

Measurement Adequacy for Monitoring Data Center Oscillations

April 2026

Kaustav Chatterjee
Jim Follum
Antos Varghese
Shuchismita Biswas
Evangelos Farantatos
Lin Zhu

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights.** Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <<https://www.ntis.gov/about>>
Online ordering: <http://www.ntis.gov>

Measurement Adequacy for Monitoring Data Center Oscillations

April 2026

Kaustav Chatterjee
Jim Follum
Antos Varghese
Shuchismita Biswas
Evangelos Farantatos
Lin Zhu

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99354

Abstract

Artificial intelligence (AI) training data centers with periodic load profiles can induce sustained grid oscillations across a wide frequency range, making accurate monitoring essential for reliable power system operation. This report evaluates the adequacy of existing measurement systems for monitoring such oscillations, focusing on phasor measurement units (PMUs) and point-on-wave (POW) measurement systems. The analysis shows that while PMUs are highly effective for monitoring low-frequency electromechanical oscillations, they have inherent limitations in accurately representing higher-frequency oscillations due to constraints imposed by reporting rates and the bandwidth of phasor estimation filters. Even when configured with higher reporting rates, the filtering inherent in the phasor estimation process can significantly attenuate oscillation magnitudes, potentially leading to underestimation of oscillatory behavior. This has important implications for compliance and performance monitoring of large loads. To address the limitations associated with PMU-based monitoring, the report examines the use of high-resolution POW measurements, which can capture oscillations across a broader frequency range. However, continuous POW monitoring introduces practical challenges related to large data volumes, communication bandwidth, and real-time data processing. For this reason, the report also discusses emerging approaches that could enable the use of POW measurements as a complementary capability alongside PMUs to improve observability of oscillations from large data center loads.

Summary

The rapid expansion of hyperscale data centers driven by artificial intelligence (AI) workloads is introducing new operational challenges for power system monitoring and stability assessment. Unlike traditional computing facilities, AI training data centers operate large clusters of graphics processing units (GPUs) executing tightly synchronized computational tasks. The repetitive nature of distributed training iterations produces periodic and persistent fluctuations in power demand that can introduce sustained forced oscillations into the electrical grid serving the load. These oscillations may span a wide frequency range, including low-frequency components in the electromechanical range as well as higher-frequency components in the sub-synchronous range.¹ As the size of AI training campuses grows to hundreds of megawatts or more, these oscillatory load behaviors have the potential to interact with existing power system dynamics, potentially exciting poorly-damped modes of oscillation or exciting mechanical torsional modes of nearby synchronous generators. Consequently, accurate monitoring of load-induced oscillations across a broad frequency spectrum is becoming increasingly important for reliable power system operation.

This report presents a comprehensive assessment of measurement system capabilities for monitoring oscillations associated with large data center loads. The analysis focuses on two primary measurement technologies used in modern power system monitoring: phasor measurement units (PMUs) and high-resolution point-on-wave (POW) measurements. PMUs form the backbone of existing wide-area monitoring systems and have been widely deployed for detecting electromechanical oscillations in the range of approximately 0.1–2 Hz. However, the oscillatory behavior produced by large AI data centers may extend beyond this traditional frequency band. This raises important questions regarding the adequacy of PMU-based monitoring for detecting and quantifying higher-frequency sub-synchronous oscillations that may arise from large data center loads.

A detailed technical evaluation of PMU-based monitoring is presented based on the performance requirements specified in the IEC/IEEE 60255-118-1 synchrophasor standard. The analysis examines how PMU reporting rates and phasor estimation filtering influence the observability of oscillations at different frequencies. Analytical derivations, simulation studies, and hardware-in-the-loop experiments demonstrate that oscillations outside the validated measurement bandwidth of PMUs can experience significant attenuation during the phasor estimation process. These effects arise because phasor estimation algorithms are designed to suppress spectral components that are sufficiently far from the nominal fundamental frequency. As a result, oscillations in higher-frequency ranges may appear with reduced amplitude in PMU measurements or may not be represented accurately in the reported phasor quantities. The report demonstrates that such attenuation can lead to underestimation of oscillation magnitudes, potentially causing monitoring systems to misread the severity of high-frequency oscillatory behavior.

To address these limitations, the report investigates the use of POW measurements as a complementary monitoring capability. POW measurements directly capture voltage and current waveforms at high sampling rates, preserving a much wider range of the signals' frequency spectra. This makes waveform-based measurements particularly well suited for detecting

¹ In power systems, electromechanical oscillations typically occur in the approximate range of 0.1–2 Hz. Colloquially, the sub-synchronous frequency range refers to components below the nominal system frequency but above the electromechanical range; for a 60 Hz system, this generally corresponds to frequencies between about 5 Hz and 59 Hz.

oscillatory phenomena across a broad frequency range, including sub-synchronous oscillations that may not be observable through phasor measurements. However, continuous (i.e., gapless) waveform monitoring introduces practical challenges associated with large data volumes, communication bandwidth constraints, and real-time processing requirements. The report therefore also examines emerging monitoring architectures designed to enable scalable use of waveform measurements. These include the use of high-speed communication networks to continuously stream high-resolution waveform measurements to the cloud, trigger-based recording of waveforms within the substation, and a hybrid approach in which analytics are applied in the substation with only summarized oscillation metrics communicated to a centralized monitoring platform.

Overall, the findings of this report highlight the complementary roles of PMU and POW measurements in monitoring oscillations induced by large data center loads. While PMUs remain highly effective for observing low-frequency grid dynamics, waveform-based monitoring provides improved visibility into higher-frequency oscillatory phenomena. Hybrid monitoring approaches that integrate both measurement types offer a promising pathway for ensuring reliable detection and assessment of oscillations associated with emerging large data center loads in modern power systems.

Acknowledgments

This work was supported by the U.S. Department of Energy (DOE) Office of Electricity. The authors gratefully acknowledge the guidance and support of DOE sponsors Sandra Jenkins and Nathaniel Horner. The authors also thank their colleagues at Pacific Northwest National Laboratory, Brett Ross, Jeff Dagle, Chris Klasen, Jason Fuller, and Kevin Schneider for their valuable discussions and comments.

The authors also acknowledge the comments and feedback from the following reviewers.

- Hamed Mohsenian-Rad, University of California, Riverside
- Kat Sico, Duke Energy
- Nadia Smith, NERC
- Yi Hu, Danovo Energy Solutions

Acronyms and Abbreviations

AESO	Alberta Electric System Operator
AEP	American Electric Power
AI	Artificial Intelligence
ATC	American Transmission Company
CPU	Central Processing Unit
DFR	Digital Fault Recorder
ERCOT	Electric Reliability Council of Texas
FE	Frequency Error
fps	Frames Per Second
GPU	Graphical Processing Unit
IEC	International Electrotechnical Commission
IEEE	Institution of Electrical and Electronic Engineers
IBR	Inverter-Based Resources
LIPA	Long Island Power Authority
NERC	North American Electric Reliability Corporation
PMU	Phasor Measurement Unit
POW	Point on Wave
ROCOF	Rate of Change of Frequency
SCADA	Supervisory Control and Data Acquisition
sps	Samples Per Second
TVE	Total Vector Error

Contents

Abstract.....	ii
Summary	iii
Acknowledgments.....	v
Acronyms and Abbreviations.....	vi
1.0 Introduction	1
1.1 Motivation and Background.....	2
1.2 Organization of the Report	5
2.0 Oscillations from Data Center Loads	6
2.1 AI Training Loads and Induced Oscillations	6
2.2 Potential Grid Impacts and Monitoring Need	7
2.3 Utility- and Operator-Proposed Compliance Requirements	8
3.0 Measurement Systems for Monitoring Large Loads.....	10
3.1 Point-on-Wave (POW) Measurements	11
3.2 Phasor Measurements.....	12
3.3 Representation of Oscillations in Phasors vs. Waveforms.....	13
4.0 Considerations for Phasor Measurement-Based Monitoring	16
4.1 Considerations for Monitoring Low- vs. High-Frequency Oscillations	16
4.2 Challenges and Limitations of Monitoring High-Frequency Oscillations.....	17
4.2.1 Impact of Reporting Rate.....	18
4.2.2 Impact of Filtering	19
4.2.3 Case Studies and Simulations	26
4.3 Implications for Monitoring Compliance with Utility-Proposed Data Center Oscillation Limits	31
4.3.1 Impact of Attenuation on Compliance Assessment	32
4.3.2 Practical Considerations for Using PMU Data in Compliance Assessment.....	34
5.0 Considerations for POW Measurement-Based Monitoring.....	35
6.0 Conclusions.....	38
7.0 References.....	40
Appendix A – Attenuation of Off-Nominal Frequency Component	A.1

Figures

Figure 1. Large load oscillation event at ERCOT observed from 30 fps PMU measurements.	3
Figure 2. Large load oscillation event at ERCOT observed from 20 samples/cycle DFR measurements..	4
Figure 3. Representative load profile of a real-world AI data center showing different periodicities of load variation.	7
Figure 4. Measurement systems for monitoring large loads.	10
Figure 5. POW and PMU measurements for a real-world disturbance (fault) event.	11
Figure 6. Representations of a 10 Hz oscillation in POW and phasor data.	14
Figure 7. Attenuation of sideband components due to low-pass filtering in phasor estimation.	19
Figure 8. Window function and magnitude response of the P-class reference filter.	23
Figure 9. Window function and magnitude response of the M-class reference filter.	23
Figure 10. Impact of output reporting rate on the attenuation characteristics of M-class reference filters.	24
Figure 11. Impact of input sample rate on attenuation characteristics of M-class reference filters.	25
Figure 12. Attenuation characteristics of the filters used in Arbiter Systems PMU.	26
Figure 13. Amplitude estimation and attenuation of a 20 Hz oscillation using the IEC/IEEE reference algorithm with M-class filter and 60 fps reporting rate.	27
Figure 14. Experimental setup for hardware-in-loop testing of commercial PMU.	28
Figure 15. Amplitude estimate of a 5 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.	28
Figure 16. Amplitude estimate of an 8 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.	29
Figure 17. Amplitude estimate of a 15 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.	30
Figure 18. Comparison of amplitude estimation for a 15 Hz oscillation reported by a generic PMU with IEC/IEEE reference algorithm and a commercial PMU, with (a) M- and (b) P-class filter designs for 60 fps reporting rate.	30
Figure 19. Amplitude estimate of a 27 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.	31
Figure 20. Comparison of amplitude estimation for a 27 Hz oscillation reported by a generic PMU with IEC/IEEE reference algorithm and a commercial PMU, with (a) M- and (b) P-class filter designs for 60 fps reporting rate.	31
Figure 21. Comparison of POW- and PMU-based monitoring in detecting 20 Hz oscillation.	33
Figure 22. Comparison of POW- and PMU-based monitoring in detecting 1 Hz oscillation.	33
Figure 23. Distributed architecture for data center oscillation monitoring using POW measurements.	37

Tables

Table 1. Window parameters for the Blackman-Harris filters used in Arbiter Systems	
PMU.....	26

1.0 Introduction

The United States is experiencing rapid growth in large-scale data centers, driven primarily by artificial intelligence (AI), cloud computing, and other digital services (NERC 2025, IEEE Spectrum 2026). Recent projections from the Lawrence Berkeley National Laboratory estimate that data centers accounted for roughly 4–5% of total U.S. electricity demand in the early 2020s (Shehabi, et al. 2024). This share is expected to increase substantially by the end of the decade, potentially reaching 6–12% of national electricity consumption under high-growth scenarios (Shehabi, et al. 2024). Much of this increase is associated with hyperscale AI training facilities, whose power requirements can exceed several hundred megawatts per campus. In some regions, developers are already planning multi-gigawatt clusters of such facilities, highlighting the scale at which AI-driven computing infrastructure is expanding.

Beyond their size, AI training data centers differ from traditional data centers in their electrical load characteristics. Conventional enterprise data centers, including corporate server farms, colocation facilities, and standard hyperscale cloud centers, generally exhibit relatively smooth and diversified demand profiles (NERC 2025). Their workload consists of many independent computing tasks whose fluctuations tend to average out over time. In contrast, AI training facilities rely on large numbers of GPUs operating in parallel during tightly synchronized distributed training cycles. Coordinated workloads across computational racks can cause power demand to rise and fall simultaneously across many servers. When aggregated across large clusters, these synchronized operations can produce coherent, periodic fluctuations in electricity demand, creating oscillatory load patterns that differ fundamentally from the smoother profiles typically observed in conventional data centers (Choukse, et al. 2025).

From a power-system perspective, this structured periodic and fluctuating demand acts as a sustained source of forced oscillations (Ko and Zhu 2025, NERC 2025). The resulting oscillations can span a broad frequency range, including low-frequency components in the 0.1–2 Hz band, higher-frequency components in the sub-synchronous 5–59 Hz band, as well as intermediate frequencies between these ranges. Unlike stochastic load variability, these oscillations are repeatable and persistent, meaning they can continuously induce power system oscillations at specific frequencies.

While sub-synchronous oscillations have also been observed in systems with large inverter-based resource (IBR) plants, the underlying mechanisms differ. Oscillations associated with IBRs are typically unintended and arise from control interactions, such as phase-locked loop (PLL) dynamics or weak-grid interactions. In contrast, oscillations induced by AI training data centers are inherent to their normal operation and arise from synchronized workload cycles. As a result, data center–induced oscillations are often persistent, repeatable, and sustained over long durations, rather than transient or condition-dependent, making them a distinct class of forced oscillations from a monitoring perspective. As the penetration of large AI data centers increases, these interactions introduce new operational and stability concerns that extend beyond traditional load growth considerations.

Depending on their frequency, these oscillations may interact with different elements of the power system. At lower frequencies, AI training load induced forced oscillations may interact with poorly-damped inter-area or regional modes, potentially amplifying oscillations across critical transmission corridors and tie lines (Ko and Zhu 2025, Biswas, et al. 2025). At higher frequencies in the sub-synchronous range, forced oscillations may align with torsional modes of

nearby synchronous generators, introducing risks of shaft stress and long-term mechanical fatigue (Zhang and Rose 2026).

