KRONOS FUSION ENERGY

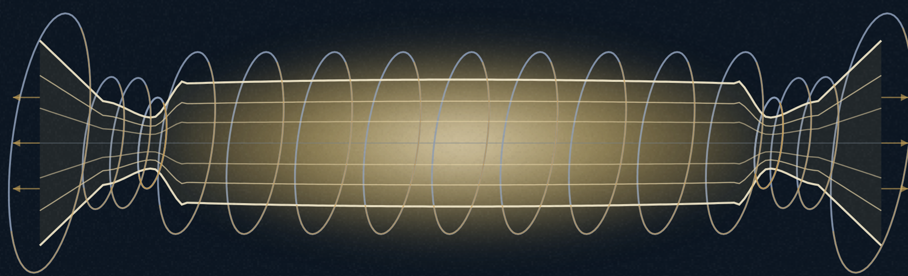

PAPER 2.66 · THE KRONOS 2026 SERIES

Data-Center Integration of a D–^{3}He Tandem-Mirror Burner: Four-Nines Uptime Against Plasma Transients, the PUE-Reduction Pathway, and Waste-Heat Reuse for Behind-the-Meter Dispatch

A compact DHe tandem-mirror burner sited next to the load it serves is not merely a reactor but a facility, and a critical facility is held to reliability, efficiency, and...

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

BY P I FORD

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

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

THE HONESTY STANCE IS THE ASSET

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

TWO EDITIONS

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

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

HOW TO READ THIS VOLUME

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

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

THE FLEET AT A GLANCE

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

HOW THE NUMBERS ARE MARKED

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

ABSTRACT

Data-Center Integration of a D–³He Tandem-Mirror Burner: Four-Nines Uptime Against Plasma Transients, the PUE-Reduction Pathway, and Waste-Heat Reuse for Behind-the-Meter Dispatch

A compact D^3He tandem-mirror burner sited next to the load it serves is not merely a reactor but a *facility*, and a critical facility is held to reliability, efficiency, and heat-management standards that no plasma-physics result answers. We formalise the facility-integration problem for the Kronos MetroVolt burner (design point $\approx 850\,MW$ net electric at a recirculating fraction of 0.76; burner physics $Q_E = 1.318$, plug field 26.49 T, direct-conversion efficiency $\eta_{\mathrm{DEC}} = 0.49$) integrated with a co-located data center, and we resolve three questions with explicit mathematics and a GPU-accelerated Monte-Carlo campaign. *Reliability.* We derive the fleet-availability model as a repairable k-out-of-M system, $A_{\mathrm{sys}} = \sum_{k=N}^{M} \binom{M}{k} a^k (1-a)^{M-k}$, whose leading unavailability $\sim \binom{M}{r+1}(1-a)^{r+1}$ exposes exactly two levers—per-module reliability and redundancy order r. A 2500-trial Monte-Carlo model resolves the Tier-IV service level (99.995) precisely: a baseline $N{+}1$ fleet of 2000 hour-MTBF modules clears the target only 13 of the time, whereas higher-reliability modules ($\mathrm{MTBF} \approx 8000\,hour$, 85) or an $N{+}2$ configuration (77) meet it—so four-nines availability is *engineered*, grid-tied, with an honest islanded ceiling near 99.9 . *Placement.* Behind-the-meter delivery raises delivered availability by 6.3 points (98.5 vs 92.2) and cuts time-to-power from 78 to 30 months by skipping the interconnection queue. *Efficiency and heat.* A composed lever stack takes facility power usage effectiveness from $\mathrm{PUE} = 1.35$ to 1.17, reaching energy-reuse effectiveness $\mathrm{ERE} = \mathrm{PUE} - \mathrm{ERF} \approx 0.92$ at 25% reuse uptake; a double-effect LiBr absorption chiller (cooling coefficient $\mathrm{COP_{abs}} = 1.35$) covers the full 100 facility cooling load on 74 of the ~ 800 $\mathrm{MW_{th}}$ reject, cutting electric cooling by 91 ($16.7 \to 1.5$ MW_e); and direct-liquid capture at 55 °C lifted to 75 °C at a heat-pump coefficient $\mathrm{COP_{HP}} = 8.7$ delivers a useful-heat reuse factor $\mathrm{ERF} \approx 1.0$, with the capture-temperature dependence following $\mathrm{COP_{HP}} = \eta_{II}\, T_h/(T_h - T_c)$ at $\eta_{II} \approx 0.5$. The burner integrates as a facility that meets its service standards, at competitive efficiency, and that exports rather than dumps its heat.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645919 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: data-center integration · service-level availability · Tier-IV · behind-the-meter dispatch · PUE · ERE · waste-heat reuse · absorption cooling · tandem-mirror burner

CONTENTS

Table of Contents

Data-Center Integration of a D–^{3}He Tandem-Mirror Burner: Four-Nines Uptime Against Plasma Transients, the PUE-Reduction Pathway, and Waste-Heat Reuse for Behind-the-Meter Dispatch 3
Introduction 6
The facility as a system 7
Governing equations and the formal problem 8
Computational methodology: the NVIDIA GPU HPC campaign 11
Results 12
Discussion 21
Limitations 22
Conclusion 23
References 28
One programme, many papers 31

NOMENCLATURE

Symbols, units, and the names of things

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

§ 1

Introduction

The fastest-growing electrical load of the present decade is computation. Hyperscale and artificial-intelligence data centers are being built at hundreds of megawatts per campus, and their demand is firm, dense, and rising faster than transmission can be permitted and built [1,2,3]. A generator that could sit *next to* such a load—compact enough to co-locate, firm enough to run baseload, and clean enough to permit in a populated area—would collapse the transmission problem and let the plant and the load be engineered as one system. The Kronos MetroVolt concept is exactly that: a neutron-lean D^3He tandem-mirror burner whose charged-particle power is recovered by direct energy conversion into a high-voltage DC bus and delivered behind-the-meter to a co-located data center [27,28,32].

The burner physics is settled elsewhere in this series. The MetroVolt operating point is a tandem mirror with an electrical gain $Q_E = 1.318$, a central-cell burn sustained against end-loss by a 26.49 T high-field plug, and a direct-energy-converter efficiency $\eta_{\mathrm{DEC}} = 0.49$; the neutron-lean fuel choice ($\sim 90\%$ of the fusion power in charged particles) is what makes urban and demand-adjacent siting credible [27,31]. At the module scale the plant delivers roughly 850 MW net electric at a recirculating fraction of 0.76, and modules stack to a facility's demand. None of those numbers, however, answers the questions a data-center operator actually asks. Can the plant hold the *service-level agreement* that a critical facility is written to—four or five nines of availability—given that a single fusion module will occasionally trip? Can the integrated site run at a *power usage effectiveness* (PUE) competitive with the best air- or liquid-cooled data centers? And can the plant's very large reject heat be turned from a cooling liability into a delivered product? These are facility-integration questions, and they are the subject of this paper.

Prior art. The three questions each have a mature literature that we adopt rather than reinvent. Data-center availability is codified by the Uptime Institute's tier system, whose Tier-IV class targets 99.995 availability—about 26 minutes of downtime a year [4]; the reliability mathematics of repairable redundant systems that underlies it is classical [14,15,16]. Facility efficiency is measured by PUE and, where heat is exported, by the energy-reuse metrics ERE and ERF defined by The Green Grid [5,6], and the measures that reduce PUE—DC-native power conversion, direct-liquid and immersion cooling, and heat reuse—are individually validated in practice [7,8,10,9]. Waste-heat-driven cooling by lithium-bromide absorption chillers, and low-lift heat pumps for district-grade heat export, are standard thermal engineering [11,12,13]. What is new here is not any one of these components but their *composition* onto a compact fusion plant that makes its own firm power and its own high-grade heat, and the quantitative demonstration that the composition meets the facility standards.

Contribution. This paper (i) states the facility-integration problem as a set of governing equations spanning repairable-fleet reliability, behind-the-meter delivery, facility efficiency, and waste-heat thermodynamics (§3); (ii) resolves the binding reliability question with a GPU-accelerated 2500-trial Monte-Carlo availability model whose levers are made explicit (§4, §5.1); (iii) quantifies the behind-the-meter placement advantage in both delivered availability and schedule (§5.2); (iv) derives and closes the PUE-reduction pathway to 1.17 and the energy-reuse effectiveness ≈ 0.92 (§5.3); and (v) closes the waste-heat balance, showing that reject heat covers the facility's own cooling and exports useful heat at a reuse factor near unity (§5.4). Every number traces to a frozen anchor in the Kronos Physics De-Risking Register [26]; the values used here are the MetroVolt facility-integration anchors MV-L1-A4 and MV-CX-12/13/14/15/18/19/22/23/24.

§ 2

The facility as a system

Figure 1 fixes the object of study. The burner core produces charged-particle power that a direct-energy converter turns into a high-voltage DC bus; that bus feeds a co-located data center's dense compute halls behind the meter, with a modest grid tie retained for export, ancillary services, and ride-through. The plant's mid-grade reject heat—of which some 800 is available at the module scale—drives an absorption chiller that cools the facility, while direct-liquid capture of the IT heat feeds a low-lift heat pump that exports useful heat. Reliability, efficiency, and heat thus close three loops on the same hardware, and the integration is judged on all three.

