KRONOS FUSION ENERGY

PAPER 2.99 · THE KRONOS 2026 SERIES

# *The Code-by-Code Verification Backbone: Role, Version, and Reference-Benchmark Provenance for Every Solver in a Multi-Product Fusion Design Campaign*

A design frozen for construction is only as credible as the solvers behind it, and a solver is only as credible as its provenance. We present the verification backbone of the…

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

BY P I FORD

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

# The Code-by-Code Verification Backbone: Role, Version, and Reference-Benchmark Provenance for Every Solver in a Multi-Product Fusion Design Campaign

A design frozen for construction is only as credible as the solvers behind it, and a solver is only as credible as its provenance. We present the verification backbone of the Kronos physics de-risking campaign, which de-risks a compact spherical-tokamak breeder (design point $Q = 3.076$, $P_{\text{fus}} = 85.04\,\text{MW}$, $I_p = 9.66\,\text{MA}$, negative triangularity $\delta = -0.30$, $B_0 = 8\,\text{T}$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,\text{m}$, aspect ratio $A = 2.5$, centrepost peak field $16.84\,\text{T}$) and a D–$^3$He tandem-mirror burner ($Q_E = 1.318$, neutron fraction $f_n = 5.44\%$, end-plug field $26.49\,\text{T}$, direct-conversion efficiency $\eta_{\text{DEC}} = 0.49$). The backbone is a catalog of **47 named solvers** across nine physics domains, each assigned a formal *role* (reference, surrogate, or gate), a pinned *version* (git commit, build environment, and data library), and a *reference benchmark* it must reproduce within a stated tolerance. We give the mathematics that makes the backbone auditable: a decomposition of total predictive error into surrogate, discretization, and model-form terms; Richardson extrapolation and the grid-convergence index for solution verification; the method-of-manufactured-solutions order test for code verification; the Monte-Carlo standard-error law for statistical convergence; the Sobol variance decomposition for sensitivity; a Minkowski-$L^1$ area validation metric; and a hash-chained provenance ledger with a tamper-evidence guarantee. We then show the backbone at work: the formal verification-and-validation suite passes **19 of 19** frozen anchors across breeder and burner; the CGYRO linear gyrokinetic reference returns a growth rate $\gamma = +0.3179$ (in $a/c_s$ units) byte-identically to its shipped benchmark; a CAD-faithful DAGMC neutronics chain reproduces the net tritium-breeding ratio to $1.014 \pm 0.004$ against the $1.026$ of the frozen recipe; and two independent density-functional-theory codes agree on a divertor-alloy formation energy to $0.002\,\text{eV/atom}$ ($+0.0147$ versus $+0.016$). Every headline number in the campaign thus carries a traceable computational provenance and an independent check behind it. All results were produced on the NVIDIA GPU HPC campaign using the cataloged community reference codes.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645866 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

**Keywords:** verification and validation · reproducibility · reference codes · computational provenance · grid-convergence index · method of manufactured solutions · Sobol sensitivity · gyrokinetics · neutronics

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) |
| $\tau_E$ | energy confinement time |
| $I_p$ | plasma current |
| $R_0$, $a$ | major radius, minor radius |
| $\kappa$, $\delta$ | elongation, triangularity |
| $\beta_N$ | 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) |

verification and validation; reproducibility; reference codes; computational provenance; grid-convergence index; method of manufactured solutions; Sobol sensitivity; gyrokinetics; neutronics

§ 1

# Introduction

## WHY A FUSION DESIGN REGISTER NEEDS A VERIFICATION BACKBONE

A compact fusion device is a coupled multiphysics system, and no single monolithic solver captures its behaviour faithfully. The Kronos campaign therefore uses the right community-standard code for each domain — gyrokinetics for core turbulence, Monte-Carlo transport for neutronics, free-boundary Grad–Shafranov for equilibrium, plane-wave density-functional theory for materials — and treats its own reduced-order and learned engines as *surrogates* that must reproduce those references before their numbers are trusted. What binds these solvers into a credible whole is not any one of them but the *backbone*: the discipline that records, for every solver, what role it plays, exactly which version produced a number, and which reference benchmark it was shown to reproduce. This paper documents that backbone and gives the mathematics that makes it auditable.

The stakes are concrete. Kronos freezes a design point for construction: a compact spherical-tokamak (ST) breeder holding $I_p = 9.66\,\mathrm{MA}$ at a major radius of only $R_0 = 1.20\,\mathrm{m}$ behind a centrepost magnet peaking at $16.84\,\mathrm{T}$, operated at negative triangularity ($\delta = -0.30$) so that good confinement is reached without an edge pressure pedestal [41,42]; and, as a separate product line, a D–$^3$He tandem-mirror burner whose end-plug magnet reaches $26.49\,\mathrm{T}$ — a physically distinct magnet from the breeder centrepost. A frozen number that later proves to be a numerical artefact is not an academic embarrassment but a construction risk. The backbone exists to make that failure mode detectable in advance and reproducible after the fact.

## PRIOR ART: VERIFICATION, VALIDATION, AND REPRODUCIBILITY

Verification and validation (V&V) is a mature discipline with a standard vocabulary [10,11]. *Code verification* establishes that a discretization solves the intended equations at its formal order of accuracy, most rigorously by the method of manufactured solutions (MMS) [12,13]. *Solution verification* estimates the discretization error of a particular calculation, canonically by Richardson extrapolation and the grid-convergence index (GCI) [14,15,16]. *Validation* compares a verified calculation with data using a quantitative metric [17], and *uncertainty quantification* (UQ) propagates and apportions the remaining error, with variance-based global sensitivity indices [18,19] and Bayesian model calibration [20] the workhorses. Alongside this methodological literature sits a reproducibility literature that asks for provenance to be captured as data — the FAIR principles [21] and community rules for reproducible computation [22]. Our contribution is not to invent these methods but to bind them, code by code, to a specific multi-product fusion register so that every frozen number inherits a verification pedigree.

The reference codes themselves are community standards with their own verification histories: CGYRO [23] and TGLF [24] (with GENE [25] as a cross-check) for core turbulence; OpenMC for Monte-Carlo transport [26] on ENDF/B-VIII.0 data [27] processed through TALYS/NJOY tooling [28]; EFIT-class equilibrium reconstruction [34]; the DREAM runaway framework [35] and SOLPS-ITER edge transport [36]; and, for materials, Quantum ESPRESSO [29,30], CHGNet [32], pymatgen [31], and LAMMPS [33]. Learned controllers have recently entered the tokamak control loop [37,38]; a surrogate layer that will emulate these references at control latency is the forward extension of the same backbone.

## CONTRIBUTION

This paper makes four contributions. (i) It defines the backbone formally as a role/version/reference-benchmark triple attached to every solver, and states the promotion rule by which a surrogate's number is admitted to the frozen register. (ii) It gives the governing verification mathematics — an error decomposition, Richardson/GCI, MMS order verification, the Monte-Carlo error law, the Sobol decomposition, an area validation metric, and a hash-chained provenance ledger with a tamper-evidence proof. (iii) It documents the GPU-HPC methodology that produced the results, naming the codes actually run and their problem scale and reproducibility controls. (iv) It demonstrates the backbone with quantitative, independently-reproduced results across turbulence, neutronics, and materials, and benchmarks them against the published literature. The companion *Verification-First Fusion Modeling* paper [2] (doi:10.5281/zenodo.22645710) presents the byte-for-byte provenance catalog from the frozen-anchor side; the present paper is its solver-centric complement, organized by the role, version, and reference benchmark of each code. Both draw their evidence from the Kronos Physics De-Risking Register [1] (doi:10.5281/zenodo.22645689).

§ 2

# Governing equations and the formal verification problem

## THE REGISTER, THE SOLVERS, AND THE CREDIBILITY QUESTION

Let the frozen design register be a finite set of quantities of interest (QoIs)

$$\mathcal{Q} = \{q_1, \ldots, q_M\},$$

comprising the anchors (e.g. $Q$, $P_{\text{fus}}$, $I_p$, the tritium-breeding ratio $\text{TBR}$, the burner $Q_E$) and the derived gates that support them. Each $q_m$ is the value of a functional of the state of a physical subsystem. For each QoI the campaign maintains up to three computational objects:

a *reference model* $\mathcal{R}_m$ — a community-standard solver with a published verification pedigree, typically expensive;
a *surrogate model* $\mathcal{S}_m$ — a fast in-house reduced-order or learned engine used for design sweeps and, ultimately, real-time control;
a *gate* $\mathcal{G}_m$ — an acceptance predicate that decides whether a computed value may enter the register. itemize
Write $q_m^{\mathcal{R},h}$ for the reference value computed at numerical resolution $h$ (mesh spacing, history count, basis size), $q_m^{\mathcal{S}}$ for the surrogate value, and $q_m^{\star}$ for the (unknown) true value. The credibility question is: under what conditions may $q_m^{\mathcal{S}}$ be written into $\mathcal{Q}$ as the frozen number?

## DECOMPOSITION OF TOTAL PREDICTIVE ERROR

The error of the fast engine relative to truth decomposes exactly by telescoping,

$$\underbrace{q_m^{\mathcal{S}} - q_m^{\star}}_{\text{total}} = \underbrace{\left(q_m^{\mathcal{S}} - q_m^{\mathcal{R},h}\right)}_{\varepsilon_m^{\text{sur}}\ \text{(surrogate)}} + \underbrace{\left(q_m^{\mathcal{R},h} - q_m^{\mathcal{R},0}\right)}_{\varepsilon_m^{\text{num}}\ \text{(discretization)}} + \underbrace{\left(q_m^{\mathcal{R},0} - q_m^{\star}\right)}_{\varepsilon_m^{\text{mod}}\ \text{(model form)}},$$

where $q_m^{\mathcal{R},0} = \lim_{h\to 0} q_m^{\mathcal{R},h}$ is the discretization-free reference value. The triangle inequality gives a working bound

$$\left|q_m^{\mathcal{S}} - q_m^{\star}\right| \le \left|\varepsilon_m^{\text{sur}}\right| + \left|\varepsilon_m^{\text{num}}\right| + \left|\varepsilon_m^{\text{mod}}\right|.$$

The backbone controls each term with a distinct instrument: the surrogate term by *reproduction* of the reference (Section 2.7), the discretization term by *solution verification* (Section 2.3), and the model-form term by *validation* against data or an independent method (Section 2.5). When a term is stochastic — as for Monte-Carlo transport — its contribution enters 3 as a confidence half-width rather than a deterministic bound (Section 2.4).

