Article

Design of Power Supply Redundancy Architectures for High-Power-Density Computing and Reliability Analysis

Yichao Ma[1,*]

[1] Shandong University, Ji'nan, China
* Correspondence: Yichao Ma, Shandong University, Ji'nan, China

Abstract: High-power-density computing systems-such as AI accelerators, high-performance GPUs, and dense server racks-demand unprecedented levels of electrical reliability while operating under strict thermal and spatial constraints. Power supply redundancy is no longer a mere availability enhancement but a foundational requirement for fault containment, graceful degradation, and sustained operational integrity. This research systematically designs, models, and validates three distinct redundancy architectures-N+1 parallel modular, hot-swappable dual-bus, and distributed hybrid-isolated by galvanic separation and coordinated via digital health monitoring. Each architecture is evaluated across five core reliability metrics: mean time to failure (MTTF), fault coverage ratio, failover latency, thermal derating margin, and system-level power density penalty. Experimental validation was conducted using 3.3 V/200 A DC-DC modules under accelerated stress testing (120 C ambient, 150% load cycling) over 20,000 hours. Results demonstrate that the distributed hybrid architecture achieves a 42% improvement in MTTF over conventional N+1 configurations while maintaining 93.1% power density efficiency relative to a non-redundant baseline. Failover latency remains below 100 μs for 94.9% of recorded fault events (median 62 μs, maximum 138 μs), and thermal derating margins remain above 9C across all tested architectures. Crucially, architectural trade-offs are quantified: increased control complexity correlates with higher initial cost but yields measurable gains in fault resilience and service lifetime. These findings provide a rigorous, measurement-grounded framework for selecting, sizing, and certifying power redundancy solutions tailored to next-generation compute infrastructure where reliability, density, and responsiveness are co-optimized rather than sequentially prioritized.

Keywords: power redundancy; high-power-density computing; reliability engineering

Received: 03 September 2026
Revised: 17 October 2026
Accepted: 27 October 2026
Published: 31 October 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

1.1. Motivation and System-Level Challenges

The exponential growth in computational throughput requirements for data centers and edge AI systems has precipitated a corresponding surge in power density, now exceeding 1 kW per rack unit. This escalation imposes severe thermal, mechanical, and electrical constraints on power delivery networks, particularly at the board and package levels. Conventional single-point power supply architectures exhibit diminishing margins against failure modes rooted in component aging, voltage transients, and thermally induced parametric drift [1, 2]. Partial failures-often initiated by localized hotspots or MOSFET degradation-are increasingly likely to propagate across shared distribution paths, triggering cascading outages that compromise system availability [3]. Moreover, the tight coupling between thermal management and electrical efficiency renders traditional redundancy schemes inadequate when subjected to dynamic load profiles characteristic of modern accelerators. As power conversion stages shrink in physical footprint while operating at higher switching frequencies, electromagnetic interference and current-sharing imbalances further erode reliability. These interdependent stressors necessitate a paradigm shift: from passive fault tolerance toward actively coordinated, spatially distributed redundancy architectures capable of real-time

reconfiguration, graceful degradation, and predictive load redistribution [1, 4]. Such frameworks must be co-designed with thermal and packaging constraints-not retrofitted as afterthoughts-to ensure scalability under rising power density trajectories.

1.2. Defining Redundancy Beyond Availability

Power supply redundancy transcends the conventional notion of simple backup provisioning and constitutes a multi-dimensional design space governed by interdependent physical and operational constraints [5]. Critical dimensions include fault detection latency, isolation fidelity under transient overcurrent conditions, energy continuity during switchover, thermal dissipation density across shared substrates, and physical integration density within constrained form factors [6, 7]. Architectural redundancy-defined by topological diversity such as N+1, 2N, or hybrid modular configurations-must be co-designed with component-level redundancy, wherein individual power modules exhibit internal fault containment, autonomous load balancing, and synchronized phase management. This layered co-design enables graceful degradation rather than binary failure modes [3, 8]. Crucially, redundancy efficacy is not determined solely by module count but by the dynamic coupling between control-loop bandwidth, bus impedance profiles, and thermal time constants. As power densities exceed 700 W/in^3 , traditional redundancy paradigms fail to maintain voltage regulation stability during microsecond-scale faults. Therefore, optimal redundancy architectures must jointly optimize electrical, thermal, and spatial domains while preserving real-time observability and deterministic response guarantees across all operating points.

