White Paper

Safeguarding against disruptive events: innovative solutions for sub-synchronous oscillations in AI data centers

Karen Cheung
Product manager
Eaton

Roger Cox
Electronics specialist
Eaton

Praveen Sutrave
Engineering manager
Eaton

Overview

Generative artificial intelligence (AI) is changing the world. For data centers, meeting the growing demand for this technology poses significant challenges and prompts a reevaluation of best practices for metering design. Now more than ever, critical infrastructure needs to be robust, reliable and intelligent.

While generative AI is creating big changes across the globe, sometimes the smallest things have the greatest impact. And by small, we are referring to a phenomenon occurring below the power system's fundamental frequency, often at a 10–12 Hz range. These sub-synchronous oscillations (SSOs) created by data centers hosting AI applications could be detrimental to continuous operation. To prevent downtime, equipment damage and potential utility penalties, it is essential to have a means to monitor and detect events that occur at sub-synchronous frequencies and provides alerts as these critical conditions arise.

EATON
Powering Business Worldwide

AI data centers' impact on energy infrastructure

Generative AI requires performance-intensive computing (PIC), which is achieved by using central processing units (CPUs) with a high core count and coprocessors such as graphical processing units (GPUs). This configuration creates the functional equivalent of having hundreds of computers working together on one problem. PIC delivers significantly greater computing performance than traditional data center infrastructure but it consumes significantly more energy as well.

According to studies published by the Electric Power Research Institute (EPRI), energy consumption needs are even more pronounced when performing certain tasks associated with AI. For example, developing and training large language models (LLMs) can consume more than 10 times the amount of energy compared to deploying or using LLMs. Further, the generation of video and images also requires significant processing power and energy usage. Even a simple web search, when augmented by AI, uses nearly 10 times as much energy as a search that is not enhanced by AI.

When these computationally heavy processes take place, a sudden swing in energy consumption occurs. For instance, a small data center might consume 30 MW and variations of +/-1 MW would not be unusual when AI is involved. This is called dynamic loading. Such swings in energy consumption impact the facility at the point of common coupling (PCC) — where the building connects to a utility — as well as affecting neighboring infrastructure, including transformers, transmission lines, generation units and other facilities. The sudden swing of energy can saturate the core of a utility transformer at the building's service entrance and lead to overheating and ferro resonant damage.

The swings of energy additionally have the potential to create further-reaching disruptions because of SSOs that occur when the energy usage characteristics are in resonance with the electrical infrastructure. The effect of this phenomenon goes beyond the impact that a severe voltage sag or localized outage may create. SSOs have been known to damage rotors, turbines and generators in the electrical infrastructure. Utilities are growing increasingly concerned about sub-synchronous resonances (SSRs) caused by SSOs because the resultant transmission issues have grid-wide impacts and may ultimately affect their ability to supply stable power to other nearby facilities.

Resonance unrelated to harmonics

When people hear the word resonance, they tend to think of harmonics. For this reason, it is understandably tempting to believe that any power quality meter capable of capturing higher orders of harmonics and establishing metrics around industry harmonic standards will monitor for SSR and mitigate concerns.

However, it’s important to understand the difference between harmonics and events that occur at sub-synchronous frequencies. Harmonics are a mathematical representation of voltage or current distortion caused by multiples of the fundamental power system frequency (60 Hz in North America) distorting the fundamental current.

A sub-synchronous frequency is one that occurs at frequencies below the power system fundamental frequency (less than 60 Hz), an area with which popular harmonics standards are unconcerned. Unlike harmonics, these SSR points typically occur between 5 Hz and 55 Hz, with AI data centers often exciting oscillations in the 2 Hz to 12 Hz range.

In the context of power systems, resonance typically refers to harmonic resonance: amplified oscillations of voltage or current as the natural frequency of a power system and its non-linear components align. (For example, a 6-pulse variable frequency drive (VFD) tends to create increased levels of the 5th, 7th, 11th and 13th harmonics, which translates to 300 Hz, 420 Hz, 660 Hz and 720 Hz.)

However, resonances are not limited to the interplay between distorted AC voltages and currents. Resonances can physically manifest between a power system and a piece of equipment, such as a generator at a utility provider. A generator shaft will typically have a natural frequency — a property of its physical design (related to properties such as the shaft’s length, diameter or material). When a power system operates on this same frequency, a mechanical resonance occurs between the power system and the generator shaft.