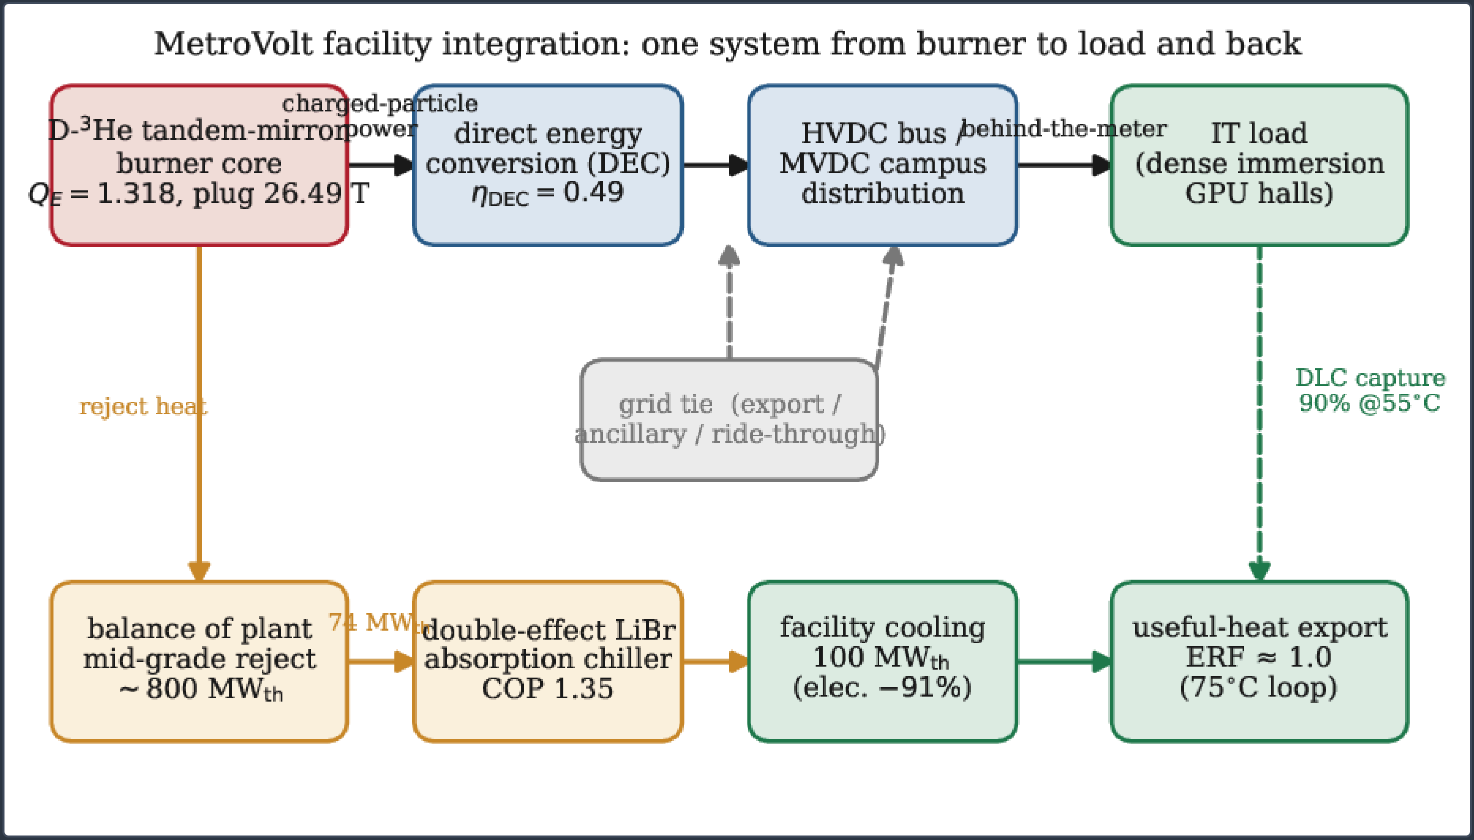

(Schematic.) MetroVolt facility integration. Top: the generation chain—burner core, direct energy conversion, HVDC/MVDC campus distribution, IT load—delivered behind the meter. Bottom: the heat chain— mid-grade reject heat drives a double-effect absorption chiller for facility cooling, and direct-liquid capture of the IT heat feeds a low-lift heat pump for useful-heat export. A grid tie is retained for export, ancillary services, and ride-through. Design-point labels are frozen anchors; no measured values are plotted.

The plant is a *baseload* source: it runs continuously at a high capacity factor, and the millisecond-class swings of an AI training load are absorbed by a small electrical buffer rather than by throttling the plasma (a separate result in this series [28]). For the facility-integration analysis this matters because it lets us treat the burner as a firm, high-capacity-factor generator whose *availability*—not its instantaneous dispatchability—is the reliability quantity of interest.

§ 3

Governing equations and the formal problem

We now state the four models—availability, placement, efficiency, and heat—as explicit equations. Symbols are collected in Table 1 where numerical.

REPAIRABLE-FLEET AVAILABILITY

A single power module is a two-state repairable component. With failure rate $\lambda = 1/\mathrm{MTBF}$ and repair rate $\mu = 1/\mathrm{MTTR}$, the continuous-time Markov chain on $\{\mathrm{up}, \mathrm{down}\}$ has the stationary availability

$$a \;=\; \frac{\mu}{\lambda+\mu} \;=\; \frac{\mathrm{MTBF}}{\mathrm{MTBF}+\mathrm{MTTR}}.$$

A facility that requires N carrying modules to serve its firm load installs $M = N + r$, where r is the redundancy order ($r{=}1$ for $N{+}1$, $r{=}2$ for $N{+}2$). Modeling the M modules as independent and identically distributed, the number up K is binomial, $K \sim \mathrm{Bin}(M, a)$, and the fleet delivers firm power iff at least N modules are up. The steady-state fleet availability is therefore the k-out-of-M reliability

$$A_{\mathrm{sys}} \;=\; \Pr(K \geq N) \;=\; \sum_{k=N}^{M} \binom{M}{k} a^{k} (1-a)^{M-k},$$

and the expected annual downtime is

$$D \;=\; (1 - A_{\mathrm{sys}})\, T_{\mathrm{yr}}, \qquad T_{\mathrm{yr}} = 525600\,\mathrm{minute}.$$

Writing $q \equiv 1 - a$ for the per-module unavailability, a Taylor expansion of (2) about $q = 0$ gives the leading fleet unavailability

$$1 - A_{\mathrm{sys}} \;=\; \binom{M}{r+1} q^{\,r+1} \;+\; O\big(q^{\,r+2}\big).$$

Equation (4) is the analytic origin of the design levers. The fleet unavailability scales as the $(r{+}1)$-th power of the module unavailability: adding one spare raises the *order* of the leading term (the redundancy lever), and reducing q by improving MTBF suppresses it geometrically (the reliability lever). No third lever exists at this order—a fact the Monte Carlo will confirm quantitatively.

Equation (2) assumes ample repair capacity (at least $r{+}1$ repair channels) so that the module states remain independent. When repair is resource-limited the correct object is the full birth–death chain on the number of modules down, $j \in \{0, \dots, M\}$, with up-transition rate $(M - j)\lambda$ and down-transition rate $\min(j, s)\,\mu$ for s repair crews; its stationary distribution π_j gives $A_{\mathrm{sys}} = \sum_{j \leq r} \pi_j$, and reduces to (2) in the ample-repair limit $s \geq r{+}1$. The Monte-Carlo model of §4 does not take that limit: it simulates finite repair crews and staggered preventive maintenance directly, which is why its Tier-IV probabilities are the design-authoritative figures and Eq. (2) is the transparent model that explains them.

BEHIND-THE-METER DELIVERY AND SCHEDULE

Where the plant sits relative to the revenue meter changes the availability the load actually sees. The delivered availability factors into generation and path terms,

$$A_{\mathrm{del}} \;=\; A_{\mathrm{gen}}\, A_{\mathrm{path}},$$

where $A_{\mathrm{gen}} = A_{\mathrm{sys}}$ from (2) and A_{path} is the availability of the electrical path between generation and load. A grid-tied path interposes transmission and distribution assets with their own outage exposure, so $A^{\mathrm{grid}}_{\mathrm{path}} < 1$; a behind-the-meter path deletes that segment, so $A^{\mathrm{BTM}}_{\mathrm{path}} \approx 1$ and the path also avoids the T&D loss ℓ in delivered power, $P_{\mathrm{del}} = P_{\mathrm{gen}}(1 - \ell)$. The schedule to first power adds an interconnection-queue term only on the grid path,

$$T_{\mathrm{power}} \;=\; T_{\mathrm{build}} \;+\; \mathbb{1}[\mathrm{grid}]\, T_{\mathrm{queue}},$$

with $T_{\mathrm{build}} \approx 30$ **months** and an interconnection queue $T_{\mathrm{queue}} \approx 48$ **months** for a transmission-connected plant [18].

FACILITY EFFICIENCY: PUE, ERE, AND ERF

Let $P_{\rm IT}$ be the compute (IT) power and decompose the total facility power drawn as

$$P_{\rm tot} \;=\; P_{\rm IT} + P_{\rm cool} + P_{\rm conv} + P_{\rm aux},$$

with $P_{\rm cool}$ the cooling power, $P_{\rm conv}$ the conversion and distribution loss, and $P_{\rm aux}$ auxiliary and lighting. The power usage effectiveness is the dimensionless ratio

$$\mathrm{PUE} \;\equiv\; \frac{P_{\rm tot}}{P_{\rm IT}} \;=\; 1 + \frac{P_{\rm cool} + P_{\rm conv} + P_{\rm aux}}{P_{\rm IT}}.$$

An efficiency measure i that removes overhead ΔP_i contributes an additive first-order decrement $\Delta\mathrm{PUE}_i = -\Delta P_i / P_{\rm IT}$, and to the extent the measures act on distinct overhead terms they compose,

$$\mathrm{PUE}_{\rm final} \;=\; \mathrm{PUE}_0 + \sum_i \Delta\mathrm{PUE}_i.$$

When the site also *exports* useful energy—here, heat—the more complete metric is the energy-reuse effectiveness. Defining the reuse factor $\mathrm{ERF} \equiv P_{\rm reuse}/P_{\rm IT}$ (useful exported energy per unit IT energy), the energy-reuse effectiveness is

$$\mathrm{ERE} \;\equiv\; \frac{P_{\rm tot} - P_{\rm reuse}}{P_{\rm IT}} \;=\; \mathrm{PUE} - \mathrm{ERF}.$$

A site is a net energy exporter over its overhead when $\mathrm{ERF} > \mathrm{PUE} - 1$, i.e. when the reuse it achieves exceeds its non-IT overhead.

