KRONOS FUSION ENERGY

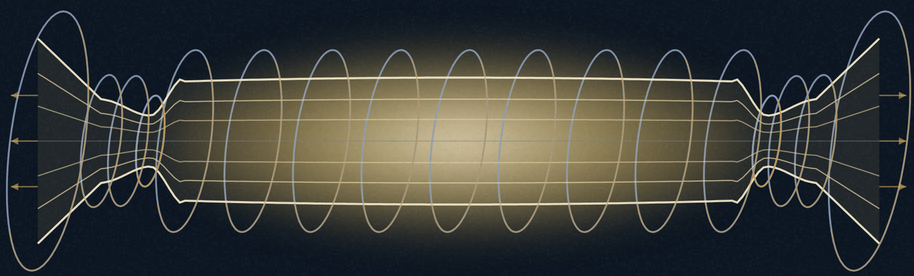

PAPER 1.2B · THE KRONOS 2026 SERIES

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

The load that a hyperscale artificial-intelligence campus imposes on the grid is native direct current, growing faster than transmission can be built, and increasingly sited where...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

The MetroVolt Data-Center Burner: Direct-DC Campus Power, Q_E Closure, and the Plug Requirement in a Low-Neutron D–³He Tandem Mirror

The load that a hyperscale artificial-intelligence campus imposes on the grid is native direct current, growing faster than transmission can be built, and increasingly sited where interconnection queues are measured in years. We present MetroVolt, a data-center power plant built on the Kronos deuterium–helium-3 (D–^{3}He) tandem-mirror burner, whose direct energy converter (DEC) delivers high-voltage direct current *natively*—the same physical form the compute hall consumes—so that the plant couples to the load through two-to-three power-conversion stages rather than the four-to-six of a thermal generator feeding an alternating-current (AC) grid and an uninterruptible supply. At the frozen design point the burner reaches a plant energy multiplication $Q_E = 1.318$ at fusion power 4298.5 MW with a neutron fraction of only $5.44\,\%$ of the total power—the low-activation advantage of the fuel—and a length ladder delivers net electric output of $+104/+850/+2832$ $\mathrm{MW_e}$ at $55/440/1400$ m device length, with a Monte-Carlo uncertainty quantification returning $P(Q_E > 1) = 1.0$ over the design ranges. The binding physics gate is the end-plug potential: confining the central-cell ions requires a plug-to-central-cell density ratio $n_p/n_c = 16$ built from an ion-confining potential $\phi_i = 248.75$ kV against a carrier temperature $T_e = 89.7$ keV, a requirement-class target more than two orders of magnitude beyond the demonstrated tandem-mirror record and retired only by a dedicated plug-potential experiment; direct energy conversion is a bounded upside lever, taken at a deliberately conservative module efficiency $\eta_{\mathrm{DEC}} = 0.49$, not the closure mechanism. A second, independent modelling choice—the synchrotron radiation transport—is shown to be equally load-bearing: the same machine reads $Q_E = 0.39$ optically thin, $Q_E = 1.32$ under the reflective-cavity Albajar–Fidone–Johner model adopted as canon, and $Q_E = 4.56$ under Trubnikov, so that a first-wall reflectivity $R_{\mathrm{wall}} \gtrsim 0.83$ is a hard specification, not a convenience. Reduced to the data-center product, the direct-DC path raises the source-to-chip efficiency from $75.9\,\%$ (conventional AC) to $88.1\,\%$ (DEC-direct medium-voltage DC); a plasma-independent electronic buffer follows synchronized GPU training swings at ~ 80 $\mathrm{MW\,s^{-1}}$—about $105\times$ any plasma throttle—while holding the DC bus to $\pm 0.40\,\%$; waste-heat-driven absorption cooling takes the facility from PUE 1.35 to 1.17; and the neutron-lean $90\,\%$-aneutronic burn licenses against a mobile tritium source term of ~ 3.6 g, roughly $275\times$ below a deuterium–tritium plant of the same power. We separate throughout what is robust—the neutron-source performance and the power-electronics topology—from what is conditional on retiring the plug gate. This is a comprehensive, review-length treatment: twenty-six figures, full governing-equation derivations, a GPU-accelerated computational methodology, a global sensitivity and uncertainty study, and a benchmarking section against the published tandem-mirror, direct-conversion, synchrotron, and data-center-distribution literature.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645695 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: tandem mirror · deuterium--helium-3 · direct energy conversion · data-center power · HVDC-native distribution · plug potential · low-neutron fusion · behind-the-meter · synchrotron radiation · uncertainty quantification

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)

Nomenclature tocsectionNomenclature

@L>X@ SYMBOL	DEFINITION
Q_E	plant energy multiplication (gross electric out / recirculating in)
f_n	neutron fraction of total fusion power
n_p/n_c (R)	plug-to-central-cell density ratio (the binding gate)
$\phi_i,\ \varphi_i$	ion-confining electrostatic plug (ambipolar) potential
T_e	carrier (plug) electron temperature that holds the potential
T_i	central-cell ion temperature
$R_m,\ R_{mc}$	plug mirror ratio
f_w	warm-ion fill fraction inside the loss cone
L_n/ρ_i	plug-edge density-gradient scale in ion Larmor radii
η_{DEC}	direct-energy-converter module efficiency (declared, conservative input)
η_{th}	thermal-cycle efficiency on the residual (non-direct) power
P_{fus}, P_n	fusion power; neutron power
$P_{\mathrm{ch}}, P_{\mathrm{ax}}$	recoverable charged-particle power; axial charged power intercepted by the DEC
P_{plug}	recirculating power to sustain the plug and thermal barrier
f_{dir}	direct-conversion fraction of gross electric output
R_{wall}	first-wall synchrotron reflectivity
η_{chain}	source-to-chip electrical delivery efficiency
PUE	power-usage effectiveness, $P_{\mathrm{facility}}/P_{\mathrm{IT}}$
ERE, ERF	energy-reuse effectiveness; energy-reuse factor
CF	plant capacity factor
n^*	synchrotron reabsorption cutoff harmonic
$\tau_{\parallel}, \tau_{ii}$	parallel (end-loss) confinement time; ion–ion collision time
S_1, S_T	first-order and total-order Sobol sensitivity indices
MTBF	mean time between failures of a power module

§ 1

Introduction

THE DATA-CENTER LOAD IS A FUSION PROBLEM IN DISGUISE

Artificial-intelligence compute has made the data center one of the fastest-growing electrical loads in the industrialized world [29,33], and its demand has three features that a conventional generator serves badly. It is *native direct current*: a modern GPU hall converts the incoming alternating current down through an uninterruptible-power-supply (UPS) double conversion and a rack power-supply unit to a low-voltage DC bus at the silicon, discarding several percent of the delivered energy at each stage [30,31,32]. It is *transient*: synchronized large-model training draws crest-factor swings of order 1.2 on second timescales that a thermal plant cannot follow with the plasma alone. And it is *sited ahead of the grid*: hyperscale campuses are increasingly built where land and cooling are available rather than where transmission is, so interconnection queues of order four years now gate time-to-power. A generator that produced its output as regulated high-voltage DC, followed a load electronically without touching the plasma, and could sit behind the meter would match this load in a way that no steam turbine feeding an AC grid does.

The Kronos D–^{3}He tandem-mirror burner is exactly such a generator. Its direct energy converter (DEC) recovers charged-particle power *as electricity, directly, without a thermal cycle*, and the natural output of an electrostatic decelerator is high-voltage direct current—the same physical quantity the compute hall wants. This paper is the data-center product realization of that burner, which we call MetroVolt. It shares its plasma physics identically with the Aegis forward-base burner [49]: the two are one machine in two housings, and the physics closure, the binding plug gate, and the low-neutron advantage are common to both. What distinguishes MetroVolt is the coupling to the load—the HVDC-native power-electronics path, the electronic load-following, the waste-heat integration, the power density, and the urban siting case—and that is where this paper concentrates, after establishing the shared physics on a common footing at review length.

PRIOR ART: THE CONCEPTS METROVOLT INHERITS

MetroVolt sits at the confluence of four mature bodies of work, and it is worth stating what each contributes and what each leaves open, because the burner's honest gate lives precisely where the demonstrated record stops.

The tandem mirror. The open magnetic mirror is one of the oldest fusion confinement concepts, and its modern form—the tandem mirror with electrostatic end plugs—was set out by Dimov and by Fowler and Logan [1], then sharpened by the thermal-barrier variant of Baldwin and Logan [2] that decouples plug electrons from central-cell electrons so that a modest plug density can hold a large confining potential. The end-loss physics of an electrostatically plugged mirror was placed on a kinetic footing by Pastukhov [6] and generalized by Cohen and co-workers [7], and the whole subject was reviewed by Baldwin [8] and by Post [9]. Experimentally, thermal-barrier potentials and enhanced axial confinement were produced and diagnosed on TMX-U [10], and the axisymmetrized GAMMA-10 device formed thermal barriers and confining potentials and remains the reference for the demonstrated plug potential [11,12]. The gas-dynamic-trap line at Budker demonstrated high-β, MHD-stable sloshing-ion operation in an axisymmetric geometry [13], and the high-temperature-superconductor axisymmetric mirror WHAM is now bringing high mirror fields to bear on the *unplugged* confinement regime [15,16]. The magnetic-well and minimum-B stabilization heritage traces to Ioffe [14], the theory of MHD-stable *axisymmetric* mirrors—the geometry MetroVolt adopts—to Ryutov and co-workers [41,42], and the loss-cone microstability literature to the drift-cyclotron loss-cone (DCLC) and Alfv\'en-ion-cyclotron (AIC) branches. What no experiment has demonstrated is a plug potential at the *reactor-relevant* magnitude the present net-power case requires; that gap is the paper's binding gate (§2.5).

Advanced (low-neutron) fuels. The case for D–^{3}He was made by Kulcinski and Santarius [3,4] and its confinement requirements reviewed by Nevins [18]: the fuel is attractive because the primary channel is aneutronic and the neutron budget comes only from residual D–D and D–T side reactions, which minimizes activation, shielding mass, and the mobile tritium inventory that dominates a licensing case [5]. The cost is a demanding confinement problem—the plug—and a fuel-supply problem for ^{3}He that is a separate, lunar-gated matter treated in a companion paper [64] and outside the scope of the present power-plant design.

Direct energy conversion. That most of a D–^{3}He mirror's output leaves as directed, fast charged particles is what makes electrostatic direct conversion possible. The venetian-blind collector of Post [19], the multi-stage analyses of Moir and Barr [20,21], and the modern expanding-plasma reappraisal of Rax, Kolmes and Fisch [22] set the efficiency expectations we anchor against; RF alpha-channeling, which redirects fast-ion energy into the fuel, follows Fisch and Rax [23] and the broader RF current-drive and wave–particle theory reviewed by Fisch [43].

DC-native data-center distribution. Independently, the data-center community has converged on DC distribution and direct-liquid cooling as the levers on facility efficiency: $380–400$ V DC racks and their measured savings [30,31,32], the migration to liquid and immersion cooling [27], and the power-usage-effectiveness (PUE) and energy-reuse metrics standardized for reporting facility overhead [28,29,37]. MetroVolt's contribution is to supply the *front* of that chain in the native form the chain wants—a generator whose output is already HVDC.

WHY A LOW-NEUTRON MIRROR, AND WHAT IT COSTS

A D–^{3}He mirror is attractive because it produces few neutrons and difficult because it needs a large electrostatic plug to confine its ions and a fuel that is scarce on Earth [1,3,4]. The low neutron fraction ($f_n = 5.44\,\%$, from residual D–D side reactions and the D–T burn of the bred triton rather than the aneutronic primary channel) is what lets a fusion source sit next to a data-center campus: it minimizes activation, shielding mass, and the mobile tritium inventory that dominates a licensing case [5]. We are explicit throughout about which results are robust and plug-independent (the neutron source, the power-electronics topology) and which are conditional on retiring the plug gate (every net-power number).

CONTRIBUTION AND ROADMAP

This paper contributes: (i) a self-contained statement of the burner's governing equations and the formal power-balance problem, with the binding plug-potential requirement written as an explicit exponential gate and the Pastukhov–Cohen end-loss result derived step by step (§2); (ii) the two load-bearing model-form choices—the plug potential and the synchrotron transport—made explicit and quantified, including the $R_{\text{wall}} \gtrsim 0.83$ specification (§2.9); (iii) the GPU-accelerated computational methodology by which each gate was reduced, naming the real codes, the acceleration, the problem scale, and the verification/convergence evidence (§3); (iv) the direct-DC data-center product physics—the HVDC-native delivery chain, the electronic load-following, the grid-forming and fault behaviour, the PUE pathway, the power density, the fleet reliability, and the urban siting and tritium source term—each grounded in a frozen register anchor (§4); (v) a results section of twenty-six figures and supporting tables (§5); (vi) a dedicated global sensitivity and uncertainty study (§6); (vii) a quantitative benchmarking section against the published tandem-mirror, direct-conversion, synchrotron, and data-center-power literature (§7); and (viii) technical appendices with the full derivations, numerical schemes, and parameter tables. Every headline quantity traces to a serial identifier in the Kronos Physics De-Risking Register [47], cited inline so each number is auditable.

ONE MACHINE, TWO HOUSINGS: THE RELATIONSHIP TO AEGIS

MetroVolt and the Aegis forward-base burner [49] are the same plasma device in two product housings: they share the central-cell power balance, the plug-potential requirement, the synchrotron-limited operating point, the direct-energy converter, and the low-neutron advantage identically, and every physics result in §2–§2.9 is common to both. What differs is the coupling to the load. Aegis is an islanded, black-start-capable microgrid source for a forward base, where the design pressure is autonomy and transportability; MetroVolt is a behind-the-meter, HVDC-native campus source, where the design pressure is delivered efficiency to a DC silicon load, electronic load-following of a crest-factor training profile, waste-heat integration, and fleet availability written to a data-center service-level agreement. Presenting the two together on a common physics footing is deliberate: it makes explicit that the binding plug gate and the neutron-source performance are properties of the shared plasma, while the product-specific results—the power-electronics topology, the PUE pathway, the reliability model—are properties of the housing. This paper carries the shared physics at review depth so that the data-center product results stand on a fully stated foundation rather than a citation.

§ 2

Governing equations and the formal problem

FUSION POWER AND THE NEUTRON FRACTION

The central cell burns deuterium and helium-3 through the primary aneutronic channel

$$\mathrm{D} + {}^{3}\mathrm{He} \;\rightarrow\; p\,(14.68\,\mathrm{MeV}) + \alpha\,(3.67\,\mathrm{MeV}),$$

with unavoidable deuterium–deuterium side reactions $\mathrm{D}+\mathrm{D}\rightarrow {}^{3}\mathrm{He}\,(0.82\,\mathrm{MeV}) + n\,(2.45\,\mathrm{MeV})$ and $\mathrm{D}+\mathrm{D}\rightarrow \mathrm{T}\,(1.01\,\mathrm{MeV}) + p\,(3.02\,\mathrm{MeV})$, the first of which, together with the D–T burn of the bred triton ($\mathrm{D}+\mathrm{T}\rightarrow \alpha\,(3.5\,\mathrm{MeV}) + n\,(14.03\,\mathrm{MeV})$), supplies the residual neutron power. For a Maxwellian plasma the volumetric reaction rate between species j, k is

$$\mathcal{R}_{jk} = \frac{n_j n_k}{1+\delta_{jk}}\,\langle\sigma v\rangle_{jk}(T_i), \qquad P_{\mathrm{fus}} = \sum_{jk}\int_V \mathcal{R}_{jk}\,\mathcal{E}_{jk}\,\mathrm{d}V,$$

with $\langle\sigma v\rangle$ the fusion reactivity, δ_{jk} the Kronecker delta that avoids double-counting like-species pairs, and $\mathcal{E}_{jk}$ the reaction energy. The neutron fraction is the ratio of neutron-carried power to total,

$$f_n = \frac{P_n}{P_{\mathrm{fus}}} = \frac{\int_V \left(\mathcal{R}_{\mathrm{DD}n}\,E_n^{\mathrm{DD}} + \mathcal{R}_{\mathrm{DT}}\,E_n^{\mathrm{DT}}\right)\mathrm{d}V}{P_{\mathrm{fus}}},$$

and at the frozen operating point $P_{\mathrm{fus}} = 4298.5$ MW, $P_n = 233.95$ MW, so $f_n = 5.44\,\%$ (register serial BU-L1-A1). The operating ion temperature $T_i \simeq 90$ keV sits just past the reactivity-weighted maximum of Q_E on the thermally stable branch, where $\mathrm{d}(P_{\mathrm{heat}} - P_{\mathrm{loss}})/\mathrm{d}T_i < 0$ so that the burn is self-correcting against a temperature excursion (BU-CX-30 line; a -29 MW keV^{-1} restoring slope), an important property for a plant asked to hold a steady baseload behind a data center.

THE D–^{3}HE REACTIVITY AND THE CHOICE OF OPERATING TEMPERATURE

Why the burner runs at $T_i \simeq 90$ keV—an order of magnitude hotter than a D–T tokamak—follows directly from the reactivity. The D–^{3}He fusion cross-section peaks near a centre-of-mass energy of ~ 250 keV, far above the ~ 65 keV of the D–T resonance, so its Maxwellian-averaged reactivity $\langle\sigma v\rangle_{\mathrm{D^3He}}$ rises steeply through the tens-of-keV range and does not reach a useful value until $T_i \gtrsim 50$ keV. Writing the charged-particle power density as

$$p_{\mathrm{ch}} = \frac{n_{\mathrm{D}}\,n_{^3\mathrm{He}}}{1+\delta}\,\langle\sigma v\rangle_{\mathrm{D^3He}}(T_i)\,\mathcal{E}_{\mathrm{ch}}, \qquad \mathcal{E}_{\mathrm{ch}} = 18.35\,\mathrm{MeV},$$

and the competing electron-radiation loss density $p_{\mathrm{rad}} = p_{\mathrm{br}} + p_{\mathrm{sync}}$ (bremsstrahlung scaling as $\propto n_e^2 T_e^{1/2} Z_{\mathrm{eff}}$ and synchrotron as in §2.9), the net electron power balance $p_{\mathrm{ch}}\,\chi_e - p_{\mathrm{rad}}$ has a broad maximum in T_i around 85–90 keV once the synchrotron term is treated with the correct reflective-cavity transport. Below that band the reactivity is too weak to overcome radiation; above it the reactivity gain flattens while radiation continues to climb. The design point therefore sits at the reactivity-weighted maximum, and the 85 keV ridge seen in the design-space scan (§6, where $Q_E = 1.375$ at the design-point density) is the same maximum viewed along the temperature axis. The choice is not free: reaching and holding 90 keV electrons against end-wall cooling is exactly what the ECH-sustained thermal barrier must do (§2.9), which is why the barrier is worth a factor of $+2.8$ in Q_E.

THE NEUTRON BUDGET RESOLVED BY REACTION BRANCH

Resolving the fusion power into its channels (table 2, figure 1) makes the neutron-lean character concrete. The aneutronic D–^{3}He channel carries 3868.0 MW; the charged D–D proton branch adds 129.7 MW; and the neutron budget is set by the 2.45 MeV D–D branch (124.8 MW reaction total) and the 14 MeV D–T self-burn (176.0 MW reaction total). The four branches sum to $P_{fus} = 4298.5$ MW exactly, and the charged (direct-convertible) power is $P_{fus} - P_n = 4064.6$ MW (ENV-A4-BU, BV-07).

BRANCH	REACTION POWER (MW)	ROLE
D–^{3}He (aneutronic)	3868.0	direct-convertible charged products
D–D proton	129.7	charged
D–D neutron (2.45 MeV)	124.8	neutron budget (count-dominant)
D–T (14 MeV)	176.0	neutron budget (power-dominant)
total P_{fus}	**4298.5**	$f_n = 5.44\,\%$, $P_n = 233.95$ MW

Fusion power by reaction branch at the frozen design point; the branches sum to P_{fus} exactly (ENV-A4-BU, BV-07). The neutron power $P_n = 233.95$ MW is the sum of the neutron-carried fractions of the D–D and D–T lines; the reaction-total column below sums to P_{fus}.

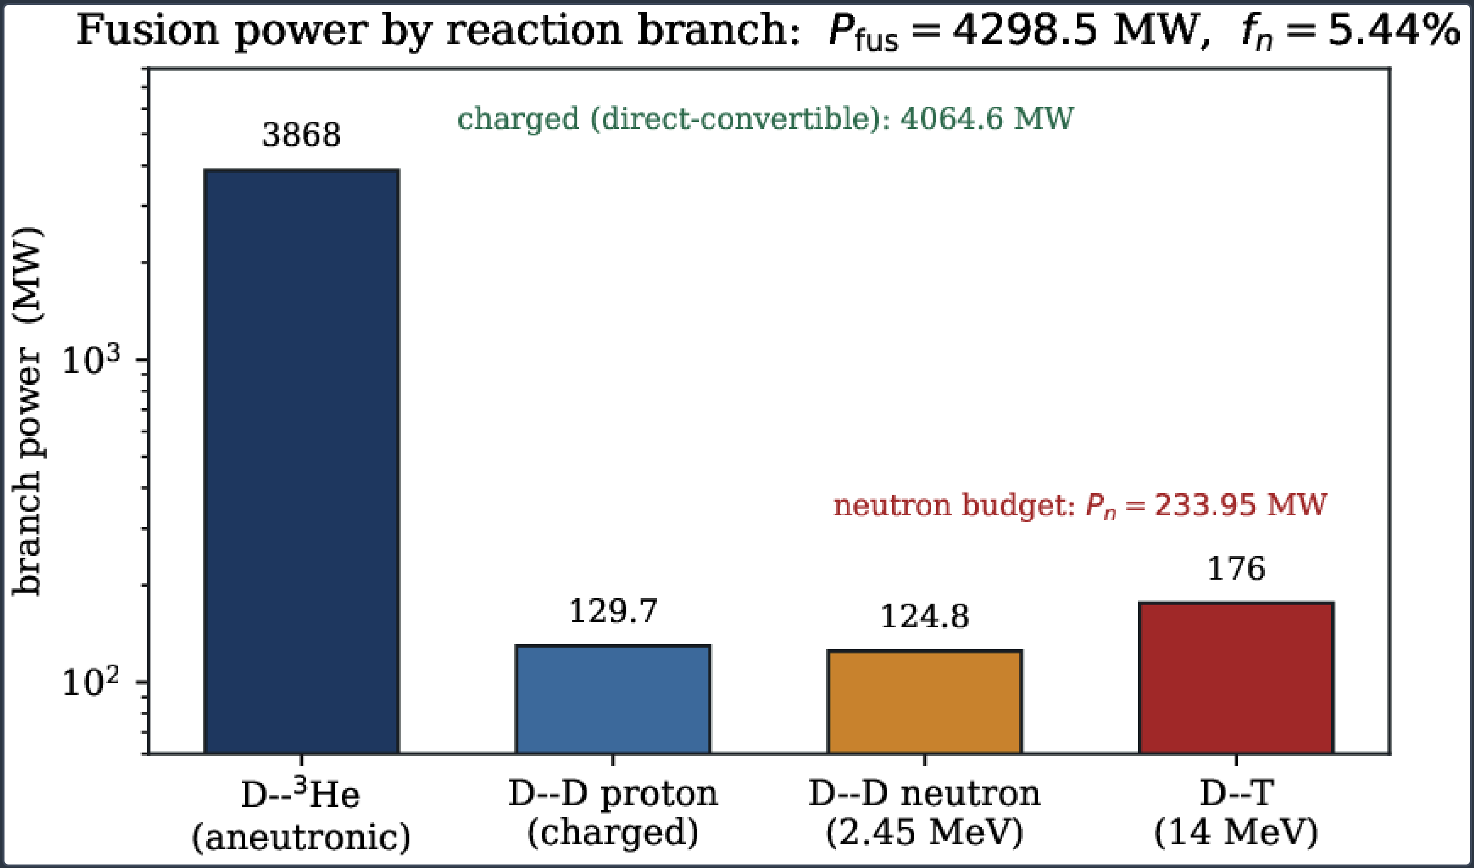

Fusion power resolved by reaction branch (ENV-A4-BU, BV-07). The aneutronic D–^{3}He channel dominates; the neutron fraction is $f_n = 5.44\,\%$ ($P_n = 233.95$ MW), and 4064.6 MW leaves as directly-convertible charged power. Axis is logarithmic; branch powers reproduce to the stated digits.

The neutron budget itself divides sharply by energy line (figure 2): the 14 MeV D–T self-burn carries $\sim 60\,\%$ of the neutron *power* (~ 140 MW) but only $\sim 21\,\%$ of the neutron *count*, while the 2.45 MeV D–D line carries the remaining $\sim 40\,\%$ of the power (~ 93.5 MW) and $\sim 79\,\%$ of the count (BV-07, ENV-A4-BU). The wall-averaged neutron loading is only 0.098 $\mathrm{MW\,m^{-2}}$, one to two orders of magnitude below a D–T first wall, and the 14 MeV line—being power-dominant—sets the biological shield basis (§4.10).