Managing these risks requires enhanced visibility into large load dynamics in near real time. System operators should not only detect these oscillations but also accurately quantify their frequency, magnitude, and persistence across a wide spectral range. This capability is particularly important in the context of emerging compliance frameworks, where oscillation limits may be specified in terms of magnitude thresholds within defined frequency bands. However, as the forcing mechanisms associated with AI training loads span from very low inter-area frequencies to sub-synchronous ranges, accurate measurement and faithful representation of oscillations across the entire band of interest present significant technical challenges. Measurement bandwidth limitations, estimation filter characteristics, and reporting rate constraints can all affect the visibility of these phenomena. This report examines these challenges in detail and systematically evaluates the capabilities and inherent limitations of existing measurement systems for monitoring AI load induced oscillations at the grid interface.

1.1 Motivation and Background

Wide-area monitoring in power systems has historically relied on synchrophasor measurements to detect and analyze low-frequency oscillations¹ (NERC 2025). Phasor measurement units (PMUs) were developed to accurately track electromechanical dynamics in the 0.1 to 2 Hz range associated with inter-area and regional modes. Over the past two decades, their performance in this regime has been thoroughly validated through extensive field deployments and operational experience, establishing PMUs as a trusted foundation for wide-area oscillation monitoring and small-signal stability assessment (Chatterjee, et al. 2023).

However, applying the same PMU technology to monitor oscillations induced by large AI-driven data center loads presents distinct challenges. The dynamic behavior associated with AI training facilities differs significantly from the traditional electromechanical phenomena for which PMUs were originally designed. As discussed in the Introduction, AI training loads can generate sustained forced oscillations with broadband spectral content extending beyond the effective low-frequency measurement range of PMUs. Consequently, PMUs may have limited ability to accurately capture higher-frequency dynamics emerging from large data center loads.

Two main factors contribute to these limitations. First, phasor estimation filters are intentionally designed to be selective around the nominal fundamental frequency. This selectivity suppresses frequency components that are widely separated from the fundamental, leading to systematic underestimation of higher-frequency oscillation magnitudes. Second, the typical PMU reporting rate imposes a Nyquist constraint on observable oscillations in the phasor measurements.² Oscillation frequencies that exceed half of the reporting rate cannot be faithfully represented and may appear at incorrect frequencies due to aliasing (Hooshyar, Farnatatos and Patel 2019).

¹ In this document, the term low-frequency oscillations refer to phasor oscillations in 0.1-2 Hz range, while higher-frequency oscillations refer to phasor oscillations in the 5-59 Hz sub-synchronous range.

² Formally, Nyquist's constraint applies to the process of sampling a continuous-time signal to create a discrete-time signal, and it does not directly extend to the reporting rate of phasor estimates computed from that signal. However, from a practical perspective, the parallels between a PMU's reporting rate and a signal's sampling rate are often straightforward to apply.

A recent field event illustrates these effects. In October 2024, the Electric Reliability Council of Texas (ERCOT) observed real power oscillations associated with ramping behavior from a large crypto-mining data center load (Gravois 2025). PMU data reported at 30 frames per second (fps) from the point of interconnection indicated an oscillation frequency of approximately 7.5 Hz with a peak-to-peak magnitude of 25 MW (see Figure 1). In contrast, high-resolution waveform data collected at 20 samples per cycle (approximately 120 samples per second (sps)) from a digital fault recorder (DFR) at the same location showed that the actual oscillation had a frequency of approximately 23 Hz and a peak-to-peak magnitude of 50 MW (see Figure 2). The limited reporting rate of the PMU resulted in aliasing of the oscillation frequency and significant attenuation of the measured magnitude¹.




Figure 1. Large load oscillation event at ERCOT observed from 30 fps PMU measurements.
Source: ERCOT (Gravois 2025).

This example demonstrates that while PMUs remain highly effective for low-frequency oscillation monitoring, they may misrepresent oscillations at higher frequencies. Understanding these limitations is essential for evaluating the adequacy of existing measurement systems and for determining how best to monitor AI load-induced oscillations across the full frequency range of concern.

This work builds on prior efforts examining measurement needs for monitoring oscillations associated with IBRs. The fundamental requirements for capturing higher-frequency sub-synchronous dynamics remain largely the same, particularly with respect to measurement bandwidth, reporting rate, and faithful magnitude representation. However, unlike IBR oscillations, which are often event-driven, the emphasis here extends to continuous compliance monitoring of persistent oscillatory behavior, requiring measurement systems that can reliably track magnitude and frequency across the full operating range over sustained durations.

¹ For detailed explanation, see Section 4.2.1



Figure 2. Large load oscillation event at ERCOT observed from 20 samples/cycle DFR measurements. The red trace in the bottom plot is three-phase power and reveals the 50 MW oscillation. Source: ERCOT (Gravois 2025).

1.2 Organization of the Report

The remainder of this report is organized as follows.

- Chapter 2.0 provides a brief overview of AI training workloads and their characteristic periodic power demand profiles. It explains how these loads can act as sustained sources of forced oscillations in the grid and discusses the potential impacts on power system stability. The section also motivates the need for compliance guidelines and real-time monitoring and summarizes emerging oscillation-related requirements proposed by utilities and system operators in North America.
- Chapter 3.0 introduces the measurement systems available at the grid interface for monitoring oscillations from large data center loads, with emphasis on PMUs and point-on-wave (POW) measurements. It reviews the key distinctions between these technologies in terms of signal representation, measurement resolution, bandwidth, and practical deployment considerations.
- Chapter 4.0 evaluates the adequacy of PMU-based monitoring for detecting oscillations induced by large data center loads. The section reviews the performance requirements specified in the IEC/IEEE 60255-118-1 standard and analyzes challenges associated with monitoring higher-frequency oscillations. It examines the effects of reporting rate and filtering on oscillation observability and presents analytical results, simulations, and hardware-in-the-loop experiments to quantify oscillation attenuation and assess implications for monitoring compliance with emerging utility requirements.
- Chapter 5.0 discusses considerations for monitoring oscillations using high-resolution POW measurements. It highlights practical challenges associated with continuous waveform-based monitoring and briefly reviews emerging approaches aimed at enabling scalable deployment.
- Chapter 6.0 summarizes the key findings of the report and presents concluding observations.

2.0 Oscillations from Data Center Loads

As introduced in Chapter 1.0, the rapid growth of large AI data centers has introduced new patterns of electrical demand that can exhibit pronounced temporal variability. This chapter examines how the operational characteristics of AI workloads can produce oscillatory load behavior and discusses the potential implications of such oscillations for power system dynamics, monitoring, and operational compliance.

2.1 AI Training Loads and Induced Oscillations

AI-centric data centers differ from conventional data centers in both workload composition and power consumption characteristics (NERC 2025). In conventional facilities, server workloads are typically dominated by central processing unit (CPU)–based applications such as web hosting, database services, and search indexing, which generally impose relatively stable computational demand. In contrast, AI-focused data centers rely heavily on graphics processing units (GPUs) optimized for large-scale parallel computation required for machine learning applications. These facilities host a mixture of workloads, including large-scale model training, fine-tuning tasks, and inference services executed across clusters of GPUs. Among these, large-scale AI training jobs typically dominate the aggregate power consumption due to their high computational intensity (Choukse, et al. 2025).

A defining characteristic of AI training workloads is their periodic power consumption pattern, which arises from the repetitive nature of mini-batch training iterations (Choukse, et al. 2025). Each iteration alternates between a high-power phase and a lower-power phase. During the high-power phase, GPUs execute large volumes of parallel computations, resulting in elevated power demand. This is followed by a lower-power phase dominated by memory transfers, communication, and synchronization across distributed compute nodes. The repeated alternation of these phases produces a sustained oscillatory load profile (Choukse, et al. 2025) (NERC 2025). The duration of the iteration cycle leads to a slower periodic component of load variation, while variability within the computational phases introduces faster fluctuations in power consumption (Ko and Zhu 2025).

From a power-system perspective, such periodic load variations can act as persistent sources of forced oscillations (NERC 2025). The multi-periodicity of AI training load profiles introduces forcing components across multiple frequency ranges relevant to power-system dynamics. The slower oscillations associated with iteration cycles typically fall in the 0.1–1 Hz range, while faster fluctuations originating within the high-power computational phases may extend into the sub-synchronous (5–59 Hz) frequency range. Figure 3 illustrates a representative load profile from a real-world AI data center, reproduced from the NERC Large Loads Task Force (LLTF) White Paper (NERC 2025), highlighting these different periodicities of load variation. Similar profiles of cyclic AI training load are also reported by American Electric Power (AEP) (ESIG 2026).




Figure 3. Representative load profile of a real-world AI data center showing different periodicities of load variation. Source: NERC (NERC 2025).

It is also important to highlight that oscillations may also emerge from non-AI conventional data centers, although the underlying mechanisms differ from periodic AI training loads. In these facilities, oscillations are typically associated with high-power electronic interfaces and their control interactions rather than cyclical computational demand. Improper controller settings, aggressive operating set points at high loading levels, or interactions among converter controls can introduce oscillations across a wide frequency range. Such behavior has been observed in recent events involving large data center loads in systems such as ERCOT (see Figures 1 and 2) and Dominion Energy (Gravois 2025, Mishra, et al. 2025).

Because these oscillatory load components, whether originating from periodic AI workloads or power-electronic control interactions, can span multiple frequency ranges relevant to power-system dynamics, they have the potential to interact with existing grid oscillation modes and other dynamic phenomena. The implications of such interactions, along with the associated monitoring needs for system operators, are discussed next.

2.2 Potential Grid Impacts and Monitoring Need

Oscillations induced by large data center loads can interact with power system dynamics across multiple frequency ranges, creating different categories of operational risk (NERC 2025, ESIG 2026). At low frequencies (typically 0.1–1 Hz), periodic load variations may interact with poorly damped inter-area and regional oscillation modes that naturally exist in large, interconnected grids. If the forcing frequency of the load is close to a system mode, the injected oscillatory power can amplify the existing oscillation through resonance, potentially increasing power swings across major transmission corridors and tie-lines (Biswas, et al. 2025, NERC 2025).

Similar forced oscillation phenomena have historically been observed from generator control system malfunctions and industrial cyclic loads, where sustained periodic disturbances propagated across wide areas, where they introduce risks of triggering protection and mitigation actions (NERC 2019). As the scale of hyperscale AI data centers increases to hundreds of megawatts or more, synchronized load behavior may introduce disturbances of comparable magnitude, increasing the likelihood of wide-area oscillatory impacts.

At higher frequencies within the sub-synchronous range (5–59 Hz), oscillatory load behavior can interact with torsional modes of nearby synchronous generators. These torsional modes correspond to the mechanical shaft dynamics of turbine–generator systems. If electrical disturbances repeatedly excite these modes, cyclic torque can accumulate in generator shafts, potentially leading to increased mechanical stress and long-term fatigue (Zhang and Rose 2026, ERCOT 2026). Historically, sub-synchronous torsional interactions have been short-lived events that subsided once mitigation actions were implemented. In contrast, oscillations produced by tightly synchronized computing workloads in large AI data center clusters may persist continuously during normal operation of the facility. Such sustained forcing can repeatedly inject energy near generator torsional frequencies over long periods, increasing the potential for fatigue accumulation. Oscillations produced by tightly synchronized computing workloads in large data center clusters could therefore create localized risks for nearby generating units.

Given these potential interactions across multiple frequency ranges, timely detection and monitoring of load-induced oscillations become essential for maintaining reliable grid operation (ESIG 2026). Accurate characterization of oscillatory load behavior allows system operators to identify emerging risks, evaluate their interaction with existing system modes, and take appropriate mitigation actions before oscillations grow to problematic levels. These operational considerations have motivated utilities and system operators to develop compliance requirements that limit load variability and oscillatory behavior from large data center facilities, as discussed in the following subsection.

2.3 Utility- and Operator-Proposed Compliance Requirements

System operators recognizing the risks associated with oscillatory behavior from data center loads are introducing compliance limits on load variability and oscillation amplitudes. These requirements are intended to constrain load-induced forced oscillations and reduce the risk of interaction with critical grid modes. Several utilities and system operators in North America have proposed preliminary compliance criteria that define acceptable limits on load variability and oscillatory power behavior. Representative examples are summarized below.

1. *American Transmission Company (ATC)*: Repetitive changes in active power must remain below 25 MW for any period shorter than 5 seconds. Additionally, oscillations with frequencies above 0.1 Hz must have peak-to-peak amplitudes not exceeding 25 MW (ATC 2025).
2. *Electric Reliability Council of Texas (ERCOT)*: Load power must not repetitively exceed a 10 MW change within a 5-second sliding time window, limiting the magnitude of short-term load fluctuations introduced into the grid (ERCOT 2026).
3. *Alberta Electric System Operator (AESO)*: Net load variability must remain below 16 kW for every 100 milliseconds. In addition, power oscillations within the sub-synchronous frequency band must be limited to ± 160 kW. These limits may be refined based on system-specific electromagnetic transient (EMT) studies, which consider factors such as

local grid strength, point-of-interconnection configuration, and facility control characteristics (AESO 2025).

4. *PSEG Long Island Power Authority (LIPA)*: Distinct limits are defined for oscillations in two sub-synchronous oscillation (SSO) bands. For the high SSO band (5–55 Hz), the power summed over any two adjacent frequency bins, computed using a 1-second FFT window and averaged over a 10-second rolling window, must not exceed 3.5 MW. For the low SSO band (1–5 Hz), a 10-second FFT window is used, and the sum of power across any two adjacent bins must remain below 10 MW, while the sum across all bins must remain below 20 MW, both averaged over a 60-second rolling window (PSEG 2026).

Assessing compliance with these requirements necessitates accurate measurement of load variability and oscillatory behavior in both the time and frequency domains. The ability to quantify these aspects depends fundamentally on the characteristics and limitations of the measurement systems used to observe power system dynamics. Understanding the capabilities and constraints of such measurement systems is therefore essential for determining whether these proposed limits can be reliably evaluated and enforced in practice. Measurement technologies and their limitations relevant to this problem are discussed in the following chapters.

3.0 Measurement Systems for Monitoring Large Loads

As outlined in Chapters 1.0 and 2.0, the rapid growth of hyperscale data centers has increased the need for measurement systems capable of capturing fast electrical dynamics at the grid interface. These loads can introduce rapid power fluctuations, oscillatory behavior, and abrupt ramping events that occur on time scales significantly faster than those traditionally monitored by supervisory control and data acquisition (SCADA) systems. Consequently, monitoring such phenomena requires measurement technologies with significantly higher temporal resolution and wider measurement bandwidth.

