Reliability-Constrained Power Utilization in AI Data Centers: A Probabilistic Framework and Reference 100-MW Case Study

Sudhir Agarwal, Ph.D.
General Reliability San Diego CA USA
generalreliability@gmail.com

***Abstract*—**Artificial-intelligence data centers combine rapidly varying computational loads with highly redundant electrical, thermal, control, and communications infrastructure. Conventional capacity planning often treats installed electrical capacity and reserve margins deterministically, while production experience demonstrates that workload diversity, power capping, redundancy sharing, and operational controls can create substantial statistical headroom. This paper proposes a reliability-constrained power-utilization framework that asks a different question: how much compute-equivalent load can a data center safely provision while satisfying specified reliability and service-risk targets? The framework combines stochastic workload modeling, component and common-mode failure states, chronological or state-sampling Monte Carlo simulation, and reliability indices analogous to loss-of-load and expected-unserved-energy measures used in electric-power planning. A transparent 100-MW reference case is developed using public literature to calibrate workload behavior and using explicitly labeled illustrative reliability assumptions where field data are unavailable. The case study demonstrates a nonlinear utilization-risk frontier: statistical headroom can be material at moderate provisioning levels, but deficiency probability rises sharply as workload tails approach degraded infrastructure states. The analysis also shows that common-mode and dependent failures can materially alter conclusions obtained from independent-component models. The proposed framework is intended as a reproducible starting point for utility, hyperscaler, colocation, and consulting studies that balance utilization, reliability, risk, and cost.

***Index Terms*—**AI data centers, availability, common-mode failure, Monte Carlo simulation, power headroom, power oversubscription, probabilistic reliability, risk, workload forecasting.

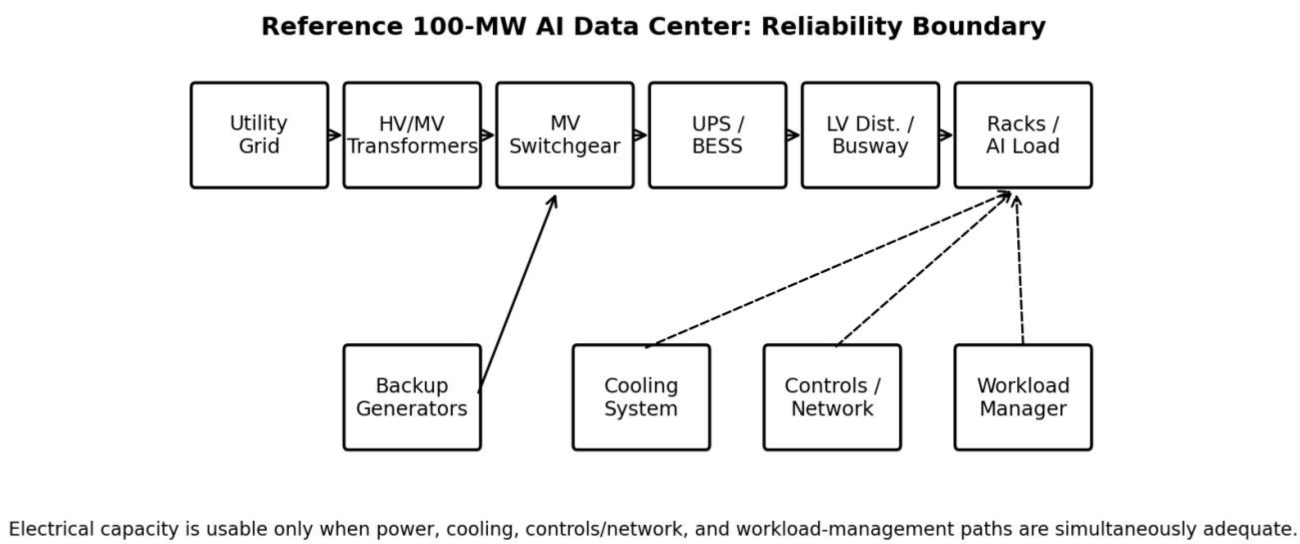

Fig. 1. Reliability boundary for the 100-MW reference AI data center. A compute workload is serviceable only when the electrical source/distribution path and its thermal, control, communications, and workload-management dependencies remain adequate.