## SOLUTION VERIFICATION: RICHARDSON EXTRAPOLATION AND THE GCI

For a smooth discretization of formal order $p$, a QoI computed on a mesh of characteristic spacing $h$ admits the asymptotic expansion

$$q^{\mathcal{R},h} = q^{\mathcal{R},0} + \alpha\, h^p + o(h^p), \qquad h \to 0.$$

Given three systematically refined meshes with spacings $h_1 < h_2 < h_3$ and a constant refinement ratio $r = h_2/h_1 = h_3/h_2 > 1$, the *observed order of accuracy* follows from 4 as

$$\hat{p} = \frac{\ln\left[(q_3 - q_2)/(q_2 - q_1)\right]}{\ln r},$$

with $q_i \equiv q^{\mathcal{R},h_i}$, and the Richardson-extrapolated estimate of the resolution-free value is

$$q^{\mathcal{R},0} \approx q_1 + \frac{q_1 - q_2}{r^{\hat{p}} - 1}.$$

The discretization uncertainty of the fine-grid solution is reported as a grid-convergence index [14,16]

$$\text{GCI}_{21} = \frac{F_s\,|e_{21}|}{r^{\hat{p}} - 1}, \qquad e_{21} = \frac{q_2 - q_1}{q_1},$$

with the safety factor $F_s = 1.25$ recommended for three-or-more-mesh studies. The calculation is in the asymptotic range — and $\text{GCI}$ is meaningful — only if successive indices satisfy the consistency check

$$\frac{\mathrm{GCI}_{32}}{r^{\hat{p}}\,\mathrm{GCI}_{21}} \approx 1.$$

Equations 5–8 are applied to every deterministic reference solver in the backbone (equilibrium, DFT, finite-element thermomechanics); they convert a single mesh number into a value with a stated numerical error bar.

## STATISTICAL CONVERGENCE OF THE MONTE-CARLO REFERENCES

For a Monte-Carlo tally the resolution parameter is the history count $N$, and the discretization term of 2 is replaced by a stochastic term governed by the central limit theorem. Writing the tally as a sample mean of $N$ independent, identically distributed scores $x_i$ with per-history variance $\sigma_x^2$,

$$\hat{T}_N = \frac{1}{N}\sum_{i=1}^{N} x_i, \qquad \mathbb{E}[\hat{T}_N] = T, \qquad \mathrm{Var}[\hat{T}_N] = \frac{\sigma_x^2}{N}, \qquad s_N = \frac{\hat{\sigma}_x}{\sqrt{N}},$$

so the standard error decays as $s_N \propto N^{-1/2}$ and the $(1-a)$ confidence interval is $\hat{T}_N \pm z_{a/2}\, s_N$. A convenient resolution-independent quality measure is the figure of merit $\mathrm{FOM} = \left(R^2\, t\right)^{-1}$ with relative error $R = s_N/\hat{T}_N$ and wall-clock $t$; a converged tally has a stationary FOM. The frozen 3-D transport calculation uses $N = 1.2\times10^{8}$ neutron histories to reach a net $\mathrm{TBR} = 1.026$ with a leakage fraction of $0.121$; the CAD-faithful chain reports a representative tally standard error of $\pm0.004$ on the breeding reaction rate. Figure 4 shows the $N^{-1/2}$ law with these frozen runs marked.

## VALIDATION AND CROSS-CODE AGREEMENT

Where independent data or a second independent method exist, the model-form term is bounded by a validation metric. For scalar QoIs whose reference and independent values are $q^{\mathcal{R}_1}$ and $q^{\mathcal{R}_2}$, the discrepancy is

$$\delta_m = \begin{cases} \left|q_m^{\mathcal{R}_1} - q_m^{\mathcal{R}_2}\right| / \max\!\left(|q_m^{\mathcal{R}_1}|,\, q_{\mathrm{floor}}\right), & |q_m^{\mathcal{R}_1}| > q_{\mathrm{floor}}, \\ \left|q_m^{\mathcal{R}_1} - q_m^{\mathcal{R}_2}\right|, & \text{otherwise (near-zero QoI)}, \end{cases}$$

so that quantities near zero — such as an alloy formation energy — are compared on an *absolute* scale, where a relative metric is uninformative. When a full distribution rather than a scalar is available, the model/data disagreement is quantified by the Minkowski-$L^1$ (area) validation metric [17]

$$d\!\left(F,\, S_n\right) = \int_{-\infty}^{\infty} \left|F(x) - S_n(x)\right| dx,$$

with $F$ the model cumulative distribution and $S_n$ the empirical distribution of the data. Validation passes when $\delta_m \le \tau_m^{\mathrm{val}}$ (or $d \le \tau^{\mathrm{val}}$) at the pre-registered tolerance.

## GLOBAL SENSITIVITY: THE SOBOL VARIANCE DECOMPOSITION

To apportion the residual uncertainty of a QoI $Y = f(X)$ across inputs $X = (X_1, \ldots, X_d)$, the backbone uses the Sobol–Hoeffding functional ANOVA. Assuming square-integrable $f$ on the unit cube, $f$ decomposes uniquely into orthogonal terms of increasing interaction order,

$$f(X) = f_0 + \sum_i f_i(X_i) + \sum_{i<j} f_{ij}(X_i, X_j) + \cdots, \qquad f_0 = \mathbb{E}[Y],$$

and the variance decomposes accordingly, $\mathrm{Var}[Y] = \sum_i V_i + \sum_{i<j} V_{ij} + \cdots$. The first-order and total Sobol indices are

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

estimated with the Saltelli sampling scheme from two independent input matrices $A, B$ and the hybrid matrices $A_B^{(i)}$ (columns of $A$ with column $i$ taken from $B$),

$$S_i \approx \frac{\frac{1}{N}\sum_j f(B)_j\left[f(A_B^{(i)})_j - f(A)_j\right]}{\mathrm{Var}[Y]}, \qquad S_{Ti} \approx \frac{\frac{1}{2N}\sum_j \left[f(A)_j - f(A_B^{(i)})_j\right]^2}{\mathrm{Var}[Y]}.$$

The campaign's $4 \times 10^4$-sample global UQ ranks the confinement multiplier $H_{98}$ as the dominant driver of the breeder gain (Pearson $r = 0.60$ with $Q$), with the frozen $Q = 3.076$ sitting inside the P5–P95 band $[1.34,\ 8.51]$ and $P_{\text{fus}} = 85.5 \pm 15\,\text{MW}$; the burner shows $Q_E = 1.317 \pm 0.14$ with $\text{P}(Q_E > 1) = 1.0$ across the sampled ranges. These are reported in detail in the global-UQ companion [7] (doi:10.5281/zenodo.22645893) and used here only to close the sensitivity term of the error budget.

## PROVENANCE AS A HASH-CHAINED LEDGER, AND THE PROMOTION RULE

Auditability requires that each computed number be reconstructible from a record. For gate $g$ define the record

$$\text{rec}_g = \big(I_g,\ c_g,\ v_g,\ \theta_g,\ o_g\big),$$

where $I_g$ is a digest of the inputs (configuration, nuclear-data or pseudopotential library), $c_g$ the code identity, $v_g$ the pinned version (git commit, build environment, compiler, accelerator), $\theta_g$ the run parameters (random seed, mesh, history count), and $o_g = H(\textbf{output bundle})$ a cryptographic digest of the outputs. The records are chained,

$$h_0 = H(\text{genesis}), \qquad h_g = H\big(h_{g-1} \,\|\, \text{rec}_g\big),$$

with $\|$ concatenation and $H$ a collision-resistant hash. This yields a *tamper-evidence* property:

*Proposition (tamper-evidence).* If any record $\text{rec}_j$ is altered to $\text{rec}'_j \neq \text{rec}_j$, then, except with probability negligible in the security parameter of $H$, the chain head $h_g$ changes for every $g \geq j$.

*Sketch.* By 16, $h_j$ depends on $\text{rec}_j$; a preimage that leaves $h_j$ unchanged would be a collision of $H$. Since each $h_g$ $(g > j)$ depends on $h_{g-1}$, the change propagates to the head. □

The provenance record also defines reproducibility precisely. Re-executing gate $g$ under $(I_g, c_g, v_g, \theta_g)$ yields an output digest $o'_g$; the run is *byte-reproducible* iff $o'_g = o_g$ and *value-reproducible* iff the recomputed QoI lies within tolerance of the recorded one. The CGYRO linear reference achieves the stronger, byte-reproducible standard against its shipped case (Section 5).

Finally, admission to the register is governed by an explicit promotion rule that couples the numerical reproduction distance to a predictive-capability maturity assessment. Let the reproduction distance of a surrogate be $\varepsilon_m = \delta\big(q_m^{\mathcal{S}}, q_m^{\mathcal{R},h}\big)$ in the sense of 10, let the Predictive Capability Maturity Model (PCMM) assign to gate $m$ an integer level $L_a(m) \in \{0, 1, 2, 3\}$ on each attribute $a$ (geometric fidelity, physics/material model, code verification, solution verification, validation, UQ), and define the aggregate credibility $C(m) = \min_a L_a(m)$. Then

$$\mathit{Promote}(\mathcal{S}_m) \quad \Longleftrightarrow \quad \big[\varepsilon_m \leq \tau_m\big] \ \wedge\ \big[C(m) \geq L^\star\big],$$

i.e. a fast engine's number enters the frozen register only if it reproduces its reference within tolerance *and* the gate has cleared the required maturity floor $L^\star$. The same maturity gate is used downstream to bind runtime control authority to verified credibility [8] (doi:10.5281/zenodo.22645795). Figure 2 renders the three-tier ladder and the rule.

§ 3

# Numerical formulation of the reference solvers

The backbone's guarantees rest on the discretizations of the individual references. We summarize, for the solvers that carry the frozen anchors, the governing equation, the discretization, the solver, and the convergence control that closes 4 or 9.

