

The power glitches AI data centers can't see



By Shannon Lippincott, marketing manager, high-speed data converters

What happens when your data center's power monitoring can't keep pace with your AI workload?

Today's largest AI deployments can span tens of thousands of graphics processing units (GPUs), producing immense compute power alongside operating conditions that change by the second. As workloads ramp up or down, server power systems must respond to increasingly dynamic demands. The result is voltage droops, overshoots and other transient events at the processor power rail, many of which can occur in nanoseconds (as shown in [Figure 1](#)).




Figure 1. A nanosecond-level transient event in a server rack

The growing impact of unseen power events

Enterprise servers, test and measurement equipment, wireless infrastructure, and other systems have long relied on telemetry to give engineers real-time information about voltage, current, power, temperature and overall system health. However, many of these systems have historically operated at millisecond-level intervals or longer – a discrepancy six orders of magnitude over the transient events now occurring at the processor power rails and the tools meant to observe them.

While nanosecond-level telemetry exists today, many of these events can occur during complex workflows and result in costly brownouts or unexpected downtime, often without ever showing up on existing telemetry systems.

In the blink of an eye (about 100 milliseconds), a server can experience several nanosecond-scale electrical events. Left undetected, they can trigger resets, retests, triage, throttling and interrupted workloads, reducing AI compute performance and utilization at a major cost. Recently, the [Uptime Institute](#) reported that 57% of major outages cost more than \$100,000, with one in five outages exceeding \$1 million.

I've heard similar stories from data center engineers: a training cluster experiences intermittent errors that stall an important run but don't occur frequently enough to reproduce on demand. Conventional telemetry tools shows stable operation where power, thermals and utilization appear fine, leaving no clues to the root cause.

High-speed detection technology has traditionally been too large, power-hungry and costly to deploy across thousands of server racks. Without it, technicians need to manually debug the situation, sometimes spending hours or days combing through hardware, firmware and system logs to isolate nanosecond-scale transients behind the problem.

How faster power monitoring improves troubleshooting

Continuous nanosecond-level monitoring captures "invisible" events the moment they occur. Providing the actionable data required for immediate root-cause analysis and preventive maintenance allows engineers to identify early warning signs before they escalate into costly system failures.

While this level of detail is important, constant streaming of high-volume data would inevitably overwhelm storage and network bandwidth. Event-based remote telemetry addresses the challenge by monitoring the environment continuously and only reporting granular, high-speed transient events, without the burden of excessive data.

According to Karthik Vasanth, TI's vice president of data converters and clocks, "The glitches technicians look for are like a needle in a haystack. Engineers program the parameters of the needle, and the chip searches the haystack."

But there's a bigger benefit here. Monitoring these nanosecond-level events, Vasanth said, is "a forerunner or precursor to the reliability of a data center." In other words, this level of visibility allows engineers to get ahead of failures instead of just reacting to them.




Figure 2. Event-based, remote telemetry enabled by TI's new nanosecond-level technology

Detecting nanosecond power events in real time

Light itself travels almost a foot in a single nanosecond, fast enough that a server can experience thousands of events within a single millisecond. Detecting an event at that speed and precision has typically exceeded the capabilities of telemetry systems built to monitor data centers.

TI's [ADC34DSO75](#) radio frequency (RF)-sampling analog-to-digital converter (ADC), with 3GSPS performance, provides 24/7 nanosecond-level, event-based remote telemetry. By leveraging advanced analog design and an integrated digital signal processor, the ADC34DSO75 processes and refines signals in real time while also reducing field-programmable gate array (FPGA) complexity and cost. It combines a signal-to-noise (SNR) ratio of 68dBFS with low phase and amplitude noise, and has half the power consumption of competing solutions.

With processors operating at <1.2V, electrical events can register as microvolts or millivolts. Identifying an event that tiny requires both timing and voltage resolution, making SNR and dynamic range as critical as raw sampling speed.

The ADC34DSO75's configurable, integrated event detection and capture enable engineers to tailor monitoring to their specific needs. Low-speed streaming to an MCU through an octal SPI helps eliminate the need for costly, power-intensive FPGAs and simplifies the overall architecture.

TI also developed a system-on-module (SOM) board ([Figure 3](#)) built around the ADC34DSO75. The eight-channel board measures eight voltages inside a processor with nanosecond resolution, using two of the data converters, each handling four channels. The board includes an MCU and flash memory that stores event data, so if power actually goes out, the information stays on the card for easy retrieval later. Designers can offload the full onboard memory in about 10mS at 200Mbps, giving them fast access to event data without disrupting operations.



Figure 3. TI's eight-channel SOM measures eight voltages inside a processor with nanosecond resolution

As the first solution to provide nanosecond-level, event-based remote telemetry, the 1-inch-by-2-inch SOM enables data center designers to achieve new levels of system performance at a fraction of the size, power and cost of competing solutions. Remote event reporting eliminates the need to send a technician on-site or connect an oscilloscope to troubleshoot an event, minimizing downtime and reducing the security risk of physical hardware access.

The future of real-time power monitoring

As AI compute density and power demands increase, the need for continuous high-resolution monitoring extends far beyond the data center. In automated test equipment, this level of monitoring prevents nanosecond glitches from skewing test validation and causing errors. In wireless infrastructure, it allows engineers to trace transient events that degrade signal quality, especially in 5G and 6G telemetry systems.

Looking ahead, combining nanosecond-level monitoring with system intelligence will enable predictive fault detection and autonomous system management. The bar for reliability keeps rising. Data center designs increasingly need to meet automotive-grade reliability standards, and in some cases, operators are targeting parts-per-billion levels of reliability. As electronics become increasingly complex, the ability to address stability challenges depends entirely on ultrafine visibility.

Additional resources

- Order the [ADC34RF72 evaluation module](#) to get started.

About the author

Shannon Lippincott is a product marketing manager for high-speed data converters at Texas Instruments. Shannon helps bring precision analog and mixed-signal technologies to market for applications spanning data centers, test and measurement, and communications infrastructure. Lippincott holds Bachelor of Science and Master of Science degrees in electrical engineering and engineering, respectively, from Arizona State University, with a focus on mixed-signal circuit design.

Trademarks

All trademarks are the property of their respective owners.

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