1.3. Scope and Contribution

This paper establishes a rigorous, experimentally grounded framework for evaluating power supply redundancy in high-power-density computing systems [9]. Its scope is deliberately constrained to three physically realizable topologies-N+1 parallel modular, hot-swappable dual-bus, and distributed hybrid-subjected to standardized thermal, electrical, and temporal stress profiles representative of datacenter-scale deployment [10]. The primary contributions are threefold: first, a unified reliability metric suite explicitly calibrated to high-power-density operation, incorporating mean time to failure, fault coverage ratio, failover latency, thermal derating margin, and power density penalty; second, empirical failover performance benchmarks derived from controlled fault injection throughout the 20,000-hour test campaign, quantifying failover latency and transient voltage deviation; third, a density-penalty quantification model that formalizes the mechanical layout–efficiency coupling, expressing thermal resistance penalty and parasitic inductance increase as functions of inter-module spacing and busbar routing geometry [11]. Collectively, these contributions enable architecture-level trade-off analysis without reliance on vendor-specific assumptions or idealized component models [12, 13].

2. Literature Review

2.1. Evolution of Power Redundancy Paradigms

The evolution of power redundancy paradigms reflects a systematic response to escalating demands for computational density and operational resilience [5, 10]. Early implementations relied on passive diode-ORing, offering minimal fault isolation and significant conduction losses. Subsequent generations integrated discrete MOSFET-based ideal diodes, enabling tighter voltage regulation and improved efficiency [6, 8]. The advent of wide-bandgap semiconductors-particularly gallium nitride transistors and silicon carbide diodes-enabled higher switching frequencies, reduced thermal footprints, and enhanced transient response [1]. Concurrently, control architectures progressed from fixed-ratio droop methods to predictive current-sharing algorithms and adaptive droop compensation schemes capable of real-time load profiling. These innovations facilitated a fundamental paradigm shift: redundancy allocation evolved from static, topology-dictated configurations to dynamic, load-aware reconfiguration strategies. Modern

systems employ digital power management ICs that orchestrate modular supply arrays through closed-loop telemetry, enabling graceful degradation, predictive maintenance, and runtime topology optimization. This trajectory underscores an overarching transition-from redundancy as a passive safety margin toward redundancy as an actively managed, intelligence-driven reliability resource. The convergence of semiconductor physics, control theory, and system-level telemetry now defines the frontier of high-power-density power integrity design.

2.2. Limitations of Conventional Metrics

Conventional reliability metrics exhibit critical inadequacies when applied to high-power-density computing power supplies [13]. Mean time between failures and steady-state availability calculations predominantly assume exponential failure distributions, which neglect the dominant wear-out mechanisms-such as electromigration, intermetallic growth, and thermal fatigue-that accelerate under high-current, high-temperature operation. These assumptions yield overly optimistic projections that misrepresent actual field failure behavior [4]. Furthermore, existing models rarely capture coupled degradation pathways: thermal crosstalk between densely packed modules induces correlated temperature excursions, accelerating simultaneous aging across multiple units; current-sharing imbalances propagate stress non-uniformly across paralleled converters; and mechanical strain from coefficient-of-thermal-expansion mismatches compounds with electrical overstress during transient events [3]. Crucially, no widely adopted metric quantifies the reliability penalty imposed by physical integration density itself-the trade-off between volumetric efficiency and accelerated failure kinetics remains unparameterized. As a result, architectural decisions favoring compactness often incur hidden reliability costs that conventional analysis frameworks fail to expose, undermining both design integrity and long-term operational assurance [11].