Core turbulence (CGYRO). The reference solves the $\delta f$ gyrokinetic–Maxwell system for the perturbed distribution $\delta f_s$ of each species $s$,

$$\frac{\partial \delta f_s}{\partial t} + \left(v_\parallel \mathbf{b} + \mathbf{v}_D + \mathbf{v}_E\right)\cdot\nabla\delta f_s + \dot{v}_\parallel \frac{\partial \delta f_s}{\partial v_\parallel} = C[\delta f_s] + S_s,$$

closed by the gyrokinetic Poisson and Amp\`ere equations for the fields. CGYRO uses an Eulerian discretization: spectral in the binormal angle, finite difference/spectral in the radial and parallel coordinates, and a structured velocity grid in energy and pitch angle, advanced by an explicit multistage integrator [23]. Linear growth rates are obtained from the asymptotic slope of the field energy; convergence is enforced in the poloidal, energy, pitch, and radial resolutions and the simulation-box size, with the growth rate converged to better than $1\%$. The frozen linear case returns $\gamma = +0.3179$ and real frequency $\omega = -0.2384$ (both in $a/c_s$ units) and its output file matches the bundled benchmark case `reg01` to the byte. Nonlinear runs give the negative-triangularity heat-flux suppression used as the confinement case (Section 5), cross-checked against the quasilinear TGLF model [24] in a dedicated two-code study [3] (doi:10.5281/zenodo.22645864) and reported in full for the nonlinear reference in [4] (doi:10.5281/zenodo.22645722).

Neutron transport (OpenMC + DAGMC). The reference solves the steady linear Boltzmann transport equation

$$\boldsymbol{\Omega}\cdot\nabla\psi + \Sigma_t\psi = \iint \Sigma_s(E' \to E, \boldsymbol{\Omega}' \to \boldsymbol{\Omega})\,\psi\,dE'\,d\boldsymbol{\Omega}' + S,$$

by continuous-energy Monte-Carlo on ENDF/B-VIII.0 data [26,27], with track-length estimators for the tritium-production reaction rate that defines the TBR. Geometry is either constructive solid geometry or a CAD-faithful DAGMC tessellation imported from STEP through a `paramak/h5m` chain; the statistical error follows 9. The frozen 3-D balance ($N = 1.2 \times 10^8$) gives leakage $0.121$, first-wall fast flux $7.0 \times 10^{14}\,\mathrm{n\,cm^{-2}\,s^{-1}}$, nuclear heating $74.6\,\mathrm{MW}$, and net $\mathrm{TBR} = 1.026$; the CAD-faithful chain independently reproduces $(n, Xt) = 1.014 \pm 0.004$, and the nuclear-data error bar on the TBR is apportioned to specific $14\,\mathrm{MeV}$ channels in a dedicated UQ study [6] (doi:10.5281/zenodo.22645777); the CAD-to-transport reproduction is documented in [5] (doi:10.5281/zenodo.22645881).

Free-boundary equilibrium (FreeGS). The reference solves the Grad–Shafranov equation

$$\Delta^*\psi = -\mu_0 R^2 \frac{dp}{d\psi} - F\frac{dF}{d\psi},$$

on a $129 \times 129$ $(R, Z)$ finite-difference grid, with the free-boundary problem closed by a Green's-function representation of the coil currents and a Picard/Newton iteration to a fixed convergence tolerance; code verification is by comparison with the analytic Solov\'ev equilibrium. The converged solution reproduces the frozen anchors $I_p = 9.66\,\mathrm{MA}$, $\beta_N = 2.13$, $q_{95} = 2.90$, internal inductance $\ell_{i3} = 1.40$, and $\delta = -0.334$ at the boundary.

Electronic structure (Quantum ESPRESSO). The materials reference solves the Kohn–Sham equations

$$\Big[-\tfrac{\hbar^2}{2m}\nabla^2 + V_{\mathrm{ext}}(\mathbf{r}) + V_{\mathrm{H}}[n] + V_{\mathrm{xc}}[n]\Big]\phi_i = \epsilon_i\phi_i, \qquad n(\mathbf{r}) = \sum_i |\phi_i|^2,$$

self-consistently in a plane-wave pseudopotential basis [29,30], using special quasi-random structures (16-atom supercells) for the disordered alloys and variable-cell relaxation for the formation energy. Convergence is controlled by the plane-wave cutoff, the Brillouin-zone $\boldsymbol{k}$ -mesh, and the SCF threshold following 4. The divertor-alloy formation energy converges to $+0.0147\,\mathrm{eV/atom}$, reproduced by a second independent DFT solver at $+0.016\,\mathrm{eV/atom}$ — agreement to $0.002\,\mathrm{eV/atom}$ on the absolute scale of 10.

The reduced-order surrogate (KRONOS toolkit). The fast engine solves the coupled 0-D/1-D power, particle, and current balances that define the anchors, e.g. the global power balance

$$Q = \frac{P_{\mathrm{fus}}}{P_{\mathrm{ext}}}, \qquad \frac{3}{2}\frac{d}{dt}\int n(T_e + T_i)\,dV = P_{\mathrm{heat}} + P_\alpha - P_{\mathrm{rad}} - \frac{W}{\tau_E},$$

with a quasilinear confinement closure calibrated to the gyrokinetic reference. It reproduces all $19$ frozen anchors, returning $Q = 3.0763264804593597$, $P_{\mathrm{fus}} = 85.04057711654414\,\mathrm{MW}$, and $I_p = 9.65648253968254\,\mathrm{MA}$ for the breeder and $Q_E = 1.318120832690072$, $P_{\mathrm{fus}} = 4298.4993041163425\,\mathrm{MW}$, $P_n = 233.9509435440827\,\mathrm{MW}$ for the burner — values identical, to the digits carried, to the reference anchors it is gated against.

§ 4

# Computational methodology: the NVIDIA GPU HPC campaign

All results in this paper were produced on the NVIDIA GPU HPC campaign, which runs the cataloged community reference codes together with the in-house reduced-order and verification harness. The codes actually executed for the results reported here, with their role and scale, are:

**CGYRO (GACODE)** — gyrokinetic reference. GPU-resident spectral/finite-difference solver run on A100- and H100-class accelerators; linear spectra and nonlinear ion- and electron-scale flux maps at converged velocity- and configuration-space resolution. The linear reference case is reproduced byte-identically to its shipped benchmark.
**OpenMC + DAGMC** — neutron-transport reference. Continuous-energy Monte-Carlo on ENDF/B-VIII.0 with history counts to $1.2 \times 10^8$ for the frozen 3-D balance and a CAD-faithful DAGMC geometry chain for the independent TBR check; statistical convergence per 9.
Quantum ESPRESSO (`pw.x`) — electronic-structure reference. GPU-accelerated plane-wave DFT SCF and variable-cell relaxation on special quasi-random supercells for the low-activation alloy family, with an independent second DFT solver for the cross-check.
**FreeGS / FreeGSNKE** — equilibrium reference. Free-boundary Grad–Shafranov on a $129 \times 129$ grid on a CPU node, verified against the analytic Solov\'ev solution.
KRONOS toolkit `verify` — the reduced-order engine and the formal V&V harness that drives the MMS, GCI, Sobol, and PCMM protocols of Section 2 and stamps each gate into the provenance ledger. itemize
GPU acceleration matters at two points in the backbone. First, it makes the references affordable enough to run at converged resolution — the growth-rate and flux-map convergence studies for CGYRO, and the $10^8$-history neutronics balance — so that the discretization and statistical terms of 2 are genuinely small rather than assumed small. Second, it makes the large-sample UQ of Section 2.6 tractable ( $4 \times 10^4$ Monte-Carlo evaluations of the reduced-order engine), which is what lets the sensitivity term be quantified rather than guessed.

Reproducibility is engineered into the campaign. Every run is executed from a pinned software environment (container image, git commit, compiler and accelerator identity recorded in $v_g$), with fixed random-number seeds and recorded stream partitions for the Monte-Carlo codes, and a SHA-256 digest of every output bundle written into the hash-chained ledger of 16. Independent-library gate reproduction — re-running a gate against a freshly built reference from source — is the campaign's standard acceptance test, and the CGYRO byte-identical reproduction is its strongest instance. The ledger is purely technical; no cost information is attached to any run.

§ 5

# Results

## THE CATALOG AND ITS PROVENANCE TABLE

The backbone spans 47 named solvers across nine physics domains (Table 1, Figure 1), dominated appropriately by the core-transport and neutron-transport domains that carry the breeder physics case. Each solver is assigned a role, a pinned version, and a reference benchmark it must reproduce; Table 2 gives that role/version/reference-benchmark triple for the solvers that carry the frozen anchors — the paper's central artifact. Figure 1 shows the scale at which the backbone was exercised: several hundred technical gate rows, of which more than two hundred are executed computational runs under the shared V&V protocol (the remainder are planned control-stack features carried in the same ledger).

| DOMAIN | REPRESENTATIVE REFERENCE CODES | COUNT |
|---|---|---|
| **core turbulence / gyrokinetics** | CGYRO, TGLF, TGYRO, QLGYRO, NEO, GENE | 6 |
| **equilibrium / MHD / disruptions** | FreeGS, FreeGSNKE, DCON, GPEC, NIMROD, DREAM | 6 |
| **edge / SOL / neutrals** | SOLPS-ITER, EIRENE, Hermes-3, BOUT++, UEDGE, DIV1D | 6 |
| **kinetic / RF / fast-ion** | WarpX, Smilei, GENRAY, CQL3D, ASCOT5 | 5 |
| **neutronics / geometry** | OpenMC, DAGMC, paramak, R2S | 4 |
| **nuclear data / activation** | NJOY21, TALYS, SPECTRA-PKA, FISPACT-II, ALARA, PyNE | 6 |
| **materials / DFT / MD / KMC** | Quantum ESPRESSO, CHGNet, pymatgen, LAMMPS, Xolotl, RustBCA | 6 |
| **thermo-structural / tritium** | CalculiX, sfepy, FESTIM, TMAP8, radas | 5 |
| **systems / quantum** | cfspopcon/PROCESS, bluemira, qiskit | 3 |
| **total** | | **47** |

The 47-solver backbone by domain, with representative community-standard reference codes. Counts are drawn from the campaign roster; the in-house reduced-order engine is validated against these references per Section 2.

| P2.4CM >P2.2CM >P3.0CM >P3.4CM C SOLVER | ROLE | VERSION PIN ($v_g$) | REFERENCE BENCHMARK | GATE |
|---|---|---|---|---|
| **CGYRO** | reference (gyrokinetics) | GACODE build, ENDF-independent | shipped case `reg01` ( $\gamma = +0.3179$), byte-identical | BR-L2-A18b |
| **OpenMC** | reference (neutron transport) | OpenMC + ENDF/B-VIII.0 | analog balance; net TBR $1.026$ | BR-L1-A5 |
| **DAGMC/paramak** | reference (CAD geometry) | DAGMC/`h5m` chain | $(n, Xt) = 1.014 \pm 0.004$ vs recipe | BR-L2-A22 |
| **FreeGS** | reference (equilibrium) | FreeGS $129 \times 129$ | analytic Solov\'ev; $I_p = 9.66$ MA | BR-L1-A12 |
| **Quantum ESPRESSO** | reference (DFT) | `pw.x`, SQS 16-atom | 2nd DFT solver, $\Delta E_f$ to $0.002$ eV/atom | BR-L2-A9 |