High-speed measurement technologies used in modern power systems can be broadly classified into two categories based on how the electrical signal is represented:

- (1) *Waveform measurements*: obtained by directly sampling three-phase voltage and current signals as point-on-wave (POW) data, and
- (2) *Phasor measurements*: derived from POW data (see Figure 4) as estimates of the magnitude, phase angle, and frequency of the underlying three-phase signals.

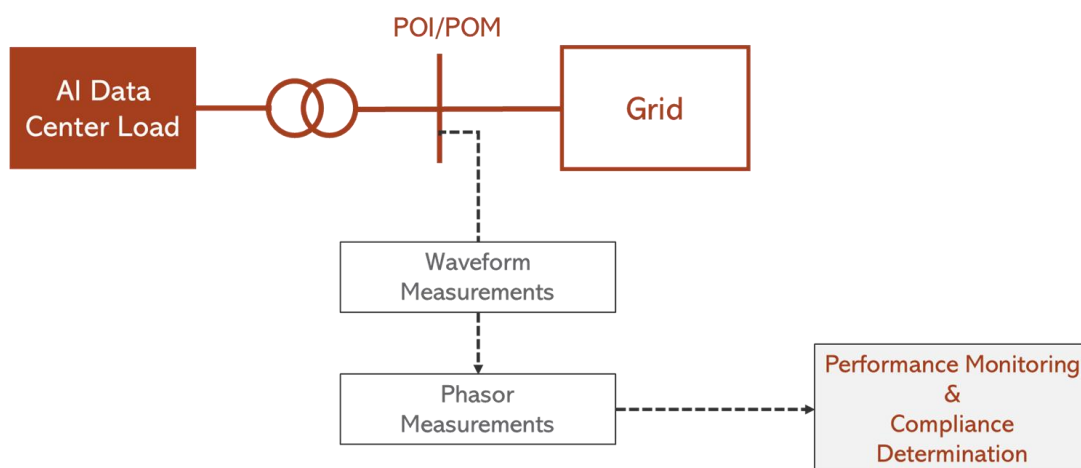


Figure 4. Measurement systems for monitoring large loads.

To highlight the differences between these measurement types, Figure 5 shows the real-world waveform and phasor measurement data for a transmission system fault, (not related to data centers) reported at sampling rates of 960 sps and 30 fps, respectively (ORNL, LLNL 2026). This example illustrates phasor measurements as representative quantities derived from the underlying waveform data through estimation over a finite time window.

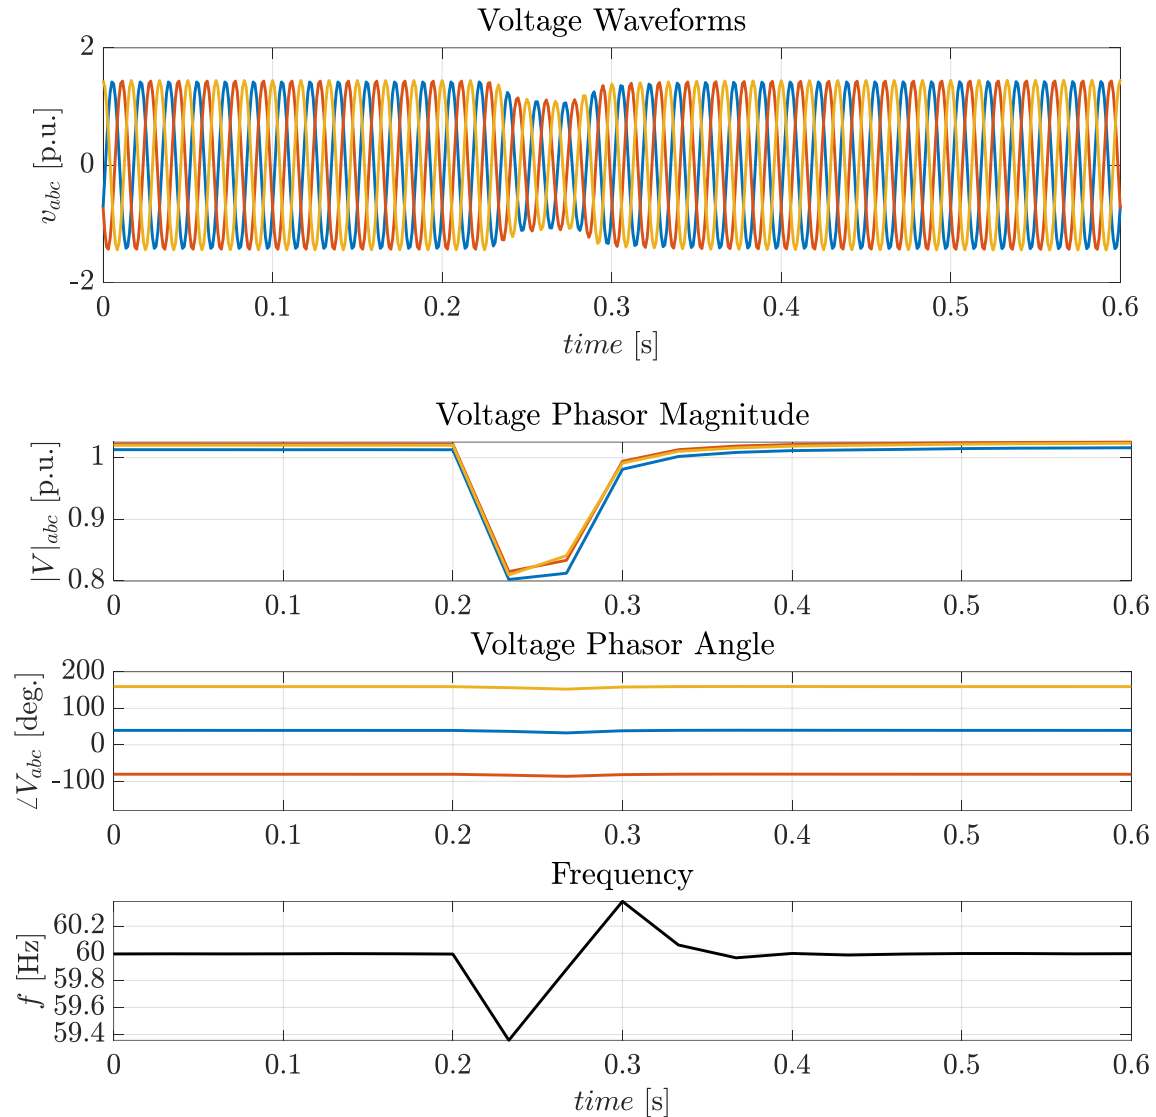


Figure 5. Waveform and phasor measurements for a real-world disturbance (fault) event.

Although both measurement types are based on the same underlying voltage and current waveforms, they differ significantly in terms of signal representation, temporal resolution, data volume, and monitoring capabilities. The following sections introduce these two measurement approaches and discuss their characteristics and practical deployment considerations.

3.1 Point-on-Wave (POW) Measurements

POW measurements are direct time-domain recordings of voltage and current signals obtained by sampling three-phase waveforms at high rates, typically ranging from several kilohertz to tens of kilohertz. Each recorded sample represents the value of the signal at a specific point in time. POW measurements preserve spectral information of the waveform up to several orders of harmonics, including inter-harmonics and other higher-frequency content. This makes POW measurements particularly well suited for monitoring fast electrical phenomena such as switching transients, inverter-driven dynamics, electromagnetic interactions, and high-frequency oscillatory behavior (Mohsenian-Rad and Xu 2023, Follum, Miller, et al. 2021). In the context of

monitoring data center loads, POW data can provide detailed insight into how rapid load variations induce fast oscillations in terminal voltages and currents that may not be observable through lower-bandwidth measurement systems.

Despite these advantages, practical deployment of continuous POW monitoring across transmission networks remains limited. High-speed waveform sampling generates very large data volumes, which impose substantial requirements on communication bandwidth, storage infrastructure, and real-time data processing capabilities (Follum, Hovsopian, et al. 2023). In operational practice, POW measurements are most commonly implemented in disturbance monitoring equipment such as digital fault recorders (DFRs) and power quality monitors. These systems typically record waveform data when triggered by disturbance events rather than continuously streaming measurements across wide-area monitoring systems. That said, developments toward continuous POW recording and monitoring are emerging and are discussed later in Section 5.0.

3.2 Phasor Measurements

Phasor measurements provide a derived representation of voltage and current waveforms through estimates of the magnitude and phase angle of their fundamental-frequency components. For example, a sinusoidal waveform

$$x(t) = \sqrt{2} X_m \cos(2\pi f_0 t + \phi) \quad (1)$$

can be represented in phasor notation as

$$\tilde{X} = X_m \angle \phi \quad (2)$$

where X_m is the phasor magnitude and ϕ is the phase angle relative to a common reference.

A phasor measurement unit (PMU) estimates the quantities X_m and ϕ , along with system's fundamental frequency f_0 , from sampled POW measurements through a signal processing procedure known as phasor estimation. Although the input waveform signals are sampled at rates consistent with POW measurements (typically several kilohertz), the estimated phasor quantities are reported at standardized reporting rates, typically 30, 60, or 120 frames per second (fps)¹. These reporting rates are significantly faster than those of traditional SCADA systems, which typically update measurements every 1–5 seconds, but slower than the underlying POW sampling rates. The phasor representation therefore reduces measurement data volume while retaining key information about the fundamental-frequency behavior of the system. This representation has enabled widespread deployment of PMUs and forms the backbone of modern wide-area monitoring systems used by most transmission utilities.

Given the extensive deployment of PMU infrastructure in recent decades, phasor measurements will continue to play a key role in monitoring system dynamics associated with emerging large electric loads such as hyperscale data centers. However, phasor quantities are not direct measurements of the instantaneous electrical waveform. Instead, they are derived quantities obtained through signal processing applied to sampled waveform data. As a result, oscillations present in the waveform may appear differently in the reported phasor quantities.

¹ Some documents use Hertz (Hz) for reporting rate; in this context, the two are equivalent and denote the number of phasor reports generated per second.

Understanding this relationship is therefore essential when interpreting oscillatory behavior observed in PMU measurements.

3.3 Representation of Oscillations in Phasors vs. Waveforms

In this section, we examine how oscillations observed in PMU measurements in the phasor domain relate to those observed in POW measurements. Although both measurement systems observe the same underlying electrical phenomenon, they characterize oscillatory behavior in fundamentally different ways. Understanding this relationship is essential for correctly interpreting measurements obtained from waveform recorders and PMUs.

Consider a periodic load variation that causes the amplitude of the voltage waveform at a bus to oscillate at frequency Δf . In general, the magnitude and angle of the phasor derived from this voltage will also oscillate and can be expressed as

$$\tilde{X}(t) = X_m \angle \phi, \quad (3)$$

where

$$X_m = \frac{X_0}{\sqrt{2}} \{ 1 + k_m \cos(2\pi \Delta f t) \},$$

$$\phi = \phi_0 \{ 1 + k_a \cos(2\pi \Delta f t) \},$$

f_0 is the nominal system frequency (e.g., 60 Hz), $\frac{X_0}{\sqrt{2}}$ is the nominal phasor magnitude, ϕ_0 is the nominal phase angle, and k_m and k_a are the modulation indices that represent the amplitude of phasor magnitude and angle oscillations, respectively.

To understand how this appears in the waveform, the instantaneous waveform $x(t)$ can be obtained from the real part of the phasor, as shown in (4).

$$\begin{aligned} x(t) &= \sqrt{2} X_m \cos(2\pi f_0 t + \phi) \\ &= X_0 \{ 1 + k_m \cos(2\pi \Delta f t) \} \cdot \cos(2\pi f_0 t + \phi) \\ &= X_0 \cos(2\pi f_0 t + \phi) \\ &\quad + \frac{k_m X_0}{2} \cos(2\pi (f_0 + \Delta f) t + \phi) \\ &\quad + \frac{k_m X_0}{2} \cos(2\pi (f_0 - \Delta f) t + \phi) \end{aligned} \quad (4)$$

Equation (4) shows that when the phasor magnitude oscillates at frequency Δf , the waveform contains two additional spectral components located at $f_0 \pm \Delta f$ frequencies. In the waveform spectrum these appear as inter-harmonic sidebands around the fundamental (or system) frequency. This establishes an important equivalence between the two measurement domains: *phasor representation maps the fundamental and its sideband frequency components in a signal waveform into a time-varying phasor oscillating at the sideband offset frequency Δf* . A much more detailed discussion of this relationship can be found in (Xu, et al. 2025).

Figure 6 illustrates this relationship using an example of a 10 Hz phasor oscillation. In the underlying POW measurements, the oscillation manifests as two inter-harmonic components located 10 Hz away from the 60 Hz fundamental, appearing at 50 Hz and 70 Hz, respectively. In

the phasor domain, the same phenomenon appears as a 10 Hz oscillation in the magnitude and phase of the estimated phasor. The corresponding frequencies are confirmed by examining the frequency spectrum of both the waveform and phasor signals.

