

Technical White Paper

Hot-Swap Solutions for High-Voltage DC-Powered AI Servers



Rakesh Panguloori
Dilip Jain

Abstract

High-availability systems such as servers, network switches, and telecommunications systems are designed for near-zero downtime throughout their useful life. Such high-availability systems often require the insertion or removal of circuit boards without powering down the entire system.

Being the first point of entry for power at the input of the tray, a hot-swap circuit provides comprehensive telemetry for system diagnostics and efficient load management. The reliability of this circuit impacts operational uptime; thus, the ability to handle system transients is critical to avoid unplanned system outages.

With the exponential growth of artificial intelligence (AI) data center server power, the distribution voltage in an IT rack is moving from 48V_{DC} to 800V_{DC}. This white paper discusses the evolution of the hot-swap architecture to support the transition to the 800V_{DC} architecture, and important considerations around the emerging system-level challenges.

Table of Contents

1 Evolution of Hot-Swap Architectures.....	2
2 Considerations Shaping Next-Generation Hot-Swaps.....	8
3 Driving the Hot-Swap FETs.....	9
4 Safety and Insulation.....	9
5 Conclusion.....	9
6 References.....	9

Trademarks

All trademarks are the property of their respective owners.

1 Evolution of Hot-Swap Architectures

As documented in [1], data center power architectures are transitioning from traditional power shelves inside the IT rack to power sidecars to meet the power demands of AI. Hot-swap technology—the ability to replace or add hardware components without shutting down a system—has evolved from basic, specialized peripheral support into a critical, high-performance feature for modern data centers and computing. The evolution of the technology spans from simple connector standards to advanced power management, driven by the need for maximum uptime and efficiency.

Early Era (Pre-USB)

Initially, hot swapping was limited to high-availability systems such as mainframe computers, where specialized, expensive hardware enabled the removal of modules without shutdown. The small computer system interface introduced hot swapping for peripherals, although it was not as reliable as today's standards.

The USB Revolution

The introduction of USB normalized hot swapping for consumer electronics, allowing the safe addition and removal of keyboards, mice, printers, cameras, and flash drives while systems were active. The evolution moved from simply powering off a slot (passive) to using active MOSFET controllers that safely manage power-on sequences, protecting both the component and the backplane.

Data Center Growth (12V Era)

As data center downtime became too expensive, hot-swapping functionality became essential to achieve high availability. Hot-swap controllers helped manage inrush current and protected the main bus during live insertion of server blades on a 12V backplane.

The first generation of hot-swap controllers, shown in Figure 1-1, used a simple power-limiting architecture to make sure that the MOSFET stayed within the safe operating area (SOA) during start-up and fault recovery. The controller consisted of a current-sense amplifier that monitored the voltage across the sense resistor to obtain current information, and a voltage-sense circuit to measure the voltage across the FET.




Figure 1-1. Power-Limiting Hot-Swap Architecture

Multiplying the output of the current-sense amplifier and voltage sense circuit gives the power loss in the FET, which must be limited to safeguard the hot-swap FET while managing inrush currents or fault currents in the system.

The second generation of hot-swap controllers, shown in Figure 1-1, offered more sophisticated features such as overtemperature protection, telemetry to use for system diagnostics, and data traffic management between trays in a rack server. But as the power level of the rack tray increases, the output capacitance also increases; this makes designing with a power limited-based controller impractical given severe SOA stress during startup, which commercial FETs cannot handle. An alternate approach [2] is to limit inrush current with the output voltage slew

rate (dv/dt) control circuit shown in Figure 1-1, which reduces the power dissipation during start-up and spreads the dissipation over a longer duration such that the FETs can safely handle the power.




Figure 1-2. Hot-swap Solution With Output dv/dt Control for a 12V Rai

To support heavy data traffic, data center operators adopted faster and more power-hungry CPUs and memories to increase data processing capability. This pushed the power requirements as high as 4kW per tray in the same form factor, thus calling for a high-density protection solution with integrated FETs (also called eFuses), as shown in Figure 1-3. TI's 12V high-current eFuses offer differentiated features [3] such as adjustable overcurrent blanking to support transient demand from the downstream computing load, black-box fault recording, active synchronization, and load sharing during start-up and steady state for unlimited scalability—all in a 4.5mm-by-5mm footprint while supporting 60A of root-mean-square (RMS) current.




