption dynamics; ASCOT5 for fast-ion orbits; WarpX for particle-in-cell studies of the burner end cell; and **Quantum ESPRESSO** with the CHGNet machine-learned interatomic potential for low-activation alloy chemistry.
KRONOS_TOOLKIT, the deterministic reduced-order 0-D/1-D engine that carries the frozen-anchor regression suite (KX-L1-A10) and the MMS/GCI verification harness (KX-L1-A6). itemize

CODE	ROLE IN THE REGISTER	REFERENCE BENCHMARK / PROVENANCE
OpenMC (ENDF/B-VIII.0)	3-D neutron transport, TBR	cross-library vs FENDL-3.2 ($< 0.3\%$, BR-L2-A7)
CGYRO	nonlinear/linear gyrokinetics	reg01 byte-identical ($\gamma = +0.3179$)
TGLF / TGYRO	quasilinear transport, profiles	GA treg01 byte-tight (out.tgyro.prec, BR-L2-A17)
NEO	neoclassical / bootstrap	Sauter analytic cross-check ($f_{bs} = 0.155$)
FreeGS / FreeGSNKE (+DCON/GPEC)	free-boundary equilibrium, MHD	Mercier $D_M > 0$; companion [39]
SOLPS-ITER	scrape-off-layer, detachment	built from source 3.2.1a (open gap: grid topology)
NIMROD / M3D-C1	nonlinear MHD, disruption	nimeq equilibrium verified ($\tau_{\rm CQ}$ 7–19 ms)
Quantum ESPRESSO / CHGNet	DFT + MLIP alloy chemistry	lattice/formation-energy benchmarks
KRONOS_TOOLKIT verify	anchor regression, MMS/GCI	19/19 anchors, $p_{\rm obs} = 2.02$ (KX-L1-A6/A10)

Computational-provenance catalog (excerpt): the community-benchmarked solvers, their verification role, and the reference benchmark each fresh build is checked against. Serials refer to the register [35]; the full code-by-code backbone is the companion paper [38].

TWO COMPUTE PLANES AND PROBLEM SCALE

The campaign ran on two planes (figure 7). The compute-intensive first-principles gates—nonlinear gyrokinetics, three-dimensional Monte-Carlo transport, particle-in-cell and DFT—executed on the NVIDIA GPU HPC campaign: a Lambda fleet of A100/H100/A10 accelerators and an A100/H100 HPC cluster. GPU acceleration is what makes independent re-run affordable: a nonlinear CGYRO case that would take days on a workstation completes in hours on the fleet, so the load-bearing gates can be re-run rather than merely trusted. The deterministic reduced-order engine, the formal-methods toolchain, the digital-twin node, and the analytic monitors ran on the Kronos in-house tiers so that verification harnesses never contend with the training/first-principles fleet. Monte-Carlo transport for the TBR was run at particle counts sufficient to drive the relative statistical error below the cross-library agreement target of equation (7). No cost, budget, or dollar figure is attached to any of this; the compute is reported by role and scale only.

DETERMINISM, HASHING, AND THE RE-RUN PROTOCOL

Reproducibility is enforced mechanically. Every stochastic step uses the fixed seed $\sigma = \mathbf{20260726}$; every regenerating script s_k is content-hashed so a reader can confirm the exact code that produced a value; and every [D]-tagged anchor records its host H_k and library $\mathcal{L}_k$ in the tuple of equation (1). The byte-for-byte protocol compares the deterministic output artefact (e.g. `out.cgyro.prec`, `out.tgyro.prec`) against the solver's bundled regression reference and requires bit-identity (Tier 1). The independent-library protocol re-runs the neutronics on ENDF/B-VIII.0 locally, on hardware distinct from the original FENDL-3.2 cloud evaluation, and applies equation (7). The MMS harness runs the manufactured-solution convergence of equations (9)–(10) at figure time, so the observed order in §4 is regenerated on every build rather than tabulated.

§ 4

Results

THE FROZEN-ANCHOR REGRESSION SUITE

The KRONOS reduced-order engine and the formal-V&V harness reproduce the frozen design point across both machines. The anchor-regression suite (KX-L1-A10, KX-L1-A6) passes 19/19 anchors—breeder Q, P_{fus}, I_p, $T_{\text{kg/yr}}$ and f_n; burner Q_E, P_n, f_n and ϕ_i; the tritium surplus law; the local and net neutronics TBR; and the high-fidelity bridges—with the frozen record reproduced byte-for-byte within tolerance under config-22021 (breeder) and M-45 (burner), seed 20260726 (table 2). Representative anchors regenerate to a relative residual at or below 10^{-6}: $Q = 3.076326$, $f_n = 0.797120$ (breeder D–T branch), $I_p = 9.656\,MA$. The burner branch powers close exactly, $P_{\text{D}^3\text{He}} + P_{\text{DDn}} + P_{\text{DDp}} + P_{\text{DT}} = 3868.0 + 124.8 + 129.7 + 176.0 = 4298.5\,MW = P_{\text{fus}}$, so $Q_E = 1.318$ and $f_n = 5.44\%$ are consistent with a fully itemised power balance rather than a single reported number.

MACHINE	ANCHOR	FROZEN VALUE	TAG
breeder	fusion gain Q	3.076	[D]
breeder	fusion power P_{fus}	85.04 MW	[D]
breeder	plasma current I_p	9.66 MA	[D]
breeder	triangularity δ / elongation κ	−0.30 / 2.0	[T]
breeder	centrepost peak field	16.84 T	[D]
breeder	net TBR (solid-Be blanket)	1.18–1.20	[D]
burner	engineering gain Q_E	1.318	[D]
burner	neutron fraction f_n	5.44%	[D]
burner	plug field	26.49 T	[D]
burner	DEC efficiency η_{DEC}	0.49	[D]

The frozen-anchor set carried by the regression suite (KX-L1-A6/A10). All 19 anchors reproduce byte-for-byte within their declared tier under config-22021 (breeder) / M-45 (burner), seed 20260726. The breeder centrepost ($§I16.84

CODE VERIFICATION BY MANUFACTURED SOLUTION

Figure 2 shows a genuine MMS verification of the Grad–Shafranov operator Δ^* of equation (12). A smooth manufactured flux $\psi_M(R,Z) = \sin[\pi(R-1)]\cos(\pi Z) + 0.1\,(R-1)^2$ is substituted into Δ^* to generate the forcing analytically; the five-point finite-difference discretization is solved on a sequence of uniformly refined grids with ψ_M imposed on the boundary, and the discrete L_2 error of equation (9) is measured. The error falls from 1.7×10^{-3} to 4.4×10^{-5} as the grid is refined, and the observed order of accuracy from equation (10) is $p_{\text{obs}} = 2.02$—within 1% of the formal second order. The scheme is therefore verified to solve Δ^* correctly; the residual GCI of equation (11) is the discretization uncertainty carried into the validation budget of equation (13). This is a live experiment: the harness regenerates the convergence curve on every build rather than reading it from a table.