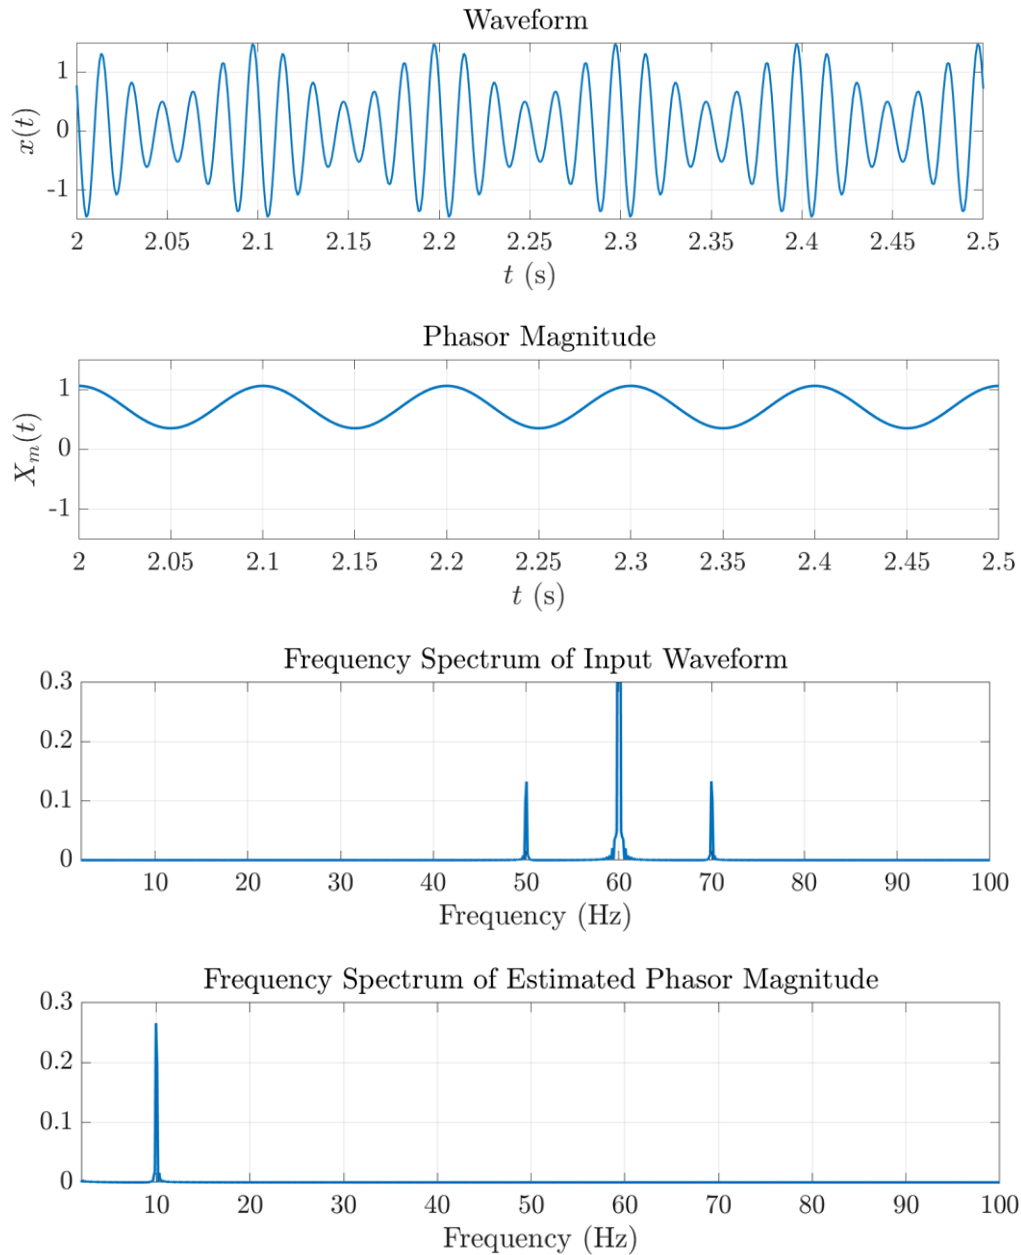


Figure 6. Representations of a 10 Hz oscillation in POW and phasor data.

This example with amplitude-modulated waveforms and inter-harmonic sidebands is representative of oscillatory phenomena observed in real-world grid events, particularly those involving IBRs as reported in (Xu, et al. 2025). The same interpretation can be extended to oscillations originating from large data center loads.

It is important to note that this example assumes an ideal phasor representation. In practice, phasor measurements are obtained through estimation algorithms that include filtering stages

designed to isolate the fundamental frequency component while suppressing other wide-band spectral components present in the waveform. These filtering operations, along with the reporting rate of PMUs, limit the range of oscillatory frequencies that can be accurately represented in phasor measurements. These considerations are discussed further in Section 4.0.

4.0 Considerations for Phasor Measurement-Based Monitoring

To assess the capabilities and limitations of phasor measurement systems for monitoring data center oscillations, it is necessary to examine the performance boundaries defined by the relevant measurement standards. The requirements for synchrophasor measurements are specified in standard IEC/IEEE 60255-118-1 (IEC/IEEE 2018). This standard establishes the framework for evaluating measurement accuracy based on the dynamic response, filtering characteristics, and reporting rates of phasor estimators. It defines compliance criteria under both steady-state and dynamic test conditions, requiring certified devices to satisfy limits on Total Vector Error (TVE), frequency error (FE), and rate-of-change-of-frequency (ROCOF) error.

Although IEC/IEEE 60255-118-1 provides rigorous performance requirements, those requirements are limited to the dynamic conditions and frequency ranges specified in the standard. As large AI-driven data center loads introduce oscillatory behavior that extends beyond traditional electromechanical dynamics, it becomes necessary to evaluate how phasor-based monitoring performs relative to the measurement bandwidth and compliance limits defined in the standard. The following subsection examines these considerations by contrasting low-frequency electromechanical oscillations with higher-frequency and sub-synchronous oscillations that may lie outside the validated dynamic measurement range.

4.1 Considerations for Monitoring Low- vs. High-Frequency Oscillations

IEC/IEEE 60255-118-1 defines two performance classes for PMUs: P-class and M-class, each intended for different operational objectives. P-class PMUs are designed for applications that prioritize fast response and low latency. In contrast, M-class PMUs are intended for applications where improved steady-state accuracy and stronger rejection of off-nominal components are emphasized. For both classes, compliance with the standard requires that the reported phasor quantities satisfy specified limits on TVE, FE, and ROCOF error under both steady-state and dynamic test conditions.

To evaluate estimator performance under dynamic conditions, the standard specifies amplitude modulation and phase modulation tests. These tests introduce sinusoidal modulation around the nominal system frequency and evaluate whether the PMU maintains compliant performance within prescribed TVE limits. From a signal perspective, these tests emulate conditions where the underlying voltage or current waveforms contain inter-harmonic sideband components around the fundamental frequency. As discussed in Section 3.3, oscillations observed in the phasor magnitude or angle arise from such sideband components present in the waveform spectrum. The modulation tests therefore provide a standardized method for assessing how well a phasor estimator can represent oscillatory behavior associated with these spectral components.

The modulation tests also define the dynamic measurement bandwidth within which PMU performance is validated. Table 4 in the standard (IEC/IEEE 2018) summarizes the minimum range of modulation frequencies over which a compliant PMU must maintain TVE within 3%. For P-class PMUs, compliant performance is required for modulation frequencies from 0.1 Hz up to the lesser of $\frac{F_s}{10}$ or 2 Hz, where F_s is the reporting rate. For M-class PMUs, the validated range extends from 0.1 Hz up to the lesser of $\frac{F_s}{5}$ or 5 Hz, also with a 3% TVE limit. These limits

establish the range of oscillation frequencies for which the PMU estimator is guaranteed to accurately represent phasor magnitude and angle variations.

An important implication of this specification is that the effective measurement bandwidth increases with reporting rate only up to a certain point. Once the reporting rate exceeds the value required to reach the upper limits of 2 Hz for P-class and 5 Hz for M-class, further increases in reporting rate do not enhance the measurement bandwidth. As a result, increasing the reporting rate from 30 fps to 60 fps, or even higher, does not increase the range of oscillation frequencies for which compliant performance is guaranteed.

Vendors may implement proprietary phasor estimation algorithms in their PMUs. However, to claim compliance with IEC/IEEE 60255-118-1, the PMU must still satisfy the dynamic performance requirements for its designated class, including the TVE limits under the specified modulation tests. Consequently, while implementations in commercial devices may vary, their effective dynamic response must remain consistent with the performance envelope defined for P-class and M-class PMUs. The standard also recognizes that some implementations may achieve bandwidths beyond these minimum 2 Hz and 5 Hz requirements. Annex I introduces an optional enhanced bandwidth (BW) classification, in which vendors may report the modulation frequency up to which TVE remains within 3%. For example, a designation such as BW11.8 indicates compliant performance with TVE below 3% up to 11.8 Hz modulation. Although this BW classification is not part of the mandatory compliance requirements and is therefore not commonly reported, it provides additional information on the achievable measurement bandwidth of commercial PMUs.

These performance boundaries have important implications for monitoring large load oscillations. As discussed in earlier sections, oscillations associated with large AI-driven data center loads can exhibit broadband spectral characteristics, producing phasor oscillations that span from very low electromechanical frequencies to higher-frequency and sub-synchronous ranges. While oscillations within the validated bandwidth can be accurately represented by PMU measurements, oscillations whose frequencies exceed approximately 2 Hz for P-class PMUs or 5 Hz for M-class PMUs fall outside the range over which estimator accuracy is guaranteed by the standard. In such cases, the reported phasor quantities may not faithfully represent the magnitude or frequency of the underlying oscillatory behavior.

Understanding how oscillatory components outside this validated bandwidth are represented in phasor measurements is therefore essential for assessing the suitability of PMU-based monitoring for large load oscillations. The following subsection examines how key elements of the phasor measurement process influence the representation of these higher-frequency oscillation components.

4.2 Challenges and Limitations of Monitoring High-Frequency Oscillations

Two primary factors determine effective performance limits of PMUs when monitoring higher-frequency oscillations.

- The first is the output reporting rate, in frames per second (fps), which determines how frequently phasor estimates are communicated and therefore limits the range of oscillation frequencies that can be correctly represented in the measurement stream.

- The second is the low-pass filtering inherent to the phasor estimation process, which suppresses spectral components that are sufficiently far from the nominal system frequency.

Together, these mechanisms define the practical boundaries within which phasor measurements can accurately represent oscillatory phenomena (Hooshyar, Faranatos and Patel 2019, Follum, Miller, et al. 2021). When oscillations occur at frequencies approaching or exceeding these limits, distortions may appear in the reported phasor quantities. These distortions can manifest as incorrect oscillation frequencies, attenuated magnitudes, or both. The following subsections examine these effects in detail.

4.2.1 Impact of Reporting Rate

In PMUs, phasor estimation is performed on waveform data at a rate much higher than the output reporting rate used for external streaming. The estimated phasor values are then down sampled before transmission, and external subscribers receive the output at a much lower rate, typically 30 or 60 frames per second (fps), referred to as the reporting rate. The time series of reported phasor measurements can therefore be interpreted as a sampled representation of the underlying phasor dynamics. Consequently, the reporting rate used for data streaming has a direct influence on how accurately the ground truth dynamics are reproduced in the measurement observations.

According to Shannon's sampling theorem, the highest oscillation frequency that can be correctly represented in the reported phasor data is limited by the Nyquist frequency, which in this case is equal to one-half of the reporting rate. When oscillations occur at frequencies higher than the Nyquist frequency associated with the reporting rate, they cannot be represented correctly in the reported phasor measurements. For this reason, most PMU implementations incorporate anti-aliasing filters that suppresses spectral components above the Nyquist frequency prior to reporting. The spectral components above the Nyquist limit that are not sufficiently attenuated by these filters are folded into the observable frequency range, producing image frequencies that differ from the true oscillation frequency.

A simple example illustrates this effect. Consider a PMU reporting at 30 fps, which corresponds to a Nyquist frequency of 15 Hz. Suppose the underlying waveform contains side-band inter-harmonics that produce phasor oscillations at 23 Hz. Since this frequency exceeds the Nyquist limit, the reported phasor data cannot represent the oscillation at its true frequency. Instead, the oscillation appears in the reported measurements at an aliased frequency, which can be calculated using the relationship $f_{alias} = |f_{true} - n F_s|$, where f_{alias} is the observed frequency in the reported phasor data, f_{true} is the true oscillation frequency present in the waveform, F_s is the reporting rate of the PMU, and n is an integer chosen such that the resulting frequency falls below the Nyquist limit. For the example above, with a reporting rate of $F_s = 30 \text{ Hz}$ ¹ and a true oscillation frequency of 23 Hz, the aliased frequency becomes 7 Hz. Consequently, the phasor measurements indicate an oscillation at 7 Hz, even though the actual oscillation present in the waveform occurs at 23 Hz. This behavior is consistent with the large load oscillation example from ERCOT discussed in Section 1.1, where oscillations observed in DFR data occurred at approximately 23 Hz, while the corresponding 30 Hz PMU measurements indicated oscillations near 7 Hz (Gravois 2025).

¹ Frames per second (fps) is the more appropriate unit for reporting rate, but Hz are used here for consistency with the other values in the equation.

There are other examples like this where aliased frequency components in PMU measurements were helpful in alerting utilities to a problem. However, such measurements are extremely limited in understanding the actual characteristics of the oscillation. Because the integer n in the equation for f_{alias} is unknown, the true frequency f_{true} cannot be ascertained from PMU measurements. Further, the anti-aliasing filters employed by the PMU will heavily attenuate the oscillation, so the apparent amplitude will be significantly smaller than the true amplitude.

This limitation is particularly relevant when monitoring oscillations associated with large data center loads, which may introduce higher-frequency components extending into sub-synchronous ranges. Increasing the reporting rate can expand the observable frequency range and reduce the likelihood of aliasing. However, even when the reporting rate is sufficiently high to avoid aliasing, the low-pass filtering inherent in the phasor estimation process may still attenuate higher-frequency components, as discussed next.

4.2.2 Impact of Filtering

In addition to the reporting-rate limitations discussed in the previous subsection, the low-pass filtering in the phasor estimation process imposes a second fundamental constraint on the ability of PMUs to represent higher-frequency oscillations. These filters are designed to extract the fundamental-frequency phasor while suppressing noise, harmonics, and other off-nominal spectral components that are sufficiently far from the nominal frequency. Such filtering improves measurement robustness and helps maintain compliance with the TVE limits specified in the IEC/IEEE 60255-118-1 standard under steady-state and dynamic conditions.

However, the same filtering mechanism also affects the ability of PMUs to accurately represent oscillatory components associated with higher-frequency phasor dynamics. As discussed in Section 3.3, oscillations observed in phasor magnitude and angle arise from inter-harmonic sideband components present in the underlying waveform spectrum. When these sideband components lie outside the passband of the estimator's filter, they are attenuated before the phasor quantities are reported (see, Figure 7). Consequently, oscillations at higher frequencies may appear with reduced magnitude or may not be observable in the phasor measurements at all.




Figure 7. Attenuation of sideband components due to low-pass filtering in phasor estimation.

To illustrate this effect, consider an input waveform containing spectral components at frequencies f_0 , $f_0 + \Delta f$, and $f_0 - \Delta f$. As discussed in Section 3.3, the sideband components $f_0 \pm \Delta f$ produce an oscillation in the estimated phasor at the frequency Δf . Even when Δf is smaller than the Nyquist frequency $f_{Nyquist}$ associated with the reporting rate, accurate

representation of this oscillation depends on whether the sideband components lie within the passband of the estimator filter. If the frequency offset exceeds the effective passband of the filter, denoted here as Δf_{pass}^{LP} , such that $\Delta f_{pass}^{LP} < \Delta f < f_{Nyquist}$, the sideband components will be attenuated by the filter during the estimation process. As a result, the corresponding phasor oscillation at frequency Δf will also appear attenuated in the reported measurements.

