Pacific Northwest
NATIONAL LABORATORY

PNNL-39487

Cloud-Based Demonstration of the Eastern Interconnection Situational Awareness Monitoring System (ESAMS)

June 2026

Jim Follum, PNNL
Xiaochuan Luo, ISO New England
Yang Chen, PJM
Neeraj Nayak, Electric Power Group
Joshua Chynoweth, Electric Power Group
Song Zhang, formerly Amazon Web Services
Tawsif Ahmad, PNNL
Shuchismita Biswas, PNNL

U.S. DEPARTMENT of ENERGY

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

DISCLAIMER

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
www.osti.gov
ph: (865) 576-8401
fox: (865) 576-5728
email: reports@osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
or (703) 605-6000
email: info@ntis.gov
Online ordering: http://www.ntis.gov

PNNL-39487

Cloud-Based Demonstration of the Eastern Interconnection Situational Awareness Monitoring System (ESAMS)

June 2026

Jim Follum, PNNL
Xiaochuan Luo, ISO New England
Yang Chen, PJM
Neeraj Nayak, Electric Power Group
Joshua Chynoweth, Electric Power Group
Song Zhang, formerly Amazon Web Services
Tawsif Ahmad, PNNL
Shuchismita Biswas, PNNL

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

Pacific Northwest National Laboratory
Richland, Washington 99354

Abstract

This report describes a cloud-based implementation and field demonstration of the Eastern Interconnection Situational Awareness and Monitoring System (ESAMS). ESAMS was developed to support the detection and source localization of forced oscillations using synchrophasor measurements from tie-lines connecting areas served by different reliability coordinators (RCs), so that RCs could better coordinate their response to wide-area events. A previous effort had identified deployment barriers associated with hosting shared situational awareness tools at a single RC. To address these barriers, ESAMS was migrated to Amazon Web Services (AWS) and evaluated in a six-month field demonstration. The Independent System Operator of New England (ISO-NE) and PJM Interconnection (PJM) streamed data to the platform using AWS Direct Connect and a site-to-site VPN, respectively. The resulting multi-utility measurement footprint enabled regional source localization across major portions of the U.S. Eastern Interconnection and supported routine identification of oscillation events. During the final three months of the trial, 24 events above 2 MW/MVAR were detected. The largest detected oscillation approached a 25 MW peak-to-peak amplitude, and the longest persisted intermittently for more than 11 hours. The demonstration also assessed operational considerations—including data transfer volumes, end-to-end latency, and cloud computing costs—and found that network and compute requirements were modest relative to typical cloud capabilities while providing performance comparable to prior on-premises deployments. Overall, the results indicate that cloud hosting can provide a practical path to shared interconnection-wide oscillation monitoring. The cloud ESAMS demonstration establishes a foundation for broader utility participation and for building future wide-area analytics that leverage measurements across organizational boundaries.

Acknowledgments

The authors gratefully acknowledge Sandra Jenkins, Director of Grid Controls within DOE's Office of Electricity, for her support of this project.

The authors are also grateful to the individuals that supported the field demonstration, especially Wenpeng Yu, Frankie Zhang, Slava Maslennikov, Davin Flatten, and Ryan Sawyer of ISO-NE and Nikki Militello and Davean Hosang of PJM. We also appreciate Tim Fritch and Jonathan Sides of TVA for their insights.

The language in this document was edited using ChatGPT (version 5.2, OpenAI) to improve grammar and clarity. All final content was reviewed for accuracy.

Acronyms and Abbreviations

AWS	Amazon Web Services
ESAMS	Eastern Interconnection Situational Awareness and Monitoring System
ISO	Independent System Operator
ISO-NE	ISO New England
NLB	Network Load Balancer
PJM	PJM Interconnection
PDC	Phasor Data Concentrator
PMU	Phasor Measurement Unit
PNNL	Pacific Northwest National Laboratory
RC	Reliability Coordinator
RDS	Relational Database Service
RTDMS	Real Time Dynamics Monitoring System
VPN	Virtual Private Network
VPC	Virtual Private Cloud

Contents

Figures

Tables

1.0 Introduction

Bulk electric power systems are massive networks of generation and transmission equipment that are operated through careful coordination among many organizations. Because these systems are interconnected, events occurring in one part of the grid can impact geographically distant regions. For example, repetitive fluctuations in the power generated or consumed by equipment can cause oscillations in the power system's frequency and power flows across an entire interconnection. These events are termed *forced oscillations* because a source is forcing the system to respond. Forced oscillations that persist at large enough amplitudes can cause reliability concerns, leading the North American Electric Reliability Corporation (NERC) to develop recommendations for oscillation monitoring and mitigation (NERC 2021).

Wide-area oscillation monitoring is enabled by synchrophasor measurements recorded by networks of phasor measurement units (PMUs). Each PMU reports voltage and current phasors (magnitude and angle), frequency, and rate-of-change-of-frequency (ROCOF). The measurements are reported at a high rate, typically 30-60 times per second in 60 Hz grids, and are synchronized to a common clock, typically using signals from a global navigation satellite system (GNSS). When collected from across a wide area, synchrophasors provide higher resolution observability of wide-area dynamics than traditional Supervisory Control and Data Acquisition (SCADA) systems.

However, streaming PMU measurements from all monitored locations in an interconnection to every reliability coordinator (RC) is generally impractical. Instead, RCs rely primarily on measurements from within their own footprint. Measurements are also shared among organizations, but providing each grid operator with a consistent interconnection-wide view would require use of the same monitoring tools with identical configurations. This would be extremely difficult to achieve in practice, especially because organizations use a variety of commercial and custom tools to monitor the grid.

The need for a common interconnection-scale view was clearly felt during the January 11, 2019, oscillation event in the U.S. Eastern Interconnection (EI). During the event, power production from a plant in Florida fluctuated every 4 seconds, resulting in 0.25 Hz oscillations that were seen across the interconnection. Though transmission system operators were aware of the oscillation, its source was not identified until post-event analysis was completed. The oscillation persisted for 18 minutes until the plant's operators took the generator offline, unaware of the oscillation's wide-area impact. Based on these findings, NERC's report on the event recommends, "RCs should consider jointly developing interconnection-wide oscillation detection and source location applications using interconnection-wide PMU and SCADA data" (NERC 2019).

To address this need, Lawrence Berkeley National Laboratory (LBNL), Pacific Northwest National Laboratory (PNNL), and Electric Power Group (EPG) worked with several RCs to develop and demonstrate the Eastern Interconnection Situational Awareness and Monitoring System (ESAMS) (Eto, et al. 2022). The tool included several analytical methods, including the ability to detect wide-area oscillations and determine their region of origination. The initial 10-month demonstration concluded in 2022 and was hosted by PJM Interconnection (PJM), with measurements streaming from six additional RCs. The demonstration was a success, but requiring a single organization to host a tool intended to provide all participating RCs with a common view of the interconnection presented a barrier to long-term deployment.

The focus of this report is an effort to address this barrier by transitioning ESAMS to a cloud environment. After deploying ESAMS to an Amazon Web Services (AWS) cloud environment, measurements from the Independent System Operator of New England (ISO-NE) and PJM were streamed to the platform to support a 6-month demonstration. During this period, many oscillations were detected and their source regions identified. Network performance, data quality, and computing costs were also evaluated. The findings from the demonstration indicate that Cloud ESAMS provides similar performance to an on-premises approach while offering several benefits. As a result, a long-term cloud-based implementation of ESAMS is under consideration. Historically, the vast majority of PMU applications have been deployed on premises, so the lessons learned from the cloud deployment of ESAMS have implications beyond oscillation monitoring.

The rest of the report is organized as follows. Chapter 2.0 provides background information on the ESAMS software and its previous deployments. The cloud-based implementation is then described in Chapter 3.0. Network performance results are presented in Chapter 4.0, followed by a review of detected oscillations in Chapter 5.0. The report's conclusions are presented in Chapter 6.0.

2.0 Background

The ESAMS project began with years of software development, followed by two on-premises demonstrations before the most recent cloud-based demonstration. These key aspects of the project's history are summarized in the following subsections.