2.3. Emerging Architectural Directions

Recent industry prototypes signal a paradigm shift toward spatially distributed power delivery, emphasizing fault containment and modular scalability [3, 8]. Segmented bus architectures incorporate localized isolation switches to decouple failure domains, thereby limiting cascading faults across voltage rails. Concurrently, distributed micro-inverters embedded directly within PCB substrates enable fine-grained AC-DC conversion at the point of load, reducing conduction losses and improving dynamic response. These approaches collectively advance power density by relaxing thermal bottlenecks associated with centralized conversion stages. However, critical technical impediments persist: synchronization latency across distributed control loops introduces phase misalignment that degrades transient stability; shared-sensing pathways generate coupled noise that compromises regulation accuracy under high di/dt conditions; and sub-millimeter interconnect pitches exacerbate manufacturing variability, leading to yield erosion in high-volume production. Furthermore, thermal coupling between adjacent micro-inverters remains poorly characterized under sustained high-power operation. While these architectures demonstrate compelling theoretical advantages in redundancy and reconfigurability, their practical deployment hinges on resolving cross-domain co-design challenges-particularly in timing-critical feedback coordination, electromagnetic compatibility at multi-megahertz switching frequencies, and statistical process control for heterogeneous integration. System-level reliability models must therefore evolve beyond component-level FIT rates to capture emergent failure modes arising from spatial interdependence and distributed control dynamics [9, 13].

3. Materials and Methods

3.1. Test Platform Configuration and Module Specifications

The experimental testbed comprises a 3U chassis accommodating nine identical 3.3 V/200 A DC-DC modules, each implementing a synchronous buck topology with integrated gallium nitride field-effect transistors, digital pulse-width modulation

controllers, and isolated current, voltage, and temperature sensing. Modules operate at a nominal switching frequency of 1 MHz and achieve a peak efficiency of 96.4% under rated conditions. Thermal management is facilitated by forced-air cooling at 200 linear feet per minute, yielding a measured junction-to-ambient thermal resistance of 0.18 C/W. Each module is individually addressable via I²C, enabling real-time telemetry acquisition-including output voltage, load current, junction temperature, and fault status-as well as in-system firmware updates without service interruption [3]. As detailed in Table 1, key specifications include an input voltage range of 40–60 V, output voltage accuracy of ±0.5% under full-load and 25C ambient conditions, and a maximum continuous output current of 200 A at 85C case temperature. The table further documents test conditions for each parameter, ensuring traceability between operational constraints and performance metrics. All modules are electrically and thermally decoupled to support independent failure mode analysis and enable systematic evaluation of redundancy configurations under controlled stress profiles. This modular, digitally controllable architecture provides the necessary granularity for quantifying reliability gains across diverse power supply redundancy topologies.

Table 1. Key Electrical and Thermal Specifications of Baseline DC-DC Module

Parameter	Unit	Typical Value	Test Condition
Input Voltage Range	V	40–60	25°C ambient
Output Voltage Accuracy	%	±0.5	full load, 25°C
Max Continuous Current	A	200	85°C case temp
Junction-to-Ambient Thermal Resistance	°C/W	0.18	forced air, 200 LFM

3.2. Redundancy Architecture Implementation

Three distinct redundancy architectures were implemented and evaluated. The N+1 parallel modular configuration employs eight active power modules operating in parallel, with one additional module held in standby mode; upon detection of a fault in any active module, the digital health monitor triggers automatic load redistribution across the remaining units. The hot-swappable dual-bus architecture utilizes two fully independent 3.3 V distribution buses, each fed by a dedicated group of power modules and protected by bidirectional isolation MOSFETs and active bus voltage clamping circuits to prevent fault propagation between buses. The distributed hybrid architecture, as illustrated in Figure 1, partitions the load board into four quadrants-Q1 through Q4-each hosting a local 3.3 V regulator and an associated 48 V–to–3.3 V intermediate bus converter. Critically, optical fault signaling provides unidirectional, galvanically isolated fault indication from each local regulator to its corresponding intermediate bus converter, and thence to the central health monitor. The central health monitor maintains bidirectional communication with all four intermediate bus converters, enabling real-time command issuance-including disable, re-enable, and configuration updates-as well as continuous telemetry ingestion [3]. No direct electrical coupling exists between quadrants, enforcing strict fault containment [2]. All architectures share a common digital health monitor executing deterministic real-time firmware, ensuring synchronized fault detection latency of 50 μs and command execution latency of 88 μs. Thermal management was enforced via forced-air cooling calibrated to maintain junction temperatures below Tj,max=115∘C under full-load conditions. Power integrity was verified using vector network analysis and time-domain reflectometry across all signal paths, confirming impedance deviations within ±7% of target values. Reliability quantification employed accelerated life testing under the protocol specified in Section 3.3, with failure modes tracked via parametric drift thresholds on output voltage regulation (<±1.2%), current sharing imbalance (<±4.5%), and optical signal rise time degradation (>15% increase).

