October 2nd, 2025

Powering the Future: MLCCs and the Evolution of AI-HPC Datacenter Infrastructure

As artificial intelligence (AI) continues to transform industries, the infrastructure supporting high-performance computing (HPC) must evolve to meet the growing demands of machine learning and deep learning workloads. At the heart of this transformation are AI servers. These are complex systems that require highly efficient power distribution networks (PDNs) and robust electronic components. Among these, **Multilayer Ceramic Capacitors (MLCCs)** have emerged as critical enablers of performance, reliability, and scalability in both server boards and datacenter systems.

Power Distribution in AI Datacenters

The PDN architecture in AI datacenters is a multi-stage system designed for scalability and efficiency. Power typically flows from a 480VAC three-phase grid through an uninterruptible power supply (UPS), then into a bank of converters that produce 48VDC. This voltage is distributed via bus bars to individual server blades, where it is stepped down to 12VDC and finally to 1VDC using onboard converters and power management integrated circuits (PMICs).

Each stage of this conversion process presents unique challenges. MLCCs are used extensively to filter noise, manage load transients, and ensure voltage stability. On the 12VDC rail, MLCCs complement aluminum polymer capacitors to buffer energy demands. On the 48VDC rail, hybrid capacitors are often used for bulk filtering, with MLCCs providing additional support.

Emerging trends in PDN design are pushing the boundaries of efficiency and integration. One such trend is the shift from horizontal to vertical power delivery, where converters are placed directly beneath the GPU to reduce parasitics and improve response time. Another is the move toward direct conversion from 48VDC to 1VDC, bypassing intermediate stages to save space and increase efficiency. Though technically challenging, this approach is expected to gain traction by 2027.

Additionally, datacenters are exploring the replacement of 480VAC distribution with ±400VDC. This change reduces heat generation and power loss, enabling higher rack-level power densities and more efficient cooling strategies.

The Role of MLCCs in AI Server Design

Modern AI servers differ significantly from traditional server architectures. While conventional servers relied on a single high-end CPU, AI servers integrate one CPU alongside multiple powerful GPUs, each containing tens of thousands of processing cores. This parallel computing model dramatically increases throughput but also drives power consumption to new extremes. Each GPU can draw over 2000 watts, with entire server cabinets consuming upwards of 100 kilowatts and trending toward 1 megawatt.

October 2nd, 2025

To manage this power efficiently, MLCCs are deployed across multiple voltage rails; 48VDC, 12VDC, and 1VDC—providing essential functions such as decoupling, filtering, and transient response. On the most critical 1VDC rail, MLCCs help stabilize voltage during rapid current swings that can reach 2000 amps in microseconds. Their low equivalent series resistance (ESR) and inductance (ESL), combined with compact form factors, make them ideal for densely packed server boards.

Advancements in MLCC technology have further enhanced their performance. Innovations such as enlarged active areas, thin dielectrics, and homogeneous $BaTiO_3$ grain structures have enabled higher capacitance and improved reliability. Simulation tools now allow engineers to model electric fields and impedance curves with precision, optimizing MLCC placement and performance.

The image below shows the methods in which capacitance can be increased when developing MLCCs. MLCC capacitance is calculated by the active area, the dielectric constant, the layer thickness, and electrode continuity. The first step is to create thinner dielectric layers, as shown in the top left image.