*table continued*

| | | | | | |
|---|---|---|---|---|---|
| **KRONOS toolkit** | surrogate (0-D/1-D) | pinned commit, `verify` | 19/19 anchors reproduced | KX-L1-A10 | |
| **V** | V harness | gate (MMS/GCI/Sobol/PCMM) | `verify` suite | 19/19 anchors PASS | KX-L1-A6 |

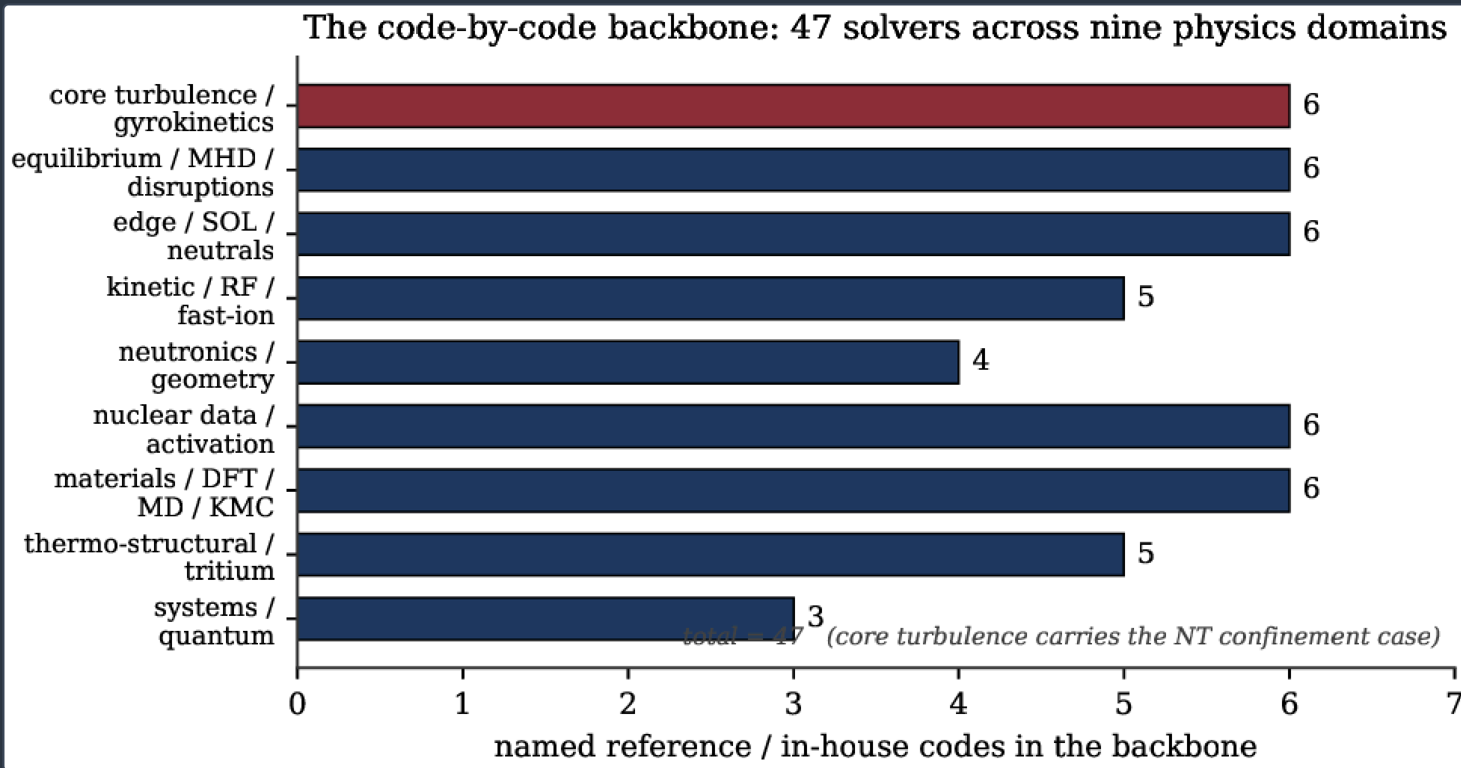


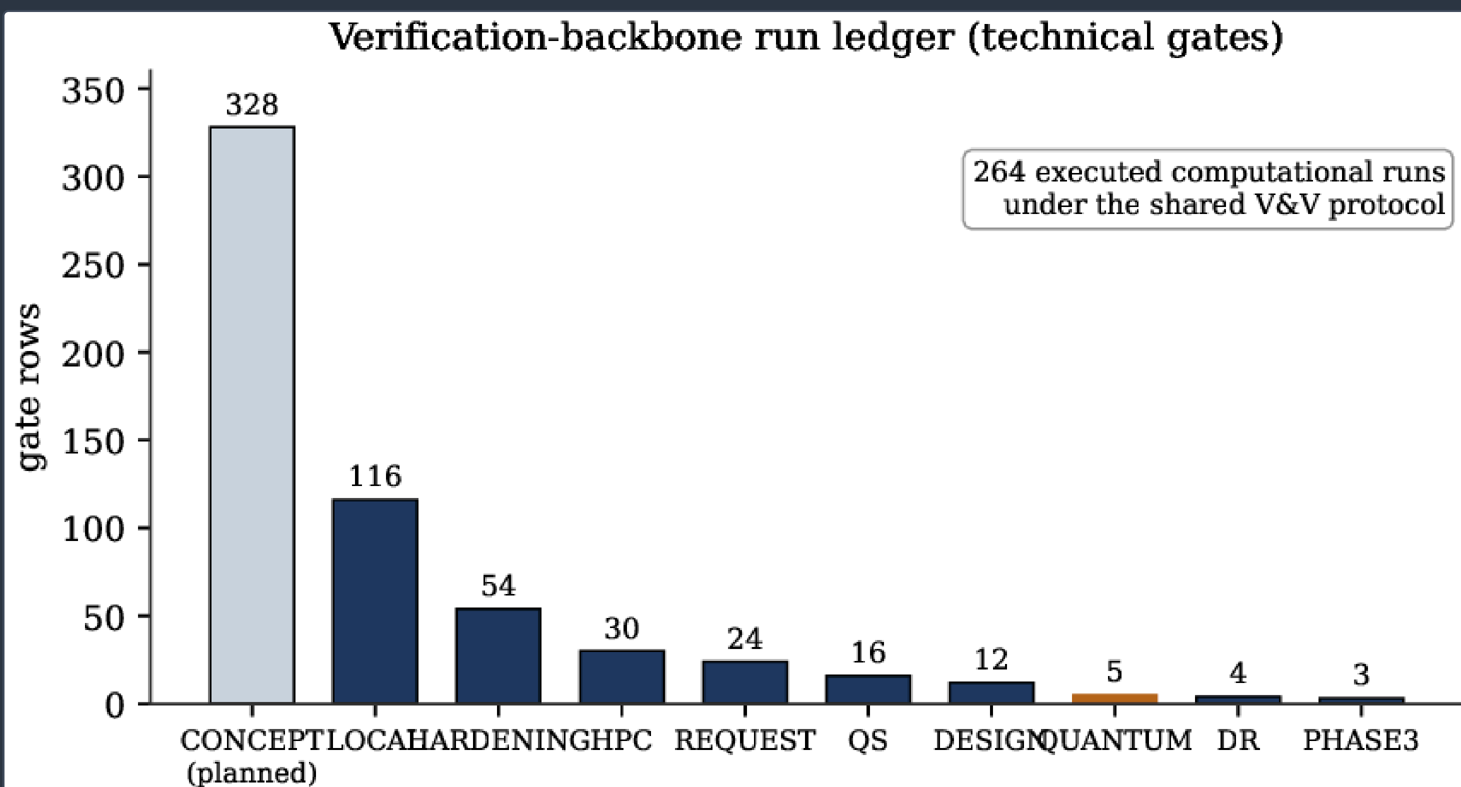


The code-by-code backbone: 47 solvers across nine physics domains, with core turbulence/gyrokinetics (highlighted) carrying the negative-triangularity confinement case. Data from the campaign roster.