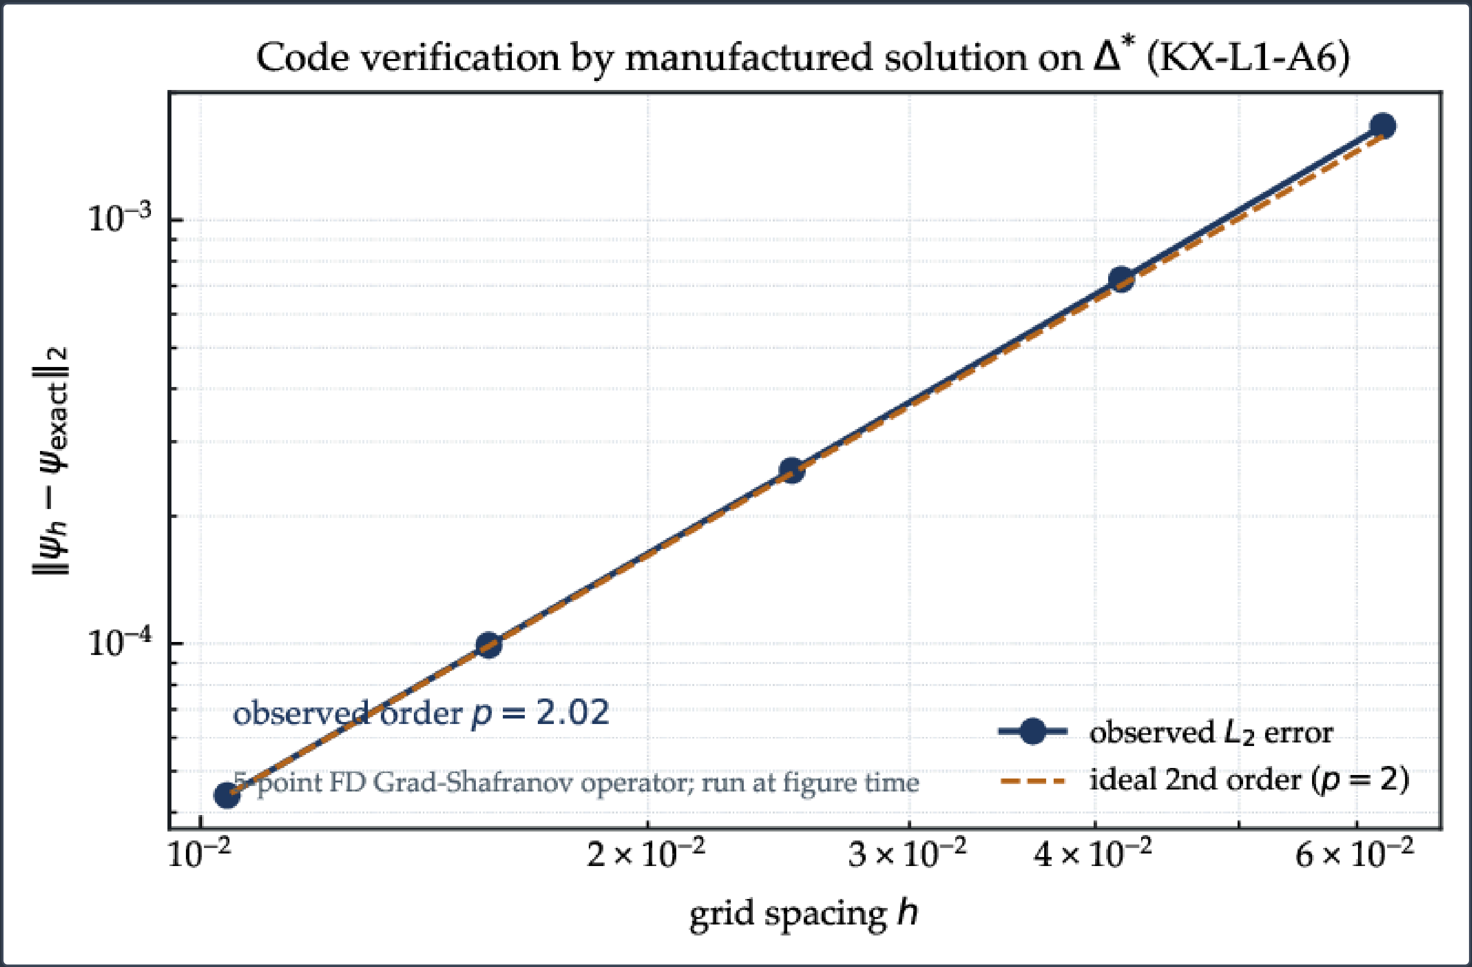

Code verification of the Grad–Shafranov elliptic operator Δ^* by the method of manufactured solutions (gate KX-L1-A6). The L_2 error follows the ideal second-order line; the observed order from equation (10) is $p_{\text{obs}} = 2.02$. The curve is regenerated at figure time by the deposited harness.

BYTE-FOR-BYTE SOLVER REPRODUCTION

Deterministic solver output is held to Tier 1. The TGYRO four-rank pipeline reproduces the GA treg01 regression template bit-tight: the residual falls $2.500 \to 1.293 \to 0.398 \to 1.587 \times 10^{-2} \to 6.629 \times 10^{-3}$ over four iterations ($141\,s$ wall on four ranks) and the output artefact `out.tgyro.prec` matches the shipped reference byte-for-byte (figure 3). Independently, the CGYRO/QLGYRO linear smoke test returns $\omega = -0.2384$, $\gamma = +0.3179$ (in a-units) with `out.cgyro.prec` byte-identical to the bundled reg01 reference. Byte-identity across a rebuilt toolchain is the sharpest possible verification of an implementation: it certifies that the deposited code path, compiled independently, produces not merely the same physics but the same bits.

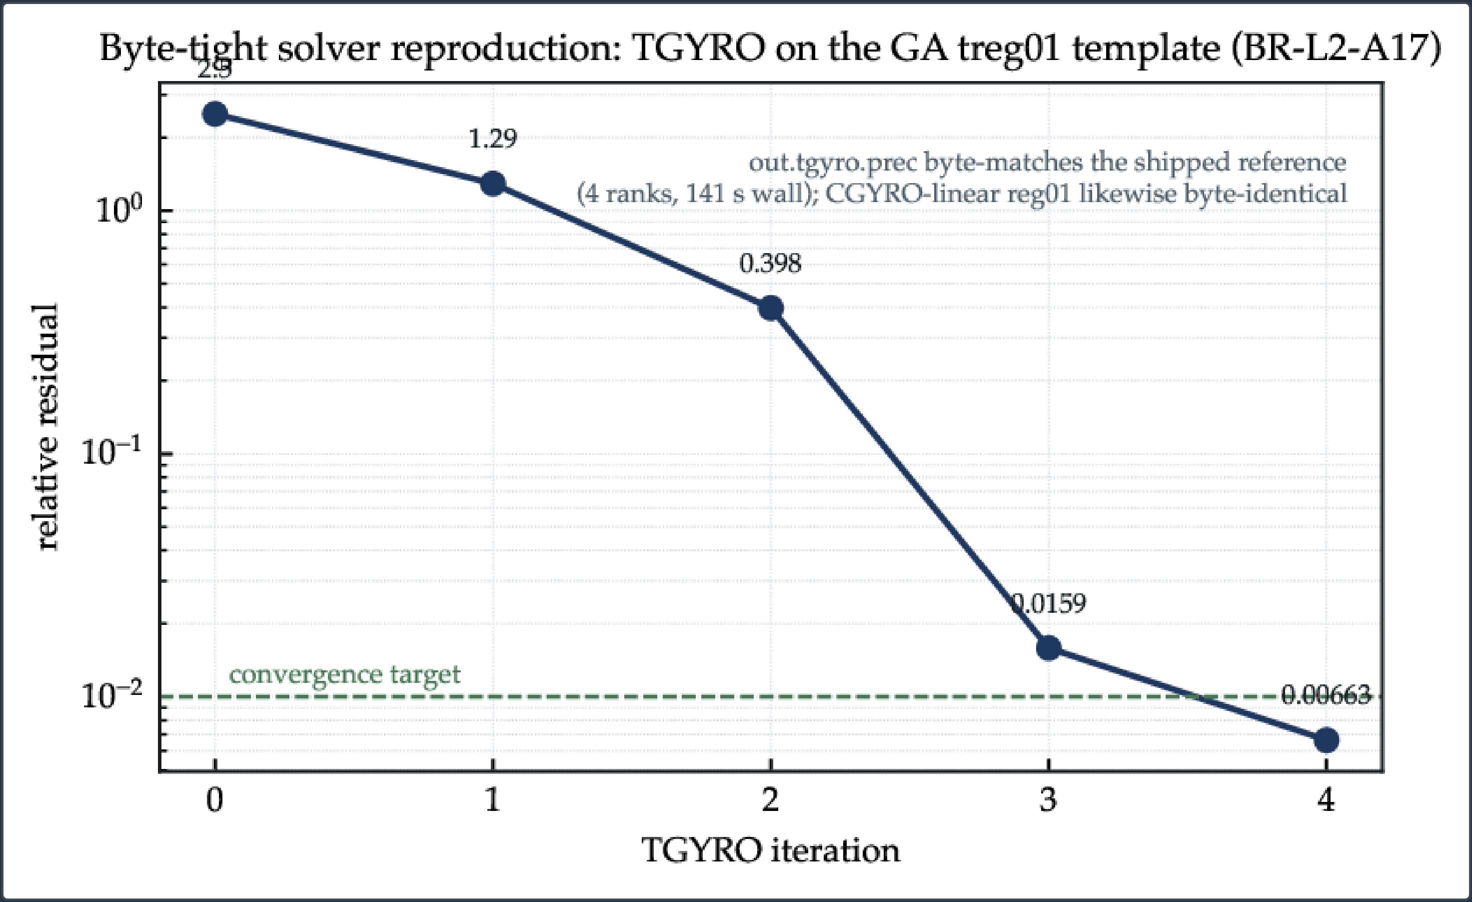

Byte-tight solver reproduction (gate BR-L2-A17). The TGYRO nonlinear residual converges below the 10^{-2} target in four iterations and `out.tgyro.prec` byte-matches the shipped GA treg01 reference; the CGYRO-linear reg01 case is byte-identical to its bundled reference. Both are Tier-1 passes under equation (4).