2.1 ESAMS Software

ESAMS is implemented as an integrated platform that manages real-time data ingestion and validation, executes automated analytics for forced oscillation detection and regional source localization, and disseminates results through reports and event-driven notifications. The subsections below describe the overall software architecture, the forced oscillation analysis application, and the data quality and availability monitoring functions that ensure meaningful results.

2.1.1 Architecture

ESAMS is a modular software platform that ingests synchrophasor data streams, performs automated data validation and conditioning, executes application analytics, and delivers results to users through a combination of periodic reporting and event-driven notifications. The platform is organized around a set of coordinated components: (a) data acquisition interfaces, (b) data quality management functions, (c) an application/analytics layer, (d) a shared data store, and (e) a reporting/notification capability. These components together provide an end-to-end path from raw measurements to actionable information on abnormal system conditions.

At the front end, ESAMS receives synchrophasor measurements in standard streaming formats (e.g., IEEE C37.118/C37.118.2). The received measurements are assembled and time-aligned and then passed through online data validation procedures before they can be used by any downstream application. These procedures are designed to detect common classes of problems encountered in operational synchrophasor deployments, including communication interruptions, message-level format or integrity errors, timestamp/order issues, problems indicated by PMU status flags, and signal values outside feasible ranges. Data that fails validation is treated as suspect and excluded from application processing so that analytics results are based on validated measurements.

Following validation, ESAMS performs additional data handling and processing to prepare signals for analysis using EPG's Real Time Dynamics Monitoring System (RTDMS) (Agarwal, et al. 2011). This includes creating derived quantities commonly required by wide-area monitoring applications (e.g., power-flow and angle-difference signals), applying conditioning steps needed for stable automated analysis, and implementing specialized checks and handling procedures developed specifically to support ESAMS applications. These functions provide consistency in the signals used by the analytics and reduce the likelihood that application outputs are driven by data artifacts rather than physical system behavior.

Validated and conditioned measurements, along with derived signals and key intermediate results, are stored in an ESAMS database. The database serves as the primary integration point between the data acquisition/conditioning layer and the application analytics. Analysis engines—implemented as independent modules—retrieve the required signals as they become available, execute the relevant detection/localization logic, and then write both summary outputs and supporting metadata (e.g., timestamps, estimated parameters, and quality/confidence information

when applicable) back to the database. This database-centered architecture supports near-real-time processing while also enabling offline follow-up, trending, and retrospective review using a consistent set of stored inputs and outputs.

On the output side, ESAMS includes a reporting and notification capability that converts application results into operator-facing products. The platform is designed to support both routine reports (e.g., daily summaries of detected conditions and data status) and near-real-time notifications for major events. In general, ESAMS uses automated screening and filtering so that notifications emphasize conditions of operational significance, while additional details are retained for reports and subsequent engineering review. This separation between analytics, persistence, and dissemination allows ESAMS to incorporate additional applications over time without requiring fundamental changes to the ingestion, data quality, or reporting backbone.

2.1.2 Oscillation Detection and Source Localization

ESAMS includes an automated capability to detect forced oscillations in synchrophasor measurements and to identify, at a regional level, the area of the interconnection most likely to contain the oscillation source. The intent is to support timely, coordinated response when an oscillation is detected—by providing a consistent interconnection-wide picture of oscillatory behavior, its estimated frequency and amplitude, where it is observed most strongly, and which region is most likely injecting dissipating energy into the rest of the system.

Forced oscillation detection operates continuously on incoming synchrophasor data using a periodogram-based detector (Follum and Pierre 2016, Follum and Tuffner 2016) to identify sustained, narrowband oscillations within a defined frequency range. For each event, ESAMS estimates the oscillation's frequency and amplitude. It then applies a configurable screening logic (e.g., amplitude thresholds and optional persistence/confirmation requirements) to determine which events warrant real-time notification versus inclusion only in daily reporting.

After detection, ESAMS performs regional source localization using the dissipating energy flow (DEF) method (Follum 2020) applied at region boundaries. DEF is based on the principle that a forced-oscillation source injects energy into the network (positive DEF), while other components and areas predominantly absorb that energy (negative DEF). Because complete observability is generally not available at the interconnection level, the source localization module in ESAMS computes DEF primarily at inter-regional tie-lines and interprets them as indicators of the direction of oscillation-energy transfer (Follum, Yin and Betzsold 2020). The source region is determined using a "max net" criterion: the region with the largest net export of dissipating energy across its monitored boundaries is identified as the source (Follum, Yin and Betzsold 2020). Several examples of the localization method's operation are provided in Chapter 5.0.

For each qualifying forced oscillation, ESAMS produces outputs suitable for both situational awareness and follow-up investigation, including: (i) event characterization (time window, estimated frequency, and amplitude), (ii) a list of locations exhibiting the highest oscillation amplitudes, (iii) the most likely source region, and (iv) supporting context such as tie-line-level DEF results. These outputs are disseminated through daily reports and real-time notifications for large events. Higher oscillation amplitude thresholds and longer observation/confirmation intervals may be used to reduce nuisance alarms.

Because DEF-based localization depends on tie-line observability, incomplete PMU coverage on inter-regional boundaries and temporary data unavailability can limit what ESAMS can observe and can make results inconclusive or, in some cases, misleading. To support appropriate

interpretation, ESAMS assigns a confidence score to the identified source region based on factors such as data availability on tie-lines, observed inter-regional energy flow patterns, and robustness of the finding, with notification-level scores summarized as Low/Medium/High (Biswas, Follum and Eto 2023).

2.1.3 Data Quality and Availability Monitoring

Because ESAMS analytics rely on continuous streams of time-synchronized PMU measurements, the application includes data quality and availability monitoring functions to identify suspect data and to summarize overall measurement health. These functions serve three primary purposes: (1) to prevent data artifacts from driving false detections or misleading forced-oscillation localization results by excluding suspect measurements from application processing, (2) to provide transparency to users regarding which PMUs and time intervals were usable versus unavailable or invalid, and (3) to inform data owners and participating entities of data quality or availability problems so they can investigate root causes and restore or improve measurement performance.

ESAMS applies a standard set of online data-checking procedures to measurement streams before they are used by downstream applications. These checks include communication-related checks (confirming data is flowing and that interfaces are not detecting transfer errors), message-level error detection (confirming conformance to expected formats and integrity), time-related validation (confirming correct ordering by timestamp), PMU status validation (confirming that status flags indicate valid, synchronized measurements), and signal-level validation (confirming measurements fall within feasible ranges). Measurements that fail these checks are flagged as suspect and are ignored by downstream applications, ensuring that results are not affected by data artifacts.

To support operational awareness and retrospective review, ESAMS produces PMU device-level data quality reports that summarize both the amount of usable data and the reasons data were unusable. In the demonstration configuration, ESAMS categorized data using PMU status flags (e.g., good samples, dropout, invalid, PMU error, synchronization error, and time error) and presented results in two complementary views: (1) a daily availability summary showing the fraction of the day in each category, and (2) an availability chronology showing when within the day each category occurred. These views help distinguish persistent issues (e.g., a PMU that is largely unavailable) from intermittent problems (e.g., repeated dropouts that may indicate communications or time-synchronization issues).

At the interconnection scale, the ESAMS demonstration experience has shown that data quality issues occur routinely even when overall availability is high. Thus, the platform is designed to explicitly account for data limitations in both analysis and interpretation. As discussed in Section 2.1.2, data availability and quality can impact source localization result accuracy, and hence data availability metrics are incorporated into application-specific confidence assessments to help users gauge the reliability of results when key measurements are missing or flagged as suspect.

2.2 Summary of Past Deployments

Prior to the cloud-based demonstration discussed in this report, two on-premises ESAMS demonstrations were performed. As discussed in the following subsections, each of these demonstrations provided valuable insights into the unique capabilities of ESAMS and the shortcomings that needed to be addressed for long-term deployment.

2.2.1 Initial Field Demonstration

The ESAMS project began in 2017 and conducted its initial field demonstration between June 2021 and March 2022, with PJM serving as the host utility (Eto, et al. 2022). The demonstration was intended to evaluate whether interconnection-scale synchrophasor data sharing and automated analytics could be operated continuously in near real time and produce actionable, operator-relevant products. In addition to assessing analytic performance, the demonstration emphasized the practical elements required for sustained operations, including establishing data-sharing agreements, managing data quality, and delivering results to participants in a form that supports operational awareness and follow-up investigation.