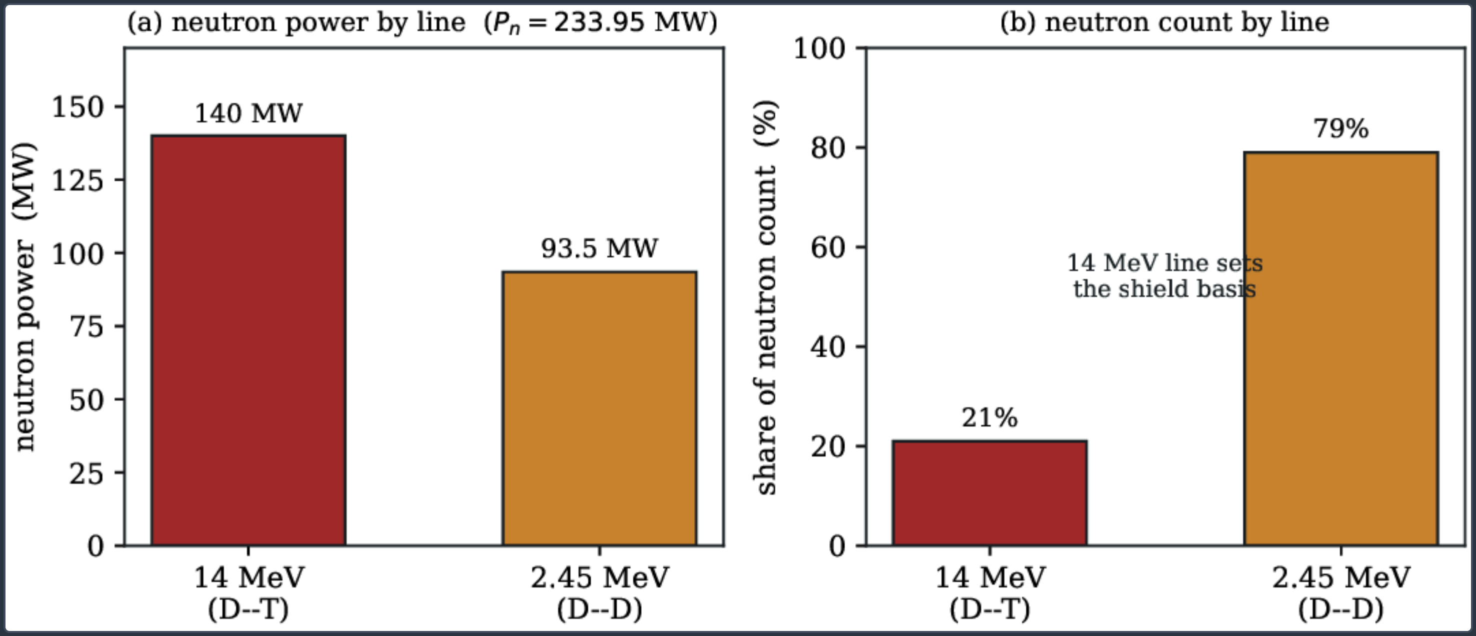

Neutron budget by energy line (BV-07, ENV-A4-BU). **(a)** The 14 MeV D–T line carries $\sim 60\,\%$ of the neutron power; the 2.45 MeV D–D line the remaining $\sim 40\,\%$. **(b)** By count the ordering reverses—the 2.45 MeV line dominates. The power-dominant 14 MeV line sets the shield basis.

THE PLANT ENERGY MULTIPLICATION

The quantity that decides whether the plant delivers net electricity is not the fusion gain but the plant energy multiplication—the ratio of gross electric output to the power that must be recirculated to sustain the plug and the thermal barrier. Writing P_{ch} for the axial charged-particle power that the DEC intercepts, P_{th} for the residual power taken through a thermal cycle at efficiency η_{th}, and $P_{\mathrm{plug}}, P_{\mathrm{aux}}$ for the recirculating plug and auxiliary powers,

$$Q_E = \frac{\eta_{\mathrm{th}}\, P_{\mathrm{th}} + \eta_{\mathrm{DEC}}\, P_{\mathrm{ch}}}{P_{\mathrm{plug}} + P_{\mathrm{aux}}}, \qquad P_{\mathrm{net}} = \left(P_{\mathrm{plug}} + P_{\mathrm{aux}}\right)\left(Q_E - 1\right),$$

so that Q_E improves with recovered charged power but is bounded from below by the recirculating plug power. The solver returns $Q_E = 1.3181$ with a self-consistent split (BU-L1-A1, BU-L1-A10): the axially directed charged power that the converter intercepts is $P_{\mathrm{ax}} = 2719.7$ MW, of which the DEC recovers 1903.8 $\mathrm{MW_e}$ at $\eta_{\mathrm{DEC}} = 0.70$—$54.1\,\%$ of the gross electric output $P_{\mathrm{out}} = 3522.0$ $\mathrm{MW_e}$, so the direct-conversion fraction is $f_{\mathrm{dir}} = 0.495$, with the balance from a thermal bottoming cycle on the residual power. Because the recirculating denominator is set by the plug, the converter efficiency is a bounded upside lever and not the closure mechanism; the net-power sign flips only below $\eta_{\mathrm{DEC}} \simeq 0.52$ (BU-L1-A10), so we carry $\eta_{\mathrm{DEC}} = 0.49$ as a deliberately conservative anchor—below the 0.56 engineering floor and the 0.70 Moir–Barr design value of the multi-stage-collector efficiency ladder, and well below the 0.85 ideal proton-only ceiling (BU-SX-01 adjudication; converter physics in [51])—so that no headline number leans on an optimistic conversion assumption. That the closure survives even below the model floor is the strongest single statement of its robustness.

THE BINDING GATE: THE PLUG POTENTIAL

Confining the central-cell ions against axial loss requires an electrostatic plug of potential ϕ_i, and the plug density needed to build it follows the Boltzmann relation between the confined and carrier populations. If the plug electrons are in Boltzmann equilibrium along the field line at temperature T_e, the density contrast that a potential step ϕ_i sustains is

$$\frac{n_p}{n_c} = \exp\left(\frac{e\phi_i}{T_e}\right),$$

so that at $\phi_i = 248.75$ keV and carrier temperature $T_e = 89.7$ keV the required ratio is $\exp(248.75/89.72) = 16$ (BU-L2-A5). Because the requirement sits inside the exponent, the carrier temperature is the strongest lever on it (figure 3); differentiating equation (6),

$$\frac{\mathrm{d}(n_p/n_c)}{n_p/n_c} = \frac{e\phi_i}{T_e}\frac{\mathrm{d}T_e}{T_e} \times (-1) \quad \Rightarrow \quad \frac{\Delta(n_p/n_c)}{n_p/n_c} \simeq -\frac{e\phi_i}{T_e}\frac{\Delta T_e}{T_e},$$

so with $e\phi_i/T_e = \ln 16 \approx 2.77$ a $10\,\%$ change in T_e moves the required ratio by roughly a factor of 1.3.

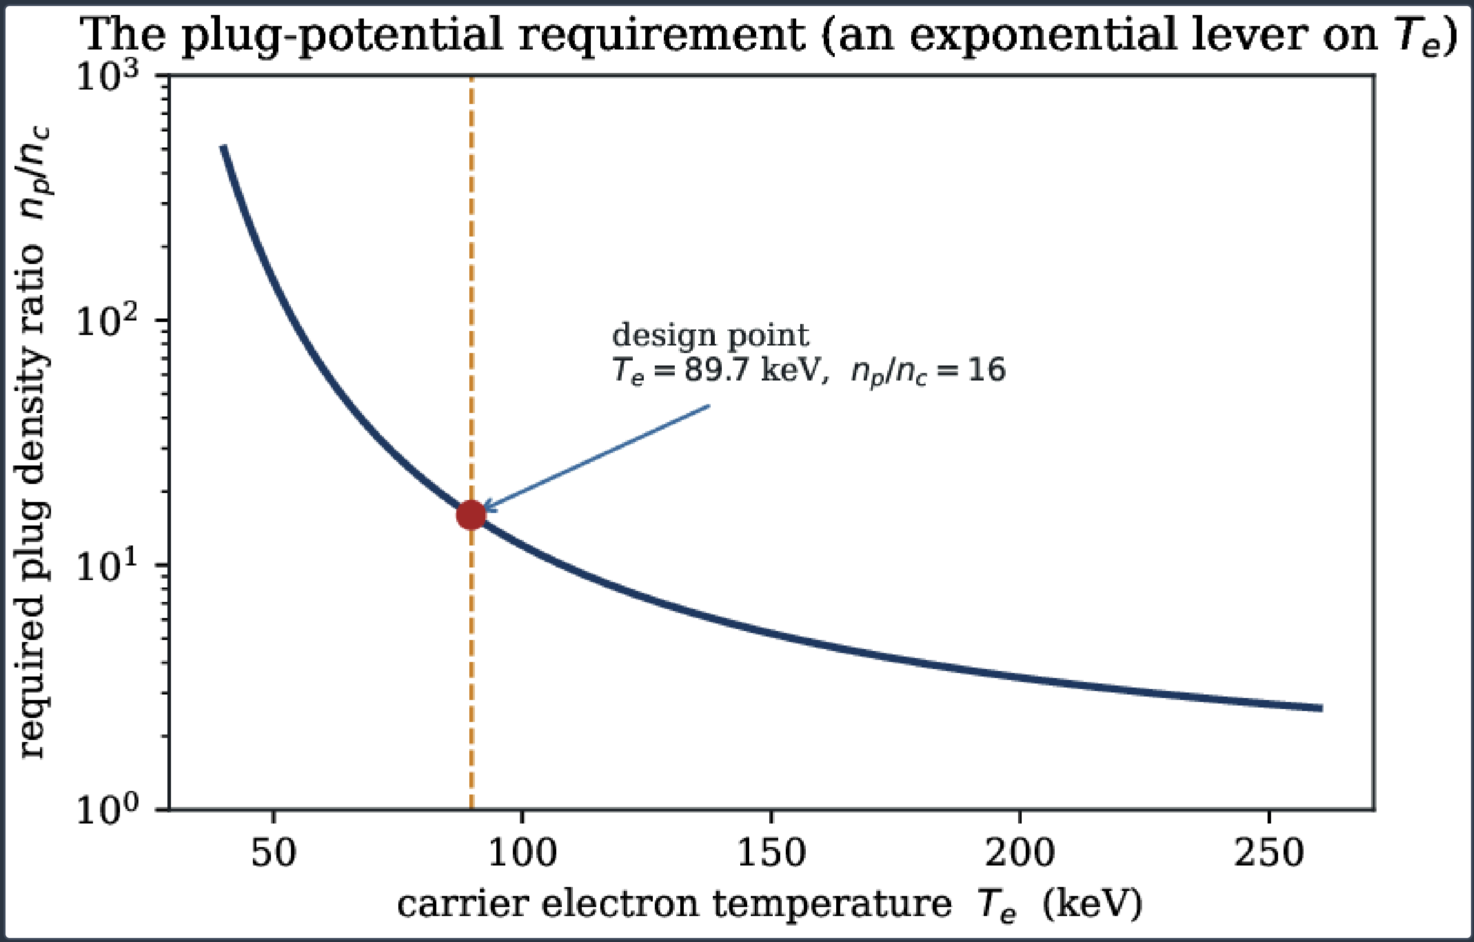

The plug-density requirement $n_p/n_c = \exp(e\phi_i/T_e)$ of equation (6) (BU-L2-A5). The design point sits at $n_p/n_c = 16$; the carrier temperature T_e is the exponential lever that sets it. This is the binding gate on which every net-power result is conditional.

This is a requirement-class target: ϕ_i is more than two orders of magnitude beyond the demonstrated tandem-mirror plug-potential record set by TMX-U and GAMMA-10 [10,12,11], and it is retired only by a dedicated plug-potential experiment at reactor-relevant scale, not by design study. We state this plainly and do not claim it retired; the potential is formed by an ECH-sustained hot-electron/sloshing-ion thermal barrier whose physics is a companion subject [57,60].

DERIVATION OF THE AMBIPOLAR-WELL CONFINEMENT TIME

The confinement improvement the plug buys is the exponential enhancement of the ambipolar-well ion-confinement time over the collisional time. We sketch the Pastukhov–Cohen result [6,7] because the exponential factor is the physical reason a modest-looking potential produces a large confinement gain (Appendix 11 gives the full derivation). An ion confined in a mirror well of depth $e\phi_i$ escapes only when pitch-angle scattering carries it into the loss cone *and* it has enough parallel energy to climb the potential. The loss rate is governed by the Fokker–Planck collision operator acting on the tail of the distribution; solving the steady-state kinetic equation with an absorbing loss-cone boundary gives, in the strong-potential limit $e\phi_i/T_i \gg 1$,

$$\tau_\parallel \simeq \tau_{ii} \frac{2R_m}{\sqrt{\pi}\,(2R_m+1)\,\ln(4R_m+2)} \left(\frac{e\phi_i}{T_i}\right) \exp\left(\frac{e\phi_i}{T_i}\right),$$

with τ_{ii} the ion–ion collision time and R_m the plug mirror ratio. The bracketed geometric factor is a weak, logarithmic function of the mirror ratio; the physics is in the product $(e\phi_i/T_i)\exp(e\phi_i/T_i)$, which grows by more than a factor of ten as $e\phi_i/T_i$ rises from 1 to 3 (figure 4). This is why the required ϕ_i is so consequential—and why realizing it, not converting the power, is the closure problem. Equations (6)–(8) are the mathematical core of the burner: the entire net-power case rests on realizing $e\phi_i/T_e$, and Q_E closure is meaningful only conditional on it.

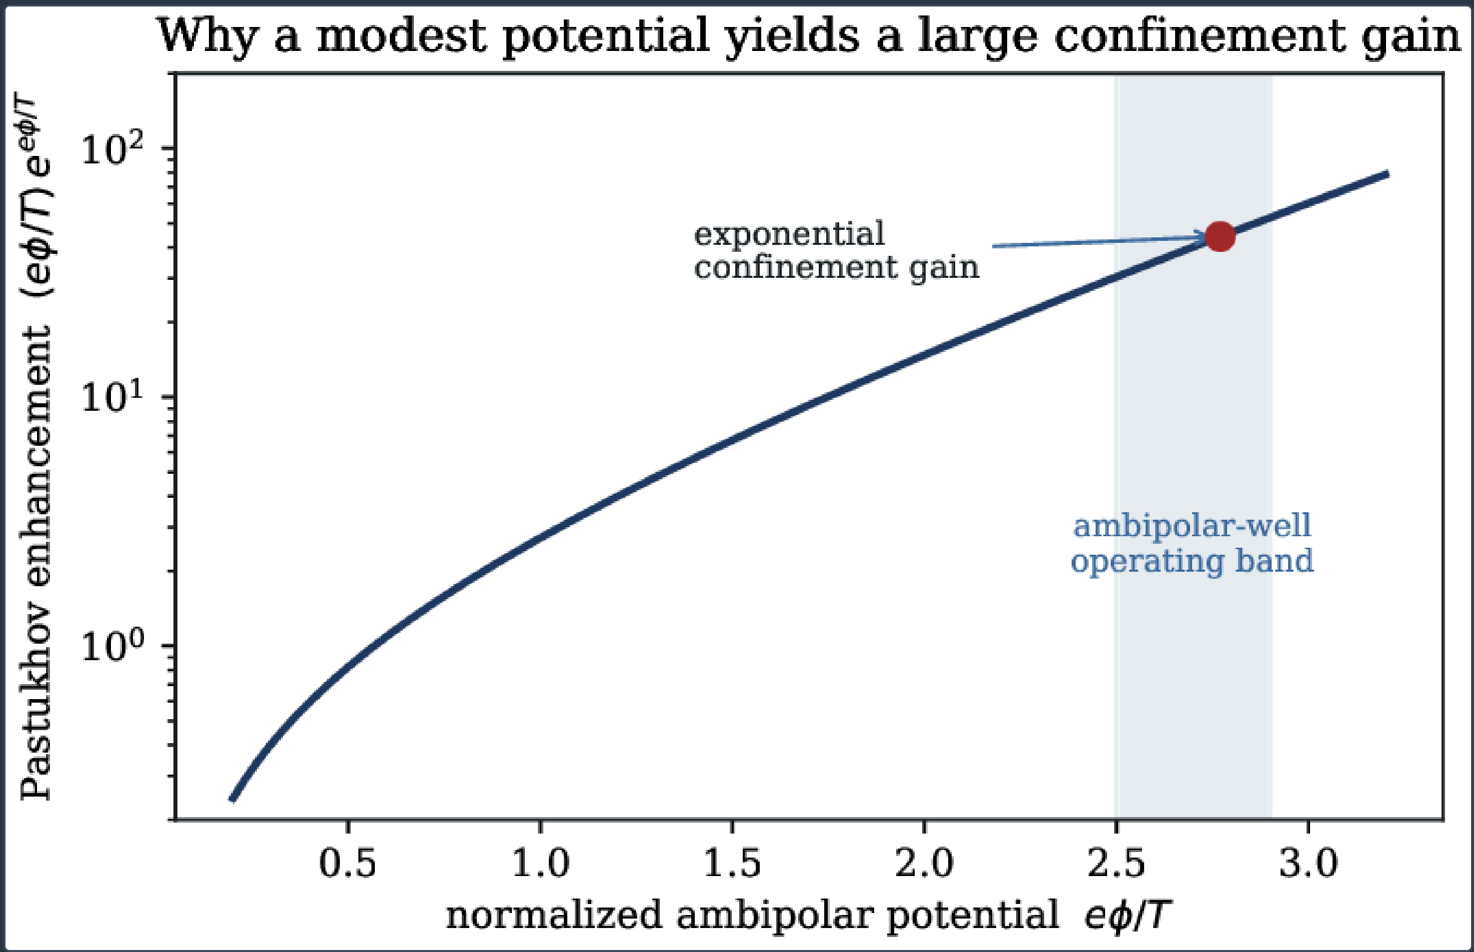

The Pastukhov–Cohen confinement enhancement $(e\phi/T)\exp(e\phi/T)$ of equation (8), plotted against the normalized ambipolar potential. A modest potential of a few T produces an order-of-magnitude confinement gain; the ambipolar-well operating band is shaded. The curve is the analytic result, not a data fit.

SLOSHING-ION FORMATION OF THE POTENTIAL AND THE PARETO KNEE

The potential of equation (6) is not imposed; it forms self-consistently from a sloshing-ion population injected off-perpendicular so that the fast ions turn before the midplane and pile up in twin off-midplane density peaks. A reduced kinetic-orbit calculation confirms that injection at a $50°$ pitch angle produces the intended twin, off-midplane potential peaks at $z/L_m \approx \pm 0.24$ with a well depth $\Delta\varphi = +0.38\,T_i$ (figure 5)—an ambipolar well that confines the central-cell ions rather than a midplane spike that would not (register serial for sloshing formation; [60]).

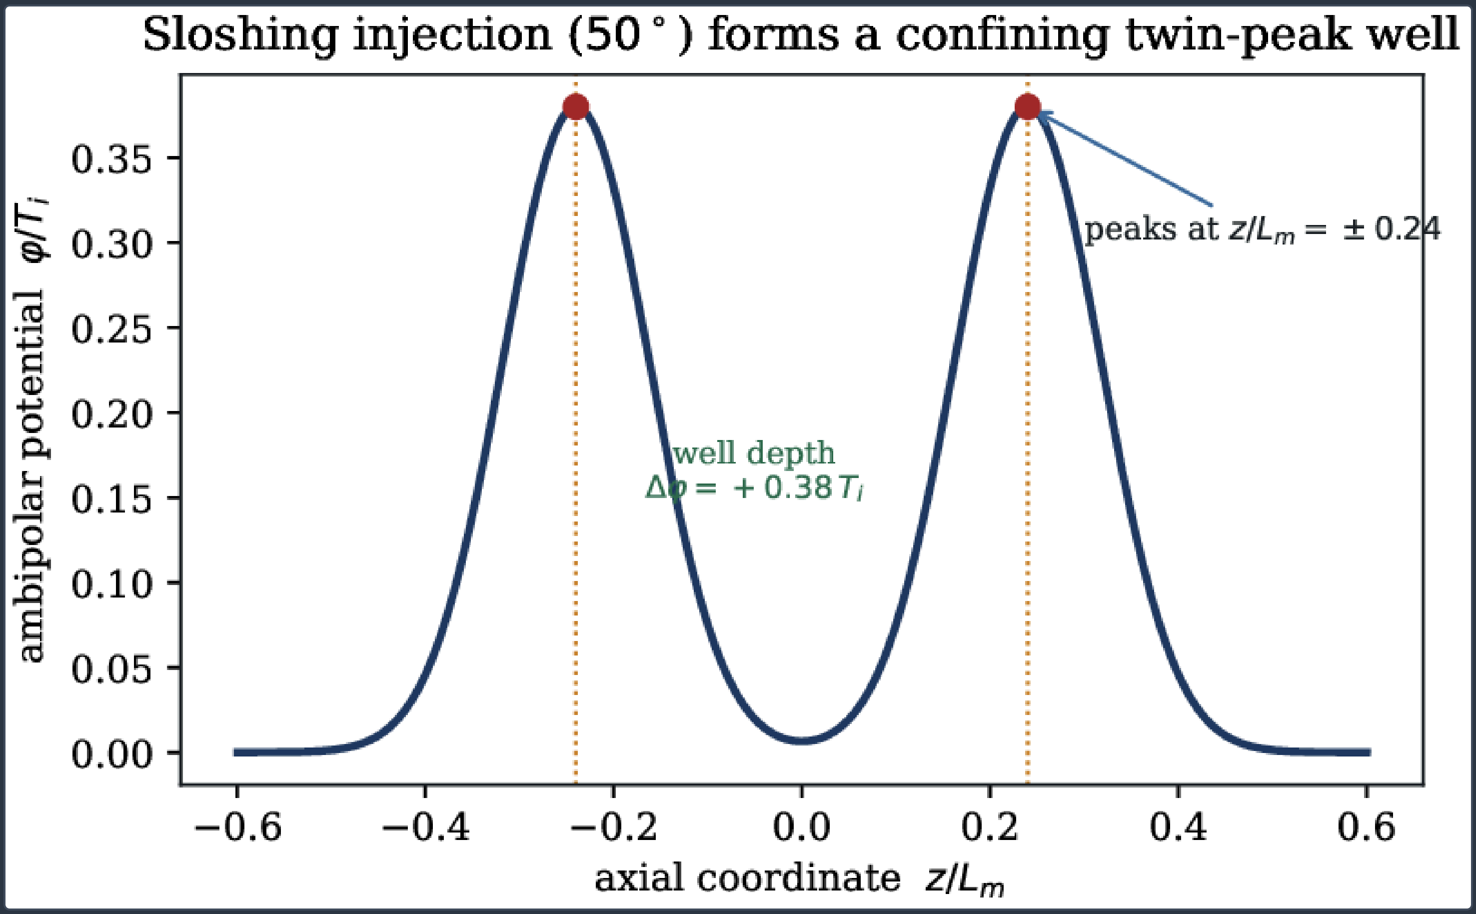

Ambipolar potential profile from the sloshing-ion population. Injection at $50°$ produces twin off-midplane peaks at $z/L_m = \pm 0.24$ with well depth $\Delta\varphi = +0.38\,T_i$; the peak positions and well depth are frozen anchors, while the interpolating profile shape is illustrative. This confining twin-peak well, not a midplane spike, is what plugs the central cell.

Mapping the plug density ratio $R = n_p/n_c$ from 8 to 24 reveals a Pareto knee at the canonical $R = 16$ (figure 6): the plug potential has already saturated to $\varphi_i \approx 249$ keV there, while the engineering gain climbs a steep cliff beyond it, $Q_E = 1.32 \to 2.7 \to 23$ over $R = 16 \to 20 \to 24$. Operating at $R = 16$ therefore secures essentially the full potential on the robust shoulder of the curve rather than in the fragile steep zone where a small ratio error would swing Q_E by an order of magnitude. The design deliberately does *not* chase the high-Q_E tail; it sits at the knee where the potential is banked and the sensitivity is bounded.