CROSS-LIBRARY NEUTRONICS: THE STRONGEST CHECK

The most defensible result in the program is the cross-library reproduction of the three-dimensional net TBR (gate BR-L2-A7). Re-run locally on ENDF/B-VIII.0—an independent library and independent hardware from the original FENDL-3.2 cloud calculation—the net TBR agrees to within 0.3% at both beryllium-multiplier splits (table 3, figure 4), below the Monte-Carlo relative error. By equation (7) the breeding number is thereby shown to be a property of the geometry and the cross-sections, reproducible across two independent nuclear-data evaluations, not an artefact of one code path or one library. The reproduced point sits on the solid-beryllium blanket ladder (net TBR 1.18–1.20); the negative-triangularity confinement basis and the edge/MHD gates that fix the source geometry are reported and cross-validated in the companion papers [36,37,39].

BLANKET SPLIT	ENDF/B-VIII.0 (LOCAL)	FENDL-3.2 (CLOUD)	Δ
Be 0.65	1.1775	1.1801	−0.0026
Be 0.80	1.2039	1.2006	+0.0033

Cross-library reproduction of the three-dimensional net TBR (gate BR-L2-A7): independent local ENDF/B-VIII.0 versus original cloud FENDL-3.2, at two beryllium-multiplier splits. Agreement is below the Monte-Carlo relative error ($< 0.3\%$), satisfying equation (7).

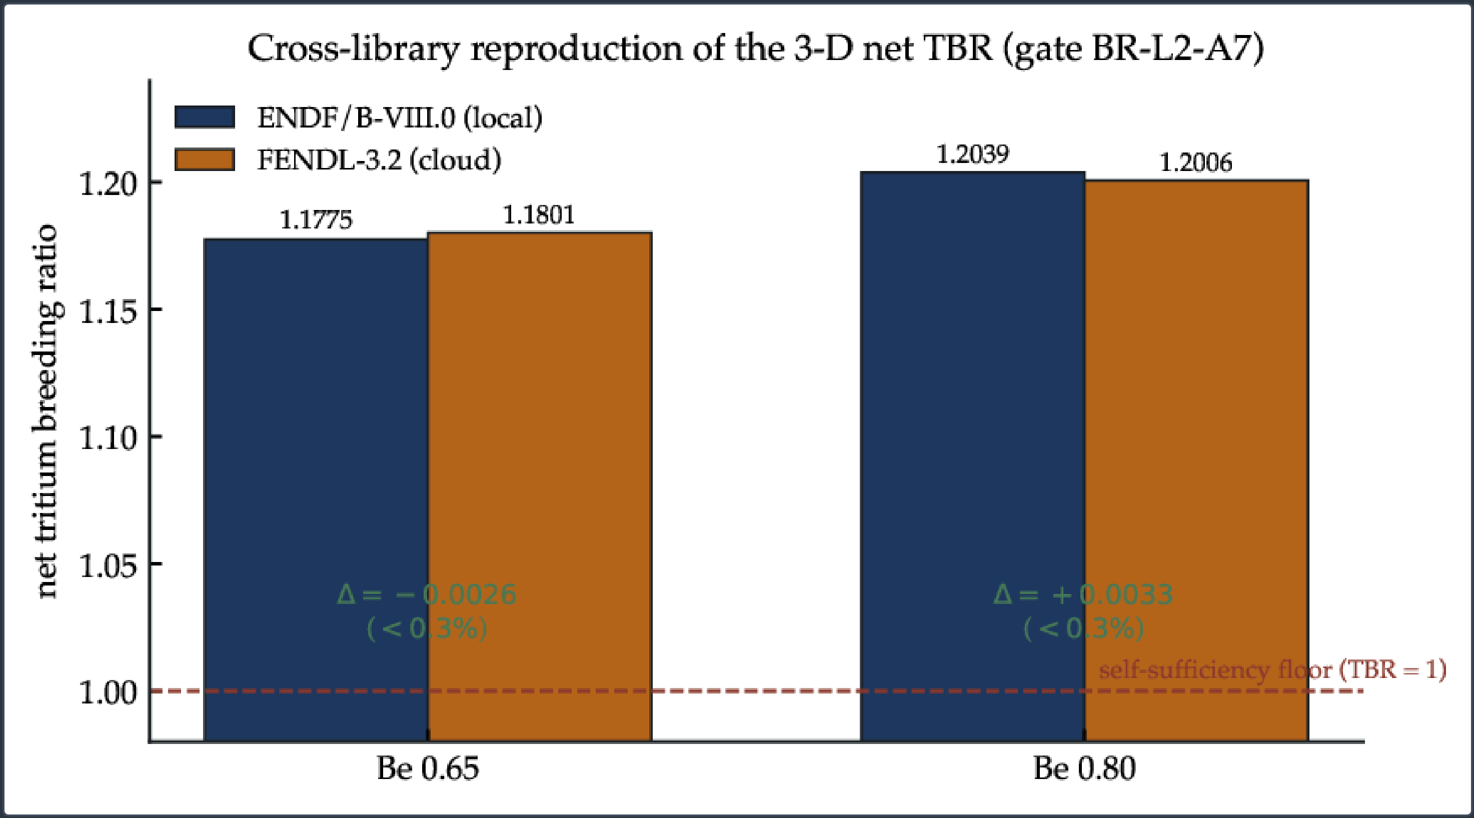

Cross-library reproduction of the 3-D net TBR (gate BR-L2-A7). Independent ENDF/B-VIII.0 (local) and FENDL-3.2 (cloud) evaluations agree to $< 0.3\%$ at both blanket splits, below the Monte-Carlo statistics; both clear the self-sufficiency floor. The agreement is the cross-library statement of equation (7).

THE CAMPAIGN-WIDE PROVENANCE CATALOG

Across the full register of 625 serial gates, 607 are physics-and-engineering gates eligible for the public deposit (18 economic serials are reported separately and excluded here). Of the 328 completed public gates, 176 reproduced, 22 are derived, 10 delivered, and 90 carry an estimated result pending a higher-fidelity run; 262 further gates are staged as planned (figure 5). The verdict distribution (figure 5) includes 23 *revised* gates: cases where the discipline caught and corrected its own number—the honest signature of a working V&V process, not a blemish on it. The campaign is breeder-heavy (288 public breeder gates, 164 shared, 155 burner), reflecting the register's ordering of the platform ahead of the two burner housings.