I. INTRODUCTION

The electric-power industry has long used reliability indices to quantify the adequacy and performance of generation, transmission, substation, and distribution systems. Deterministic criteria remain essential for secure planning and operation, but probabilistic methods add a second question: not only whether equipment can carry a specified loading, but what risk accompanies that loading and how the risk changes as operating margin is consumed. The rapid buildout of AI data centers makes this question unusually valuable because power is simultaneously a hard infrastructure constraint, a reliability dependency, and a large capital cost.

Modern data centers are not passive loads. They are coupled cyber-physical systems whose useful output depends on the simultaneous availability of grid supply, transformers, switchgear, uninterruptible power supplies (UPS), batteries, generators, low-voltage distribution, cooling, networking, controls, and computational workloads. Barroso et al. describe warehouse-scale computing as a computer composed of the data-center system itself [1]. Recent power-system work extends this perspective to the electrical interface from the grid point of interconnection to the GPU [2].

Two trends motivate a reliability-constrained utilization framework. First, measured and production workload data show that the electrical demand of computing systems is stochastic and frequently below aggregate nameplate power. Google reported production deployments in which coordinated power sharing and priority-aware capping enabled power oversubscription of 25% or higher while preserving workload availability and performance [3]. Second, infrastructure failure and degradation cannot be represented adequately by independent component availability alone. Meta has emphasized zero-notice power-loss readiness, bounded tolerable impacts, recovery time, and defense-in-depth testing across mechanical, electrical, storage, compute, and orchestration layers [4].

The objective of this paper is therefore to formulate and demonstrate a framework for estimating reliability-constrained usable power: the maximum compute-equivalent provisioning or electrical utilization that can be supported while specified risk limits remain satisfied. The contribution is not a claim that every data center contains unused capacity. Rather, it is a method for determining whether headroom exists, what limits it, how dependent failures alter it, and what operational or capital actions could change the frontier.

II. RELATIONSHIP TO PRIOR WORK

A. Data-Center Electrical Reliability

Reliability assessment of data-center power architectures is established in the literature. Shrestha et al. compared 480-V AC and 380-V DC architectures using Monte Carlo simulation and different UPS redundancy levels, showing that architecture and component count materially affect system reliability [5]. Related work has used Markov, reliability-block, fault-tree, and Monte Carlo methods to represent UPS, PDU, PSU, breaker, battery, and generator states. Yu et al. recently developed lifecycle probability-of-power-insufficiency and energy-insufficiency measures for different data-center tiers and evaluated battery participation in grid flexibility services [6]. These studies establish that data-center reliability can be quantified using methods closely related to those used in bulk-power adequacy studies.

B. Workload Power and Oversubscription

The computing literature approaches the same system from the demand side. Google has demonstrated medium-voltage power sharing domains and a fast workload-aware power-capping service across production clusters [3]. Public Google PowerData traces now expose power-utilization information for 57 power domains, providing a basis for reproducible oversubscription research [7]. More recently, Vercellino et al. measured H100-based training, fine-tuning, and inference workloads at 0.1-s resolution and released the resulting power profiles, then used them in event-driven whole-facility simulations [8]. Such datasets make it possible to model demand distributions rather than assume coincident nameplate load.

C. Dynamic Stability and Large AI Loads

Reliability must also be distinguished from dynamic stability. Sun et al. documented resonance events in data centers and developed impedance-based models and design criteria for converter-dominated data-center power systems [9], [10]. Choukse et al. showed that large synchronous AI training jobs can create substantial power swings and proposed software, GPU-hardware, and facility-level stabilization approaches [11]. These phenomena are not captured by annual availability indices alone. A complete framework therefore requires multiple timescales: steady-state utilization and adequacy, component outage and restoration processes, and sub-second-to-minute dynamic constraints.