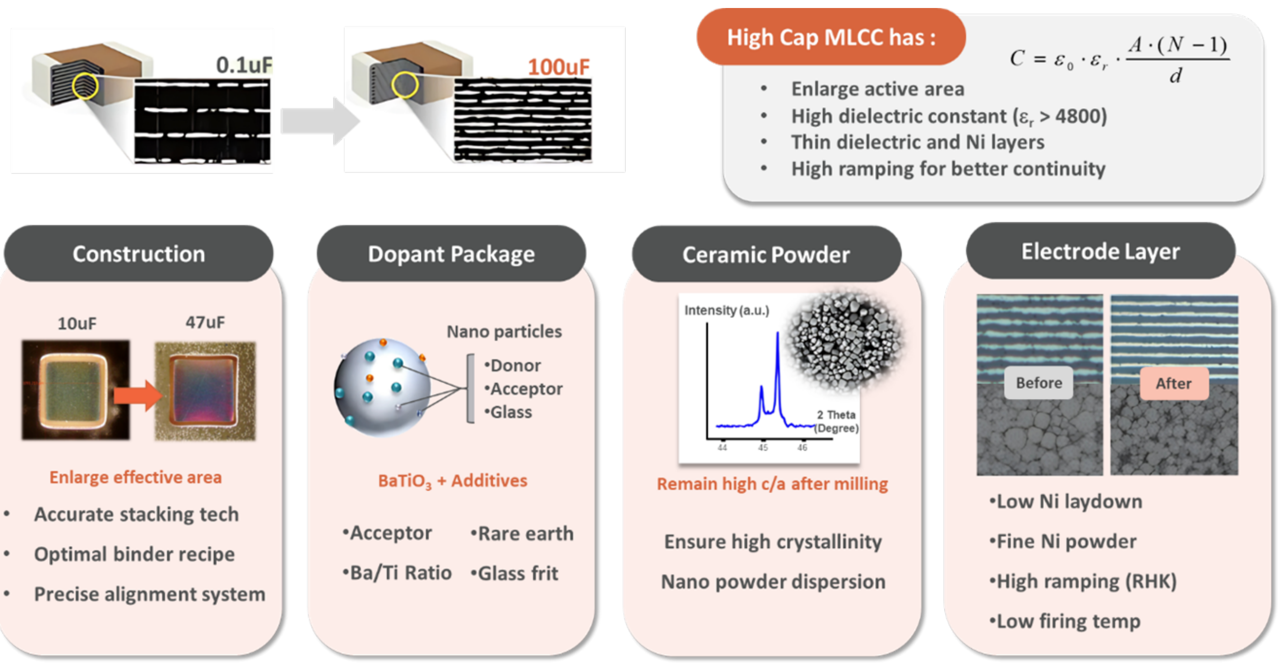

October 2nd, 2025

MLCCs Beyond the Server Board

MLCCs are not limited to server boards—they play a vital role across the entire datacenter ecosystem. In power conversion circuits such as DC-link, snubber, and resonant stages, MLCCs offer superior performance at high switching frequencies, often replacing traditional film and electrolytic capacitors. Advanced MLCC products like KC-Link and KONNEKT technologies provide high ripple current capability and voltage stability, making them ideal for demanding applications.

MLCCs also complement solid and hybrid polymer capacitors in bulk filtering and energy buffering roles. In AC input stages and climate control systems, safety and film capacitors are used alongside MLCCs to ensure reliable operation under varying environmental conditions.

The Infrastructure Behind AI-HPC

A modern AI datacenter comprises several interconnected systems, each contributing to overall performance and reliability:

1. Main Power Distribution: Includes switchgear, transformers, PDUs, and busways to route power efficiently. MLCCs are used in control and monitoring circuits within switchgear and PDUs to suppress electrical noise, stabilize voltage, and ensure reliable signal integrity for power routing and protection systems.
2. Generators: Provide backup power during outages, typically activated via automatic transfer switches. MLCCs are integrated into generator control modules and automatic transfer switches to filter high-frequency noise and support transient voltage suppression during power transitions.
3. Advanced Power Controllers (APC): Manage smart UPS systems, rack-level power distribution, and environmental monitoring. MLCCs play a key role in APCs by providing decoupling and filtering in digital control boards, ensuring stable operation of smart UPS systems and precise environmental monitoring.
4. Network Power Chillers: Remove heat generated by servers using chillers, cooling towers, and heat exchangers. MLCCs are used in the motor drive and control electronics of chillers and heat exchangers to manage switching noise and improve power quality in high-frequency inverter circuits.
5. UPS Batteries: Bridge the gap between power loss and generator startup, protecting sensitive equipment. MLCCs are found in battery management systems (BMS) and inverter circuits, where they help regulate voltage, suppress EMI, and support fast switching during backup power activation.
6. Climate Control: Maintains optimal temperature and humidity using CRAC/CRAH units and aisle containment strategies. MLCCs are found in CRAC/CRAH unit controllers and fan drive electronics to filter power supply noise and ensure stable operation under varying thermal loads and environmental conditions.
7. Servers: Include compute, storage, and networking nodes, supported by MLCCs, VRMs, and heatsinks for stability and thermal management. MLCCs are extensively deployed across server boards on CPU, GPU, and memory power rails for decoupling, transient response, and noise suppression, enabling stable high-performance computing