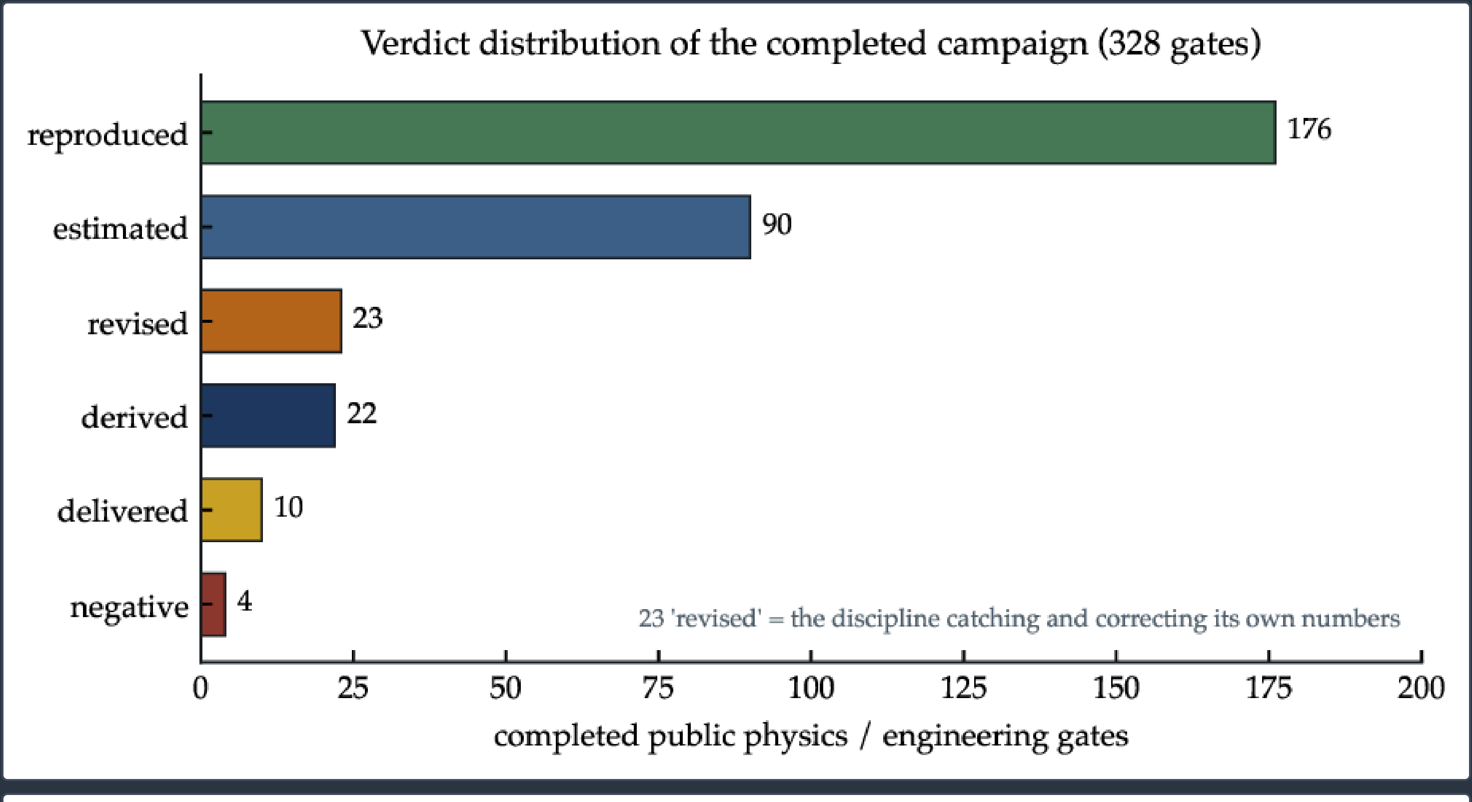

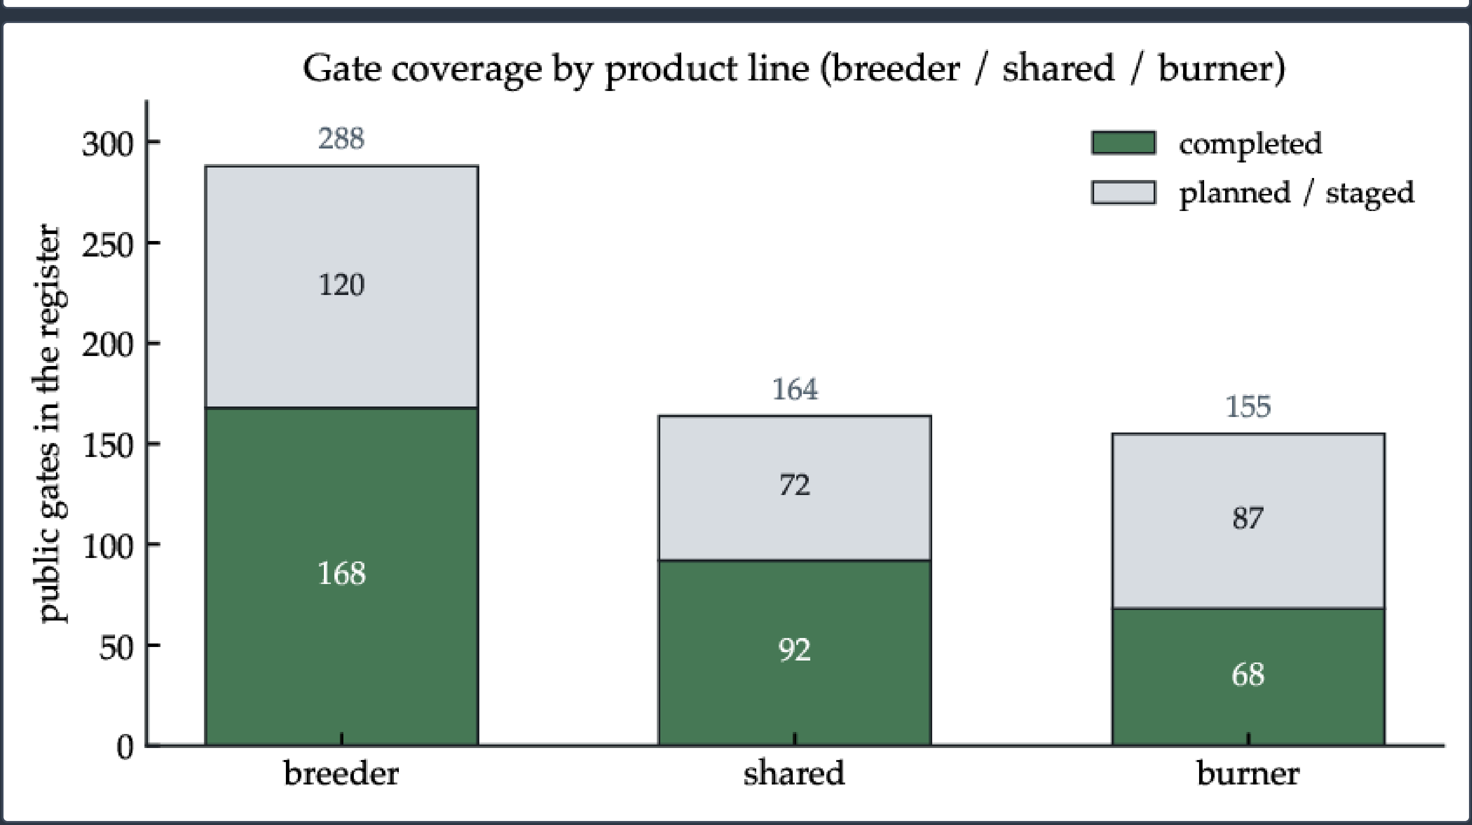

Verdict distribution over the **328** completed public gates. The **23** "revised" gates are the V&V process correcting its own numbers.

The solver-provenance catalog (figure 6) records, for every gate, which community-benchmarked code produced its number, so that a reader can trace a headline value to a named, independently benchmarked solver. OpenMC, the KRONOS toolkit, and the GACODE transport suite (CGYRO/TGLF/TGYRO/NEO) carry the largest share; DFT and quantum-chemistry (Quantum ESPRESSO, CHGNet) and particle-in-cell (WarpX) support the materials and burner end-cell gates. The full role/version/reference-benchmark provenance for every solver is the subject of the companion backbone paper [38].