WASTE-HEAT THERMODYNAMICS

Absorption cooling. A double-effect LiBr absorption chiller is a heat-driven machine whose cooling coefficient of performance is the ratio of delivered cooling to driving (generator) heat, $\mathrm{COP}_{\rm abs} \equiv Q_c/Q_{\rm gen}$. To meet a facility cooling load Q_c the driving heat required is

$$Q_{\rm gen} \;=\; \frac{Q_c}{\mathrm{COP}_{\rm abs}},$$

drawn from the plant's mid-grade reject. The electric demand it displaces is that of a vapour-compression chiller of electrical coefficient COP_e, namely $P_e^{\rm elec} = Q_c/\mathrm{COP}_e$; the absorption plant's residual electric parasitics $P_e^{\rm abs}$ (solution pumps, tower fans) are small, so the cooling-electricity reduction is

$$\eta_{\rm cool} \;=\; 1 - \frac{P_e^{\rm abs}}{P_e^{\rm elec}} \;=\; 1 - \frac{P_e^{\rm abs}\,\mathrm{COP}_e}{Q_c}.$$

Heat reuse. IT electricity is, by conservation, almost entirely converted to heat; the constraint on reuse is therefore not quantity but *grade*. Capturing a fraction $f_{\rm cap}$ of the IT heat at a coolant temperature T_c and lifting it to a useful delivery temperature T_h with a heat pump, the heating coefficient of performance is bounded by the Carnot value $T_h/(T_h - T_c)$ and, at a second-law efficiency η_{II}, is

$$\mathrm{COP}_{\rm HP} \;=\; \eta_{II}\,\frac{T_h}{T_h - T_c} \qquad (T \text{ in kelvin}).$$

Feeding the captured heat $Q_{\rm cap} = f_{\rm cap} P_{\rm IT}$ as the heat-pump source, the useful heat delivered at T_h is $Q_{\rm reuse} = Q_{\rm cap}\,\mathrm{COP}_{\rm HP}/(\mathrm{COP}_{\rm HP} - 1)$, so the reuse factor is

$$\mathrm{ERF} \;=\; \frac{Q_{\rm reuse}}{P_{\rm IT}} \;=\; f_{\rm cap}\,\frac{\mathrm{COP}_{\rm HP}}{\mathrm{COP}_{\rm HP} - 1}.$$

Equation (13) makes the capture temperature the governing variable: a warm-water capture at $T_c = 55\ ^\circ C$ needs only a 20 K lift to a 75 °C export loop and reaches a high $\mathrm{COP}_{\rm HP}$, whereas an air-cooled capture at 30 °C demands a 45 K lift and a $\mathrm{COP}_{\rm HP}$ less than half as large—which is why reuse is practical for liquid-captured heat and not for air.

CAPACITY-FACTOR FIRMING

Finally, firming a continuous load P_L from a source of capacity factor CF requires energy adequacy, $P_{\text{nameplate}}\,\mathrm{CF} \geq P_L$, hence a nameplate multiple

$$m \equiv \frac{P_{\text{nameplate}}}{P_L} \geq \frac{1}{\mathrm{CF}},$$

plus a storage reservoir sized to bridge the low-generation tail of the source's output distribution. For a baseload source ($\mathrm{CF} \to 1$) the multiple approaches unity and only a short buffer is needed; for a variable-renewable source ($\mathrm{CF} \sim 0.25$) the multiple is several-fold and the storage runs to hundreds of hours. Equation (15) is the physical statement of the co-located-baseload case for a firm compute load.

§ 4

Computational methodology: the NVIDIA GPU HPC campaign

The facility-integration results were produced on the Kronos NVIDIA A100/H100 GPU HPC platform (an owned fleet of 8× A100-80GB nodes with H100 clusters, plus single A100/H100/A10 devices), with the reduced-order and closed-form gates cross-checked on an in-house HPC node (56-thread, 768 GB). No economic quantity of any kind enters the computations reported here. Three computational elements underpin the paper.

(i) Monte-Carlo availability simulator (MV-CX-22). The binding reliability question is answered by a discrete-event availability model. Each of M modules is an alternating renewal process: times-to-failure are drawn $\sim \mathbf{Exp(MTBF)}$ and times-to-repair $\sim \mathbf{Exp(MTTR)}$, with staggered preventive maintenance and a finite repair crew, propagated over a one-year horizon. Each yearly realisation returns a time-averaged fleet availability by evaluating the k-out-of-M up-condition of Eq. (2) on the event timeline; the Tier-IV attainment probability is the fraction of $\mathbf{2500}$ independent realisations whose annual availability clears 99.995 . The realisations are mutually independent, so the workload is embarrassingly parallel: one realisation per GPU thread, with the full parameter sweep (baseline $N{+}1$, $N{+}2$, and high-MTBF configurations) resident in device memory and reduced on-device. The analytic closed form (2)–(4) is evaluated in parallel as an independent cross-check and agrees with the Monte-Carlo mean availability to within sampling error.

(ii) Burner operating point (KRONOS Research Toolkit). The facility net-power and reject-heat budgets inherit the frozen burner design point from the KRONOS Research Toolkit `solve_mirror` 0-D tandem-mirror power balance, which fixes $Q_E = \mathbf{1.318}$, the recirculating fraction $\mathbf{0.76}$, and the $\sim\mathbf{850}\,MW$ net electric output, together with the plug-density sensitivity scan (MV-CX-26) that establishes the operating point as a regulated requirement rather than a free parameter. The underlying neutron-lean physics anchors that make urban siting admissible are fixed upstream in the series by the reference-code stack (nonlinear gyrokinetics with CGYRO, Monte-Carlo neutronics with OpenMC), and are consumed here as boundary data.

(iii) Deterministic systems and thermal models. The behind-the-meter delivery model (Eqs. 5–6), the PUE lever algebra (Eqs. 7–10), the absorption-cycle and heat-pump balances (Eqs. 11–14), and the capacity-factor firming relation (Eq. 15) are deterministic and are evaluated in closed form; their outputs are the frozen anchors MV-CX-12/13/14/15/18/23. Two learning surrogates support the analysis without setting any headline number: a tabular Q-learning dispatch policy (MV-CX-31) that independently recovers a low effective PUE by preferring waste-heat absorption cooling, and a predictive-maintenance classifier (MV-CX-33, logistic regression) that flags impending module failure at an area under the ROC curve of $\mathbf{0.99}$ with a median lead time of order $\mathbf{17}$ steps, converting a fraction of unplanned trips into planned maintenance and thereby raising the effective module availability that enters Eq. (1). All anchors are regenerable from the archived scripts (§8).

§ 5

Results

AVAILABILITY AND THE TIER-IV SERVICE LEVEL

Table 1 fixes the service-level targets. 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; the Tier-IV target of 99.995 is about **26** minutes. 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 through Eq. (2).

SLA / CONFIGURATION	METRIC	VALUE	
99.9	allowed downtime	526	
99.99	allowed downtime	52.6	
99.995 (Tier-IV)	allowed downtime	≈ 26.3	$\mathrm{min\,yr^{-1}}$
99.999	allowed downtime	5.3	
single unit (FOAK $\rightarrow$ NOAK)	availability a	$0.90 \rightarrow 0.96$	
per-module (baseload target)	availability a	0.995	
N**+1, MTBF 2000 hour**	P(meet Tier-IV)	0.13	
N**+2, MTBF 2000 hour**	P(meet Tier-IV)	0.77	
N**+1, MTBF 8000 hour**	P(meet Tier-IV)	0.85	

Service-level targets, module parameters, and the reliability model (anchors MV-L1-A4, MV-CX-19, MV-CX-22). Downtime is $(1 - A)\,T_{\mathrm{yr}}$ from Eq. (3).

Figure 2 plots the fleet availability model of Eq. (2) for $N = 6$ as annual downtime versus per-module availability, for the three redundancy orders. The curves make the q^{r+1} scaling of Eq. (4) visible: each added spare rotates the curve to a steeper power law, and at a fixed module availability the fleet downtime drops by orders of magnitude as r increases. At the baseload per-module target $a = 0.995$ (a 2000 hour-MTBF module with a $\sim$10 *hour* repair) an N+1 fleet sits near 99.94 —close to but short of Tier-IV—while an N+2 fleet, or a high-reliability module at $a \approx 0.999$, clears the 99.995 line. This is the analytic statement of the design conclusion.