The project start-up consisted of two parallel efforts. First, PJM led the negotiation of data-sharing agreements with participating utilities and the research team. Agreements were completed in two stages: an initial stage with ISO-NE, Midcontinent Independent System Operator (MISO), New York Independent System Operator (NYISO), and PJM, followed by a second stage that added Southern Company, Southwest Power Pool (SPP), and Tennessee Valley Authority (TVA). Second, the project team developed and tested the ESAMS software and data systems needed for continuous processing of streaming synchrophasor measurements. Although the underlying algorithms were based on prior research, significant customization was required to support uninterrupted, real-time operation. Field testing of the core application algorithms began in June 2020 in the PJM hosting environment.

ESAMS was developed as a wide-area situational awareness platform intended to produce reports of abnormal system conditions using streaming PMU data. The demonstration evaluated multiple analytics, reflecting both earlier research priorities and operational feedback from participating entities. These included:

- **Forced oscillation detection and source localization:** Continuous detection of sustained, narrowband forced oscillations, with estimation of oscillation frequency, amplitude, and source region.
- **Voltage angle pair monitoring:** Identification and tracking of wide-area phase angle differences to indicate wide-area stress and potential precursors to severe system conditions.
- **Anomaly detection:** Automated detection of unusual or atypical combinations of measurements that may indicate abnormal system behavior.
- **Mode metering:** Continuous tracking of key inter-area electromechanical modes of oscillation using the mode meter in RTDMS, producing updated estimates of mode frequency and damping ratio every 15 seconds.
- **Ringdown oscillation analysis:** Recording of transient oscillations following large grid events to support subsequent manual analysis to identify excited modes and estimate damping, supporting corroboration and calibration of mode-meter findings.

Across the portfolio of analytics explored, the forced oscillation analysis capability was recognized as particularly impactful. This emphasis was driven by the fact that forced oscillations can propagate over large geographic footprints, be observed strongly far from their source, and require rapid cross-entity coordination to identify the responsible area and mitigate the driving mechanism. To strengthen this capability, the project incorporated additional synchrophasor data

streams to improve coverage for oscillation monitoring and later enabled real-time notifications for major forced oscillations. Compared to routine summary reporting, these forced-oscillation notifications were more actionable because they are delivered while the oscillation is ongoing and include preliminary information needed to initiate coordinated investigation (e.g., likely source region, oscillation frequency, and the highest-amplitude locations). This notification capability was intended to help RCs quickly determine lead responsibility and reduce potentially conflicting interpretations that can arise from independent local observations.

The field demonstration established several notable successes:

- **Feasibility of interconnection-scale synchrophasor data sharing for real-time analytics.**
  The project demonstrated that multiple organizations could stream synchrophasor data to a common host environment and use those measurements to generate shared analytic outputs. Participating utilities did not only provide data; they also helped identify priorities, supplied information on PMU configurations, followed up on data issues, and conducted independent analyses to corroborate and extend ESAMS findings.

- **Translation of research algorithms into continuous, operational processing.**
  ESAMS demonstrated that analytics originally developed through offline research could be customized for continuous real-time processing. The platform integrated data ingestion, data quality checks, database storage, analytic execution, and routine dissemination of results.

- **Demonstrated value of regular dissemination and information-sharing.**
  A key outcome of the project was the routine sharing of findings—both among participants and with broader industry groups—supporting increased awareness of synchrophasor-based monitoring opportunities and practical implementation considerations.

The demonstration also surfaced challenges that shaped priorities for subsequent work:

- **Data quality management as a central operational requirement.**
  Significant effort was required to address data quality issues encountered in real-time streaming environments. The need to delay the start of the demonstration to develop improved methods underscored that data quality and availability are not ancillary concerns; they are prerequisites for reliable automated analytics and for maintaining user confidence in results.

- **Customization and integration effort for sustained real-time operation.**
  Even when research-grade algorithms exist, substantial engineering and integration effort is needed to translate them to production-grade software that can run continuously in operation environments, manage edge cases, and produce stable outputs suitable for regular reporting and notification.

- **Inter-organizational coordination and operationalization.**
  The demonstration highlighted that realizing the full value of wide-area awareness requires not only analytics, but also sustained processes for communication, interpretation, and action across organizational boundaries.

The demonstration concluded that synchrophasors can provide real-time, wide-area situational awareness of grid phenomena, while also recognizing that further work is needed to integrate that

awareness into routine operations. As the initial demonstration concluded, the expected next steps included continued refinement and operationalization of the most impactful analytics—particularly forced oscillation monitoring—along with further maturation of data quality practices, reporting/notification workflows, and sustained participant engagement. The overall trajectory anticipated building from the initial demonstration toward a capability that could be maintained over the long term and more directly integrated into day-to-day utility operational processes.

2.2.2 Southern Company Field Demonstration

The Southern Company field demonstration (2024-25) focused on how ESAMS's oscillation monitoring capabilities could be applied within an RC's territory and across key operating boundaries. A central theme of the deployment was the complementary nature of the two oscillation detection approaches available through ESAMS: RTDMS, which can also be deployed as a standalone tool, and the periodogram detector in PNNL's oscillation analysis engine. Rather than treating these tools as competing approaches, the demonstration used the primary ESAMS periodogram detector to baseline and better understand the performance of the RMS-energy detector (Donnelly, et al. 2015) in RTDMS.

The two algorithms are designed for different but overlapping monitoring needs. RTDMS is oriented toward analysis within an operating area and is especially effective at detecting large oscillations that may pose reliability concerns, including oscillations that are not purely sinusoidal. It does not require oscillation frequency or amplitude to be time invariant and identifies the general range of an oscillation's frequency. By contrast, PNNL's periodogram detector is intended to provide a common view across operating areas and performs best for sustained, narrowband oscillations with relatively stable frequency and amplitude. It estimates a specific oscillation frequency, uses automatically determined dynamic thresholds, and can detect smaller sinusoidal oscillations that are useful for system monitoring, setup, and testing. Detecting small oscillation can also help diagnose persistent equipment configuration issues before they become major reliability concerns. Together, these capabilities provide broader value than either tool would provide independently.

During the deployment, ESAMS was configured to monitor Southern Company's borders with neighboring systems and boundaries between three Southern Company transmission service providers. This configuration allowed the project team to evaluate how oscillation notifications could support both wide-area awareness and more localized operational decision-making. It also provided a practical setting for comparing results from RTDMS's RMS-energy detector with those from PNNL's oscillation analysis engine under real operating conditions.

One important outcome was improving the process of setting oscillation detection thresholds in RTDMS. Prior to this work, default RTDMS thresholds produced numerous false alarms in Southern Company. To support a meaningful comparison between the detectors, the project developed new methods for setting thresholds in the RMS-energy detector (Follum, Biswas, et al. 2025), and Southern Company deployed the updated thresholds in RTDMS. The demonstration also included detailed reviews of 21 oscillation events, which supported the thresholding work by providing real operating examples for evaluating detector performance, understanding alarm behavior, and improving the practical usefulness of RTDMS notifications.

In addition, PNNL developed operator guidance to address a need identified during the initial ESAMS field demonstration: ensuring that oscillation notifications are actionable for participants with different monitoring capabilities. The guidance described recommended actions for both the suspected source area and other affected reliability coordinators, with recommendations tailored

to the tools available to each entity, including oscillation source localization tools, wide-area monitoring systems, or SCADA-only visibility (Biswas and Follum 2024). Overall, the Southern Company deployment showed that combining RTDMS and PNNL's periodogram detector within ESAMS strengthens oscillation monitoring by pairing robust local detection with interconnection-scale situational awareness.