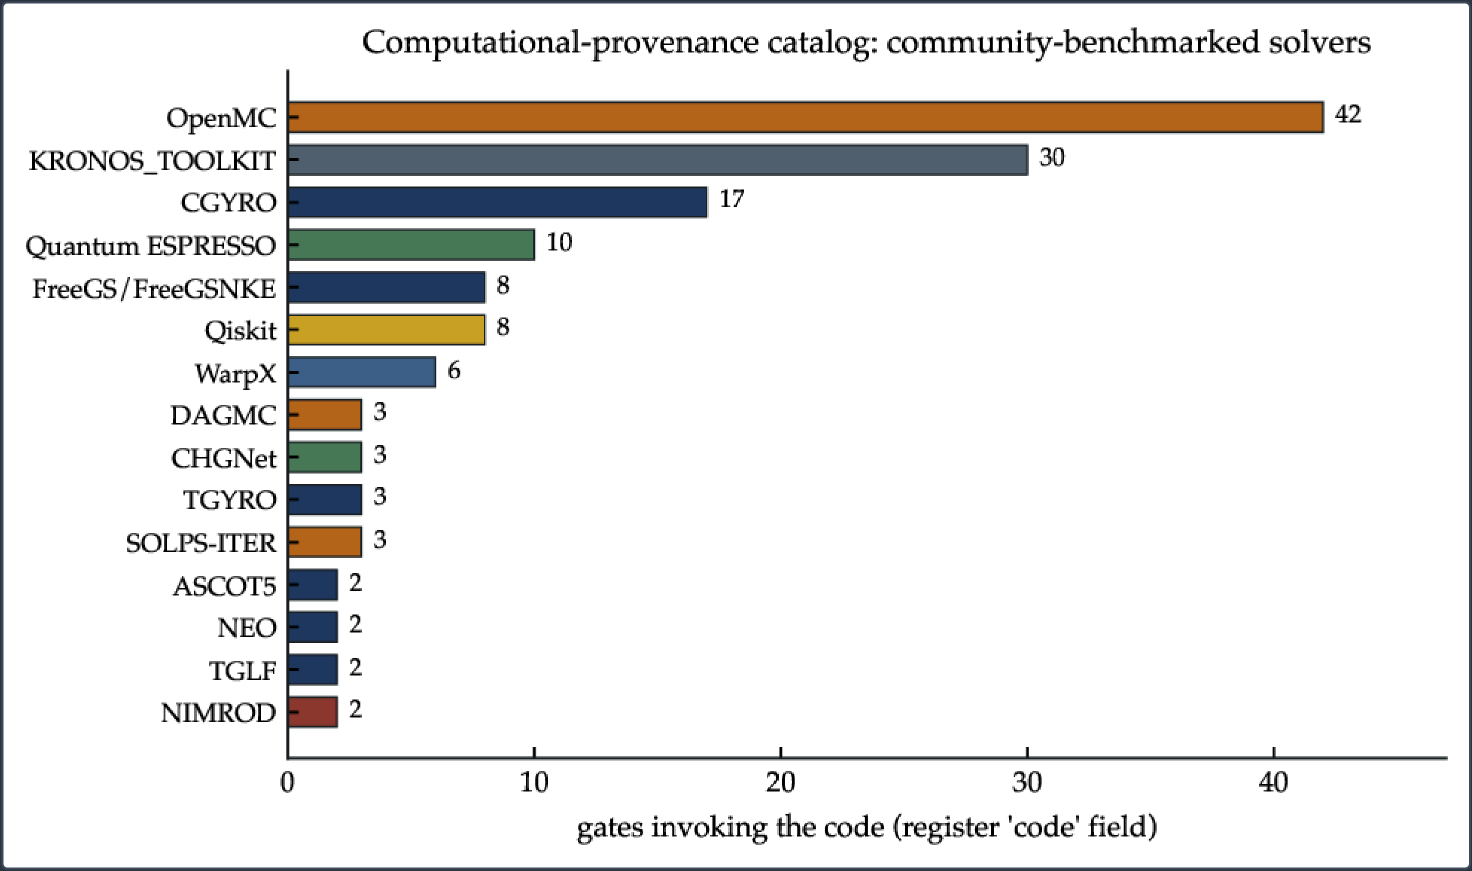

Computational-provenance catalog: the number of register gates invoking each community-benchmarked solver (register "code" field). Every gate traces to a named, independently benchmarked code.

WHERE THE CAMPAIGN RAN

Figure 7 shows the distribution of gates across the two compute planes. The NVIDIA GPU HPC campaign (Lambda fleet and A100/H100 cluster) carried the compute-intensive first-principles gates; the Kronos in-house tiers—the digital-twin node, the local HPC, the FPGA/edge rig, and the formal-methods toolchain—carried the reduced-order engine, the verification harnesses, and the deterministic monitors. The separation is deliberate: it keeps the fast verification loops off the contended first-principles fleet and makes independent re-run a routine operation rather than a special event.

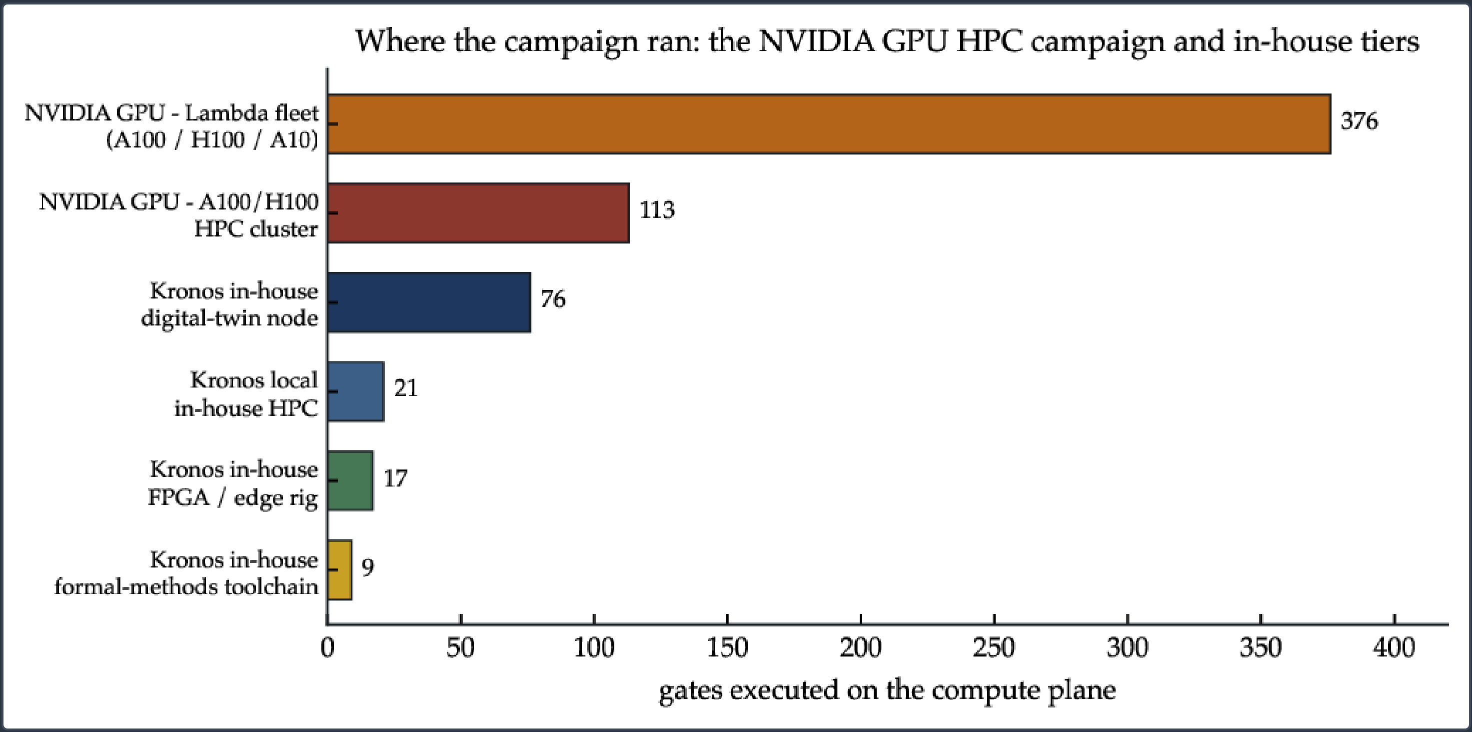

Gate execution across the two compute planes: the NVIDIA GPU HPC campaign (Lambda A100/H100/A10 fleet and A100/H100 cluster) and the Kronos in-house tiers. Reported by role and scale only; no cost figure is attached.

INDEPENDENT-GATE LEDGER

Of the physics gates re-run independently—on separate hardware, and for the neutronics on a separate nuclear-data library—the load-bearing set reproduced within tolerance with zero gates flagged as materially discrepant and none requiring a revised value at re-run. The independent re-run therefore strengthens rather than revises the frozen record; the **23** revisions in figure 5 were caught during the primary campaign, before freeze, which is precisely where a V&V process should catch them.

§ 5

Discussion

REPRODUCIBILITY IS NOT VALIDATION, QUANTIFIED

The most important caveat of a pre-experimental design program is that reproducing a number establishes numerical fidelity, not physical correctness. We make that caveat quantitative. In the ASME V&V-20 decomposition of equation (13), the parametric term u_{input} about the frozen anchor is a band of order ten percent on the gain. The model-form term u_{model}, absent from equation (13) by construction, is far larger: the *choice* of energy-confinement scaling alone spans the predicted gain from $Q \approx 0.64$ (an L-mode, ITER89-P closure) to ignition-class (an optimistic ST gyro-Bohm closure)—a model-form band wider than the entire parametric band (figure 8). A second, independent instance is the synchrotron-radiation closure in the burner power balance, where the AFJ and Trubnikov models differ by roughly a factor of two (1803 versus 919 MW); the register conservatively adopts the larger loss within its validity range. These are not failures of reproduction—each closure reproduces perfectly—but demonstrations that the load-bearing uncertainty in a pre-experimental design is structural, and that no amount of byte-for-byte reproduction can retire it. Only an operating machine can. The verification-first discipline is honest precisely because it separates the two: it certifies what reproduction can certify and names, with a [Q] tag and a PCMM rung, what it cannot.