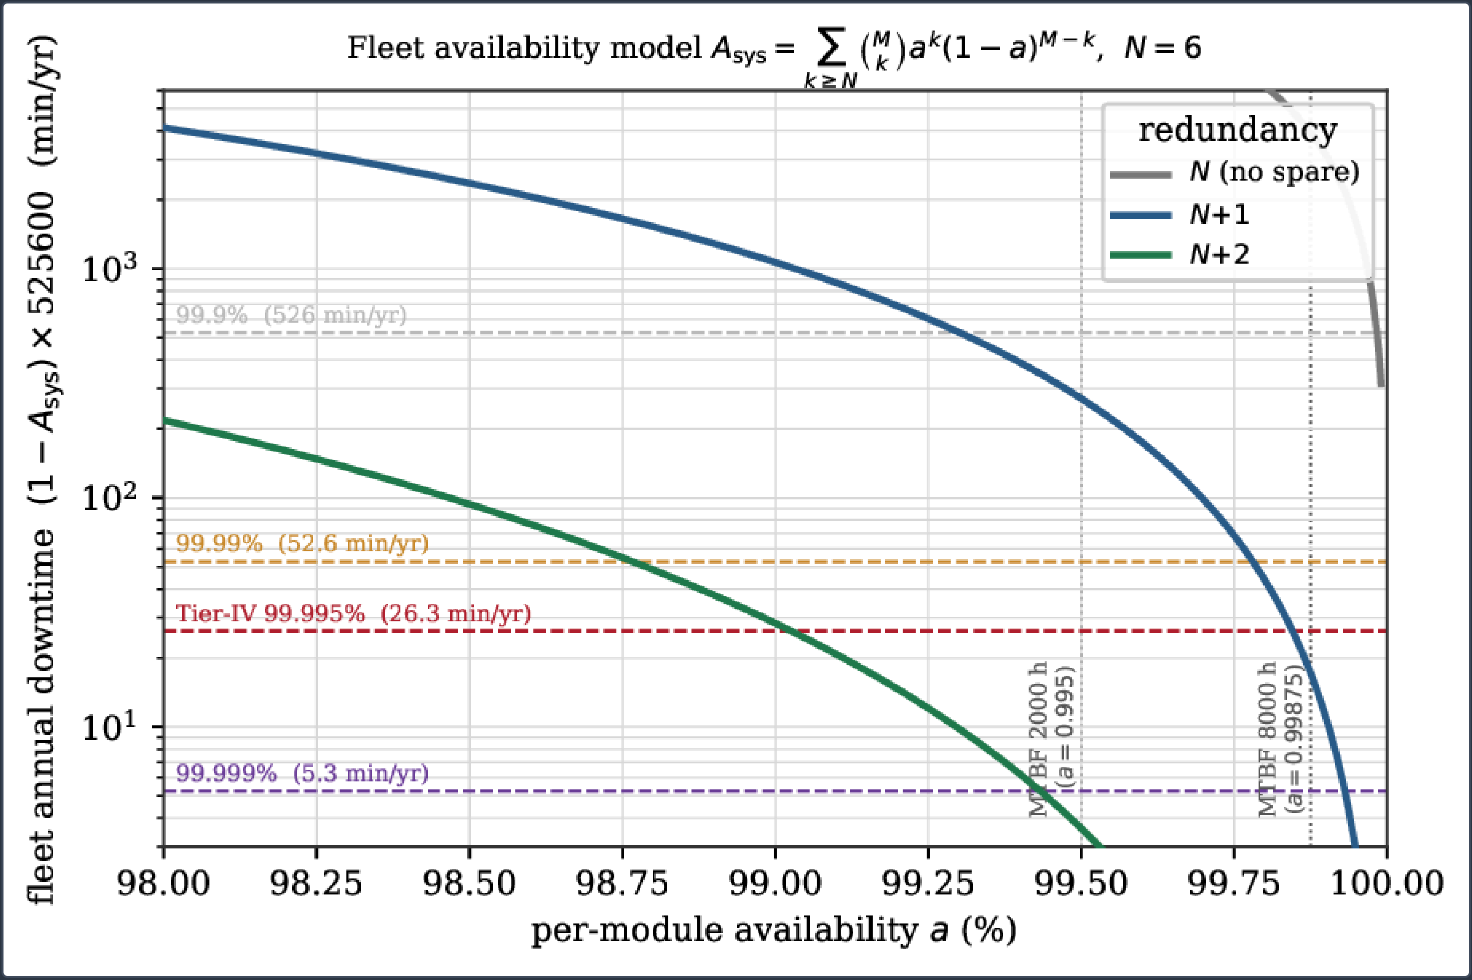

Fleet availability model (Eq. 2, $N = 6$): annual downtime versus per-module availability for redundancy orders $r = 0, 1, 2$. Dashed horizontals are the SLA downtime budgets (Table 1); dotted verticals mark the 2000 hour and 8000 hour MTBF module points. Added redundancy steepens the power law (Eq. 4); higher module MTBF slides left along it. Both routes cross the Tier-IV line.

The Monte-Carlo model resolves the probabilities the closed form only bounds (Figure 3). A baseline $N+1$ fleet built from 2000 hour-MTBF modules meets the Tier-IV target only 13 of the time—it misses. Raising module reliability to $\mathbf{MTBF} \approx 8000\,\textit{hour}$ lifts the probability to 85 , and adding redundancy ($N+2$ at the baseline MTBF) reaches 77 . The two levers of Eq. (4) are therefore not merely asymptotic: they are the physical knobs that move a real fleet across the line. 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 (§7).

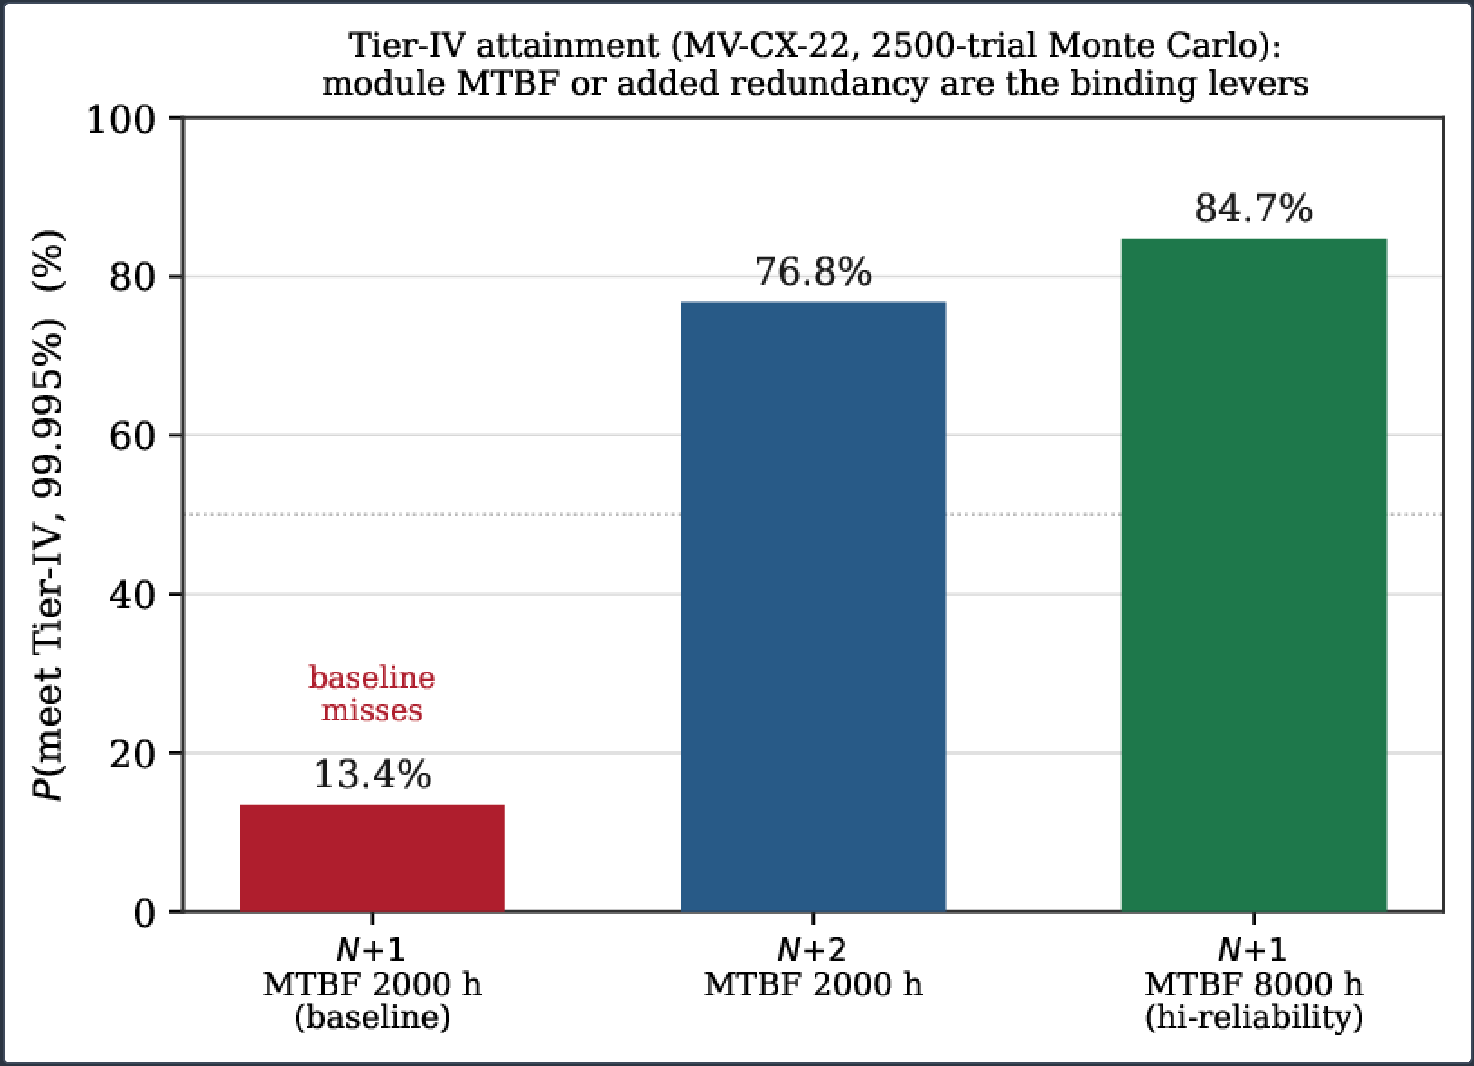

Tier-IV attainment probability from the **2500**-trial Monte-Carlo availability model (MV-CX-22). A baseline $N{+}1$ fleet misses; higher-reliability modules or $N{+}2$ redundancy clear the 99.995 target. Module MTBF and redundancy order are the binding levers.

BEHIND-THE-METER PLACEMENT

Where the plant sits relative to the utility meter is itself an availability lever (Figure 4). Delivering power behind the meter—directly to the co-located load, per Eq. (5) with $A^{\mathrm{BTM}}_{\mathrm{path}} \approx 1$— avoids transmission and distribution outage exposure and loss and raises delivered availability by **6.3** points, **98.5** against 92.2 for a grid-tied path (MV-CX-18). It also collapses the schedule of Eq. (6): behind-the-meter time-to-power is **30** months against **78** for a grid interconnection, skipping a roughly 48 months queue [18]. A modest grid tie is still wanted—for export, for ancillary services, and for the islanded-to-grid ride-through that lifts the reliability ceiling to four nines—so the design answer is a hybrid: behind-the-meter primary delivery with a grid tie for redundancy and export. That hybrid is what makes the Tier-IV result of §5.1 achievable in practice.

