KRONOS FUSION ENERGY

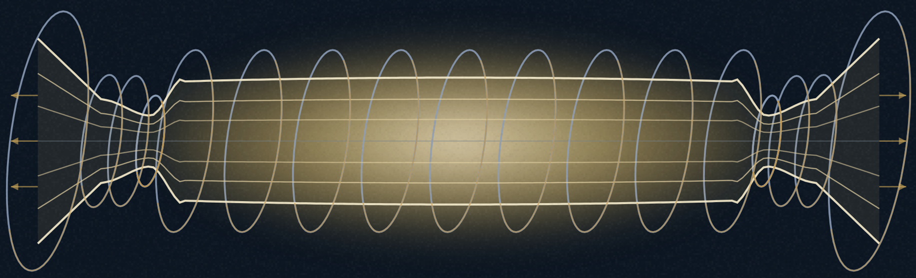

PAPER 2.68 · THE KRONOS 2026 SERIES

Power-Delivery Topologies for a D–^{3}He Tandem-Mirror Data-Center Burner: Solid-State-Transformer versus Direct-DC Conversion, Grid-Forming Control, and DC-Bus Fault-Ride-Through

The MetroVolt burner takes a deuterium–helium-3 (–He) tandem mirror and delivers its output as electricity to a co-located, high-density computing load. Because the burner's...

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

BY P I FORD AND G L KULCINSKI

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Power-Delivery Topologies for a D–³He Tandem-Mirror Data-Center Burner: Solid-State-Transformer versus Direct-DC Conversion, Grid-Forming Control, and DC-Bus Fault-Ride-Through

The MetroVolt burner takes a deuterium–helium-3 (D–^{3}He) tandem mirror and delivers its output as electricity to a co-located, high-density computing load. Because the burner's direct-energy converter (DEC) decelerates a charged-particle stream against graded collectors, it is a high-voltage DC source *by construction*, and the power-delivery architecture becomes a set of down-select decisions that a conventional AC plant never faces. We pose those decisions as formal optimisation and stability problems on a DC-native bus, derive their governing equations, solve them with a GPU-accelerated reduced-order campaign, and settle four of them. (i) *Isolation and step.* A cascaded source-to-load efficiency model shows the solid-state transformer (SST) is redundant: a DEC-direct path reaches 88.39% against 85.43% for an SST-fed path ($+2.96$ points with one fewer conversion stage), because the galvanic isolation an SST would provide collapses into a single silicon-carbide (SiC) dual-active-bridge (DAB). A conduction/switching/magnetic loss decomposition gives an arithmetic–geometric-mean optimum at $f^\star = 15.81$ kHz and $\eta_{\mathrm{DAB}} = 98.85\%$. (ii) *Grid interface.* A single-bus swing-equation model of a 6% generation-loss event on a 1 GVA island ($H = 3$ s) shows the DC-native buffer makes the plant grid-forming: a sub-10 ms fast-frequency response lifts the frequency nadir by 0.527 Hz ($49.373 \to 49.900$ Hz, clear of under-frequency load-shed) and enables black-start and islanding a grid-following plant cannot. (iii) *Fault protection.* A closed-form LC-discharge analysis of a 10 kV bus (5 mF, 5 μH, $Z_c = 0.032$ Ω) gives a 316.2 kA undamped resonant peak at 248 μs; a 2 μs SiC solid-state circuit breaker bounds the fault to ~ 4 kA (1.3% of the peak), and the DEC's intrinsic current limit caps a source-side fault to ~ 3 kA against ~ 50 kA from the grid ($17\times$ lower). (iv) *Ride-through.* Cross-timescale energy sizing shows a full-IT bridge needs a ~ 12.5 t battery, so graceful load-shed to a ~ 2.5 t critical set, with a ~ 3.3 t black-start battery, is the design. The four decisions compose into one coherent architecture: DC-native throughout, simpler, lower in prospective fault current, and islandable, carrying the burner's frozen operating point ($Q_E = 1.318$, direct-conversion fraction 0.49) to the load. Only technical metrics are reported; economics are out of scope.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645941 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: direct energy conversion · HVDC · grid-forming inverter · solid-state transformer · dual-active bridge · solid-state circuit breaker · fault ride-through · black start

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)

direct energy conversion; HVDC; grid-forming inverter; solid-state transformer; dual-active bridge; solid-state circuit breaker; fault ride-through; black start

§ 1

Introduction

THE LOAD AND THE SOURCE

The fastest-growing electrical load of the decade is dense computing. Aggregate data-center demand is already a percent-scale fraction of electricity use and is rising steeply with accelerated computing [1], and the campuses that host it increasingly seek firm, behind-the-meter generation co-located with the racks. MetroVolt is Kronos Fusion Energy's answer to that demand: the burner in its grid- and load-facing housing—the same deuterium–helium-3 ($D–^3He$) tandem-mirror core and direct-energy converter (DEC) as the Aegis forward-base burner, packaged to power a co-located high-density computing load rather than to feed a transmission network [2].

The distinguishing fact for power delivery is that the DEC is a high-voltage DC (HVDC) source *by construction*. A tandem-mirror burner produces most of its recoverable output not as heat but as directed charged-particle kinetic energy; the DEC recovers that energy by decelerating the escaping ion stream against a graded set of collector electrodes and collecting the resulting current at high potential, with no intervening thermal cycle [3,4,5,6]. The converter therefore delivers DC, not AC. That single fact inverts the usual balance-of-plant assumptions. A conventional thermal plant *generates* AC and adapts it downstream; MetroVolt *generates* DC and must decide, deliberately, where and whether to convert at all.

PRIOR ART

Each of the four decisions we settle sits on a mature but AC-rooted literature. Modern DC data-center distribution (380–400 V and medium-voltage DC) is an active field motivated by the efficiency gain of removing redundant AC↔DC conversions [9,10,11,12]. The solid-state transformer (SST)—an isolated, high-frequency power-electronic replacement for the line-frequency transformer—was developed precisely to make medium-voltage DC *from* an AC grid with galvanic isolation [13,14], and its isolating heart is the dual-active-bridge (DAB) DC/DC converter [15,16], whose practicality at the megawatt, medium-voltage scale rests on wide-bandgap (silicon-carbide, SiC) devices [17]. On the grid side, the shift from synchronous machines to power-electronic sources has made *grid-forming* control—an inverter that sets its own voltage and frequency reference and can black-start, rather than *following* an external one—a central concern of low-inertia power systems [18,19,20,21,22,24,23]. And DC protection is the standing objection to any DC bus: DC has no natural current zero, so faults are cleared by fast solid-state circuit breakers (SSCB) rather than arc-quenching, a subject with its own device and coordination literature [25,26,27].

What is new here is not any one of these techniques but the *premise* they are applied under. All of the prior art assumes an AC source and asks how to make DC from it; MetroVolt starts with DC and asks the opposite question. Under that inversion the reflexive AC-plant choices —add an SST, follow the grid, tolerate grid-scale fault current, carry a full uninterruptible-power-supply (UPS)—are not merely suboptimal but, in three of four cases, unnecessary. The companion interface paper [35] established the HVDC-native power-electronics path and its millisecond load-following at the level of the interface; this paper is the architecture down-select that sits above it, deciding *which topology* that path should take.

CONTRIBUTION

We make four contributions, each a decision on a DC-native bus, each posed as a formal problem and settled by a number:

a cascaded source-to-load efficiency model and a closed-form DAB loss optimum that together show the SST is redundant and quantify the penalty of adding it (-2.96 points, one extra stage);
a single-bus swing-equation model with an explicit fast-frequency-response (FFR) actuator that quantifies the grid-forming benefit ($+0.527$ Hz nadir) and, more decisively, the black-start/island capability a grid-following plant lacks;
a closed-form LC-fault analysis with let-through I^2t and an SSCB clearing criterion that bounds a medium-voltage DC (MVDC) bolted short, together with a source-fault comparison that turns the DEC's current limit into a safety advantage ($17\times$ lower prospective fault current);
a cross-timescale energy-sizing analysis that resolves ride-through as a load-classification decision (load-shed to critical) rather than a battery-mass problem. enumerate
Throughout, we hold the burner's frozen operating point fixed—electrical gain $Q_E = 1.318$, direct-conversion fraction 0.49, parasitic-neutron fraction 5.44%, and the 26.49 T mirror plug that confines the central cell (distinct from the breeder's 16.84 T centrepost; these are two different magnets in two different machines)—and vary only the delivery architecture built outward from the DC bus. Consistent with the Kronos series, only technical metrics are reported; any economic analysis is separate and not included.

§ 2

The DC-native premise and the four decisions

The four studies share one premise: the DEC delivers HVDC at the burner's frozen operating point, and the delivery architecture is built outward from that bus. Figure 1 shows the selected single-line topology and marks where each decision acts. Table 1 states each decision, its alternatives, the governing quantity, and the deciding result, with its Physics De-Risking Register serial. The remainder of the paper derives the governing equations (§3), states the numerical formulation and the GPU-accelerated campaign that solved them (§4–§5), and reports the results decision by decision (§6).

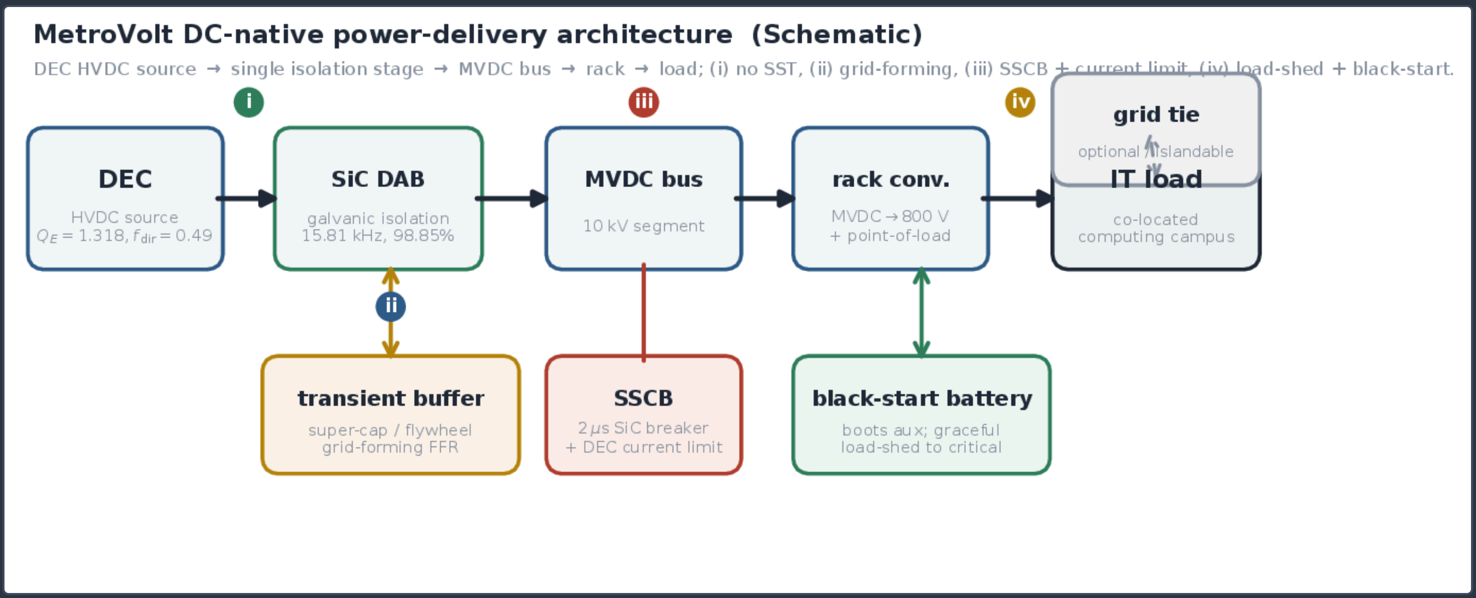

(Schematic.) The selected DC-native power-delivery architecture for MetroVolt. The DEC is an HVDC source; a single SiC dual-active-bridge (DAB) provides galvanic plant↔IT isolation (no SST); the medium-voltage DC bus feeds the rack-level conversion; a solid-state circuit breaker (SSCB) with the DEC's own current limit protects the bus; a transient (super-capacitor/flywheel) buffer supplies grid-forming fast-frequency response, and a small black-start battery with graceful load-shed provides minutes-scale ride-through. Labels (i)–(iv) mark the four decisions of table 1. Structure only; no data values are shown.