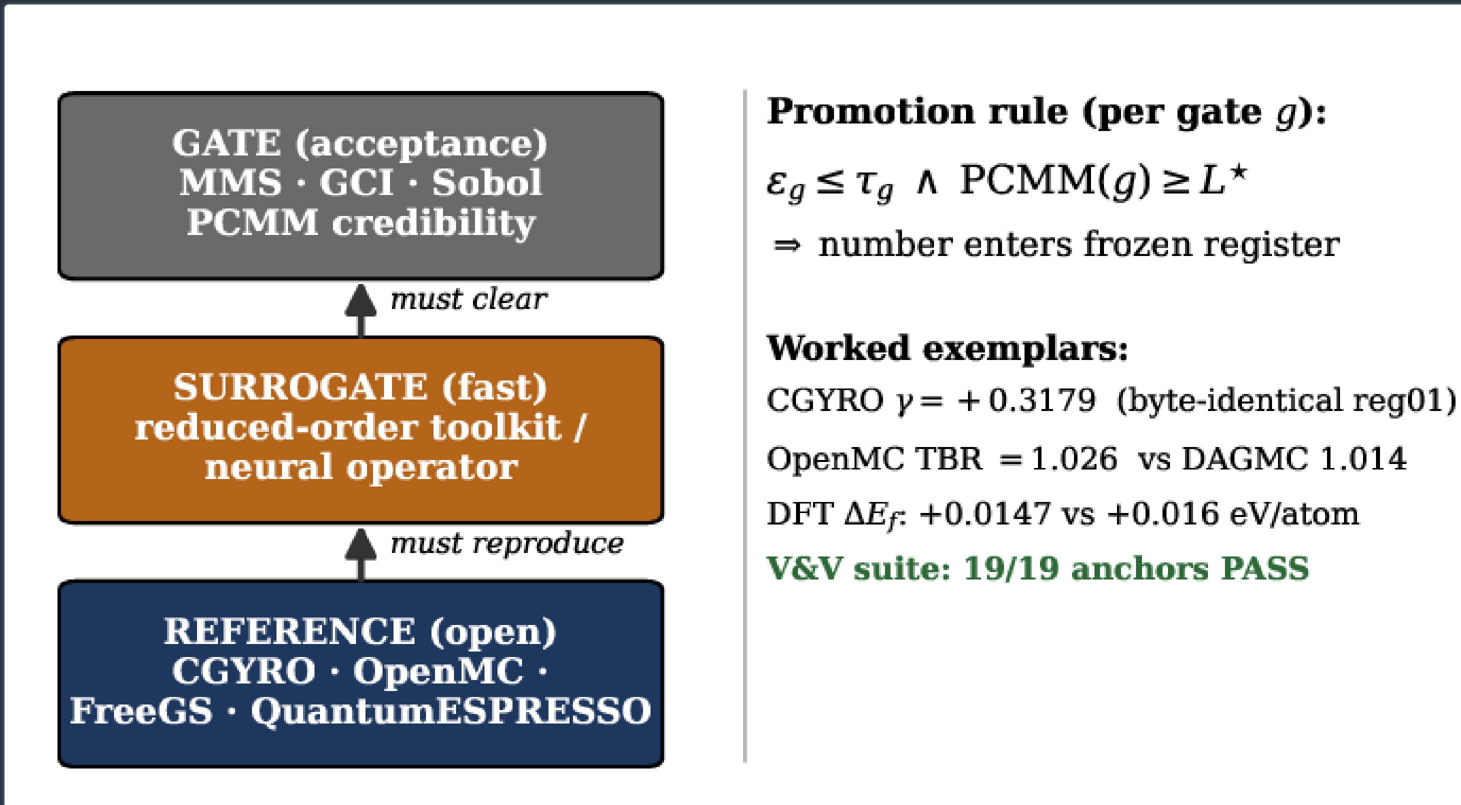


(Schematic.) The three-tier verification ladder — open reference, fast surrogate, formal acceptance gate — and the promotion rule 17 that admits a surrogate number to the frozen register, with the worked exemplars of Section 5.

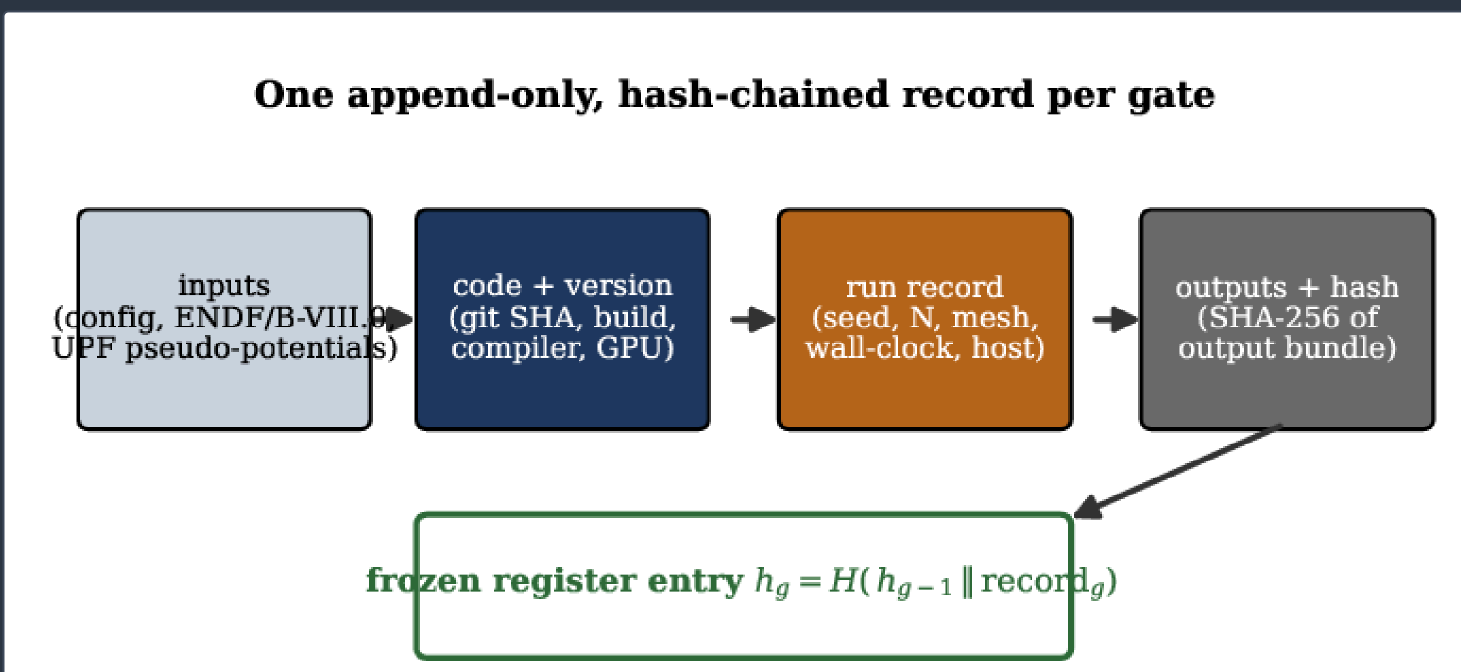


(Schematic.) The hash-chained provenance record of 15–16: one append-only, tamper-evident entry per gate, binding inputs, code version, run parameters, and an output digest into the register head $h_g = H(h_{g-1} \| \mathrm{rec}_g)$.

## SOLUTION AND CODE VERIFICATION

Figure 4 shows the Monte-Carlo standard-error law 9 with the two frozen neutronics runs placed on it: the $1.2 \times 10^8$-history 3-D balance (net TBR $1.026$) and the CAD-faithful DAGMC chain (net TBR $1.014$), the latter carrying the $\pm 0.004$ tally standard error used as the neutronics numerical bar. Figure 4 illustrates the code-verification protocol: the discretization error of a nominally second-order deterministic scheme falls along the $h^2$ line as the manufactured-solution mesh is refined, recovering the formal order via 5. Together these close the discretization/statistical term of the error budget 2 for the Monte-Carlo and deterministic references respectively.