III. PROBLEM DEFINITION

Let C_inst denote installed electrical capacity at the chosen reliability boundary, D(t) the stochastic IT electrical demand, and C_av(t) the capacity available after equipment outages, maintenance, deratings, support-system states, and operating constraints. A capacity-deficiency event occurs when D(t) exceeds C_av(t).

$$I(t) = 1 \;\; \text{if} \;\; D(t) > C\_av(t), \;\; \text{else} \;\; 0 \quad (1)$$

For a provisioning level P, the capacity-deficiency probability is estimated as

$$P\_def(P) = Pr[D\_P(t) > C\_av(t)] \quad (2)$$

An energy-severity metric analogous to expected unserved energy (EUE) is defined here as expected unserved compute energy (EUCE):

$$EUCE(P) = E[\, max(D\_P(t) - C\_av(t), 0) \,] \times 8760 \quad (3)$$

A time-severity index analogous to loss-of-load expectation is the loss-of-compute expectation (LOCE):

$$LOCE(P) = P\_def(P) \times 8760 \;\; hours/year \quad (4)$$

The reliability-constrained usable capacity is then

$$P^* = max\, P \;\; \text{subject to} \;\; P\_def(P) \le \varepsilon\_p, \;\; EUCE(P) \le \varepsilon\_e, \;\; \text{and operating constraints} \quad (5)$$

The thresholds ε_p and ε_e are policy choices, not universal constants. A hyperscaler may establish separate targets by workload class, fault domain, service-level objective, or consequence category. This is analogous to defining different reliability criteria for bulk-power adequacy, transmission security, and customer interruption performance.

IV. REFERENCE 100-MW AI DATA CENTER

The reference case represents a 100-MW IT-capacity facility or fault domain. It is intentionally generic and is not intended to reproduce a specific hyperscaler design. The electrical path contains utility/grid supply, five 25-MW transformer blocks, MV switchgear, an 11-by-10-MW UPS plant, LV distribution, and racks. Backup generation is represented by eleven 10-MW units. Cooling is represented by eleven 10-MW-equivalent modules so that the thermal system can become a capacity constraint rather than an external perfect dependency. Network/control and workload-management functions are represented explicitly as support states.

Model element	Reference assumption
IT design boundary	100 MW; provisioning P varied from 80 to 150 MW compute-equivalent nameplate.
Workload power	Normalized stochastic factor 0.2 + 0.6·Beta(5,3); chosen to represent substantial diversity/tail behavior consistent with public AI workload studies, not a site-specific fit.
Transformers	5 × 25 MW; per-state unavailability 0.002; one unit in planned maintenance 1% of sampled states.
UPS	11 × 10 MW; module unavailability 0.0015.
Cooling	11 × 10 MW-equivalent; module unavailability 0.002.
Grid / generation	Grid unavailability 0.0003; during grid outage, 11 × 10 MW generators with 0.02 unit failure probability.
Common-mode state	Illustrative low-frequency control/network derating to 60 MW plus rarer zero-capacity state.
Simulation	5,000,000 independent operating-state samples; annual indices obtained by state-probability scaling.

TABLE I. Reference-case assumptions. Failure and maintenance parameters are illustrative engineering assumptions for framework demonstration; they must be replaced by site-specific or validated fleet statistics before operational use.

The reliability parameters in Table I are deliberately transparent. They are not presented as industry-average failure rates. Public literature provides architecture examples and broad reliability data, but hyperscaler component failure, maintenance, and common-mode statistics are rarely disclosed at the resolution required for a calibrated model. The correct engineering practice is therefore to separate public evidence from assumed inputs and perform sensitivity analysis until field data become available.