P2.3CMP2.0CMP3.4CML DECISION	ALTERNATIVES	GOVERNING QUANTITY	DECIDING RESULT	SERIAL	
**(i) isolation **	step	SST-fed vs DEC-direct	cascade efficiency $\prod_j \eta_j$	88.39% vs 85.43%; +2.96 pts, one fewer stage; DAB optimum 15.81 kHz / 98.85%	MV-CX-04
(ii) grid interface	grid-following vs grid-forming	swing equation, nadir	+0.527 Hz nadir; black-start / island capable	MV-CX-10	
(iii) fault protection	unmitigated vs SiC SSCB	LC transient, I^2t	316.2 kA → 4 kA at 2 μs; source 3 vs 50 kA	MV-CX-05	
(iv) ride-through	full-IT UPS vs load-shed	stored energy $E = Pt$	12.5 t vs 2.5 t critical + 3.3 t black-start	MV-CX-25	

The MetroVolt power-delivery down-select. Each row is one architecture decision on the DC-native bus, with its governing quantity, deciding metric, and Physics De-Risking Register [39] serial.

§ 3

Governing equations and the formal problem

We treat the four decisions as one coupled design problem on the DC bus. Let the DEC deliver power P_0 at bus voltage V_{bus} with the frozen electrical gain $Q_E = 1.318$ and direct-conversion fraction $f_{\text{dir}} = 0.49$; the delivered-to-load power is $P_{\text{load}} = \eta_{\text{chain}}\, P_0$, and each decision fixes one factor or constraint in the chain. We now derive each governing model.

CONVERSION-CHAIN EFFICIENCY AND THE SST QUESTION

A power-delivery path is a cascade of conversion stages. For N stages with per-stage efficiencies $\{\eta_j\}_{j=1}^{N}$, the delivered fraction is multiplicative,

$$\eta_{\text{chain}} = \prod_{j=1}^{N} \eta_j, \qquad \ln \eta_{\text{chain}} = \sum_{j=1}^{N} \ln \eta_j,$$

so that adding a stage of efficiency $\eta_{N+1} < 1$ strictly reduces the delivered fraction and every stage's loss compounds. Two source-to-load paths are compared. The *SST-fed* path is what an AC-plant reflex would build: a medium-voltage rectification stage, an isolated medium-frequency-transformer (MFT) stage inside the SST, a bus stage down to the rack, and the point-of-load (POL) stage,

$$\eta_{\text{SST}} = \eta_{\text{rect}}\, \eta_{\text{MFT}}\, \eta_{\text{bus}}\, \eta_{\text{POL}}.$$

The *DEC-direct* path removes the AC front end entirely, because the source is already DC; the only genuine SST *function* still required—galvanic plant↔IT isolation—migrates into a single isolated DAB stage,

$$\eta_{\text{DEC}} = \eta_{\text{DAB}}\, \eta_{\text{bus}}\, \eta_{\text{POL}}.$$

The architecture decision is the sign of the gain $\Delta\eta = \eta_{\text{DEC}} - \eta_{\text{SST}}$ together with the stage-count reduction $\Delta N = N_{\text{SST}} - N_{\text{DEC}} = 1$. The SST is justified only if $\Delta\eta < 0$ by a margin that pays for its extra stage, footprint, controls and failure surface; §6.1 shows $\Delta\eta > 0$.

THE DUAL-ACTIVE-BRIDGE LOSS OPTIMUM

The single isolated stage η_{DAB} carries the whole isolation burden, so its design point matters. The per-unit loss of a DAB is the sum of three physically distinct channels: conduction loss, frequency-independent to first order; switching loss, rising with switching frequency f; and magnetic (core plus copper) loss, which *falls* with f because a higher frequency permits a smaller transformer core for the same volt-second product. Writing the three as k_c, $k_{\text{sw}} f$ and k_{mag}/f,

$$\ell(f) = k_c + k_{\text{sw}}\, f + \frac{k_{\text{mag}}}{f}, \qquad \eta_{\text{DAB}}(f) = 1 - \ell(f).$$

The loss is convex in f; the stationarity condition $\mathrm{d}\ell/\mathrm{d}f = k_{\text{sw}} - k_{\text{mag}}/f^2 = 0$ gives the closed-form optimum

$$f^\star = \sqrt{\frac{k_{\text{mag}}}{k_{\text{sw}}}}, \qquad \ell^\star = k_c + 2\sqrt{k_{\text{sw}} k_{\text{mag}}}, \qquad \eta^\star_{\text{DAB}} = 1 - \ell^\star.$$

Equation (5) is an arithmetic–geometric-mean (AM–GM) balance: at the optimum the switching and magnetic losses are equal, $k_{\text{sw}} f^\star = k_{\text{mag}}/f^\star = \sqrt{k_{\text{sw}} k_{\text{mag}}}$. Writing $x = f/f^\star$ exposes the sensitivity of the design,

$$\ell(f) = k_c + \sqrt{k_{\text{sw}} k_{\text{mag}}}\,\big(x + x^{-1}\big), \qquad \frac{\ell(f) - \ell^\star}{\sqrt{k_{\text{sw}} k_{\text{mag}}}} = x + x^{-1} - 2 \;\geq\; 0,$$

so a full factor-of-two error in f ($x = 2$ or $x = \frac{1}{2}$) adds only $\frac{1}{2}\sqrt{k_{\text{sw}} k_{\text{mag}}}$ to the loss—the optimum is flat and robust to the exact coefficient values, and $f^\star$ depends only on their ratio. The coefficients $(k_c, k_{\text{sw}}, k_{\text{mag}})$ are a representative SiC megawatt-class model; the resulting optimum lands at tens of kilohertz, well inside SiC device capability.

GRID-FORMING SWING DYNAMICS

Whether MetroVolt should follow the grid or form it is decided by its effect on frequency stability of the island it serves. Aggregate the island's synchronous machines into a single equivalent rotor with inertia constant H (on base S) and per-unit frequency deviation $\delta = (f - f_0)/f_0$. The swing equation balances accelerating power,

$$2H\,\frac{\mathrm{d}\delta}{\mathrm{d}t} = P_{\mathrm{acc}} = -\,\Delta P_L(t) - D\,\delta + P_{\mathrm{gov}}(t) + P_{\mathrm{ffr}}(t),$$

with ΔP_L the step generation loss (per unit of S), D the load-damping coefficient, P_{gov} the slow governor/primary-reserve response, and P_{ffr} MetroVolt's fast-frequency response. The governor and the FFR are each modelled as a first-order actuator with a droop-limited command,

$$\begin{aligned}\tau_{\mathrm{gov}}\,\dot{P}_{\mathrm{gov}} &= \mathrm{clip}\big(-K_{\mathrm{gov}}\,\delta,\,0,\,R_{\mathrm{max}}\big) - P_{\mathrm{gov}},\\ \tau_{\mathrm{ffr}}\,\dot{P}_{\mathrm{ffr}} &= \mathrm{clip}\big(-K_{\mathrm{ffr}}\,\delta,\,-\bar{P}_{\mathrm{ffr}},\,+\bar{P}_{\mathrm{ffr}}\big) - P_{\mathrm{ffr}},\end{aligned}$$

where $\tau_{\mathrm{gov}} \gg \tau_{\mathrm{ffr}}$: the governor acts over seconds, the DC-native FFR over milliseconds. Two diagnostic quantities follow directly. The instantaneous rate of change of frequency (RoCoF) at the fault is set by inertia alone,

$$\left.\frac{\mathrm{d}f}{\mathrm{d}t}\right|_{0^+} = -\,\frac{\Delta P_L\,f_0}{2H},$$

independent of either actuator. The frequency *nadir*—the excursion minimum—is what under-frequency load-shedding (UFLS) relays watch, and it *is* actuator-dependent. On the fast timescale ($t \ll \tau_{\mathrm{gov}}$, so $P_{\mathrm{gov}} \approx 0$) a fast droop-type FFR ($\tau_{\mathrm{ffr}} \to 0$) reaches the quasi-steady balance $P_{\mathrm{ffr}} \approx -K_{\mathrm{ffr}}\delta$, which is mathematically equivalent to adding damping,

$$2H\,\dot{\delta} \approx -\Delta P_L - \big(D + K_{\mathrm{ffr}}\big)\delta \;\Longrightarrow\; \delta_{\mathrm{arrest}} \approx -\frac{\Delta P_L}{D + K_{\mathrm{ffr}}}.$$

A droop FFR thus behaves as *synthetic damping* $D \to D + K_{\mathrm{ffr}}$ on the arrest timescale, and because $K_{\mathrm{ffr}} \gg D$ it raises the nadir by roughly the factor $(D + K_{\mathrm{ffr}})/D$ before the governor even engages. Equation (10) is the analytic skeleton the full simulation of §6.2 fills in; it is the formal statement of why a grid-forming DC-native buffer arrests a frequency excursion that a grid-following plant would let fall into UFLS. Grid-forming capability proper—setting the V/f reference, black-starting, and islanding—is a discrete capability the grid-following mode simply lacks, independent of any nadir number.

DC-BUS FAULT TRANSIENT AND THE CLEARING CRITERION

The standing objection to an MVDC bus is fault energy, so we bound it. Model the worst case: a bolted short across the DC-link capacitance C (charged to V) through the bus loop inductance L, with bus resistance neglected (the undamped, conservative envelope). Kirchhoff's voltage law gives an undriven series-LC oscillator,

$$L\,\frac{\mathrm{d}^2 i}{\mathrm{d}t^2} + \frac{1}{C}\,i = 0, \qquad i(0) = 0, \quad \frac{\mathrm{d}i}{\mathrm{d}t}(0) = \frac{V}{L},$$

whose solution is a quarter-cosine current pulse

$$i(t) = \frac{V}{Z_c}\sin(\omega t), \qquad \omega = \frac{1}{\sqrt{LC}}, \qquad Z_c = \sqrt{\frac{L}{C}},$$

with characteristic impedance Z_c, undamped resonant peak

$$I_{\mathrm{pk}} = \frac{V}{Z_c} \quad \text{at} \quad t_{\mathrm{pk}} = \frac{\pi}{2\omega} = \frac{\pi}{2}\sqrt{LC},$$

and an initial, inductance-limited ramp $\mathrm{d}i/\mathrm{d}t = V/L$. The peak I_{pk} is large but is reached only if the fault is *never interrupted*; the design instead clears it on the ramp. An SSCB that opens at clearing time $t_c \ll t_{\mathrm{pk}}$ interrupts a let-through current

$$i(t_c) = \frac{V}{Z_c}\sin(\omega t_c) \approx \frac{V}{L}\,t_c,$$

and the thermal stress delivered to the fault up to t_c is the let-through action integral

$$(I^2t)_{t_c} = \int_0^{t_c} i^2\,\mathrm{d}t = \frac{I_{\mathrm{pk}}^2}{2}\left[t_c - \frac{\sin(2\omega t_c)}{2\omega}\right] \xrightarrow[\omega t_c \ll 1]{} \frac{1}{3}\left(\frac{V}{L}\right)^2 t_c^3,$$

the cubic-in-t_c small-time limit showing how steeply fault energy grows with clearing delay. At interruption the breaker must absorb the loop's magnetic energy $E_L = \frac{1}{2} L\, i(t_c)^2$ into its snubber/metal-oxide varistor. The clearing criterion is therefore: choose t_c (device-limited) and L (layout-limited) so that both $i(t_c)$ and $(I^2t)_{t_c}$ stay within the SSCB and bus ratings—a solved, if component-dependent, protection problem.