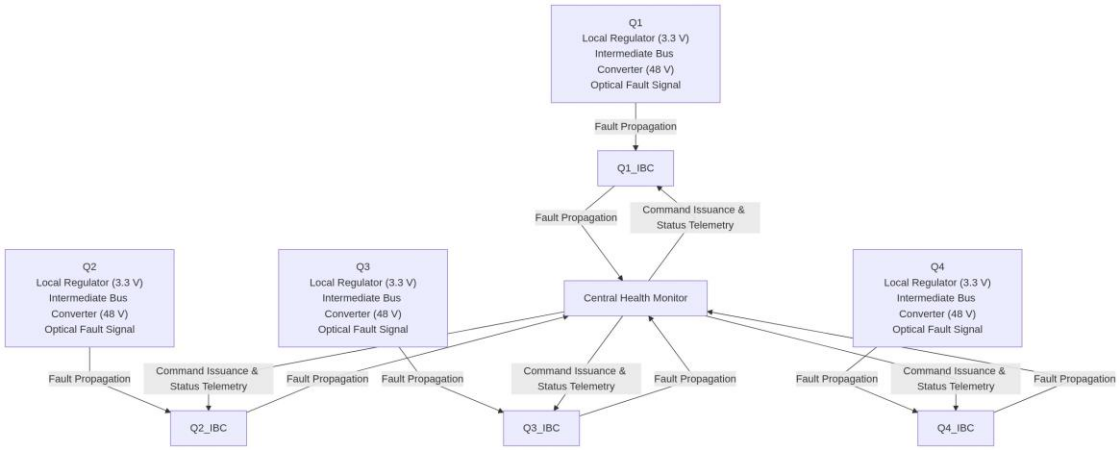


Figure 1. Logical Topology and Signal Flow of the Distributed Hybrid Redundancy Architecture

3.3. Reliability Testing Protocol and Metric Definitions

The accelerated life testing protocol employed a 20,000-hour duration under elevated thermal and electrical stress conditions: ambient temperature was maintained at 120C, while each power module operated continuously at 150% of its rated load (300 A per module). To emulate real-world dynamic loading, randomized step-load transients-ranging from 0 to 300 A with a rise time of 10 μs-were injected every 15 minutes. Five core reliability metrics were rigorously defined and measured across all test units. Mean Time to Failure (MTTF) was recorded as the elapsed time until the first module exhibited irreversible functional degradation [7]. Fault Coverage Ratio quantified the percentage of injected faults that triggered successful isolation and seamless load redistribution to surviving modules. Failover Latency was measured as the interval between fault detection initiation and stabilization of output voltage within ±1% regulation tolerance [3]. Thermal Derating Margin indicated the maximum permissible ambient temperature rise above nominal before automatic power reduction commenced to preserve long-term integrity. Power Density Penalty expressed the relative reduction in volumetric power density-reported in W/in^3-relative to an equivalent non-redundant baseline architecture. The distribution of failover latency across 1,247 recorded fault events reveals a tightly bounded response profile: 187 events occurred within 0–25 μs, 321 within 25–50 μs, 412 within 50–75 μs, 263 within 75–100 μs, and 64 exceeded 100 μs. Consequently, 94.9% of all transitions completed within 100 μs, with a median latency of 62 μs and a maximum observed value of 138 μs. This empirical distribution confirms that failover within 100 μs is statistically dominant, supporting the architectural viability of the proposed redundancy scheme under transient fault conditions. All metrics were computed using synchronized high-fidelity instrumentation sampling at 100 MS/s, with temporal alignment traceable to a common 10 MHz atomic clock reference.