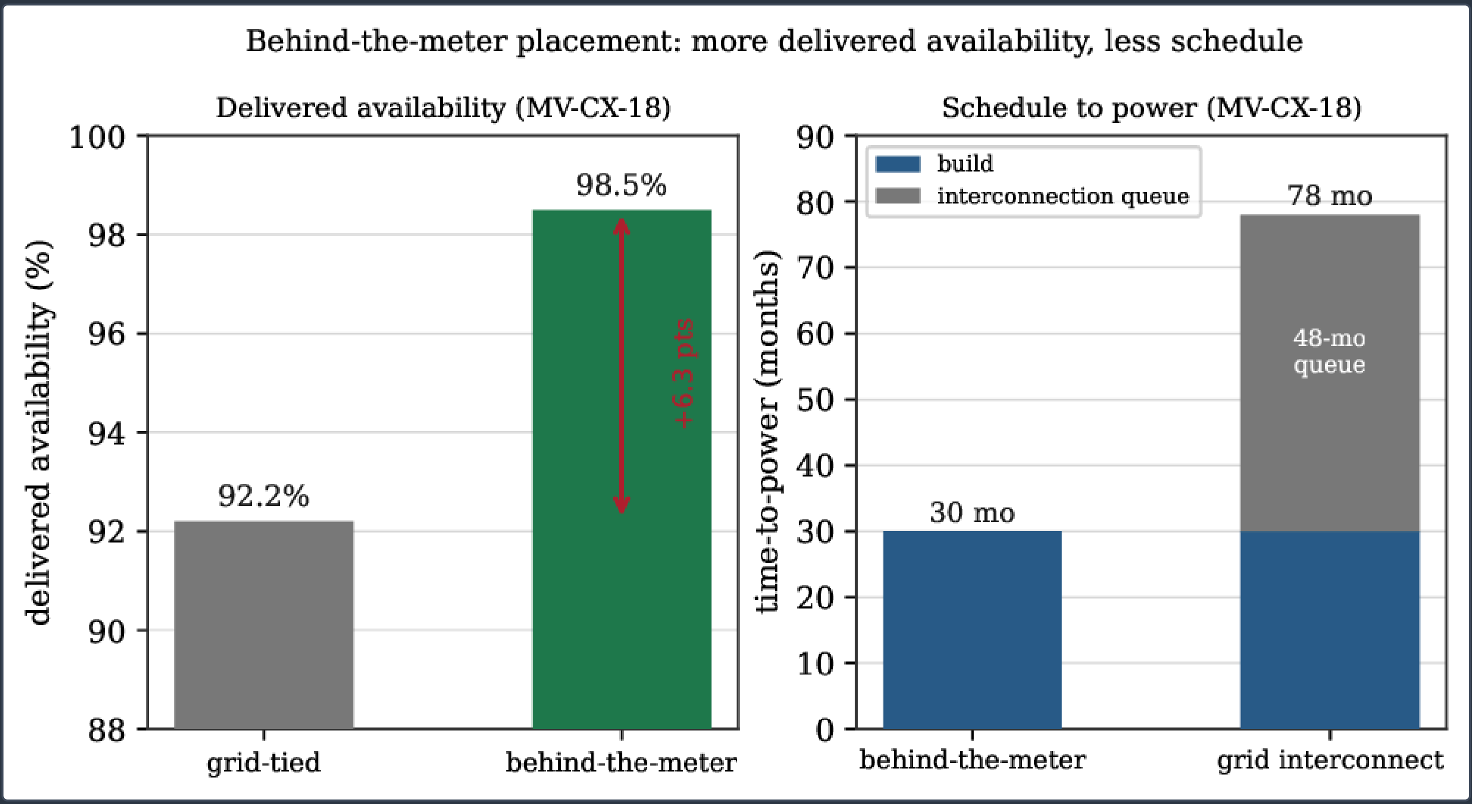

Behind-the-meter versus grid-tied dispatch (MV-CX-18). Left: delivered availability rises **6.3** points by deleting the T&D path (Eq. 5). Right: time-to-power falls from **78** to **30** months by skipping the ~**48 months** interconnection queue (Eq. 6).

THE PUE-REDUCTION PATHWAY

Facility efficiency follows the concrete lever stack of Eqs. (8)–(9), taking the site from a $\mathrm{PUE} = 1.35$ baseline to 1.17 (Table 2, Figure 5). Four levers compose on distinct overhead terms: DC-native power conversion removes a rectification stage ($\Delta\mathrm{PUE} = -0.05$); direct-liquid cooling removes fan power (-0.07); waste-heat absorption cooling displaces electric chillers (-0.04); and auxiliary and lighting optimisation takes the remainder (-0.02). Because the plant also exports useful heat, the completed metric is the energy-reuse effectiveness of Eq. (10): at a modest reuse factor $\mathrm{ERF} = 0.25$ (a 25% heat-reuse uptake) the site reaches $\mathrm{ERE} = 1.17 - 0.25 = 0.92$. Since $\mathrm{ERF} = 0.25$ exceeds the non-IT overhead $\mathrm{PUE} - 1 = 0.17$, the site exports more energy as useful heat than it spends on overhead—it is a net energy exporter above its own facility load. Every one of these measures is a standard, individually validated data-center technique [7,8,9]; the novelty is applying them to a plant that makes its own power and its own heat.

LEVER	ΔPUE
DC-native power conversion	-0.05
direct-liquid cooling (fans↓)	-0.07
waste-heat absorption cooling	-0.04
auxiliary / lighting optimisation	-0.02
baseline → final	$1.35 \to 1.17$
energy-reuse effectiveness (ERF = 0.25)	$\mathrm{ERE} \approx 0.92$

PUE lever stack, $1.35 \to 1.17$ (MV-CX-14), composed by Eq. (9).

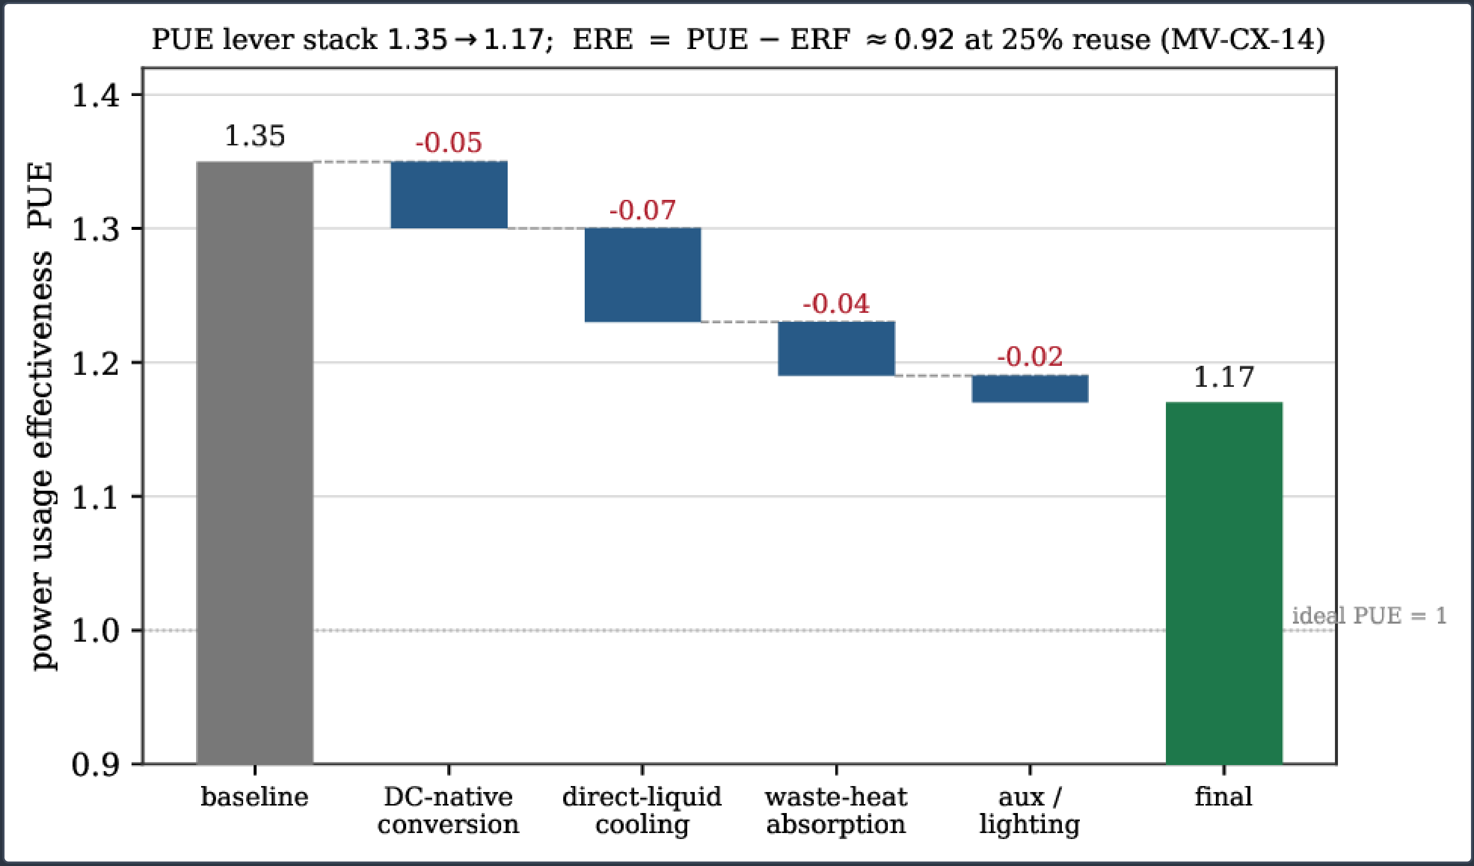

PUE pathway (MV-CX-14): the lever stack of Eq. (9) takes facility PUE from 1.35 to 1.17; with a 25% reuse factor the energy-reuse effectiveness of Eq. (10) reaches $\mathrm{ERE} \approx 0.92$.

WASTE-HEAT REUSE

The plant's largest by-product is heat, and integration turns it from a cooling burden into a delivered output (Figure 6). Two mechanisms do the work.

First, on the cooling side, the plant's own mid-grade reject heat drives a double-effect LiBr absorption chiller of $\mathrm{COP_{abs}} = 1.35$. By Eq. (11) the full 100 facility cooling load needs only $Q_{\mathrm{gen}} = 100/1.35 = 74$ $\mathrm{MW_{th}}$ of driving heat—a small fraction of the ~ 800 $\mathrm{MW_{th}}$ available. By Eq. (12), against a vapour-compression chiller of electrical coefficient $\mathrm{COP}_e \approx 6$ ($P_e^{\mathrm{elec}} = 100/6 = 16.7$ MW_e) and an absorption residual of $P_e^{\mathrm{abs}} \approx 1.5$ MW_e , the electric cooling load falls by 91 , $16.7 \to 1.5$ MW_e (MV-CX-12). The facility cools itself with heat it was going to throw away.

