Carrier-Grade Infrastructure Solution Brief: Deploying Redundant Power Architecture Design

MARKET POSITIONING

The Redundant Power Architecture Design represents a foundational engineering discipline for carrier-grade network infrastructure, ensuring continuous service availability in the most demanding telecom environments. As service providers migrate toward 5G transport, edge computing, and ultra-reliable low-latency communication (URLLC) use cases, the power subsystem has transitioned from a passive support element to a critical determinant of overall system availability. This solution brief defines the architectural principles, redundancy models, and operational specifications for implementing a carrier-class redundant power architecture across core, aggregation, and edge routing platforms.

Modern network operators face stringent Service Level Agreements (SLAs) that mandate 99.999% availability, translating to less than 5.26 minutes of unplanned downtime per year. The power architecture directly governs the achievability of this target. A properly designed redundant power scheme eliminates single points of failure across AC/DC conversion, power distribution, and load sharing, while providing deterministic failover behavior under transient grid anomalies, brownouts, and component degradation.

The Redundant Power Architecture Design supports both N+1 and 1+1 redundancy topologies, enabling operators to align capital expenditure with service criticality. For core routing nodes, 1+1 active-active redundancy with independent feed paths is recommended; for aggregation and edge platforms, N+1 shared redundancy offers an optimal balance of resilience and efficiency.

HIGH-AVAILABILITY REDUNDANCY

The architecture implements a multi-tier redundancy model spanning power entry, conversion, distribution, and monitoring. Dual independent power feeds (Feed A and Feed B) are routed from separate upstream power distribution units (PDUs), each protected by dedicated circuit breakers. These feeds terminate at hot-swappable power supply units (PSUs) operating in

active-active load-sharing mode, with automatic current balancing to within ± 5% across all installed modules.

In the event of a PSU failure, the remaining units instantly assume full load without service interruption. The failover transition is completed within 10 milliseconds, well below the hold-up time provided by internal bulk capacitors, ensuring zero packet loss during power events. Each PSU is equipped with independent over-voltage, under-voltage, over-current, and thermal protection circuits. Front-panel LED indicators and SNMP traps provide real-time status per PSU, enabling proactive replacement before redundant capacity is exhausted.

For DC-powered deployments, the architecture supports -48V DC nominal input with a wide operating range of -40V to -72V, compliant with ETSI EN 300 132-2 and Telcordia GR-1089-CORE. AC-powered variants support 100-240V AC, 50/60 Hz, with power factor correction (PFC) exceeding 0.98. The power distribution board (PDB) provides per-slot current monitoring and electronic fuse (eFuse) protection, isolating faults to the affected line card without propagating to the backplane.

PROTOCOL INTEROPERABILITY

The redundant power subsystem integrates with industry-standard

management protocols including SNMPv3, NETCONF/YANG, and Redfish. Power telemetry — including input voltage, input current, output voltage, output current, and PSU temperature—is exposed via standard MIBs and YANG models, enabling seamless integration with existing Network Management Systems (NMS) and Data Center Infrastructure Management (DCIM) platforms. The architecture supports PMBus 1.3 for internal PSU communication, allowing granular control over output voltage trimming, fan speed, and fault logging.

DETAILED PARAMETERS

The following table summarizes the key technical parameters of the Redundant Power Architecture Design.

Parameter	**Specification**
Form Factor	1RU / 2RU Chassis
Redundancy Model	1+1 Active-Active / N+1 Shared
Power Input (AC)	100-240V AC, 50/60 Hz, PFC > 0.98
Power Input (DC)	-48V DC nominal, -40V to -72V range
PSU Output Power	1200W per module (AC), 1500W per module (DC)
Load Sharing Accuracy	±5% across all installed PSUs
Failover Time	< 10 ms (zero packet loss)

Hold-Up Time	> 20 ms at full load
Protection	OVP, UVP, OCP, OTP, eFuse per slot
Management	SNMPv3, NETCONF/YANG, Redfish, PMBus 1.3
MTBF (PSU)	> 500,000 hours at 25°C
MTBF (System)	> 1,000,000 hours with redundancy
MTTR	< 2 minutes (hot-swap)
Compliance	ETSI EN 300 132-2, Telcordia GR-1089-CORE

LIFECYCLE ASSURANCE (MTBF)

Reliability engineering is central to the Redundant Power Architecture Design. Each PSU is rated for a Mean Time Between Failures (MTBF) exceeding 500,000 hours at 25 ° C ambient, calculated per Telcordia SR-332 Issue 4. The system-level MTBF, accounting for redundancy and load sharing, exceeds 1,000,000 hours. Power supplies are hot-swappable and support online replacement without service interruption, reducing Mean Time To Repair (MTTR) to under 2 minutes. A comprehensive sparing strategy is recommended, with a minimum of one spare PSU per deployed chassis and a 5-year lifecycle support commitment including firmware maintenance and hardware replacement

availability.

TARGET NETWORK TOPOLOGIES

The Redundant Power Architecture Design is optimized for deployment in core routing PoPs, metro aggregation sites, MEC (Multi-access Edge Computing) facilities, and submarine cable landing stations. It is equally suited for enterprise campus core, data center spine, and colocation environments where power availability is contractually guaranteed. The architecture supports both isolated and common-battery DC plants, as well as dual-corded AC rack configurations.