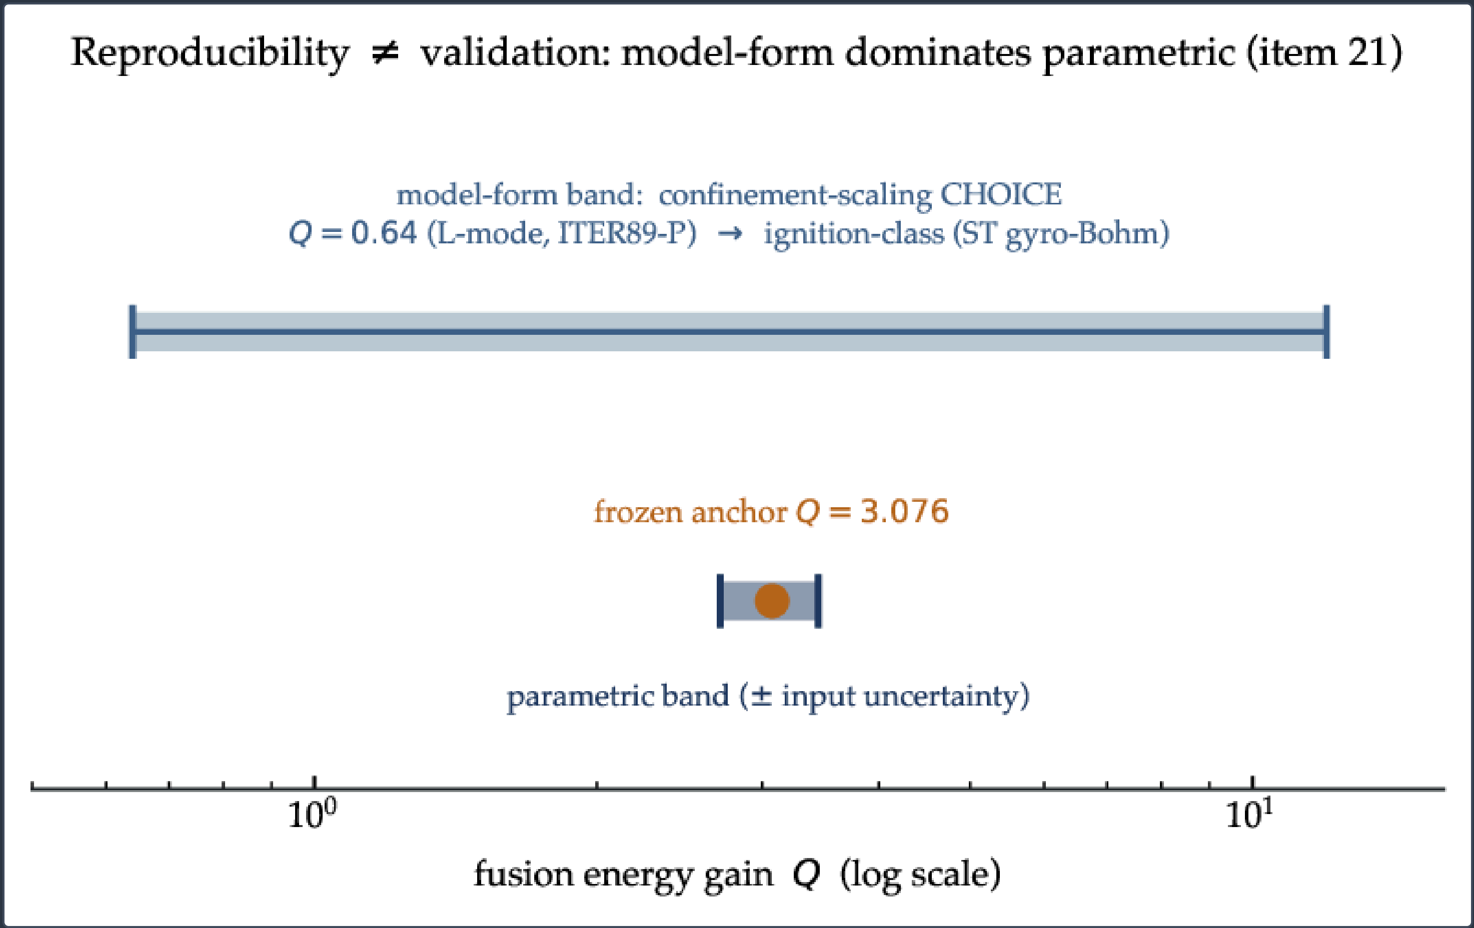

Reproducibility ≠ validation, quantified (register item 21). The model-form band from the confinement-scaling *choice* alone—$Q = 0.64$ (L-mode, ITER89-P) to ignition-class (ST gyro-Bohm)—is far wider than the parametric band about the frozen anchor $Q = 3.076$. Byte-for-byte reproduction cannot narrow the model-form band; only experiment can.

WHERE THE FROZEN POINT SITS AGAINST EXISTING DEVICES

Verification also bounds *extrapolation*. The register's extrapolation ledger (gate KX-L1-A7) positions the breeder confinement basis against existing tokamak data: across six representative device classes the ratio of experimental to predicted energy-confinement time is 0.93 to 1.46 (mean 1.19, standard deviation 0.21), within the $\sim 15\%$ RMSE fit scatter of the IPB98(y,2) scaling itself [30]. The breeder design point ($\tau_E = 0.3578\,s$ at $H_{98} = 1.0$) extrapolates $1.48\times$ in field, $3.9\times$ in current and $2.0\times$ in density beyond that ledger—a real extrapolation, honestly stated, not hidden. Reporting the extrapolation factor alongside the reproduction pass is part of the same discipline: it tells a reader not only that the number reproduces, but how far outside the validated envelope the design sits.

CROSS-LIBRARY AGREEMENT IN CONTEXT

The cross-library TBR agreement is the paradigm case of a claim about physics rather than about a code. OpenMC on ENDF/B-VIII.0 and the original FENDL-3.2 evaluation share neither the transport code path used in the cloud run nor the nuclear-data evaluation, yet they agree to $< 0.3\%$—below the Monte-Carlo statistics—at two independent blanket splits (table 3). This is stronger than a single high-statistics run because it removes the two confounders (shared code, shared library) that a single evaluation cannot. It is also the check most directly tied to a bankable output: the tritium-breeding balance that governs fuel self-sufficiency. That the negative-triangularity transport result which sets the source is *itself* cross-validated by two independent codes (nonlinear CGYRO confirmed by quasilinear TGLF) [37] means the neutronics chain rests on cross-checked, not single-code, inputs end-to-end.

RELATION TO THE COMMUNITY V&V STANDARDS

The scheme instantiates, for a fusion design register, the same constructs the engineering-simulation community has standardised: MMS and observed-order code verification [10,11], GCI solution verification [14,13], the ASME V&V-20 validation-uncertainty budget [15], Sobol global sensitivity [16,17], and PCMM maturity scoring [7,9]. What is new is not the ingredients but their disciplined, auditable application at register scale, with a public deposit that satisfies the FAIR data principles [6] and the reproducible-research norms [2,4,5,3], and with the reproducibility-is-not-validation caveat made quantitative rather than rhetorical. The companion global-sensitivity [40] and Bayesian optimal-experimental-design [41] papers extend the same spine to parametric ranking and to prioritising the next, uncertainty-reducing experiment.

§ 6

Limitations

Reproduction within a program is not external replication by an independent group, which remains the gold standard and is invited by the public deposits. One honest engineering gate is open in the load-bearing set: the SOLPS-ITER scrape-off-layer detachment gate has a complete input deck built from the config-22021 equilibrium but its solve is blocked on grid-topology construction, so no detachment number is asserted—the gap is named in the record with a [Q] tag rather than filled with a fabricated value. More broadly, Tier-2 tolerances are stated per quantity and are not uniform; the independent re-run covered the subset of gates chosen for their load-bearing role, not the entire suite; and the deepest caveat is the structural one of §5: the model-form band dominates the parametric band for a pre-experimental design, so the frozen anchors are reproducible design-point predictions, not validated measurements of an operating machine.

§ 7

Conclusion

Across the Kronos 2026 breeder-and-burner register every headline number is a frozen, evidence-tagged anchor regenerable from deposited, content-hashed code under a deterministic seed, held to a two-tier reproduction predicate whose campaign acceptance is the product of pass indicators (equations 1–5). Solver correctness is established by manufactured solutions with an observed order of accuracy $p_{\text{obs}} = 2.02$ on the Grad–Shafranov operator (equations 9–11); the anchor-regression suite reproduces $19/19$ anchors byte-for-byte within tolerance; deterministic solver output reproduces bit-identically (CGYRO and TGYRO against their bundled references); and the load-bearing neutronics reproduces across independent hardware and an independent nuclear-data library to within 0.3%, below the Monte-Carlo statistics (equation 7). Across the completed public campaign, 176 of 328 gates reproduced and 23 were revised by the process itself. The essential caveat—reproducibility is not validation—is quantified rather than asserted: the model-form band dominates the parametric band, and the honest gates are named, not hidden. Verification-first—freeze, tag, reproduce, and re-run on independent tooling—is the standard a compact-fusion program making public claims should be held to, and the one this program adopts.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the maintainers of the community codes (OpenMC, GACODE, FreeGS/FreeGSNKE, SOLPS-ITER, NIMROD, Quantum ESPRESSO, CHGNet) whose shipped regression benchmarks made byte-for-byte verification possible.