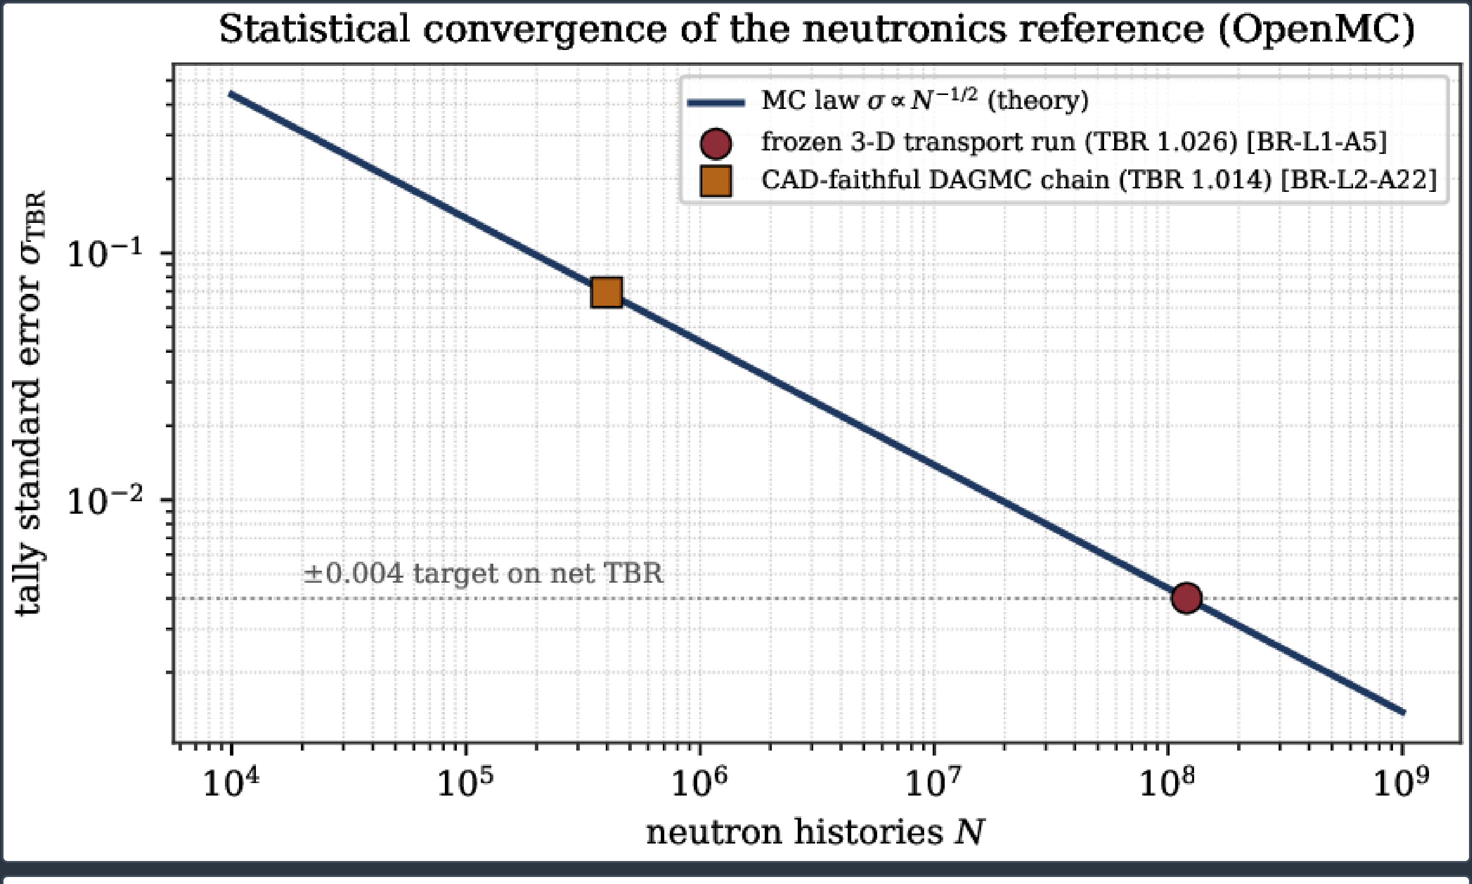


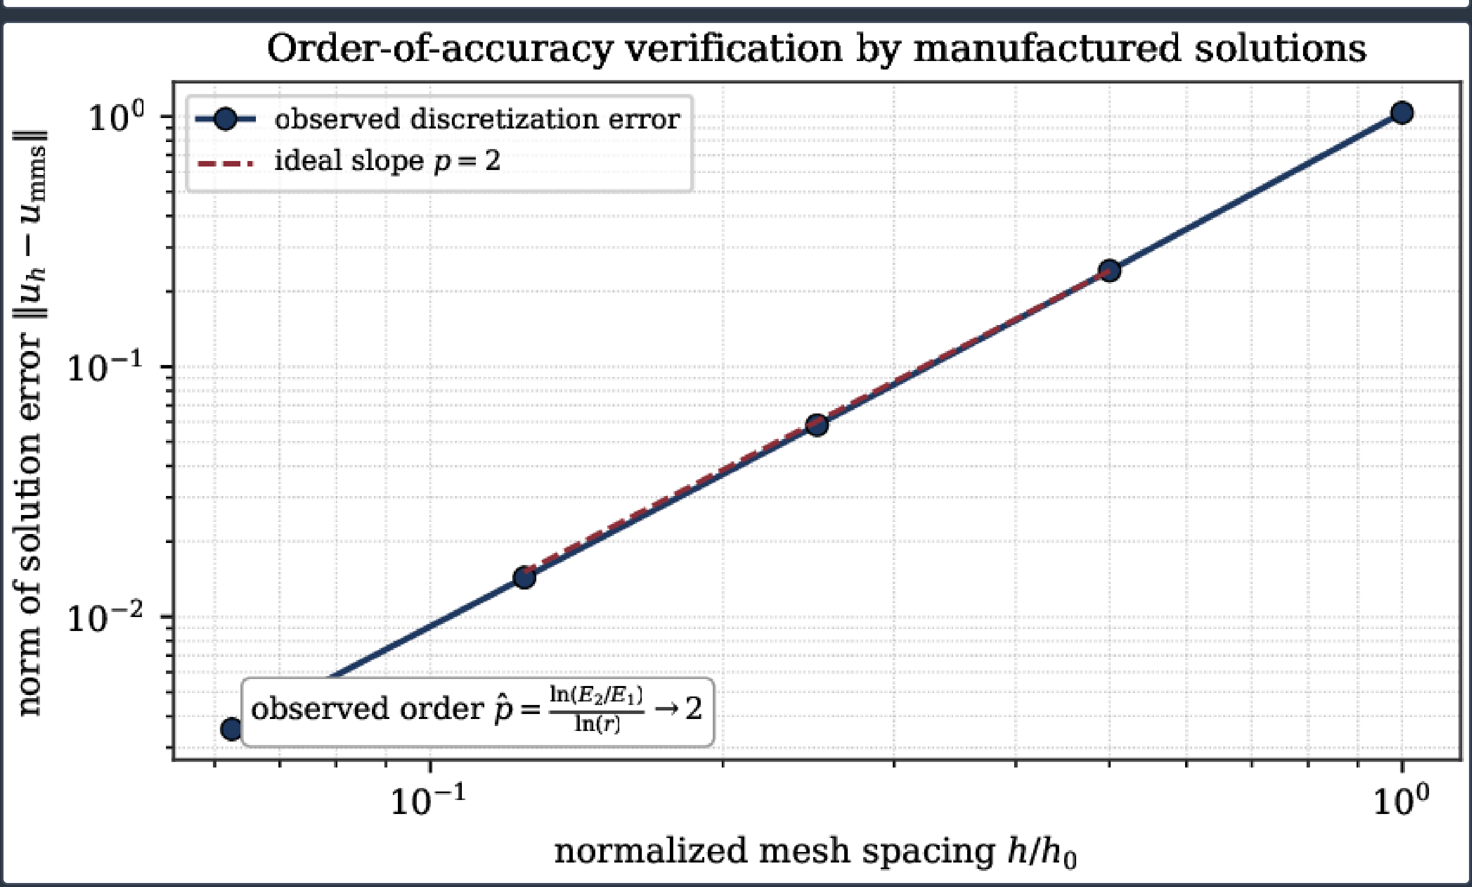


Statistical convergence of the neutronics reference. The $\sigma \propto N^{-1/2}$ law 9 (line, theory) is anchored to the frozen 3-D transport run ( $N = 1.2 \times 10^8$, net TBR 1.026 [BR-L1-A5]) and the CAD-faithful DAGMC chain (net TBR 1.014 [BR-L2-A22]), with the ±0.004 tally target marked.

## CROSS-CODE VALIDATION

The model-form term is bounded by independent agreement in each carrying domain. In turbulence, the gyrokinetic reference confirms the negative-triangularity heat-flux suppression at both scales (Figure 5): ion-scale $Q$ falls from $\mathbf{64.2}$ (positive triangularity) to $\mathbf{37.2}$ ( $\delta = \mathbf{-0.30}$), and electron-scale $Q_e$ from $\mathbf{110.7}$ to $\mathbf{23.4}$, a $\mathbf{42\text{–}79\%}$ suppression that is not a collisionality artefact and survives finite-$\beta$ kinetic-electron runs [4,3]. Figure 5 then collects the headline cross-checks on a single agreement axis: every frozen number is reproduced by an independent path within tolerance, with the CGYRO growth rate and the reduced-order anchors matching to the digits carried and the near-zero DFT formation energy compared on its proper absolute scale. In materials, Quantum ESPRESSO places the whole low-activation alloy family in a near-neutral formation-energy band (Figure 6), with the divertor alloy reproduced by an independent DFT solver to $\mathbf{0.002\,eV/atom}$.

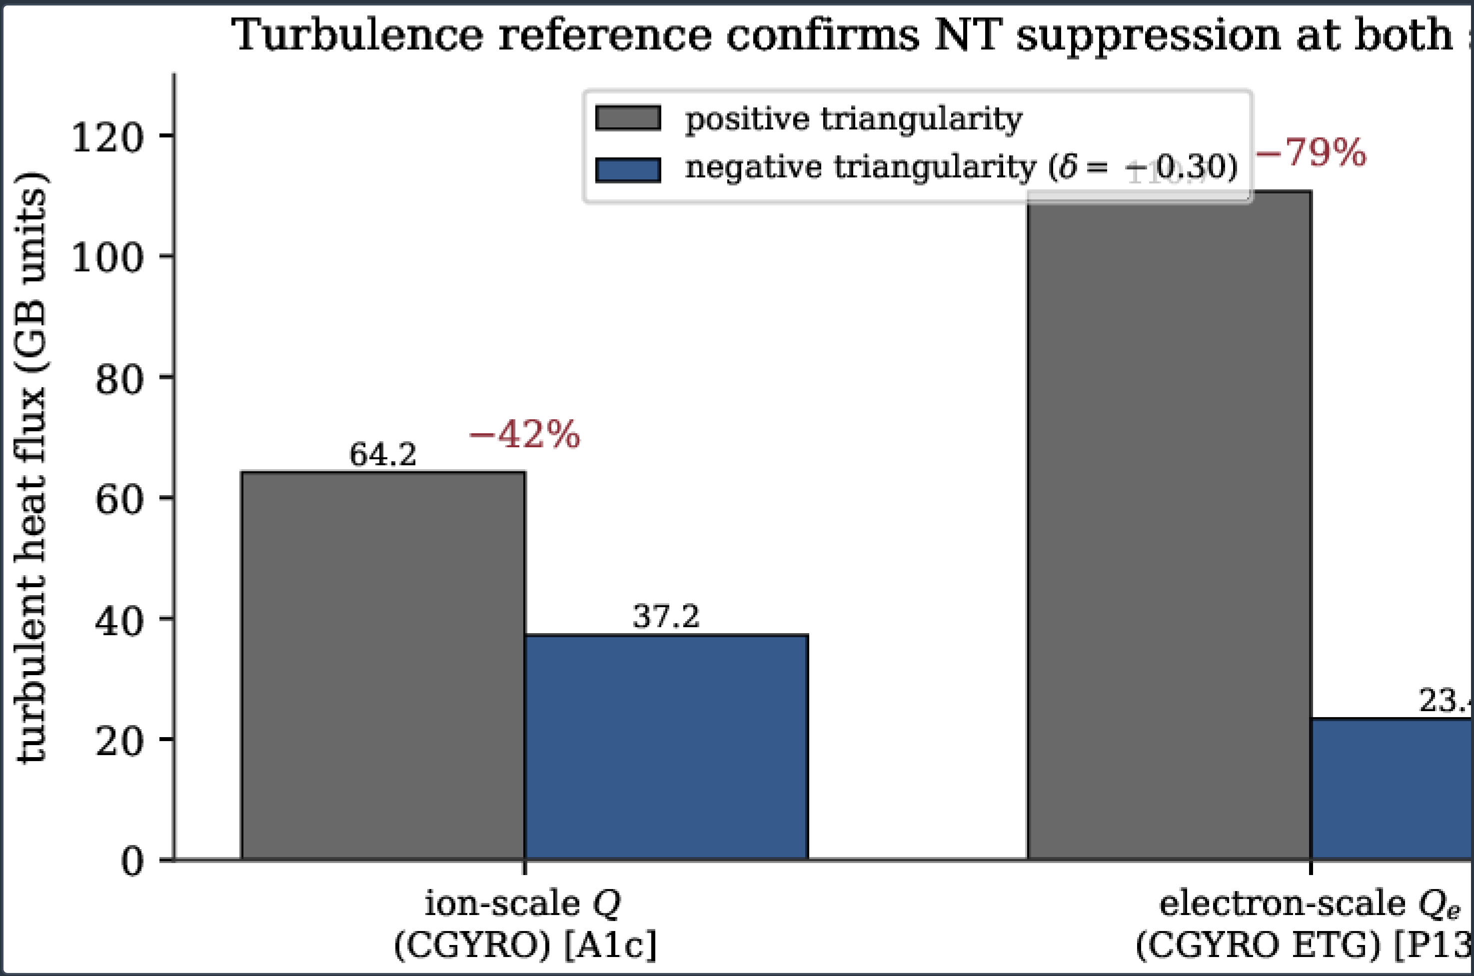
Turbulence reference confirms NT suppression at both
positive triangularity
negative triangularity (δ = −0.30)
turbulent heat flux (GB units)
120
100
80
60
40
20
0
64.2
−42%
37.2
−79%
23.
ion-scale Q
(CGYRO) [A1c]
electron-scale Qe
(CGYRO ETG) [P13