October 2nd, 2025

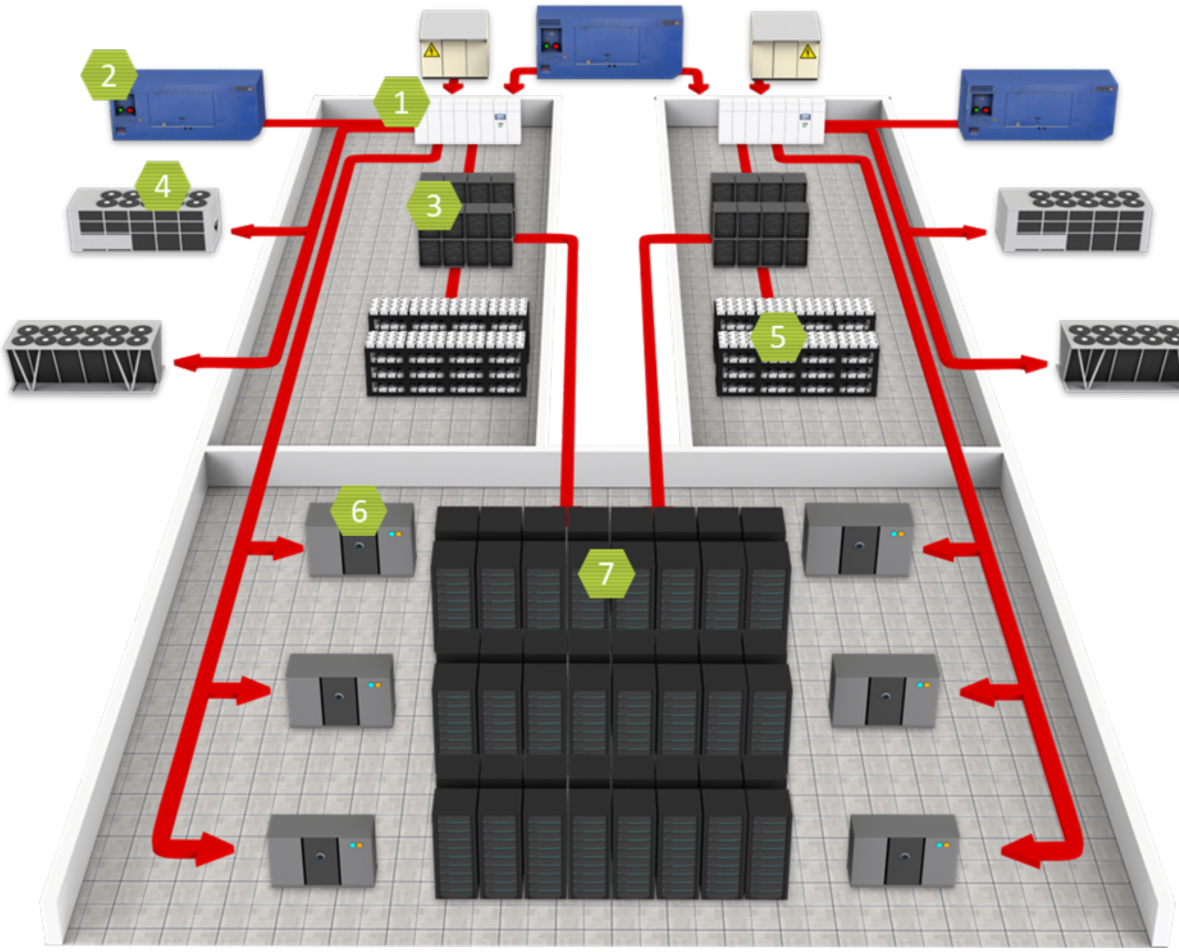

The synergy between MLCCs and AI datacenter infrastructure is central to enabling scalable, efficient, and reliable computing environments. As AI workloads intensify and power architectures evolve, MLCCs, especially high-capacitance variants, will remain indispensable in managing power integrity, thermal performance, and electromagnetic noise suppression.

October 2nd, 2025

Conclusion

As AI and HPC technologies continue to push the boundaries of performance, the supporting infrastructure must evolve in tandem. MLCCs have proven to be foundational components in this transformation—enabling precise power management, enhancing system reliability, and supporting the extreme demands of modern AI servers. From board-level decoupling to rack-level power conversion, their versatility and performance are unmatched. Looking ahead, innovations in capacitor design and PDN architecture will be critical to sustaining the growth of AI datacenters. By embracing these advancements, the industry can build more efficient, scalable, and resilient computing environments that meet the challenges of tomorrow.

Author

Jake Michels - Associate Technical PM, Ceramic PBU, YAGEO Group
Joshua Chen - Technical PM, MLCC PBU, YAGEO Group
Harley Le - PM Engineer, MLCC PBU, YAGEO Group