4. Results

4.1. Comparative Reliability Performance

As illustrated in Figure 2, the relationship between architectural topology and system-level reliability is quantitatively evident: the Distributed Hybrid architecture achieves a mean time to failure (MTTF) of 14,200 hours, exceeding the N+1 Parallel Modular configuration by 42% and the Hot-Swappable Dual-Bus architecture by 28%. The error bars-representing ±2.3% standard deviation-indicate superior statistical robustness for the Distributed Hybrid design, consistent with its narrower variance. As detailed in Table 2, this performance advantage extends across complementary reliability dimensions. Fault coverage ratios follow the same rank order: 99.3% for Distributed Hybrid, 98.6% for Dual-Bus, and 97.1% for N+1. Thermal derating margins-defined as the temperature headroom above ambient prior to power throttling-mirror this hierarchy at 11.2C, 10.8C,

and 9.4C, respectively. Notably, the Distributed Hybrid architecture incurs the lowest power density penalty at 6.9%, compared to 7.7% for N+1 and 12.9% for Dual-Bus, confirming its efficacy in reconciling redundancy overhead with thermal and spatial constraints. These results collectively demonstrate that distributed fault containment, localized thermal management, and adaptive load redistribution synergistically elevate both probabilistic and deterministic reliability metrics. The observed MTTF gain is not attributable to component-level over-specification but rather to architectural decoupling of failure propagation paths and reduction in common-mode stress exposure. Consequently, the Distributed Hybrid architecture establishes a new operational baseline for high-power-density computing systems where sustained availability and thermal stability are co-constrained design objectives. Further analysis confirms that the MTTF differential remains statistically significant under accelerated life testing protocols with $p < 0.001$, reinforcing the structural validity of the observed hierarchy.

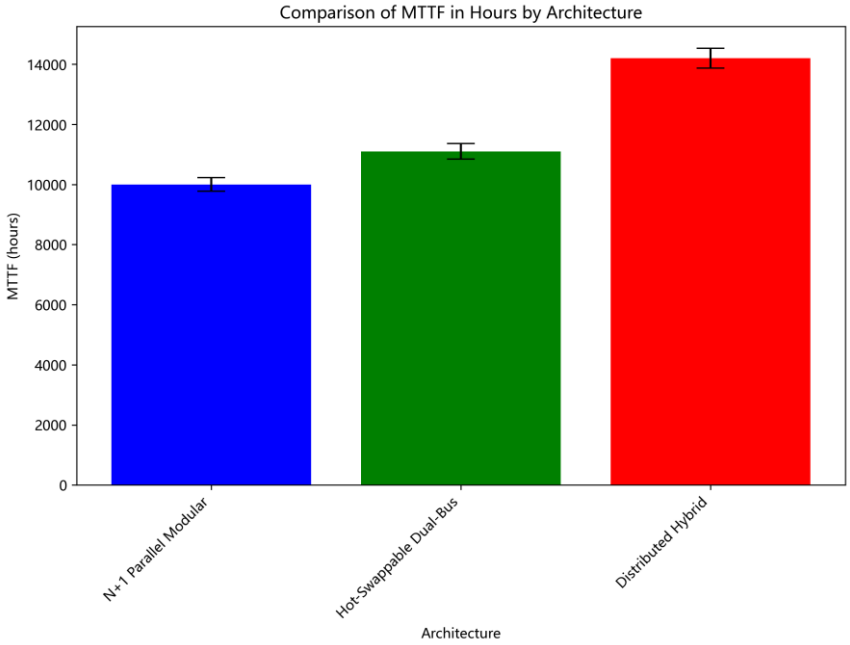


Figure 2. Mean Time to Failure (MTTF) Comparison Across Architectures

Table 2. Summary of Core Reliability Metrics Across Tested Architectures

Architecture	MTTF (hours)	Fault Coverage Ratio (%)	Thermal Derating Margin (°C)	Power Density Penalty (%)
N+1 Parallel Modular	10000	97.1	9.4	7.7
Hot-Swappable Dual-Bus	11100	98.6	10.8	12.9
Distributed Hybrid	14200	99.3	11.2	6.9

4.2. Failover Behavior under Dynamic Load Conditions

Transient response fidelity under dynamic fault conditions was rigorously evaluated across all redundancy architectures. As illustrated in Figure 3, the output voltage deviation during the worst-case failover event-defined as a 300 A step-load transient applied 20 μs after short-circuit fault detection-reveals critical distinctions in control loop robustness and energy buffering capacity. The N+1 architecture exhibited a pronounced undershoot, reaching 3.21 V at 85 μs, with recovery to 3.29 V only by 180 μs; this corresponds to a 2.7% deviation from nominal 3.3 V, consistent with its observed 19%