3.0 Cloud ESAMS

Following the initial ESAMS field demonstrations, it became clear that sustaining an interconnection-scale monitoring capability over the long term presents challenges that are not limited to the analytics themselves. ESAMS was intentionally designed to complement—rather than replace—each operator's existing SCADA, Energy Management System (EMS), and local synchrophasor tools by providing advisory, wide-area information that a single entity cannot obtain from within its own footprint. The demonstration showed that this added interconnection-wide perspective is valuable, particularly for forced oscillations where timely coordination among RCs can be important. It also showed that a prototype system can deliver standardized analysis products and provide oscillation-event notifications in near real time, supporting more efficient communication across organizational boundaries.

At the same time, the demonstration highlighted a practical barrier to transitioning ESAMS from a time-bounded project to an enduring capability: long-term hosting at a single organization can create an outsized operational burden and can complicate expansion to additional participants. A single host must provision and maintain the computing environment, operate and secure the software stack, manage backups and patching, monitor system health, and support onboarding and troubleshooting for a diverse set of data providers. In an interconnection-wide context, that host is also responsible for keeping pace with changing data feeds, evolving network configurations, and shifting operational priorities. Even when the analytics are stable, the surrounding "platform work" required to keep a continuously running real-time system healthy and trustworthy is substantial. During the demonstration, considerable effort was devoted to making the system robust to real-world issues such as data quality problems, intermittent outages, and the practicalities of delivering results as routine reports and event-driven notifications. These lessons reinforced that any future deployment would need to minimize operational friction while remaining scalable as more measurements and participants are added.

In considering a path forward, the project team identified cloud hosting as a potential approach to address these platform and sustainability challenges. The motivation to evaluate a cloud-based deployment was not driven by a desire to change what ESAMS does, but rather by the need to find a deployment model that could make ESAMS easier to operate, maintain, and scale. The potential advantages of cloud-hosting are discussed in the following subsection. The rest of the chapter then explains how Cloud ESAMS was configured for the field demonstration and the network design that provided measurements to the system.

3.1 Anticipated Benefits of a Cloud-Based ESAMS Deployment

In reviewing findings from the on-premises demonstrations, the project team identified several potential benefits of transitioning to a cloud-based implementation, as summarized next.

- **Elastic scalability to match interconnection monitoring needs:** A primary motivation for exploring cloud deployment is the elastic scalability that cloud infrastructure can provide. Interconnection-scale monitoring leveraging data across organizations is inherently subject to growth and variability in both data volume and computational needs. As additional utilities participate, the number of monitored tie-lines and PMU signal inputs increases. In a traditional on-premises hosting model, accommodating growth often requires advance capacity planning, procurement cycles, and periodic hardware refreshes. This can make it difficult to expand coverage quickly or to experiment with alternative configurations (for example,

enabling additional monitored signals or widening the set of boundary tie-lines used for localization) without creating concern that the platform will be under-provisioned.

In contrast, cloud deployment provides a way to size infrastructure to current needs while retaining the ability to scale as participation expands or as additional analytics are incorporated. This is valuable to interconnection-scale monitoring where elastic scaling reduces the friction associated with adding new PMUs, new analytic capabilities, and onboarding new participating entities.

- **Reduced operational overhead:** A second major motivation is reducing the operational burden associated with platform maintenance. Field demonstrations have shown that running ESAMS as a continuous near-real-time service requires round-the-clock attention to the health of the data pipeline, the availability of storage and compute resources, and the reliability of reporting and notification mechanisms. In addition, any long-running operational system must address routine—but time-consuming—tasks such as server maintenance, operating system and application patching, vulnerability management, certificate management, automated backups, log retention, monitoring/alerting, and periodic upgrades.

  Cloud hosting can allow many of these functions to be implemented using managed services and standardized operational practices, rather than custom one-off solutions maintained by a single host organization. This “offloading” is not about avoiding responsibility for reliability or security; rather, it is about reducing the amount of effort that must be devoted to infrastructure fundamentals so that the participating utilities can focus on higher-value domain-specific tasks such as root cause analysis of detected oscillations and implementing mitigation actions. Minimizing the day-to-day platform maintenance burden is an important enabler for sustained use and long-term deployment.

- **Resilient deployment and recovery for a continuously running service:** Cloud deployment is expected to result in improved resilience in deployment practices and in recovery from failures. The success of synchrophasor-based monitoring platforms relies heavily on continuous availability. If a critical event goes undetected due to platform component failures – user confidence in the tool itself can be significantly eroded. As a result, ESAMS benefits from a cloud deployment approach that supports high availability, controlled upgrades, and rapid recovery from component failures.

  Cloud environments commonly support deployment patterns such as blue/green releases, where a new version of a service can be deployed in parallel with the existing version, validated, and then switched into production with the ability to roll back quickly if needed. For ESAMS, this is valuable to accommodate the addition of monitored signals and refinements to the underlying algorithms. Being able to introduce these changes without extended downtime—and with a clear rollback option—reduces operational risk and encourages continuous improvement rather than infrequent, high-risk upgrades.

  In addition, cloud infrastructure can be designed for redundancy across multiple availability zones, enabling rapid recovery from localized failures without requiring a single organization to maintain duplicate physical infrastructure. This aligns with the operational expectation that wide-area monitoring systems should remain available during stressed conditions, including conditions that may coincide with storms or other disruptions that can affect a single data center.

- **Predictable, usage-based cost structure:** Another motivation is the ability to better align costs with actual usage. Interconnection-wide monitoring systems can incur costs in proportion to the number of participating PMUs, the amount of retained data, the computational intensity of analytics, and the frequency and richness of reporting. In an on-premises model, cost is often dominated by fixed infrastructure investments and staffing commitments that must be sized for peak needs and maintained even when usage is lower. This can complicate budgeting and can create barriers to incremental growth.

  Cloud deployment, on the other hand, offers a model where infrastructure costs scale with consumption. This does not eliminate the need for cost discipline—cloud resources can be mis-sized just as on-premises infrastructure can—but it enables more transparent accounting of what drives cost (data volume, retention, compute cycles, network egress, etc.). For a multi-entity effort, this transparency can be helpful in planning long-term sustainment and in evaluating the marginal cost of adding additional PMU streams, analytics capabilities, or extending data retention for after-action review.

  Security and compliance considerations for sensitive grid information: The ESAMS application is intended to complement utilities' internal monitoring capabilities by providing a shared interconnection-wide view. Thus, for long-term deployment, the participants must be confident that shared data and derived results are protected appropriately, access is controlled, and auditing is possible. The mature security controls and compliance frameworks in modern enterprise cloud environments can offer this confidence, and hence are well-suited alternatives for on-premises deployments.

  Cloud platforms provide a broad set of security mechanisms that can be applied systematically: identity and access management, network segmentation, encryption at rest and in transit, centralized logging and monitoring, key management, and security posture management tooling. They also provide documented processes and certifications that organizations can map to internal security requirements and applicable reliability-sector guidance. For a multi-entity deployment, these standardized security capabilities can be an enabler because they reduce the need to design bespoke controls at a single host site and can support consistent security baselines as the system evolves.

The objective of the cloud ESAMS demonstration was to evaluate if this deployment model could realize these benefits and offer comparable performance to an on-premises alternative, minimizing the infrastructure maintenance burden on participating organizations.

3.2 Cloud ESAMS Configuration

Using the networks described in Section 3.3, synchrophasor measurements were streamed from ISO-NE and PJM to the cloud-based ESAMS deployment. Oscillation detection was performed on active and reactive power measurements from 67 transmission lines monitored from 20 substations. Once an oscillation was detected, regional source localization was performed based on measurements from 19 substations covering 22 tie lines between regions. As depicted in Figure 1, tie-line measurements from ISO-NE provided observability of dissipating energy flow to neighboring utilities New Brunswick System Operator (NBSO) and NYISO. Measurements from PJM provided observability of PJM's borders with NYISO and the rest of the EI south and west of PJM. Though the configuration involved only a modest number of measurements, the results presented in Chapter 5.0 indicate that it provided reliable oscillation detection and source localization.