The challenge of sub-synchronous events

Research has shown that SSR, one of many types of events that occur at frequencies below the power system fundamental frequency, can cause a generator shaft to physically oscillate. This puts torque on the generator shaft, leading to generator damage and potential loss of power supply from the utility.

Existing literature tends to talk about SSOs in the context of microgrids, renewable energy sources (exacerbated by the variability of wind and solar power), series transmission on electrical utility lines and use of power electronics in certain applications. In decades past, it was incumbent on the grid owner to manage and attempt to dampen these oscillations to protect power generation infrastructure and ensure stable power supply to the grid.

As technology changes, new sources of SSOs are introduced. A data center hosting an AI application becomes a hotspot for these low frequency events because of the high density of computing driven by high-volume usage of GPUs and microprocessors.

This creates a need to detect events that occur at sub-synchronous frequencies to manage AI workloads and alert data center operations teams when action is needed. Power quality meters (sometimes simply called power meters) are an essential part of electrical distribution equipment like switchgear, switchboards and power distribution units (PDUs), providing actionable information around detected voltage sags, harmonics and transients. Given their integral role in providing data center operations teams with actionable insights around power quality events, these meters are a natural choice for managing and detecting events occurring at sub-synchronous frequencies. It is crucial to select a power quality meter with sufficient sensitivity and processing power to detect and characterize such events, as indicated by the manufacturer’s technical data and product specifications.

Detection vs mitigation of sub-synchronous events

When it comes to metering, there is no adage more common than, “What gets measured gets managed.” This is a principle that is as true in management as it is in power monitoring and event detection.

The ability to detect sub-synchronous events is critical, as it provides key visibility to what is happening and allows data center operations teams to take real-time actions (such as load shedding or dynamic scaling of GPUs) when their loads approach these resonance points. This could prevent the severe consequences of damage to critical infrastructure and supply interruptions.

Sub-synchronous event detection is a smart alternative to other strategies that flatten energy consumption. This helps data centers get their best capacity utilization and avoid unnecessary intervention. It helps to determine the most effective mitigation and validate that an implemented solution was effective.

Sub-synchronous event detection provides long-term robustness as well. Systems change over time as loads and usages change. By choosing a meter that offers event detection, data center personnel can ensure their mitigation continues to be effective as time passes and their needs change.

Sub-synchronous event detection

Different meter manufacturers offer different strategies for detecting sub-synchronous events but some implementations are more effective and efficient than others.

Here are some common approaches manufacturers may take, along with the advantages and limitations of each. These differences should be taken into account when choosing a meter to ensure your power monitoring system design is robust, reliable and intelligent.

Approach one: Flicker

Flicker (sometimes called perceptibility) is a type of sub-synchronous event defined by industry standards IEC 61000-4-15 and IEEE 1453. Historically, the ability to measure flicker was one of the hallmarks of a capable power quality meter because of the sensitivity and precision required. Advances in technology have made the capability to measure flicker common in power quality meters.

However, flicker is much less relevant in today’s market. The curves for flicker severity were derived based on studies of how small voltage fluctuations impacted 60 W incandescent light bulbs (which are largely obsolete today) and how noticeable or perceptible this effect was.

Because flicker measurement capabilities are so ubiquitous, some manufacturers may claim this as a means of detecting the low-frequency voltage artifacts that may cause damaging SSOs. This is a tempting prospect, particularly at data centers that may already have meters installed, or where designs were previously standardized.

This approach has major shortcomings. While monitoring flicker may make it possible to detect a change in the power system, this provides no further details and cannot be used to characterize an event or determine mitigation. This approach cannot differentiate between resonance and other types of sub-synchronous activity, resulting in false positives.

Approach two: High fidelity data “waveform” streaming

Another approach that is gaining traction is the ability to stream metered data over communications protocols. (Some meter manufacturers are calling this high-fidelity data “waveform,” drawing a parallel to full-resolution waveform captures that are typical for power quality meters.) The idea behind this philosophy is that a power quality meter should focus on measuring and communicating data as quickly as possible so a computer or server can process the data and make the determination if the power system is approaching a potentially damaging resonance frequency, essentially a “stream and catch” architecture.

This approach of streaming and catching poses challenges to critical applications. The biggest of these challenges is the computing power and data storage such a solution would require. This becomes a costly prospect for any organization.

From the perspective of solution robustness, this solution falls short. With most data streaming, especially at higher resolutions, options for backing up the data are limited, so any data payloads missed due to an interruption or communications network latency are lost.