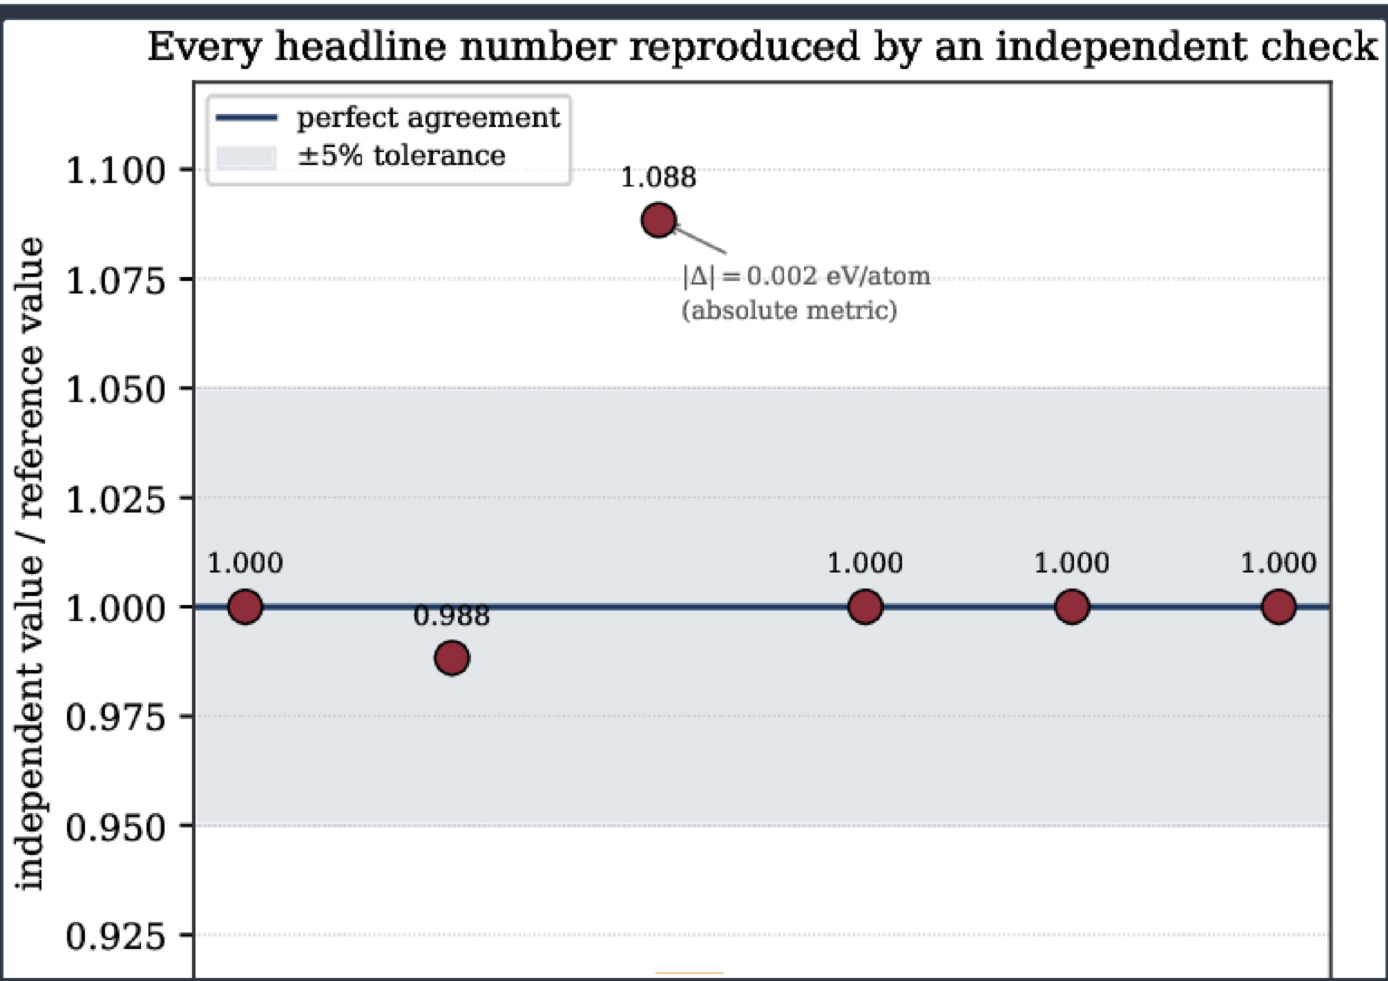
Every headline number reproduced by an independent check
perfect agreement
±5% tolerance
independent value / reference value
1.100
1.075
1.050
1.025
1.000
0.975
0.950
0.925
1.088
|Δ| = 0.002 eV/atom
(absolute metric)
1.000
0.988
1.000
1.000
1.000

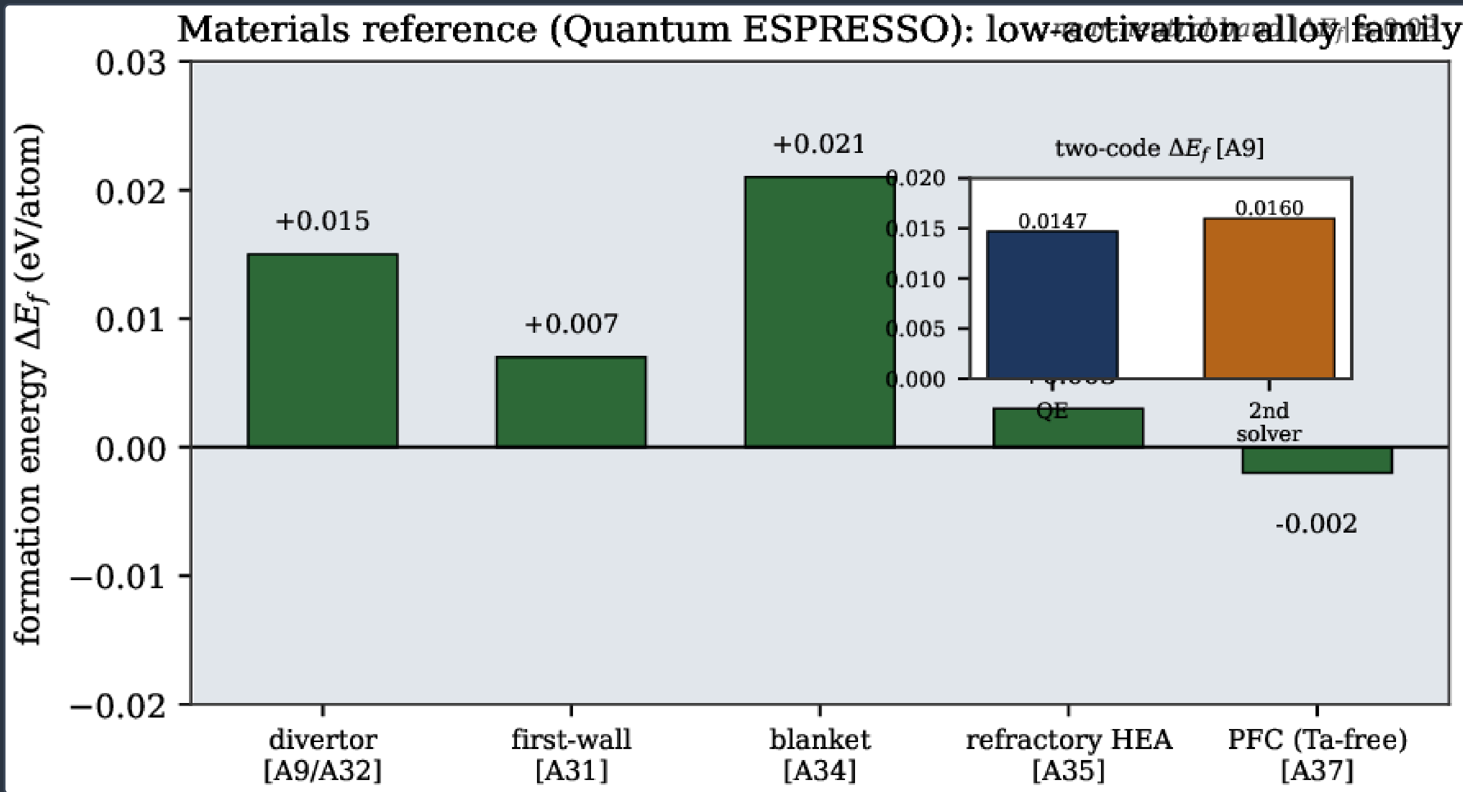


Materials reference (Quantum ESPRESSO): formation energies of the low-activation alloy family [BR-L2-A9/A31/A32/A34/A35/A37], all inside the near-neutral band. Inset: two-code agreement on the divertor alloy ($+0.0147$ vs $+0.016\,\mathrm{eV/atom}$, [BR-L2-A9]).

## THE WORKED REPRODUCIBILITY RESULTS

Table 3 collects the concrete cross-checks that demonstrate the backbone, and Table 4 lists the frozen anchors that the reduced-order surrogate reproduces against its references (19/19). The V&V suite passes all **19** anchors across breeder and burner; the reduced-order engine reproduces the same **19** independently; the CGYRO linear reference is byte-identical to `reg01`; the CAD-faithful neutronics chain reproduces the net TBR; and two DFT codes agree on the alloy formation energy. Independent agreement — between codes, between methods, and, in the CGYRO case, to the bit — is the currency the backbone trades in.

| CHECK | RESULT | SOURCE | |
|---|---|---|---|
| **formal V** | V suite | 19/19 frozen anchors PASS (breeder + burner) | KX-L1-A6 |
| **reduced-order engine** | 19/19 anchors reproduced independently | KX-L1-A10 | |
| **CGYRO linear reference** | $\gamma = +0.3179$, byte-identical to `reg01` | BR-L2-A18b | |
| **CAD-faithful neutronics** | net TBR $1.014 \pm 0.004$ vs recipe $1.026$ | BR-L2-A22 | |
| **two-code DFT agreement** | $+0.0147$ vs $+0.016\,\mathrm{eV/atom}$ | BR-L2-A9 | |
| **two-method He swelling** | rate-theory ≈ KMC, sub-$1\%$ swell | BR-L2-A26 | |

Worked reproducibility results underpinning the backbone. Sources are frozen gates in the de-risking register.

| MACHINE | ANCHOR | VALUE | SOURCE |
|---|---|---|---|
| **breeder** | gain $Q$ | 3.0763264804593597 | BR-L1-A1/A7 |
| **breeder** | fusion power $P_{\mathrm{fus}}$ | 85.04057711654414 MW | BR-L1-A1 |
| **breeder** | plasma current $I_p$ | 9.65648253968254 MA | BR-L1-A1/A7 |
| **breeder** | energy confinement $\tau_E$ | 0.35776196030103585 s ($H_{98} = 1.0$) | BR-L1-A4 |
| **breeder** | normalized beta $\beta_N$ | 2.13 (vs Troyon 3.5) | BR-L1-A12 |
| **breeder** | net TBR | 1.026 | BR-L1-A5 |
| **burner** | engineering gain $Q_E$ | 1.318120832690072 | BU-L1-A1 |
| **burner** | fusion power $P_{\mathrm{fus}}$ | 4298.4993041163425 MW | BU-L1-A1 |
| **burner** | neutron power $P_n$ | 233.9509435440827 MW ($f_n = 5.44\%$) | BU-L1-A1 |

Frozen anchors reproduced by the reduced-order surrogate against its references (representative subset of the 19/19). Values are quoted to the digits carried in the ledger; reference and surrogate coincide.