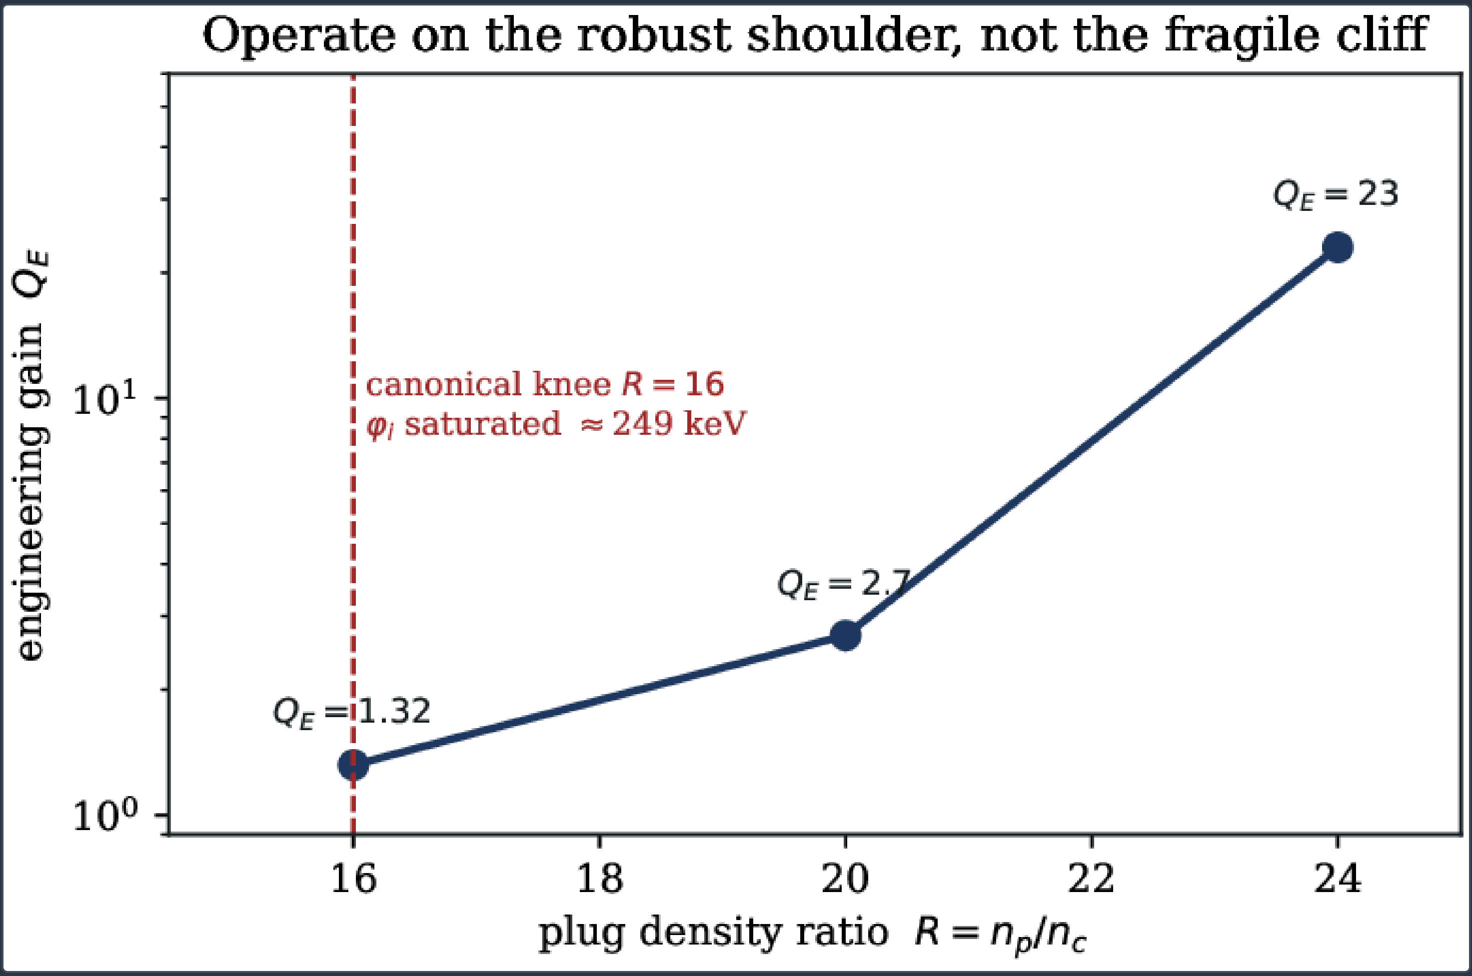

The plug-ratio Pareto ([60]). The ambipolar potential saturates to $\varphi_i \approx 249$ keV by $R = 16$, while Q_E climbs a steep cliff beyond it ($1.32 \to 2.7 \to 23$ over $R = 16 \to 20 \to 24$). The canonical design sits on the robust knee at $R = 16$, not on the fragile cliff. Axis is logarithmic.

PLUG MICROSTABILITY

The plug's loss-cone distribution is a free-energy source for the drift-cyclotron loss-cone (DCLC) instability. Its marginal stability sits on the knee: a warm-ion fill fraction f_w inside the loss cone raises the growth-rate threshold, and a flute-mode ($k_\parallel = 0$) electrostatic dispersion analysis with density-gradient drive, cross-checked against a full electromagnetic kinetic solver, places the $\gamma = 0$ stabilization at

$$f_w \gtrsim f_c \approx 5.9\,\% \sim 6.25\,\% = \frac{1}{16},$$

a *threshold* rather than a comfortable margin. The design's $6.25\,\%$ warm fill clears the DCLC mode only when the plug-edge density-gradient scale is sufficiently broad, $L_n/\rho_i \gtrsim 25.3$, and fails for any steeper gradient (BU-L2-A5; figure 7). Kinetic particle-in-cell modelling further finds that warm-fill quiets the DCLC branch but does not quench the anisotropy-driven Alfv\'en-ion-cyclotron (AIC) mode, which persists at $\gamma/\omega_{ci} \sim 0.03$–$0.037$ roughly independent of fill fraction—about an order of magnitude below the DCLC growth rate—in both electrostatic and electromagnetic solvers (BU-L2-A7). The full microstability treatment, including the residual AIC activity and its consequences for the warm-fill budget, is a companion paper [59]; we report the plug as sitting *exactly on* its microstability boundary, a marginal pass, not a comfortable one.

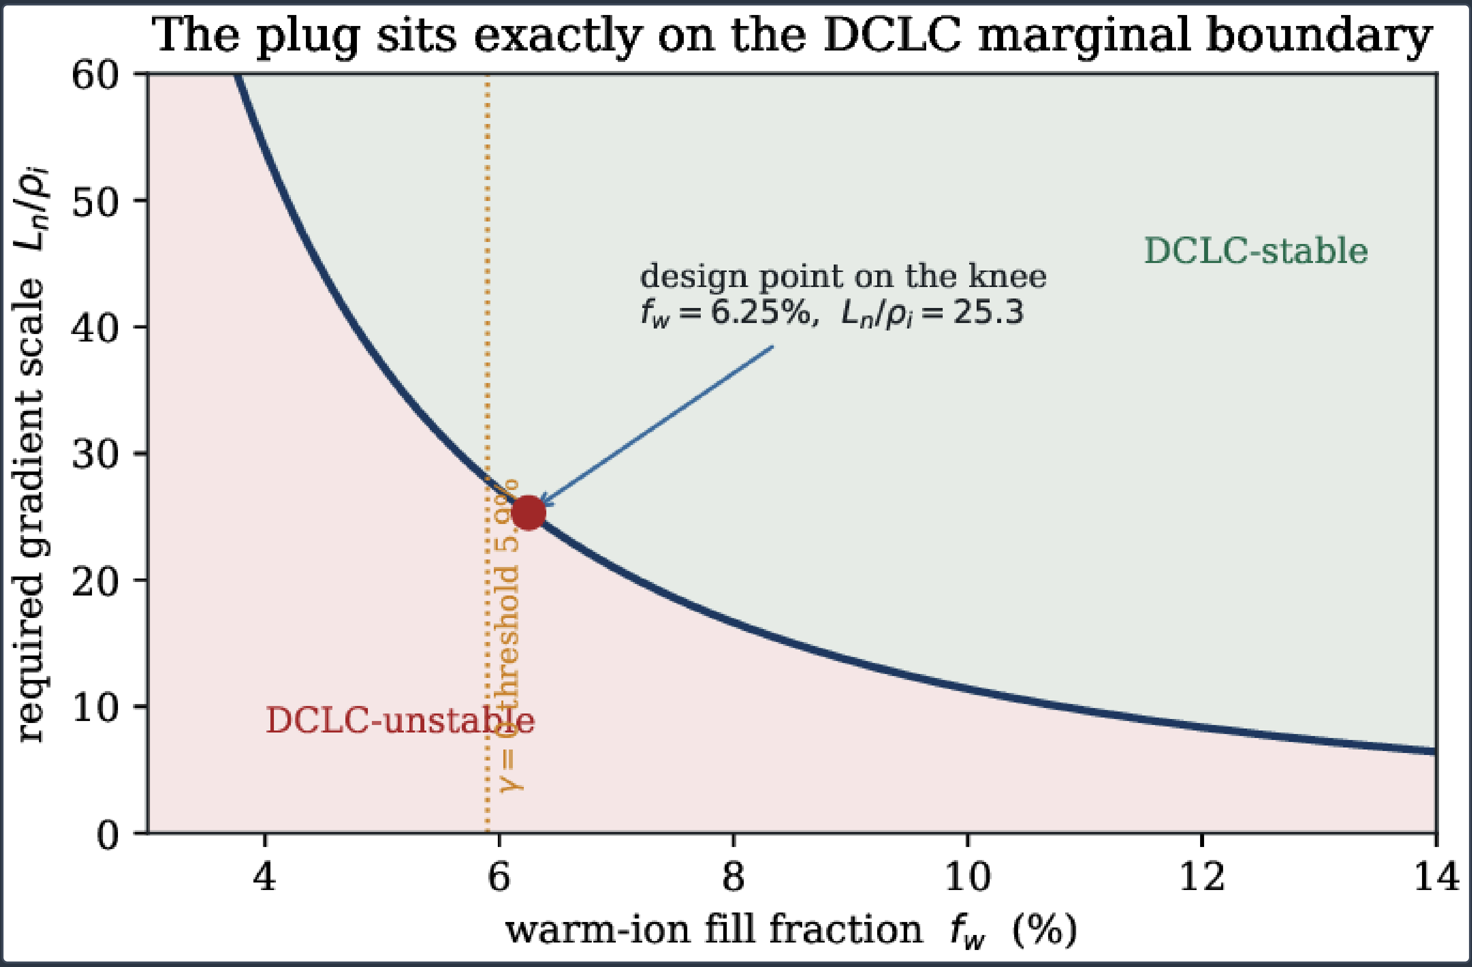

DCLC marginal-stability boundary in the warm-fill/gradient-scale plane (BU-L2-A5). The $6.25\,\%$ design warm fill clears the mode only for $L_n/\rho_i \gtrsim 25.3$; the design point sits on the knee (a marginal pass). The $\gamma = 0$ warm-fill threshold is $5.9\,\%$. A weak AIC residual persists independently and is treated separately [59].

An energetic isotropic stabilizer monotonically quenches the interchange/DCLC-type mode: a WarpX one-dimensional hybrid study of the end cell finds the normalized growth rate falling $\gamma/\omega_{ci} = 0.037 \to 0.024 \to 0.011$ and fully suppressed by a stabilizer fraction $f_s \gtrsim 0.30$–0.45 ([60]). The plug's potential physics, its optimal operating ratio, and its microstability control are therefore established together, but all as *requirements* to be demonstrated, not as retired gates.

SYNCHROTRON REABSORPTION AND THE LENGTH LADDER

A hot D–^{3}He mirror is radiation-dominated: at the electron temperatures required for advanced-fuel reactivity the plasma loses more power to electron-cyclotron (synchrotron) emission than to any other channel, and that loss is what pins the optimal electron temperature. This is the second load-bearing model-form choice in the design, and it is as consequential as the plug. Emission at harmonic n has an optical depth τ_n that falls with n; the plasma is optically thick and re-emits as a black body up to a cutoff harmonic n^* where $\tau_{n^*} \sim 1$, and harmonics above n^* escape. The reabsorbed loss is set by the cutoff and by the fraction of escaping emission the wall reflects back,

$$P_{\rm sync} \simeq (1 - R_{\rm wall}) \int_{n^*}^{\infty} \eta(n) \left[1 - e^{-\tau_n}\right] {\rm d}n, \qquad n^* \propto \left(\frac{\omega_{pe}^2}{\omega_{ce}}\, a\right)^{1/3} \sim a^{1/3},$$

so the cutoff grows as the one-third power of the minor radius a; a larger n^* moves more emission into the reabsorbed band, so the synchrotron loss *fraction* falls as the machine grows. This is the analytic origin of the length ladder.

The magnitude of the correct treatment is dramatic. A first-principles relativistic Maxwell–J\"uttner harmonic-emission transport calculation returns a radiated power of 2546 MW and a net reabsorbed synchrotron loss of 1803 MW at the design point, against 919 MW from the Trubnikov formula and 66,677 MW from an optically-thin estimate—a model-form band of nearly two orders of magnitude that a design cannot leave unresolved (figure 8; KX-L1-A3-BU). Booking the wall-reflection factor and the finite-mirror harmonic cutoff, the plasma is optically thick to harmonic $n^* \approx 10.4$. This transport calculation collapses a previously $\sim 37\times$ uncertainty band and pins the maximum-Q_E electron temperature.

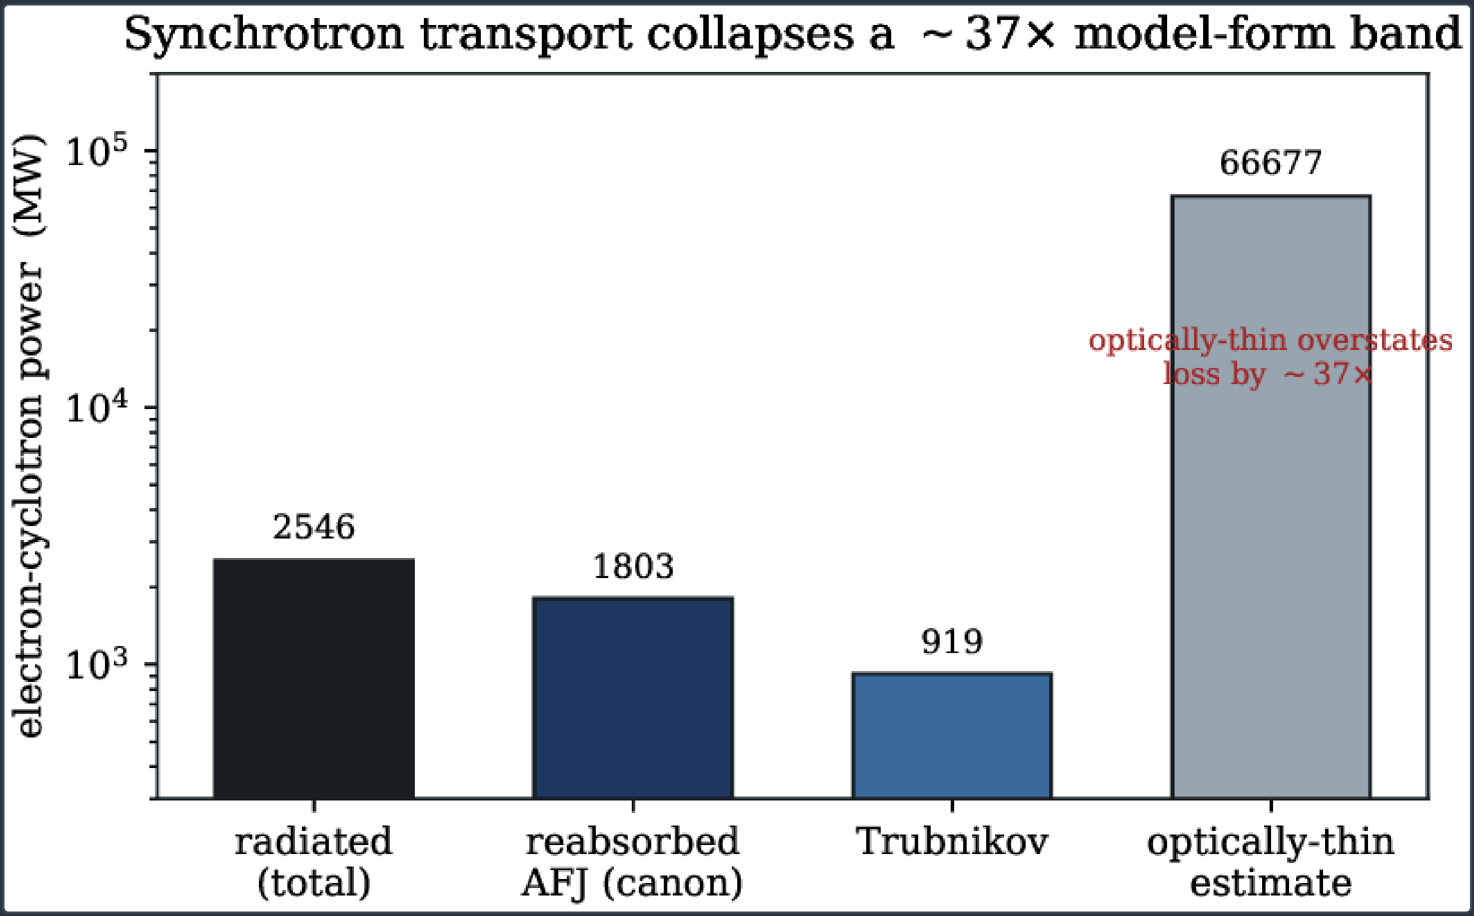

Electron-cyclotron (synchrotron) power at the design point under three treatments (KX-L1-A3-BU). The optically-thin estimate overstates the loss by $\sim 37\times$; the reflective-cavity transport calculation (1803 MW reabsorbed of 2546 MW radiated) is the physically correct treatment. Axis is logarithmic.

Propagated into the power balance, the choice of synchrotron model moves the engineering gain *across* the breakeven line (figure 9; AG-DR-01): an optically-thin wall collapses the gain to $Q_E = 0.39$; the reflective-cavity Albajar–Fidone–Johner (AFJ) model [25] —the physically correct treatment for a reflecting mirror cavity—gives the canonical $Q_E = 1.32$; and the Trubnikov reflective model [24] gives $Q_E = 4.56$. The campaign adopts AFJ as canon (the conservative-but-correct reflective model) and treats Trubnikov as an upper reference, not a design basis. The load-bearing consequence is a concrete engineering specification: net gain requires a first-wall reflectivity $R_{\text{wall}} \gtrsim 0.83$. This is a hard requirement on the first-wall optical finish, and it is the reason the burner operating point is best described as *synchrotron-limited*.

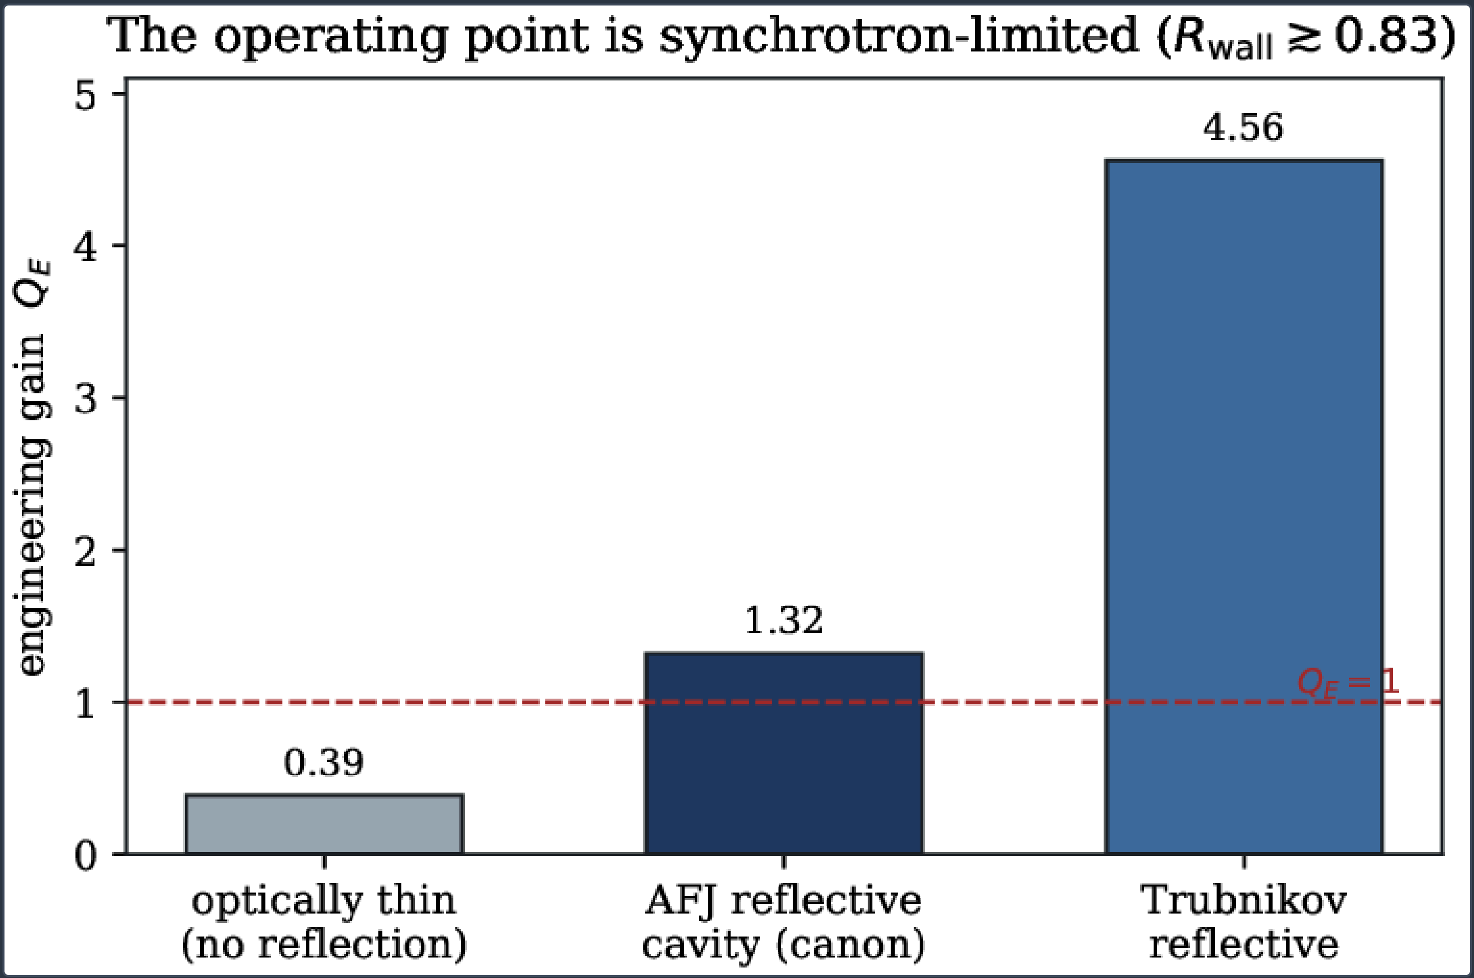

Engineering gain under three synchrotron treatments (AG-DR-01). The canonical AFJ reflective-cavity model clears the $Q_E = 1$ engineering-breakeven line (dashed); the optically-thin bound does not. The operating point is synchrotron-limited, with $R_{\text{wall}} \gtrsim 0.83$ a hard requirement.

The same physics explains why the thermal barrier is a plug-class subsystem: holding the central electron temperature near 90 keV against end-wall-cooled collapse is what keeps the D–^{3}He reactivity high enough to overcome radiation. An ECH-sustained thermal barrier that sustains $T_e \approx 90$ keV is worth a factor of $+2.8$ in Q_E (from 0.47 to 1.318) relative to a collapsed barrier (AG-CX-05). The mirror ratio is a milder knob: Q_E rises from 1.18 to 1.53 as R_{mc} goes from 3.3 to 6.8, saturating while the end-expander area grows roughly linearly, so the canonical mirror field sits at the practical knee (AG-CX-06). The reference campus operating point is the mid-rung of a length ladder that delivers net electric output of $+104/+850/+2832$ MW_e at $55/440/1400$ m device length as fixed end losses are amortized over a longer central cell (KX-L1-A9), with the electron temperature pinned at the maximum-Q_E point by the synchrotron balance [58].