This observation highlights two key aspects that must be carefully considered when assessing the ability of PMUs to capture higher-frequency oscillations. First, it is necessary to evaluate the frequency response of the filters to determine the passband range. This determines the frequency range around the nominal system frequency over which oscillatory components can be preserved during phasor estimation. Second, it is necessary to quantify the attenuation introduced by the filters on higher-frequency oscillations as a function of their frequency. These are discussed next.

4.2.2.1 Attenuation of High-Frequency Oscillations Due to Filtering

To demonstrate the impact of filtering, we present a simple analytical example based on the oscillation representation in Section 3.3. The results in this section will be applied to specific filter designs from IEC/IEEE standard and commercial PMUs in Sections 4.2.2.2 and 4.2.2.3, respectively.

Consider the waveform signal in (5) containing a fundamental component and two inter-harmonic sidebands, reproduced from the signal model discussed earlier in (4).

$$\begin{aligned} x(t) = & X_0 \cos(2\pi f_0 t + \phi) \\ & + \frac{k_m X_0}{2} \cos(2\pi (f_0 + \Delta f) t + \phi) \\ & + \frac{k_m X_0}{2} \cos(2\pi (f_0 - \Delta f) t + \phi) \end{aligned} \quad (5)$$

We show that the magnitudes of the sideband components in (5) are attenuated during phasor estimation, which results in the oscillations in the estimated phasors to be attenuated.

Based on the analytical derivations provided in Appendix A, it can be shown that phasor estimation (using the IEC/IEEE reference algorithm) attenuates each frequency component at $f_0 \pm \Delta f$ in (5) by the normalized magnitude response of the filter, denoted $\mathcal{W}(\Delta f)$. Thus, the filtering operation in the phasor estimator affects each frequency component differently. For window function $W[k]$, with sample index k , this magnitude response is given by

$$\mathcal{W}(\Delta f) = \frac{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k] \cdot e^{-j \frac{2\pi \Delta f}{f_s} k}}{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k]} \quad (6)$$

The filters are designed such that the response at $\Delta f = 0$ satisfies $\mathcal{W}(0) = 1$. Consequently, the fundamental component at frequency f_0 is preserved without attenuation during the phasor estimation process. In contrast, the sideband components at $f_0 \pm \Delta f$ are attenuated according to $\mathcal{W}(\Delta f)$. When Δf lies within the passband of the estimator filter, $\mathcal{W}(\Delta f) \approx 1$ and the attenuation is negligible. As Δf moves outside the passband, the magnitude response decreases rapidly, leading to attenuation of the corresponding phasor oscillations.

The attenuation of the sideband magnitudes directly impacts the amplitude of the oscillation observed in the estimated phasor quantities. In particular, the signal in (5) may be rewritten as

$$x(t) = \sqrt{2} X_m^{true} \cdot \cos(2\pi f_0 t + \phi) \quad (7)$$

where

$$\begin{aligned} X_m^{true} &= \frac{X_0}{\sqrt{2}} + \frac{k_m X_0}{2\sqrt{2}} \cos(2\pi \Delta f t) \\ &+ \frac{k_m X_0}{2\sqrt{2}} \cos(2\pi \Delta f t) = \frac{X_0}{\sqrt{2}} \{ 1 + k_m \cos(2\pi \Delta f t) \} \end{aligned} \quad (8)$$

is the true magnitude of the phasor. Note that, X_m^{true} oscillates at Δf and has a peak-to-peak amplitude $\Delta X_m^{true} = \sqrt{2} k_m X_0$. Due to filtering during phasor estimation, the sidebands will be attenuated. As a result, the estimated magnitude of the phasor becomes

$$\begin{aligned} X_m^{estm} &= \frac{X_0}{\sqrt{2}} + W(\Delta f) \frac{k_m X_0}{2\sqrt{2}} \cos(2\pi \Delta f t) \\ &+ W(\Delta f) \frac{k_m X_0}{2\sqrt{2}} \cos(2\pi \Delta f t) \\ &= \frac{X_0}{\sqrt{2}} \{ 1 + W(\Delta f) k_m \cos(2\pi \Delta f t) \} \end{aligned} \quad (9)$$

The estimated phasor magnitude X_m^{estm} also oscillates at Δf , but with a reduced peak-to-peak amplitude $\Delta X_m^{estm} = \sqrt{2} W(\Delta f) k_m X_0$.

This result shows that phasor estimation using the IEC/IEEE reference algorithm attenuates a Δf -frequency phasor oscillation by a factor equal to the magnitude response of the estimator's filter. In commercial PMUs that employ proprietary algorithms different from the IEC/IEEE reference algorithm, the exact attenuation expression may vary depending on the underlying implementation; however, the attenuation is still fundamentally governed by the magnitude response of the effective filter.

In the following subsection, we examine the design and frequency responses of the reference filters defined in the IEC/IEEE 60255-118-1 standard, along with several commonly used filter structures implemented in commercial PMUs. In Section 4.2.3.2, we then evaluate the attenuation behavior of a commercial PMU to illustrate the practical impact of filtering and compare its performance with that of the IEC/IEEE reference algorithm.

4.2.2.2 Frequency Response and Attenuation Characteristics of P- and M-class Reference Filters in the IEC/IEEE Standard

The IEC/IEEE 60255-118-1 standard specifies two reference filter designs used for phasor estimation: (1) a triangular window associated with the P-class estimator and (2) a sinc–Hamming composite window associated with the M-class estimator (IEC/IEEE 2018). These filters represent reference implementations defined in the standard for evaluating estimator behavior and verifying compliance with the dynamic performance requirements of each PMU class.

It is important to note that PMU vendors may implement proprietary estimator filters that differ from these reference designs. However, regardless of the internal implementation, any PMU

claiming compliance with the standard must satisfy the dynamic performance requirements associated with its designated class (P-class or M-class), including the TVE limits under the modulation tests described in Section 4.1. Consequently, while the exact filter kernels implemented in commercial devices may vary, their effective dynamic response must remain consistent with the performance envelope defined for the respective class.

In this section, the reference filter designs from the standard are used as representative examples to demonstrate how estimator filtering influences the attenuation of phasor oscillations. The resulting analysis therefore illustrates the expected attenuation behavior associated with the P-class and M-class performance requirements.

The reference filter designs in the standard are as follows (IEC/IEEE 2018).

1. The P-class reference filter uses a triangular window defined over a symmetric window of length N . For a filter order N and sample index $k \in \{-\frac{N}{2}, \dots, \frac{N}{2}\}$, the filter window is given by (10).

$$W[k] = 1 - \frac{2k}{N + 2} \quad (10)$$

The standard specifies the filter order as $N = 2 \left(\frac{f_s}{f_0} - 1 \right)$, where f_s is the sampling frequency of the input waveform data and f_0 is the nominal frequency. Therefore, $N = 306$ for $f_s = 960$ sps and $f_0 = 60$ Hz, which corresponds to approximately 2 cycles at the fundamental frequency.¹

2. The M-class reference filter uses a composite window formed by multiplying a sinc(\cdot) function with a Hamming window of length N . For sample index $k \in \{-\frac{N}{2}, \dots, \frac{N}{2}\}$, the filter window is given by (11).

$$W[k] = \begin{cases} \frac{\sin(\Omega_c k)}{\Omega_c k} \left(0.54 + 0.46 \cos \frac{2\pi k}{N} \right), & k \neq 0 \\ 1, & k = 0 \end{cases} \quad (11)$$

In (11), $\Omega_c = 2\pi \frac{f_r}{f_s}$ where f_r is the filter reference frequency. Unlike the P-class, the filter order N in M-class is determined as a function of the reporting rate F_s , along with the input sample frequency f_s . The reference design in the standard specifies that for:

- $F_s = 30$ fps and $f_s = 960$ sps, the parameters should be $f_r = 5.02$ and $N = 306$;
- $F_s = 60$ fps and $f_s = 960$ sps, the parameters should be $f_r = 8.19$ and $N = 164$;
- $F_s = 120$ fps and $f_s = 960$ sps, the parameters should be $f_r = 16.25$ and $N = 70$.

¹ The filter order is the window length. In terms of cycles, the window length for a P-class filter is approximately $\frac{N \cdot f_0}{f_s} \approx 2$ cycles at the fundamental frequency.

The window length in cycles can be determined from N using the expression in the footnote 1. For 30, 60 and 120-fps designs, these are approximately 19, 10, and 4 cycles respectively.

For the same waveform sampling frequency and reporting rate, the P-class filter order is significantly smaller than that of the M-class filter. Consequently, the P-class estimator provides faster dynamic response and lower latency. In contrast, the M-class filter provides stronger attenuation of spectral components located away from the nominal system frequency.

Because of these structural differences, the two filters exhibit significantly different passband characteristics and attenuation behavior. The magnitude responses of the P- and M-class reference filters for $f_s = 960$ sps and $F_s = 60$ fps are examined in Figure 8 and Figure 9 respectively. These responses characterize how oscillatory components at different frequency offsets from the nominal system frequency are attenuated during the phasor estimation process. Both filter designs exhibit strong attenuation for frequency offsets greater than approximately 10 Hz. The P-class filter, however, has a comparatively narrower passband extending up to about 2 Hz, within which the filter gain remains close to unity. In contrast, the M-class filter has a wider passband extending to approximately 5 Hz where the gain remains near unity, but it exhibits significantly sharper attenuation beyond this range compared to the P-class filter.




Figure 8. Window function and magnitude response of the P-class reference filter.




Figure 9. Window function and magnitude response of the M-class reference filter.

As discussed in Section 4.2.2.1, the magnitude response of these filters determines their attenuation characteristics. A phasor oscillation at frequency Δf is attenuated by a factor of $\mathcal{W}(\Delta f)$. Therefore, to understand how filter designs impact attenuation, it is important to explore how different parameters in the filter's window function influence the filter's magnitude response $\mathcal{W}(\Delta f)$. We study two key design factors – (1) the reporting rate F_s of the output and (2) the sampling rate f_s of the input waveform.

Figure 10 studies the impact of reporting rate on the magnitude response and attenuation characteristics of the M-class filters described in page 22. The filter with a lower reporting rate has a narrower passband and a sharper attenuation behavior. Observe that an oscillation of 10 Hz is attenuated by nearly 40% for a filter associated with a 30 fps report rate. The same oscillation is attenuated much less, approximately 10%, when a 60 fps report rate is used. If the report rate is further increased to 120 fps, the filter gain at 10 Hz is nearly unity and there is almost no attenuation. Nonetheless, at higher frequencies, the attenuation is still significant.




Figure 10. Impact of output reporting rate on the attenuation characteristics of M-class reference filters.

In Figure 11, the impact of input sample rate on attenuation characteristics is studied, for a fixed filter order (i.e., same window length). Two designs of the M-class filter for $F_s = 60$ fps are evaluated corresponding to $f_s = 960$ sps (i.e., 16 samples/cycle) and $f_s = 7680$ sps (i.e., 128 samples/cycle), for $N = 164$. Increasing f_s slows the rate of attenuation for higher frequencies.




Figure 11. Impact of input sample rate on attenuation characteristics of M-class reference filters.

4.2.2.3 Frequency Response and Attenuation Characteristics of Filters Used in Commercial PMUs

We also analyze the attenuation characteristics of some common filters used in commercial PMUs. The Blackman–Harris family of window functions are employed for phasor estimation in the Arbiter Systems PMU Model 1133A (Arbiter Systems 2007). The parametric form of this filter family is given by

$$W[k] = a_0 + a_1 \cos\left(\frac{2\pi k}{N}\right) + a_2 \cos\left(\frac{4\pi k}{N}\right) + a_3 \cos\left(\frac{6\pi k}{N}\right), \quad (12)$$

for $k = -\frac{N}{2}, \dots, \frac{N}{2}$, where coefficients a_0, a_1, a_2 and a_3 determine the specific window type within the Blackman–Harris family. The parameter choices for respective window types are presented in Table 1. Arbiter 1133A allows the user to select between these filter windows. The filters are designed for an input sample rate $f_s = 10240$ sps and a window size of $N = 1024$ samples.

The magnitude response and the attenuation characteristics of these filters are compared with the reference P-class filter for the same f_s in Figure 12. The Blackman-Harris filter windows have relatively slower attenuation behavior compared to the reference P-class filter. Among these, the flat-top window has the best performance in this regard with nearly zero dB attenuation up to 10 Hz. However, beyond this range, the attenuation is significant. For most other filters in Table 1 the passband is much narrower, and magnitude drops appreciably beyond 2 Hz.

Table 1. Window parameters for the Blackman-Harris filters used in Arbiter Systems PMU.

Window	a_0	a_1	a_2	a_3
Hann	0.50	0.50	0	0
Hamming	0.54	0.46	0	0
Blackman	0.42	0.50	0	0.08
Nuttall	0.355768	0.487396	0.144232	0.012604
Flat-Top	0.2810639	0.5208972	0	0.1980399




Figure 12. Attenuation characteristics of the filters used in Arbiter Systems PMU.

4.2.3 Case Studies and Simulations

To demonstrate the impact of filtering on oscillation attenuation, this section presents simulation and experimental case studies evaluating the performance of phasor estimation algorithms under high-frequency oscillations.

4.2.3.1 Phasor Estimation using the IEC/IEEE Reference Algorithm

We synthesize an amplitude modulated three-phase waveform signal with a 60 Hz fundamental and two sidebands at 40 Hz and 80 Hz. This corresponds to a 20 Hz phasor oscillation with a peak-to-peak (p-p) amplitude of 0.25 p.u. Phasors are estimated from the synthesized waveforms using the reference algorithm specified in IEC/IEEE 60255-118-1, configured with an M-class filter design with an input sample rate $f_s = 960$ sps and an output reporting rate $F_s = 60$ fps, corresponding to a filter window length $N = 164$ samples (~ 10 cycles). The results are

shown in Figure 13. The estimated phasor magnitude exhibits significant attenuation of the oscillatory component. Specifically, the original 0.25 p.u. peak-to-peak oscillation is reduced by approximately 80%, resulting in an observed oscillation magnitude of about 0.05 p.u. Although the oscillation frequency is preserved in the phasor estimates, the magnitude is substantially attenuated. This attenuation is consistent with the magnitude response of the reference M-class filter.