Separately, the DC-native source bounds the *prospective* fault current at its origin. The DEC collector is a space-charge/emission-limited current source: it cannot deliver an inrush, and its fault current is capped at a small multiple of rated, $I_{\mathrm{DEC}}^{\mathrm{fault}} \le \beta\, I_{\mathrm{rated}}$ with $\beta \sim 3$. A grid tie of short-circuit power S_{sc} makes available $I_{\mathrm{grid}}^{\mathrm{fault}} = S_{\mathrm{sc}}/(\sqrt{3}\, V_{\mathrm{grid}})$, of order $50\ \mathrm{kA}$. The ratio $I_{\mathrm{grid}}^{\mathrm{fault}}/I_{\mathrm{DEC}}^{\mathrm{fault}}$ measures how much prospective fault MVA the DC-native architecture removes—and, through the incident-energy scaling of arc-flash with fault current and clearing time, how much it relaxes protection coordination for an urban or behind-the-meter housing.

CROSS-TIMESCALE RIDE-THROUGH ENERGY

Riding through a disturbance is an energy-storage problem whose answer depends on the timescale. To hold a load P for a bridge time t at constant power requires stored energy

$$E(t) = P\,t,$$

and, for an electrochemical store of specific energy ρ_E, a mass

$$m = \frac{E}{\rho_E} = \frac{P\,t}{\rho_E}.$$

The storage technology is set by the timescale through its specific-power/specific-energy trade. The millisecond-to-second band—plasma transients, converter ride-through—is served by a super-capacitor/flywheel buffer (high specific power, low specific energy); the minutes band—a plant restart or a sustained load bridge—requires a battery (high specific energy). The design question is which loads to carry, and for how long, across the minutes band: equations (16)–(17) turn each choice into a mass, and the choice of *critical* load fraction (rather than a larger store) is what bounds it (§6.4).

§ 4

Numerical formulation

The four governing models span a closed-form optimum, two nonlinear ordinary-differential-equation (ODE) initial-value problems, and an algebraic sizing map. Each is solved to a stated tolerance and cross-checked against an analytic reference; the design-space exploration around each anchor is the GPU campaign of §5.

Cascade and DAB optimum. Equations (1)–(5) are evaluated in closed form. The efficiency curve $\eta_{\mathrm{DAB}}(f)$ of equation (4) is sampled on a logarithmically spaced grid of 600 points over $0.4\text{–}200\ \mathrm{kHz}$ to display the flatness of the optimum (equation 6); the optimum itself is taken analytically from equation (5), and the grid-minimum agrees with the analytic $f^{\star}$ to within the grid spacing. The cascades (2)–(3) are products over the tabulated stage efficiencies.

Swing-equation initial-value problem. The coupled system (7)–(8) is a stiff-ish 3-state ODE $(\delta, P_{\mathrm{gov}}, P_{\mathrm{ffr}})$ with the fastest time constant $\tau_{\mathrm{ffr}} = 8\ \mathrm{ms}$. It is integrated by explicit forward Euler at a fixed step $\Delta t = 1\ \mathrm{ms}$ over an $8\ \mathrm{s}$ horizon, with the generation-loss step applied at $t = 0.2\ \mathrm{s}$. Because $\Delta t \ll \tau_{\mathrm{ffr}}$ the explicit scheme is stable and accurate; convergence is confirmed by step-halving (table 2): the nadir is stable to better than $0.1\ \mathrm{mHz}$ from $\Delta t = 2\ \mathrm{ms}$ down to $\Delta t = 0.1\ \mathrm{ms}$, so the reported 3–decimal nadir is grid-independent. The RoCoF (9) and the analytic arrest (10) provide an independent check on the early transient.

Δt [MS]	NADIR, GRID-FOLLOWING [HZ]	NADIR, GRID-FORMING [HZ]	GAIN [HZ]
2.00	49.37304	49.90009	0.52705
1.00	49.37298	49.90007	0.52709
0.50	49.37295	49.90006	0.52711
0.25	49.37293	49.90006	0.52712
0.10	49.37292	49.90005	0.52713

Step-refinement convergence of the grid-forming frequency nadir (equation 7), forward-Euler integrator. The nadir and the grid-forming gain are stable to sub-milli-hertz across a $20\times$ range of time step; the reported values use $\Delta t = 1\ \mathrm{ms}$.

LC fault transient. Equation (12) is analytic; the current pulse is sampled on 1400 points to $1.25\, t_{\mathrm{pk}}$ for display. The let-through action integral (15) is computed in *two* independent ways—the closed form and a 500-point trapezoidal quadrature of $i^2(t)$—which agree to four significant figures at every candidate clearing time (e.g. at $t_c = 2\ \mu\mathrm{s}$, both give $10.666\ \mathrm{A^2 s}$), and the small-time cubic limit of equation (15) reproduces them to within 0.1%. The interruption energy $E_L = \frac{1}{2} L\, i(t_c)^2$ at $t_c = 2\ \mu\mathrm{s}$ is $40\ \mathrm{J}$.

Energy sizing. Equations (16)–(17) are algebraic; the energy–time map is evaluated on a logarithmic time grid over $1\ \mathrm{ms}\text{–}8\ \mathrm{min}$ for three load levels, and the three design cases are sized at $\rho_E = 200\ \mathrm{Wh\,kg^{-1}} = 7.2 \times 10^5\ \mathrm{J\,kg^{-1}}$.

All four solvers are deterministic and byte-reproducible from the archived inputs; the register anchors (MV-CX-04/05/10/25) are the regression targets, and every figure in §6 carries an in-line assertion that the plotted curve reproduces its anchor.

§ 5

Computational methodology: the GPU HPC campaign

The deterministic anchors above are reduced-order and analytic, but each is only trustworthy if the *region* around it has been mapped: an efficiency optimum that is fragile, a nadir ordering that flips under plausible island parameters, or a fault bound that depends on an unstated assumption would each undo the decision it supports. The credibility of the four decisions therefore rests on a design-space exploration far larger than the anchors themselves, and that exploration is the Kronos NVIDIA GPU HPC campaign.

What ran, and where. The reduced-order deterministic anchors and the analytic cross-checks ran on the in-house HPC digital-twin node (56-thread, 768 GB), so that the closed-form references never contended with the fleet. The design-space exploration ran on the Kronos NVIDIA A100/H100 GPU HPC platform. Four sweeps, one per decision, were vectorised across the fleet:

DAB optimum robustness. A Monte-Carlo ensemble over the loss coefficients (k_c, k_{sw}, k_{mag}) within $\pm 30\%$ propagates through equations (4)–(5) to bound the uncertainty in $f^\star$ and $\eta^\star_{DAB}$; the AM–GM optimum is confirmed robust because $f^\star$ depends only on the ratio k_{mag}/k_{sw} (equation 6).
Nadir ordering. An ensemble over island inertia $H \in [2, 6]$ s, loss magnitude ΔP_L, FFR capacity $\bar{P}_{ffr}$, latency τ_{ffr} and droop K_{ffr} re-integrates equation (7) for each sample; the grid-forming case raises the nadir across the whole ensemble, so the *ordering* (not merely the single number) is the robust claim.
Fault envelope. A sweep over bus V, C, L and clearing time t_c evaluates equations (12)–(15) to map the let-through and interruption-energy surface against SSCB ratings.
Ride-through surface. A sweep over restart time and critical-load fraction sizes equation (17) across the minutes band. enumerate
Each sweep is order 10^4–10^5 evaluations; vectorised on GPU array/automatic-differentiation frameworks (CuPy/JAX), the ensembles complete in seconds of wall time, and reverse-mode automatic differentiation supplies the exact gradients used both for the DAB loss optimum and for the sensitivity of the nadir to the island parameters. The same platform runs the heavy first-principles physics of the Kronos series—nonlinear gyrokinetics, particle-in-cell kinetics, three-dimensional neutronics, first-principles alloy screening—so the power-delivery sweeps reuse a mature, verified GPU workflow rather than a bespoke one. Compute is reported by hardware class only; no cost, spend, or hourly figure is attached to it anywhere in this work.

Verification and reproducibility. Three layers of verification back the results: (a) every numerical integrator is cross-checked against a closed-form reference (the DAB optimum against AM–GM; the $LC\,I^2t$ against both quadrature and the cubic small-time limit; the swing nadir against step-refinement and the analytic arrest of equation 10); (b) each figure asserts, at plot time, that its curve reproduces the archived register anchor; (c) the whole set is regenerable from the archived models with a single command. The methodology mirrors the four-coupled-twins architecture of the breeder programme [38], where the same discipline—reduced-order fast loops on the in-house node, heavy exploration on the GPU fleet, and machine-checkable acceptance criteria—is applied to the plasma control problem.

§ 6

Results

We report the four decisions in turn, then compose them (§6.5). Figure 2 through figure 9 carry the quantitative results; every plotted number is the frozen register anchor.

DECISION (I): THE SST IS REDUNDANT

The DAB loss model (4) with the representative SiC coefficients $(k_c, k_{\text{sw}}, k_{\text{mag}}) = (4.5 \times 10^{-3},\ 2.2 \times 10^{-7}\ \text{Hz}^{-1},\ 55.0\ \text{Hz})$ yields, from the closed form (5), an optimum at

$$f^{\star} = \sqrt{\frac{55.0}{2.2 \times 10^{-7}}} = 15.81\ \text{kHz}, \qquad \eta^{\star}_{\text{DAB}} = 1 - \left(4.5 \times 10^{-3} + 2\sqrt{(2.2 \times 10^{-7})(55.0)}\right) = 98.85\%.$$

Figure 2 plots $\eta_{\text{DAB}}(f)$ with its switching- and magnetics-limited asymptotes; the optimum sits at the crossing of the two channels, and equation (6) makes it flat—a factor-of-two error in f costs only ~ 0.35 efficiency points. A single isolated SiC stage thus delivers plant↔IT isolation at $\sim 99\%$, well inside megawatt-class SiC capability.