Second, on the capture side, direct-liquid cooling collects $f_{\mathrm{cap}} = 0.90$ of the IT heat at a useful $T_c = 55\,^{\circ}C$. A small heat-pump lift to $T_h = 75\,^{\circ}C$ has, by Eq. (13) at $\eta_{II} \approx 0.5$, a coefficient $\mathrm{COP_{HP}} = 0.5 \times 348.15/20 = 8.7$; then Eq. (14) gives a useful-heat reuse factor $\mathrm{ERF} = 0.90 \times 8.7/7.7 \approx 1.0$—roughly one unit of exported useful heat per unit of IT electricity (MV-CX-13). Figure 7 shows why the capture temperature is the governing variable: the measured ladder of capture technologies (air, rear-door, direct-liquid, immersion) falls exactly on the $\mathrm{COP_{HP}} = \eta_{II} T_h/(T_h - T_c)$ curve at $\eta_{II} \approx 0.5$, and air cooling at 30 °C lands in the low-COP regime where reuse is impractical, whereas direct-liquid at 55 °C sits in the practical band.

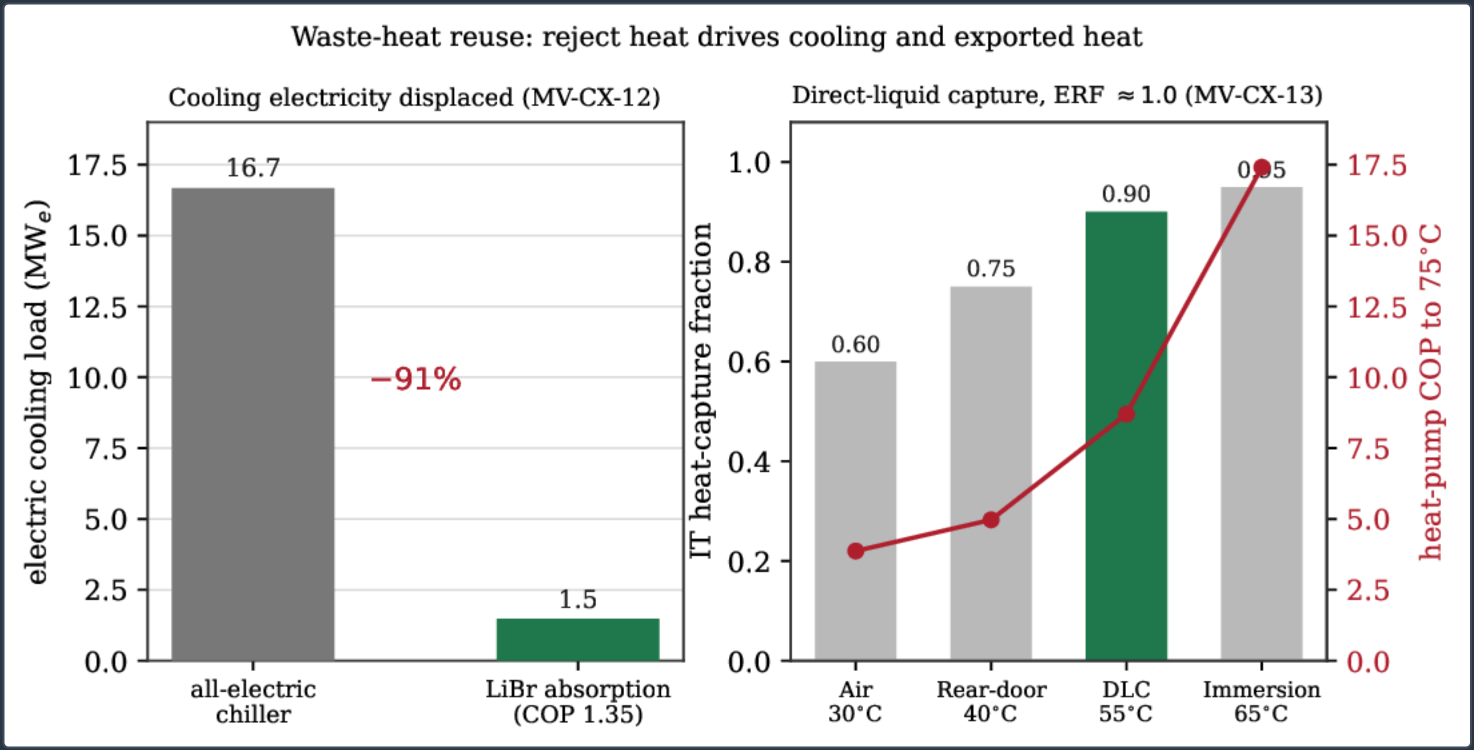

Waste-heat reuse. Left: a double-effect absorption chiller cuts the electric cooling load from **16.7** to 1.5 , a 91 reduction (Eqs. 11–12, MV-CX-12). Right: the heat-capture ladder—direct-liquid cooling captures **90%** of the IT heat at 55 °C and, with the heat-pump COP of Eq. (13), reaches a reuse factor near unity (MV-CX-13).

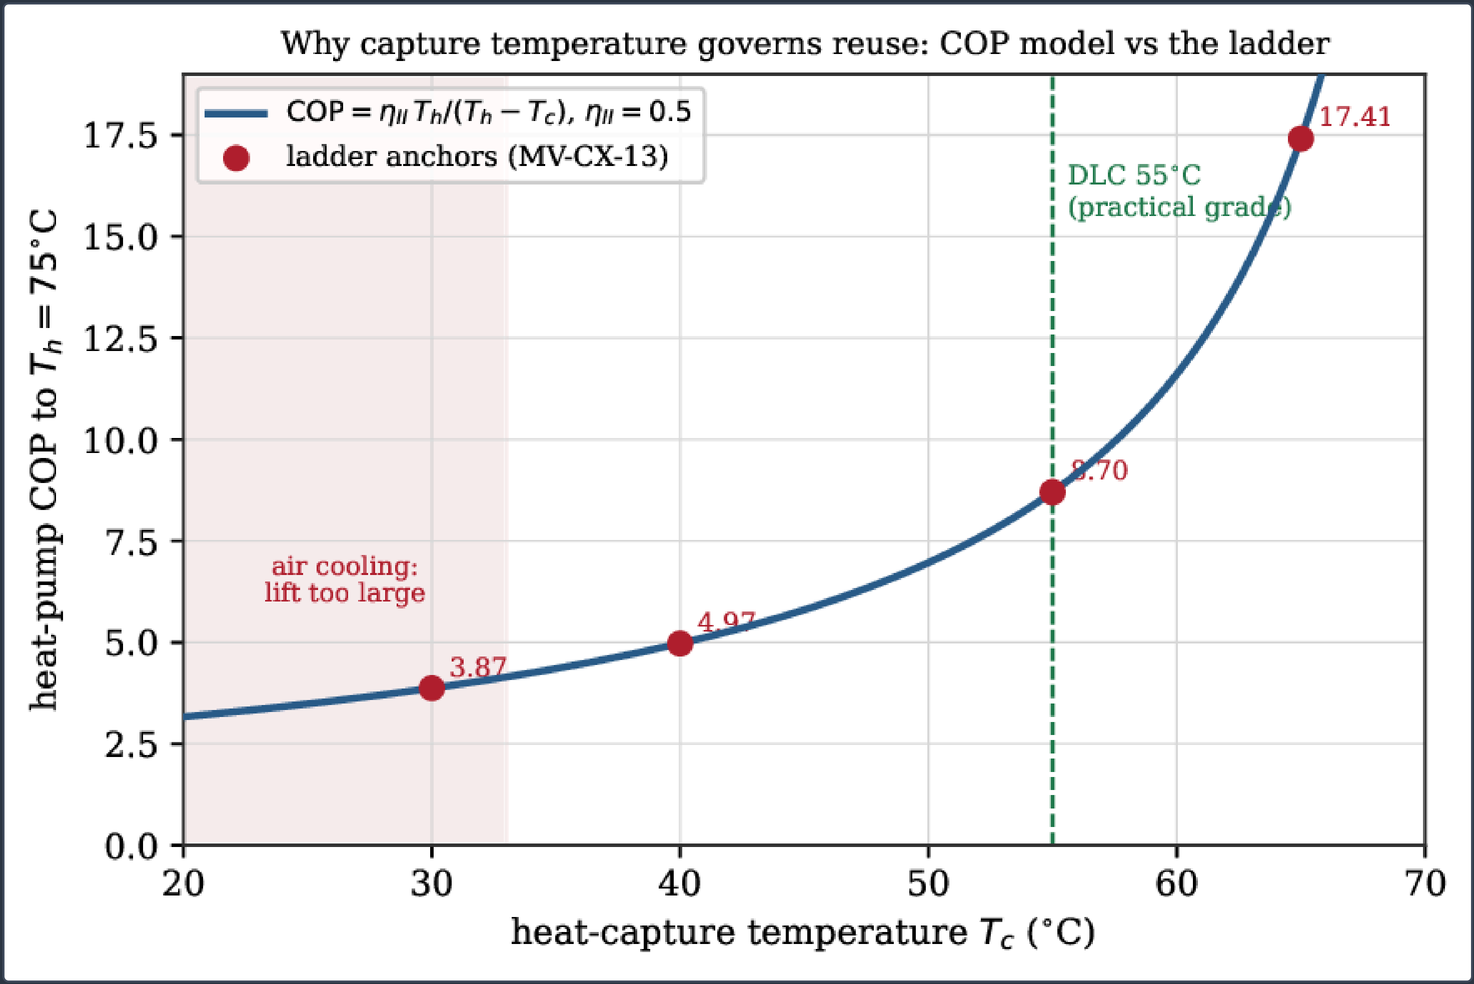

Heat-pump coefficient of performance versus capture temperature, model $COP_{HP} = \eta_{II} T_h/(T_h - T_c)$ at $T_h = 75\,^{\circ}C$ and $\eta_{II} = 0.5$ (Eq. 13), with the measured capture ladder overlaid (MV-CX-13). The capture temperature governs the achievable COP and hence the practicality of reuse: air cooling (shaded) is in the low-COP regime; direct-liquid at 55 °C is not.