Figure 13. Amplitude estimation and attenuation of a 20 Hz oscillation using the IEC/IEEE reference algorithm with M-class filter and 60 fps reporting rate.

4.2.3.2 Phasor Estimation using Commercial-Grade PMUs

In this section, we demonstrate the limitations of PMUs in capturing high-frequency oscillations through hardware-in-the-loop (HIL) testing of a commercial PMU. The experimental setup is shown in Figure 14.

An OPAL-RT real-time simulator is used to generate analog waveforms for phasor estimation. The simulator also acts as a PMU data client to receive the synchrophasor data that is streamed by the commercial PMU. Both the OPAL-RT simulator and the PMU are synchronized using a substation clock that receives precise time through a global navigation satellite system (GNSS) antenna. The analog waveform signals, generated by the OPAL-RT simulator, are amplified using a voltage amplifier to the voltage levels suitable for input to the commercial PMU. The commercial PMU streams the synchrophasor data, complying with the IEEE C37.118.2 standard format (IEEE 2024), over TCP/IP to the OPAL-RT simulator where the data is logged for analysis (Hooshyar, Haddadi, et al. 2022).

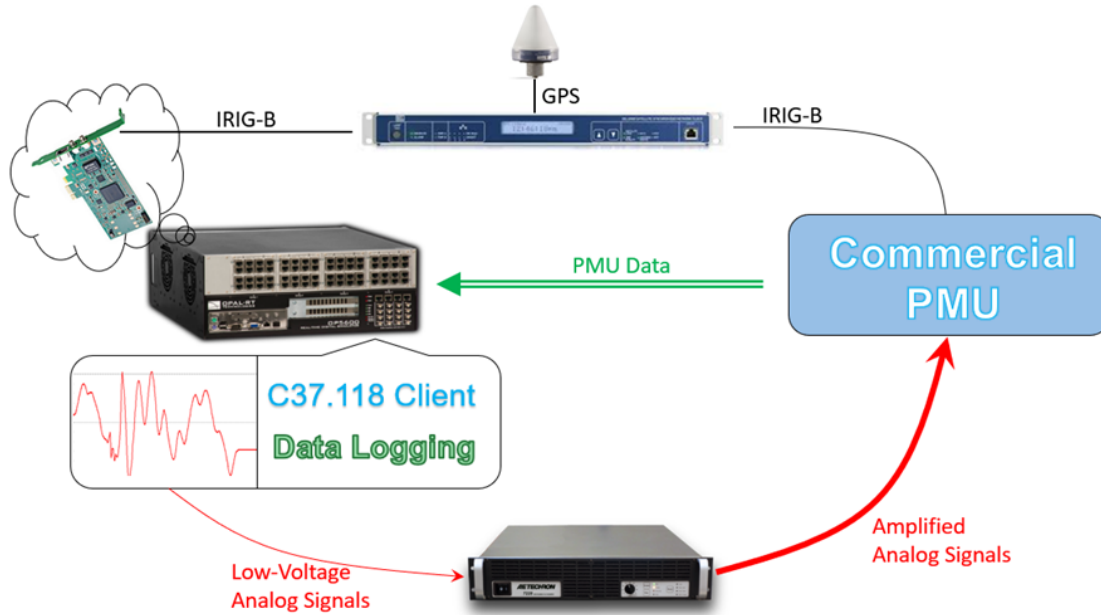


Figure 14. Experimental setup for hardware-in-loop testing of commercial PMU.

Several cases are investigated with varying oscillation frequency. In all cases the PMU is configured as both P- and M-class, and two reporting rates of 30 and 60 fps were tested. When configured as M-class, the window length of the filter was set to six cycles.

- 1) *5 Hz oscillation*: First, we analyze a case of 5 Hz oscillation with 0.1 p.u. magnitude. The waveform signal with a 60 Hz fundamental was amplitude modulated at 5 Hz to create sidebands at 55 Hz and 65 Hz. Figure 15 shows the results from the FFT analysis on the PMU-reported phasor magnitude outputs for P- and M-class configurations with two different reporting rates. It is shown that the PMU output, in this case, preserves the 5 Hz oscillation without aliasing or attenuation. The percentage value in the figures for each class represents the percentage of the true oscillation magnitude that was preserved after phasor estimation.




Figure 15. Amplitude estimate of a 5 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.

- 2) *8 Hz oscillation*: Next, we analyze an 8 Hz oscillation. In this case it is shown in Figure 16 that the PMU performance becomes more dependent on both reporting rate and PMU filter class, with some attenuation in the output signal. In particular, for 60 fps reporting rate and for both P- and M-class PMU configurations, there is a 9% attenuation in oscillation magnitude. For 30 fps reporting rate, when the PMU is configured as M-class the output has higher attenuation compared to the case when the PMU is configured as P-class where the attenuation is minimal.




Figure 16. Amplitude estimate of an 8 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.

- 3) *15 Hz oscillation*: In the next test case, we increase the oscillation frequency to 15 Hz. The results for the P- and M-class configurations are shown in Figure 17. When the PMU is configured with 30 fps reporting rate, the oscillation is not observed in the reported output, as expected considering the Nyquist limit. For 60 fps reporting rate, there is noticeable signal attenuation for both P- and M-class PMU configurations, larger in the M-class case. For this case, we also compare the outputs of the commercial PMU with a generic PMU model following the reference phasor estimation algorithm in the IEC/IEEE standard, shown in Figure 18. Compared to the commercial reference algorithm, the attenuation is higher in the commercial PMU for both PMU classes.

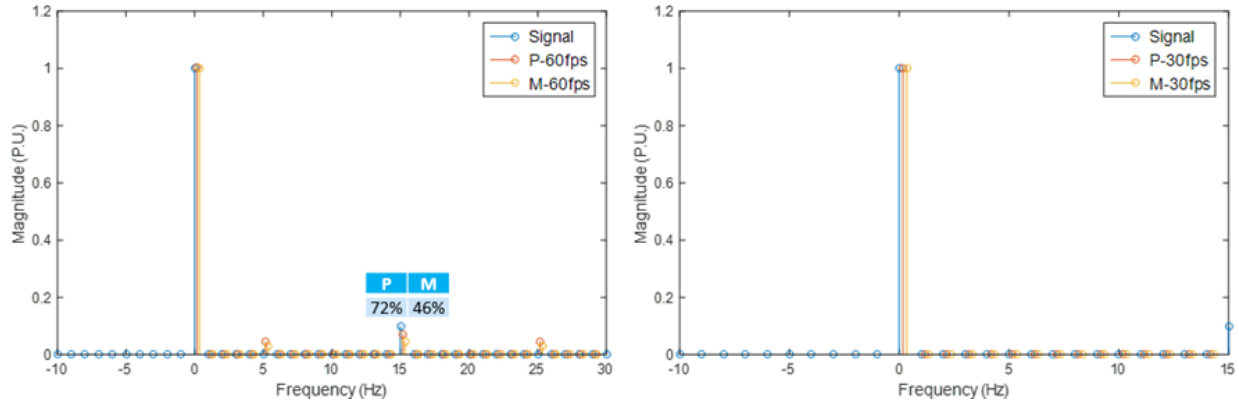


Figure 17. Amplitude estimate of a 15 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.




Figure 18. Comparison of amplitude estimation for a 15 Hz oscillation reported by a generic PMU with IEC/IEEE reference algorithm and a commercial PMU, with (a) M- and (b) P-class filter designs for 60 fps reporting rate.

- 4) *27 Hz Oscillation*: Lastly, we evaluate the performance for a 27 Hz oscillation. The results are shown in Figure 19. For 30 fps reporting rate, the oscillation is not observed in the outputs, as expected based on the Nyquist limit. For 60 fps reporting rate configuration the PMU output signal has significant attenuation. The attenuation for the P-class configuration is 72%, while for the M-class configuration the attenuation is nearly 100%. In this case too, a comparison is made with the output of a generic PMU with IEC/IEEE reference algorithm, and the results are shown in Figure 20. Similar to the previous case, the IEEE PMU model shows less attenuation for both PMU class configurations.

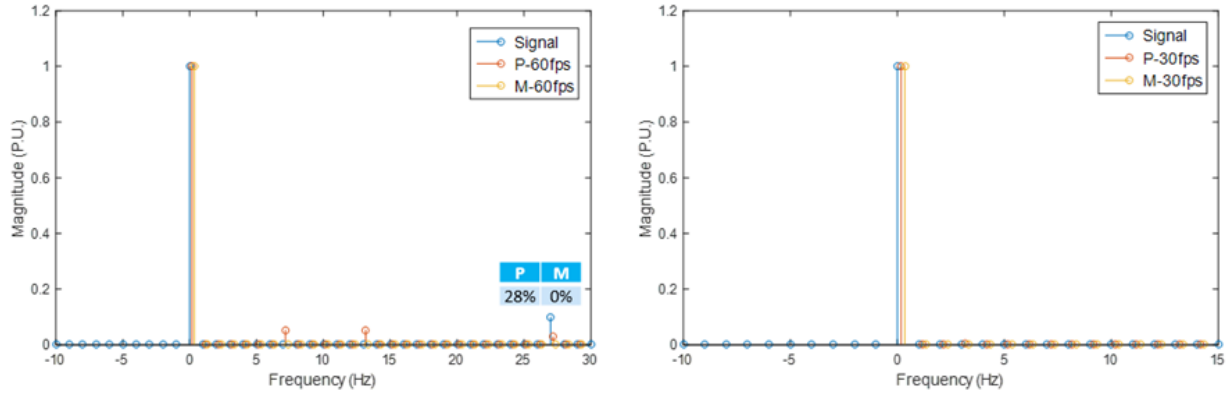


Figure 19. Amplitude estimate of a 27 Hz oscillation reported by a commercial PMU with P- and M-class filter designs for (a) 60 fps and (b) 30 fps reporting rates.




Figure 20. Comparison of amplitude estimation for a 27 Hz oscillation reported by a generic PMU with IEC/IEEE reference algorithm and a commercial PMU, with (a) M- and (b) P-class filter designs for 60 fps reporting rate.

These case studies point to the following conclusions. PMUs are highly effective in monitoring low-frequency oscillations but their ability to capture higher frequency oscillations is fundamentally constrained by the PMU's reporting rate and the attenuation characteristics of the filters internal to the phasor estimation process. This has profound implications in monitoring compliance of data center loads to utility-specified oscillation limits.

4.3 Implications for Monitoring Compliance with Utility-Proposed Data Center Oscillation Limits

The case studies presented in Section 4.2.3 demonstrate the limitations of PMUs in accurately representing oscillations at higher frequencies. In this section, we examine the implications of these limitations for monitoring compliance with the utility-proposed oscillation limits introduced in Section 2.3. As discussed previously, oscillations whose frequencies fall outside the effective passband of phasor estimation filters experience attenuation in their estimated amplitudes. Consequently, high-amplitude oscillations may appear significantly smaller in the phasor domain than in the underlying waveform measurements. When compliance assessments rely solely on these attenuated phasor estimates, oscillatory behavior that exceeds the actual limits may nevertheless appear to satisfy the specified thresholds, thereby providing an overly optimistic assessment of compliance.

4.3.1 Impact of Attenuation on Compliance Assessment


To illustrate this effect, we consider an example based on the criteria proposed by the PSEG Long Island Power Authority (LIPA) (PSEG 2026). LIPA specifies a detection criterion for oscillations in the 5–55 Hz band. In this approach, instantaneous power is first computed from POW voltage and current measurements. A Fast Fourier Transform (FFT) is then applied to a 1-second sliding window, and the zero-to-peak magnitude of power (in MW) for each frequency bin is obtained. The magnitude summed over two adjacent bins is subsequently averaged over the preceding 10-second period and compared against a threshold of 3.5 MW. The LIPA criterion is defined using power derived from POW measurements with a minimum sampling rate of 600 sps.

To demonstrate the impact of phasor filtering, the same detection criterion is implemented with phasor estimates derived from the POW data. Note that (PSEG 2026) does not allow for the use of PMU data; the PMU implementation is for illustration only, and the results are not intended to evaluate the method in (PSEG 2026).

For the input POW data, we synthesize an amplitude-modulated waveform designed to emulate a 20 Hz oscillation. The true oscillation amplitude of 4.8 MW exceeds the 3.5 MW detection threshold. Phasors are then estimated from the sampled waveform data using the reference phasor estimation algorithm specified in the IEC/IEEE 60255-118-1 standard, employing a 120 fps P-class filter.

Using this setup, we compare the detection performance obtained when the LIPA criterion is applied to instantaneous power from POW measurements versus power computed from estimated phasors. The results are shown in Figure 21. For the POW measurements, the power component at 20 Hz computed using LIPA’s approach exceeds the 3.5 MW threshold, and the oscillation is successfully detected. However, when the same analysis is performed using the active power derived from phasors, the oscillation magnitude is attenuated due to the filtering characteristics of the phasor estimator. As the 10-second averaging window slides across the signal, the estimated oscillation magnitude gradually increases but converges to a value that remains below the detection threshold. As a result, an oscillation that is clearly detectable with POW data fails to be detected when the monitoring system relies solely on PMU data, even when phasors are reported at 120 fps.

In contrast, when the oscillation frequency is reduced to 1 Hz, the estimated oscillation magnitudes obtained from both waveform and phasor measurements are essentially identical (see Figure 22). This occurs because 1 Hz lies within the effective passband of the 120 fps P-class estimation filter, meaning that the oscillation amplitude is not attenuated during phasor estimation. Under these conditions, the LIPA detection criterion produces the same outcome regardless of whether waveform measurements or phasor estimates are used.



10

Figure 21. Comparison of POW- and PMU-based monitoring in detecting 20 Hz oscillation.




Figure 22. Comparison of POW- and PMU-based monitoring in detecting 1 Hz oscillation.

These results highlight an important implication for monitoring compliance with utility-specified oscillation limits. While PMU-based monitoring may provide reliable compliance assessment for low-frequency electromechanical oscillations, its effectiveness can be significantly challenged when evaluating higher-frequency oscillatory phenomena, particularly those approaching or exceeding the passband limits of phasor estimation filters. Consequently, compliance monitoring frameworks that rely exclusively on phasor measurements may underestimate the severity of high-frequency oscillations and may therefore report compliance even when the true waveform behavior exceeds the specified limits.