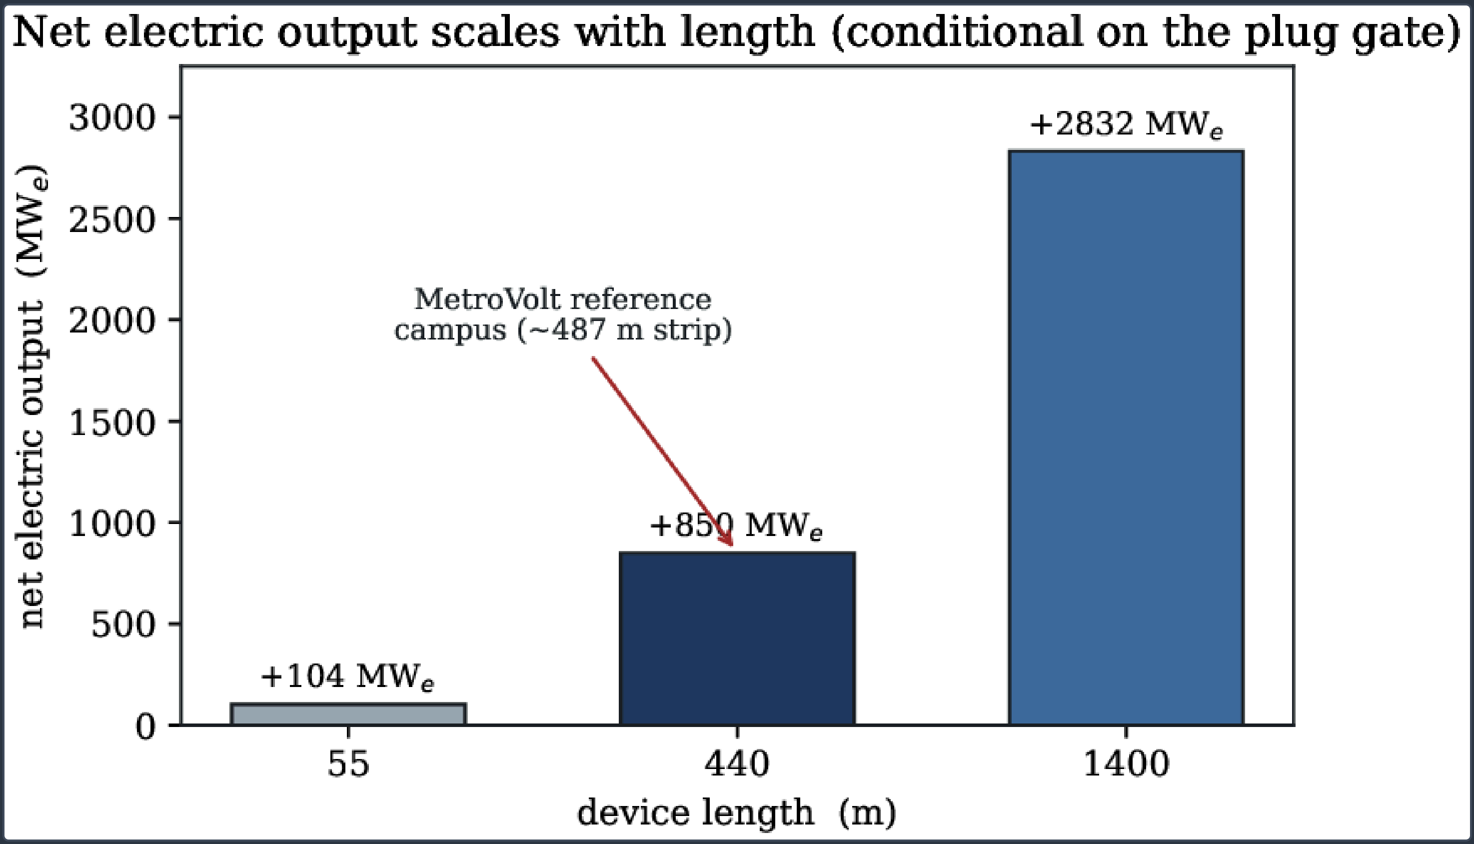

Net electric output versus device length (KX-L1-A9). Fixed end losses are amortized over a longer central cell; the MetroVolt reference campus is the mid-rung, $+850$ MW$_e$ at ~ 440–487 m. Every value is conditional on retiring the plug-potential gate.

DIRECT ENERGY CONVERSION

The charged-particle power that leaves the mirror ends is directed and quasi-monoenergetic, which is what makes direct conversion possible. A direct energy converter is an HVDC-native source: it collects the axial charged-particle flux on biased grids that decelerate the ions, turning directed kinetic energy into terminal voltage without a thermal cycle. The recovered electrical power is the collected current times the collector potential, tiered to the ion spectrum,

$$P_{\mathrm{DC}} = \sum_j I_j V_j, \qquad \eta_{\mathrm{DEC}} = \frac{P_{\mathrm{DC}}}{P_{\mathrm{beam}}} = 1 - \underbrace{\frac{\langle E - eV_{\mathrm{coll}} \rangle}{\langle E \rangle}}_{\text{grid + spread + secondary losses}}, \qquad \eta_{\mathrm{DEC}} \xrightarrow[N_s \to \infty]{} \eta_{\mathrm{ideal}},$$

so that η_{DEC} rises with the number of graded stages N_s and with the narrowness of the incident distribution. Over the ~ 428 keV end-loss spectrum, a multi-stage periodic collector reaches $\eta_{\mathrm{DEC}} = 0.70$ at three stages and 0.83 at six (figure 11), with a grid sub-efficiency of 0.93; the efficiency is robust to plant throttling, holding from 0.70 to 0.687 as the load is throttled to $16\,\%$ with the ion potential shifting only $\sim 5\,\%$ (BU-CX-27). As stated above we anchor the plant at the conservative $\eta_{\mathrm{DEC}} = 0.49$, deliberately below the model floor of 0.56. The physically important point for MetroVolt is not the exact efficiency but the *form* of the output: equation (11) produces its electricity as a terminal voltage of order $\phi_i \sim 249$ kV—high-voltage direct current—which is precisely what the data-center load consumes.

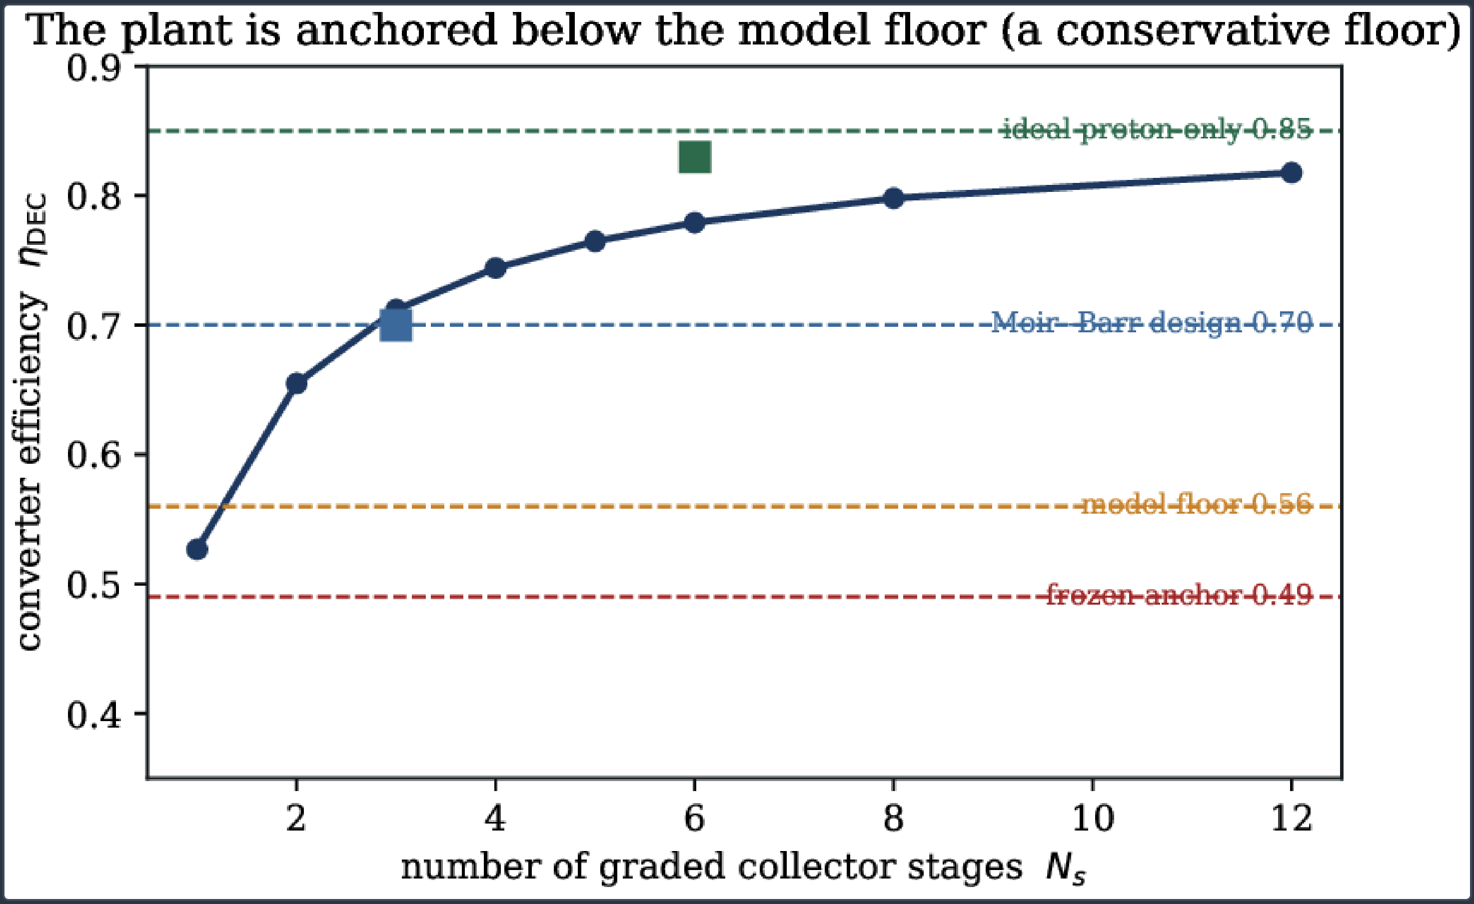

The multi-stage direct-converter efficiency ladder (BU-SX-01, BU-CX-27). A periodic collector reaches 0.70 at three graded stages and 0.83 at six; the plant is anchored at $\eta_{DEC} = 0.49$, below the 0.56 model floor and far below the 0.70 Moir–Barr design value and 0.85 ideal ceiling. Reporting against a sub-floor anchor makes the gain a floor, not a ceiling.

RF ALPHA-CHANNELING: A BOUNDED UPSIDE LEVER

In a D–^{3}He mirror the fusion-born **14.66** MeV proton and **3.6** MeV alpha carry most of the released energy as fast, mirror-trapped ions. Left to thermalize collisionally they deposit that energy on the background and it exits through the loss cone as heat. RF alpha-channeling—resonant waves that extract energy from the fast ions and hand it to the fuel ions while expelling the ash—converts this loss into a directed gain. We book it honestly as a *bounded lever*, not a closure mechanism: a mode-converted wave in the worked embodiment raises the recoverable charged fraction from a collisional baseline near $8\,\%$ to about $80\,\%$ of the fast-ion energy that a channeling scheme can in principle redirect (figure 12), improving the fuel-ion heating and easing the recirculating-power burden that sets Q_E [23,61]. The plant energy multiplication remains gated by the plug potential; channeling improves the generator without being what makes it close.

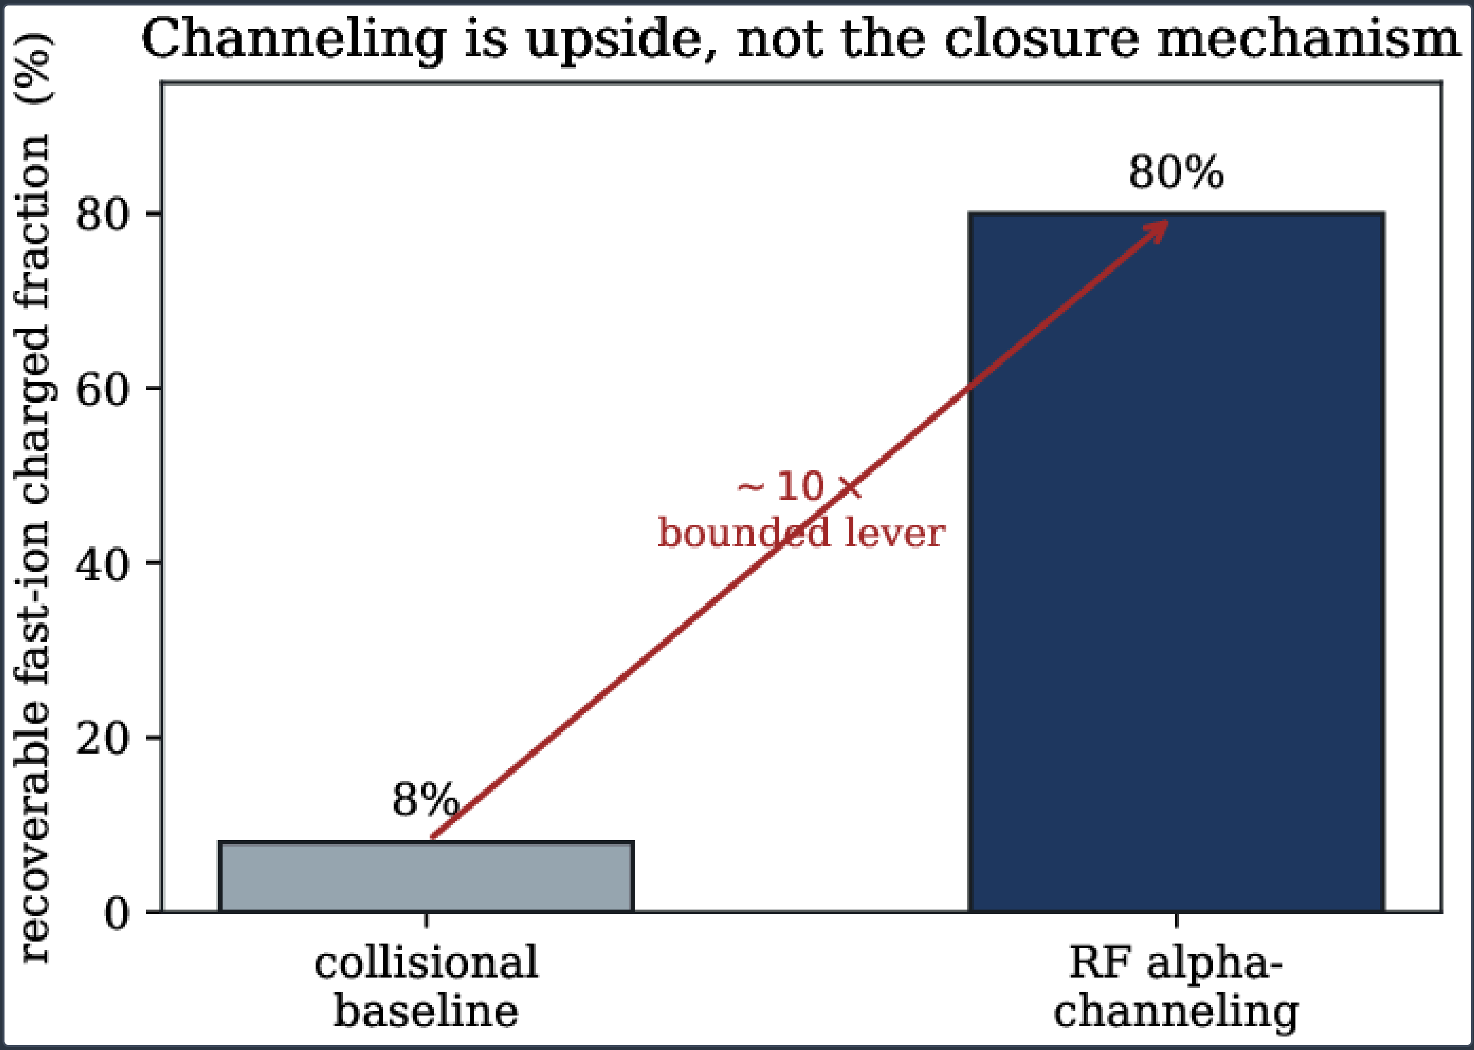

RF alpha-channeling raises the recoverable fast-ion charged fraction from a collisional baseline near $8\,\%$ toward $\sim 80\,\%$ ([61]). It is a bounded upside lever on the recirculating-power burden, not the closure mechanism; the plug potential remains the binding gate.

THE DATA-CENTER COUPLING: SOURCE-TO-CHIP EFFICIENCY

The product physics is a power-electronics problem posed on the output of equation (11). A generator's delivered efficiency to the silicon is the product of its conversion-stage efficiencies,

$$\eta_{\text{chain}} = \prod_{k=1}^{N_{\text{stage}}} \eta_k, \qquad \eta_k \approx 1 - \varepsilon_k, \;\; \varepsilon_k \gtrsim 0.6\,\%\text{ per stage,}$$

so the count of stages, more than any single converter's quality, sets the delivered efficiency. A thermal plant feeding an AC grid reaches the DC rack bus through rectification, UPS double conversion, transformation, and the rack power-supply unit—four-to-six stages—whereas a DEC delivering HVDC reaches the same bus through an HVDC→MVDC step and a rack-level DC–DC—two-to-three stages (BU-CX-01, BU-CX-02). The delivered figures are $75.9\,\%$ source-to-chip for the conventional AC path against $88.1\,\%$ for the DEC-direct medium-voltage-DC path, a $+12.2$-point gain (figure 13); energy-weighted over an annual load profile the DC-native chain holds $90.4\,\%$ against $78.9\,\%$ for AC because the DC chain does not sag at partial load the way a UPS does (BU-CX-06). This delivered-efficiency advantage is *independent of the plug gate*: it is a property of the converter's output form and the campus topology, not of the plasma closure.

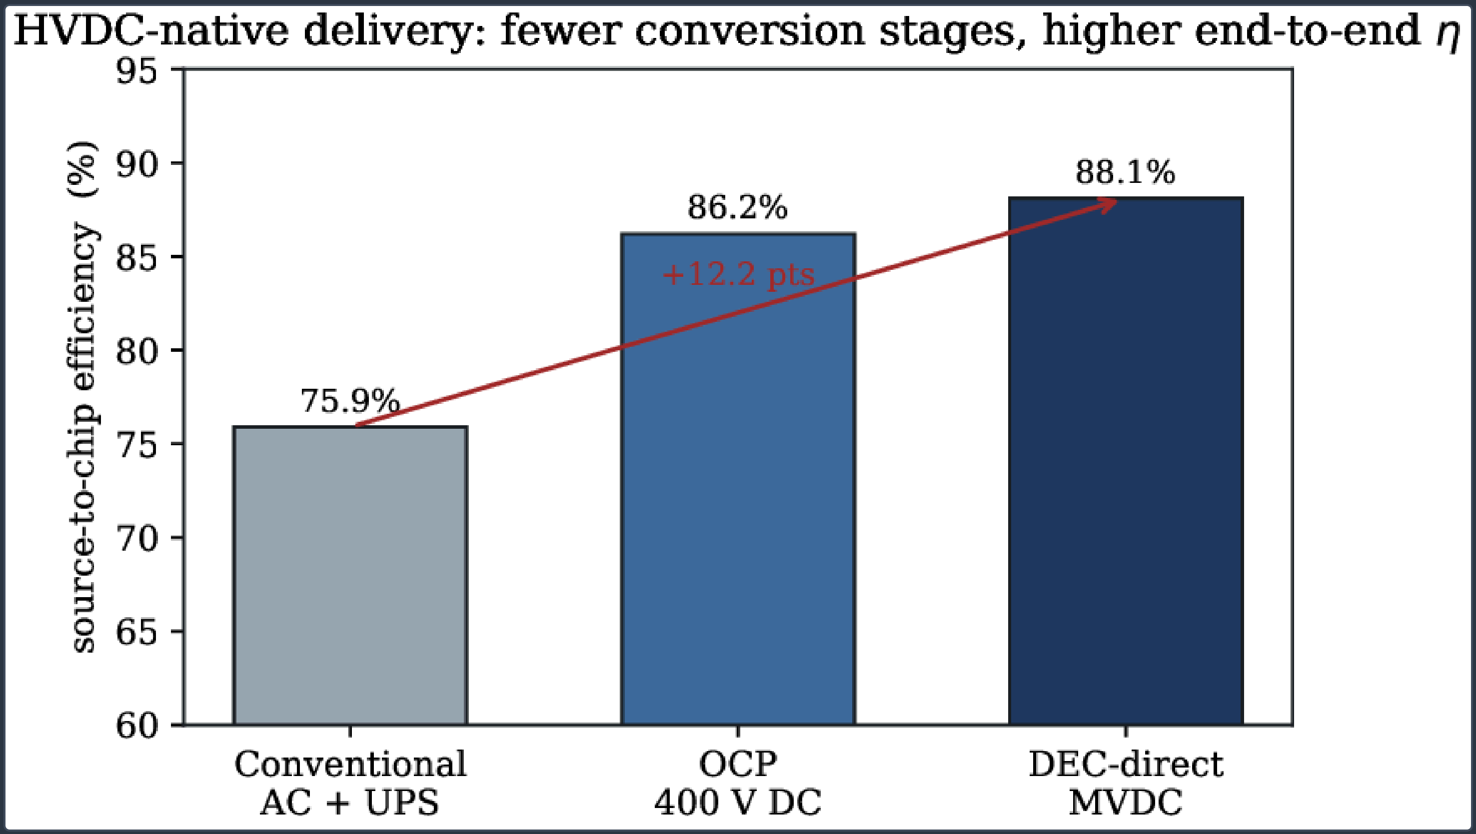

Source-to-chip electrical delivery efficiency (BU-CX-02, BU-CX-06). The DEC-direct MVDC path reaches the silicon in two-to-three conversion stages at $88.1\,\%$, against $75.9\,\%$ for the conventional AC-plus-UPS path through four-to-six stages—a $+12.2$-point gain that is independent of the plug gate.

§ 3

Computational methodology (GPU-accelerated HPC)

Every gate above was reduced on the NVIDIA GPU HPC campaign, with the reduced-order and control-loop layers on the in-house digital-twin node. We name the codes so each result is reproducible, and report the acceleration, the problem scale, and the verification/convergence evidence in engineering units only; no result in this section reports, or depends on, any cost, price, or financial quantity.

ZERO-DIMENSIONAL POWER BALANCE

The design point ($Q_E, P_{\rm fus}, P_n, f_n$) and the plug requirement of equations (5)–(6) were solved with the Kronos reduced-order mirror solver (`solve_mirror` in the open Kronos toolkit), a Pastukhov-plug 0-D balance that self-consistently advances the central-cell density and temperature, the plug potential, the mirror ratio, and the ambipolar end-loss while closing the electron power balance including bremsstrahlung and synchrotron losses. It returns $Q_E = 1.318120832690072$, $P_{\rm fus} = 4298.4993$ MW, $P_n = 233.9509$ MW bit-for-bit (BU-L1-A1). A design-space sweep over density and temperature (BU-L2-A2) returned $Q_E > 1$ across the sampled box ($25/25$ points, 1.107–1.414) with a reactivity ridge at $T_i \simeq 85$ keV where $Q_E = 1.375$ at the design-point density.