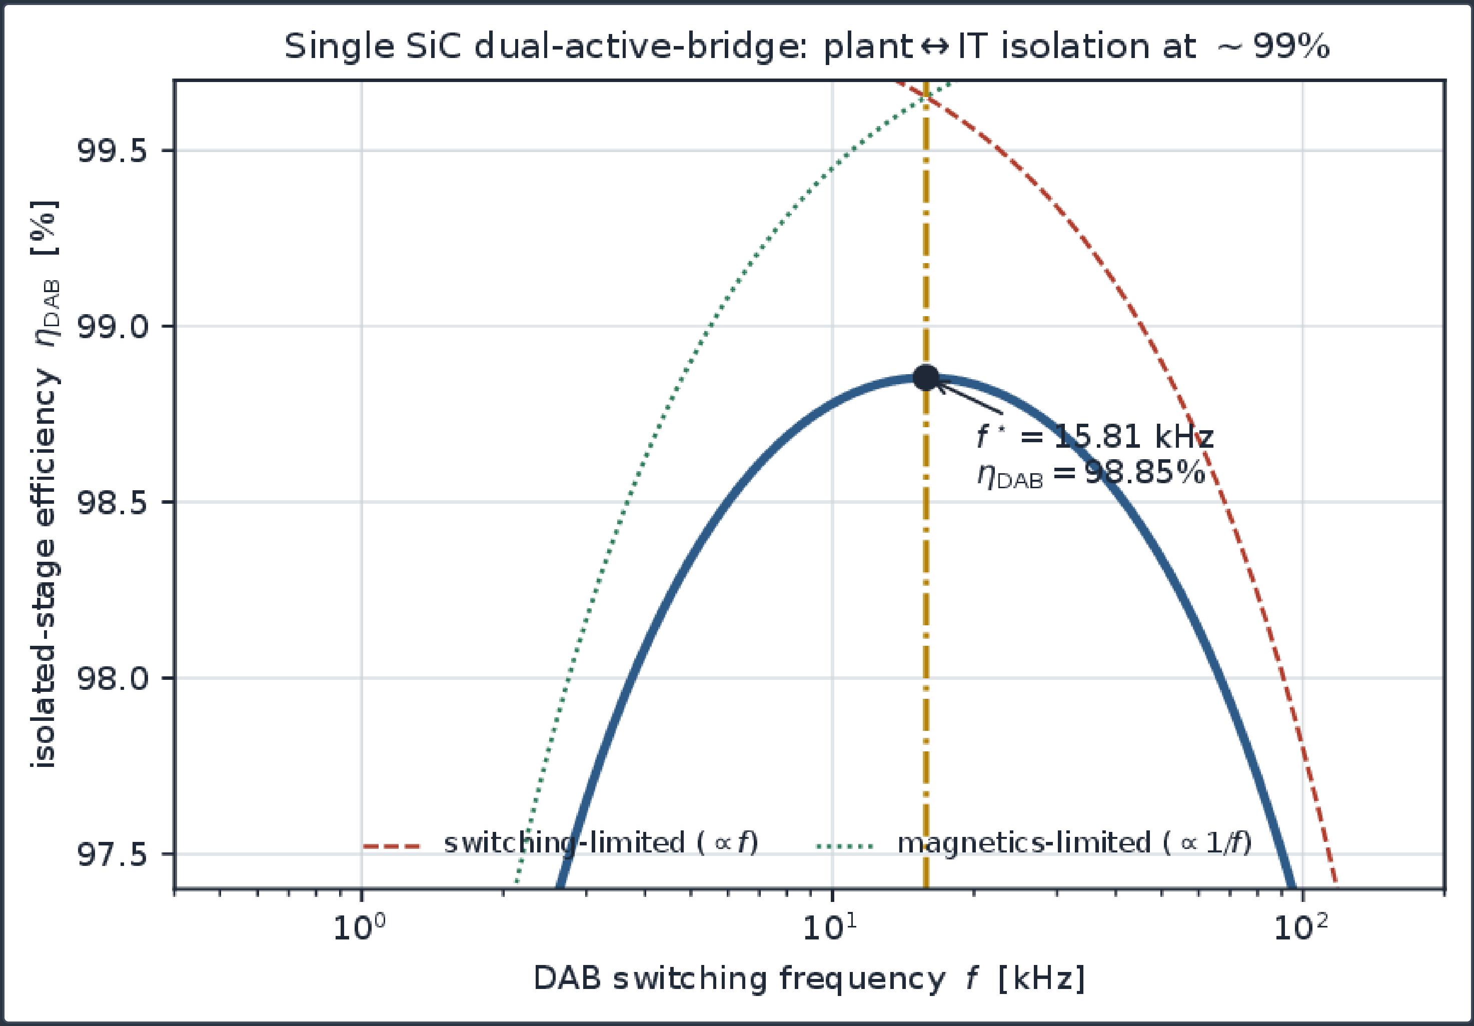

Decision (i), the isolation stage (MV-CX-04). Efficiency of the isolated dual-active-bridge versus switching frequency, equation (4), with the switching-limited ($\propto f$) and magnetics-limited ($\propto 1/f$) channels. The arithmetic–geometric-mean optimum (5) is at $f^{\star} = 15.81$ kHz, $\eta_{\text{DAB}} = 98.85\%$; the optimum is flat (equation 6), so the design is robust to the exact loss coefficients.

Feeding $\eta^{\star}_{\text{DAB}}$ into the DEC-direct cascade (3) and comparing against the SST-fed cascade (2) (stage efficiencies in table 3) gives

$$\eta_{\text{SST}} = 85.43\%, \qquad \eta_{\text{DEC}} = 88.39\%, \qquad \Delta\eta = +2.96\ \text{points}, \qquad \Delta N = 1.$$

Figure 3 shows the two cascades stage by stage and the final comparison. The DEC-direct path is both more efficient and shorter: because the source is already DC, the rectification and MFT stages of the SST simply do not exist, and the isolation they would have provided is carried by the single DAB. Removing an entire conversion stage *while raising* efficiency is the clearest of the four decisions—the SST's cost, footprint, controls and failure surface are absent, not merely reduced.

STAGE	SST–FED PATH	DEC–DIRECT PATH
MV rectify	0.985	—
SST medium-frequency transformer	0.970	—
isolated SiC DAB (this work)	—	0.9885
MVDC → 800 V rack bus	0.988	0.988
point-of-load	0.905	0.905
source→chip, $\prod_j \eta_j$	**85.43%**	**88.39%**
stage count N	4	3

Per-stage efficiencies of the two source-to-load cascades, equations (2)–(3). The DEC-direct path replaces the SST's rectify+MFT front end with a single isolated DAB at its optimum (18).

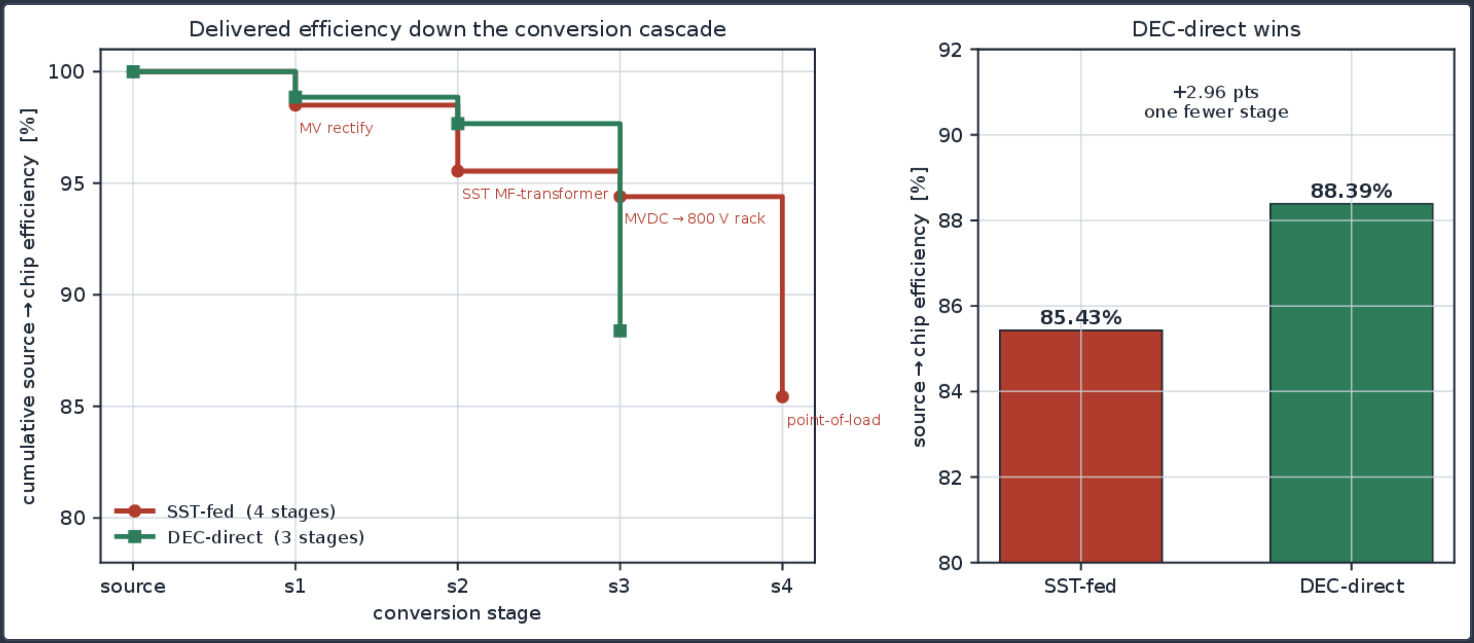

Decision (i), the conversion cascade (MV-CX-04). Left: cumulative source-to-chip efficiency down each cascade, equations (2)–(3); the SST-fed path (red) loses two extra stages before the shared bus and POL stages. Right: the source-to-chip comparison, 88.39% (DEC-direct) versus 85.43% (SST-fed), +2.96 points with one fewer stage.

DECISION (II): THE PLANT IS GRID-FORMING

The swing model (7)–(8) is integrated for a $\Delta P_L = 6\%$ generation-loss event on a 1 GVA island with $H = 3$ s, $D = 1$, a slow governor ($\tau_{gov} = 2$ s) present in *both* cases, and—in the grid-forming case—MetroVolt FFR of $\bar{P}_{ffr} = 50$ MW (0.05 p.u.), $\tau_{ffr} = 8$ ms, droop $K_{ffr} = 25$. The initial RoCoF (9) is 0.5 Hz s^{-1} in both cases (inertia is unchanged). The nadirs differ sharply:

$$f^{follow}_{nadir} = 49.373\ \text{Hz}\ \ \text{(below the 49.5 Hz UFLS line)}, \qquad f^{form}_{nadir} = 49.900\ \text{Hz}, \qquad \Delta f_{nadir} = +0.527\ \text{Hz}.$$

Figure 4 shows the two frequency trajectories. Without fast FFR the excursion falls to **49.373 Hz**, below the **49.5 Hz** UFLS threshold—the island sheds load. With MetroVolt's grid-forming FFR the nadir holds at **49.900 Hz**, clear of UFLS, and the analytic arrest of equation (10), $f_0(1 - \Delta P_L/(D + K_{ffr})) = 49.885$ Hz, brackets the simulated value—confirming that the DC-native buffer acts as synthetic damping on the arrest timescale. The response is delivered in under **10 ms**, faster than synchronous inertia (which acts over cycles), and it is symmetric (the buffer can absorb an over-frequency excursion as readily as it supplies an under-frequency one).