Figure 1-3. A Stackable High-Current eFuse Solution for a 12V Rail

Cloud Computing (48V Era)

With the growing demands of cloud computing applications, enterprise servers have become power-intensive, as these servers simultaneously process a large amount of data and storage. The steady-state power rating of each server motherboard is now 5kW to 6kW, which makes power distribution on a 12V rail ineffective. Moving from a 12V to 48V server architecture reduces current by 75% for the same power. This reduction drops resistive power losses (I^2R) by a factor of 16, shrinks copper bus-bar and cable requirements by up to 90%, and significantly lowers rack heat dissipation.

The graphics processing units (GPUs) form the predominant load in a 48V server system, where the load amplitude, slew rate and frequency of transient loads are three to four times higher when compared to general

servers. This requires a few circuit enhancements, as discussed in [4], when using TI's first-generation LM5066I hot-swap controller to support high-power 48V designs.

TI's next-generation hot-swap controller, the LM5066H2, offers a dual gate-drive architecture (Figure 1-4) that enables a single SOA FET to handle power stresses during start-up and fault conditions. Combined with multiple low drain-to-source on-resistance ($R_{DS(ON)}$) FETs for normal load-current operation, the LM5066H2 helps reduce total solution size by 40% for >4kW designs.




Figure 1-4. Hot-Swap Solution for 48V Rack Server Power Distribution.

The industry's first 48V eFuse, TI's TPS1685, enables enhanced FET protection and ease of scalability to realize high-density high-power designs. Figure 1-5 illustrates the circuit configuration with 48V eFuses.




Figure 1-5. A Stackable High-Current eFuse Solution for a 48V Rail.

AI Computing ($\pm 400V_{DC}$ or $800V_{DC}$ Era)

The surge in AI processing power has driven a transition from 12V to 48V architectures in server racks to reduce power losses and improve efficiency. AI-based IT racks use hundreds of processors to rapidly process the volume of computations necessary for AI to function best, thus demanding rack power in the range of 100kW to 1MW in the near future [5]. Distributing this amount of power on a 48V rail results in currents exceeding 20kA, requiring heavy and expensive bus bars that are impractical to implement and necessitating a higher $\pm 400V_{DC}$ or $800V_{DC}$ bus voltage for power distribution in AI IT server racks. The low-voltage DC architecture requires

advanced hot-swap controllers to handle higher-voltage inrush currents and provide accurate telemetry, while taking appropriate safety and insulation measures.

Leading hyperscalers and cloud service providers are currently assessing several facility-level architectures such as $\pm 400V_{DC}$ (two wire), $\pm 400V_{DC}$ (three wire), and $800V_{DC}$. These architectures directly influence hot-swap protection requirements.

Figure 1-6, Figure 1-7, Figure 1-8, and Figure 1-9 show the various architectures [6], along with the IT rack-side hot-swap protection sockets.




Figure 1-6. Three-Wire $\pm 400V_{DC}$ Power Distribution System




Figure 1-7. Two-Wire $\pm 400V_{DC}$ Power Distribution System




Figure 1-8. Two-Wire $800V_{DC}$ Power Distribution System




Figure 1-9. Two-Wire 800V_{DC} Power Distribution System With High-Resistance Grounding Protective Earth

2 Considerations Shaping Next-Generation Hot-Swaps

As AI models and scientific simulations grow in complexity, a single processor often is not enough. The IT tray needs to accommodate several high-power CPUs and GPUs. Since this is the highest-value real estate in the rack, the limited available space to accommodate the hot-swap circuit and intermediate bus converter calls for a compact 800V hot-swap solution. A discrete implementation [7] comprises isolated gate drivers, current- and voltage-sense amplifiers, digital isolators, a multiwinding auxiliary supply and a microcontroller.

TI's 800V [LMI800PS](#) hot-swap controller has all of the traditional hot-swap controller elements, along with integrated power and signal isolation to enable a more power dense hot-swap solution.

[Figure 2-1](#) shows the LMI800PS-based hot-swap solution for a $\pm 400V_{DC}$ two-wire system architecture.