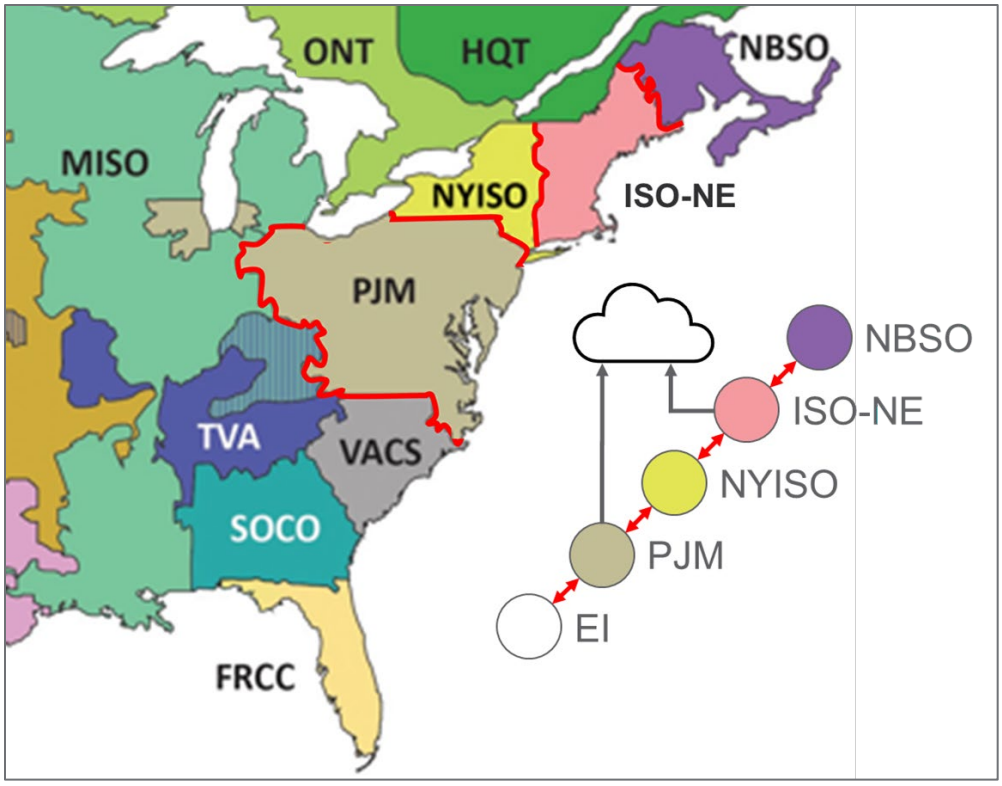

Figure 1. Summary of the Cloud ESAMS configuration, with measurements from PJM and ISO-NE streaming to the cloud and providing visibility of dissipating energy flow between five regions.

3.3 Network Design

The ESAMS cloud network architecture illustrated in Figure 2 is designed to provide secure, scalable, and highly available connectivity between geographically distributed Phasor Data Concentrators (PDCs) and the centralized ESAMS application environment hosted on AWS. The design prioritizes network isolation, encrypted transit, and deterministic traffic routing to meet the operational and security requirements of the ESAMS program.

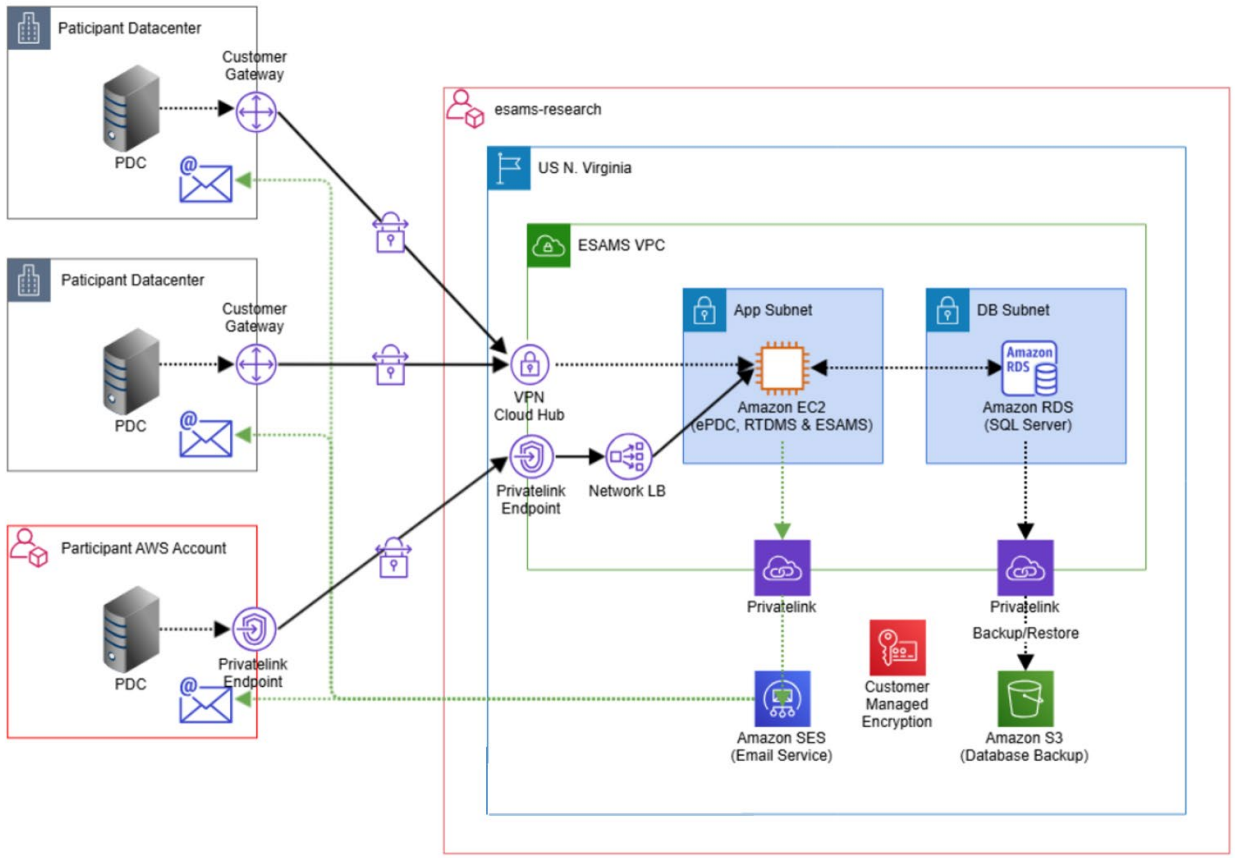

Figure 2. Summary of the ESAMS Architecture.

3.3.1 ESAMS Virtual Private Cloud

At its core, the network topology is organized around a dedicated ESAMS Virtual Private Cloud (VPC) deployed in the AWS US East (N. Virginia) region under a dedicated AWS account for the ESAMS project. This VPC is logically segmented into an Application Subnet, hosting the Amazon Elastic Compute Cloud (EC2) instances that run the ePDC, RTDMS, and ESAMS application tiers, and a Database Subnet, which contains the Amazon Relational Database Service (RDS) instance running SQL Server. This subnet separation enforces a layered security posture, ensuring that database resources are not directly reachable from external network paths.

PDCs connect to the ESAMS VPC through two distinct ingress patterns depending on their network posture. PDCs operating from on-premises environments with no dedicated private connectivity such as AWS Direct Connect (DCX) establish encrypted Internet Protocol Security (IPsec) tunnels via Customer Gateways, terminating at a VPN Cloud Hub that aggregates multiple site-to-site connections into a single managed entry point. For PDCs operating in environments where DCX is available, connectivity is established through AWS PrivateLink endpoints, enabling private, high-throughput cross-account connectivity without traversing the public internet. In both cases, inbound traffic is directed through a Network Load Balancer (NLB) before reaching the application tier, providing connection-level load distribution and enabling health-based routing.

A key architectural consideration in the ESAMS design is the handling of dual-stack (IPv4/IPv6) User Datagram Protocol (UDP) traffic. The NLB is configured in dual-stack mode to accept both IPv4 and IPv6 inbound streams from PDCs. However, because AWS requires that registered

targets for a dual-stack NLB be addressed using IPv6, a dedicated Nginx proxy server sits between the NLB and the ESAMS application server within the App Subnet. The proxy receives IPv6 UDP traffic from the NLB and performs protocol translation—converting the IPv6 UDP streams to IPv4 UDP—before forwarding the data to the ESAMS application server. This design allows the application tier to continue operating over IPv4 without modification while fully supporting IPv6-sourced PDC connections at the network ingress layer.