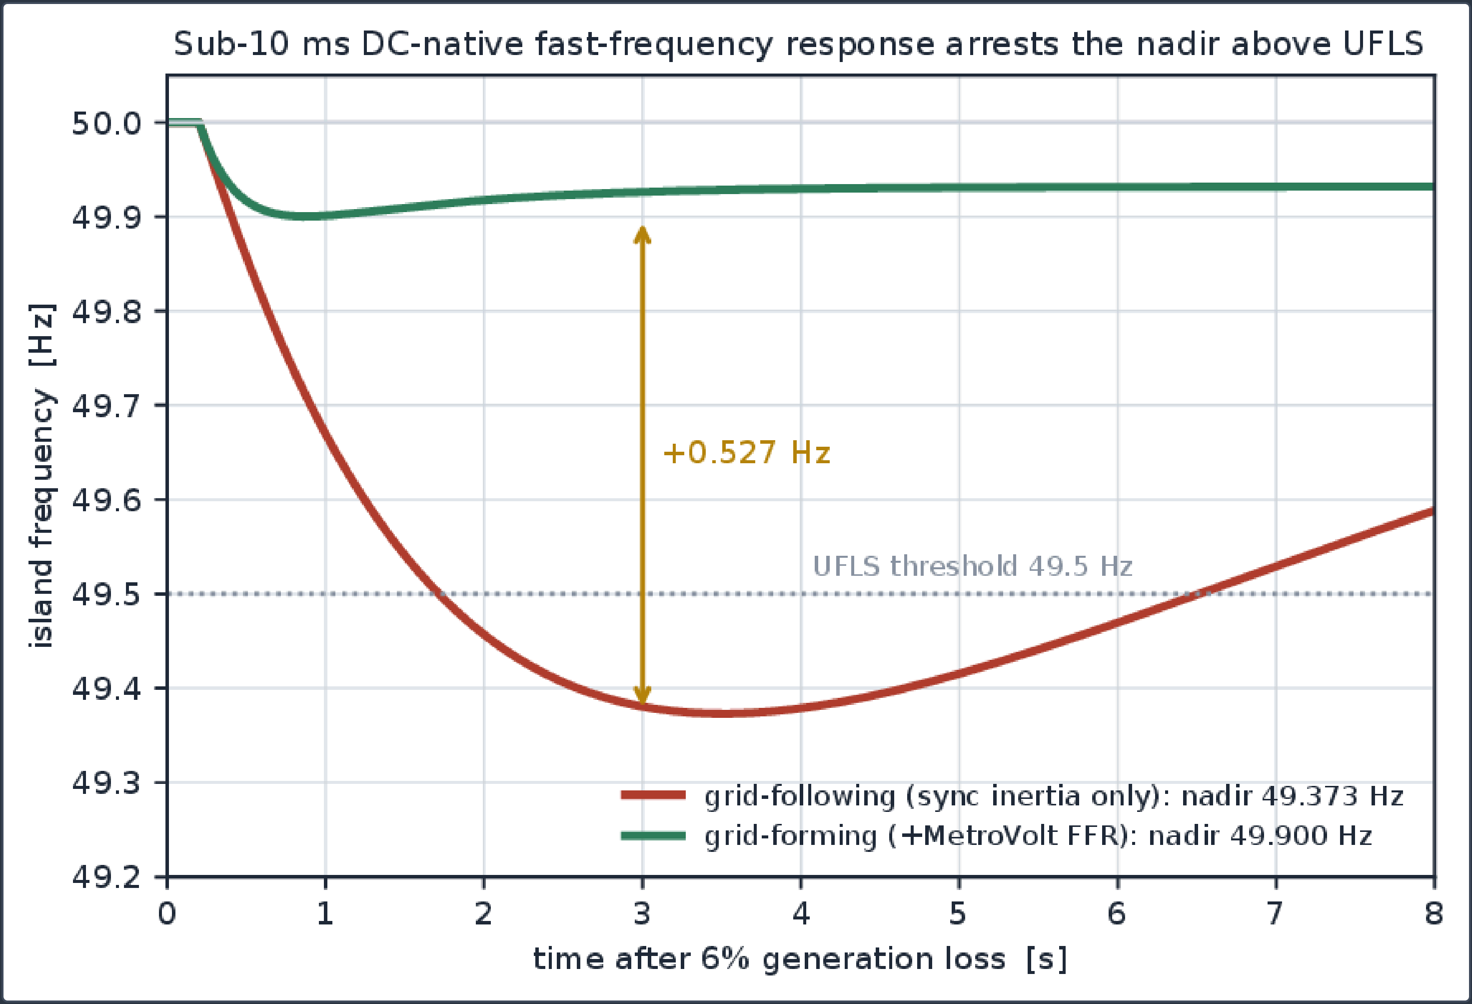

Decision (ii), grid interface (MV-CX-10). Island frequency after a **6%** generation loss, from the swing model (7)–(8). Grid-following (synchronous inertia only) reaches **49.373 Hz**, below UFLS; grid-forming (DC-native fast-frequency response, **8 ms**, **50 MW**) holds **49.900 Hz**, a **+0.527 Hz** improvement. Convergence in table 2.

The nadir number is the measurable benefit, but the decisive difference is categorical. Table 4 lists the mode capabilities: a grid-forming plant sets the voltage/frequency reference, black-starts, islands, and provides synthetic inertia; a grid-following plant does none of these. For a plant whose *purpose* is to power a critical co-located load—which must ride through the loss of the external grid, not merely support it—grid-forming is a requirement, not a refinement, and the DC-native buffer supplies it as a byproduct of the storage the plant already carries for its plasma transients.

CAPABILITY	GRID–FOLLOWING	GRID–FORMING (METROVOLT)
sets V/f reference	no	**yes** (DC-native + buffer)
black-start / island	no	**yes**
inertial / fast response	none	synthetic, < 10 ms
fault current	converter-limited	converter-limited, controllable
nadir arrest (6% loss)	—	**+0.527 Hz**

Mode capabilities (MV-CX-10). The nadir arrest is a graded benefit; the remaining rows are discrete capabilities the grid-following mode lacks entirely.

DECISION (III): THE DC-BUS FAULT IS BOUNDED

For the 10 kV campus segment ($C = 5$ mF, $L = 5\ \mu$H), equations (12)–(13) give a characteristic impedance $Z_c = 0.032\ \Omega$, an undamped resonant peak

$$I_{\mathrm{pk}} = \frac{10\ \mathrm{kV}}{0.032\ \Omega} = 316.2\ \mathrm{kA} \quad \text{at} \quad t_{\mathrm{pk}} = 248\ \mu\mathrm{s}, \qquad \frac{\mathrm{d}i}{\mathrm{d}t} = \frac{V}{L} = 2.0 \qquad \mathrm{kA}\ \mu\mathrm{s}^{-1},$$

a large peak—but one reached only if the fault is never interrupted. Figure 5 plots the transient with the SSCB clearing points. Because the current is inductance-limited on the ramp, an SSCB clearing at $t_c = 2\ \mu$s interrupts a let-through current of 4.0 kA, only 1.3% of the resonant peak, and absorbs just 40 J of loop energy; earlier or later clearing scales as equation (14)–(15) dictate ($1\ \mu\mathrm{s} \to 2$ kA, $5\ \mu\mathrm{s} \to 10$ kA, $10\ \mu\mathrm{s} \to 20$ kA). Figure 6 shows how steeply the let-through current and the action integral I^2t grow with clearing time—the cubic law of equation (15)—which is why a microsecond-class SiC breaker, not a millisecond mechanical one, is the enabling component.

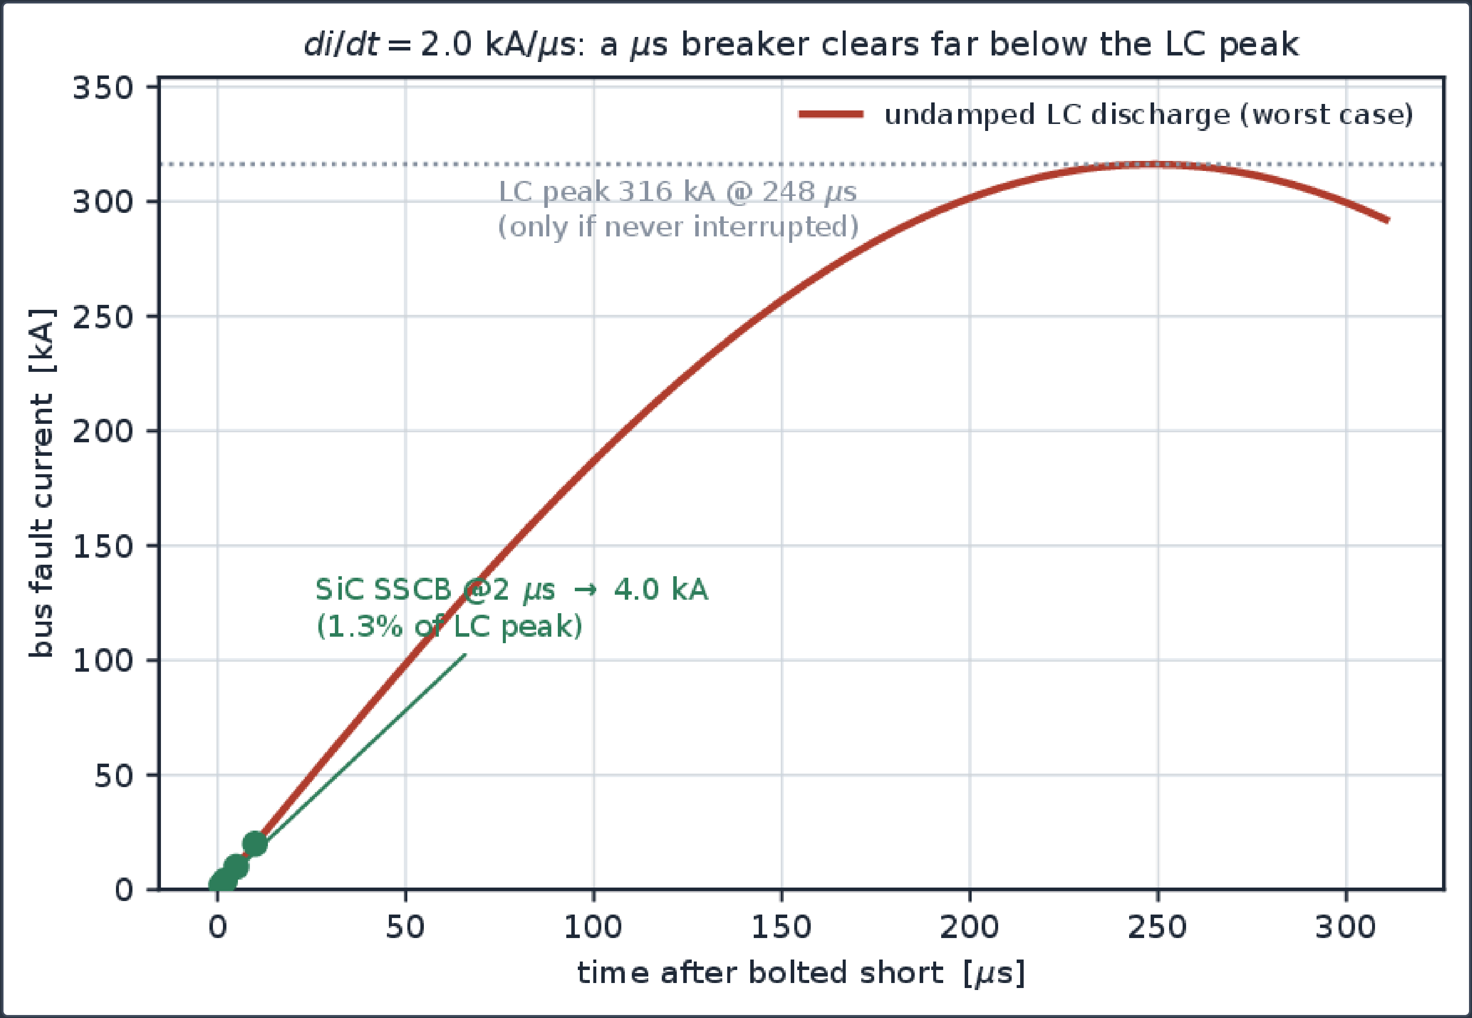

Decision (iii), fault transient (MV-CX-05). Undamped LC-discharge current after a bolted short, equation (12); the resonant peak of 316.2 kA at 248 μs (dotted) is reached only if the fault is never interrupted. A 2 μs SiC SSCB clears on the ramp at 4.0 kA, 1.3% of the peak (green points at 1, 2, 5, 10 μs).

