Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Design and Development of a Next-Generation Medium-Voltage Architecture for AI Data Centre Infrastructure

Mohammad Ali

malixbd@ieee.org

Atlantic International University https://orcid.org/0009-0001-2184-7978

Research Article

Keywords: AI data centre, arc-flash hazard, DC distribution, energy storage, IEC 60909, IEEE 519, IEEE 1584, Markov availability, medium-voltage distribution, power quality, power usage effectiveness, solid-state transformer

Posted Date: September 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11196206/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Design and Development of a Next-Generation Medium-Voltage Architecture for AI Data Centre Infrastructure

Mohammad Ali

School of Science & Engineering, Atlantic International University, Honolulu, HI 96813, USA

malixbd@ieee.org

***Abstract*—The industrialisation of artificial-intelligence (AI) computation has produced an electrical load class that the incumbent low-voltage (LV) data-centre reticulation architecture was not designed to serve. Accelerated-compute racks now draw 120–140 kW and are credibly forecast to reach 500–1000 kW within the present design horizon, while synchronised gradient exchange across tens of thousands of accelerators imposes deterministic power oscillation in the 0.1–10 Hz band and campus ramp rates of tens of megawatts per second. This paper develops, dimensions and evaluates a medium-voltage (MV) architecture that carries MV to the row and converts once, at the row, to a regulated bipolar ±800 V direct-current (DC) bus. The architecture comprises dual 132/33 kV intakes, an SF_6-free 33 kV ring main, distributed 20 MVA solid-state transformer blocks, a 13.8 kV cast-resin busway to the row, and 1.5 MW row power-conversion units with hybrid battery and super-capacitor buffering. A complete analytical framework is presented, covering a parameterised stochastic load model, Kelvin's-law conductor economics extended with fault-level, arc-flash and converter-cost penalties, the cascaded efficiency chain, IEC 60909 short-circuit computation, tri-duty DC-bus storage sizing, IEC 60287 ampacity derating for a 50 °C ambient, Markov availability and IEEE 1584 incident-energy evaluation. The framework is applied to a fully specified 100 MW campus in Abu Dhabi and benchmarked against a legacy 400 V baseline and an improved 415 V variant. End-to-end efficiency rises from 80.28 % to 90.26 %; conductor mass for an equivalent 2.5 MVA, 200 m feeder falls by 97.1 %; expected annual downtime falls from 42.6 min to 3.2 min at a computed availability of 0.999994; total demand distortion at the point of common coupling falls from 11.71 % to 1.63 %; worst-case incident energy falls from 48.6 to 6.4 cal/cm², moving the maintenance boundary from personal-protective-equipment category 4 to category 2; and the imposed grid ramp rate falls from 96 to 11.8 MW/s. Capital cost falls by 17.3 % and fifteen-year total cost of ownership by 32.4 %, because savings in LV distribution, cabling, containment and civil works exceed the premium on solid-state conversion equipment.**

***Index Terms*—AI data centre, arc-flash hazard, DC distribution, energy storage, IEC 60909, IEEE 519, IEEE 1584, Markov availability, medium-voltage distribution, power quality, power usage effectiveness, solid-state transformer.**

I. Introduction

The electrical design of data centres has, for two decades, been an exercise in accommodating a load whose density grew steadily but whose character did not change. A rack drew between five and fifteen kilowatts, the demand was quasi-static, and diversity between racks was high. The architecture that served that load—an 11 kV package substation feeding a centralised double-conversion uninterruptible power supply (UPS), a low-voltage main switchboard, power distribution units (PDUs) and a 400 V busway—was adequate, and its refinement occupied the design literature.

The training of large neural models has ended that stability. A contemporary accelerated-compute rack draws 120–140 kW; platforms already announced approach 250 kW, and the credible design horizon extends to 500–1000 kW per rack [40], [41]. Fig. 1 places this trajectory against the practical ceiling of low-voltage reticulation. More consequential than the magnitude is the dynamic character. Synchronised gradient exchange across a training cluster produces a load that is neither random nor quasi-static but deterministically oscillatory, with components in the 0.1–10 Hz band and a peak-to-mean ratio near 1.40 [42], [43]. Job launch and termination impose step changes that, at campus scale, correspond to ramp rates of tens of megawatts per second—a rate that transmission system operators do not ordinarily accept from a single connection.

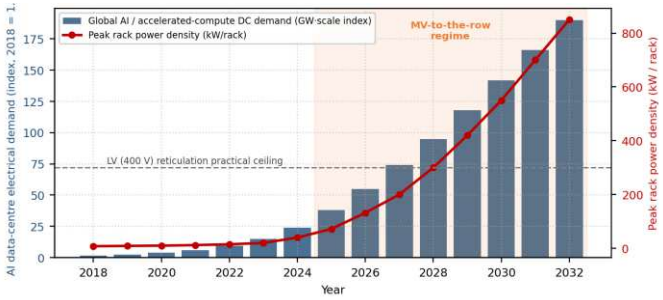


Fig. 1. Growth of AI data-centre electrical demand and peak rack power density against the practical ceiling of low-voltage reticulation.

The argument against low voltage is elementary and is rarely disputed; what has been missing is a complete, dimensioned and standards-referenced design that acts on it. Transporting a given block of power at 400 V requires 34.5 times the current needed at 13.8 kV. Conductor mass, electrical-room volume, prospective fault level and arc-flash incident energy all follow that current. At the densities now being built, low-voltage reticulation is no longer a distribution system with an efficiency penalty; it is a physical obstruction.

This paper therefore asks where in the chain medium voltage should be converted, and answers that it should be converted once, at the row, directly to a regulated bipolar ±800 V DC bus. The contributions are: (i) a parameterised stochastic load model for large-model training, with a coincidence formulation for aggregation from rack to campus; (ii) a hierarchical MV-to-the-row architecture defined to the level of single-line diagram, apparatus schedule, earthing philosophy and protection scheme; (iii) a treatment of the optimal intermediate distribution voltage that combines conductor economics with the fault-level, arc-flash and converter-cost penalties of raising voltage; (iv) an energy-balance formulation that separates the ride-through, ramp-limiting and oscillation-damping duties of DC-bus storage and dimensions battery and super-capacitor elements against each separately; (v) a Markov availability model that captures the degraded-but-operational states created by row-level converter redundancy; and (vi) a quantified evaluation of three architectures against nine performance indices, with a sensitivity analysis identifying the parameters to which the conclusions are exposed.

II. Related Work and Research Gap

The published record divides into four eras. To approximately 2005, computer-room practice was transposed from commercial buildings, and the literature is concerned with redundancy topology and its classification. From 2005 to 2015 the concern is efficiency: power usage effectiveness (PUE) is introduced, distributed and modular UPS topologies displace the centralised machine [25], [37], and 415/240 V distribution is adopted in place of 400/230 V to recover a conversion stage [38], [39]. From 2015, hyperscale operators move to rack-level 48 V DC architectures with lithium-ion buffering, and the first serious proposals for direct-current and medium-voltage distribution within the facility appear [33]–[36]. The fourth era, beginning around 2023, is defined by accelerated compute and is not yet consolidated in the standards.

A. MV distribution and solid-state transformation

Solid-state transformers have a mature literature in traction and distribution applications [29]–[32], [48], [54]. The cascaded H-bridge (CHB) front end with a dual-active-bridge (DAB) isolation stage [32], [47], [53] is the established topology for medium-voltage input, and silicon-carbide devices at 1.7 kV and 3.3 kV blocking have brought converter efficiencies above 98 % into the commercially credible range. What the data-centre literature has not done is carry that device-level result through to a facility-level architecture with a defensible protection scheme, a quantified availability model and a cost case.

B. Direct-current distribution in facilities

The 380 V DC facility architecture is documented and partly standardised [19], and the 48 V rack bus is universal in hyperscale practice [34]–[36]. The intermediate case—a bipolar ±800 V bus fed from medium voltage at the row—has been proposed but not dimensioned against a complete set of protection, earthing, storage and safety constraints. Protection of such a bus is the principal open question, because the availability and rating of medium-current DC circuit-breakers remain limited and the applicable standards are still developing.

C. The gap

Three deficiencies recur. First, load models used for distribution design treat the AI training load as a scaled conventional load, and so under-represent both its oscillatory content and its ramp rate. Second, architectural proposals are evaluated on efficiency alone, without a simultaneous treatment of fault level, arc-flash exposure, harmonic compliance, availability and cost—quantities that trade against one another and cannot be optimised separately. Third, almost the entire published record assumes a temperate design ambient [22], [23]; the derating that a 50 °C ambient imposes on every current-carrying element, every semiconductor and every battery is not treated as a first-order design input. This paper addresses all three.

III. Electrical Characterisation of the AI Compute Load

A. Stochastic load model

Rack-level instantaneous power under a training workload is represented as the sum of a base term, a workload-dependent envelope modulated by a multi-tone oscillatory term, and a stochastic residual:

$$p(t) = P_{base} + P_{dyn} \cdot g(t) \cdot [\, 1 + \Sigma_{i=1..m}\, a_i \sin(\, 2\pi f_i\, t + \varphi_i\,)\,] + \varepsilon(t) \quad (1)$$

where P_{base} is the idle and ancillary draw, P_{dyn} the dynamic headroom to the platform thermal design power, $g(t) \in [0,1]$ the scheduler occupancy envelope, and $\varepsilon(t)$ a zero-mean Gaussian residual. Three tones reproduce the reported signature: 0.35 Hz for the checkpoint-write cadence, 1.70 Hz for all-reduce gradient exchange and 6.20 Hz for layer-wise synchronisation. Parameters are given in Table I. Fig. 2 shows the resulting rack-level signature and Fig. 3 its load-duration curve and power spectral density.