The end-to-end traffic flow therefore follows a five-stage path: (1) originating at the PDC, (2) traversing the Customer Gateway or PrivateLink endpoint, (3) passing through the Virtual Private Network (VPN) Cloud Hub or NLB, (4) entering the ESAMS VPC where the Nginx proxy performs IPv6-to-IPv4 translation, and (5) terminating at the EC2 application layer. The network supports both IPv4 and IPv6 UDP transport, accommodating the protocol requirements of ESAMS data feeds and real-time telemetry streams.

Supporting services are integrated into the network design through additional PrivateLink connections. Amazon Simple Email Service (SES) provides outbound email notification capability, while Amazon Simple Storage Service (S3) is used for database backup and restore operations via a dedicated PrivateLink Backup/Restore path. Customer Managed Encryption (CME) is applied to data at rest within the RDS instance, ensuring that encryption key custody remains with the ESAMS program rather than delegated entirely to AWS-managed keys.

This network design reflects a defense-in-depth approach: external connectivity is constrained to well-defined, authenticated ingress paths; internal VPC traffic is segmented by function; and all data flows—whether in transit or at rest—are subject to encryption controls. The sections that follow describe the network design for each participating grid operator, detailing their respective connectivity patterns, PDC configurations, and integration points with the ESAMS cloud environment.

3.3.2 ISO-NE Network Design

ISO-NE is the host operator for the Cloud ESAMS demonstration and the primary connectivity reference implementation for the program. ISO-NE contributes synchrophasor measurement data from its footprint covering the six-state New England transmission system, with observable regions spanning New Brunswick, ISO-NE (including DC ties to Hydro-Québec TransÉnergie), NYISO/Ontario, and PJM boundary interfaces. The ISO-NE network design leverages dedicated private connectivity to provide a low-latency, high-reliability path between the ISO-NE Participant Datacenter and the ESAMS VPC, as shown in Figure 3.

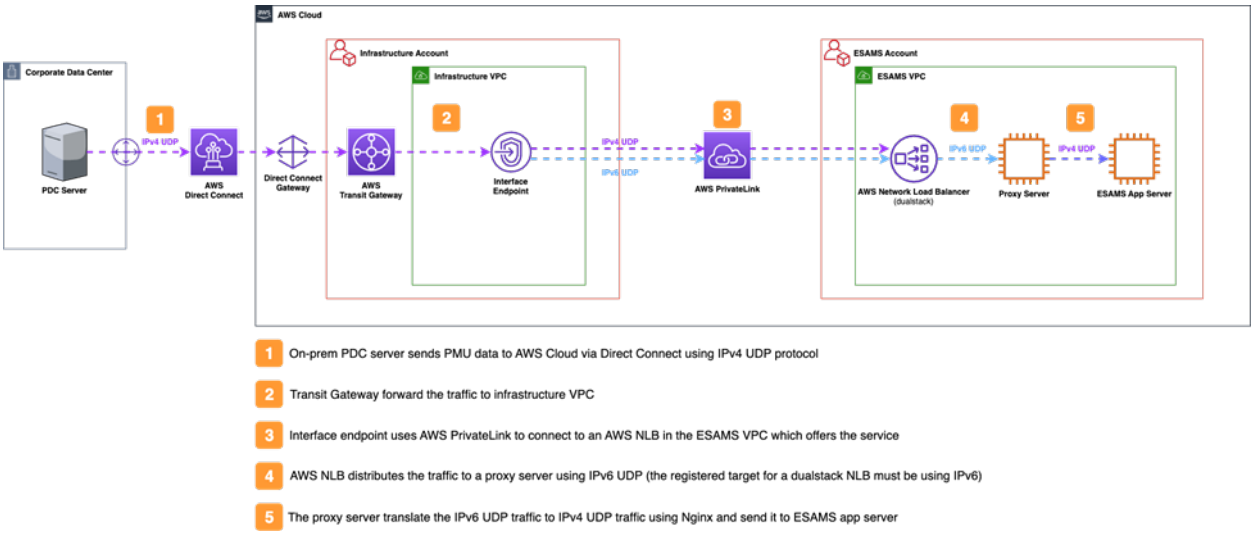

Figure 3. Dedicated connection from ISO-NE on-premises environment to AWS cloud.

Connectivity Model

ISO-NE PDCs are deployed within the ISO-NE operational domain and aggregate high-resolution phasor measurements from PMUs distributed across the New England transmission grid. Connectivity from the ISO-NE environment to the ESAMS VPC is established via AWS Direct Connect, providing a dedicated private circuit that bypasses the public internet entirely. Direct Connect is the preferred connectivity model for ISO-NE given its established AWS relationship as the hosting organization for the Cloud ESAMS demonstration.

Within the ESAMS VPC, traffic arriving via Direct Connect is received at a PrivateLink endpoint and routed through the NLB to the application tier. This path ensures that synchrophasor data streams—which are latency-sensitive and require consistent throughput—are delivered with minimal jitter and predictable latency characteristics.

Data Flow

Phasor data originates from PMUs installed at substations and generation facilities across the ISO-NE footprint. Measurements are forwarded via IEEE C37.118 protocol over UDP to the ISO-NE PDC, which aggregates the streams and re-transmits them toward the ESAMS cloud. The end-to-end path from field measurement to cloud ePDC—including transit through the Transmission Owner's PDC, ISO-NE's PDC, and the cloud ePDC ingestion layer—has been measured at approximately 3.74 seconds during the demonstration period.

The six-stage traffic path proceeds as follows:

1. *PMU (Substation / Generation Facility)* — Captures high-resolution phasor measurements at the field level and forwards them to the ISO-NE PDC.

2. *PDC (ISO-NE Participant Datacenter)* — Aggregates PMU streams from across the ISO-NE footprint and initiates outbound data flow toward ESAMS over IPv4 or IPv6 UDP.

3. *AWS Direct Connect / PrivateLink Endpoint* — Provides the dedicated private path from the ISO-NE environment into the AWS network fabric, terminating at the PrivateLink endpoint in the ESAMS VPC.

4. *Network Load Balancer (ESAMS VPC)* — Operates in dual-stack mode, receiving inbound IPv4 and IPv6 UDP streams and distributing connections to registered IPv6 targets in the App Subnet.

5. *Nginx Proxy Server (App Subnet)* — Receives IPv6 UDP traffic from the NLB and translates it to IPv4 UDP before forwarding to the ESAMS application server. This translation layer is required because the ESAMS application server operates over IPv4, while the dual-stack NLB requires its registered targets to be addressed via IPv6.

6. *Amazon EC2 — ESAMS Application Server (App Subnet)* — Receives IPv4 UDP streams from the Nginx proxy and hosts the ePDC, RTDMS, and ESAMS application components that process and store incoming phasor data. Processed results are written to Amazon RDS (DB Subnet) in the SQL Server database for analysis, reporting, and event notification.

Security and Isolation

All ISO-NE data flows traverse the Direct Connect / PrivateLink path, which is scoped to the ESAMS VPC endpoint service and does not permit lateral movement to other AWS resources. Network access controls at the VPC level restrict inbound traffic to the expected source Internet Protocol (IP) ranges associated with the ISO-NE PDC environment. Data at rest in RDS is protected by Customer Managed Encryption keys under ESAMS program control.

3.3.3 PJM Network Design

PJM participates in the Cloud ESAMS demonstration as a major contributing grid operator, providing synchrophasor measurement data from its footprint spanning 13 states and the District of Columbia. The PJM network design uses site-to-site VPN connectivity to deliver phasor data streams to the ESAMS VPC, as shown in Figure 4.