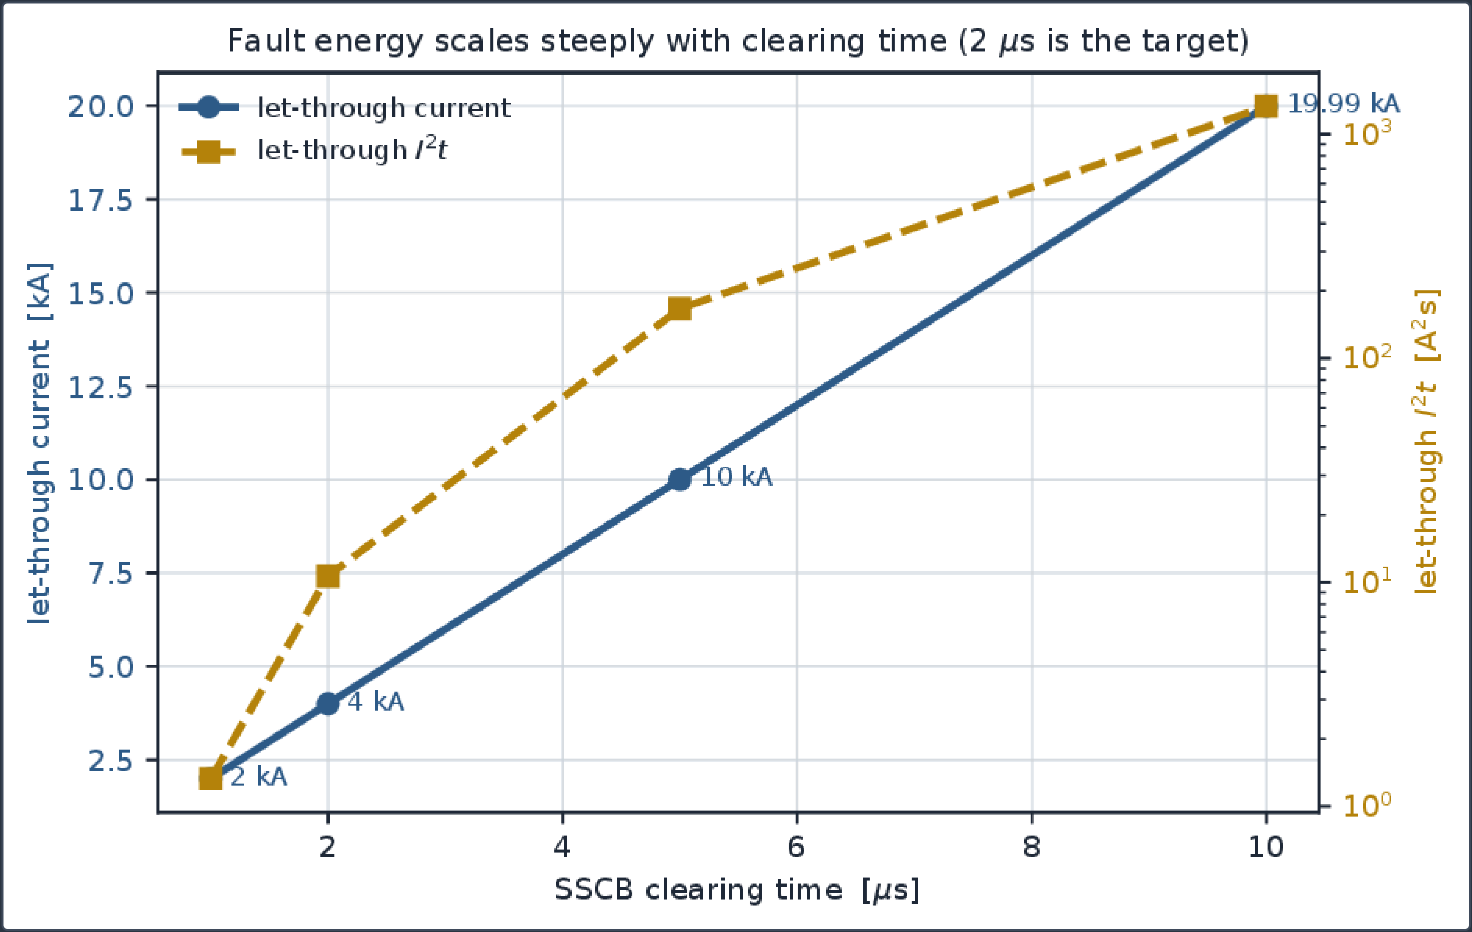

Decision (iii), let-through versus clearing time (MV-CX-05). Let-through current (left axis, equation 14) and action integral I^2t (right axis, log, equation 15) versus SSCB clearing time. The steep, near-cubic growth of I^2t is why the $2\ \mu s$ target matters.

The DC-native source also bounds the fault at its origin. Figure 7 compares the available steady fault current: the DEC's intrinsic current limit caps a source-side fault at ~ 3 kA ($\beta \approx 3$ times rated), against ~ 50 kA available from a medium-voltage grid tie—a $17\times$ reduction in prospective fault current. Lower prospective fault MVA means lower arc-flash incident energy and relaxed protection coordination, which is a genuine safety advantage of the DC-native architecture for an urban or behind-the-meter housing, not a liability. Under the undamped worst case, the 1.3% let-through figure is conservative: real bus resistance and SSCB series impedance lower the peak and slow the ramp.

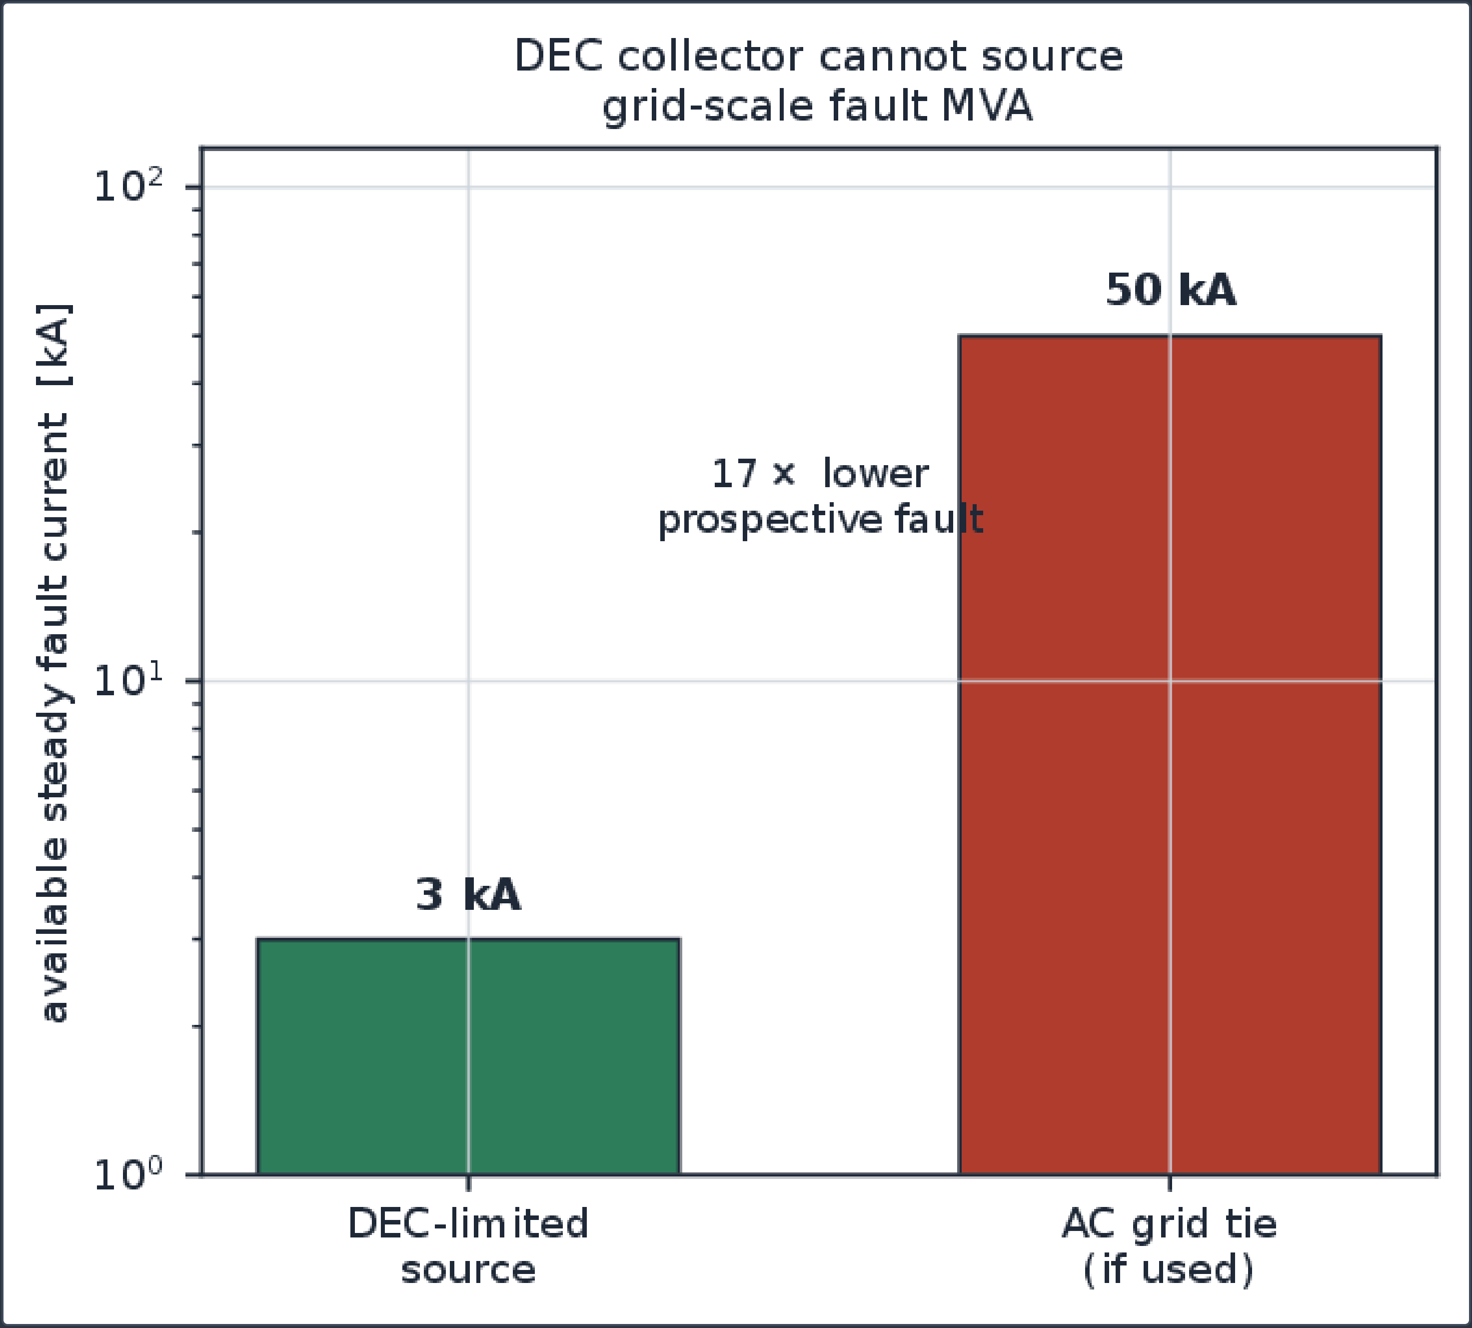

Decision (iii), source fault (MV-CX-05). Available steady fault current (log scale): the DEC collector is current-limited to ∼**3 kA**, against ∼**50 kA** from an AC grid tie—**17×** lower prospective fault current, hence lower arc-flash energy and easier protection.

DECISION (IV): RIDE-THROUGH IS A LOAD-CLASSIFICATION CHOICE

Figure 8 maps the stored energy (16) required to bridge three load levels—full IT (**30 MW**), a critical set (**6 MW**, **20%**), and plant auxiliaries (**5 MW**)—across timescales from **1 ms** to **8 min**. The millisecond-to-second band is covered by the super-capacitor/flywheel transient buffer the plant already carries; the minutes band crosses into the battery regime. Sizing the three design cases from equation (17) at $\rho_E = 200\ \mathrm{Wh\,kg^{-1}}$ (figure 9) gives

$$\begin{aligned} \text{full-IT, 5 min:} \quad & E = 30\ \mathrm{MW} \times 300\ \mathrm{s} = 9000\ \mathrm{MJ} \;\Rightarrow\; m = 12.5\ \mathrm{t}, \\ \text{critical set, 5 min:} \quad & E = 6\ \mathrm{MW} \times 300\ \mathrm{s} = 1800\ \mathrm{MJ} \;\Rightarrow\; m = 2.5\ \mathrm{t}, \\ \text{black-start aux, 8 min:} \quad & E = 5\ \mathrm{MW} \times 480\ \mathrm{s} = 2400\ \mathrm{MJ} \;\Rightarrow\; m = 3.3\ \mathrm{t}. \end{aligned}$$

A full-IT bridge would demand a $\sim$**12.5 t** battery, disproportionate to the need. The design therefore does *not* carry a full-IT UPS: it gracefully sheds to the $\sim$**2.5 t** critical set (a standard data-center UPS room) while non-critical compute checkpoints or pauses, and it keeps a small $\sim$**3.3 t** black-start battery to boot the plant auxiliaries. On that basis MetroVolt is genuinely black-start capable and islandable: the aux battery energizes the DEC bus, the plant comes online, and the IT load ramps behind it, with no external grid required—the same black-start capability the islanded forward-base burner relies on [37].