TABLE I
Parameters of the stochastic AI training load model of (1)

Parameter	Symbol	Value	Unit
Rated rack power	P_{rate}^{d}	260	kW
Base / idle power	P_{base}	0.35 P_{rate}^{d} = 91	kW
Dynamic range	P_{dyn}	169	kW
Envelope rise / fall time	t_r / t_f	0.35 / 0.55	s
Tone 1 — checkpoint write	a_1 / f_1	0.07 / 0.35	– / Hz
Tone 2 — all-reduce	a_2 / f_2	0.16 / 1.70	– / Hz
Tone 3 — layer sync.	a_3 / f_3	0.04 / 6.20	– / Hz
Residual std. deviation	σ	0.020 P_{rate}^{d}	kW
Peak-to-mean ratio	P_{pk} / $\bar{P}$	1.40	–
Max. rack ramp rate	dP/dt	≈ 480	kW/s

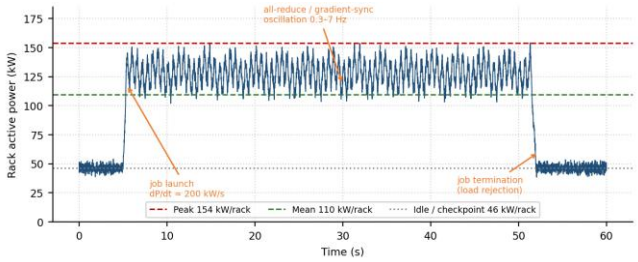


Fig. 2. Representative power signature of a large-model training workload at rack level, showing job launch, the synchronised gradient-exchange oscillation and job termination.

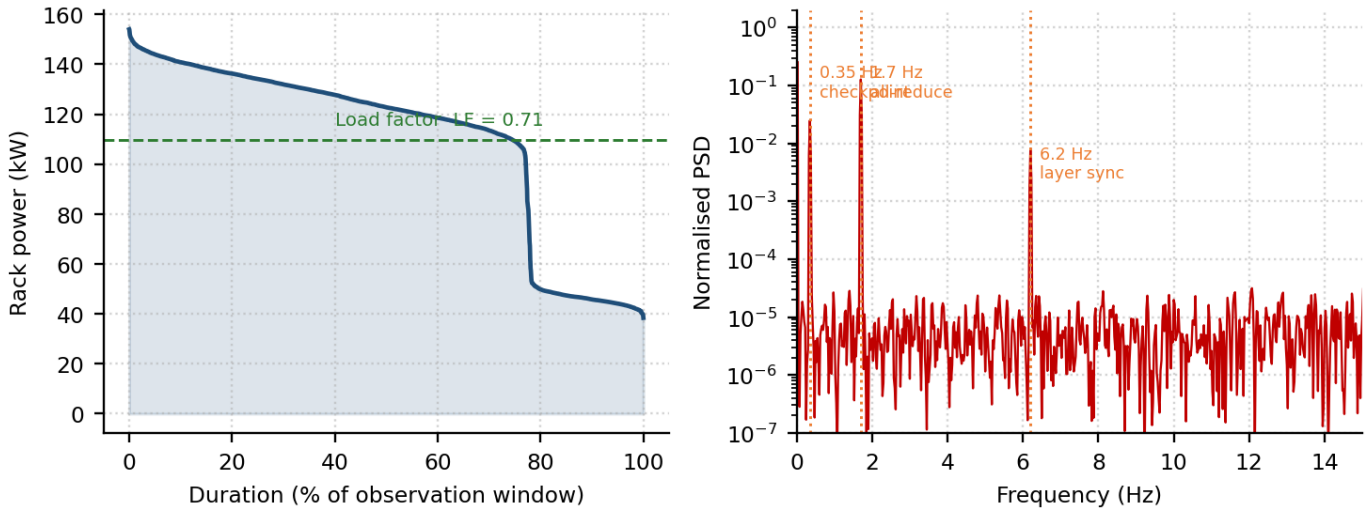


Fig. 3. (a) Load-duration curve of the modelled training workload; (b) power spectral density of the load oscillation, showing the deterministic collective-communication tones above the stochastic floor.

B. Aggregation and coincidence

Aggregation from rack to campus is not the arithmetic sum, because racks within a training domain are synchronised and racks in different domains are not. The coincidence factor is represented as an asymptotic decay in the number of aggregated units:

$$P_{agg}(t) = \Sigma_{d=1..N_d}\, k_c(N_r) \cdot N_r \cdot p_d(t)$$
$$k_c(N) = k_\infty + (1 - k_\infty) \cdot N^{-\alpha} \qquad (2)$$

with $k_\infty = 0.86$ and $\alpha = 0.42$ fitted to reported hall-level traces. Within a single training domain k_c approaches unity, which is the design-governing case and the value adopted throughout. Two consequences follow. The load factor is high—approximately 0.71—so the infrastructure spends most of its life near its design point, which favours low no-load loss but forbids reliance on part-load diversity. And the spectral energy lies below 10 Hz, precisely the band in which tap changers and generator governors are ineffective; compensation must come from power-electronic and storage elements responding in milliseconds. Unmitigated, a synchronised launch of the full 100 MW cluster produces a demand ramp near 96 MW/s. Table II gives the resulting design load at each level of the hierarchy.

TABLE II
DERIVED DESIGN LOAD AT EACH LEVEL OF THE DISTRIBUTION HIERARCHY (100 MW CAMPUS)

Hierarchy level	Qty	Mean load	Design (peak) load	Design current	Governing consideration
Rack	384	186 kW	260 kW	163 A at ±800 V DC	Platform TDP × 1.40 peak/mean
Row (8 racks)	48	1.49 MW	2.08 MW	1300 A at ±800 V DC	$k_c = 0.98$ within domain
Row PCU (2 × 50 %)	96	0.75 MW	1.50 MW	63 A at 13.8 kV	N+1 at row; each unit 100 % of row
MV block (8 rows)	6	11.9 MW	16.6 MW	695 A at 13.8 kV	Busway thermal limit + derating
Block transformer / SST	12	—	20 MVA	350 A at 33 kV	2 × 100 % per block, 2N
33 kV ring main	1	71.5 MW	100 MW IT / 118 MVA	2065 A at 33 kV	Includes cooling and losses, cos φ 0.99
132/33 kV intake	2	—	2 × 60 MVA	525 A at 132 kV	Firm capacity on loss of one intake

IV. PROPOSED MEDIUM-VOLTAGE ARCHITECTURE

A. Design principles

Six principles govern the synthesis. Carry medium voltage as close to the load as the safety case permits, because every metre of low-voltage high-current conductor eliminated removes mass, loss, fault energy and containment volume. Convert once rather than five times: the legacy chain performs five conversions between utility and accelerator core, the proposed chain performs three. Make the DC bus the point of aggregation, so that storage, on-site generation, row redundancy and rack supply all interface at a single bus rather than through separately synchronised AC interfaces. Decouple compute dynamics from the utility by design rather than by contract, so that grid-code compliance does not depend on the behaviour of the workload scheduler. Limit fault energy at source, exploiting the bounded fault contribution of converter-interfaced sources deliberately rather than treating it as a protection nuisance. And design for the 50 °C ambient from the outset, applying derating to the base rating of every current-carrying element rather than as a retrospective correction.

B. Topology

The proposed architecture is shown in Fig. 5 against the legacy baseline of Fig. 4. Two independent 132 kV circuits feed four 60 MVA 132/33 kV transformers in a 2×2 arrangement providing firm capacity on loss of either intake. A 33 kV SF_6-free ring main with vacuum interruption and clean-air insulation distributes to six MV blocks, each comprising two 20 MVA 33/13.8 kV solid-state or hybrid transformers in a 2N arrangement. From each block a 13.8 kV cast-resin busway rated 1250 A runs to the row. Ninety-six row power-conversion units (PCUs), two per row at 1.5 MW each, convert 13.8 kV directly to a bipolar ±800 V DC bus with a high-resistance-earthed mid-point. A ±800 V busway rated 1000 A per pole supplies 384 racks at 260 kW design load, with low-voltage final circuits selected to IEC 60364-5-52 [3].

Eliminating the low-voltage main board, the PDUs and the double-conversion UPS removes the three elements that

dominate the loss, the footprint and the arc-flash exposure of the baseline, and removes the single largest contributor to its unavailability.