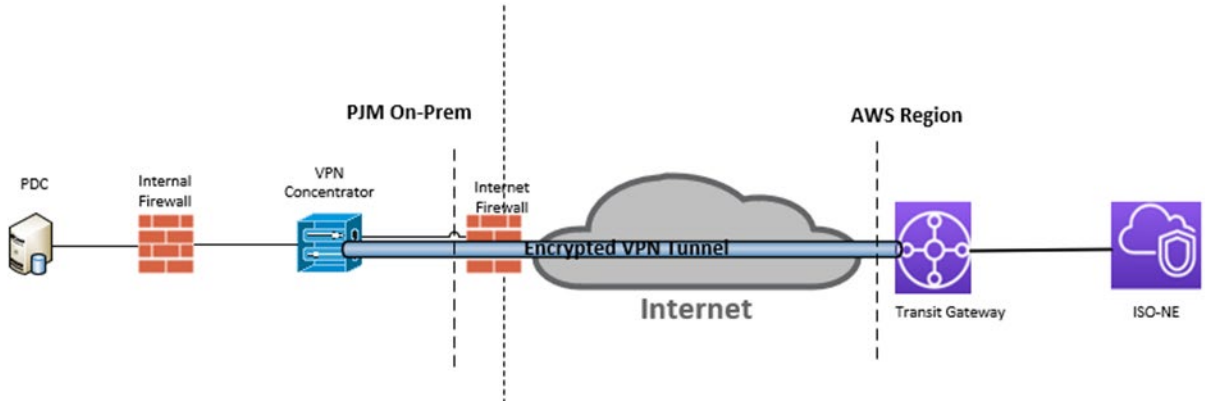

Figure 4. Site-to-site VPN from PJM on-premises environment to AWS cloud

Connectivity Model

PJM PDCs are hosted within PJM-managed Participant Datacenters and aggregate phasor measurements from PMUs distributed across PJM's footprint. As PJM does not have dedicated Direct Connect connectivity to the ESAMS hosting environment, connectivity to the ESAMS VPC

is established via encrypted site-to-site IPsec VPN tunnels. Each PDC site is configured with a Customer Gateway that terminates the VPN tunnel at the AWS VPN Cloud Hub. The Cloud Hub aggregates the incoming VPN connections from PJM PDC sites and routes traffic into the ESAMS VPC through the NLB.

This VPN-based model provides encrypted connectivity for all PJM sites and is well-suited to the data volumes involved in the demonstration—network ingress to the ESAMS server runs at approximately 2.5 MB every 5 minutes across the combined PMU feeds between ISO-NE and PJM, covering tie-line PMU data from both system operators.

Data Flow

PJM phasor data originates from PMUs across its footprint and is forwarded to regional PDCs that aggregate the measurements before transmission to ESAMS. The six-stage traffic path proceeds as follows:

1. *PMU (Substation / Generation Facility)* — Captures high-resolution phasor measurements at the field level and forwards them to the PJM PDC.

2. *PDC (PJM Participant Datacenter)* — Aggregates PMU streams from the PJM footprint and initiates outbound data flow toward ESAMS over IPv4 or IPv6 UDP.

3. *Customer Gateway* — Each PJM PDC site terminates an IPsec VPN tunnel at a Customer Gateway, providing an encrypted egress path from the PJM environment.

4. *VPN Cloud Hub / NLB* — The Cloud Hub aggregates VPN connections from PJM PDC sites and routes traffic into the ESAMS VPC. The NLB operates in dual-stack mode and distributes inbound UDP streams to registered IPv6 targets in the App Subnet.

5. *Nginx Proxy Server (App Subnet)* — same as the description in Section 3.2.2.

6. *Amazon EC2 — ESAMS Application Server (App Subnet)* — same as the description in Section 3.2.2.

Scale and Redundancy Considerations

Given the volume of PMU data generated across the PJM footprint, the network design for PJM places particular emphasis on throughput capacity and path redundancy. The NLB can be configured to distribute data payload across multiple EC2 instances in the App Subnet, ensuring that no single instance becomes a bottleneck during peak data ingestion periods. VPN tunnels for on-premises PDC sites are provisioned with redundant tunnel endpoints to provide failover in the event of a tunnel disruption. For PrivateLink-connected sites, AWS's underlying network redundancy provides inherent path resilience without requiring additional configuration.

Security and Isolation

PJM data flows are encrypted end-to-end via IPsec across the site-to-site VPN tunnels. Within the ESAMS VPC, the traffic path is the same as for all participants: inbound streams are received at the NLB, passed through the Nginx proxy for protocol translation, and delivered to the application server over IPv4. VPC security groups and network Access Control Lists (ACLs)

restrict inbound access to authorized PJM source ranges. All data at rest in RDS is protected by Customer Managed Encryption keys under ESAMS program control.

4.0 Network Performance Results

Along with the accuracy of the oscillation analysis modules, a major objective of the cloud ESAMS demonstration was evaluating whether such a hosting model could deliver performance comparable to on-premises deployment. Low-latency reliable performance is critical for near-real-time applications intended to aid decision-making during wide-area grid disturbances. In this chapter, some of the network performance parameters recorded during the cloud demonstration are presented.

4.1 Latency

Data latency is a critical performance parameter for synchrophasor data delivery. For the ISO-NE PMU stream, the measured end-to-end latency was approximately 3.74 seconds. This value represents the total elapsed time from the moment a measurement is captured at the PMU device in the substation to the point at which it is received and processed by the cloud-hosted ePDC.

It is important to note that this reported latency reflects a multi-stage data transmission and aggregation pipeline, rather than direct point-to-point communication. Specifically, the data path includes:

1. Transmission from the PMU device to the Transmission Owner's PDC (e.g., Eversource or National Grid)
2. Forwarding from the Transmission Owner's PDC to the ISO New England on-premises PDC
3. Final delivery from the ISO New England PDC to the cloud-based ePDC instance

Each intermediate PDC introduces buffering, alignment, and processing delays, which are necessary to perform time synchronization, data quality checks, and stream aggregation across multiple PMU sources. These functions are essential for ensuring data consistency and reliability, but they also contribute significantly to the overall latency.

As a result, most of the observed 3.74-second latency is attributable to intermediate processing and buffering within the PDC hierarchy, rather than network transmission delay. This is further supported by measurements of the cloud segment latency, which indicate that the transit time from on-premises PDCs (including ISO-NE and PJM) to the cloud environment is only approximately 150 – 200 milliseconds. This relatively low value demonstrates that the wide-area network connectivity between on-premises systems and the cloud infrastructure is efficient and not a performance bottleneck.

From a system design perspective, this distinction is important. The current latency profile reflects a deliberate trade-off between data quality/consistency and real-time responsiveness, with upstream PDC layers prioritizing synchronized and validated data streams before forwarding them to downstream applications.

These observations indicate that expanding the ESAMS deployment to additional regions or RCs would not significantly impact latency performance within the cloud segment, and that any overall latency changes would be dominated by upstream PDC architecture and configuration rather than network scalability constraints.

4.2 Network Traffic

As illustrated in Figure 5, the ESAMS cloud application exhibits a consistently low and stable network traffic profile. The observed data shows that inbound traffic to the application instance remains at approximately 2.5 – 2.8 MB every 5 minutes, while outbound traffic from the application instance is slightly higher, at approximately 4.5 – 5.0 MB every 5 minutes. These values are stable over time, indicating predictable and well-controlled data flow characteristics.

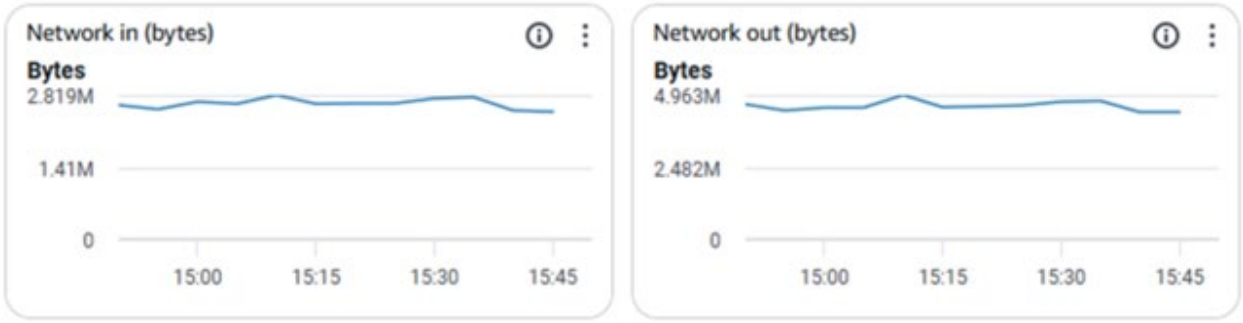

Figure 5. Screenshot of the network monitoring application running on the AWS server hosting ESAMS.