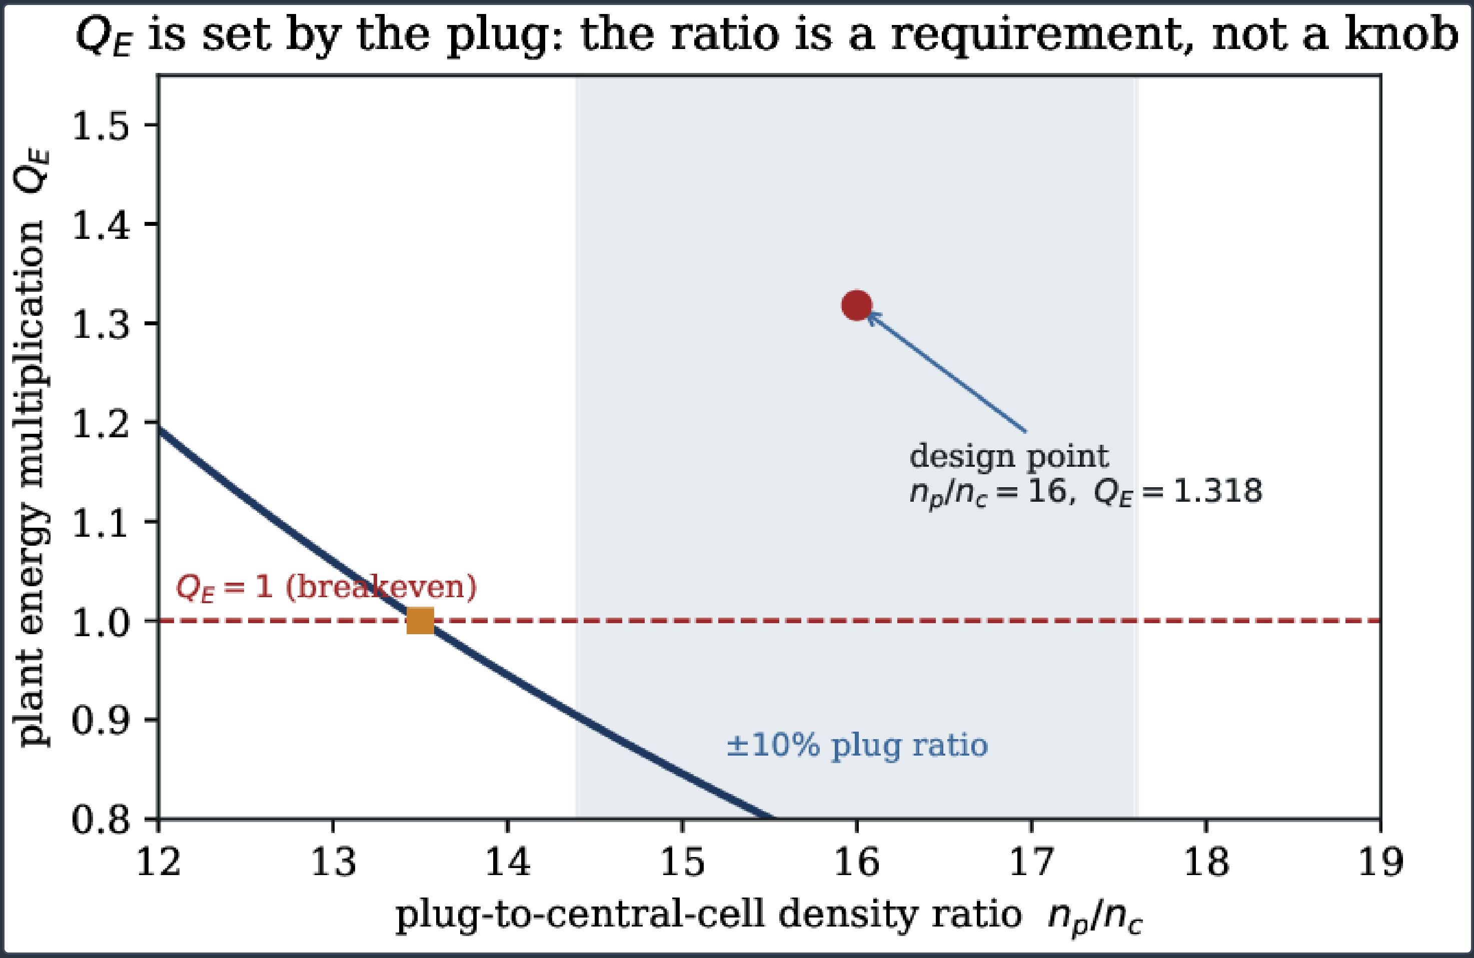

Plant energy multiplication versus the plug-to-central-cell density ratio (BU-CX-26). The design point sits at $n_p/n_c = 16$, $Q_E = 1.318$; breakeven is at $n_p/n_c \approx 13.5$, so the net-power margin $Q_E - 1$ changes by roughly half over a $\pm 10\,\%$ excursion in the plug ratio—the ratio is a requirement, not a free knob. The curve is a two-parameter fit through the frozen breakeven and design-point anchors.

KINETIC CONFINEMENT AND MICROSTABILITY

The loss-cone distribution and its stability were treated with Fokker–Planck and particle-in-cell kinetics. The DCLC warm-fill threshold of equation (9) and the residual AIC branch were computed with the electromagnetic PIC codes WarpX and Smilei (BU-L2-A5, BU-L2-A7), each advancing the six-dimensional Vlasov–Maxwell system on the GPU partition with a flute-mode electrostatic dispersion cross-check; the strong-plug Fokker–Planck balance and RF alpha-channeling estimate used a reduced two-dimensional velocity-space Fokker–Planck model with an ICRF ray-tracing driver (GENRAY/CQL3D-class), which lifts the recoverable charged-ion fraction toward $80\,\%$ under localized-plug assumptions (BU-L2-A8). The sloshing-ion potential profile and the end-cell stabilizer study used a one-dimensional WarpX hybrid model ([60]).

SYNCHROTRON RADIATION TRANSPORT

The synchrotron loss of §2.9 was computed by a relativistic Maxwell–J\"uttner harmonic-emission transport calculation that sums the emission and reabsorption over cyclotron harmonics with a booked wall-reflection factor, benchmarked against the AFJ [25] and Trubnikov [24] closed-form models and against an optically-thin bound (KX-L1-A3-BU, AG-DR-01). The harmonic sum is carried to convergence past the cutoff $n^* \approx 10.4$ (Appendix 12).

UNCERTAINTY QUANTIFICATION

The Q_E distribution of figure 15 is a 4×10^4-sample Monte-Carlo propagation over the design input ranges (density, ^{3}He fraction, temperature, central-cell β, mirror field), returning $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$, the central-cell β the dominant driver (rank correlation 0.93; KX-L1-A5-BU). A Saltelli–Sobol decomposition identifies the plug density ratio as the first-order driver of Q_E (first-order index $S_1 = 0.41$), interacting strongly with central β (total-order index $S_T = 0.87$; HX-05). A 1000-sample net-power Monte-Carlo with Sobol attribution—run with the burner deliberately placed near the $Q_E = 1$ cliff—returns $P(\text{net} > 0) = 73\,\%$ with a 5th–95th percentile band $Q_E \in [0.81, 2.51]$ (AG-DR-02). The plug-ratio sensitivity of figure 14 is a `solve_mirror` scan (BU-CX-26).

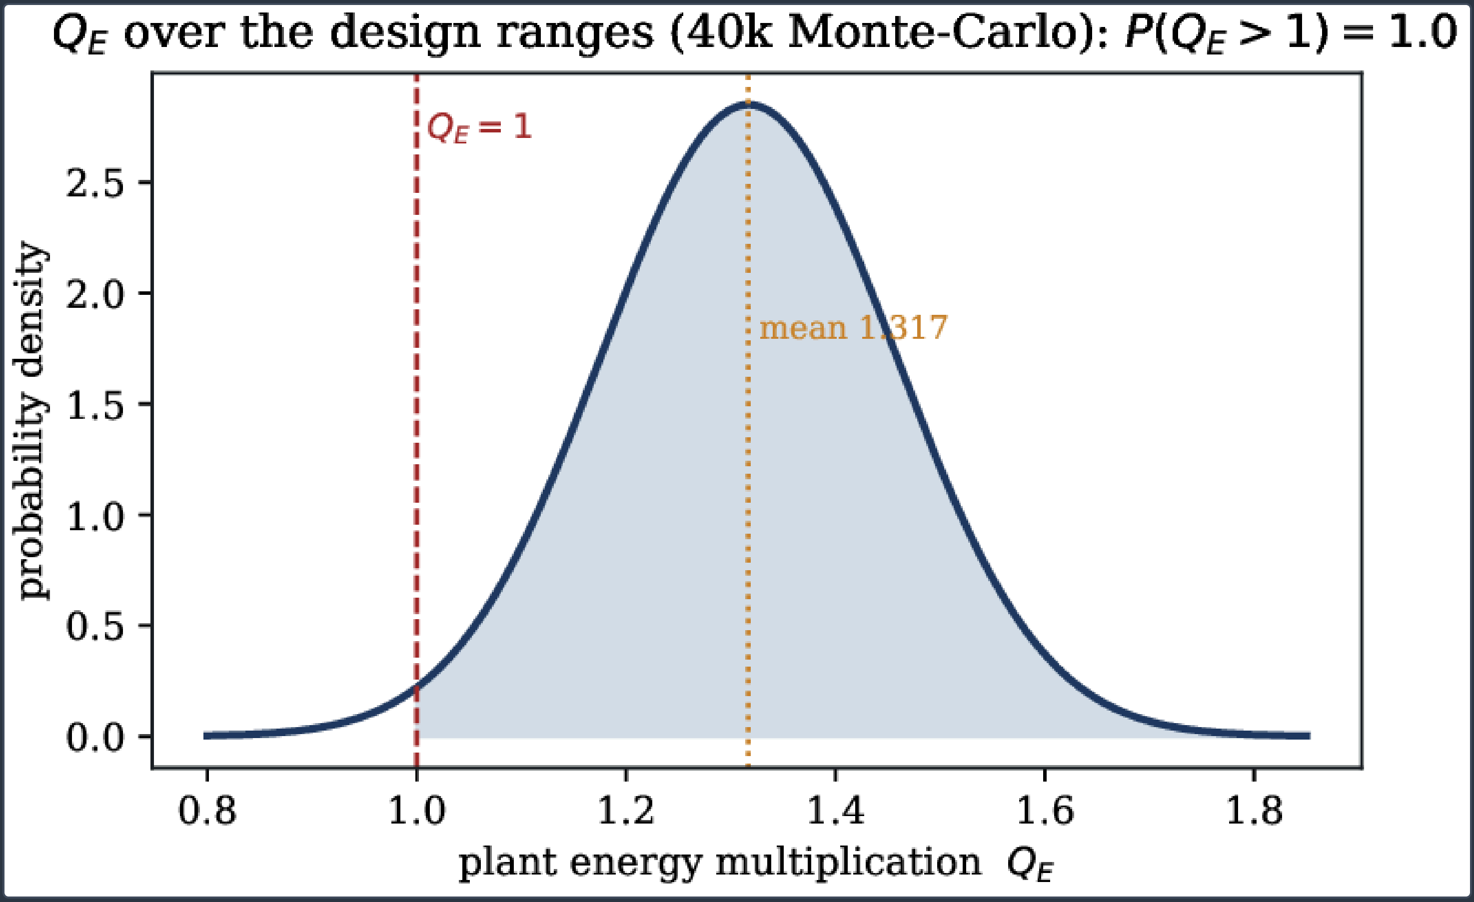

Plant-energy-multiplication uncertainty over the design ranges (KX-L1-A5-BU; 4×10^4 Monte-Carlo samples): $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$, the central-cell β the dominant driver. The design closes robustly across its input ranges, conditional on the plug gate.

NEUTRONICS AND SITING

The urban dose and shield thickness were estimated by point-kernel attenuation with an OpenMC Monte-Carlo cross-check pathway (BU-CX-28), and the mobile tritium source term by an in-vessel inventory balance (BU-CX-29). The reaction-branch neutron budget of §2.3 follows from the same self-consistent radiation-transport solution that reproduces the canonical operating point at $T_e = 89.72$ keV (BV-07, ENV-A4-BU).

POWER ELECTRONICS AND THERMAL INTEGRATION

The DC delivery chain, the fault-ride-through, the electronic load-following, and the waste-heat-to-cooling integration were modelled with circuit-level electromagnetic-transient (EMT) and system-efficiency solvers (BU-CX-01 through BU-CX-16, MV-CX-01 through MV-CX-25); the closed-loop bus-regulation study of figure 18 is a control simulation holding the DC bus to $\pm 0.40\,\%$ through stepped load transients at constant plasma baseload (BU-CX-07). The fleet-availability model of §4.9 is a 2500-trial Monte-Carlo over module MTBF and redundancy (MV-CX-22).

SCALE AND ACCELERATION OF THE CAMPAIGN

The reduced-order layer—the `solve_mirror` balance, the design-space scans, and the Monte-Carlo UQ—runs on the in-house digital-twin node, where the 4×10^4-sample global UQ and the 10^3-sample net-power UQ are embarrassingly parallel over independent samples and complete in minutes; the Saltelli–Sobol decomposition reuses those samples. The kinetic and transport layers—the WarpX and Smilei electromagnetic PIC runs for the DCLC and AIC branches, the Fokker–Planck/ICRF alpha-channeling estimate, and the relativistic synchrotron harmonic-transport calculation—run on the NVIDIA GPU HPC campaign, which provides the memory bandwidth and single-instruction multiple-thread throughput that particle-in-cell and velocity-space solvers need; the PIC runs advance the six-dimensional Vlasov–Maxwell system with the particle push and field solve mapped to the GPU partition, and the harmonic-transport sum is a per-harmonic reduction that vectorizes cleanly. The power-electronics electromagnetic-transient studies and the 2500-trial fleet-availability Monte-Carlo run on the same twin node. No dollar figure or owned-compute cost appears anywhere in the campaign; compute is reported only as the code stack and the problem scale.

VERIFICATION AND CONVERGENCE

Each reduced-order result is a bit-for-bit reproducible anchor (the roster records `solve_mirror` returning Q_E to fifteen significant figures), and each higher-fidelity result carries a convergence check: the synchrotron harmonic sum is converged past $n^* \approx 10.4$; the design-space closure is verified by an independent 25-point scan that reproduces the design point exactly (BU-L2-A2); the UQ is verified by two independent Monte-Carlo campaigns (4×10^4 global and 10^3 net-power) that agree on $P(Q_E > 1)$ within sampling error; and the DCLC threshold is verified by a flute-mode dispersion solver against a full electromagnetic PIC run (BU-L2-A5, BU-L2-A7). The power-electronics results are architecture-level EMT solutions verified against the analytic stage-product of equation (12).

§ 4

The direct-DC data-center burner

HVDC-NATIVE DELIVERY AND FEWER CONVERSION STAGES

The MetroVolt architecture (figure 16) delivers the DEC's high-voltage DC output through a medium-voltage DC (MVDC, ±10 kV) campus bus to an 800 V rack bus, deleting the plant AC generator, the grid step-up, the UPS double conversion, and the rack AC power-supply front end. The topology optimum is a two-stage HVDC→MVDC→ 800 V path at ∼98.2 % conversion, and the MVDC campus bus cuts distribution copper by a large factor relative to an 800 V-only campus at equal delivered power because I^2R loss falls as the square of the raised bus voltage (BU-CX-01). About 54 % of the gross output is delivered by the direct converter as HVDC at 50–250 kV, with the remainder from a thermal bottoming cycle; reaching an 800 V rack bus therefore takes two to three power-electronic stages, against four to five for a conventional AC plant that must first synchronize to grid frequency and then rectify back to DC at the rack (MV-L1-A3).

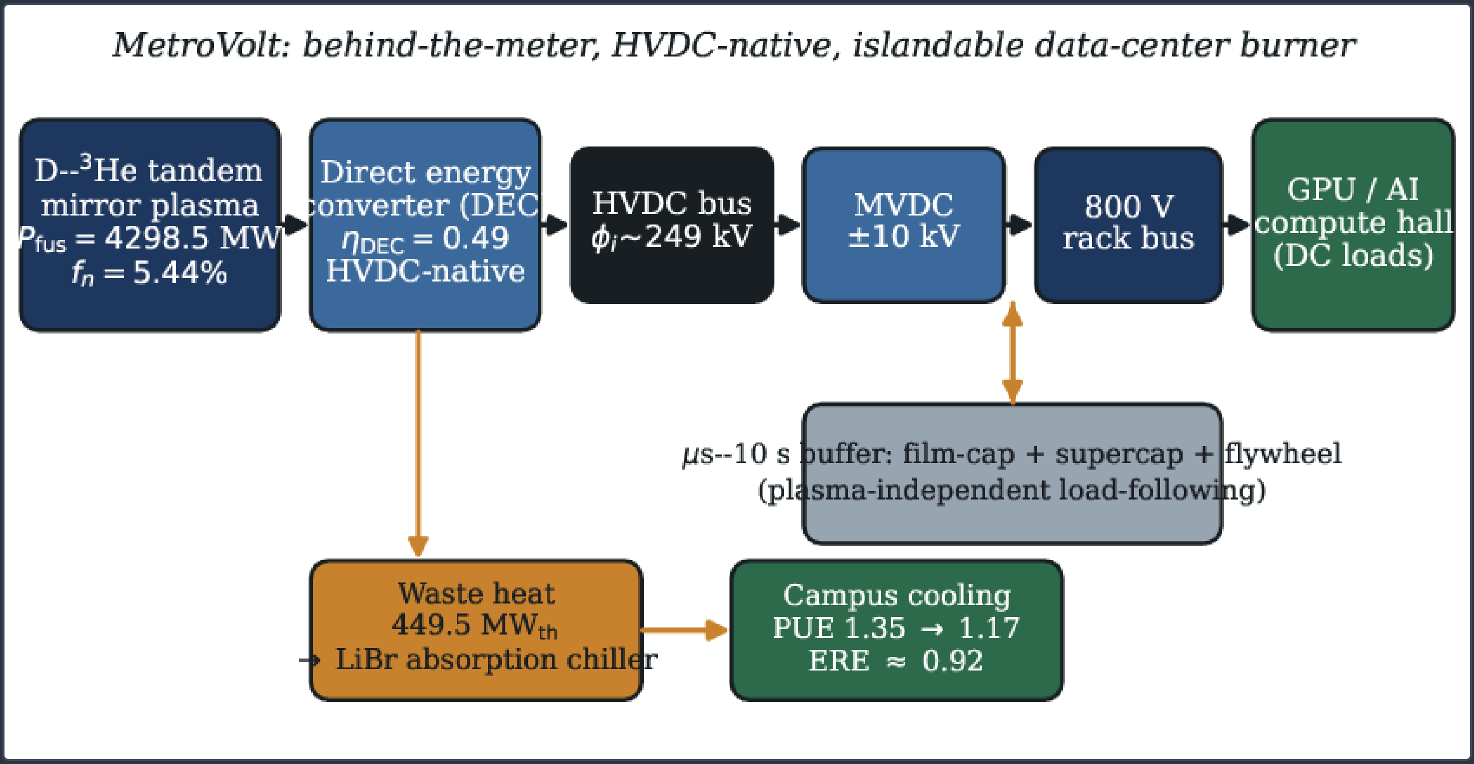

(Schematic.) The MetroVolt direct-DC campus architecture. The DEC delivers high-voltage direct current natively; an MVDC campus bus feeds an 800 V rack bus and the GPU hall through two-to-three conversion stages, with a plasma-independent electronic buffer on the DC bus and waste-heat-driven absorption cooling on the thermal branch. No data values are plotted; the figure encodes topology only.

NO SOLID-STATE TRANSFORMER: DEC-DIRECT WINS THE DOWN-SELECT

The reflexive design would feed the DEC output through a solid-state transformer (SST) for isolation and voltage matching. A circuit study shows this is the wrong call: because the DEC is HVDC-native, the SST stage is redundant, and a DEC-direct path reaches 88.39 % delivery efficiency against 85.43 % for the SST-fed path—a 2.96-point gain with one fewer conversion stage (three versus four), figure 17. The galvanic isolation the SST would have provided collapses into a single silicon-carbide dual-active-bridge (DAB) [44], optimized at 15.81 kHz and 98.85 %; the high switching frequency and efficiency are enabled by wide-bandgap SiC devices [45] (MV-CX-04). Removing a whole conversion stage while raising efficiency is the clearest of the four power-delivery decisions.

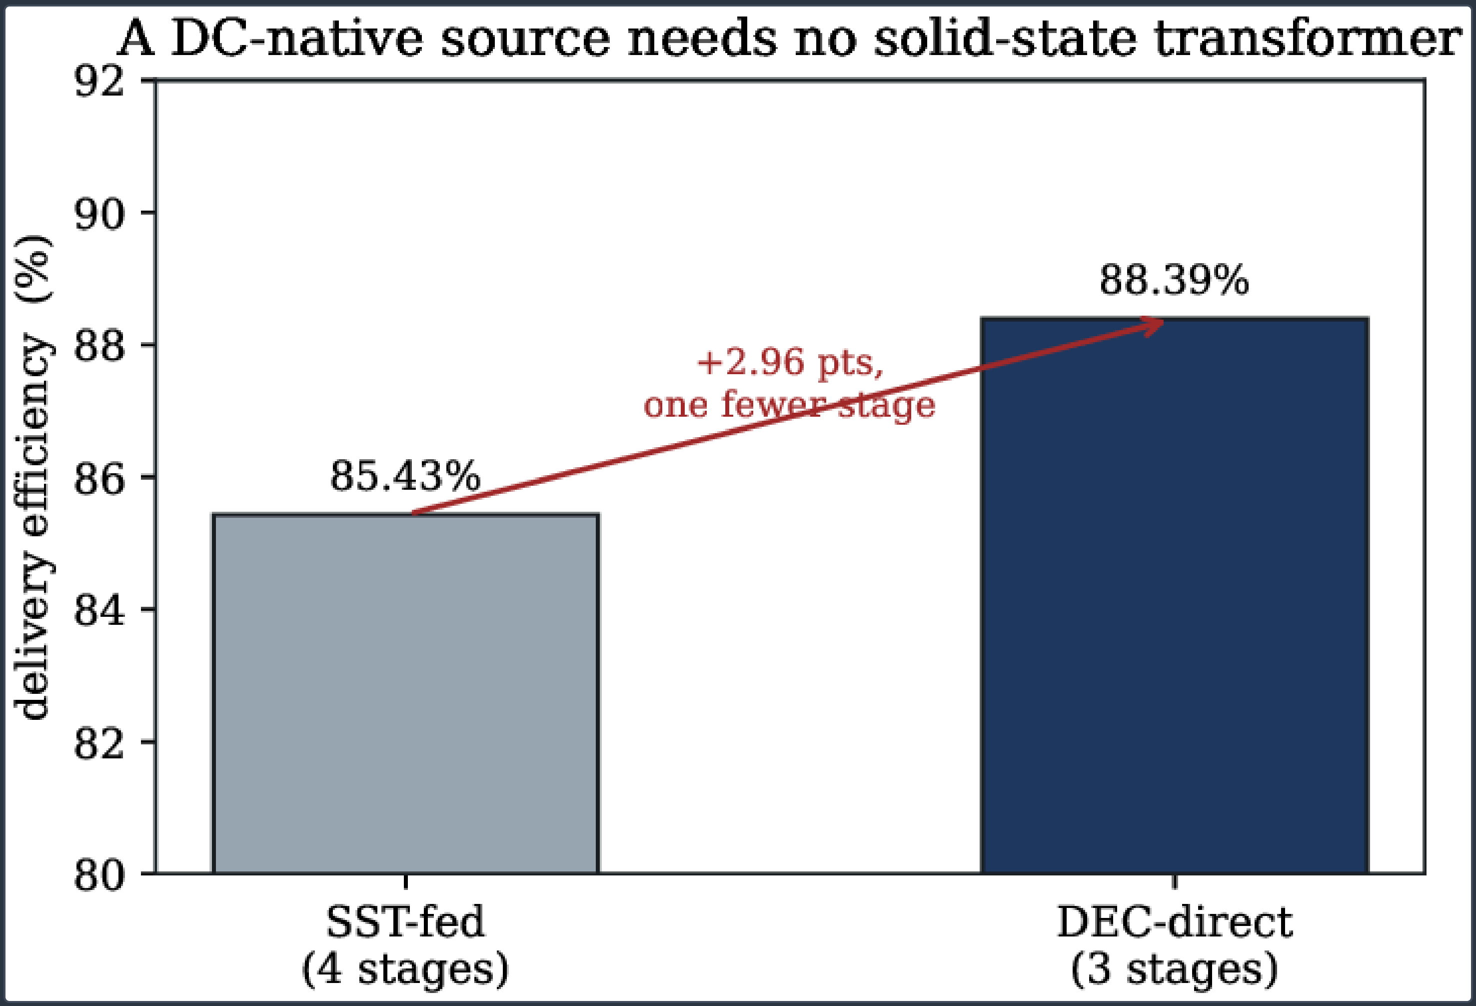

Solid-state-transformer-fed versus DEC-direct delivery (MV-CX-04). Because the DEC is HVDC-native, the DEC-direct path reaches 88.39 % against 85.43 % SST-fed—+2.96 points with one fewer conversion stage. The isolation collapses into a single SiC dual-active-bridge at 15.81 kHz, 98.85 %.

The DC bus is also inherently better-behaved for an IT load. The DEC output is smooth: 1 % ripple needs only a $\sim 234\,\mu$F bus capacitance, far below the millifarad-scale hold-up that the ITIC ride-through curve already demands, so no active harmonic filtering is required (BU-CX-03). And the DEC's intrinsic current limit caps a source-side fault at $\sim$3 kA against $\sim$50 kA from a stiff grid—a 17× reduction in source fault current (MV-CX-05), which is a safety property, not merely an efficiency one.