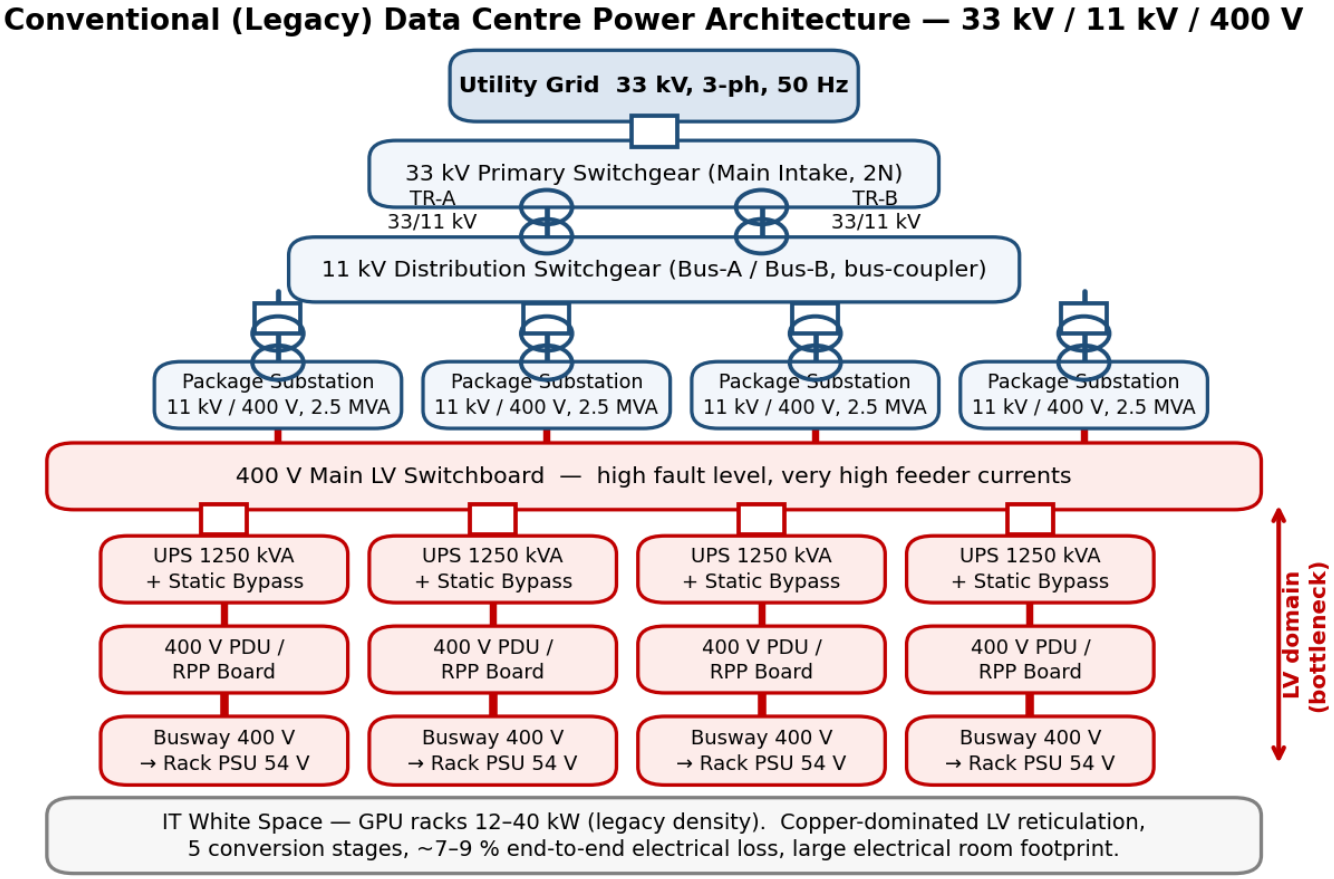


Fig. 4. Legacy low-voltage data-centre power architecture used as the comparison baseline, showing the five-stage conversion chain from the utility intake to the accelerator core.

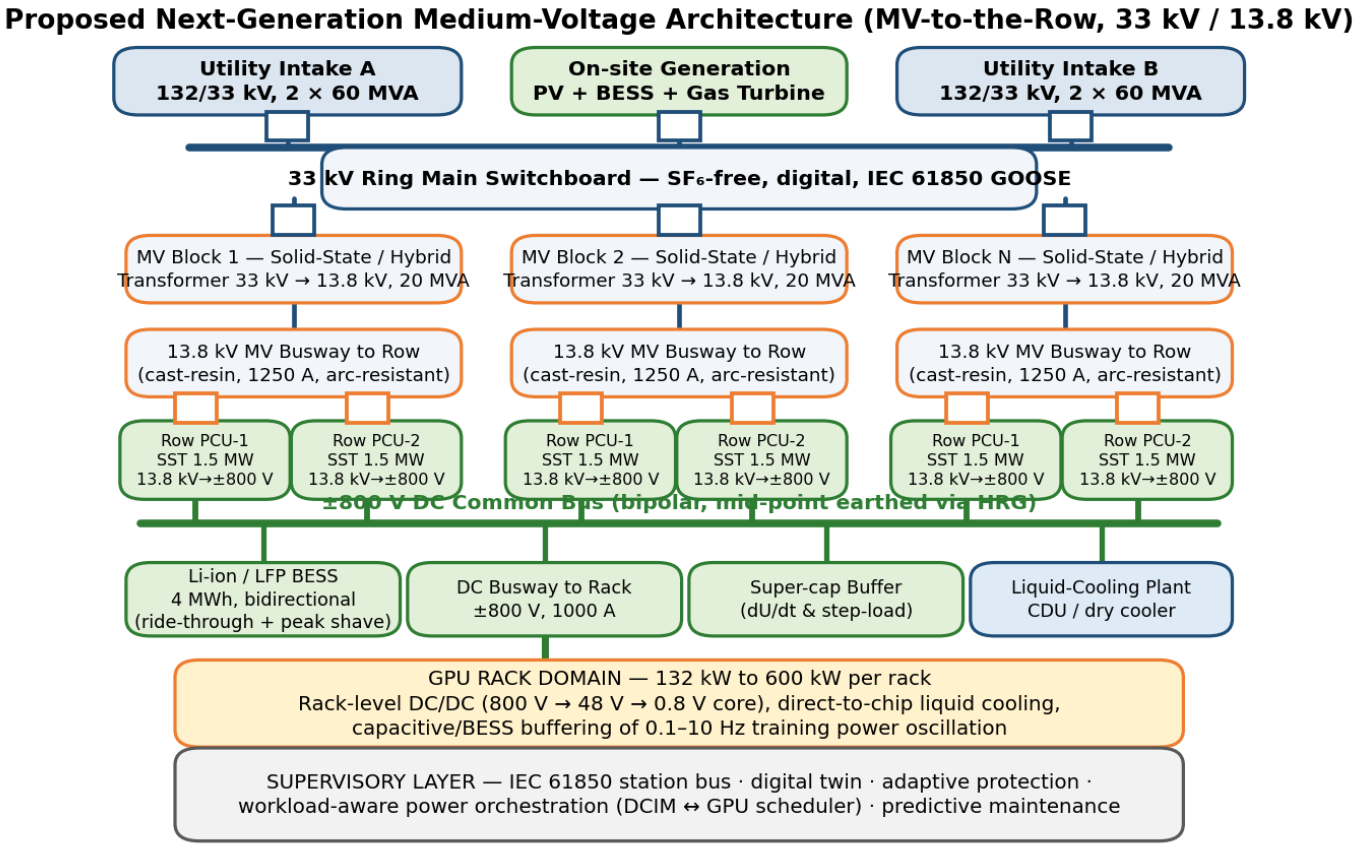


Fig. 5. Proposed next-generation medium-voltage architecture: 33 kV ring main, distributed medium-voltage blocks, 13.8 kV busway to the row, and row-level conversion to a bipolar ±800 V direct-current bus.

C. Row power-conversion unit

The row PCU is the technical heart of the architecture; its internal topology is shown in Fig. 6 and its specification in Table III. A CHB front end uses six cells per phase at 1.7 kV silicon-carbide blocking voltage, giving a thirteen-level line-to-line waveform and a per-cell DC link near 1.2 kV. Each cell feeds a DAB stage switching at 20 kHz into a nanocrystalline-core medium-frequency transformer, and the outputs are combined onto the ±800 V bus. A seventh cell per phase is fitted against six required, so that a single cell failure degrades rather than fails the unit. The control loop enforces a 2.0 p.u. fault-current limit within 100 µs, which bounds the DC fault energy and is the property on which the arc-flash result of Section VII-E rests.

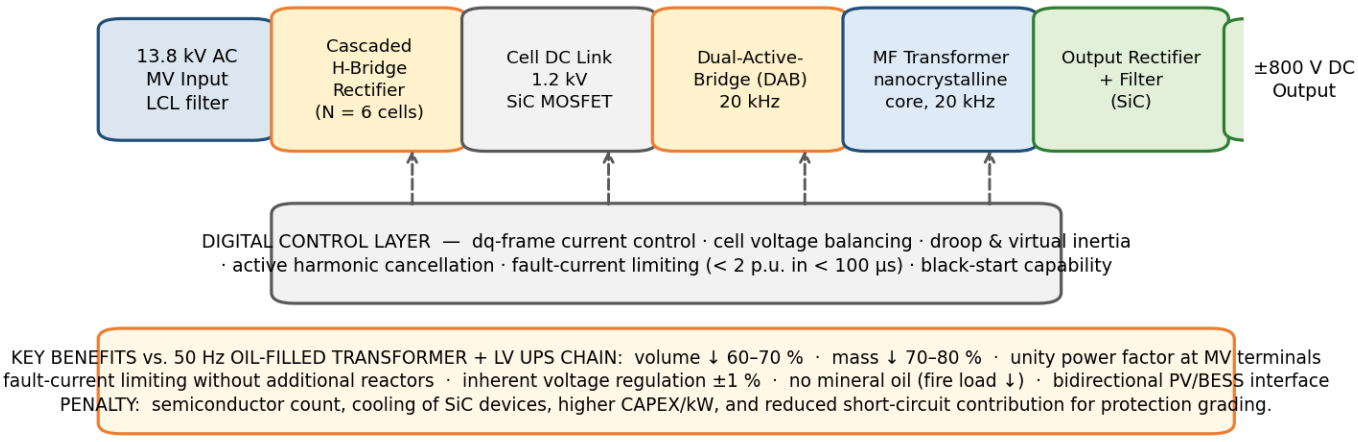


Fig. 6. Internal topology of the solid-state transformer power-conversion cell adopted for the row-level power-conversion unit: cascaded H-bridge front end, medium-frequency isolation and bipolar direct-current output stage.