The transferring of data to a computer or server for analysis creates a challenging bottleneck. The throughput limitations of communications networks restrict the amount of data that can be provided. This requires decimation of data before transmission. Further, communicating the data to a computer or server for post processing adds a latency, delaying critical decision-making.

The drawbacks of this approach make full deployment at scale challenging. The practical challenges prevent this from being the preferred solution, though it would be a potential alternative for redundant monitoring of critical circuits and systems if used in conjunction with edge-based processing and event detection.

Approach three: Edge-based processing and event detection

Edge-based processing is an approach that evaluates metered data directly on-board the metering device. It allows data to be continuously monitored for event patterns as it is being captured, making this one of the most efficient means of detecting damaging sub-synchronous events. This holds significant potential and is the recommended approach for critical, high-precision applications.

Here are some reasons why edge-based processing is the recommended method of sub-synchronous event detection.

1. **Reduced latency on critical alerts:** When milliseconds matter, it becomes crucial to minimize delays and system latencies. A power quality meter suited for event detection will be able to process measurements in real time, as data is being captured. This removes the communications bottleneck and reduces latency. The meter can detect a sub-synchronous event, analyze for resonance and send an alert immediately.
2. **Higher resolution data gives better detection accuracy:** When it comes to frequency analysis, the amount of data being analyzed matters. A typical power quality meter can sample at least 10x–20x faster (possibly more, depending on the communications protocol) than what is practical to provide over communications. Having access to this volume of data means the meter can more accurately identify the frequency components in the signal. The results are less likely to be influenced by electrical signal noise. Resonance frequencies can be determined at a higher resolution, making it easier to target a narrow band of critical values and reducing the number of false positives.
3. **Reduced cost by eliminating superfluous computing:** An edge-based approach puts the computing power where it is most needed — right at each critical node. A multi-core processor within each power quality meter provides dedicated computing and data processing . This is a time-tested concept since most power quality meters already provide a level of data processing to capture transients and characterize sag severity. Selecting a meter capable of sub-synchronous event detection makes it unnecessary to invest in additional infrastructure to perform the same analysis again.
4. **Option for redundancy in critical areas:** Edge-based processing does not preclude the possibility of additional processing by another power monitoring system or other software. For the most sensitive or critical applications, where redundancy is highly desired, there are multiple options for a power quality meter to support another application with management of potential SSOs. Meters can still provide data streaming in parallel with event detection; using an application for redundancy instead of a primary method of detection means continuous data processing is no longer required . A redundant application may even coordinate its analysis with the power quality meter's detection alarm and waveform captures, prioritizing computing where and when it is needed.
5. **Targeted insights:** By measuring flicker in accordance with IEC 61000-4-15, power quality meters already collect the necessary voltage data to detect low-frequency disturbances . However, analyzing coincident current and power data is more insightful, as significant changes in current and power result in much smaller voltage variations. Spectral analysis of voltage, current and power can go beyond basic detection by highlighting additional downstream issues.

Key takeaways

While the world explores the potential of AI, data centers are being called upon to examine the impact such massive power consumption is having on the surrounding communities as well as the utility that supplies them.

The unique characteristics of AI usage create a significant risk of damaging the electrical infrastructure and the results can be devastating. This creates a need for a detection method that is efficient, cost-effective and scalable.

While several potential methods were discussed in this paper, Eaton recommends using an edge-based processing approach, analyzing high-resolution meter data directly on a metering device. This provides real-time analysis and alerts on the largest amount of data possible, ensuring timely and accurate results for actionable data center management.

Beyond the boundary — other applications for sub-synchronous event detection

With affordable options for sub-synchronous event detection, monitoring for damaging sub-synchronous events need not be reserved for the point of common coupling.

While sub-synchronous events can damage generator shafts, other motors and rotating devices may be affected as well. Data centers with a liquid-cooled design use motors to pump and circulate cooling water through the system. Air-cooled systems may use fans to circulate the air. Fans, HVAC compressors and backup generators are all motors. Much like a generator at a utility, all these rotating devices have a shaft with a natural frequency that might be impacted if resonance causes mechanical (torsional) oscillations.

While this paper is focused on grid-level effects caused by a particular facility, the concepts could be extended to a smaller at-risk area by choosing a different coupling point.

The concept of a point of common coupling considers an imaginary boundary drawn around each facility. The point at which each facility joins (or couples) to the grid is defined as a convenient standard for evaluating the power system within the boundary and whether it has an impact on infrastructure beyond the boundary.