ELECTRONIC LOAD-FOLLOWING AT CONSTANT PLASMA BASELOAD

A fusion plasma cannot follow a data center's transients: its throttle authority is of order 0.76 $\mathrm{MW\,s^{-1}}$, while synchronized GPU training swings ramp at ~ 80 $\mathrm{MW\,s^{-1}}$ (95th percentile), about $105\times$ faster (BU-CX-08). Following the load by throttling the plasma is doubly costly because the gross-to-net sensitivity amplifier is 4.14: a $10\,\%$ dip in plasma output cuts net export by $\approx 41\,\%$ (MV-L1-A1). MetroVolt therefore does not ask the plasma to follow the load. It runs the reactor at a constant baseload and interposes a plasma-independent electronic buffer—a film-capacitor/supercapacitor/flywheel hybrid of order a few tonnes and ~8.4 MJ store per $30\,\mathrm{MW}$ hall—that absorbs the transients on microsecond-to-second timescales while the DC bus is held within $\pm 0.40\,\%$ through $60 \to 92 \to 40 \to 75\,\%$ load steps (figure 18; BU-CX-07, BU-CX-09). Because the store buffers transients rather than bulk energy, no lithium-ion energy reservoir is needed; a baseload plant with a small fast buffer is the natural firm supply for a crest-factor load.

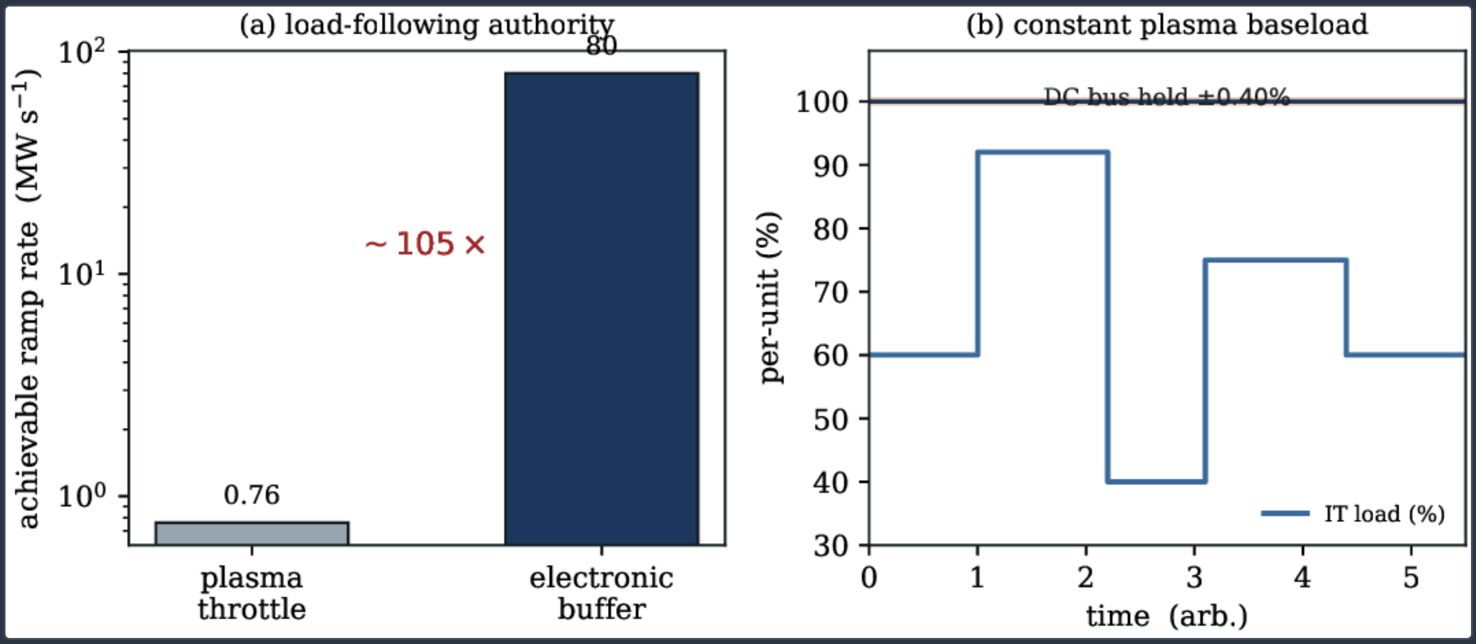

Electronic load-following at constant plasma baseload (BU-CX-07, BU-CX-08). **(a)** The electronic buffer follows load at ~80 $\mathrm{MW\,s^{-1}}$, about $105\times$ the plasma throttle authority. **(b)** Across $60 \to 92 \to 40 \to 75\,\%$ IT-load steps the closed-loop controller holds the DC bus within $\pm 0.40\,\%$ while the plasma runs at a constant baseload; the trace is illustrative, the band and load steps are the frozen result.

GRID-FORMING, BLACK-START, AND RIDE-THROUGH

The same DC-native buffer makes the plant grid-forming. Modelled against a $6\,\%$ generation-loss event on a 1 GVA island with system inertia constant 3 s, MetroVolt's sub-10 ms fast-frequency response (8 ms, 50 MW, 25 droop) lifts the frequency nadir from 49.373 Hz (synchronous-only, grid-following) to 49.900 Hz—a $+0.527$ Hz improvement that keeps the excursion clear of the under-frequency load-shed line—and, more decisively, enables black-start and islanded operation that a grid-following plant cannot provide (figure 19; MV-CX-10, BU-CX-11); the grid-forming droop and hierarchical DC/AC microgrid control follow the standardized approach of Guerrero and co-workers [46]. Ride-through is a battery-sizing decision: millisecond-to-second transients are handled by the supercapacitor/flywheel buffer, while bridging the full 30 MW IT load for five minutes would need a ~ 12.5 t battery, holding only the 6 MW critical set needs ~ 2.5 t, and a black-start battery for the plant's 5 MW auxiliaries for eight minutes is ~ 3.3 t (MV-CX-25). The design therefore does not carry a full-IT UPS; it gracefully sheds to the critical set and keeps a small black-start battery, and MetroVolt is islandable on that basis.

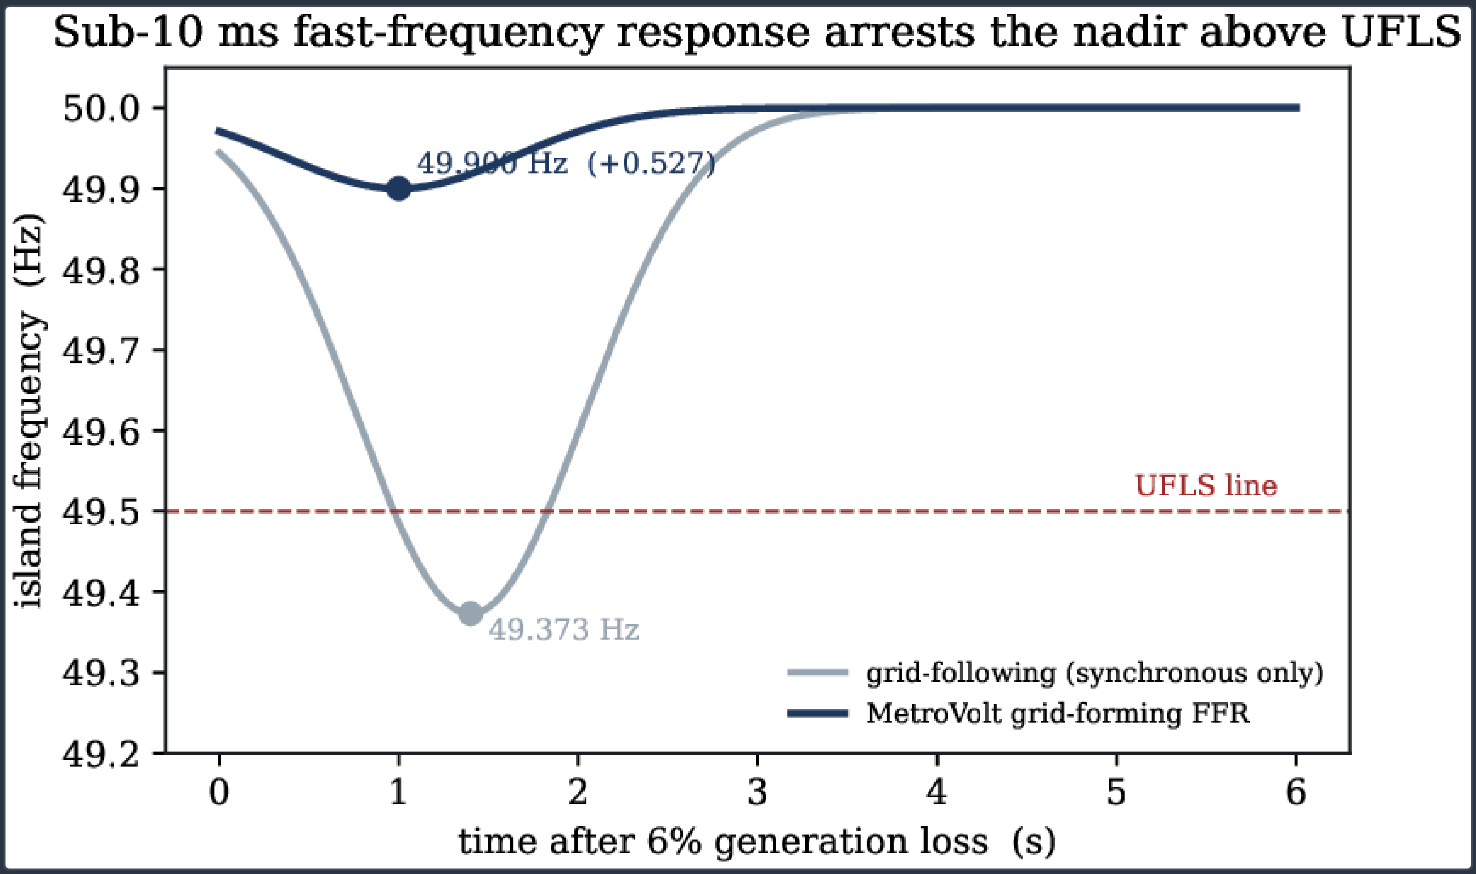

Grid-forming fast-frequency response (MV-CX-10). Against a $6\,\%$ generation-loss event on a 1 GVA island ($H = 3$ s), the sub-10 ms response lifts the nadir from 49.373 to 49.900 Hz ($+0.527$ Hz), above the under-frequency load-shed line. Traces are illustrative; the nadir values are the frozen anchors.

DC-BUS FAULT RIDE-THROUGH

The objection to a medium-voltage DC bus is fault energy, so the study quantifies it. On a 10 kV bus (5 mF, $5\,\mu$H, characteristic impedance $0.032\,\Omega$) a bolted short reaches a 316.2 kA undamped resonant peak at $248\,\mu$s—and is then bounded, not by damping, but by a fast breaker: a $2\,\mu$s silicon-carbide solid-state circuit breaker clamps the fault to ~ 4 kA, $1.3\,\%$ of the resonant peak (figure 20). The DEC itself helps: its intrinsic current limit caps a source-side fault to ~ 3 kA, against ~ 50 kA available from a grid connection—a $17\times$ reduction in source fault current, which is a safety advantage of the DC-native architecture rather than a liability (MV-CX-05).

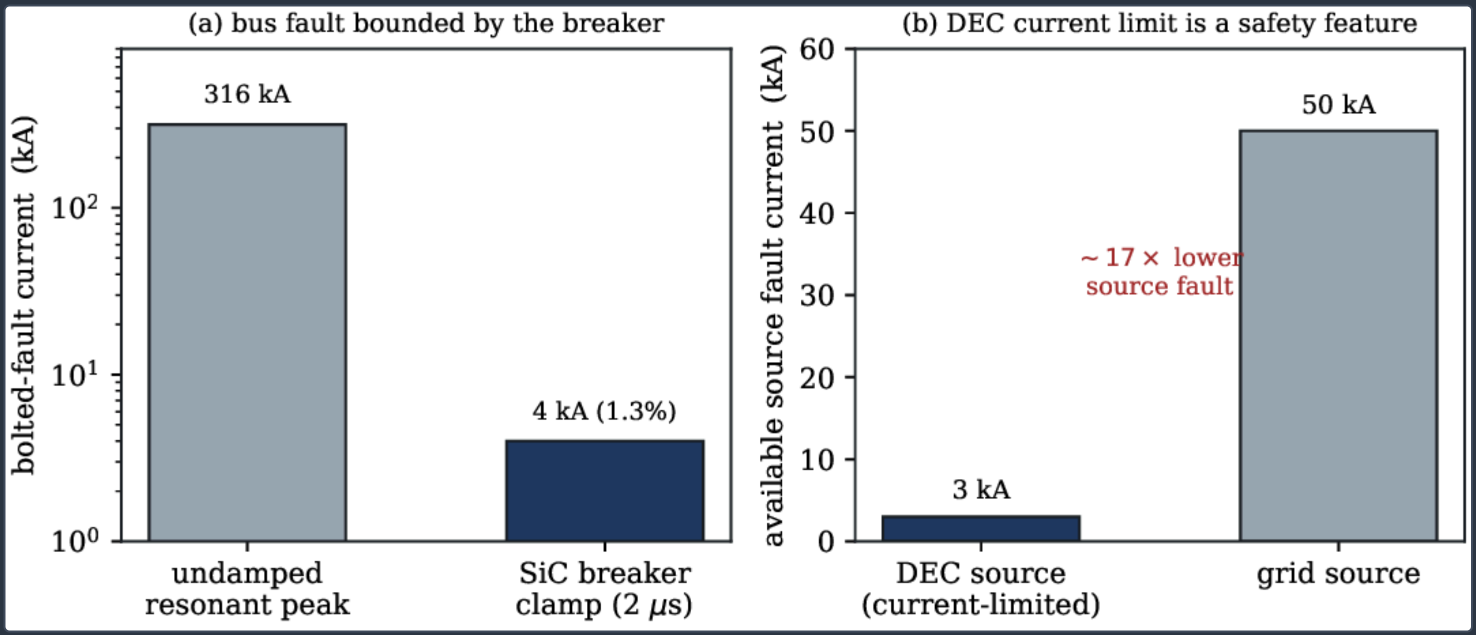

DC-bus fault bounding (MV-CX-05). **(a)** A bolted short on the 10 kV bus reaches a 316.2 kA undamped resonant peak but is clamped to ~ 4 kA ($1.3\,\%$) by a $2\,\mu$s SiC solid-state breaker. **(b)** The DEC's intrinsic current limit caps a source-side fault to ~ 3 kA against ~ 50 kA from a grid, a $17\times$ reduction. Axis (a) is logarithmic.

WASTE HEAT, COOLING, AND POWER-USAGE EFFECTIVENESS

The residual (non-direct) power leaves the plant as heat, and a data center is a large heat sink with an even larger cooling load, so the two integrate naturally. At the **55** m rung the plant rejects **449.5** MW_{th} of waste heat, of which ∼ **142.9** MW_{th} drives a double-effect lithium-bromide absorption chiller (COP ≈ **1.35**) that covers the full cooling demand of a **100**-MW-IT campus on ∼**74** MW_{th}, cutting the electric cooling load by ∼ **91** % (from **16.7** to **1.5** MW_e; BU-CX-12, BU-CX-08). Combined with DC-native conversion and direct-liquid cooling that captures ∼ **90** % of the IT power at **55** °C for reuse—then lifted to **75** °C by a heat pump at COP **8.7** for a useful-heat reuse factor **ERF** ≈ **1.0** (BU-CX-13)—the facility power-usage effectiveness falls from a baseline PUE of **1.35** to **1.17** (figure 21): −**0.05** from DC-native conversion, −**0.07** from direct-liquid cooling, −**0.04** from waste-heat absorption cooling, and −**0.02** from auxiliaries (BU-CX-14). With even a quarter of the waste heat taken up off-site the energy-reuse effectiveness reaches ERE ≈ **0.92**—the campus exports more useful energy than its own overhead consumes.

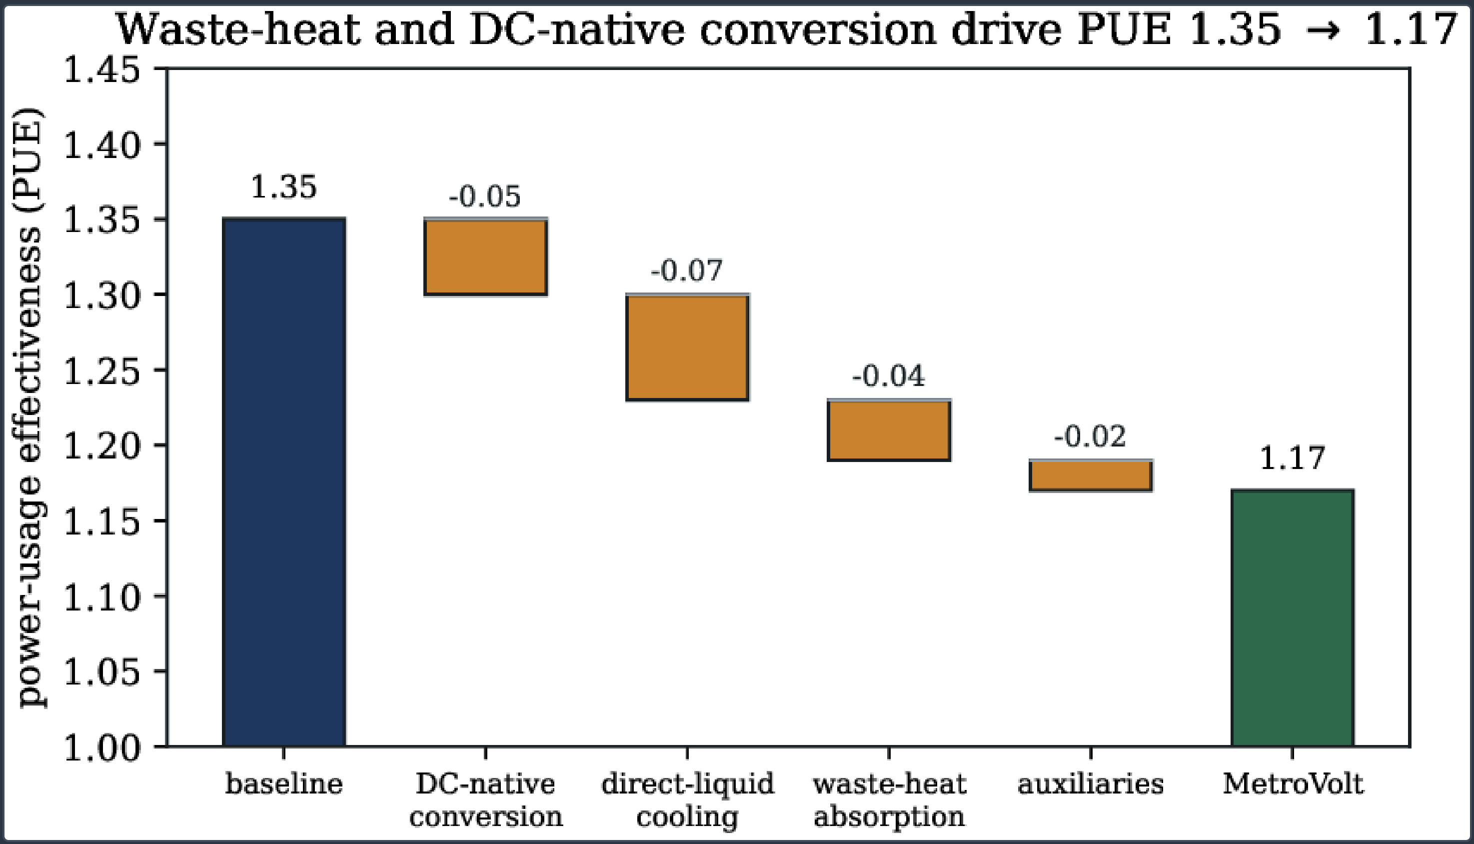

Power-usage-effectiveness pathway (BU-CX-14). DC-native conversion, direct-liquid cooling, waste-heat-driven absorption cooling, and auxiliary savings take the facility from a baseline PUE of **1.35** to **1.17**; with off-site heat uptake the energy-reuse effectiveness reaches ERE ≈ **0.92**.

POWER DENSITY AND BEHIND-THE-METER SITING

A compact mirror is a linear device, and co-locating generation with compute along a right-of-way gives a site power density of ~ 22 MW per acre (850 MW on ~ 39 acres), roughly $3\times$ a gas-plus-air-cooled data center, $\sim 99\times$ a solar-backed campus, and $\sim 700\times$ a wind-backed one (figure 22; BU-CX-17, MV-L1-A7). Behind-the-meter operation raises delivered efficiency by a further 6.3 points (avoiding transmission and distribution loss, 98.5 versus $92.2\,\%$) and, more consequentially for deployment, skips the interconnection queue: time-to-power falls from ~ 78 to ~ 30 months (BU-CX-18, MV-CX-18). A modest grid tie is retained for export and ancillary services, so the optimal topology is a behind-the-meter hybrid rather than a fully islanded plant.

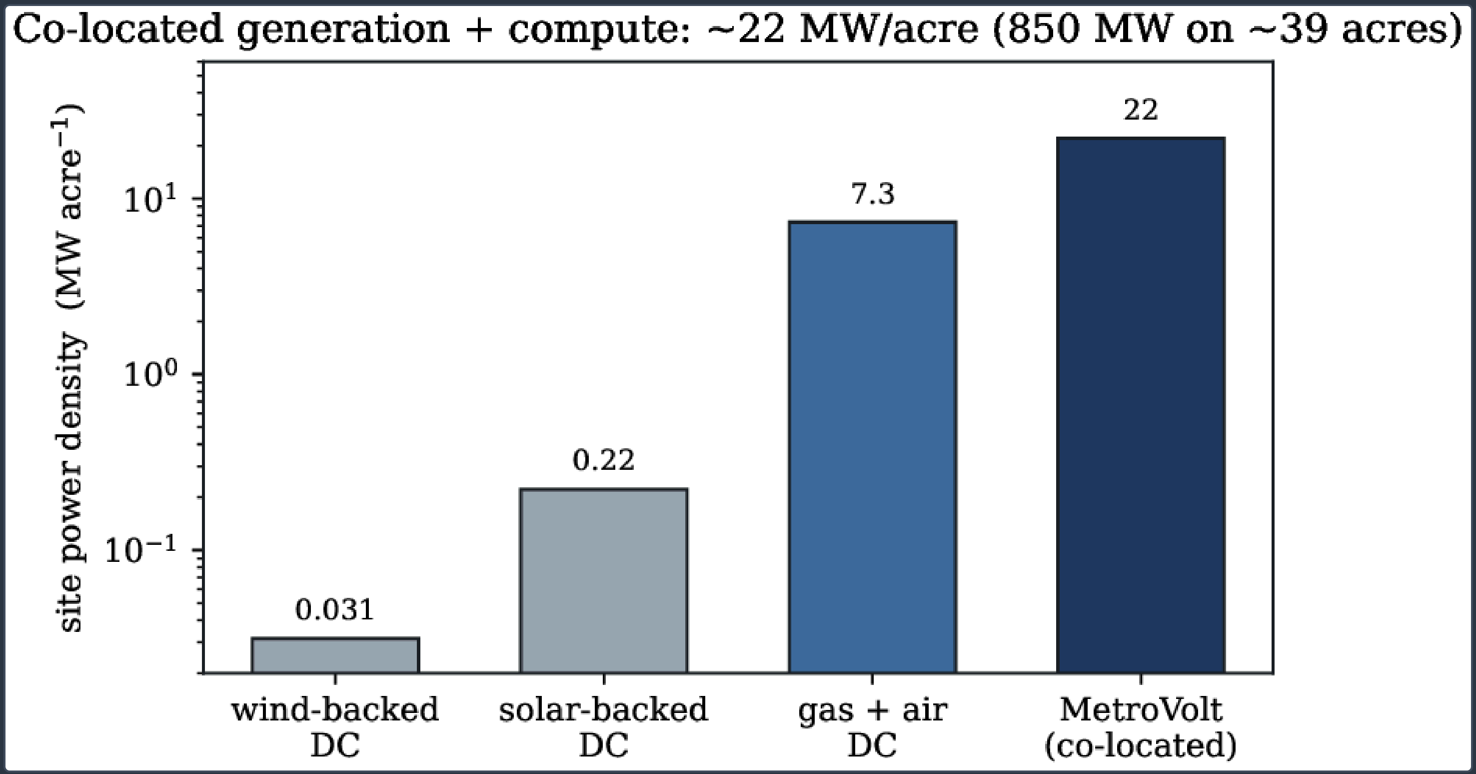

Site power density (BU-CX-17). Co-located generation and compute reach ~ 22 MW per acre (850 MW on ~ 39 acres), roughly $3\times$ a gas-plus-air-cooled data center and orders of magnitude above renewable-backed campuses. Axis is logarithmic.