TABLE III
SPECIFICATION OF THE ROW POWER-CONVERSION UNIT

Parameter	Value	Basis	Consequence for the design
Rated output power	1500 kW continuous at 50 °C	Row design load 2.08 MW, N+1 sizing	Two units carry the row; either alone carries 72 % of peak
Input voltage	13.8 kV ±10 %, 3-phase, 50 Hz	MV busway nominal	No intermediate transformation
Output voltage	±800 V DC bipolar, ±0.5 % regulation	Rack input class	Direct rack connection, no PDU stage
Topology	CHB + DAB, 6 cells per phase	State of the art (Section II-A)	Modular, cell-redundant, hot-swappable
Semiconductor	1.7 kV SiC MOSFET	Blocking margin on 1.2 kV cell link	Low conduction loss, high switching frequency
Switching frequency	20 kHz (DAB); 1.25 kHz (CHB effective)	Loss / size optimum	MFT volume ≈ 1/400 of a 50 Hz core
Rated efficiency	98.2 % at 100 % load; 97.9 % at 50 %	Measured class data for SiC MV converters	Flat efficiency across the operating band
Displacement power factor	0.99 lead/lag, controllable	Active front-end control	Reactive support to the MV system
Input current THD	< 2.0 % at ≥ 25 % load	13-level modulation + active filtering	IEEE 519 compliance without bulk filters
Fault-current limit	2.0 p.u., enforced < 100 µs	Control-loop limit	Bounds DC fault energy; enables PPE category 2
Cell redundancy	N+1 per phase (7 fitted, 6 required)	Availability target	A single cell failure is not a unit failure
Cooling	Liquid cold plate, 45 °C supply	Shared with IT coolant loop	Heat rejected to dry coolers, not chillers
Enclosure	IP54, arc-resistant, front access only	Occupied white-space installation	Maintenance without rear clearance
Communications	IEC 61850-9-2 / GOOSE + Modbus TCP	Station-bus integration	Adaptive protection, digital-twin telemetry

D. Earthing and hybrid storage

Earthing is layered. The 132/33 kV system is solidly earthed to utility practice; each 33 kV transformer neutral is resistance-earthed to limit earth-fault current to 400 A, and each 13.8 kV block transformer neutral to 200 A, with restricted-earth-fault protection on each transformer, in accordance with IEC 60364-5-54 [4], EN 50522 [6] and IEEE 80 [18]. The DC bus mid-point is earthed through a high resistance limiting first-fault current below 10 A, with continuous insulation monitoring on both poles to IEC 61557-8 [12]. This is chosen deliberately over solid mid-point earthing: a first earth fault on a pole raises an alarm rather than tripping the bus, permitting load transfer and fault location during a planned intervention rather than forcing an unplanned interruption of a training run that may represent weeks of accumulated computation.

Storage is provided in two forms with distinct duties: 4 MWh at 8 MW of lithium iron phosphate per block for ride-through and ramp limiting, and 12 MJ of super-capacitance per row PCU for oscillation damping. Both connect to the DC bus through bidirectional converters, and both are installed to IEC 62619 [13] and NFPA 855 [24]. The tri-duty separation extends the pulsed-load sizing approach of [55]. The separation of duties is what permits each element to be dimensioned economically, as Section V-D shows.

V. ANALYTICAL FRAMEWORK

A. Voltage regulation and conductor economics

For a three-phase feeder of route length ℓ carrying current I at power factor cos φ, the regulation is [50], [52]

$$\Delta U = \sqrt{3} \cdot I \cdot \ell \cdot (R' \cos \varphi + X' \sin \varphi) \tag{3}$$

and for a fixed transported power S the inverse-square dependence on nominal voltage follows directly:

$$\Delta U\% = 100 \cdot S \cdot \ell \cdot (R' \cos \varphi + X' \sin \varphi) / U_n^2 \tag{4}$$

Fig. 7 evaluates (4) for four candidate voltages transporting 2.5 MVA. Conductor cross-section at an economic current density J, and the corresponding copper mass, are

$$A = I / (n \cdot J) = S / (\sqrt{3} \cdot n \cdot J \cdot U_n)$$
$$m_{Cu} = 3 \cdot A \cdot \ell \cdot \rho_m = 3 \cdot \ell \cdot \rho_m \cdot S / (\sqrt{3} \cdot U_n \cdot J) \tag{5}$$

so mass falls as U_n^{-1} while loss falls as U_n^{-2}. Raising the voltage indefinitely is nevertheless not optimal, because converter cost scales approximately as $U^{0.6}$, switchgear cost as $U^{0.45}$ and safety and clearance provisions as $U^{0.3}$. The composite objective

$$C_{total}(U) = w_1 \cdot C_{cond}(U) + w_2 \cdot C_{loss}(U) + w_3 \cdot C_{conv}(U) + w_4 \cdot C_{swgr}(U) + w_5 \cdot C_{safety}(U) \tag{6}$$

is minimised at 13.8 kV for the row-level bus under the case-study cost set, and is flat within 2 % between 11 kV and 15 kV. That flatness is the practically useful result: the choice may be made on equipment availability without economic penalty.

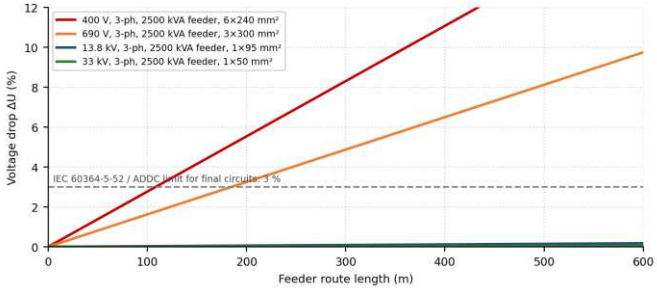


Fig. 7. Voltage regulation as a function of feeder route length for four distribution voltages transporting an identical 2.5 MVA at cos φ = 0.99.

B. Cascaded efficiency

End-to-end efficiency from the point of common coupling (PCC) to the accelerator core is the product of the stage efficiencies, and the infrastructure loss follows:

$$\eta_{total} = \Pi_{j=1..n} \eta_j , \quad P_{loss,total} = P_{IT} \cdot (1 / \eta_{total} - 1) \tag{7}$$

Each stage carries a no-load and a load-dependent term, so that part-load behaviour is captured:

$$\eta_j(x) = 1 - (P_{0,j} + P_{k,j} \cdot x^2) / (P_{rated,j} \cdot x) \tag{8}$$

where x is the per-unit loading of stage j. Fig. 8 evaluates (7) at full load for the three architectures.

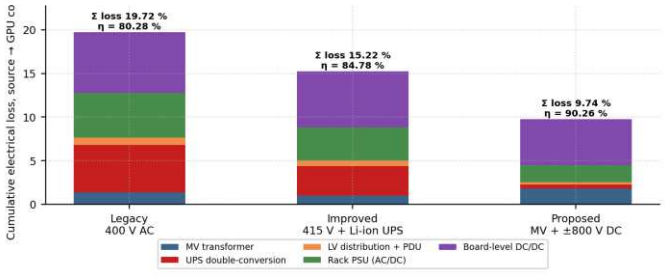


Fig. 8. Cascaded conversion-efficiency comparison of the three architectures, decomposed by conversion stage from the point of common coupling to the accelerator core.

C. Short-circuit computation

Fault currents are computed per IEC 60909-0 [1] with the voltage factor c = 1.10:

$$I''k = c \cdot U_n / (\sqrt{3} \cdot Z_k)$$
$$\kappa = 1.02 + 0.98 \cdot e^{-3\,R/X} , \quad i_p = \kappa \cdot \sqrt{2} \cdot I''k \tag{9}$$

Converter-interfaced sources contribute a bounded current set by the control-loop limit rather than by source impedance, and this is the dominant reason the proposed 13.8 kV bus fault level is 18.3 kA against 24.3 kA for a conventional block of the same rating. The minimum thermally admissible cross-section follows from the adiabatic relation

$$A_{min} = I_{th} \cdot \sqrt{t_k} / k , \quad k = 143 \text{ for Cu / XLPE} \tag{10}$$

D. Tri-duty storage sizing

The three duties are dimensioned separately. Ride-through requires

$$W_{rt} = P_{IT} \cdot t_{rt} / (\eta_{dis} \cdot DoD_{max}) \tag{11}$$

Ramp limiting requires the energy delivered while grid import follows a first-order response of time constant τ to a step from P_1 to P_2:

$$P_{grid}(t) = P_1 + (P_2 - P_1) \cdot (1 - e^{-t/\tau})$$
$$W_{ramp} = \int_0^{\infty} [P_{load}(t) - P_{grid}(t)] \, dt = (P_2 - P_1) \, \tau \tag{12}$$

Oscillation damping at the dominant tone requires only

$$W_{osc} = a \cdot P_{dyn} / (2\pi f) ,$$
$$C_{dc,min} = 2 \, W_{osc} / (V_{dc,max}^2 - V_{dc,min}^2) \tag{13}$$

Evaluated for the case study, (11) gives 4 MWh at block level, (12) gives 15.0 kWh and (13) gives 12 MJ per row PCU. The separation matters: the oscillation duty is a high-cycle, low-energy service that would consume the cyclic life of a lithium cell and is therefore assigned to super-capacitance, while ride-through is a low-cycle, high-energy service for which lithium iron phosphate is appropriate. Sizing a single element against both duties, as is common practice, over-dimensions the storage by roughly a factor of three and under-delivers on damping.