V. STOCHASTIC MODELING METHOD

A. Workload Model

For each sampled operating state, normalized workload power w is drawn from a bounded beta distribution. Provisioned compute-equivalent power P scales the electrical demand as $D = P \cdot w$. The beta distribution is used only as an illustrative surrogate for an empirical trace. In an implementation, the preferred approach is to fit or directly resample public or customer power traces, preserve autocorrelation, and condition demand on workload class, time of day, temperature, scheduler state, and job synchronization.

The need to preserve tail behavior is critical. Average utilization can be modest while the P99 or P99.9 demand approaches the electrical boundary. Extreme-value or peaks-over-threshold methods can therefore be added when sufficient data exist. Bayesian updating can revise tail and correlation parameters as operating evidence accumulates.

B. Capacity-State Model

Component states are sampled and converted to subsystem capacities. Available IT power is the minimum simultaneous capability across source, transformer, UPS, cooling, and control/network constraints. This “minimum-capacity” abstraction is appropriate for a first adequacy model. A detailed implementation would replace it with topology-aware power-flow and switching logic, including feeder paths, bus sections, protection, transfer schemes, battery state of charge, generator start sequences, fuel, and restoration times.

C. Dependent and Common-Mode Failures

Independent failure assumptions can materially overstate the value of redundancy. Common-mode mechanisms include shared switchgear or controls, common UPS firmware, common battery strings or DC controls, generator fuel and starting systems, chilled-water loops, network/configuration faults, software orchestration, shared utility contingencies, and environmental hazards. The reference case therefore includes a small-probability common-mode derating state. In production work, dependencies should be represented by conditional failure probabilities, common-cause factors, copulas, Bayesian networks, multi-state Markov models, or event-driven scenario logic depending on available evidence.

D. Chronological Extension

The present calculation uses state sampling to expose the utilization-risk relationship with minimal assumptions. A production model should be chronological. Sequential Monte Carlo simulation can preserve outage duration, repair, maintenance, battery state of charge, generator start/run failure, workload autocorrelation, weather, and restoration. This is especially important when the consequence depends on event duration rather than instantaneous capacity alone.

VI. ILLUSTRATIVE RESULTS

Fig. 2 compares mean and P99 demand as compute-equivalent provisioning is increased. Because the reference workload is diversified, the P99 electrical demand remains below the 100-MW design boundary until provisioning approaches approximately 130 MW. This illustrates the statistical origin of headroom. It does not establish that 30% oversubscription is safe: infrastructure degradation and tail-risk constraints determine the usable limit.

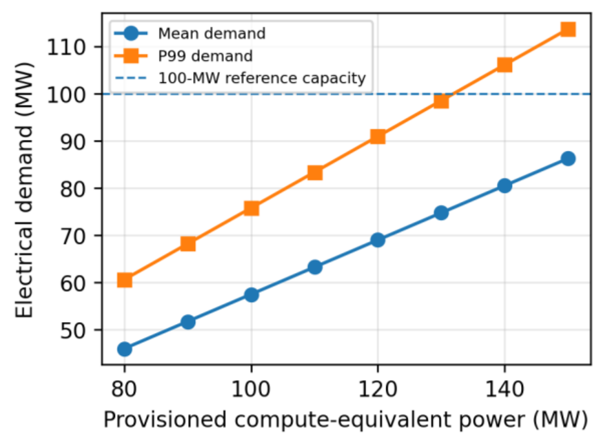

Fig. 2. Illustrative workload headroom. Provisioned compute-equivalent power exceeds the electrical boundary before the workload P99 reaches it because not all compute resources draw peak power simultaneously.

Fig. 3 shows the estimated capacity-deficiency probability with and without the illustrative common-mode state. Risk remains low across a range of provisioning levels and then rises nonlinearly as the workload tail overlaps N+1/degraded capacity states. The common-mode state is most visible at lower and moderate provisioning levels, where independent N+1 redundancy otherwise masks much of the component-outage risk.