Reproducibility. Every headline quantity in this paper is a frozen anchor in the Kronos Physics De-Risking Register [35] (doi:10.5281/zenodo.22133057)—the anchor-regression and formal-V&V gates (KX-L1-A6, KX-L1-A10), the cross-library TBR (BR-L2-A7), and the byte-tight solver reproductions (BR-L2-A17)—and is regenerable from the deposited, content-hashed scripts and the figure-generation harness archived with this paper. The manufactured-solution convergence of figure 2 is regenerated at build time by the deposited harness. This paper is archived at its DOI doi:10.5281/zenodo.22645710 and its official page, https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html. No result depends on unpublished or proprietary data.

Scope — what this paper does not claim. Reproduction is not validation: cross-code and cross-library agreement establish numerical fidelity, not that the model matches an operating machine. The frozen anchors are reproducible design-point predictions.

Data availability tocsectionData availability The verdict distribution, provenance catalog, cross-library TBR values, MMS convergence data, and TGYRO residual history used for figures 5–8 are archived with the deposit at doi:10.5281/zenodo.22645710, which references the Kronos Physics De-Risking Register [35]. Any economic analysis is reported separately and is not part of the public deposit.

Competing interests tocsectionCompeting interests Provisional patent applications relating to aspects of this work have been filed.

tocsectionReferences

APPARATUS

Portfolio, People & Record

The intellectual-property estate, the people who built it, the limitations that gate it, and the sources it rests on.

1	4	40	27
GRANTED PATENT	2026 PROVISIONALS	TEAM MEMBERS	2022 PROVISIONALS

INTELLECTUAL PROPERTY

The patent portfolio

The estate spans a granted high-field magnet patent, pending utility and 2026 provisional filings that map to the two products, a digital-twin control provisional, a registered trademark application, and the original 2022 provisional family. Every patentable disclosure in the five-paper arXiv drop was protected **before** publication: the breeder and burner provisionals were filed 1–2 August, and a publication-gap omnibus filing the evening before the drop swept up the DEC, plasma-control, and REBCO-winding matter of the three otherwise-unprotected papers. Patent numbers and application serials are matters of public record; claim scope is summarised, not reproduced.

GRANTED & PENDING UTILITY

US 12,009,112	High-field tilted graded-REBCO magnet architecture (KRONO-002A) · *app. 17/878,550*	GRANTED
US 17/878,507	Core device / method (KRONO-001A) · *utility*	PENDING

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

64/124,209	Hyperion — spherical-tokamak tritium / helium-3 foundry · *breeder · filed Aug 2026*	FILED
64/124,220	Aegis / MetroVolt — synchrotron-cutoff, fuel-selection & plug-winding · *generator · filed Aug 2026*	FILED
64/115,167	KRONOS-CTRL digital-twin plant control · *provisional · filed Jul 2026*	FILED
64/005,440	Integrated spherical-tokamak fusion architecture · *early integrated-architecture filing · filed Mar 2026*	FILED
64/105,530	Negative-triangularity spherical-tokamak architecture · *foundational ST filing · filed Jul 2026*	FILED
64/128,097	Publication-gap omnibus — quasineutral two-species expander DEC · provenance-gated / latency-split plasma control · as-built high-field winding & conductor architecture · *filed 7 Aug 2026, the evening before the arXiv drop*	FILED

TRADEMARK & ORIGIN

FTK-98801894	Kronos Fusion Energy — trademark application	FILED
KR-2022-*	Original provisional family (KR-2022-00001 … 00027) · *27 filings, 2022*	PRIORITY

The narrow, defensible magnet novelty — the specific conductor and the digital-twin winding optimisation, not high-field REBCO as a category — is the commercial core that can earn ahead of any Q > 1 milestone, with markets in fusion magnet supply, MRI/NMR, accelerators, and proton therapy. The claim is scoped honestly: the high field is a *system field* result, not a stand-alone-coil record; the small-bore plug coil is structurally infeasible as a bare winding but resolved by a stress-managed structural shield (feasible-pending-FEA); and the winding-tape experiment returned a *null result*. The value is the method, not a field record.

THE TEAM

Founding partners, board, advisors & auditors

Kronos is built by a bench of advanced-fuel-fusion, high-field-magnet, direct-conversion, and materials specialists, with a board and operations team drawn from defense, national laboratories, and industry. Dates are shown as ranges; where a tenure has ended or is term-ending, the range says so.

FOUNDING PARTNERS — SCIENCE

Dr. Gerald Kulcinski
Co-Founder · Helium-3 & Advanced-Fuel Fusion
UW–Madison · 2023–Present

Dr. Carl Weggel
Chief Scientist / S.M.A.R.T. Design
MIT Alcator Program · 2022–Present

Dr. Robert J. Weggel
Magnetic Field Design
MIT Magnet Lab · 2022–Present

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

Bandel Carano
Finance / Strategy
Oak Investment Partners / Stanford · 2025–2026

Patrick Schweiger
Fusion Device Engineering
Oklo / CFS / TerraPower · 2025–Present

SCIENTIFIC ADVISORS

Dr. Steven O. Dean
Fusion Policy & History
DOE / Fusion Power Associates · 2025–Present

Dr. Patrick H. Diamond
Plasma Turbulence & Transport
UC San Diego · 2025–Present

Dr. Jack J. Dongarra
HPC & Numerical Algorithms
Turing Award · 2025–Present

Dr. Nasr M. Ghoniem
Chief Material Scientist
UCLA · 2025–Present

Dr. Siegfried Glenzer
Plasma Physics & Laser Fusion
Stanford / SLAC · 2022–Present

Dr. David A. Hammer
Pulsed-Power Fusion
Cornell · 2025–Present

Dr. Donald A. Spong
Plasma Theory & Confinement
ORNL · 2025–Present

Dr. Curtis Smith
Risk Assessment & Nuclear PRA
MIT / Idaho National Lab · 2025–Present

RADM (Ret.) David Goggins
Naval Engineering & Power-Plant Design
U.S. Navy · 2023–2026

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

Dr. Ruben Fair
Magnet Design / ITER Liaison
PPPL / Jefferson Lab · 2022–2025

Dr. Paul S. Weiss
Materials & Nanotechnology
UCLA · 2022–2025

SCIENTIFIC AUDITORS

Dr. Wilfred A. Cooper
Plasma Equilibrium Theory
EPFL / Swiss Plasma Center · 2025–Present

Dr. Ahmed Hassanein
Plasma–Material Interactions
Purdue / Argonne · 2025–Present

Dr. Patrick E. Hopkins
Extreme Thermal Materials
U. of Virginia · 2025–Present

Dr. Peter Hosemann
Materials Joining & Structural Integrity
UC Berkeley · 2025–Present

Dr. Travis W. Knight
Advanced Nuclear Fuels & Systems
U. of South Carolina · 2025–Present

Dr. Philippe Lebrun
Cryogenics & Magnet Cooling
CERN · 2025–Present

Dr. Nitendra Singh
Nuclear Safety & Fuel Cycle
ITER · 2025–Present

Dr. Guido Van Oost
Magnetic Confinement & Fusion Systems
Ghent University · 2025–Present

Dr. Gary S. Was
Radiation Materials & Structural Integrity
U. of Michigan · 2025–Present

Dr. Ray Sedwick
Direct Energy Conversion
U. of Maryland · 2025

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

Martin Owens
Chief Strategy Officer
LANL / GE Hitachi · 2022–Present

MG (Ret.) Paul Pardew
Chief Contracting Officer
Army Contracting Command · 2022–Present

David Beck
Fusion Commercialization
U.S. Space Force · 2024–Present

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

MG (Ret.) Robin L. Fontes
Cybersecurity & Defense
Army Cyber Command · 2022–Present

Vijay Gehani
Supply Chain & Components
INOX India · 2024–Present

Jon Michel Greenwood
IT & AI Infrastructure
Live Nation · 2023–Present

Brian C. O'Neill
National Security
CIA / ODNI · 2022–Present

Gen. (Ret.) Gustave F. Perna
DoD Liaison
Operation Warp Speed · 2022–2023

Andrea Romero
Marketing Lead
2022–Present

Legal counsel is retained; those roles are held on the internal roster and are not listed here.