4.3.2 Practical Considerations for Using PMU Data in Compliance Assessment

The limitations highlighted above do not necessarily preclude the use of PMU data for oscillation monitoring. If the magnitude response of the phasor estimation filter is known, the estimated oscillation amplitudes could be adjusted to compensate for the expected attenuation at a given frequency. In such a scenario, the oscillation magnitude obtained from PMU data may be corrected using the known frequency response of the estimator, thereby providing a more accurate representation of the underlying waveform oscillation. However, in practice, implementing such corrections is challenging. In most commercially deployed PMUs, the phasor estimation algorithms and associated filter designs are proprietary. As a result, the frequency response characteristics of the estimator are generally not available to system operators. Without this information, it is difficult to analytically determine the attenuation experienced by oscillations at specific frequencies, limiting the ability to directly compensate PMU measurements for filter-induced magnitude reduction.

In an appropriate laboratory environment, it may be possible to empirically determine a PMU's attenuation characteristics by playing in waveform data with oscillations at a range of frequencies. Such comparisons can provide approximate correction factors that describe how oscillation magnitudes observed in PMU data relate to the underlying waveform oscillations. These empirically derived relationships may then be used to refine detection thresholds or guide the interpretation of compliance assessments when PMU data is the only measurement available during routine operation. For example, if it is known that oscillations in a particular frequency band are systematically attenuated by a certain factor in PMU measurements, compliance thresholds could be adjusted accordingly or monitoring systems could incorporate correction factors when evaluating oscillation magnitudes.

It is also worth noting that some PMU vendors may provide dynamic compliance beyond the minimum bandwidth specified in IEC/IEEE 60255-118-1. In such cases, the achievable measurement bandwidth is typically documented in device manuals or technical specifications (Vizimax 2026). Some vendors also report an extended bandwidth classification, consistent with the optional enhanced bandwidth (BW) concept described in Annex I of the standard, indicating the modulation frequency up to which TVE remains within specified limits. Although this information is not uniformly reported and is not part of the mandatory compliance requirements, it can provide useful insight into the practical dynamic performance of specific PMU implementations.

While POW-based monitoring provides the most accurate representation of high-frequency oscillatory behavior, waveform data are not universally available across transmission networks. In contrast, PMUs remain the backbone of wide-area monitoring infrastructure and provide time-synchronized measurements at many locations in the grid. Consequently, understanding and accounting for the bandwidth and attenuation characteristics of PMU measurements is essential for interpreting oscillation measurements correctly. At the same time, the limitations identified in this section highlight the value of high-resolution POW measurements as a complementary monitoring capability, particularly for detecting and characterizing high-frequency oscillations. The use of POW-based measurements alongside PMU data therefore represents an important direction for improving the observability of large load oscillations, which is explored further in the next chapter.

5.0 Considerations for POW Measurement-Based Monitoring

The analysis in the previous chapter highlighted the limitations of PMU-based monitoring for accurately representing higher-frequency oscillations and motivated the need for using POW measurements. While POW measurements provide significantly higher temporal resolution and can accurately capture oscillatory behavior across a broader frequency range, their adoption for continuous monitoring introduces practical challenges related to data acquisition, transmission, and storage (Silverstein and Follum 2020, Follum, Hovsapien, et al. 2023). This chapter discusses these barriers and examines the practical considerations associated with implementing POW-based monitoring for large data center loads.

Although the power industry has long utilized instruments such as power quality (PQ) meters and flicker analyzers, these devices were primarily designed to monitor voltage disturbances, harmonic distortion, and flicker phenomena. Consequently, they generally do not provide the spectral resolution required to estimate oscillatory components within narrow frequency bands needed to detect dynamic oscillations associated with large and rapidly varying loads. Without significant redesign of their measurement and signal processing capabilities, many of these instruments are not well suited for resource performance monitoring.

Recognizing this gap, the IEEE vision paper on synchronized waveform measurements identifies three schemes for handling high-resolution POW data (Xu, et al. 2022). The first scheme, and the most commonly deployed, relies on event-triggered recording of waveform measurements locally at a substation or point of measurement. In this approach, high-resolution voltage and current waveforms are stored locally and retrieved on demand but are not streamed continuously to a central location (Xu, et al. 2022). DFRs are a common example of devices operating in this mode. This paradigm has been widely used for post-event analysis of faults, switching events, and other transient disturbances, and is effective in limiting communication and storage requirements. However, event-triggered recording may not be well suited for monitoring oscillations associated with large and rapidly varying loads such as data centers. Load-induced oscillations may persist for extended periods and may initially occur below traditional disturbance trigger thresholds. As a result, relying solely on event-triggered recordings may fail to provide the continuous visibility required to detect sustained oscillatory behavior or evaluate compliance with oscillation limits.

The second scheme involves continuous recording and streaming of waveform data to a centralized location, either through gapless recording or periodic multi-snapshot capture (Xu, et al. 2022). Although continuous streaming provides the most complete visibility for compliance monitoring, it introduces significant practical challenges. POW recorders typically sample voltage and current signals at rates ranging from several hundred to several thousand samples per second. Continuous streaming of such measurements from multiple monitoring locations can therefore generate extremely large data volumes (Follum, Miller, et al. 2021). Existing communication infrastructure deployed across utilities and system operators is generally not designed to support such high-volume real-time data transfer (Follum, Hovsapien, et al. 2023). Upgrading communication and storage systems to support continuous waveform streaming would require substantial investments.

A real-world operational experience illustrates the scale of this challenge. A pilot study conducted by Dominion Energy reported that continuous streaming of waveform data from 66 substations sampled at 16 samples per cycle generated approximately 3.9 TB of data per day.

Even with 1 Gbps communication links, several hours were required to transport the data to centralized servers (Purcell, et al. 2024). These results highlight the practical challenges associated with large-scale centralized streaming of waveform measurements for routine monitoring applications.

More recent efforts have explored cloud-hosted data ingestion platforms designed to handle large volumes of waveform measurements (Aminifar 2025). These approaches demonstrate the feasibility of centralized streaming using advanced cloud infrastructure and high-speed communication networks, but they still require substantial data management capabilities and introduce additional architectural complexity. Beyond communication and storage constraints, real-time analysis of high-resolution waveform data also presents computational challenges. Monitoring large numbers of measurement locations would require utilities to process high volumes of incoming data continuously to detect oscillatory behavior, evaluate compliance with operational limits, and maintain system-wide situational awareness.

The third scheme adopts a hybrid architecture in which POW measurements are recorded continuously at the measurement location and waveform segments are streamed only when disturbances are detected (Xu, et al. 2022). In this configuration, streaming need not occur in real time, and delays may be acceptable depending on the end-use application and monitoring objective. Prior work has also demonstrated that with the recently standardized IEEE 2664 streaming telemetry transport protocol (STTP), short-duration waveform snapshots of disturbance events can be packetized and interleaved with existing phasor data streams using available bandwidth, reducing additional communication requirements for POW streaming (Chatterjee and Follum 2025) (Saroare, Hasnat and Ahmed 2026). These studies were performed with a focus on faults, but the concept could be expanded to address oscillations.

For example, recent developments have focused on a broader set of distributed waveform analytics in which POW measurements are recorded continuously and analyzed locally at or near the point of measurement (see Figure 23). In these approaches, high-resolution waveform measurements sampled at rates on the order of several thousand samples per second can be streamed to local analytics engines, where real-time processing is performed without requiring transmission of raw waveform data to centralized servers (Bestebreuer 2025, Lackner and Carroll 2025). This enables continuous monitoring of oscillatory behavior while avoiding the communication and storage burdens associated with centralized waveform streaming.

Several emerging implementations support streaming of high-resolution voltage and current waveform measurements to local processing engines that perform real-time oscillation detection (Bestebreuer 2025, Cheung, Cox and Suttrave 2025). These systems enable continuous monitoring of oscillation magnitude within configurable frequency bands and support persistence-based detection logic and configurable thresholds. Detected oscillatory behavior may trigger local alarms and generate summary metrics including oscillation frequency, magnitude, and duration. Because detection is performed locally, only low-bandwidth summary information needs to be communicated, while optional waveform capture can be retained locally for detailed analysis.

Local analytics engines may compute instantaneous or half-cycle active power from voltage and current waveforms and apply spectral techniques such as FFT or narrow-band filtering to detect oscillatory components. Continuous tracking of oscillation magnitude within selected frequency bands enables detection of both sustained and intermittent oscillatory behavior. These

approaches allow monitoring of oscillations that may not exceed traditional disturbance thresholds and therefore may not be captured using event-triggered recording.

Distributed monitoring architectures also support centralized aggregation of locally detected oscillation metrics (see Figure 23). Summary observations from multiple monitoring locations can be combined to provide system-wide situational awareness, track oscillation activity over time, and identify propagation of oscillatory disturbances (Lackner and Carroll 2025). Because only derived metrics are transmitted, these architectures can operate over existing communication infrastructure without requiring continuous streaming of waveform data.

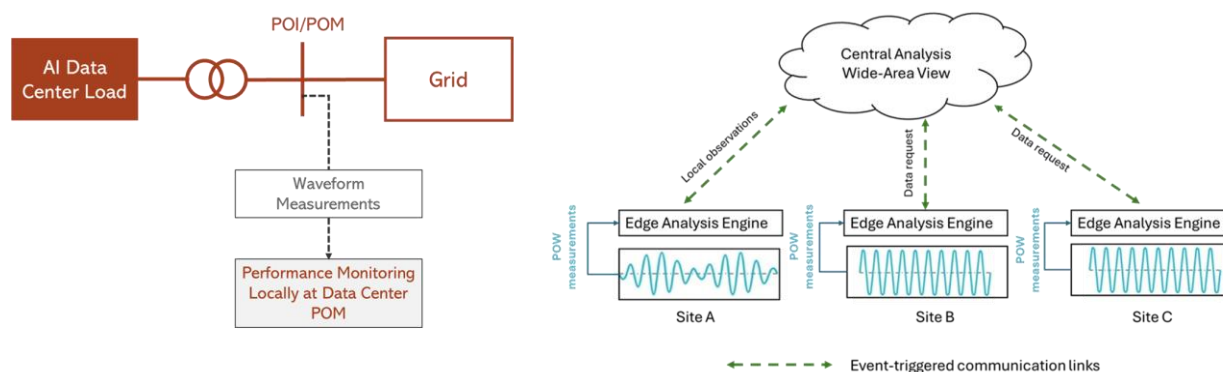


Figure 23. Distributed architecture for data center oscillation monitoring using POW measurements.

Similar monitoring principles may also be applied within data center facilities. High-resolution measurements at the point of interconnection and internal feeder levels can be analyzed locally using detection algorithms designed to be consistent with utility-side monitoring methodologies. This consistency allows compliance assessments performed within the facility to align with those applied at the grid interface. Data center operators can therefore identify oscillatory behavior locally and implement mitigation measures such as workload management or cluster-level controls before oscillations propagate to the transmission system.

By combining local waveform-based analysis with centralized reporting, hybrid monitoring architecture provides a practical pathway for utilizing POW measurements in large-scale monitoring systems. Although these approaches are still evolving and continue to be validated through research and pilot deployments, they offer a promising framework for enabling accurate detection of high-frequency oscillations while mitigating communication and storage requirements associated with continuous centralized streaming of waveform data.

All three schemes described above are technically viable and represent different tradeoffs between visibility, infrastructure requirements, and implementation complexity. Event-triggered recording minimizes communication requirements but provides limited continuous visibility required for compliance monitoring. Continuous centralized streaming offers the most complete observability but requires substantial communication and data management infrastructure. Hybrid and distributed analytics approaches provide a middle ground by enabling continuous local analysis with selective communication of summary results. With ongoing research and product development in this area, utilities and data center developers now have access to practical solutions for monitoring load-induced oscillations using POW measurements.

6.0 Conclusions

This report examined the adequacy of existing measurement systems for monitoring oscillations induced by large data center loads, with particular focus on phasor measurement units (PMUs) and point-on-wave (POW) measurements. The analysis shows that PMUs remain highly effective for monitoring low-frequency electromechanical oscillations but have fundamental limitations when applied to higher-frequency oscillations that may arise from large AI training loads.

The effective measurement bandwidth of PMU-based monitoring is constrained by both reporting-rate limitations and the filtering inherent in phasor estimation algorithms. According to the IEC/IEEE 60255-118-1 standard, the dynamic performance of PMUs is validated primarily for oscillation frequencies up to approximately 2 Hz for P-class devices and up to about 5 Hz for M-class devices. Oscillations beyond this range fall outside the frequency band for which accurate amplitude tracking is required. As a result, PMUs are generally not well suited for monitoring oscillations significantly higher than approximately 5 Hz.

Even when PMUs operate at higher reporting rates, the filtering used in phasor estimation can attenuate inter-harmonic sideband components associated with higher-frequency oscillations. Because oscillations in the phasor domain originate from these sideband components in the waveform spectrum, attenuation of the sidebands leads directly to attenuation of the estimated oscillation magnitude. Analytical derivations, simulations, and hardware-in-the-loop experiments presented in this report demonstrate that oscillations in the sub-synchronous frequency range can experience substantial attenuation in PMU measurements.

This attenuation has important implications for monitoring compliance with emerging utility-proposed limits on oscillatory load behavior. When compliance assessments rely solely on PMU measurements, attenuated oscillation amplitudes may lead to an overly optimistic interpretation of system behavior. In such cases, a critical oscillation could be incorrectly assessed as benign, potentially delaying mitigation actions or obscuring the true dynamic impact of the load. These findings highlight the importance of understanding PMU measurement bandwidth and estimator characteristics when interpreting oscillatory behavior in phasor data.

High-resolution POW measurements provide a complementary capability that is better suited for monitoring higher-frequency oscillations. Because POW measurements directly record the instantaneous voltage and current waveforms, they preserve the full spectral content of the signal and can accurately capture oscillatory behavior across a much broader frequency range. This makes POW-based monitoring particularly valuable for detecting oscillations in the sub-synchronous frequency band that may not be faithfully represented in phasor measurements.