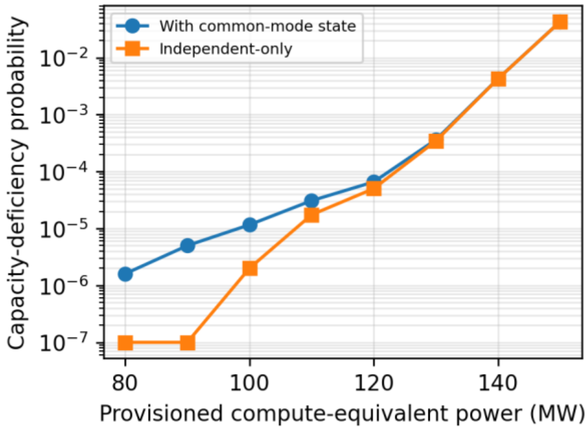

Fig. 3. Illustrative capacity-deficiency probability. Explicit common-mode states increase risk relative to an independent-component model.

P99 D (MW)	LOCE (h/yr)	EUCE (MWh/yr)
75.8	0.102	1.01
83.3	0.27	2.25
90.9	0.576	4.97
98.5	3.17	12.7
106.1	37.4	128
113.7	377	1.34e+03

TABLE II. Illustrative state-sampling results. These values demonstrate framework behavior and are not predictions for any actual facility.

Table II makes the tradeoff explicit. In the reference assumptions, increasing provisioning from 100 to 120 MW increases P99 demand from about 75.8 to 90.9 MW while LOCE remains below one expected hour per year. At 130 MW, the P99 demand approaches the 100-MW boundary and LOCE increases to several hours per year. At 140 MW, the workload tail frequently overlaps degraded 100-MW and 110-MW infrastructure states, causing a sharp increase in both frequency and energy severity. This “knee” is the feature a reliability-constrained optimization should locate.

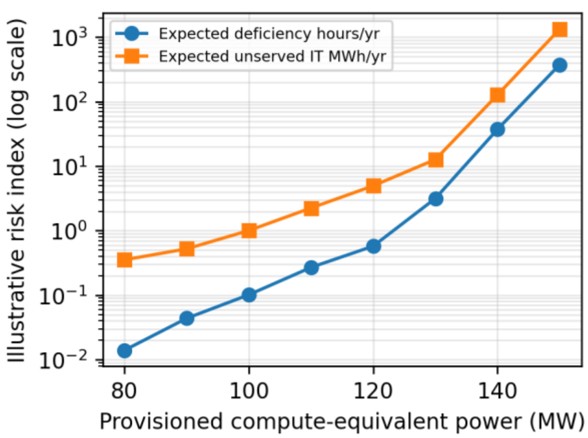

Fig. 4. Illustrative reliability indices versus provisioned compute-equivalent power. The nonlinear rise identifies a risk-utilization frontier rather than a single deterministic capacity number.

VII. FROM RELIABILITY INDICES TO A RISK-UTILIZATION MATRIX

A useful operational product is a risk-utilization matrix for major components and fault domains. For each transformer, breaker, UPS block, generator group, cooling train, network/control dependency, and compute cluster, the matrix would contain present loading, stochastic demand exposure, available redundancy, failure and repair statistics, common-mode dependencies, consequence severity, and the incremental economic value of additional utilization. Such a matrix provides more information than a single facility availability number because it identifies which constraint is active and which mitigation has the greatest value.

The same principle used in power-system planning applies: reliability indices are not ends in themselves; they support decisions. If an additional 5 MW of provisioning increases annual expected unserved compute only slightly, the incremental revenue or deferred capital may justify the risk. If the increase crosses a common-mode or thermal threshold, the same 5 MW may be unjustified. This is a risk-benefit-cost decision, ideally with explicit uncertainty bands rather than point estimates.