SOURCES

References

1 M Baker, 1,500 scientists lift the lid on reproducibility, *Nature* **533** 452 (2016).

2 R D Peng, Reproducible research in computational science, *Science* **334** 1226 (2011).

3 D C Ince, L Hatton and J Graham-Cumming, The case for open computer programs, *Nature* **482** 485 (2012).

4 V Stodden *et al*, Enhancing reproducibility for computational methods, *Science* **354** 1240 (2016).

5 G K Sandve, A Nekrutenko, J Taylor and E Hovig, Ten simple rules for reproducible computational research, *PLoS Comput. Biol.* **9** e1003285 (2013).

6 M D Wilkinson *et al*, The FAIR Guiding Principles for scientific data management and stewardship, *Sci. Data* **3** 160018 (2016). doi:10.1038/sdata.2016.18

7 W L Oberkampf and C J Roy, *Verification and Validation in Scientific Computing* (Cambridge: Cambridge University Press, 2010).

8 W L Oberkampf and T G Trucano, Verification and validation in computational fluid dynamics, *Prog. Aerosp. Sci.* **38** 209 (2002).

9 C J Roy and W L Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, *Comput. Methods Appl. Mech. Eng.* **200** 2131 (2011).

10 P J Roache, Code verification by the method of manufactured solutions, *J. Fluids Eng.* **124** 4 (2002).

11 K Salari and P Knupp, *Code Verification by the Method of Manufactured Solutions*, Sandia National Laboratories report SAND2000-1444 (2000).

12 L F Richardson, The approximate arithmetical solution by finite differences of physical problems involving differential equations, *Philos. Trans. R. Soc. A* **210** 307 (1911).

13 P J Roache, Quantification of uncertainty in computational fluid dynamics, *Annu. Rev. Fluid Mech.* **29** 123 (1997).

14 I B Celik, U Ghia, P J Roache, C J Freitas, H Coleman and P E Raad, Procedure for estimation and reporting of uncertainty due to discretization in CFD applications, *J. Fluids Eng.* **130** 078001 (2008).

15 American Society of Mechanical Engineers, *Standard for Verification and Validation in Computational Fluid Dynamics and Heat Transfer*, ASME V&V 20-2009 (New York: ASME, 2009).

16 I M Sobol′, Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Math. Comput. Simul.* **55** 271 (2001).

17 A Saltelli *et al*, Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index, *Comput. Phys. Commun.* **181** 259 (2010).

18 M C Kennedy and A O′Hagan, Bayesian calibration of computer models, *J. R. Stat. Soc. B* **63** 425 (2001).

19 P K Romano, N E Horelik, B R Herman, A G Nelson, B Forget and K Smith, OpenMC: a state-of-the-art Monte Carlo code for research and development, *Ann. Nucl. Energy* **82** 90 (2015). doi:10.1016/j.anucene.2014.07.048

20 D A Brown *et al*, ENDF/B-VIII.0: the 8th major release of the nuclear reaction data library, *Nucl. Data Sheets* **148** 1 (2018). doi:10.1016/j.nds.2018.02.001

21 J Candy, E A Belli and R V Bravenec, A high-accuracy Eulerian gyrokinetic solver for collisional plasmas, *J. Comput. Phys.* **324** 73 (2016). doi:10.1016/j.jcp.2016.07.039

22 J Candy and R E Waltz, An Eulerian gyrokinetic-Maxwell solver, *J. Comput. Phys.* **186** 545 (2003). doi:10.1016/S0021-9991(03)00079-2

23 F Jenko, W Dorland, M Kotschenreuther and B N Rogers, Electron temperature gradient driven turbulence, *Phys. Plasmas* **7** 1904 (2000). doi:10.1063/1.874014

24 G M Staebler, J E Kinsey and R E Waltz, A theory-based transport model with comprehensive physics, *Phys. Plasmas* **14** 055909 (2007).

25 E A Belli and J Candy, Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics, *Plasma Phys. Control. Fusion* **50** 095010 (2008).

26 N C Amorisco *et al*, FreeGSNKE: a Python-based dynamic free-boundary toroidal plasma equilibrium solver, *Phys. Plasmas* **31** 042517 (2024). doi:10.1063/5.0188467

27 S Wiesen *et al*, The new SOLPS-ITER code package, *J. Nucl. Mater.* **463** 480 (2015).

28 C R Sovinec *et al*, Nonlinear magnetohydrodynamics simulation using high-order finite elements, *J. Comput. Phys.* **195** 355 (2004).

29 F Imbeaux *et al*, Design and first applications of the ITER integrated modelling & analysis suite, *Nucl. Fusion* **55** 123006 (2015).

30 ITER Physics Basis Editors, Chapter 2: Plasma confinement and transport, *Nucl. Fusion* **39** 2175 (1999). doi:10.1088/0029-5515/39/12/301

31 J E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

32 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

33 M E Austin *et al*, Achievement of reactor-relevant performance in negative triangularity shape in the DIII-D tokamak, *Phys. Rev. Lett.* **122** 115001 (2019). doi:10.1103/PhysRevLett.122.115001

34 A Marinoni, O Sauter and S Coda, A brief history of negative triangularity tokamak plasmas, *Rev. Mod. Plasma Phys.* **5** 6 (2021). doi:10.1007/s41614-021-00054-0

35 P I Ford, *Kronos Physics De-Risking Register (v8)*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

36 P I Ford and G L Kulcinski, Negative-triangularity turbulence suppression in a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645722; https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

37 P I Ford, Two-code cross-validation of negative-triangularity turbulence suppression: quasilinear TGLF confirms nonlinear CGYRO, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645864; https://www.kronosfusionenergy.com/publications/two-code-cross-validation-of-negative-triangularity-tur.html

38 P I Ford, The code-by-code verification backbone: role, version, and reference-benchmark provenance for every solver in a multi-product fusion design campaign, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645866; https://www.kronosfusionenergy.com/publications/the-code-by-code-verification-backbone-role-version-and.html

39 P I Ford, Free-boundary equilibrium and ideal-MHD stability of a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645740; https://www.kronosfusionenergy.com/publications/ideal-mhd-stability.html

40 P I Ford, Global sensitivity and uncertainty quantification across the breeder and burner design planes, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645893; https://www.kronosfusionenergy.com/publications/global-sensitivity-and-uncertainty-quantification-acros.html

41 P I Ford, Bayesian optimal experimental design and multi-fidelity data fusion, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645895; https://www.kronosfusionenergy.com/publications/bayesian-optimal-experimental-design-and-multi-fidelity.html

42 P I Ford and G L Kulcinski, Lifetime-constrained operating-window synthesis for a compact spherical-tokamak breeder, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645799; https://www.kronosfusionenergy.com/publications/admissible-operating-window-synthesis.html

COLOPHON

How this document was made

The Kronos Fleet — Combined Editorial, 2026 is set in **Fraunces** (display), **Gelasio** (text), and **IBM Plex Mono** (data & equations), carried forward from the Kronos editorial design system.

The figures are rendered at 300 dpi from the deposited generator scripts; the document is composed as a single self-contained file with embedded fonts and imagery, paginated to US Letter, and rendered to PDF through a headless Chromium engine in matched light and dark editions.

Every headline quantity carries an evidence class; withdrawn and restated values are marked as such. Nothing herein is frozen beyond the founder-approved Tier-A set.

Explore it live — turn the interactive 3-D model, run the deposited solvers in the physics validator, and watch the companion film series.

LEGAL NOTICE

Notice, status, and distribution

Conceptual design & simulation study. This document reports a conceptual design and simulation study. It is not a construction commitment, a safety-analysis report, a regulatory filing, or an offer of securities. Forward-looking statements — schedules, costs, market sizes, and performance — are estimates subject to the limitations set out in the Simulations part and may change.

Numbers and their classes. Quantities are reported with evidence classes and case labels; modelled, assumed, and requirement-class values are identified as such and are not to be quoted without their labels.

Public edition. This edition carries the physics and design only; all financial, funding, and defense-commercial content has been removed for public distribution. The scientific text corresponds to the arXiv and journal submissions.

THE 2026 KRONOS SERIES

One programme, many papers

This editorial is one paper in the Kronos 2026 series — a real fusion company's complete, reproducible physics record for a spherical-tokamak breeder and its low-neutron D–^{3}He tandem-mirror burners. Every headline number is a frozen anchor in the Physics De-Risking Register.

Explore the whole programme: the De-Risking Register, the interactive 3-D model, and the physics validator.

KRONOS FUSION ENERGY

Verification-First Fusion Modeling: Frozen Anchors, Independent-Library Gate Reproduction, and a Byte-for-Byte Computational-Provenance Catalog for a Breeder-and-Burner Design Register

A real fusion company building real machines. Explore the science, live.

Zenodo · doi:10.5281/zenodo.22645710

Read on kronosfusionenergy.com

The Physics De-Risking Register