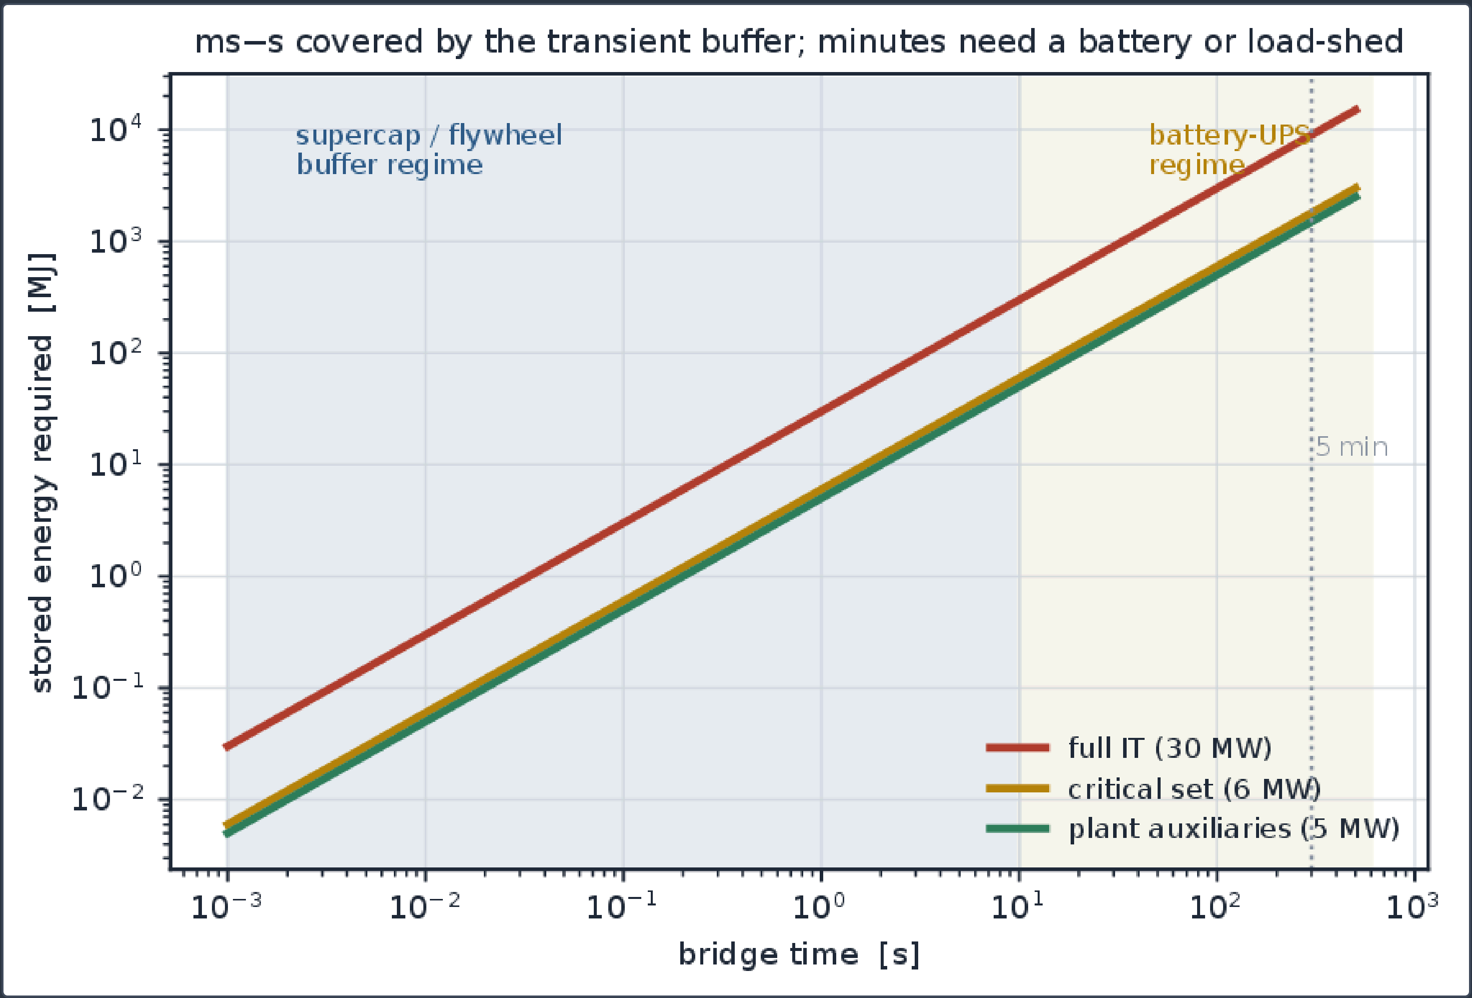

Decision (iv), energy–time map (MV-CX-25). Stored energy $E = Pt$ required to bridge full-IT, critical-set, and plant-auxiliary loads across timescales. The ms–s band (blue) is covered by the transient buffer; the minutes band (gold) needs a battery. The dotted line marks the **5 min** ride-through sizing point.

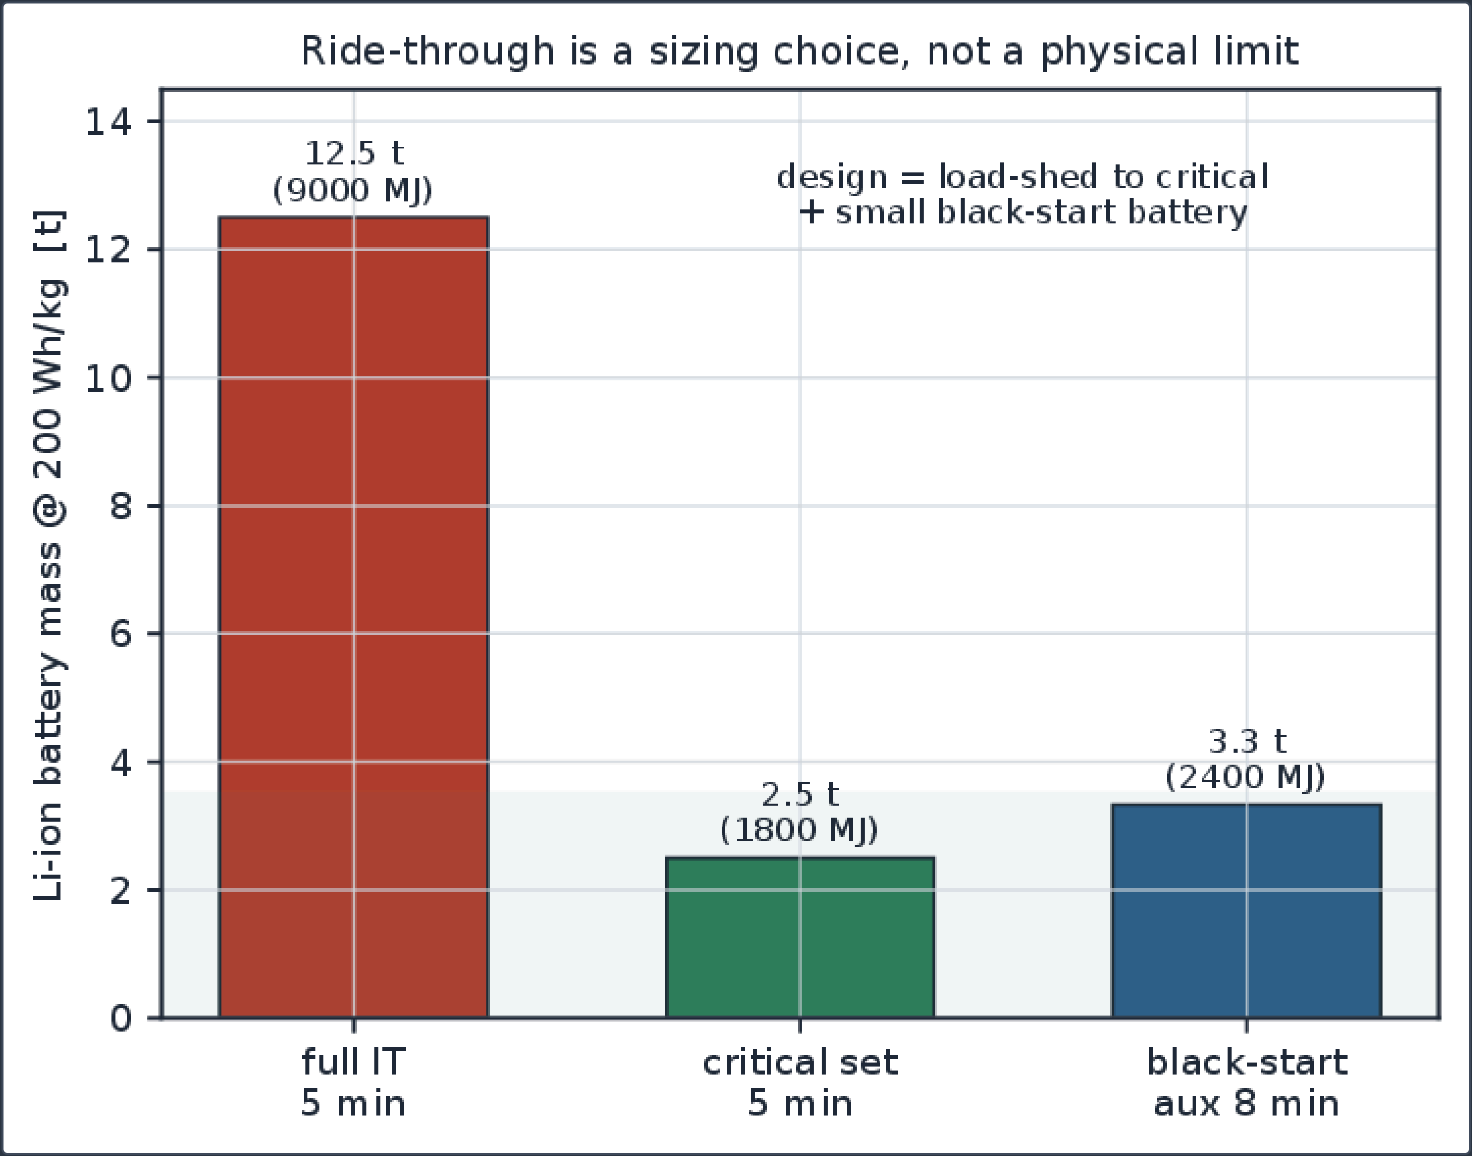

Decision (iv), battery sizing (MV-CX-25). Li-ion mass at $200\ \mathrm{Wh\,kg^{-1}}$ for the three bridge cases, equations (22)–(22). The design (green band) load-sheds to the 2.5 t critical set and carries a 3.3 t black-start battery rather than a 12.5 t full-IT UPS.

THE COMPOSED ARCHITECTURE

The four decisions compose into one coherent, DC-native delivery architecture (figure 1): a DEC-direct HVDC bus with a single SiC dual-active-bridge for isolation (no SST); grid-forming control off the DC buffer; a $2\ \mu$s SiC solid-state breaker backed by the DEC's own current limit for fault protection; and a small black-start battery with graceful load-shed for ride-through. The elements reinforce one another: the same DC buffer that makes the plant grid-forming (decision ii) also bridges the fast transients the breaker and load-shed scheme are sized around (decisions iii, iv), and the single DAB that removes the SST (decision i) is also the galvanic barrier the fault study assumes. The architecture preserves the burner's frozen operating point ($Q_E = 1.318,\ f_{\mathrm{dir}} = 0.49$) and adds no conversion stage the DC-native source does not require.