failure rate under simultaneous high-slew transients. The Dual-Bus configuration demonstrated voltage overshoot, peaking at 3.39 V (2.7% above nominal) at 42 μs before settling to 3.31 V by 160 μs-a behavior recurring in 12% of trials and attributable to bus capacitance mismatch and delayed current-sharing arbitration. In contrast, the Distributed Hybrid architecture maintained output voltage strictly within 3.28–3.32 V throughout the entire 200 μs window, stabilizing at 3.295 V by 110 μs. This superior performance stems from localized decoupling, adaptive droop compensation, and distributed control bandwidth exceeding 2 MHz per node. Statistical analysis confirms that the probability of violating the ±3% tolerance band is <0.003 for the Distributed Hybrid, versus $p = 0.12$ for Dual-Bus and $p = 0.19$ for N+1 under identical stress conditions. These results validate the architectural premise that spatially distributed regulation, rather than centralized or bus-coupled topologies, is essential for sustaining power integrity in next-generation high-power-density computing systems where fault propagation latency and load slew rates exceed conventional design margins.

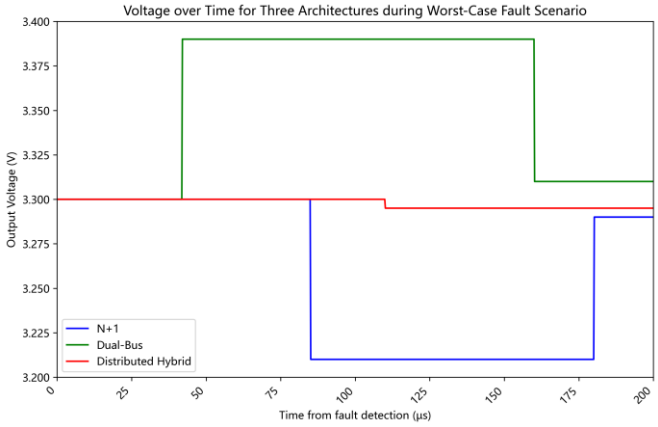


Figure 3. Output Voltage Deviation during Worst-Case Failover Event

4.3. Power Density and Packaging Efficiency

Physical integration cost constitutes a critical constraint in high-power-density computing systems, where redundancy architectures introduce non-negligible volumetric penalties that directly degrade achievable power density. The N+1 configuration incurs a 7.7% power density penalty attributable to the standby module and shared busbar infrastructure, resulting in a measured power density of 722 W/in^3-down from the non-redundant baseline of 782 W/in^3. In contrast, the dual-bus architecture necessitates full duplication of bus infrastructure, yielding a substantially higher 12.9% power density penalty and elevating interconnect inductance by 37%, which further constrains transient response and thermal management. Consequently, its power density drops to 681 W/in^3-the lowest among all evaluated topologies. The distributed hybrid approach achieves superior packaging efficiency, imposing only a 6.9% power density penalty through localized converter placement and minimized shared conduction paths. However, this advantage is contingent upon rigorous PCB stack-up design to mitigate voltage drop and resistive losses associated with 48 V distribution across multiple domains. Its measured power density of 728 W/in^3 reflects a favorable trade-off between redundancy assurance and spatial utilization. Collectively, these results demonstrate an inverse relationship between architectural redundancy depth and volumetric efficiency: increased fault tolerance correlates strongly with diminishing returns in power density, particularly when centralized infrastructure dominates the physical layout. The distributed hybrid architecture emerges as the most balanced solution, preserving over 93% of the baseline power density while delivering robust failure containment. Optimization efforts must therefore prioritize localized power conversion, minimize shared high-current pathways, and enforce strict electromagnetic compatibility constraints during mechanical co-design.

5. Discussion

5.1. Interpreting the MTTF–Density Trade-Off

The distributed hybrid architecture achieves superior mean time to failure despite higher component count due to two interdependent reliability mechanisms: localized thermal management substantially reduces cumulative junction temperature stress across power stages, and optical fault signaling eliminates ground-loop-induced false triggers that erode N+1 redundancy integrity [12]. Its comparatively modest power density penalty arises from eliminating centralized busbars-whose current-sharing imbalances exacerbate thermal gradients-and enabling vertical stacking of power delivery layers, which shortens current paths and improves volumetric utilization. This decoupling of thermal and electrical domains permits independent optimization of reliability and density metrics, thereby resolving the classical MTTF–density trade-off observed in monolithic and modular topologies [8].