Figure 2-1. LMI800PS Hot-Swap Solution

Universal Isolated Hot-Swap Controller Architecture

As shown in [Figure 1-6](#), [Figure 1-7](#), [Figure 1-8](#), and [Figure 1-9](#), the hot-swap circuit can exist both on the high and low side of the input supply rails based on the system architectures and grounding schemes. The LMI800PS controller enables usage on both rails by floating the local-chip GND and providing the necessary signals and bias power across the isolation barrier.

3 Driving the Hot-Swap FETs

For 800V_{DC} systems, the main choice for the power FET is either 1.2kV drain-to-source voltage (V_{DS}) rated silicon carbide (SiC) or SiC-cascoded junction FETs. SiC FETs require an 18V gate drive to fully enhance the FET for lowest ON-resistance; SiC-cascoded junction FETs require a 12V gate drive. In a hot-swap application, the FETs typically operate in linear mode during the inrush current-limiting phase. It is important to design the hot-swap circuit appropriately to keep the SOA stress during the inrush phase within the FET's SOA limits. Protecting the FET during fault scenarios such as a power-up into an output short circuit also requires protection schemes. Overload protection cannot put the switch in an active mode because that leads to rapid overheating.

The LMI800PS's configurable gate-drive feature and lower power-limit thresholds (down to 200W) with configurable timers enable flexibility for system designers to develop a reliable hot-swap solution.

4 Safety and Insulation

Isolating the control and indication signals (management signals) from the high-voltage section and enabling the system to communicate with the hot-swap controller is important. The insulation for these signals must comply with reinforced insulation for Pollution Degree 2, 10% to 90% humidity noncondensing (per safety standard International Electrotechnical Commission [IEC] 62368-1) and clearance and creepage distances for up to 3,050m of altitude (5,000m preferably).

The LMI800PS device package—with ≥ 8 mm of creepage between the high- and low-voltage domain, a 5kVrms isolation withstand voltage (V_{ISO}) rating and 10kV of surge capability—meets the overall safety requirements.

One of those safety requirements is to discharge the output bulk capacitors to <120 V within 2s when the IT tray disconnects from the rack, or when the output disconnects. Placing a FET with series discharge resistors in parallel to the output capacitors can enable this feature. The LMI800PS features a discharge FET control output signal, simplifying output discharge design.

5 Conclusion

The surge in AI processing power has driven a critical shift toward 800V_{DC} power delivery in AI server racks to accommodate the increasing power requirements of advanced AI processors and support high compute density. This shift creates new challenges for system protection and monitoring, particularly during hot-swap events.

TI is at the forefront of development with the LMI800PS high-voltage hot-swap controller, which effectively manages inrush currents, offers comprehensive telemetry for system diagnostics and safety, and maintains reliable operation.

6 References

1. [Diablo 400 Project: Rack and Power](#) (Version 0.5.2). (2025, May 30). Open Compute Project: Austin, Texas.
2. Huntington, J., and Tu, M. (n.d.). [800VDC Architecture for Next Generation AI Infrastructure](#). Nvidia white paper. Accessed Sept. 28, 2026.
3. McDonald, B. (2025, November). [Data Centers Evolve to Meet AI's Massive Power Needs](#) (Literature No. SSZTDB4). Texas Instruments technical article.
4. Pal, A., and Panguloori, R. (2026, March). [Floating Ground Hot-Swap Architecture to Enable Robust Protection on 800V or \$\pm 400\$ V DC Power Distribution](#) (Literature No. STDA029). Texas Instruments.
5. Pal, A., & Panguloori, R. (2025, Q2). [Addressing the Protection Challenges of 48V AI Servers Using Hot-Swap Controllers](#) (Literature No. SLYT863). Texas Instruments Analog Design Journal.
6. Pal, A., and Sagar, S. (2025, January). [Enhancing Server Power Design Using Advanced Features of TI's Smart eFuses](#) (Literature No. SLVAFR8). Texas Instruments.
7. Rogachev, A. (2014, November). [Robust Hot Swap Design](#) (Literature No. SLVA673 Revision A). Texas Instruments.

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you fully indemnify TI and its representatives against any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to [TI's Terms of Sale](#), [TI's General Quality Guidelines](#), or other applicable terms available either on [ti.com](#) or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products. Unless TI explicitly designates a product as custom or customer-specified, TI products are standard, catalog, general purpose devices.

TI objects to and rejects any additional or different terms you may propose.

Copyright © 2026, Texas Instruments Incorporated

Last updated 10/2025