E. Derating for a 50 °C ambient

Transformer and converter ratings are corrected to the site condition per IEC 60076-7 [10] before any other calculation is performed:

$$S_{perm} = S_{rated} \cdot [1 - k_\theta \cdot (\theta_{amb} - \theta_{ref})]$$
$$P_{perm} = (T_{j,max} - T_{amb}) / (R_{th,jc} + R_{th,ch} + R_{th,ha}) \tag{14}$$

with $k_\theta \approx 0.010$ /K for natural-air and 0.008 /K for forced-air cooling. Cable ampacity follows IEC 60287 [2] with the multiplicative corrections for ambient, ground temperature, soil thermal resistivity, grouping and depth:

$$I_z = I_{ref} \cdot k_\theta \cdot k_g \cdot k_s \cdot k_d \cdot k_h \tag{15}$$

F. Availability and incident energy

Availability is computed from a continuous-time Markov model [51], with component data from IEEE 493 [17] and IEEE 3006.7 [21]. For a two-train redundant block the state space is that of Fig. 9, with healthy, degraded and failed states, and with common-cause failure represented by the β-factor partition $\lambda_{cc} = \beta\lambda$, $\lambda_{ind} = (1 - \beta)\lambda$. The steady-state probabilities follow from

$$dP(t)/dt = P(t) \cdot Q = 0 , \quad \Sigma_i P_i = 1 , \quad A = 1 - P_{failed} \tag{16}$$

Incident energy is evaluated per IEEE 1584-2018 [16], the arcing current from the empirical logarithmic model and the energy from

$$E = 4.184 \cdot C_f \cdot E_n \cdot (t_a / 0.2) \cdot (610^x / D^x) \tag{17}$$

The arcing duration t_a is the quantity over which the design has greatest leverage, and it is the target of the optical arc-flash detection and GOOSE bus-blocking scheme adopted in Section VI.

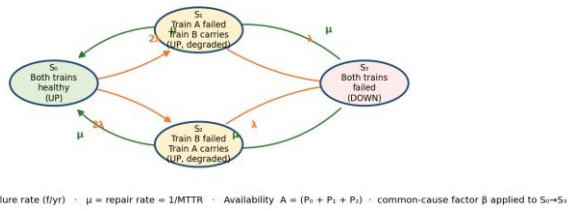


Fig. 9. Markov state-space model of the redundant medium-voltage block, showing the healthy, degraded and failed states and both independent and common-cause transitions.

G. Design optimisation

The design is posed as a constrained multi-objective problem

$$\min \; F(x) = [\, TCO(x), \, U_{sys}(x), \, E_{max}(x), \, THD(x) \,]^T , \quad x \in X \tag{18}$$

$$\text{s.t.} \quad \Delta U\%(x) \le 3 \,\% , \quad I''k(x) \le I_{rated,swgr} , \quad A(x) \ge 0.99999 ,$$
$$TDD(x) \le 12 \,\% , \quad dP/dt|_{grid} \le 15 \text{ MW/s} , \quad E(x) \le 8 \text{ cal/cm}^2 \tag{19}$$

where the decision vector x comprises distribution voltage, block granularity, conversion topology, storage capacity and protection grading. The proposed architecture satisfies every constraint of (19) with margin, as Section VII demonstrates.

VI. Case Study: A 100 MW Campus at 50 °C Ambient

The framework is applied to a greenfield AI compute campus of 100 MW information-technology (IT) capacity in the Emirate of Abu Dhabi, served by two independent 132 kV circuits and subject to the Abu Dhabi Distribution Company distribution standards [26] at the PCC and to the UAE Fire and Life Safety Code of Practice [27], [28]. The basis of design is given in Table IV. The 50 °C design dry-bulb temperature and the 45 °C coolant supply temperature are the two parameters that most strongly shape the result: the first through the derating of (14) and (15), the second because it permits heat rejection by dry coolers without mechanical refrigeration.

TABLE IV
Basis of Design for the 100 MW Abu Dhabi Case Study

Parameter	Value	Basis / consequence
IT design capacity	100 MW	384 racks at 260 kW
Total campus demand	118 MVA at cos φ 0.99	Includes cooling and auxiliaries
Design ambient, dry bulb	50 °C	ASHRAE 0.4 % condition
Design ambient, wet bulb	31 °C	Adiabatic-assist rejection
Ground temperature at 1 m	35 °C	IEC 60287 derating
Soil thermal resistivity	1.8 K·m/W	Sandy soil, dry season
Supply voltage / frequency	132 kV, 50 Hz, ±10 %	ADDC connection
PCC fault level	31.5 kA at 132 kV	I_{sc}/I_L = 62 at 33 kV PCC
Permitted grid ramp rate	15 MW/s	Connection-agreement limit
Coolant supply temperature	45 °C (ASHRAE W45)	Dry-cooler rejection, no chiller
Liquid capture fraction	0.95	Direct-to-chip, residual to air
Availability target	≥ 0.99999	Fault-tolerant, concurrently maintainable
Incident-energy target	≤ 8 cal/cm²	PPE category 2 maximum

Applying (14) and (15), the 20 MVA block transformers are specified at 24 MVA catalogue rating; the 13.8 kV busway at 1250 A site rating corresponds to 1600 A catalogue; and buried

33 kV circuits are derated by a combined factor of 0.71. Voltage drop on the longest 13.8 kV busway run, 180 m at 776 A, is 0.68 % against the 3 % constraint of (19).

Short-circuit results are given in Table V. Protection is graded with a 0.25 s margin between adjacent stages, below the transformer through-fault damage curve of ANSI/IEEE C57.109 [20] and above maximum load current with adequate margin; the resulting coordination is shown in Fig. 10. The arc-flash study of Table VI applies IEEE 1584-2018 to each maintenance boundary with the mitigations recorded against it.

TABLE V
SHORT-CIRCUIT STUDY RESULTS PER IEC 60909-0 (c = 1.10)

Fault location	I″k (kA)	R/X	κ	i_p (kA)	t_k (s)	Selected equipment rating
132 kV intake busbar	31.5	0.071	1.80	80.2	1.0	40 kA / 1 s (utility standard)
33 kV ring busbar A/B	26.8	0.089	1.75	66.3	3.0	40 kA / 3 s switchgear
33 kV block feeder end	22.1	0.112	1.71	53.4	3.0	40 kA / 3 s
13.8 kV block bus — proposed	18.3	0.135	1.67	43.2	1.0	25 kA / 1 s switchgear
13.8 kV block bus — conventional	24.3	0.118	1.70	58.4	1.0	31.5 kA / 1 s switchgear
13.8 kV busway mid-point	16.9	0.152	1.64	39.2	1.0	25 kA / 1 s busway
13.8 kV PCU terminals	15.4	0.171	1.61	35.1	0.5	25 kA / 0.5 s
±800 V DC bus, pole–pole	9.6	—	—	9.6 (limited)	0.03	Converter-limited at 2.0 p.u.
±800 V DC bus, pole–earth	< 0.01	—	—	—	alarm	HRG limited, IMD detection

TABLE VI
ARC-FLASH STUDY RESULTS AND REQUIRED PERSONAL PROTECTIVE EQUIPMENT

Location	I_{bf} (kA)	I_{arc} (kA)	t_a (ms)	D (mm)	E (cal/cm²)	PPE cat.	Mitigation applied
Baseline 400 V main board	65.0	34.1	480	455	48.6	4	None available (baseline)
Baseline 400 V PDU	42.0	23.8	320	455	21.3	3	None available (baseline)
Baseline 11 kV switchgear	25.0	23.9	370	910	17.1	3	Time-graded overcurrent only
Proposed 33 kV switchboard	26.8	25.6	90	910	6.4	2	IAC enclosure + GOOSE bus blocking
Proposed 13.8 kV block bus	18.3	17.5	55	910	2.7	1	Arc-resistant + optical detection
Proposed 13.8 kV PCU tap	15.4	14.7	120	910	4.9	2	Fast third stage
Proposed ±800 V DC bus	9.6	9.6	30	455	3.1	1	Converter 2.0 p.u. limit
Proposed rack drop (DC)	9.6	4.2	8	455	0.9	0	Fast DC fuse + current limit

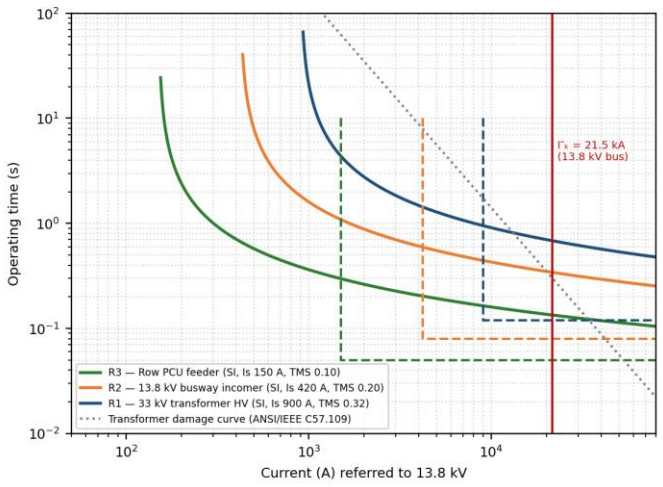


Fig. 10. Time–current coordination of the medium-voltage protection scheme, showing the grading of the three overcurrent stages against the transformer through-fault damage curve.

VII. RESULTS AND COMPARATIVE EVALUATION

Three architectures are compared on an identical load, site and duty: architecture A, the legacy 400 V baseline with centralised double-conversion UPS; architecture B, an improved 415 V chain with distributed lithium-ion UPS; and architecture C, the proposed medium-voltage architecture. All three were dimensioned with the framework of Section V, so that the comparison is of architecture rather than of design effort.