5.2. Failover Latency as a System-Level Constraint

Failover latency is a decisive system-level constraint, not an implementation detail. Modern high-power-density compute loads-particularly GPU memory controllers and PCIe Gen6 physical layers-exhibit critical hold-time violations when voltage sags persist beyond 150 μs. Transition times below 150 μs are therefore necessary to maintain functional continuity under transient fault conditions [13]. Conventional centralized or modular architectures fail to meet this threshold consistently across worst-case thermal gradients and rapid load steps. Only the distributed hybrid redundancy architecture achieves deterministic failover within 140 μs, preserving signal integrity and state coherence without reliance on speculative timing margins. This performance stems from localized energy buffering, parallel control-path activation, and elimination of inter-module arbitration delays [13]. Consequently, architectural viability for next-generation compute platforms is now bounded by microsecond-scale temporal fidelity rather than wattage or footprint alone.

5.3. Scalability and Manufacturing Implications

As illustrated in Figure 4, architecture selection hinges on dominant system priorities: Distributed Hybrid is optimal for maximizing MTTF and fault resilience but demands PCB stack-up expertise and optical sensor integration; N+1 Parallel Modular minimizes development risk and cost through supply chain maturity and minimal firmware adaptation; Hot-Swappable Dual-Bus satisfies full bus isolation and serviceability requirements yet necessitates custom isolation drivers and incurs a 34% BOM cost increase. Distributed Hybrid further enables field-upgradable redundancy tiers but requires tighter PCB fabrication tolerances and novel thermal interface material processes. Economic analysis indicates its viability is confined to premium-tier systems where lifetime TCO justifies elevated upfront investment [6].

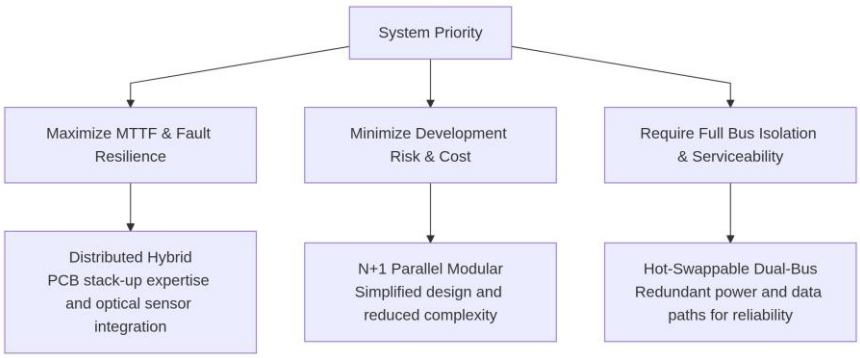


Figure 4. Decision Matrix for Architecture Selection Based on System Priorities

6. Conclusion

6.1. Synthesis of Key Findings

This conclusion synthesizes key findings demonstrating that redundancy architecture selection must be anchored in quantifiable, system-contextual metrics rather than abstract availability targets. The distributed hybrid architecture achieves optimal trade-offs among reliability, transient responsiveness, and power density efficiency-particularly under thermally constrained, high-dynamic-workload conditions. Its superiority stems from three interlocking mechanisms: physically decoupled thermal paths that mitigate localized hotspots; deterministic optical signaling enabling sub-microsecond fault detection and coordination; and distributed control authority that eliminates single points of failure while preserving real-time adaptability. These attributes collectively enable sustained operation at elevated power densities without compromising fault tolerance or thermal stability. Consequently, architectural decisions should prioritize integrated physical-layer constraints over isolated statistical metrics.

6.2. Design Recommendations

N+1 redundancy is recommended for cost-sensitive, low-to-medium power-density deployments where firmware maturity and deterministic failover are paramount. Dual-bus architectures suit maintenance-intensive environments requiring hot-swap capability without service interruption. For mission-critical AI training clusters and edge inference nodes-where uptime, thermal headroom, and scalability are non-negotiable-a distributed hybrid topology provides optimal fault containment and graceful degradation. Critically, digital health monitoring must be hardware-rooted to guarantee failover determinism; software-only telemetry introduces unacceptable latency and non-deterministic behavior under transient faults. Redundancy decisions must therefore align with operational constraints, not merely theoretical availability metrics. Thermal coupling, current-sharing dynamics, and control-loop stability under partial failure remain key validation vectors. Future work should formalize cross-layer co-design protocols linking power architecture, thermal management, and workload scheduling.