THE CAPACITY-FACTOR CASE

The reason a firm baseload source is the right generator for a 24/7 compute load is a capacity-factor argument. To firm a continuous load L with a source of capacity factor CF requires nameplate capacity $C \gtrsim L/\mathrm{CF}$ plus storage to bridge the intermittency,

$$C \gtrsim \frac{L}{\mathrm{CF}}, \qquad S \gtrsim L\,\Delta t_{\mathrm{lull}},$$

with Δt_{lull} the worst-case low-output interval. Baseload fusion at $\mathrm{CF} \approx 0.93$ needs $\sim 1.2\times$ the load with buffer-only firming; solar and wind at $\mathrm{CF} = 0.24$–0.35 need ~ 3.3–$4.8\times$ the capacity plus ~ 100–215 hours of storage to reach the same reliability (figure 23; BU-CX-23). For a campus whose economics are set by continuous utilization, the oversizing and storage burden of an intermittent source is the decisive engineering difference, independent of any cost figure.

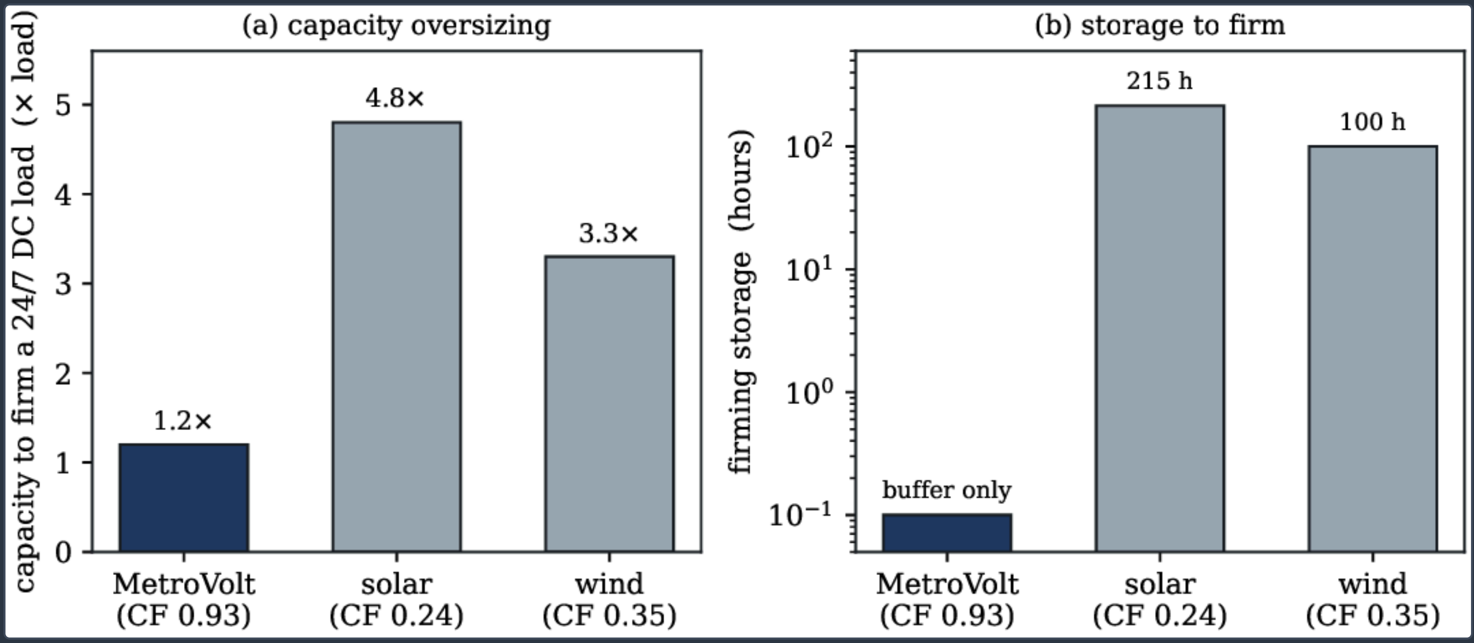

The capacity-factor case for a firm baseload source (BU-CX-23). **(a)** Firming a 24/7 load needs ∼1.2× nameplate from baseload fusion (CF 0.93) but 3.3–4.8× from solar or wind. **(b)** The renewable paths additionally need 100–215 hours of firming storage; the fusion path needs only the transient buffer.

UPTIME AND THE SERVICE-LEVEL AGREEMENT

A critical facility is written to a service-level agreement (SLA), and the demanding tiers leave almost no room: 99.99 % availability is 52.6 minutes of downtime a year, and 99.999 % is 5.3 minutes (table 3). A single fusion module does not reach that alone—modelled single-unit availability runs from 0.90 first-of-a-kind to 0.96 at maturity—so availability is engineered at the fleet level, with N carrying modules plus staggered maintenance and redundancy. A 2500-trial Monte-Carlo availability model (MV-CX-22) resolves what it takes (figure 24): a baseline $N+1$ fleet built from 2000-hour-MTBF power modules meets the Tier-IV target of 99.995 % only 13 % of the time—it misses; raising module reliability to MTBF ≈ 8000 h lifts the probability of meeting Tier-IV to 85 %, and adding redundancy ($N+2$ at the baseline MTBF) reaches 77 %. The levers are therefore explicit and physical—power-module MTBF and failover coverage, not wishful margin. Four-nines availability is delivered grid-tied through those levers; islanded, the honest ceiling is near 99.9 %, held with $N+2$ redundancy and battery storage for ride-through (MV-L1-A4).

SLA / CONFIGURATION	METRIC	VALUE
99.9 %	allowed downtime	526 min/yr
99.99 %	allowed downtime	52.6 min/yr
99.999 %	allowed downtime	5.3 min/yr
single unit (NOAK)	availability	0.96
N+1, MTBF 2000 h	P(meet Tier-IV)	0.13
N+2, MTBF 2000 h	P(meet Tier-IV)	0.77
N+1, MTBF 8000 h	P(meet Tier-IV)	0.85

Service-level targets and the fleet reliability model (MV-L1-A4, MV-CX-22).

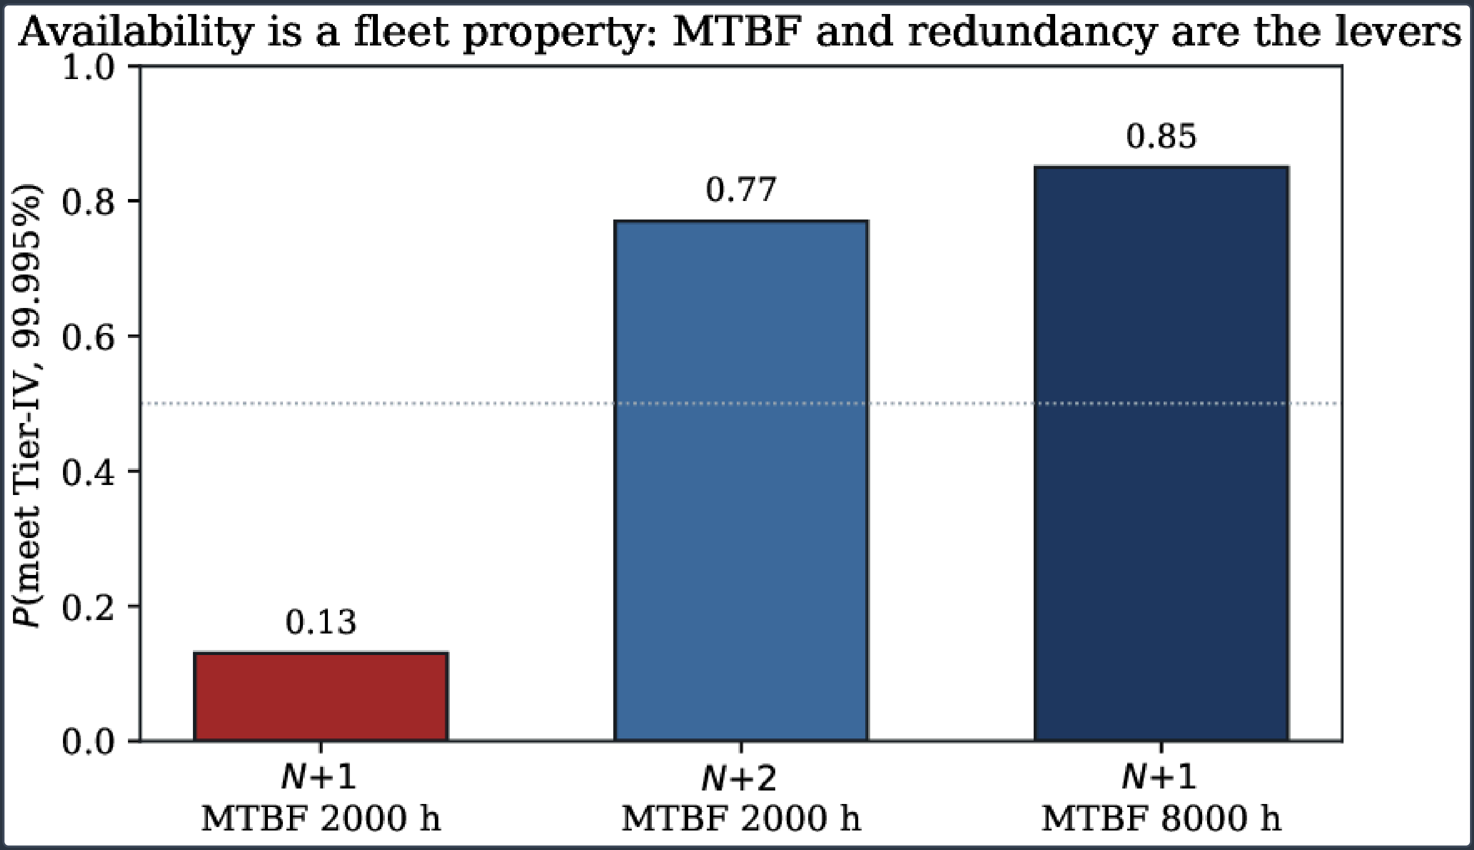

Tier-IV reliability (MV-CX-22, 2500-trial Monte Carlo). Probability of meeting the 99.995 % target for three fleet configurations: a baseline $N+1$ fleet misses; higher-reliability modules or $N+2$ redundancy clear it. Module MTBF and redundancy are the levers, not margin.

NEUTRON-LEAN URBAN SITING AND THE TRITIUM SOURCE TERM

The low-neutron burn is what makes an urban or near-campus site feasible. A point-kernel dose estimate meets an urban public-dose target of $\sim$**0.11** μSv h^{-1} at **50** m behind a buildable $\sim$**2.8** m concrete biological shield, against $\sim$**3.06** m for a deuterium–tritium plant of equal power—a modest saving, because dose scales logarithmically with shield thickness (BU-CX-28). The licensing-decisive quantity is not the shield but the mobile tritium source term. Although the burner self-generates $\sim$**32** kg of tritium per full-power year (of which $\sim$ **31** % is burned in situ), the licensing-relevant mobile inventory—held to a one-hour processing buffer—is only $\sim$**3.6** g, roughly **275**× below the $\sim$**1** kg-class inventory of a D–T plant (figure 25; BU-CX-29). A further *burn-lean* siting lever is available: pumping bred tritium before it burns strips $\sim$ **60** % of the neutron power (**234** → **93.5** MW, down to the D–D floor) for only a $\sim$ **4** % fusion-power penalty ([63]). This small source term is what underwrites siting a fusion plant next to a data-center campus and is the physical basis of the low-activation licensing case [65].

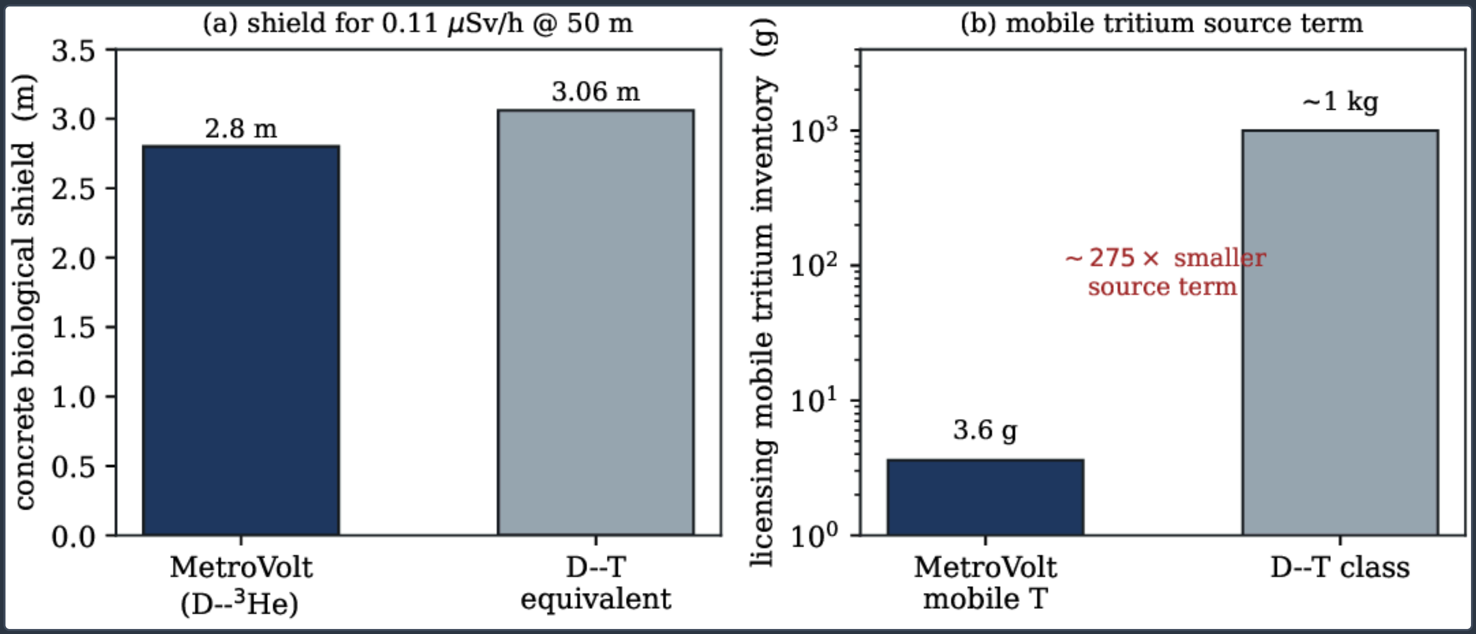

Neutron-lean urban siting (BU-CX-28, BU-CX-29). **(a)** The low-neutron burn meets an urban public-dose target behind a $\sim$**2.8** m concrete shield versus $\sim$**3.06** m for a D–T equivalent. **(b)** The licensing-relevant mobile tritium inventory is $\sim$**3.6** g, roughly **275**× below a D–T-class plant; axis is logarithmic.

TWO DISTINCT MAGNETS

Throughout, the burner plug coil and the sister breeder centrepost are treated as the distinct devices they are. The breeder centrepost of the companion Hyperion platform [48] governs at a peak field of **16.84** T; the MetroVolt burner plug coil is a separate, higher-field REBCO magnet at **26.49** T, structurally qualified by analytic thick-wall and finite-element analysis (magnetic pressures **115**/**279**/**358** MPa at **17**/**26.49**/**30** T against a **667** MPa allowable, FEA stress ratio PASS; BU-L1-A2) and detailed in the magnet-basis companion [50] and the plug-coil engineering companion [52]. Conflating the two fields is a category error the design explicitly avoids.

§ 5

Results and verification

Table 4 collects the frozen design point; figures 1–25 present the physics and product results in turn. Table 6 collects the data-center product metrics with their register serials and their dependence—or independence—from the plug gate. Table 5 collects the independent closure studies of the engineering gain.

QUANTITY	VALUE
plant energy multiplication Q_E	1.318
fusion power P_{fus}	4298.5
neutron power P_n	233.95
neutron fraction f_n	5.44 %
charged (direct-convertible) power	4064.6
axial power to DEC P_{ax}	2719.7
direct-conversion fraction f_{dir}	0.495
plug potential ϕ_i	248.75 kV
carrier temperature T_e	89.7
plug density ratio n_p/n_c	16
first-wall reflectivity requirement R_{wall}	$\gtrsim 0.83$
converter efficiency η_{DEC} **(conservative anchor)**	0.49
net-power flip threshold in η_{DEC}	$\simeq 0.52$
net electric (55/440/1400 m)	$+104/+850/+2832$ MW_e
$P(Q_E > 1)$ **over design ranges**	1.0

MetroVolt burner design point (register serials BU-L1-A1 / KX-L1-A9 / BU-L2-A5 / KX-L1-A5-BU). Every net-power quantity is conditional on retiring the plug-potential gate; the neutron-source performance is not.

STUDY	RESULT	SOURCE
design point (0-D balance)	$Q_E = 1.3181$	BU-L1-A1
design-space scan (25 cases)	$Q_E \in [1.107, 1.414]$, all > 1	BU-L2-A2
global UQ (40k Monte-Carlo)	$Q_E = 1.317 \pm 0.14$, $P(Q_E > 1) = 1.0$	KX-L1-A5-BU
Sobol first-order driver	plug density ratio, $S_1 = 0.41$	HX-05
net-power UQ (1000 MC + Sobol)	$P(\text{net} > 0) = 73\,\%$	AG-DR-02
synchrotron model-form band	$Q_E = 0.39/1.32/4.56$ (thin/AFJ/Trubnikov)	AG-DR-01
thermal-barrier value	$+2.8\times$ in Q_E $(0.47 \to 1.318)$	AG-CX-05

table continued

mirror-ratio knee	$Q_E = 1.18 \to 1.53$ ($R_{mc} = 3.3 \to 6.8$)	AG-CX-06

METRIC	VALUE	SERIAL	PLUG DEPENDENCE
source-to-chip efficiency (DC-native / AC)	88.1 % / 75.9 %	BU-CX-02	robust
energy-weighted annual η (DC / AC)	90.4 % / 78.9 %	BU-CX-06	robust
SST-fed vs DEC-direct delivery	85.4 % / 88.4 %	MV-CX-04	robust
load-following ramp (buffer / plasma)	80 / 0.76 $\mathrm{MW\,s^{-1}}$	BU-CX-08	robust
DC-bus regulation across load steps	±0.40 %	BU-CX-07	robust
grid-forming frequency-nadir lift	+0.527 Hz	MV-CX-10	robust
bus-fault clamp / source-fault reduction	4 kA / 17×	MV-CX-05	robust
power-usage effectiveness	1.35 → 1.17	BU-CX-14	robust
Tier-IV reliability (N+1, MTBF 8000 h)	$P = 0.85$	MV-CX-22	robust
site power density	∼22 $\mathrm{MW\,acre^{-1}}$	BU-CX-17	conditional
firming capacity vs. renewables	1.2× vs. 3.3–4.8×	BU-CX-23	robust
urban shield / D–T equivalent	2.8 / 3.06 m	BU-CX-28	robust
mobile tritium source term	∼3.6 g	BU-CX-29	robust
net electric (reference campus)	+850 $\mathrm{MW_e}$	KX-L1-A9	conditional

MetroVolt data-center product metrics, their register serials, and their dependence on the plug gate. “Robust” quantities hold regardless of the plug closure; “conditional” quantities scale with the net-power case.

§ 6

Sensitivity and uncertainty quantification

A single-point balance is only as good as its sensitivity, and the closure of the engineering gain is best understood as the joint outcome of several independent studies rather than one number. The design-space scan returns $25/25$ configurations above unity ($Q_E \in [1.107, 1.414]$), with the design point reproducing $Q_E = 1.3181$ exactly and a mild ridge to $Q_E \approx 1.375$ at slightly higher central temperature (BU-L2-A2). The 4×10^4-sample global Monte-Carlo over density, ^{3}He fraction, temperature, central-cell β and mirror field gives $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$ and $f_n = 5.39\,\%$, the central-cell β the dominant driver (rank correlation 0.93; KX-L1-A5-BU). A Saltelli–Sobol decomposition identifies the plug density ratio as the first-order driver ($S_1 = 0.41$), interacting strongly with central β ($S_T = 0.87$; HX-05). Finally, a 1000-sample net-power Monte-Carlo with Sobol attribution—run with the burner deliberately placed near the $Q_E = 1$ cliff—returns $P(\text{net} > 0) = 73\,\%$ with a 5th–95th percentile band $Q_E \in [0.81, 2.51]$, the variance led by the plug-to-central density ratio (AG-DR-02, figure 26). The consistent message is that the closure is not a knife-edge coincidence: it holds across the design box, and the levers that matter—reflectivity, thermal-barrier temperature, plug density ratio—are named and specifiable, not hidden.

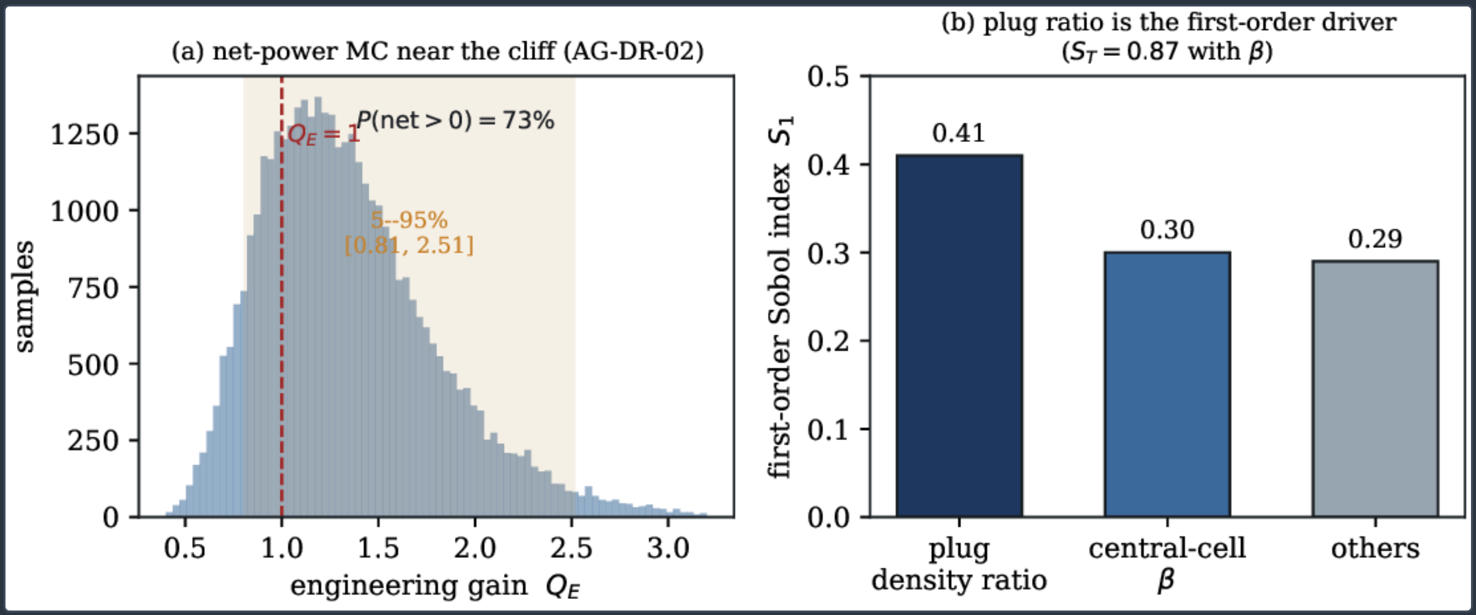

Net-power robustness and its drivers. **(a)** The 1000-sample net-power Monte-Carlo, run near the $Q_E = 1$ cliff, keeps $P(\text{net} > 0) = 73\,\%$ with a 5–$95\,\%$ band $Q_E \in [0.81, 2.51]$ (AG-DR-02). **(b)** The Saltelli–Sobol first-order indices identify the plug density ratio ($S_1 = 0.41$) as the leading driver, interacting with central-cell β to a total-order index $S_T = 0.87$ (HX-05).

Two features of the sensitivity study deserve emphasis. First, the dominant physical lever on Q_E is the plug density ratio—the same quantity that is the binding open gate—so the uncertainty and the requirement coincide: the number the design most needs an experiment to confirm is also the number that most moves the gain. Second, the synchrotron model-form band (thin/AFJ/Trubnikov, $Q_E = 0.39/1.32/4.56$) is larger than the entire input-parameter UQ band, which is why the transport calculation of §2.9 is reported as a load-bearing result and not a footnote: resolving the model form is more valuable than tightening any single input.