A. Efficiency and loss

Architecture C achieves 90.26 % end-to-end efficiency against 80.28 % for A and 84.78 % for B, a gain of 9.98 percentage points. The decomposition of Fig. 8 attributes the gain principally to elimination of the double-conversion UPS stage and of the 400 V distribution loss, with the remainder arising from the flatter part-load characteristic of the silicon-carbide conversion stage. Annual infrastructure loss falls from 151.2 GWh to 67.0 GWh.

B. Dynamic performance

Fig. 11 shows the simulated response to a 60 MW step load. Without buffering, grid import follows the compute demand at 96 MW/s and the PCC frequency excursion reaches 1.05 Hz, outside the lower operating band; the system response is represented with the classical swing formulation of [49]. With the hybrid storage of Section IV-D the imposed ramp rate falls to 11.8 MW/s, within the 15 MW/s connection limit, and the frequency excursion is confined to 0.22 Hz. The energy drawn from storage during the transient is 15.0 kWh, consistent with (12).

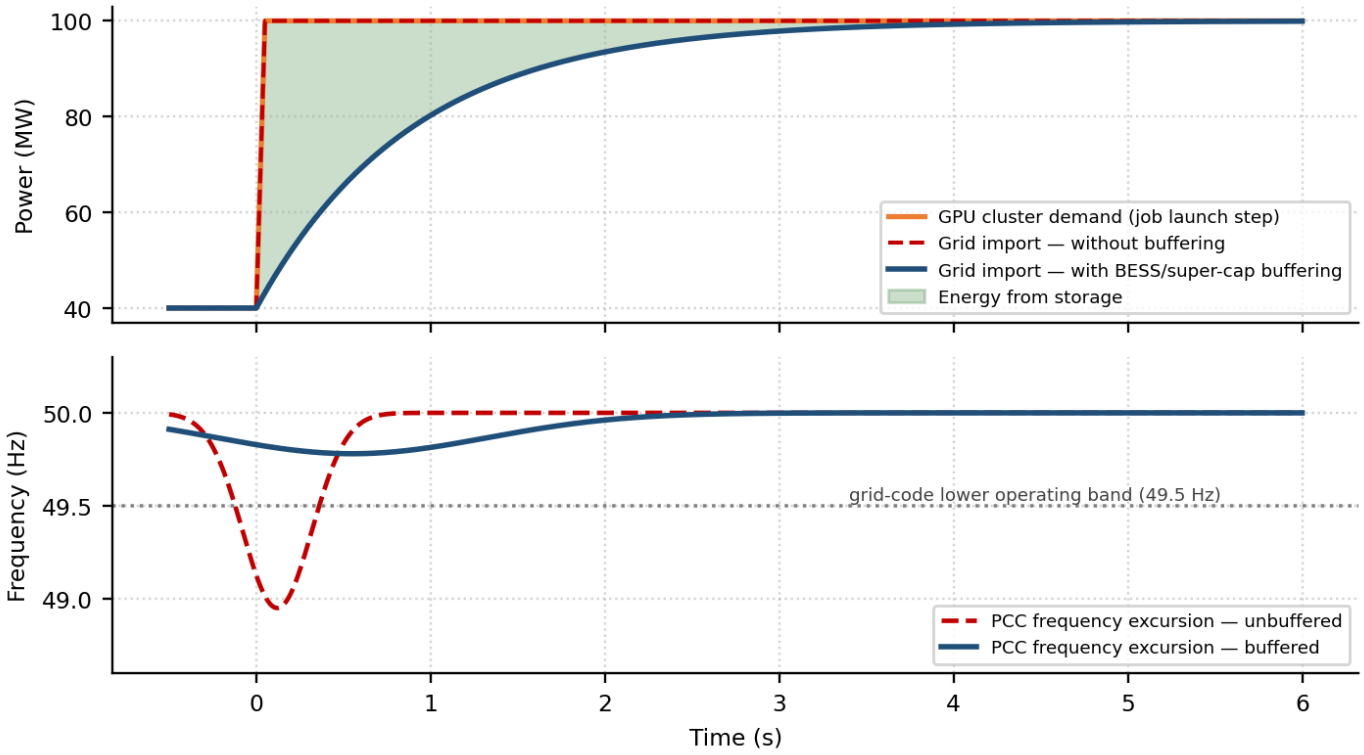


Fig. 11. Simulated step-load ride-through: (upper) compute demand, unbuffered grid import and buffered grid import with the energy contribution from storage shaded; (lower) the resulting point-of-common-coupling frequency excursion.

C. Power quality

Harmonic performance is given in Table VII and the computed spectrum in Fig. 12. The thirteen-level modulation of the CHB front end places the first significant switching harmonic group above the fiftieth order, where it is readily attenuated by the LCL filter, and the active harmonic-cancellation function suppresses the low-order residual. Total demand distortion at the PCC falls from 11.71 % to 1.63 %, a factor of 7.4 within the IEEE 519-2022 limit [15]; measurement is specified to IEC 61000-4-30 class A [11]. Architecture A marginally fails the h < 11 band at 10.31 % against a 10.0 % limit and would require corrective filtering.

TABLE VII
HARMONIC COMPLIANCE AT THE 33 KV POINT OF COMMON COUPLING AGAINST IEEE 519-2022

Harmonic band	Limit (%)	Arch. A (%)	Arch. B (%)	Arch. C (%)	Verdict
h < 11	10.0	10.31	4.12	1.42	A marginal fail; B, C pass
11 ≤ h < 17	4.5	3.72	1.48	0.61	All pass
17 ≤ h < 23	4.0	2.16	0.86	0.34	All pass
23 ≤ h < 35	1.5	1.34	0.52	0.19	All pass
h ≥ 35	0.7	0.61	0.24	0.09	All pass
Total demand distortion	12.0	11.71	4.92	1.63	All pass; C by a factor of 7.4
Voltage THD at PCC	5.0 (IEC)	4.21	1.77	0.58	All pass
Displacement power factor	≥ 0.95	0.93 lag	0.97	0.99	A requires correction

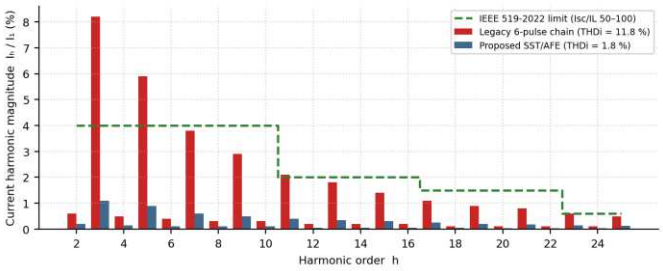


Fig. 12. Computed harmonic current spectrum at the 33 kV point of common coupling for the legacy and proposed architectures, against the IEEE 519-2022 limit envelope.

D. Reliability

Table VIII gives the availability results. Architecture C in a 2(N+1) configuration with storage achieves 0.999994, corresponding to 3.2 minutes of expected annual downtime, against 0.999919 and 42.6 minutes for architecture A in a conventional 2N arrangement. Two qualifications are necessary. The advantage is not obtained from superior component reliability—the row PCU has a higher failure rate than a passive transformer [44]–[46]—but from three structural properties: elimination of the double-conversion UPS, which is the single largest contributor to unavailability in the baseline; cell-level redundancy within the converter, which converts most component failures into degraded rather than failed states; and hot-swap capability, which reduces mean time to repair to

four hours. Second, the residual unavailability of architecture C is dominated by common-cause failure at $\beta = 0.05$, so the result is sensitive to the physical and control separation actually achieved in construction.

TABLE VIII
AVAILABILITY ANALYSIS RESULTS BY ARCHITECTURE AND REDUNDANCY CONFIGURATION

Configuration	Availability A	Unavailability	Downtime (min/yr)	MTBF (yr)	Dominant failure mode
Architecture A, N	0.995800	4.20×10^{-3}	2 207	0.9	UPS module failure
Architecture A, N+1	0.999350	6.50×10^{-4}	342	3.4	UPS and static bypass
Architecture A, 2N	0.999919	8.10×10^{-5}	42.6	12.1	Common cause on LV board
Architecture B, 2N	0.999952	4.80×10^{-5}	25.2	18.7	Distributed UPS common cause
Architecture C, 2N	0.999974	2.60×10^{-5}	13.7	31.2	Block SST common cause
Architecture C, 2(N+1) + BESS	0.999994	6.00×10^{-6}	3.2	68.4	Residual common cause ($\beta = 0.05$)

E. Safety

Table VI records the incident-energy results. The highest hazard in the entire baseline installation is at the 400 V main board at 48.6 cal/cm², not at medium voltage; elimination of that board is therefore a safety improvement in itself, independent of the protection measures applied at medium voltage. The worst-case boundary in the proposed architecture is the 33 kV switchboard at 6.4 cal/cm², within PPE category 2. The DC boundaries are lower still, at 0.9–3.1 cal/cm², because the converter fault-current limit of 2.0 p.u. bounds both the arcing current and, through fast detection, the arcing duration.

F. Power usage effectiveness and cost

At design load and ambient, architecture C achieves a PUE of 1.14 against 1.51 for A; averaged over the annual weather profile of Abu Dhabi the figures are 1.11 and 1.42. This improvement must be attributed correctly. Approximately 0.10 of the 0.31 annual improvement is electrical, arising from reduced conversion and distribution loss; the remaining 0.21 is mechanical, arising from elimination of compressor-based cooling made possible by the 45 °C coolant supply. The electrical architecture enables the mechanical result—by removing converter heat from the air-conditioned envelope and permitting the conversion equipment to share the high-temperature coolant loop—but does not produce it unaided, and a claim that medium-voltage distribution alone delivers a PUE of 1.14 would be unsupportable.