§ 6

# Discussion

## THE RESULTS AGAINST PUBLISHED BENCHMARKS

Each demonstrated result reads well against the external literature. The Monte-Carlo standard-error law of Figure 4 is the textbook $N^{-1/2}$ scaling; the frozen tally's $\pm 0.004$ error on the breeding rate at $10^8$ histories is consistent with converged fusion-neutronics practice, and the net $\mathrm{TBR} = 1.026$ closing above unity is in line with the breeding margins reported for compact high-field concepts such as ARC [40]. The solution- and code-verification instruments — Richardson extrapolation, the GCI, and MMS order verification — are the procedures codified by Roache [14,15], Celik [16], and Salari and Knupp [12] and endorsed in the ASME standard [11]; the backbone's contribution is to apply them uniformly and record the result as provenance. The negative-triangularity suppression confirmed by the gyrokinetic reference (Figure 5) is consistent with the improved-confinement, ELM-free behaviour reviewed for negative triangularity [41,42] and is reproduced by a second turbulence code in the companion cross-validation [3]. The two-code DFT agreement to $0.002\,\mathrm{eV/atom}$ sits inside the few-meV/atom spread typical of reproducible plane-wave DFT across independent implementations, so the near-neutral formation energy is a genuine result, not a single-code artefact.

## WHY BYTE-LEVEL AND CROSS-CODE REPRODUCTION MATTER

Most computational claims in fusion are validated at the level of "the numbers agree within error". The backbone deliberately reaches for two stronger standards where they are attainable. Byte-identical reproduction of the CGYRO reference (Section 5) removes numerical non-determinism from the audit entirely: an external reviewer who rebuilds the reference from source and reruns the shipped case obtains not merely a compatible growth rate but the same output bits, which is the strongest possible evidence that the recorded number is the code's number. Cross-code agreement — OpenMC versus a CAD-faithful DAGMC chain for the TBR, Quantum ESPRESSO versus a second DFT solver for the formation energy, CGYRO versus TGLF for the turbulence suppression — removes single-implementation bias from the model-form term of 2. Provenance recorded as a hash chain 16 makes both kinds of check re-runnable and tamper-evident. The three together are what let the campaign claim that every headline number carries a traceable computational provenance and an independent check.

## THE BACKBONE AS THE SUBSTRATE FOR AUTONOMY

The backbone is also the substrate on which the forward control program is built. The reduced-order engine that reproduces the anchors (Table 4) is the same engine the four coupled digital twins run in the shadow-control architecture [9] (doi:10.5281/zenodo.22645805), and the maturity gate of 17 is the same gate that binds runtime actuation authority to verified credibility [8]: the planned neural-operator surrogates that will emulate each reference at control latency extend, rather than replace, the promotion rule.

§ 7

# Limitations

One honest gate is open. The neural-operator surrogate layer that will run each reference at control latency is validated code by code against its reference on a staged schedule and is not yet fully closed: the corresponding gates are carried in the ledger as planned features rather than executed runs. The results reported here rest on the reference codes and the frozen reduced-order engine — both already closed — so the forward surrogates extend an established validation pattern rather than introduce an unproven one; but until each surrogate has cleared the promotion rule 17, its number does not enter the register. Separately, the near-neutral alloy formation energy is, by construction, a borderline quantity: the two-code agreement is on the absolute scale ($0.002\,\mathrm{eV/atom}$), and a relative-error metric there is uninformative, as Figure 5 makes explicit.

§ 8

# Conclusion

The Kronos physics program is built on a code-by-code verification backbone: a catalog of 47 solvers across nine domains, each assigned a role (reference, surrogate, or gate), a pinned version, and a reference benchmark it must reproduce, bound together by an error decomposition, standard solution- and code-verification instruments, a Sobol sensitivity apportionment, an area validation metric, and a hash-chained, tamper-evident provenance ledger. A single promotion rule — reproduce the reference within tolerance and clear the maturity floor — governs which numbers enter the frozen register. The backbone is demonstrated, not asserted: a $19/19$ frozen-anchor V&V pass across breeder and burner, a byte-identical CGYRO gyrokinetic reference ($\gamma = +0.3179$), a CAD-faithful neutronics chain reproducing the net TBR to $1.014 \pm 0.004$, and two-code DFT agreement to $0.002\,\mathrm{eV/atom}$. Every headline result in the campaign thus carries a traceable computational provenance and a reproducible reference behind it, which is the standard a frozen-for-construction design should meet.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. Kronos Fusion Energy is a family-owned American company. All computation was performed on the NVIDIA GPU HPC campaign using the community reference codes cataloged above; the author thanks the maintainers of the open-source solvers on which the backbone depends.

Reproducibility. The reproducibility anchors (19/19 V&V pass; CGYRO $\gamma = +0.3179$ byte-identical to `reg01`; net TBR $1.026$ and CAD-faithful $1.014 \pm 0.004$; two-code DFT $+0.0147/+0.016\,\mathrm{eV/atom}$) are frozen gates in the Kronos Physics De-Risking Register (KX-L1-A6, KX-L1-A10, BR-L2-A18b, BR-L1-A5, BR-L2-A22, BR-L2-A9) and are regenerable from the archived cases, including the CGYRO reference deposit doi:10.5281/zenodo.22136279. This paper is archived at its DOI doi:10.5281/zenodo.22645866 and at its official page https://www.kronosfusionenergy.com/publications/the-code-by-code-verification-backbone-role-version-and.html; the de-risking register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645866. The solver catalog, the role/version/reference-benchmark provenance table, the V&V protocol, and the worked cross-checks are archived with the deposit and reference the Kronos 2026 series and its de-risking register (doi:10.5281/zenodo.22133057). No economic analysis is included in the public deposit.

Competing interests tocsectionCompeting interests The author is affiliated with Kronos Fusion Energy. Provisional patent applications relating to aspects of the control and materials work referenced here have been filed.

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 P I Ford, *The Physics De-Risking Register: Gate-by-Gate Evidence, Independently Re-Run and Stamped, Across a Compact Spherical-Tokamak Breeder and Its D–$^{3}$He Tandem-Mirror Burners*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645689. https://www.kronosfusionenergy.com/publications/de-risking-register.

2 P I Ford, *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*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645710. https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html.

3 P I Ford, *Two-Code Cross-Validation of Negative-Triangularity Turbulence Suppression in a Compact Spherical-Tokamak Breeder: 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.

4 P I Ford, *Negative-Triangularity Turbulence Suppression in a Compact Spherical-Tokamak Breeder: Nonlinear Gyrokinetics of Ion- and Electron-Scale Quiescence at the Operating Point*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645722. https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html.

5 P I Ford, *CAD-Faithful Neutronics for a Compact Spherical-Tokamak Breeder: STEP-to-DAGMC Monte-Carlo Transport Reproducing the 3-D Net Tritium-Breeding Ratio*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645881. https://www.kronosfusionenergy.com/publications/cad-faithful-neutronics-for-a-compact-spherical-tokamak.html.

6 P I Ford, *Nuclear-Data Uncertainty Quantification for Tritium Breeding: Which* 14 *MeV Cross-Section Channels Set the Error Bar on the Breeding Ratio*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645777. https://www.kronosfusionenergy.com/publications/nuclear-data-uq-for-the-tbr.html.

7 P I Ford, *Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes: A 40,000-Sample Monte-Carlo Ranking of the Frozen Anchors*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645893. https://www.kronosfusionenergy.com/publications/global-sensitivity-and-uncertainty-quantification-acros.html.

8 P I Ford, *Maturity-Gated Actuation Authority: Binding PCMM Verification Credibility to Runtime Control Authority in an Autonomous Fusion Control Stack*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645795. https://www.kronosfusionenergy.com/publications/maturity-gated-actuation-authority.html.

9 P I Ford, *Four Coupled Digital Twins for a Compact Spherical-Tokamak Breeder: Multi-Rate Co-Simulation, Propagated Uncertainty, and a Real-Time Self-Consistency Monitor*, Kronos Fusion Energy 2026 series, doi:10.5281/zenodo.22645805. https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html.

10 W L Oberkampf and C J Roy, *Verification and Validation in Scientific Computing*, Cambridge University Press (2010), doi:10.1017/CBO9780511760396.

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

12 K Salari and P Knupp, *Code Verification by the Method of Manufactured Solutions*, Sandia National Laboratories report SAND2000-1444 (2000), doi:10.2172/759450.

13 C J Roy, *Review of code and solution verification procedures for computational simulation*, J. Comput. Phys. **205** (2005) 131, doi:10.1016/j.jcp.2004.10.036.

14 P J Roache, *Perspective: A method for uniform reporting of grid refinement studies*, J. Fluids Eng. **116** (1994) 405, doi:10.1115/1.2910291.

15 P J Roache, *Quantification of uncertainty in computational fluid dynamics*, Annu. Rev. Fluid Mech. **29** (1997) 123, doi:10.1146/annurev.fluid.29.1.123.

16 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** (2008) 078001, doi:10.1115/1.2960953.

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

# The Code-by-Code Verification Backbone: Role, Version, and Reference-Benchmark Provenance for Every Solver in a Multi-Product Fusion Design Campaign

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

Zenodo · doi:10.5281/zenodo.22645866

Read on kronosfusionenergy.com

The Physics De-Risking Register