By choosing a different boundary, like a particular area of a facility being fed by the same breaker, one could monitor for impacts caused by that area and determine if it has an impact on the rest of the building. Depending on the facility, additional boundaries and coupling points should be considered to manage the impact of AI loads. Metering for SSR at these boundaries would provide data that could be used as an input to a control system managing the loads and make changes when they approach resonance conditions.

Conversely, larger boundaries could be drawn around a campus or a region with multiple data center campuses to assess the overall impact of the data centers on the grid. This approach provides additional verification between data centers and their utility provider, further ensuring the health and stability of the grid.

Eaton's PXQ: the edge-based power quality monitoring solution for high-performance applications

Eaton's Power Xpert quality event analysis system (PXQ) is designed for edge-based metering and monitoring. It features multiple cores that process three-phase voltages and currents at a high sampling rate of up to 1,024 samples per cycle (equivalent to 61,440 samples per second for 60 Hz power systems and 51,200 samples per second for 50 Hz power systems). This high-resolution sampling allows for precise intermediate processing, where RMS voltage, current and power are updated every half cycle. This is ideal for subsequent flicker and sub-synchronous event analysis, covering frequencies from 0.1 Hz to 55 Hz with a resolution of 0.06 Hz, updated every second.

The images below provide an example of the edge-based processing PXQ performs every second, resulting in the example waveforms and corresponding spectral analysis. Figure 1 shows the captured waveform data for a load current with a sustained 2 Hz rectangular modulation and the corresponding spectral analysis. Figure 2 shows the captured waveform data for 36kW of total power with 2 Hz and 5 Hz oscillations, along with the corresponding spectral analysis. These images demonstrate the previously described capability to capture and detect these low frequency oscillations that contribute to SSO.

In addition to providing monitoring and detection of SSR, PXQ also detects and records sags, swells, transients and harmonics, automatically characterizing events by severity and providing time-based insights to streamline root cause analysis. With precise time stamping, ANSI (revenue-grade) accuracy and the option for high-resolution data streaming, PXQ becomes an indispensable tool for driving operational value and ensuring reliable data center performance.

PXQ integrates seamlessly with Eaton's Brightlayer portfolio, providing a holistic approach to intelligent power management and energy optimization. For assistance with implementing this solution, Eaton also provides comprehensive turnkey services, including consulting, design, deployment and commissioning.

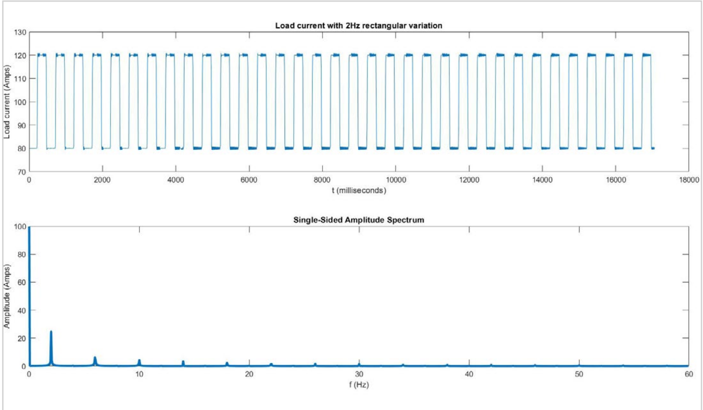

Figure 1

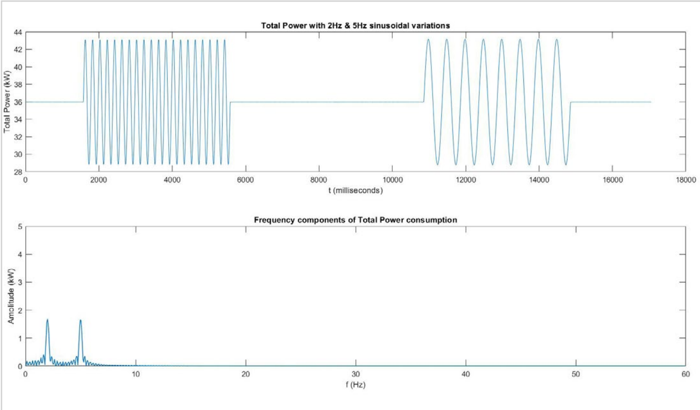

Figure 2

For more information, visit
Eaton.com/preventsso

Eaton
1000 Eaton Boulevard
Cleveland, OH 44122
United States
Eaton.com

Printed in USA
Publication No. WP150003EN
August 2025

Powering Business Worldwide

Follow us on social media to get the latest product and support information.