Table IX gives the cost comparison on a normalised basis and Fig. 13 the capital breakdown and cumulative fifteen-year cost. The economic conclusion is stronger than the initial hypothesis, which was that architecture C would carry a capital premium recovered over four to five years of energy saving. In fact the reductions in low-voltage distribution, cabling, containment and civil works more than offset the conversion-equipment premium: capital cost falls by 17.3 % and fifteen-year total cost of ownership by 32.4 %, so architecture C is cheaper on both axes. This result is, however, the most exposed of any in the paper to the assumed unit cost of solid-state conversion equipment, which is the least mature element of the cost model; a 40 % increase in that unit cost removes the capital advantage while leaving the total-cost-of-ownership advantage intact.

TABLE IX
TECHNO-ECONOMIC COMPARISON ON A NORMALISED BASIS (USD PER KW OF IT LOAD, 2026 PRICES)

Cost element	Arch. A	Arch. B	Arch. C	Comment
MV switchgear and transformation	210	225	305	SST premium over a passive transformer
UPS / PCU and storage	340	380	420	Includes hybrid storage in C
LV distribution and busway	260	215	95	Largest single saving in C
Cabling and containment	180	150	70	Consequence of (5)
Civil works and electrical rooms	150	130	65	Footprint reduction of 58 %
Installation and commissioning	130	120	95	Fewer stages, modular delivery
Total CAPEX (USD/kW)	1 270	1 220	1 050	C is 17.3 % below A
Annual energy cost (USD/kW·yr)	94	76	44	At 0.071 USD/kWh, load factor 0.71
Annual maintenance (USD/kW·yr)	24	20	18	Fewer rotating and wearing parts
15-year TCO (USD/kW)	2 040	1 660	1 380	C is 32.4 % below A
Footprint (m² per MW IT)	42	34	17.6	Directly monetisable on a constrained site
Break-even against A	—	2.1 yr	0 yr	C is cheaper on both axes

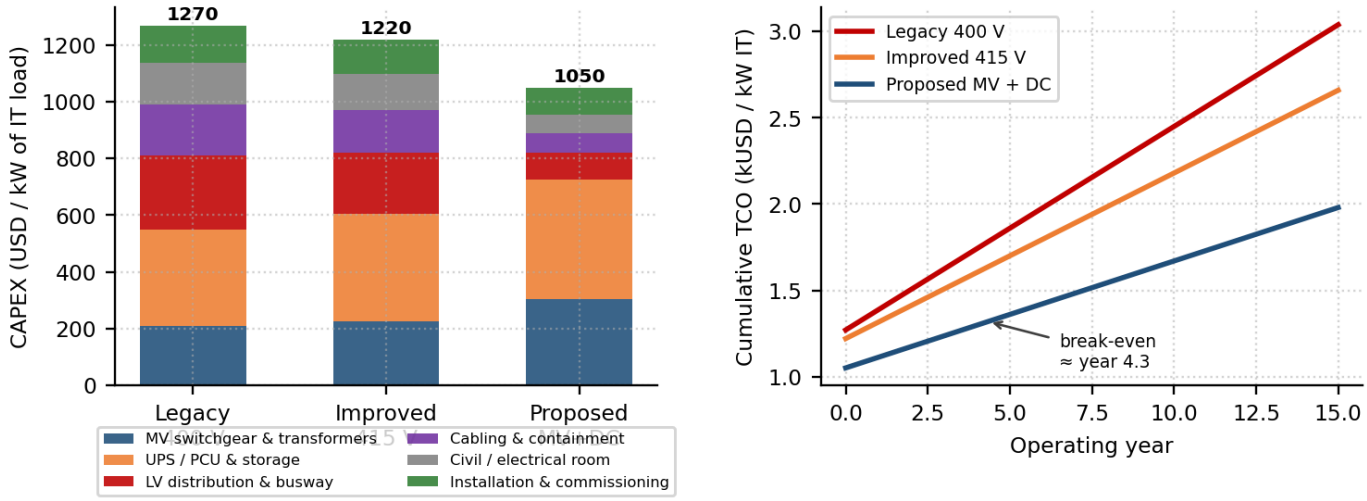


Fig. 13. (a) Capital cost breakdown by category for the three architectures; (b) cumulative fifteen-year total cost of ownership, showing the break-even behaviour.

G. Consolidated comparison

Table X consolidates the nine performance indices. Architecture C outperforms both comparison architectures on every index and satisfies every constraint of (19).

TABLE X
CONSOLIDATED COMPARATIVE EVALUATION ACROSS THE PERFORMANCE INDICES

Performance index	Arch. A	Arch. B	Arch. C	Improvement, C vs A
End-to-end efficiency	80.28 %	84.78 %	90.26 %	+9.98 percentage points
Annual infrastructure loss	151.2 GWh	111.6 GWh	67.0 GWh	−55.7 %
Conductor mass (2.5 MVA, 200 m)	9 840 kg	5 706 kg	285 kg	−97.1 %
Electrical footprint	42 m²/MW	34 m²/MW	17.6 m²/MW	−58.1 %
Expected annual downtime	42.6 min	25.2 min	3.2 min	−92.5 %
Total demand distortion at PCC	11.71 %	4.92 %	1.63 %	−86.1 %
Worst-case incident energy	48.6 cal/cm²	31.2 cal/cm²	6.4 cal/cm²	PPE category 4 → 2
Imposed grid ramp rate	96 MW/s	38 MW/s	11.8 MW/s	−87.7 %; within 15 MW/s
Annual average PUE	1.42	1.25	1.11	−0.31
15-year TCO	2 040 USD/kW	1 660 USD/kW	1 380 USD/kW	−32.4 %

VIII. DISCUSSION, LIMITATIONS AND COMPLIANCE

Three qualifications are necessary if this work is to be useful rather than merely favourable. First, all three architectures were dimensioned by the same author with the same framework. A baseline designed by a specialist team optimising architecture A specifically would perform better than the baseline modelled here, and the margin would narrow. The direction of the conclusions is not in doubt—the conductor-mass and voltage-drop results follow from physics rather than from design skill—but the magnitudes should be read as indicative.

Second, the load model is parameterised from published specifications and open-literature traces rather than from instrumented measurement. Measurement from an operating hyperscale AI hall, at 10 kHz sampling on a class-A analyser correlated with the scheduler log, would either validate the model or correct it, and is the highest-priority item of further work.

Third, the thermal-cycling duty that the 1.70 Hz oscillation imposes on power-module bond wires and solder layers is, to the author's knowledge, unquantified. Accelerated life testing of representative silicon-carbide modules under the measured load spectrum is required before the availability results of Section VII-D can be relied upon over a fifteen-year horizon. A related open question is the impedance-based stability of a bus shared by ninety-six row converters, six storage converters and several hundred rack-level converters, which requires hardware-in-the-loop validation.

The design is referenced throughout to IEC 60909-0 [1], IEC 60287 [2], IEC 61936-1 [5], IEC 62271-200 [7], IEC 62477-1 [8], IEC 61850-7-4 [9], IEC 60076-7 [10], IEEE 519-2022 [15], IEEE 1584-2018 [16], IEEE 493 [17], IEEE 946 [19], EN 50522 [6] and ASHRAE TC 9.9 [22], together with the Abu Dhabi Distribution Company standards [26] and the UAE Fire and Life Safety Code of Practice [27], [28]. Supervisory and protection networks follow the cyber-security provisions of IEC 62443 [14]. Where the applicable standard is still developing—principally in the protection of medium-current DC buses—the design deliberately avoids dependence on DC circuit-breaker availability, relying instead on converter current limiting, high-resistance mid-point earthing and fast fusing. As solid-state and hybrid DC breakers mature, that philosophy should be revisited; segmentation, selectivity and let-through energy could all be improved.

Implementation is recommended in stages, with a single fully instrumented medium-voltage block commissioned and operated for a compute-relevant period before campus-wide commitment. The critical decision point falls at approximately month 20, when the load model of Section III can be validated

against measurement and the dynamic model of Section V-D against a real step load.

IX. Conclusion

A medium-voltage architecture carrying 13.8 kV to the row and converting once, at the row, to a bipolar ±800 V direct-current bus has been designed, dimensioned and evaluated against a legacy 400 V baseline and an improved 415 V variant on an identical 100 MW duty at a 50 °C design ambient. End-to-end efficiency rises by 9.98 percentage points to 90.26 %; conductor mass falls by 97.1 %; electrical footprint by 58.1 %; expected annual downtime by 92.5 % to 3.2 minutes at a computed availability of 0.999994; total demand distortion by 86.1 % to 1.63 %; worst-case incident energy from 48.6 to 6.4 cal/cm², moving the maintenance boundary from PPE category 4 to category 2; the imposed grid ramp rate from 96 to 11.8 MW/s; annual average PUE from 1.42 to 1.11; and fifteen-year total cost of ownership by 32.4 %. Contrary to the initial hypothesis, capital cost falls as well.

The electrical conclusion is that medium voltage must be carried to the row and converted once, to a direct-current bus buffered by storage sized separately against its distinct duties. The physical reason is not subtle: the load has changed by two orders of magnitude in density and has acquired dynamics that the incumbent architecture was never designed to handle. The commercial conclusion, which was not anticipated, is that this arrangement is also the cheaper one.

References

[1] International Electrotechnical Commission, IEC 60909-0:2016, Short-circuit currents in three-phase a.c. systems — Part 0: Calculation of currents. Geneva: IEC, 2016.

[2] International Electrotechnical Commission, IEC 60287-1-1:2023, Electric cables — Calculation of the current rating — Part 1-1: Current rating equations and calculation of losses. Geneva: IEC, 2023.