§ 7

Discussion

Efficiency, against the DC-distribution literature. The $+2.96$-point gain from deleting the SST is consistent in direction and magnitude with the case for DC data-center distribution, where removing redundant AC↔DC conversions is repeatedly found to recover a few points of end-to-end efficiency [9,10,11]. MetroVolt takes that argument to its limit: the source itself is DC, so the front-end conversions that DC-distribution studies still must perform (they start from an AC grid) are absent entirely, and the isolation those studies assign to an SST [13,14] collapses into the single DAB whose 98.85% optimum (18) is squarely in the range reported for SiC megawatt-class DABs [15,16,17]. The novelty is the premise, not the converter.

Grid-forming, against the low-inertia-systems literature. The finding that a fast, droop-type converter response raises the frequency nadir and can hold an island above UFLS is exactly the benefit the grid-forming literature attributes to inverter-based resources in low-inertia grids [18,19,20,21,22,23,24]. Our equation (10)—a droop FFR acting as synthetic damping $D \to D + K_{\text{ffr}}$—is the reduced form of the synchronverter/virtual-synchronous-machine equivalence established by D'Arco and Suul [20,19]. What is specific to MetroVolt is that the storage which supplies the FFR is *already present* for plasma-transient ride-through, so grid-forming is obtained without a dedicated asset; and that the plant's raison d'\^etre—firm power to a critical co-located load—makes the discrete capabilities (black-start, islanding) the operative benefit, beyond the graded nadir number.

Protection, against the DC-breaker literature. The bolted-short peak of $316\ \text{kA}$ is the standard DC-bus objection [25,27], and the standard answer is a fast solid-state breaker [26]. Our contribution is quantitative rather than conceptual: the closed-form LC/I^2t analysis (12)–(15) shows that a microsecond-class SiC SSCB clears on the inductance-limited ramp at $\sim 1\%$ of the resonant peak, and the source-fault comparison turns the DEC's space-charge current limit [6,4] into a $17\times$ reduction in prospective fault current—a safety property that a grid-tied AC plant, drawing on grid-scale short-circuit MVA, cannot claim. Together these make MVDC protection a solved-if-component-dependent problem for this housing, with the DC-native source improving rather than worsening the fault case.

Ride-through, against data-center UPS practice. The energy-sizing result reframes ride-through from a battery-mass problem into a load-classification decision. The $2.5\ \text{t}$ critical-set UPS is an ordinary data-center UPS room, whereas the $12.5\ \text{t}$ full-IT bridge is not proportionate; graceful degradation to a checkpointed critical set is the standard high-availability posture, and it dovetails with the four-nines-uptime and workload-dispatch analysis of the data-center-integration companion [36]. Black-start at $3.3\ \text{t}$ is cheap because it need only boot the plant, not carry the load.

The through-line. Every decision here went *against* the conventional AC-plant choice—no SST, grid-forming rather than following, a solid-state breaker rather than grid-scale fault tolerance, load-shed rather than a full-IT battery—and each went that way because the DC bus and the DEC's current limit make it the better-performing option, not merely the cheaper one (which we do not assess). The result is an architecture that is simpler (fewer stages), safer (lower prospective fault current), and more capable (black-start, islanding) than the AC-mimicking alternative. A DC-native generator, in short, should be delivered as one.

§ 8

Limitations

The one honest open item is the minutes-scale full-IT ride-through. Holding the entire **30 MW** IT load through a five-minute event would require a $\sim$**12.5 t** battery (equation 22), which the design does not carry. This is resolved by graceful load-shed to the $\sim$**2.5 t** critical set rather than by a larger store—a design choice about which workloads are critical (a workload-orchestration decision), not a physical limit on the architecture. The full-IT UPS mass scales linearly with the assumed plant restart time, so a faster warm restart shrinks it and a cold start lengthens it; the plasma black-start trajectory that sets that time is a plant-physics item outside this power-delivery study. Two further caveats bound the scope. The DAB loss coefficients are a representative SiC model, not a measured converter, so while $f^{\star}$ is robust (equation 6) the absolute η_{DAB} carries $\sim\pm$**0.3** points; and the swing-equation nadir uses lumped, representative island parameters, so the robust claim is the *ordering* (fast FFR raises the nadir above UFLS), confirmed across the GPU ensemble of §5, rather than the exact **0.527 Hz**. A full grid-forming stability certification under weak-grid, multi-inverter conditions requires an impedance-based small-signal study beyond the present scope.

§ 9

Conclusion

MetroVolt's power delivery is DC-native throughout, and four architecture decisions follow from that single premise. (i) A DEC-direct HVDC path beats a solid-state-transformer path by 2.96 points at 88.39% with one fewer stage, with isolation carried by a single SiC dual-active-bridge at its arithmetic–geometric-mean optimum (15.81 kHz, 98.85%). (ii) The DC-native buffer makes the plant grid-forming, lifting a 6%-generation-loss frequency nadir by 0.527 Hz ($49.373 \to 49.900$ Hz, clear of UFLS) and enabling the black-start and islanding a grid-following plant cannot provide. (iii) A $2\ \mu$s SiC solid-state breaker bounds a 316 kA undamped resonant fault to ~ 4 kA (1.3% of the peak), while the DEC's current limit caps a source fault to ~ 3 kA against ~ 50 kA from the grid—a $17\times$ safety margin. (iv) Ride-through is met by graceful load-shed to a ~ 2.5 t critical set with a ~ 3.3 t black-start battery, rather than a ~ 12.5 t full-IT UPS. The four decisions compose into a single simpler, safer, islandable architecture that carries the burner's frozen operating point ($Q_E = 1.318$, $f_{\mathrm{dir}} = 0.49$) to the load. Each decision was posed as a formal problem, solved on a GPU-accelerated reduced-order campaign, and anchored to the Physics De-Risking Register; the architecture is the public, technical basis for the MetroVolt balance of plant.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. Circuit and control computations—the deterministic anchors on an in-house HPC node, and the design-space and uncertainty sweeps on the Kronos NVIDIA GPU HPC campaign—produced the results reported here. We thank the Kronos de-risking team for the register anchors and independent re-runs.

Reproducibility. The four power-delivery results are register anchors—MV-CX-04 (SST versus DEC-direct and the DAB optimum), MV-CX-10 (grid-forming versus grid-following), MV-CX-05 (DC-bus fault ride-through) and MV-CX-25 (black-start and UPS bridge)—in the Kronos Physics De-Risking Register, and are regenerable from the archived circuit and control models by a single command; each figure asserts, at plot time, that its curve reproduces its anchor. This paper is archived at doi:10.5281/zenodo.22645941 and its official version is hosted at https://www.kronosfusionenergy.com/publications/power-delivery-topologies-for-a-d-he-tandem-mirror-data.html. The de-risking register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability The circuit- and control-model inputs and the run artifacts (results.json for MV-CX-04/05/10/25) are archived with the deposit at doi:10.5281/zenodo.22645941 and reference the Kronos Physics De-Risking Register doi:10.5281/zenodo.22133057, the MetroVolt burner design paper [2] and the HVDC-native interface paper [35]. Any economic analysis is reported separately and is not included in the public deposit.

Competing interests tocsectionCompeting interests The authors are affiliated with Kronos Fusion Energy. Provisional patent applications relating to aspects of this work have been filed. No other competing interests are declared.

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 E Masanet, A Shehabi, N Lei, S Smith and J Koomey, Recalibrating global data center energy-use estimates, *Science* **367** 984 (2020). doi:10.1126/science.aba3758

2 P I Ford *et al*, *The MetroVolt Data-Center Burner: Direct-DC Campus Power, Q_E Closure, and the Plug Requirement in a Low-Neutron D–^{3}He Tandem Mirror*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645695; https://www.kronosfusionenergy.com/publications/metrovolt-burner.html

3 R W Moir and W L Barr, Venetian-blind direct energy converter for fusion reactors, *Nucl. Fusion* **13** 35 (1973). doi:10.1088/0029-5515/13/1/005

4 R F Post, The magnetic mirror approach to fusion, *Nucl. Fusion* **27** 1579 (1987). doi:10.1088/0029-5515/27/10/001

5 P I Ford *et al*, *Direct Energy Conversion for a D–^{3}He Tandem-Mirror Burner: Recovering Charged-Particle Power Directly Without a Thermal Cycle*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645700; https://www.kronosfusionenergy.com/publications/direct-energy-conversion.html

6 P I Ford *et al*, *Direct-Energy-Converter Collector Physics in a D–^{3}He Tandem-Mirror Burner: Multi-Stage Deceleration, Secondary Emission, and the Space-Charge-Limited Grid Optimum*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645858; https://www.kronosfusionenergy.com/publications/direct-energy-converter-collector-physics-in-a-d-he-tan.html

7 T K Fowler, R W Moir and T C Simonen, A new simpler way to obtain high fusion power gain in tandem mirrors, *Nucl. Fusion* **57** 056014 (2017).

8 D Endrizzi *et al*, Physics basis for the Wisconsin HTS Axisymmetric Mirror (WHAM), *J. Plasma Phys.* **89** 975890501 (2023).

9 M Ton, B Fortenbery and W Tschudi, *DC Power for Improved Data Center Efficiency*, Lawrence Berkeley National Laboratory report (2008).

10 G AlLee and W Tschudi, Edison redux: 380 Vdc brings reliability and efficiency to sustainable data centers, *IEEE Power Energy Mag.* **10**(6) 50 (2012).

11 T Dragicevi\'c, X Lu, J C Vasquez and J M Guerrero, DC microgrids—Part I: a review of control strategies and stabilization techniques, *IEEE Trans. Power Electron.* **31** 4876 (2016).

12 T Dragicevi\'c, X Lu, J C Vasquez and J M Guerrero, DC microgrids—Part II: a review of power architectures, applications, and standardization issues, *IEEE Trans. Power Electron.* **31** 3528 (2016).

13 X She, A Q Huang and R Burgos, Review of solid-state transformer technologies and their application in power distribution systems, *IEEE J. Emerg. Sel. Topics Power Electron.* **1** 186 (2013).

14 J E Huber and J W Kolar, Applicability of solid-state transformers in today's and future distribution grids, *IEEE Trans. Smart Grid* **10** 317 (2019).

15 R W De Doncker, D M Divan and M H Kheraluwala, A three-phase soft-switched high-power-density DC/DC converter for high-power applications, *IEEE Trans. Ind. Appl.* **27** 63 (1991).

16 B Zhao, Q Song, W Liu and Y Sun, Overview of dual-active-bridge isolated bidirectional DC–DC converter for high-frequency-link power-conversion system, *IEEE Trans. Power Electron.* **29** 4091 (2014).

17 J Mill\'an, P Godignon, X Perpi\ n\`a, A P\'erez-Tom\'as and J Rebollo, A survey of wide bandgap power semiconductor devices, *IEEE Trans. Power Electron.* **29** 2155 (2014).

18 J Rocabert, A Luna, F Blaabjerg and P Rodr\'iguez, Control of power converters in AC microgrids, *IEEE Trans. Power Electron.* **27** 4734 (2012).

19 Q-C Zhong and G Weiss, Synchronverters: inverters that mimic synchronous generators, *IEEE Trans. Ind. Electron.* **58** 1259 (2011).

20 S D'Arco and J A Suul, Equivalence of virtual synchronous machines and frequency-droop control in inverter-based systems, *IEEE Trans. Smart Grid* **5** 394 (2014).

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

Power-Delivery Topologies for a D–^{3}He Tandem-Mirror Data-Center Burner: Solid-State-Transformer versus Direct-DC Conversion, Grid-Forming Control, and DC-Bus Fault-Ride-Through

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

Zenodo · doi:10.5281/zenodo.22645941

Read on kronosfusionenergy.com

The Physics De-Risking Register