However, widespread deployment of continuous POW monitoring presents practical challenges. High-speed waveform sampling generates extremely large volumes of data, which places significant demands on communication bandwidth, storage infrastructure, and real-time data processing capabilities. Multiple promising solutions and architectures are emerging to address these challenges. These include implementations where POW data is streamed to cloud for centralized analysis. Other solutions include combining local waveform analysis with selective communication summary metrics, alarms, or event data to centralized platform for wide-area situational awareness. These architectures provide a scalable pathway for incorporating POW

measurements into operational monitoring systems while effectively managing the communication and storage burdens associated with continuous waveform streaming.

Overall, the findings of this report indicate that PMU-based monitoring will remain an essential component of wide-area measurement systems for observing low-frequency grid dynamics. At the same time, the limitations of phasor measurements in capturing higher-frequency oscillations highlight the need for complementary waveform-based monitoring capabilities when evaluating oscillatory behavior associated with large AI-driven data center loads. Continued development of hybrid monitoring architectures that integrate PMU and POW measurements will therefore play an important role in ensuring reliable and accurate monitoring of emerging large electric loads in modern power systems.

7.0 References

- AESO. 2025. *AESO Connection Requirements for Transmission-Connected Data Centres*. August. aesoengage.aeso.ca/49634/widgets/209340/documents/157140.
- Aminifar, F. 2025. "Grid Operator Analytics and Assessment Tools for Inverter-Based Resources Dominated Grid (GOAAT-IBR)." https://www.naspi.org/sites/default/files/2025-11/02_11_01_Aminifar_Hart_Quanta_20250924.pdf.
- Arbiter Systems. 2007. *Model 1133A Power Sentinel™*. <https://www.arbiter.com/catalog/product/model-1133a-power-sentinel.php>.
- ATC. 2025. *Transmission System Planning Criteria*. August. <https://cdn.misoenergy.org/ATC%20TO%20Planning%20Criteria108210.pdf>.
- Bestebreuer, Jared. 2025. *Exploring Approaches to Time-Series Streaming and Event Triggers*. www.naspi.org/sites/default/files/2025-11/02_12_01_Bestebreuer_SEL_20250924.pdf.
- Biswas, Shuchismita, Antos C Varghese, Kaustav Chatterjee, Sameer Nekkallapu, Brett Ross, and Jim Follum. 2025. *Evaluating the risk to bulk power system reliability from large load induced oscillations*. <https://www.techrxiv.org/users/956566/articles/1325575>.
- Chatterjee, K, and J Follum. 2025. "Streaming Event-Triggered Point-on-Wave Measurements for Disturbance Monitoring and Analysis." *IEEE Access* 13. <https://ieeexplore.ieee.org/abstract/document/11205344>.
- Chatterjee, Kaustav, Shuchismita Biswas, Jim Follum, and Slaven Kincic. 2023. *Online Monitoring Applications Enabled by Phasor Measurement Units: Technical Assistance to the Power Sectors of Southeast Asia*. Richland: Pacific Northwest National Laboratory.
- Cheung, K, R Cox, and P Sutrave. 2025. *Safeguarding against disruptive events: innovative solutions for sub-synchronous oscillations in AI data centers*. www.eaton.com/content/dam/eaton/markets/data-center/eaton-wp-sub-synchronous-oscillations-event-management-wp150003en-us.pdf.
- Choukse, Esha, B Warriar, S Heath, L Belmont, A Zhao, Hassan A Khan, B Harry, M Kappel, R Hewett, and K Datta. 2025. *Power Stabilization for AI Training Datacenters*. August. <https://arxiv.org/abs/2508.14318>.
- ERCOT. 2026. *Large Load Power Variation Requirement Consideration*. February. <https://www.ercot.com/files/docs/2026/02/19/ERCOT-LEL-SSO-Power-Variation-Consideration.pdf>.
- ESIG. 2026. "Large Load Disturbance Events." www.esig.energy/wp-content/uploads/2026/03/ESIG-Large-Loads-Disturbance-Events-report-2026.pdf.
- ESIG. 2026. *Large Load Interconnection Performance Requirements*. February. <https://www.esig.energy/reports-briefs/large-load-interconnection-performance-requirements/>.
- Follum, J, L Miller, P Etingov, H Kirkham, A Riepnies, X Fan, and E Ellwein. 2021. *Phasors or Waveforms: Considerations for Choosing Measurements to Match Your Application*. Richland: PNNL. <https://www.naspi.org/documents/phasors-or-waveforms-considerations-choosing-measurements-match-your-application>.
- Follum, J, R Hovsapian, N.M Stenvig, Y Agalgaonkar, K Chatterjee, K Mahapatra, A Riepnies, A Wilson, S Chanda, and S Granda. 2023. *Advanced Measurements for Resilient Integration of Inverter-Based Resources*. Richland: PNNL. <https://www.pnnl.gov/publications/advanced-measurements-resilient-integration-inverter-based-resources-progress-matrix>.
- Gravois, Patrick. 2025. *Large Load Oscillation Event*. ERCOT, March.

- Hooshyar, H, A Haddadi, E Farantatos, and M Patel. 2022. "Investigation of PMU Limitations in Monitoring Fast Dynamics Through Real-Time Hardware-In-The-Loop Experiments." *International Conference on Smart Grid Synchronized Measurements and Analytics*. IEEE.
- Hooshyar, H, E Faranatos, and M Patel. 2019. *Impact of PMU Signal Processing on Synchrophasor Measurements, Monitoring and Control Applications*. EPRI.
- IEC/IEEE. 2018. 60255-118-1-2018 - IEEE/IEC International Standard - Measuring relays and protection equipment - Part 118-1: Synchrophasor for power systems - Measurements.
- IEEE. 2024. C37.118.2-2024 - IEEE Standard for Synchrophasor Data Transfer for Power Systems. <https://ieeexplore.ieee.org/document/10794635>.
- IEEE Spectrum. 2026. *The Data Center Boom Is Concentrated in the U.S.* <https://spectrum.ieee.org/data-center-growth>.
- Ko, Min-Seung, and Hao Zhu. 2025. *Wide-area power system oscillations from large-scale AI workloads*. <https://arxiv.org/pdf/2508.16457>.
- Lackner, C, and R Carroll. 2025. *Distributed Waveform Analytics in the WaveApps Platform*. https://powaconference.org/presentations/session_5/Distributed%20Waveform%20Analytics%20in%20the%20Wave%20Apps%20Platform.pdf.
- Mishra, C, L Vanfretti, J Delaree, T.J Purcell, and K Jones. 2025. "Understanding the inception of 14.7 Hz oscillations emerging from a data center." *Sustainable Energy, Grids and Networks*.
- Mohsenian-Rad, H, and W Xu. 2023. "Synchro-Waveforms: A Window to the Future of Power Systems Data Analytics." *IEEE Power and Energy Magazine*. <https://ieeexplore.ieee.org/document/10226356>.
- Mohsenian-Rad, Hamed, Jhi-Young Joo, Michael Balestrieri, Lakshan Piyasinghe, Shuchismita Biswas, Kaustav Chatterjee, Hamed Valizadeh-Haghi, et al. 2024. *Synchro-Waveform Measurements and Data Analytics in Power Systems*. IEEE. <https://iee-synchrowaveform.engr.ucr.edu/sites/default/files/2025-06/pes-tr127.pdf>.
- NERC. 2025. *Characteristics and Risks of Emerging Large Loads*. www.nerc.com/globalassets/who-we-are/standing-committees/rstc/whitepaper-characteristics-and-risks-of-emerging-large-loads.pdf.
- NERC. 2019. *Eastern Interconnection Oscillation Disturbance*. www.nerc.com/globalassets/our-work/reports/event-reports/january_11_oscillation_event_report.pdf.
- NERC. 2025. *Power System Oscillation Monitoring and Mitigation*. www.nerc.com/globalassets/our-work/technical-reference-documents/trd-oscillation_analysis_for_monitoring_and_mitigation---draft.pdf.
- ORNL, LLNL. 2026. "Grid Event Signature Library." <https://gesl.ornl.gov/>.
- PSEG. 2026. *Performance Requirements for Large Loads Connected to the LIPA System*. February. <https://www.psegliny.com/aboutpseglongisland/legalandregulatory>.
- Purcell, T J, F Wai, J Johnston, M Till, and K Jones. 2024. *Continuous Oscillography Waveforms from the Substation to the Cloud*. April. www.naspi.org/sites/default/files/2024-05/D2_S09_P02_Purcell_Dominion_20240417.pdf.
- Saroare, Md K, Md A Hasnat, and Md R Ahmed. 2026. *GridStream: A Hardware-Efficient Framework for Bandwidth-Constrained Point-on-Wave Disturbance Monitoring*. <https://doi.org/10.36227/techrxiv.176834178.83076178/v1>.
- Shehabi, Arman, Alex Newkirk, Sarah J Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. *2024 United States Data Center Energy Usage Report*. Berkeley: Lawrence Berkeley National Laboratory.
- Silverstein, A, and J Follum. 2020. *High-Resolution, Time-Synchronized Grid Monitoring Devices*. NASPI. <https://www.naspi.org/documents/high-resolution-time-synchronized-grid-monitoring-devices>.

- Vizimax. 2026. *Phasor Measurement Unit (PMU) Datasheet: PMU010000-SP-en*.
<https://www.vizimax.com/wp-content/uploads/2022/03/pmu010000-sp-en-20260304-phasor-measurement-unit-datasheet.pdf>.
- Xu, W, J Yong, H Marquez, and C Li. 2025. "Interharmonic Power—A New Concept for Power System Oscillation Source Location." *IEEE Transactions on Power Systems*.
- Xu, W., Z. Huang, X. Xie, and C. Li. 2022. "Synchronized Waveforms – A Frontier of Data-Based Power System and Apparatus Monitoring, Protection, and Control." *IEEE Transactions on Power Delivery* 3-17.
- Zhang, Jimmy, and Jonathan Rose. 2026. *Large Electronic Load Sub-Synchronous Oscillation (LEL-SSO) Power Variation Challenges and Discussion*. ERCOT.
https://www.ercot.com/files/docs/2026/01/21/ERCOT_Electranix-LEL-SSO-Power-Variation-Criteria_LLWG-20260122-post.pdf.

Appendix A – Attenuation of Off-Nominal Frequency Component

This appendix describes how an off-nominal sideband frequency component in the waveform spectrum is attenuated during phasor estimation using the reference algorithm specified in the IEC/IEEE 60255-118-1 standard (IEC/IEEE 2018).

Consider an inter-harmonic sideband component separated from the nominal fundamental frequency f_0 by an offset Δf , as shown in (13).

$$x(t) = X_0 \cdot \cos(2\pi (f_0 + \Delta f) t + \phi) \quad (13)$$

Discretizing this with a sampling frequency f_s yields

$$\begin{aligned} x[n] &= X_0 \cdot \cos\left(\frac{2\pi (f_0 + \Delta f)}{f_s} n + \phi\right) \\ &= \frac{X_0}{2} \left(e^{j(\Omega_0 + \Delta\Omega)n + \phi} + e^{-j(\Omega_0 + \Delta\Omega)n + \phi} \right) \end{aligned} \quad (14)$$

where, $\Omega_0 = \frac{2\pi f_0}{f_s}$ and $\Delta\Omega = \frac{2\pi \Delta f}{f_s}$.

The phasor estimation of this component using the reference algorithm is expressed as,

$$\begin{aligned} \tilde{X}[i] &= \frac{\sqrt{2}}{G} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} x[i+k] \cdot W[k] \cdot e^{-j(i+k)\frac{2\pi f_0}{f_s}} \\ &= \frac{\sqrt{2}}{G} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} \frac{X_0}{2} \left(e^{j(\Omega_0 + \Delta\Omega)(i+k) + \phi} + e^{-j(\Omega_0 + \Delta\Omega)(i+k) + \phi} \right) \cdot W[k] \cdot e^{-j(i+k)\Omega_0} \\ &= \frac{X_0}{\sqrt{2}} e^{j\phi} e^{j\Delta\Omega i} \frac{1}{G} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k] e^{j\Delta\Omega k} + \frac{X_0}{\sqrt{2}} e^{-j((2\Omega_0 + \Delta\Omega)i + \phi)} \frac{1}{G} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k] e^{j(2\Omega_0 + \Delta\Omega)k} \end{aligned} \quad (15)$$

where,

$$G = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k] \quad (16)$$

is a normalization constant and $W[k]$ is the window corresponding to the estimation filter.

Define

$$\mathcal{W}(f) = \frac{1}{G} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}} W[k] e^{-j \frac{2\pi f}{f_s} k} \quad (17)$$

as the discrete-time Fourier transform (DTFT) of the filter $W[k]$ normalized by the constant G . This quantity corresponds to the frequency response of the estimation filter.

Following this, (15) can be expressed as

$$\tilde{X}[i] = \frac{X_0}{\sqrt{2}} e^{j\phi} e^{j\Delta\Omega i} \mathcal{W}(\Delta f) + \frac{X_0}{\sqrt{2}} e^{-j((2\Omega_0 + \Delta\Omega)i + \phi)} \mathcal{W}(\Delta f + 2f_0). \quad (18)$$

From the magnitude response plots of the reference filters in Figure 8 and Figure 9, it is evident that

$$\mathcal{W}(\Delta f) \gg \mathcal{W}(\Delta f + 2f_0). \quad (19)$$

Therefore, the estimated phasor may be approximated as

$$\tilde{X}[i] \approx \frac{X_0}{\sqrt{2}} e^{j\phi} e^{j\Delta\Omega i} \mathcal{W}(\Delta f), \quad (20)$$

where the magnitude of the estimated phasor is

$$|\tilde{X}| = \frac{X_0}{\sqrt{2}} \mathcal{W}(\Delta f). \quad (21)$$

Observe that, for the signal model in (13), the true magnitude of the sinusoid is

$$|X^{true}| = \frac{X_0}{\sqrt{2}}. \quad (22)$$

Phasor estimation thus attenuates the magnitude of the sinusoid by

$$\frac{|\tilde{X}|}{|X^{true}|} = \mathcal{W}(\Delta f). \quad (23)$$

Therefore, the filtering inherent in phasor estimation attenuates the magnitude of any off-nominal frequency component at $f_0 + \Delta f$ by a factor equal to the magnitude response of the filter evaluated at the offset frequency Δf .



NASPI

A DOE-EPRI Joint Initiative

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99354

1-888-375-PNNL (7665)

www.pnnl.gov

EPRI

ELECTRIC POWER
RESEARCH INSTITUTE



Pacific Northwest
NATIONAL LABORATORY



U.S. DEPARTMENT
of **ENERGY**