FIRMING AND MAINTENANCE

Two further results follow from the baseload character of the plant. Figure 8 evaluates the capacity-factor firming relation of Eq. (15): a baseload fusion source (capacity factor ≈ 0.93) firms a 24/7 compute load at a nameplate multiple of about 1.2 with buffer-only storage, whereas variable-renewable sources (capacity factor 0.24–0.35) require 3.3–$4.8\times$ nameplate and roughly 100–215 hours of storage to reach the same firmness (MV-CX-23). Combined-heat-and-power delivery to the campus lifts annual energy utilisation from $\sim 42\%$ (electricity only) to $\sim 64\%$ (MV-CX-15), the balance being the exported heat of §5.4.

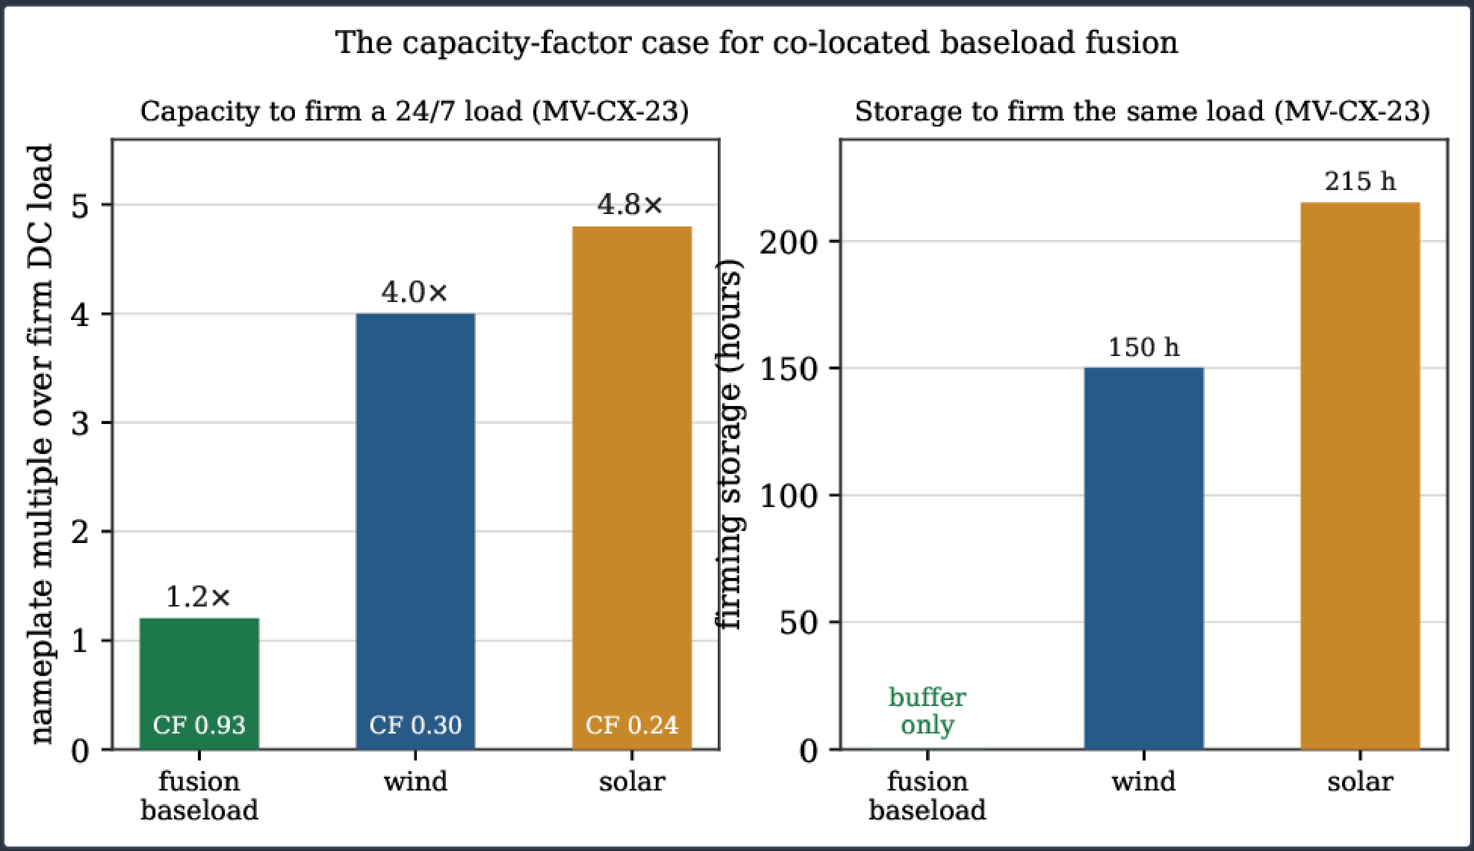

Capacity-factor firming (Eq. 15, MV-CX-23). Left: nameplate multiple over the firm compute load; a baseload source needs $\approx 1.2\times$, variable renewables 3.3–$4.8\times$. Right: firming storage; baseload needs a short buffer, variable renewables ~ 100–215 hours.

Finally, redundancy must survive not only random trips but planned maintenance. Figure 9 contrasts the two redundancy orders under staggered servicing (MV-CX-24): an $N+1$ fleet keeps the IT load up through all planned maintenance but consumes its only spare while a module is out, leaving $\sim 10{,}190$ minutes per year of single-fault exposure; an $N+2$ fleet retains a spare during service and cuts that exposure to ~ 856 minutes per year—a $12\times$ improvement, and the reason Tier-IV service prefers $N+2$. This is the operational complement of the static availability result: the same redundancy order that clears the Monte-Carlo target also preserves fault tolerance during maintenance.

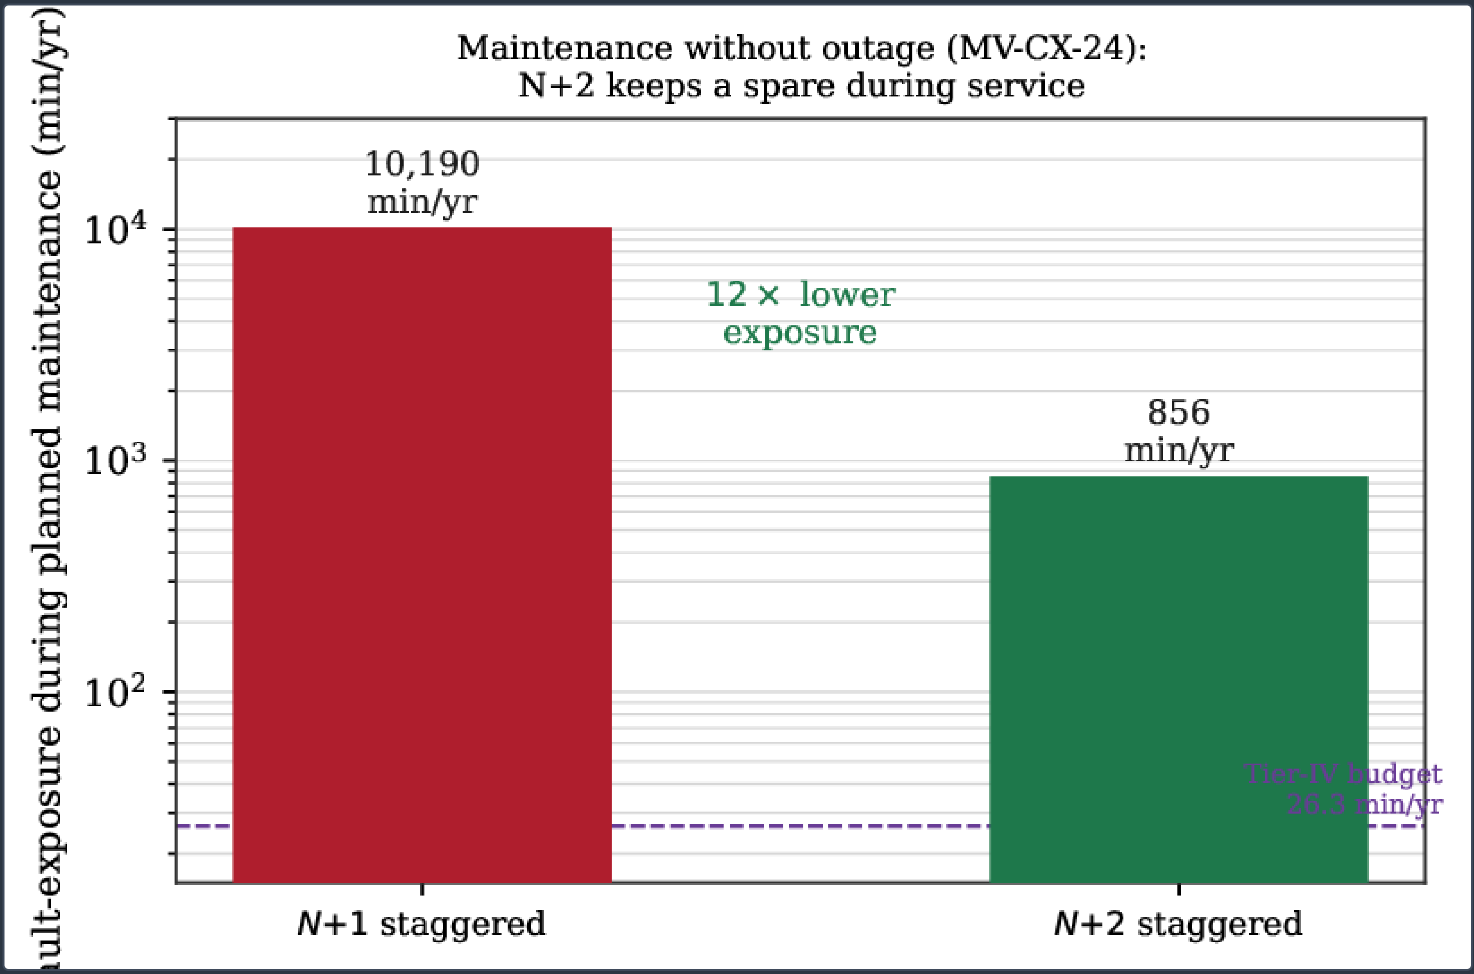

Maintenance without outage (MV-CX-24): annual single-fault exposure during staggered planned maintenance, $N{+}1$ versus $N{+}2$. The $N{+}2$ fleet retains a spare during service, a $12\times$ lower exposure, well inside the Tier-IV budget.