VIII. ROLE OF AI AND STATISTICAL METHODS

AI is not the reliability model by itself. Its near-term value is in accelerating model construction, data cleaning, anomaly detection, trace classification, scenario generation, report automation, and optimization. Statistical methods remain central. Time-series forecasting is needed for demand trajectories; probabilistic forecasting provides distributions rather than point estimates; extreme-value methods characterize rare peaks; Bayesian models update uncertainty as evidence accumulates; stochastic processes model failures and repairs; causal inference can distinguish correlation from intervention effects; and uncertainty quantification communicates confidence in the final headroom estimate.

Machine-learning surrogates can also accelerate repeated power-flow, thermal, or dynamic calculations inside Monte Carlo simulation, provided their error is bounded and monitored. For high-consequence operational decisions, verification and validation should remain explicit: training-data provenance, out-of-distribution detection, backtesting, calibration, human review, and fallback rules are part of the engineering system, not optional governance layers.

IX. GRID INTERACTION AND MULTI-TIMESCALE CONSTRAINTS

A data center that increases utilization also changes its relationship with the external grid. Large AI loads can trip, ramp, oscillate, throttle, or migrate. Distributed training and geographic workload rebalancing can correlate load changes across facilities. Therefore, a site-level headroom result should be screened against interconnection limits, voltage/reactive capability, protection, dynamic performance, transmission contingencies, and system restoration requirements. Facility reliability and grid reliability are bidirectional: the grid affects data-center availability, while data-center behavior can affect the grid and other customers.

This motivates a hierarchy of constraints. Millisecond-to-second limits address converter stability, protection, ride-through, and frequency/voltage response. Seconds-to-minutes limits address workload ramps, UPS/BESS response, generator starting, and cooling transients. Hours-to-years address outage adequacy, maintenance, fuel, workload growth, and investment. A credible utilization decision must satisfy all relevant layers rather than optimizing annual adequacy in isolation.

X. DATA REQUIREMENTS FOR A FIELD STUDY

A field implementation should replace the illustrative parameters with at least twelve months of synchronized operating data where possible. Electrical measurements should include MW, Mvar, voltage, current, breaker/switch status, UPS loading, battery SOC, generator state, and event logs. Infrastructure data should include equipment topology, ratings, maintenance schedules, failure/repair histories, and protection/transfer logic. Thermal data should include cooling-plant loading, temperatures, flows, and degraded states. Compute data should include rack/cluster utilization, workload type, scheduler actions, throttling, migration, and service class. Grid data should include outage/disturbance histories, interconnection limits, and relevant contingency conditions.

The analysis should then proceed in stages: (1) empirical workload characterization; (2) subsystem reliability model; (3) dependency/common-mode model; (4) chronological Monte Carlo and calibration; (5) risk-utilization curve; (6) sensitivity and dominant-risk analysis; (7) mitigation and optimization; and (8) backtesting against actual events. The result should include confidence intervals and a clear separation between observed data, inferred parameters, and engineering assumptions.

XI. LIMITATIONS AND RESEARCH AGENDA

The numerical results in this draft are illustrative. The workload distribution is a synthetic surrogate calibrated qualitatively to public AI workload evidence, and the component unavailabilities are not fleet statistics. The state-sampling model does not preserve temporal autocorrelation, outage duration, battery energy, generator start delay, topology-specific switching, thermal transients, or electrical dynamic stability. Consequently, the reported LOCE and EUCE values should be interpreted only as demonstrations of the framework.

The next research steps are to replace the surrogate workload with public NLR and Google traces, introduce chronological failure/repair sequences, model battery and generator transitions, represent tiered electrical topologies explicitly, and add correlated workload and common-cause models. A further extension should translate unserved IT MWh into workload-level consequences such as delayed training, lost inference service, GPU-hours, or economic value. These steps would allow direct comparison of capital reinforcement, operational flexibility, and accepted risk.

XII. CONCLUSION

The central engineering question for power-constrained AI infrastructure is not simply how much equipment is installed, but how much useful compute can be supported at an explicitly accepted level of risk. Public evidence shows that workload diversity and power controls can create meaningful statistical headroom, while reliability studies show that redundancy, support-system failures, and common-mode dependencies can remove it. Bringing the two perspectives together yields a reliability-constrained utilization problem.