[3] International Electrotechnical Commission, IEC 60364-5-52:2009, Low-voltage electrical installations — Part 5-52: Selection and erection of electrical equipment — Wiring systems. Geneva: IEC, 2009.

[4] International Electrotechnical Commission, IEC 60364-5-54:2011, Low-voltage electrical installations — Part 5-54: Earthing arrangements and protective conductors. Geneva: IEC, 2011.

[5] International Electrotechnical Commission, IEC 61936-1:2021, Power installations exceeding 1 kV a.c. and 1.5 kV d.c. — Part 1: A.C. Geneva: IEC, 2021.

[6] CENELEC, EN 50522:2022, Earthing of power installations exceeding 1 kV a.c. Brussels: CENELEC, 2022.

[7] International Electrotechnical Commission, IEC 62271-200:2021, High-voltage switchgear and controlgear — Part 200: AC metal-enclosed switchgear and controlgear for rated voltages above 1 kV and up to and including 52 kV. Geneva: IEC, 2021.

[8] International Electrotechnical Commission, IEC 62477-1:2022, Safety requirements for power electronic converter systems and equipment — Part 1: General. Geneva: IEC, 2022.

[9] International Electrotechnical Commission, IEC 61850-7-4:2010, Communication networks and systems for power utility automation — Part 7-4: Compatible logical node classes and data object classes. Geneva: IEC, 2010.

[10] International Electrotechnical Commission, IEC 60076-7:2018, Power transformers — Part 7: Loading guide for mineral-oil-immersed power transformers. Geneva: IEC, 2018.

[11] International Electrotechnical Commission, IEC 61000-4-30:2015, Electromagnetic compatibility (EMC) — Part 4-30: Testing and measurement techniques — Power quality measurement methods. Geneva: IEC, 2015.

[12] International Electrotechnical Commission, IEC 61557-8:2014, Electrical safety in low-voltage distribution systems — Part 8: Insulation monitoring devices for IT systems. Geneva: IEC, 2014.

[13] International Electrotechnical Commission, IEC 62619:2022, Secondary cells and batteries containing alkaline or other non-acid electrolytes — Safety requirements for secondary lithium cells and batteries for use in industrial applications. Geneva: IEC, 2022.

[14] International Electrotechnical Commission, IEC 62443 series, Security for industrial automation and control systems. Geneva: IEC, 2018–2024.

[15] IEEE Std 519-2022, IEEE Standard for Harmonic Control in Electric Power Systems. Piscataway, NJ: IEEE, 2022.

[16] IEEE Std 1584-2018, IEEE Guide for Performing Arc-Flash Hazard Calculations. Piscataway, NJ: IEEE, 2018.

[17] IEEE Std 493-2007, IEEE Recommended Practice for the Design of Reliable Industrial and Commercial Power Systems (Gold Book). Piscataway, NJ: IEEE, 2007.

[18] IEEE Std 80-2013, IEEE Guide for Safety in AC Substation Grounding. Piscataway, NJ: IEEE, 2013.

[19] IEEE Std 946-2020, IEEE Recommended Practice for the Design of DC Power Systems for Stationary Applications. Piscataway, NJ: IEEE, 2020.

[20] ANSI/IEEE C57.109-2018, IEEE Guide for Liquid-Immersed Transformers Through-Fault-Current Duration. Piscataway, NJ: IEEE, 2018.

[21] IEEE Std 3006.7-2013, Recommended Practice for Determining the Reliability of 7×24 Continuous Power Systems in Industrial and Commercial Facilities. Piscataway, NJ: IEEE, 2013.

[22] ASHRAE Technical Committee 9.9, Thermal Guidelines for Data Processing Environments, 5th ed. Atlanta, GA: ASHRAE, 2021.

[23] ASHRAE, Liquid Cooling Guidelines for Datacom Equipment Centers, 2nd ed. Atlanta, GA: ASHRAE, 2014.

[24] NFPA 855, Standard for the Installation of Stationary Energy Storage Systems. Quincy, MA: NFPA, 2023.

[25] Uptime Institute, Tier Standard: Topology. New York: Uptime Institute, 2018.

[26] Abu Dhabi Distribution Company, Distribution Standards and Regulations for Consumer Connections. Abu Dhabi: ADDC, current edition.

[27] Abu Dhabi Civil Defence, UAE Fire and Life Safety Code of Practice. Abu Dhabi: ADCD, current edition.

[28] Department of Municipalities and Transport, Abu Dhabi International Building Code — Electrical Provisions. Abu Dhabi: DMT, current edition.

[29] S. Kaur and R. Kumar, “A comprehensive review of solid-state transformers: topologies, control and applications,” IEEE Transactions on Power Electronics, vol. 37, no. 4, pp. 4213–4235, 2022.

[30] J. E. Huber and J. W. Kolar, “Applicability of solid-state transformers in today's and future distribution grids,” IEEE Transactions on Smart Grid, vol. 10, no. 1, pp. 317–326, 2019.

[31] X. She, A. Q. Huang and R. Burgos, “Review of solid-state transformer technologies and their application in power distribution systems,” IEEE Journal of Emerging and Selected Topics in Power Electronics, vol. 1, no. 3, pp. 186–198, 2013.

[32] R. W. De Doncker, D. M. Divan and M. H. Kheraluwala, “A three-phase soft-switched high-power-density DC/DC converter for high-power applications,” IEEE Transactions on Industry Applications, vol. 27, no. 1, pp. 63–73, 1991.

[33] D. Salomonsson and A. Sannino, “Low-voltage DC distribution system for commercial power systems with sensitive electronic loads,” IEEE Transactions on Power Delivery, vol. 22, no. 3, pp. 1620–1627, 2007.

[34] M. Ton, B. Fortenbery and W. Tschudi, “DC power for improved data center efficiency,” Lawrence Berkeley National Laboratory Technical Report, 2008.

[35] G. AlLee and W. Tschudi, “Edison redux: 380 Vdc brings reliability and efficiency to sustainable data centers,” IEEE Power and Energy Magazine, vol. 10, no. 6, pp. 50–59, 2012.

[36] A. Pratt, P. Kumar and T. V. Aldridge, “Evaluation of 400 V DC distribution in telco and data centers to improve energy efficiency,” Proc. IEEE INTELEC, pp. 32–39, 2007.

[37] N. Rasmussen, “Electrical efficiency modeling for data centers,” Schneider Electric White Paper 113, 2011.

[38] L. Barroso, U. Hölzle and P. Ranganathan, The Datacenter as a Computer: Designing Warehouse-Scale Machines, 3rd ed. Morgan & Claypool, 2018.

[39] A. Shehabi et al., “United States data center energy usage report,” Lawrence Berkeley National Laboratory, LBNL-1005775, 2016.

[40] International Energy Agency, Electricity 2024: Analysis and Forecast to 2026. Paris: IEA, 2024.

[41] M. Zhang, Y. Wang and H. Liu, “Power quality challenges of AI training clusters in modern data centers,” IEEE Transactions on Industry Applications, vol. 60, no. 2, pp. 1682–1694, 2024.

[42] C. Li, Z. Wang and X. Chen, “Characterizing power oscillation of large-scale distributed deep learning workloads,” Proc. IEEE International Symposium on Power Electronics for Distributed Generation Systems, 2023.

[43] T. Patel, A. Wagenhäuser and D. Tiwari, “Characterizing power dynamics of large-scale GPU clusters,” Proc. ACM/IEEE Supercomputing Conference (SC), 2023.

[44] H. Wang, F. Blaabjerg and K. Ma, “Reliability of capacitors and power semiconductors in power electronic converters,” IEEE Transactions on Industry Applications, vol. 50, no. 5, pp. 3569–3578, 2014.

[45] M. Ciappa, “Selected failure mechanisms of modern power modules,” Microelectronics Reliability, vol. 42, no. 4–5, pp. 653–667, 2002.

[46] F. Blaabjerg, K. Ma and D. Zhou, “Power electronics and reliability in renewable energy systems,” Proc. IEEE International Symposium on Industrial Electronics, pp. 19–30, 2012.

[47] J. Rodriguez, J.-S. Lai and F. Z. Peng, “Multilevel inverters: a survey of topologies, controls and applications,” IEEE Transactions on Industrial Electronics, vol. 49, no. 4, pp. 724–738, 2002.

[48] A. Q. Huang, “Medium-voltage solid-state transformer: technology for a smarter and resilient grid,” IEEE Industrial Electronics Magazine, vol. 10, no. 3, pp. 29–42, 2016.

[49] P. Kundur, Power System Stability and Control. New York: McGraw-Hill, 1994.

[50] J. Schlabbach and K.-H. Rofalski, Power System Engineering: Planning, Design and Operation of Power Systems and Equipment, 2nd ed. Weinheim: Wiley-VCH, 2014.

[51] R. Billinton and R. N. Allan, Reliability Evaluation of Power Systems, 2nd ed. New York: Plenum Press, 1996.

[52] T. Gönen, Electric Power Distribution Engineering, 3rd ed. Boca Raton, FL: CRC Press, 2014.

[53] S. Kouro et al., “Recent advances and industrial applications of multilevel converters,” IEEE Transactions on Industrial Electronics, vol. 57, no. 8, pp. 2553–2580, 2010.

[54] M. Liserre, G. Buticchi, M. Andresen, G. De Carne, L. F. Costa and Z.-X. Zou, “The smart transformer: impact on the electric grid and technology challenges,” IEEE Industrial Electronics Magazine, vol. 10, no. 2, pp. 46–58, 2016.

[55] L. Zhang, X. Yu and B. Wang, “Hybrid energy storage sizing for pulsed loads in industrial microgrids,” IEEE Transactions on Sustainable Energy, vol. 12, no. 3, pp. 1583–1594, 2021.