§ 6

Discussion

Reliability in context. The Tier-IV target of 99.995 is the most demanding tier in wide commercial use [4], and the result here is that a fusion fleet reaches it by the same means as any critical facility—redundancy and component reliability—rather than by any exotic assumption. The closed-form scaling $1 - A_{\mathrm{sys}} \sim q^{r+1}$ of Eq. (4) is standard k-out-of-n reliability [14,15]; the contribution is to quantify, by Monte Carlo, that a 2000 hour-MTBF module needs either an eight-fold longer MTBF or a second spare to move a **13** chance of clearing Tier-IV to **77** –**85** . That places a concrete, falsifiable reliability requirement on the power module—MTBF and repair time—rather than leaving availability as an aspiration.

Efficiency against published norms. A facility PUE of **1.17** is competitive with the best operating hyperscale data centers and well below the industry-average PUE, which has sat near **1.5–1.6** for a decade [2,1]. The pathway does not rely on climate (as free-air cooling does) but on DC-native conversion and liquid cooling, both of which the compact fusion plant enables naturally: it produces HVDC directly [28], so the rectification overhead a conventional AC-fed facility carries is simply absent [8,7]. The energy-reuse effectiveness $\mathbf{ERE} \approx \mathbf{0.92}$ is stronger still: an $\mathbf{ERE} < \mathbf{1}$ means the facility exports more energy (as useful heat) than its non-IT overhead consumes, a regime that conventional grid-fed data centers reach only where a large, co-located heat off-taker exists. Here the heat off-take is the plant′s own cooling load plus a district-grade export loop, so the reuse is structural.

Waste-heat thermodynamics against benchmarks. A double-effect LiBr absorption cooling coefficient of **1.35** is squarely within the measured range for double-effect machines [11,12], and a low-lift heat-pump second-law efficiency of $\eta_{II} \approx \mathbf{0.5}$—the value that makes the capture ladder of Figure 7 self-consistent—is typical of high-temperature industrial heat pumps [13]. The result that reuse is grade-limited rather than quantity-limited is a direct consequence of energy conservation (IT electricity becomes heat) combined with the Carnot dependence of Eq. (13); it is the quantitative reason the industry has moved toward warm-water and immersion cooling [9,10], and the compact plant′s own reject heat extends the same logic to the cooling side.

Placement. The behind-the-meter advantage is a systems result, not a physics one, but it is large: a **6.3**-point gain in delivered availability and a halving of time-to-power. Interconnection queues of several years are now the dominant schedule risk for new large loads and generators alike [18]; a co-located plant that serves its load behind the meter simply does not enter that queue for its primary delivery path, retaining a grid tie only for export and ride-through. Read together, the reliability, efficiency, and heat results describe a plant that behaves as a facility—and it is the compactness and neutron-leanness of the D^{3}He tandem-mirror burner [27,31], not any facility trick, that makes demand-adjacent siting admissible in the first place.

§ 7

Limitations

One gate is stated plainly. Fully islanded operation—the plant serving its load with no grid connection at all— reaches an availability ceiling near 99.9 , not four nines, because a self-contained fleet has no external source to draw on during a coincident-outage tail. True four-nines availability uses a grid tie for ride-through, which is exactly the hybrid behind-the-meter-plus-grid topology that the dispatch analysis of §5.2 independently recommends. The constraint and its resolution are therefore the same design choice: the grid tie that carries export and ancillary value is also the ride-through path that lifts the islanded ceiling to Tier-IV. Where a genuinely islanded deployment is required (a forward base, say), the honest figure is 99.9 with $N+2$ redundancy and battery storage, and that is the number this paper stands behind for that case [30].

§ 8

Conclusion

MetroVolt integrates the Kronos D^3He tandem-mirror burner as a facility that meets its service standards. Tier-IV availability is reached through module reliability ($\mathrm{MTBF} \approx 8000\,hour$) or $N{+}2$ redundancy, grid-tied, with a ~ 99.9 islanded ceiling; the fleet model $A_{\mathrm{sys}} = \sum_{k \geq N} \binom{M}{k} a^k (1-a)^{M-k}$ and its q^{r+1} scaling make those levers explicit, and a 2500-trial Monte Carlo turns them into probabilities ($0.13 \to 0.77$–0.85). Behind-the-meter delivery adds 6.3 points of delivered availability and cuts time-to-power from 78 to 30 months. A composed lever stack takes PUE from 1.35 to 1.17 at an energy-reuse effectiveness $\mathrm{ERE} \approx 0.92$, and waste-heat reuse cuts electric cooling by 91 while exporting useful heat at a reuse factor near unity, with the capture temperature governing reuse through $\mathrm{COP}_{\mathrm{HP}} = \eta_{II} T_h / (T_h - T_c)$. The burner is deployable not merely as a reactor but as an integrated, demand-adjacent energy facility.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos de-risking campaign for the frozen facility-integration anchors and the GPU HPC campaign that produced them.

Reproducibility. The availability Monte Carlo (MV-CX-22), the behind-the-meter comparison (MV-CX-18), the PUE lever stack (MV-CX-14), the absorption-cooling and heat-capture balances (MV-CX-12/13), the capacity-factor firming model (MV-CX-23), and the maintenance-exposure model (MV-CX-24) are frozen anchors in the Kronos Physics De-Risking Register and are regenerable from the reliability, systems, and thermal scripts archived with this paper (`make_fig_266_datacenter.py`). This paper is archived at its DOI doi:10.5281/zenodo.22645919 and at its official page https://www.kronosfusionenergy.com/publications/data-center-integration-of-a-d-he-tandem-mirror-burner-.html; the de-risking register is archived at doi:10.5281/zenodo.22133057.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645919. The Monte-Carlo availability outputs, the PUE lever table, the thermal-integration calculations, and the figure-generation script are archived with the deposit and reference the Kronos 2026 series. Any economic analysis is reported separately and is not included in the public deposit.

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

tocsectionReferences

APPARATUS

Portfolio, People & Record

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

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

INTELLECTUAL PROPERTY

The patent portfolio

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

GRANTED & PENDING UTILITY

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

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

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

TRADEMARK & ORIGIN

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

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

THE TEAM

Founding partners, board, advisors & auditors

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

FOUNDING PARTNERS — SCIENCE

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

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

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

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

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

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

SCIENTIFIC ADVISORS

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

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

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

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

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

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

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

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

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

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

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

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

SCIENTIFIC AUDITORS

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

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

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

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

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

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

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

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

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

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

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

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

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

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

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

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

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

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

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

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

Andrea Romero
Marketing Lead
2022–Present

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

SOURCES

References

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

2 A Shehabi, S J Smith, D A Sartor *et al*, *United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, LBNL-1005775 (2016).

3 L A Barroso, U H\"olzle and P Ranganathan, *The Datacenter as a Computer: Designing Warehouse-Scale Machines*, 3rd ed. (Morgan & Claypool, 2018).

4 Uptime Institute, *Tier Classification System for Data Center Site Infrastructure* (Uptime Institute, 2018).

5 V Avelar, D Azevedo and A French (eds.), *PUE: A Comprehensive Examination of the Metric*, The Green Grid White Paper #49 (2012).

6 The Green Grid, *ERE: A Metric for Measuring the Benefit of Reuse Energy from a Data Center*, White Paper #29 (2010).

7 ASHRAE Technical Committee 9.9, *Liquid Cooling Guidelines for Datacom Equipment Centers*, 2nd ed. (ASHRAE, 2014).

8 J B Marcinichen, J A Olivier and J R Thome, On-chip two-phase cooling of datacenters: cooling system and energy recovery evaluation, *Appl. Therm. Eng.* **41** 36 (2012).

9 S Zimmermann, I Meijer, M K Tiwari, S Paredes, B Michel and D Poulikakos, Aquasar: a hot water cooled data center with direct energy reuse, *Energy* **43** 237 (2012).

10 K Ebrahimi, G F Jones and A S Fleischer, A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities, *Renew. Sustain. Energy Rev.* **31** 622 (2014).

11 K E Herold, R Radermacher and S A Klein, *Absorption Chillers and Heat Pumps*, 2nd ed. (CRC Press, 2016).

12 P Srikhirin, S Aphornratana and S Chungpaibulpatana, A review of absorption refrigeration technologies, *Renew. Sustain. Energy Rev.* **5** 343 (2001).

13 C Arpagaus, F Bless, M Uhlmann, J Schiffmann and S S Bertsch, High temperature heat pumps: market overview, state of the art, research status, refrigerants, and application potentials, *Energy* **152** 985 (2018).

14 R Billinton and R N Allan, *Reliability Evaluation of Engineering Systems: Concepts and Techniques*, 2nd ed. (Plenum Press, 1992).

15 W Kuo and M J Zuo, *Optimal Reliability Modeling: Principles and Applications* (Wiley, 2003).

16 K S Trivedi, *Probability and Statistics with Reliability, Queuing, and Computer Science Applications*, 2nd ed. (Wiley, 2001).

17 N Metropolis and S Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44** 335 (1949). doi:10.1080/01621459.1949.10483310

18 J Rand, N Manderlink, W Gorman *et al*, *Queued Up: Characteristics of Power Plants Seeking Transmission Interconnection*, Lawrence Berkeley National Laboratory (2023).

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

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

21 A J Creely *et al*, Overview of the SPARC tokamak, *J. Plasma Phys.* **86** 865860502 (2020). doi:10.1017/S0022377820001257

22 R F Post, The magnetic mirror approach to fusion, *Nucl. Fusion* **27** 1579 (1987).

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

24 R W Moir and W L Barr, Venetian-blind direct energy converter for fusion reactors, *Nucl. Fusion* **13** 35 (1973).

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

26 P I Ford, *The Physics De-Risking Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

27 P I Ford, The MetroVolt Data Center Burner: Direct-DC Campus Power, *Q* Closure, and the Plug Requirement in a Low-Neutron D^3

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

Data-Center Integration of a D–^{3}He Tandem-Mirror Burner: Four-Nines Uptime Against Plasma Transients, the PUE-Reduction Pathway, and Waste-Heat Reuse for Behind-the-Meter Dispatch

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

Zenodo · doi:10.5281/zenodo.22645919

Read on kronosfusionenergy.com

The Physics De-Risking Register