The proposed framework extends familiar power-system reliability concepts to the full data-center system. It quantifies capacity-deficiency probability, expected unserved compute energy, and loss-of-compute expectation as functions of provisioning; incorporates independent and dependent failures; and identifies a nonlinear risk-utilization frontier. With site data, the method can support decisions on additional provisioning, redundancy, maintenance, batteries, generators, cooling, workload flexibility, and capital investment. The intended outcome is not maximum loading at any cost, but the best balance of utilization, reliability, risk, and economics.

REFERENCES

[1] L. A. Barroso, U. Hölzle, and P. Ranganathan, The Data Center as a Computer: Designing Warehouse-Scale Machines, 4th ed. Cham, Switzerland: Springer, 2022.

[2] S. Shah, V. Gevorgian, P. Koralewicz, R. Wallen, B. Mather, B. Kroposki, and J. Kurtz, “Data Center Power Systems: Architectures, Impact on Grid Reliability, Modeling Considerations, and Megawatt-Scale Hardware Testing,” National Laboratory of the Rockies, Mar. 2026.

[3] V. Sakalkar et al., “Data Center Power Oversubscription with a Medium Voltage Power Plane and Priority-Aware Capping,” in Proc. ASPLOS, 2020, pp. 497–511, doi: 10.1145/3373376.3378533.

[4] Meta Engineering, “Lights Out, Systems On: Validating Instant Power Loss Readiness,” Jun. 3, 2026. [Online]. Available: engineering.fb.com/2026/06/03/data-center-engineering/lights-out-systems-on-validating-instant-power-loss-readiness/

[5] B. R. Shrestha, U. Tamrakar, T. M. Hansen, B. P. Bhattarai, S. James, and R. Tonkoski, “Efficiency and Reliability Analyses of AC and 380 V DC Distribution in Data Centers,” IEEE Access, vol. 6, 2018, doi: 10.1109/ACCESS.2018.2877354.

[6] Y. Yu, K. Shan, H. Li, and S. Wang, “Quantitative impacts of energy system flexibility provision and design/operation on data center reliability and the key influential parameters,” Applied Energy, vol. 415, Art. no. 127904, 2026, doi: 10.1016/j.apenergy.2026.127904.

[7] Google, “PowerData 2019 traces,” Google Cluster Data, released 2024. [Online]. Available: github.com/google/cluster-data/blob/master/PowerData2019.md

[8] R. Vercellino et al., “Measurement of Generative AI Workload Power Profiles for Whole-Facility Data Center Infrastructure Planning,” arXiv:2604.07345, Apr. 2026.

[9] J. Sun, M. Xu, M. Cespedes, and M. Kauffman, “Data Center Power System Stability—Part I: Power Supply Impedance Modeling,” CSEE Journal of Power and Energy Systems, vol. 8, no. 2, pp. 403–419, Mar. 2022.

[10] J. Sun, M. Mihret, M. Cespedes, D. Wong, and M. Kauffman, “Data Center Power System Stability—Part II: System Modeling and Analysis,” CSEE Journal of Power and Energy Systems, vol. 8, no. 2, pp. 420–432, Mar. 2022.

[11] E. Choukse et al., “Power Stabilization for AI Training Datacenters,” arXiv:2508.14318v2, Aug. 2025.

[12] Y. Yu, K. Shan, H. Li, and S. Wang, “Reliability and economic impacts of utilizing battery energy storage in data centers for energy flexibility services in smart grids,” Energy Conversion and Management, vol. 339, Art. no. 119951, 2025, doi: 10.1016/j.enconman.2025.119951.

[13] Uptime Institute, “Tier Standard: Topology,” data-center infrastructure tier definitions. [Online]. Available: uptimeinstitute.com/tiers.

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