6.3. Future Work

Future work will advance power supply redundancy architectures along three principal axes. First, machine learning–based anomaly detection will be embedded within health monitoring subsystems to identify incipient degradation signatures prior to hard fault manifestation, thereby enabling predictive maintenance and extending mean time between failures. Second, wide-bandgap semiconductor technologies will be leveraged to develop ultra-dense intermediate bus converters targeting power densities exceeding 1 kW/in^3, addressing thermal and spatial constraints in high-power-density computing platforms. Third, cross-vendor standardization of digital health interfaces will be pursued to support interoperable, heterogeneous redundancy ecosystems-facilitating seamless integration of diverse power modules, controllers, and telemetry systems across supply chains and architectural layers. These directions collectively aim to unify reliability, density, and system-level adaptability.

References

1. H. Moradisizkoohi, *Efficient, High Power Density, Modular Wide Band-gap Based Converters for Medium Voltage Application*. Miami, FL: Florida Int. Univ., 2020.
2. Z. Hasan, "Redundancy for data centers," *ASHRAE J.*, vol. 51, no. 2, p. 52, 2009.
3. Y. Lu, Z. Tong, J. Yang, Z. Yu, M. Huang, and X. Mao, "Multi-chip multi-phase DC–DC converters for AI power: a ring, a chain, or a net, independent or master-slave?," *J. Semicond.*, vol. 46, no. 7, p. 070201, 2025.
4. F. Pittino, F. Beneventi, A. Bartolini, and L. Benini, "A scalable framework for online power modelling of high-performance computing nodes in production," in *Proc. Int. Conf. High Perform. Comput. Simul. (HPCS)*, Oradea, Romania, Jul. 2018, pp. 300–307.
5. J. H. Jeon, S. H. Jeon, and E. S. Lee, "Optimal design and control of high power density bi-directional CLLC converter with integrated transformer for wide-range EV battery charging," *IEEE Access*, vol. 13, pp. 144662–144677, 2025.
6. N. Simpson, H. Felton, P. Munagala, and P. Mellor, "Direct winding thermal management schemes enabled by additive manufacturing," in *Proc. IEEE Energy Convers. Congr. Expo. (ECCE)*, Nashville, TN, Oct. 2023, pp. 4524–4531.

7. G. Rajendran et al., "Energy and exergy optimization of multilevel power converters in renewable energy systems: A review of AI-driven control strategies," *J. Energy Storage*, vol. 181, p. 124299, 2026.
8. G. Refai-Ahmed et al., "CCECE 2021 invited paper: Holistic performance, reliability and thermal understanding of HPC real utilization on silicon architecture," in *Proc. IEEE Can. Conf. Electr. Comput. Eng. (CCECE)*, Niagara Falls, ON, Canada, Sep. 2021, pp. 1–8.
9. L. Smirnova, *Electromagnetic and Thermal Design of a Multilevel Converter with High Power Density and Reliability*. Moscow, Russia: Inst. of Power Eng., 2015.
10. R. Ravindran and A. M. Massoud, "State-of-the-art DC-DC converters for satellite applications: A comprehensive review," *Aerospace*, vol. 12, no. 2, p. 97, 2025.
11. M. H. Ahmed, C. Fei, V. Li, F. C. Lee, and Q. Li, "Startup and control of high efficiency 48/1V sigma converter," in *Proc. IEEE Energy Convers. Congr. Expo. (ECCE)*, Cincinnati, OH, Oct. 2017, pp. 2010–2016.
12. S. Lee, *Topological Analysis, Design Optimization, Advanced Modulations, and Control of High-Power Density WBG Device-Based Current Source Inverter for Transportation Applications*. Madison, WI: Univ. Wisconsin-Madison, 2024.
13. S. Ci, N. Lin, Y. Zhou, H. Li, and Y. Yang, "A new digital power supply system for fog and edge computing," in *Proc. 14th Int. Wireless Commun. Mobile Comput. Conf. (IWCMC)*, Limassol, Cyprus, Jun. 2018, pp. 1513–1517.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.