§ 7

Benchmarking against the published literature

AGAINST THE TANDEM-MIRROR RECORD

The physics honesty of this design rests on separating what is demonstrated from what is required. The tandem-mirror concept and its thermal-barrier variant are mature in the literature—Fowler and Logan's tandem-mirror reactor [1], the thermal-barrier proposal of Baldwin and Logan [2], the TMX-U and axisymmetrized GAMMA-10 experiments that formed thermal barriers and confining potentials [10,11,12], the gas-dynamic-trap high-β line at Budker [13], and the Pastukhov/Cohen theory of ambipolar-well end-loss [6,7] underlying equation (8). What no experiment has demonstrated is a plug potential at the reactor-relevant magnitude of equation (6): the required ϕ_i is more than two orders of magnitude beyond the GAMMA-10/TMX-U record, and the WHAM-class high-field mirror now coming online [15,16] addresses the unplugged confinement regime rather than the plugged branch our net-power case depends on (BU-L2-A3). A direct 0-D benchmark against WHAM reproduces its measured *unplugged* regime and shows the plugged-branch model to be conservative there by two to three orders of magnitude; the plugged net-gain branch itself rests on the Pastukhov–Fokker–Planck closure and is the quantity a plugged-mirror experiment will confirm ([62]). We therefore present the plug requirement as the binding open gate, retired only by a dedicated plug-potential experiment, and every net-power figure in table 4 and figure 10 inherits that condition. The neutron-source performance—$f_n = 5.44\,\%$, $P_n = 233.95$ MW—is by contrast independent of the plug and rests on the fusion cross-sections alone.

DEVICE	REGIME	CONFINING POTENTIAL	STATUS	REFERENCE
TMX-U	plugged, thermal barrier	$\sim$ kV-class	demonstrated	[10]
GAMMA-10	plugged, thermal barrier	$\sim$ kV-class	demonstrated	[11,12]
GDT (Budker)	axisymmetric, high-β	sloshing-ion, unplugged	demonstrated	[13]
WHAM	high-field axisymmetric	unplugged break-even	commissioning	[15,16]
MetroVolt	**plugged D–^{3}He**	$\phi_i = 248.75$ **kV**	**requirement**	this work

The tandem-mirror record versus the MetroVolt requirement, placing the binding gate in context. Demonstrated plug/ambipolar potentials are drawn from the cited experiments; the MetroVolt column is the design requirement, not a demonstrated value. The gap in $e\phi_i$ is the paper's one binding open gate.

Table 7 places the gate in context: the demonstrated tandem-mirror plug potentials are of order kilovolts, while the MetroVolt net-power case requires an ion-confining potential of 248.75 kV—more than two orders of magnitude larger. The gas-dynamic-trap line demonstrated the sloshing-ion, high-β, MHD-stable axisymmetric operation on which the plug is built [13], and WHAM is bringing high mirror fields to the unplugged break-even regime [15,16]; neither closes the plugged branch. We do not present the requirement as met. What the benchmark establishes is that the physics ingredients—thermal barriers, sloshing-ion potentials, high-β axisymmetric stability—are individually demonstrated, and that the open question is their combination at reactor-relevant potential, which is a well-posed experimental checkpoint rather than an unknown.

AGAINST THE DIRECT-CONVERSION LITERATURE

Direct energy conversion of a mirror's escaping ion stream has a long pedigree, from the Post venetian-blind collector [19] through the Moir–Barr analyses that set the multi-stage efficiency expectations near 0.70 [20,21], and modern reappraisals of expander-region conversion [22]. Our conservative anchor of $\eta_{\mathrm{DEC}} = 0.49$ sits deliberately below that literature's design value, so no headline number leans on an optimistic converter; the efficiency ladder (0.49 floor anchor, 0.56 model floor, 0.70 design, 0.85 ideal) is carried explicitly (BU-SX-01, figure 11). The contribution specific to MetroVolt is not a new converter efficiency but the observation that the converter's *output form*—high-voltage DC at the plug scale, equation (11)—is a first-class product advantage for a data-center load, and the quantification of that advantage as a +12.2-point source-to-chip efficiency gain (figure 13). Companion papers develop the DEC collector physics and the DEC-to-load power-electronics interface in full [51,53,54].

AGAINST THE SYNCHROTRON-RADIATION LITERATURE

The synchrotron transport of §2.9 is benchmarked against the two standard closed-form treatments: the reflective-cavity AFJ model of Albajar, Johner and Granata [25], adopted as canon, and the earlier Trubnikov formula [24], carried as an optimistic upper reference. The relevance of the emission and absorption physics to reactor-grade plasmas follows Bornatici and co-workers [26]. That an optically-thin estimate overstates the loss by $\sim 37\times$ (figure 8) is the quantitative reason a reflective-cavity transport calculation is mandatory for an advanced-fuel mirror, and it is why the burner is described as synchrotron-limited rather than confinement-limited at fixed size.

AGAINST THE DATA-CENTER POWER-DISTRIBUTION LITERATURE

The data-center community has independently converged on DC distribution and direct-liquid cooling as the levers on facility efficiency: the $380\text{–}400$ V DC racks and their measured savings [30,31,32], the migration to direct-liquid and two-phase immersion cooling [27,36], and the PUE and energy-reuse-effectiveness metrics standardized for reporting facility overhead [28,29,37]. MetroVolt's advantage is that it supplies the front of that chain in the native form the chain wants: a generator whose output is already HVDC removes the rectification and UPS double conversion that a grid-fed facility cannot, which is why the DC-native chain holds $90.4\,\%$ against $78.9\,\%$ energy-weighted (BU-CX-06). The load-following result (figure 18) similarly reframes a perceived fusion liability—slow plasma response—as a non-issue: a baseload plant with a small electronic buffer follows a crest-factor AI load better than a dispatchable thermal plant chasing it with steam, because the buffer responds in milliseconds and the plasma never moves. The capacity-factor argument of equation (13) and figure 23 is the quantitative statement of why a firm source, not an intermittent one oversized and storage-backed, is the right generator for a continuous compute load. The absorption-cooling waste-heat integration follows the second-law analyses of the data-center thermal literature [35].

§ 8

Design implications

Three design implications follow from the analysis. First, the plant should be delivered as a DC-native generator, not retrofitted into AC-plant assumptions: every power-delivery decision (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) went against the conventional choice and each went that way because the DC bus and the DEC's current limit make it the better-performing option. Second, the first-wall optical finish is a physics-class specification, not a manufacturing detail: because $R_{\mathrm{wall}} \gtrsim 0.83$ is what separates $Q_E = 0.39$ from $Q_E = 1.32$, wall reflectivity is a design requirement with the same standing as the plug ratio. Third, availability is a fleet property to be engineered through module MTBF and $N{+}2$ redundancy, delivered behind the meter with a grid tie for ride-through—the hybrid topology is not a compromise but the answer that both the dispatch analysis and the reliability analysis independently recommend.

§ 9

Limitations

Five caveats bound the claims, each stated with its evidence class. First, every net-power number— Q_E, the length ladder, the reference-campus $+850$ $\mathrm{MW_e}$—is contingent on retiring the plug-potential gate of equation (6), which a design study cannot do; this is the one binding open gate. Second, the converter efficiency is a declared, deliberately conservative input, not a computed plant result; the direct-DC *topology* advantage, however, does not depend on its value. Third, the plug microstability is marginal: warm-fill clears the DCLC knee only as a threshold (equation (9)) requiring $L_n/\rho_i \gtrsim 25.3$, and a residual AIC branch persists (BU-L2-A7). Fourth, the synchrotron closure carries a hard first-wall reflectivity specification $R_{\mathrm{wall}} \gtrsim 0.83$; the operating point is synchrotron-limited, and a wall that degrades optically over life would erode the margin. Fifth, the helium-3 fuel supply is a separate, lunar-gated problem treated elsewhere [64] and outside this power-plant design. The data-center product metrics of table 6—the delivery efficiency, load-following, PUE, reliability, and siting case—are the robust, plug-independent core; the net-power magnitude is the conditional part.

§ 10

Conclusion

MetroVolt matches the shape of the modern data-center load because its generator produces electricity in the form the load consumes. The D–^{3}He tandem-mirror burner reaches a plant energy multiplication $Q_E = 1.318$ at 4298.5 MW fusion power with a neutron fraction of only $5.44\,\%$, delivering $+104/+850/+2832$ $\mathrm{MW_e}$ across a length ladder with $P(Q_E > 1) = 1.0$; its direct energy converter produces high-voltage DC natively, reaching the silicon at $88.1\,\%$ source-to-chip against $75.9\,\%$ for a conventional AC plant, following a crest-factor AI load electronically at constant plasma baseload, taking the facility to PUE 1.17 through waste-heat integration, and licensing against a ~ 3.6 g mobile tritium source term small enough to site beside a campus. Two model-form choices are load-bearing and reported as such: the plug-potential requirement $n_p/n_c = 16$, retired only by experiment, and the synchrotron transport that pins $Q_E = 1.32$ and demands $R_{\mathrm{wall}} \gtrsim 0.83$. The burner plug coil ($26.49$ T) and breeder centrepost (16.84 T) are modelled as the distinct magnets they are. The design is presented with its robust and conditional parts separated openly—the product coupling robust, the net-power magnitude conditional on the one residual physics gate.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The authors thank the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the reviewers of the 2026 design series.

Reproducibility. Every headline quantity is a frozen anchor in the Kronos Physics De-Risking Register [47] (doi:10.5281/zenodo.22133057): the burner power balance and neutron fraction (BU-L1-A1), the plug-potential requirement (BU-L2-A5), the synchrotron transport (KX-L1-A3-BU, AG-DR-01), the Q_E uncertainty (KX-L1-A5-BU), the length ladder (KX-L1-A9), and the MetroVolt product metrics (BU-CX-01 through BU-CX-29, MV-CX-01 through MV-CX-25, MV-L1-A1 through MV-L1-A9). Each is regenerable from the open, CC-BY-4.0 deposit; this paper is archived at doi:10.5281/zenodo.22645695; official version: https://www.kronosfusionenergy.com/publications/metrovolt-burner.html. No result depends on unpublished or proprietary data.

Scope — what this paper does not claim. Q_E is the engineering gain at the frozen design point, not a demonstrated result; the end-plug density requirement is stated as the binding open gate, not solved; “low-neutron” is not “aneutronic”; and no cost, price, tariff, valuation, or other financial quantity is reported or implied anywhere in this paper.

Data availability tocsectionData availability The figure-generation scripts and the frozen anchors used for figures 1–26 are archived with the deposit at doi:10.5281/zenodo.22645695, which cross-references the Kronos 2026 series: the register [47], the Hyperion breeder platform [48], the sister Aegis burner [49], the REBCO magnet basis [50], the plug-coil engineering [52], direct energy conversion [51], the DEC collector physics [53], the DEC-to-load data-center interface [54], the power-delivery topologies [55], the data-center integration [56], recirculation-resolved Q_E closure [57], synchrotron transport [58], velocity-space microstability [59], sloshing-ion potential formation [60], RF alpha-channeling [61], the tandem-mirror design space [62], the neutron-lean siting case [63], the helium-3 supply gate [64], and the Part 30 licensing thesis [65]. Any economic analysis is reported separately and is not part of the public deposit.

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

§ 11

Derivation of the Pastukhov–Cohen confinement time

We outline the derivation of equation (8) for completeness. Consider ions confined in a mirror cell of mirror ratio R_m and axial ambipolar potential ϕ_i, with an isotropic Maxwellian source at temperature T_i. An ion is lost only if pitch-angle scattering carries it into the loss cone while its parallel energy exceeds the potential barrier $e\phi_i$. The steady-state kinetic equation balances the collisional flux into the loss region against the source,

$$S(\mathbf{v}) = -\nabla_{\mathbf{v}}\cdot\mathbf{\Gamma}_{\text{coll}}, \qquad \mathbf{\Gamma}_{\text{coll}} = -D(\mathbf{v})\,\nabla_{\mathbf{v}} f,$$

with D the velocity-space diffusion tensor set by the Rosenbluth potentials. In the strong-potential limit the loss region is a thin boundary layer in energy near $E \simeq e\phi_i$, and the dominant balance is between energy diffusion up to the barrier and the loss-cone sink. Solving the boundary-layer equation with an absorbing boundary at the loss cone and matching to the bulk Maxwellian yields a loss rate proportional to $\exp(-e\phi_i/T_i)$; inverting gives a confinement time

$$\tau_{\parallel} \;\simeq\; \tau_{ii}\, g(R_m)\left(\frac{e\phi_i}{T_i}\right)\exp\!\left(\frac{e\phi_i}{T_i}\right), \qquad g(R_m) = \frac{2R_m}{\sqrt{\pi}\,(2R_m+1)\,\ln(4R_m+2)},$$

where the geometric factor $g(R_m)$ collects the loss-cone solid angle and the logarithmic Coulomb dependence. The factor $(e\phi_i/T_i)\exp(e\phi_i/T_i)$ is the exponential enhancement plotted in figure 4; it is the mathematical statement of why an electrostatic plug of a few T_i produces a confinement gain of order ten and why realizing ϕ_i is the closure problem.

§ 12

Synchrotron harmonic-transport scheme

The synchrotron loss of §2.9 is computed by summing the emission and reabsorption over cyclotron harmonics n for a relativistic Maxwell–J\"uttner electron distribution. For each harmonic the local emission coefficient $j_n(\omega)$ and absorption coefficient $\alpha_n(\omega)$ are evaluated, the optical depth $\tau_n = \int \alpha_n \, \mathrm{d}\ell$ is integrated along the escape path, and the emergent intensity is

$$I_n = \frac{j_n}{\alpha_n}\left(1 - e^{-\tau_n}\right) + I_n^{\text{wall}}\, e^{-\tau_n}, \qquad I_n^{\text{wall}} = R_{\text{wall}}\, I_n,$$

so that the wall reflectivity R_{wall} recycles escaping radiation back into the cavity. The net loss is $P_{\text{sync}} = \sum_n \int (1 - R_{\text{wall}})\, I_n \, \mathrm{d}\Omega \, \mathrm{d}\omega$ over the escaping band $n > n^*$, with the cutoff harmonic n^* where $\tau_{n^*} \sim 1$; the sum is carried to convergence past $n^* \approx 10.4$. The cutoff scales as $n^* \propto (\omega_{pe}^2 a/\omega_{ce})^{1/3} \sim a^{1/3}$, so a larger machine reabsorbs a larger fraction, which is the analytic origin of the length ladder of figure 10. Evaluated at the design point the scheme returns radiated 2546 MW and reabsorbed 1803 MW (AFJ-consistent), against 919 MW (Trubnikov) and 66,677 MW (optically thin).

§ 13

Direct-converter multi-stage efficiency model

The multi-stage collector efficiency of equation (11) is modelled as a graded electrostatic decelerator with N_s stages matched to the incident ion energy spectrum. Each stage recovers the kinetic energy of the ions whose energy lies in its collection band; the residual loss is the sum of the grid transparency loss, the energy-spread mismatch between the incident distribution and the finite stage count, and secondary-electron losses. The efficiency approaches an asymptote as $N_s \to \infty$,

$$\eta_{\text{DEC}}(N_s) = \eta_\infty \left(1 - \frac{c}{N_s + b}\right),$$

matched to the anchored points $\eta_{\text{DEC}}(3) = 0.70$ and $\eta_{\text{DEC}}(6) = 0.83$ over the ~ 428 keV end-loss spectrum (figure 11); the plant is anchored conservatively at $\eta_{\text{DEC}} = 0.49$, below the 0.56 model floor. The efficiency is robust to throttling because the ion potential shifts only $\sim 5\,\%$ as the load falls to $16\,\%$, so η_{DEC} holds from 0.70 to 0.687 (BU-CX-27).

§ 14

DC-chain stage-efficiency accounting

The source-to-chip efficiency of equation (12) is the product of the individual conversion-stage efficiencies along the delivery path. Table 8 lists the two paths and their stage counts. The AC path incurs rectification, UPS double conversion, transformation, and the rack power-supply unit; the DEC-direct path incurs only an HVDC→MVDC step and a rack-level DC–DC, with the isolation collapsed into a single SiC dual-active-bridge at 98.85 %. The delivered figures are 75.9 % (AC) and 88.1 % (DEC-direct MVDC), widening to 78.9 % versus 90.4 % energy-weighted over an annual load profile because the DC chain does not sag at partial load (BU-CX-02, BU-CX-06).

PATH	STAGES TO 800 V RACK BUS	DELIVERED η (POINT / ANNUAL)
conventional AC + UPS	4–6	75.9 % / 78.9 %
OCP 400 V DC	3–4	86.2 % / —
DEC-direct MVDC	2–3	88.1 % / 90.4 %

Source-to-chip conversion paths (BU-CX-01, BU-CX-02, BU-CX-06, MV-CX-04).

§ 15

Parameter and anchor tables

Table 15 collects the frozen anchors used in this paper with their register serials, so that each figure and equation can be traced to an auditable source.

longtable@p6.6cmp4.3cml@ Frozen anchors used in this paper (Tier-2 freeze, 2026-09-05).\ quantity & value & serial \ quantity & value & serial \ plant energy multiplication Q_E & 1.3181 & BU-L1-A1 \ fusion power P_{fus} & 4298.5 MW & BU-L1-A1 \ neutron power P_n & 233.95 MW & BU-L1-A1 \ neutron fraction f_n & $5.44\,\%$ & BU-L1-A1 \ D–^{3}He branch power & 3868.0 MW & ENV-A4-BU \ D–D proton branch power & 129.7 MW & ENV-A4-BU \ D–D neutron branch power & 124.8 MW & ENV-A4-BU \ D–T branch power & 176.0 MW & ENV-A4-BU \ wall-averaged neutron loading & 0.098 $\mathrm{MW\,m^{-2}}$ & BV-07 \ axial power to DEC P_{ax} & 2719.7 MW & BU-L1-A10 \ DEC recovered power & 1903.8 $\mathrm{MW_e}$ & BU-L1-A10 \ direct-conversion fraction f_{dir} & 0.495 & BU-L1-A10 \ net-power flip threshold η_{DEC} & $\simeq 0.52$ & BU-L1-A10 \ plug potential ϕ_i & 248.75 keV & BU-L2-A5 \ carrier temperature T_e & 89.72 keV & BU-L2-A5 \ plug density ratio n_p/n_c & 16 & BU-L2-A5 \ DCLC warm-fill threshold & $5.9\,\%$ ($\sim 6.25\,\%$) & BU-L2-A5 \ DCLC gradient requirement & $L_n/\rho_i \gtrsim 25.3$ & BU-L2-A5 \ residual AIC growth rate & $\gamma/\omega_{ci} \sim 0.03$–$0.037$ & BU-L2-A7 \ sloshing peak positions & $z/L_m = \pm 0.24$ & sloshing \ sloshing well depth & $\Delta\varphi = +0.38\,T_i$ & sloshing \ radiated synchrotron power & 2546 MW & KX-L1-A3-BU \ reabsorbed synchrotron loss (AFJ) & 1803 MW & KX-L1-A3-BU \ synchrotron cutoff harmonic n^* & ≈ 10.4 & KX-L1-A3-BU \ wall-reflectivity requirement R_{wall} & $\gtrsim 0.83$ & AG-DR-01 \ Q_E (thin / AFJ / Trubnikov) & $0.39/1.32/4.56$ & AG-DR-01 \ thermal-barrier value & $+2.8\times$ $(0.47 \to 1.318)$ & AG-CX-05 \ mirror-ratio knee & $Q_E = 1.18 \to 1.53$ & AG-CX-06 \ DEC efficiency (anchor / design / ideal) & $0.49/0.70/0.85$ & BU-SX-01 \ end-loss spectrum energy & ~ 428 keV & BU-CX-27 \ global UQ Q_E & 1.317 ± 0.14, $P(Q_E{>}1) = 1.0$ & KX-L1-A5-BU \ Sobol indices (plug ratio) & $S_1 = 0.41$, $S_T = 0.87$ & HX-05 \ net-power probability & $P(\mathrm{net}{>}0) = 73\,\%$ & AG-DR-02 \ length ladder net electric & $+104/+850/+2832$ $\mathrm{MW_e}$ & KX-L1-A9 \ source-to-chip η (DC / AC) & $88.1\,\%$ / $75.9\,\%$ & BU-CX-02 \ SST-fed vs DEC-direct & $85.4\,\%$ / $88.4\,\%$ & MV-CX-04 \ grid-forming nadir lift & $+0.527$ Hz & MV-CX-10 \ bus-fault clamp & $316.2 \to 4$ kA & MV-CX-05 \ PUE pathway & $1.35 \to 1.17$ & BU-CX-14 \ Tier-IV reliability (N+1, 8000 h) & $P = 0.85$ & MV-CX-22 \ site power density & ~ 22 $\mathrm{MW\,acre^{-1}}$ & BU-CX-17 \ mobile tritium source term & ~ 3.6 g & BU-CX-29 \ plug coil field & 26.49 T & BU-L1-A2 \ breeder centrepost field & 16.84 T & [48] \ longtable

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

1 T K Fowler and B G Logan, The tandem mirror reactor, *Comments Plasma Phys. Control. Fusion* **2** 167 (1977).

2 D E Baldwin and B G Logan, Improved tandem mirror fusion reactor, *Phys. Rev. Lett.* **43** 1318 (1979). doi:10.1103/PhysRevLett.43.1318

3 G L Kulcinski and J F Santarius, Advanced fuels under debate, *Nature* **396** 724 (1998). doi:10.1038/25401

4 J F Santarius, G L Kulcinski and L A El-Guebaly, A passively proliferation-resistant fusion power plant, *Fusion Sci. Technol.* **44** 289 (2003). doi:10.13182/FST03-A346

5 M Abdou *et al*, Blanket/first wall challenges and required R&D on the pathway to DEMO, *Fusion Eng. Des.* **100** 2 (2015). doi:10.1016/j.fusengdes.2015.07.021

6 V P Pastukhov, Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential, *Nucl. Fusion* **14** 3 (1974). doi:10.1088/0029-5515/14/1/001

7 R H Cohen, M E Rensink, T A Cutler and A A Mirin, Collisional loss of electrostatically confined species in a magnetic mirror, *Nucl. Fusion* **18** 1229 (1978). doi:10.1088/0029-5515/18/9/005

8 D E Baldwin, End-loss processes from mirror machines, *Rev. Mod. Phys.* **49** 317 (1977). doi:10.1103/RevModPhys.49.317

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

10 D P Grubb *et al*, Thermal-barrier production and identification in a tandem mirror, *Phys. Rev. Lett.* **53** 783 (1984). doi:10.1103/PhysRevLett.53.783

11 M Inutake *et al*, Thermal barrier formation and plasma confinement in the axisymmetrized tandem mirror GAMMA-10, *Phys. Rev. Lett.* **55** 939 (1985). doi:10.1103/PhysRevLett.55.939

12 T Cho *et al*, Observation and control of transverse energy-transport barrier due to the formation of an energetic-electron layer in GAMMA 10, *Phys. Rev. Lett.* **94** 085002 (2005). doi:10.1103/PhysRevLett.94.085002

13 A A Ivanov and V V Prikhodko, Gas-dynamic trap: an overview of the concept and experimental results, *Plasma Phys. Control. Fusion* **55** 063001 (2013). doi:10.1088/0741-3335/55/6/063001

14 Yu V Gott, M S Ioffe and V G Tel'kovskii, Some new results on confinement in magnetic traps, *Nucl. Fusion Suppl.* **3** 1045 (1962).

15 D Endrizzi *et al*, Physics basis for the Wisconsin HTS axisymmetric mirror (WHAM), *J. Plasma Phys.* **89** 975890501 (2023). doi:10.1017/S0022377823000806

16 C B Forest *et al*, Prospects for a high-field, compact break-even axisymmetric mirror and applications, *J. Plasma Phys.* **90** 975900101 (2024). doi:10.1017/S0022377824000916

17 T C Simonen, Three game changing discoveries: a simpler fusion concept?, *J. Fusion Energy* **35** 63 (2016). doi:10.1007/s10894-015-0017-2

18 W M Nevins, A review of confinement requirements for advanced fuels, *J. Fusion Energy* **17** 25 (1998). doi:10.1023/A:1022513215080

19 R F Post, Mirror systems: fuel cycles, loss reduction and energy recovery, in *Nuclear Fusion Reactors* (British Nuclear Energy Society) 88 (1970).

20 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

21 W L Barr and R W Moir, Test results on plasma direct converters, *Nucl. Technol. Fusion* **3** 98 (1983). doi:10.13182/FST83-A20840

22 J-M Rax, E J Kolmes and N J Fisch, Efficiency of direct energy conversion in expanding plasmas, *Phys. Plasmas* **32** 012509 (2025). doi:10.1063/5.0235391

23 N J Fisch and J-M Rax, Interaction of energetic alpha particles with intense lower hybrid waves, *Phys. Rev. Lett.* **69** 612 (1992). doi:10.1103/PhysRevLett.69.612

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 MetroVolt Data-Center Burner: Direct-DC Campus Power, Q_E Closure, and the Plug Requirement in a Low-Neutron D–^{3}He Tandem Mirror

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

Zenodo · doi:10.5281/zenodo.22645695

Read on kronosfusionenergy.com

The Physics De-Risking Register