The relatively low inbound traffic volume is primarily a result of the system design choice to transmit only tie-line PMU data to the cloud environment. By limiting data ingestion to only essential measurement points, the system avoids unnecessary data transfer overhead, thereby minimizing bandwidth requirements and improving overall efficiency.

The outbound traffic from the ESAMS application instance is moderately higher than inbound traffic. This difference reflects the system's data processing and enrichment functions. In addition to forwarding raw PMU measurements, the ESAMS applications generate and transmit derived data products, including calculated power flows, oscillation detection and analysis results, and other intermediate or aggregated datasets. These value-added outputs naturally increase the data volume being written to downstream systems, such as the Amazon RDS database, which explains the observed asymmetry between inbound and outbound traffic.

Importantly, the absolute magnitude of both inbound and outbound traffic remains very small in the context of typical cloud network capacity, indicating that network bandwidth is not a limiting factor for system performance. The steady traffic pattern also suggests that the system is operating under stable load conditions without significant spikes or variability, which is desirable for both performance predictability and cost control.

From a scalability perspective, this efficient network utilization has important implications. Because the system only transmits a targeted subset of PMU data and processes it efficiently, expanding ESAMS to include additional RCs within the Eastern Interconnection would result in only a modest increase in total network traffic. The projected growth in bandwidth demand is expected to scale approximately linearly with the number of monitored tie-lines, but given the low baseline, the overall impact on network infrastructure and associated costs would remain minimal.

In summary, the ESAMS architecture demonstrates a highly efficient and scalable network design, characterized by low bandwidth consumption, stable traffic patterns, and effective data prioritization. These attributes support the system's ability to scale regionally without imposing significant additional network or operational burdens.

4.3 Financial Performance

As illustrated in Figure 6, the system demonstrates a stable and predictable cloud cost profile over time. The average monthly operating cost is approximately $1,200, with most expenses attributed to Amazon RDS, which accounts for roughly $800 per month.

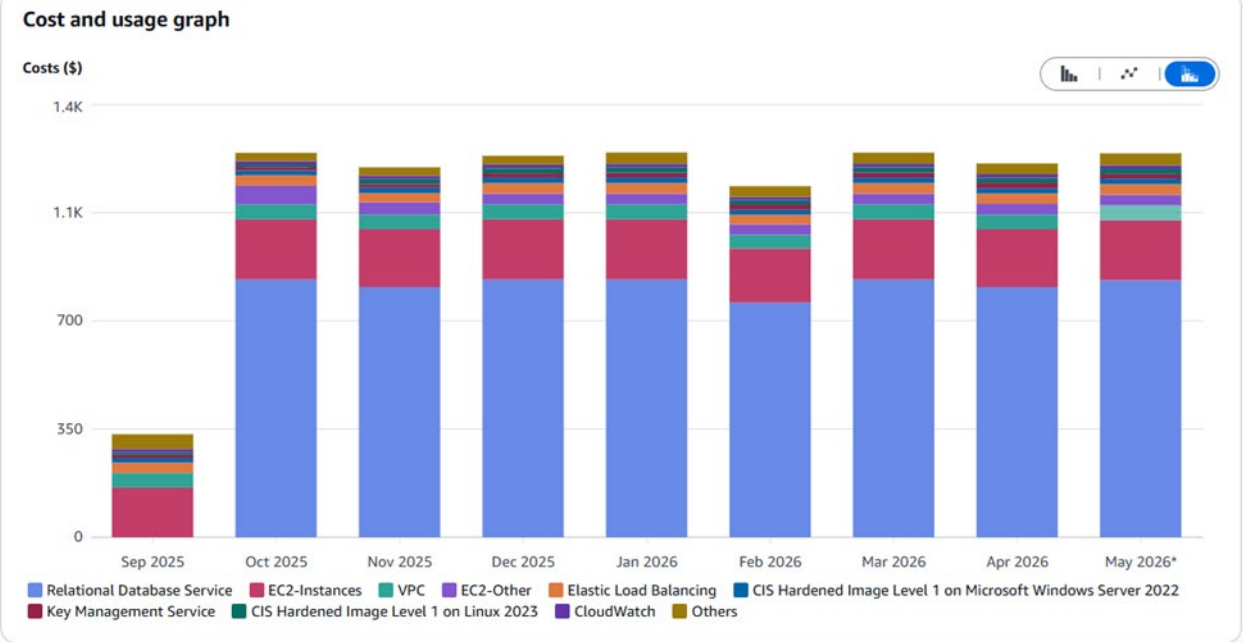

Figure 6. Screenshot of AWS Cost Explorer showing a breakdown of costs associated with the Cloud ESAMS demonstration.

Amazon RDS is a fully managed database service that significantly reduces operational overhead by providing 24/7 infrastructure management, including automated backups, patch management, and built-in high availability features. These capabilities eliminate the need for a dedicated database administrator (DBA) to manage routine maintenance and ensure system reliability. Additionally, the reported RDS cost includes the Microsoft SQL Server licensing fees, which would otherwise represent a separate expense in a self-managed deployment. As such, the RDS cost reflects both infrastructure and enterprise-grade database licensing, contributing to its relatively large share of the overall spending.

The remaining cloud costs are comparatively modest. The two Amazon EC2 instances, which support application and compute workloads, collectively account for approximately $200 per month, indicating efficient resource sizing and utilization. Virtual Private Cloud (VPC)-related charges, including networking and data transfer components, contribute an additional $50 per month. Other minor cost categories, such as monitoring, security services, and ancillary AWS components, represent only a small fraction of the total expenditure, as shown in the chart.

Overall, the current cloud operating cost is highly cost-effective, particularly when considering the level of reliability, scalability, and managed service capabilities provided by AWS. The architecture demonstrates a balanced allocation of resources, with most of the cost concentrated in critical, high-value services such as the database layer.

Looking ahead, if the ESAMS platform were scaled to support all RCs across the Eastern Interconnection, the total cloud cost is projected to increase only modestly, to approximately

$2,000 per month. This relatively small increase reflects economies of scale inherent in cloud-based architectures. Furthermore, because this cost would be shared across multiple participating organizations, the per-entity financial burden would remain low, reinforcing the economic viability and scalability of the solution.

4.4 Data Quality and Availability

Data quality and availability are fundamental prerequisites for reliable wide area monitoring and situational awareness. Throughout the Cloud ESAMS demonstration, data quality was continuously monitored using the ePDC input monitoring interface, ESAMS's built-in data validation functions, and RTDMS's PMU performance reporting capability. As described in Section 2.1.3, ESAMS incorporates built-in data quality monitoring capabilities that continuously assess PMU health, communication status, synchronization status, data completeness, and measurement validity. These functions operate independently of the oscillation analytics and provide visibility into the quality of incoming data streams before they are utilized by any downstream applications. During the cloud demonstration, these capabilities were used to evaluate both aggregate data availability from ISO-NE and PJM as well as the performance of individual PMUs contributing measurements to the cloud environment. The monitoring results demonstrated that the cloud platform itself introduced no degradation to data quality or measurement availability. Any data quality issues identified during the demonstration were traced to upstream sources including individual PMUs in the field and the on-premises PDC infrastructure at participating RCs before measurements were transferred to the cloud environment.

Incoming Input Stream Availability Monitoring

The ePDC continuously monitored incoming PMU streams from both ISO-NE and PJM and maintained statistics on expected versus received measurements. The ePDC Input Monitoring interface provided real-time visibility into data stream health, tracking expected samples, received samples, missing samples, and the resulting data availability percentage for each input source. Figure 7 shows the ePDC input monitoring interface used during the demonstration. The interface provides real-time visibility into the health of each incoming PMU feed, including expected samples, received samples, missing samples, communication status, and overall availability percentage.

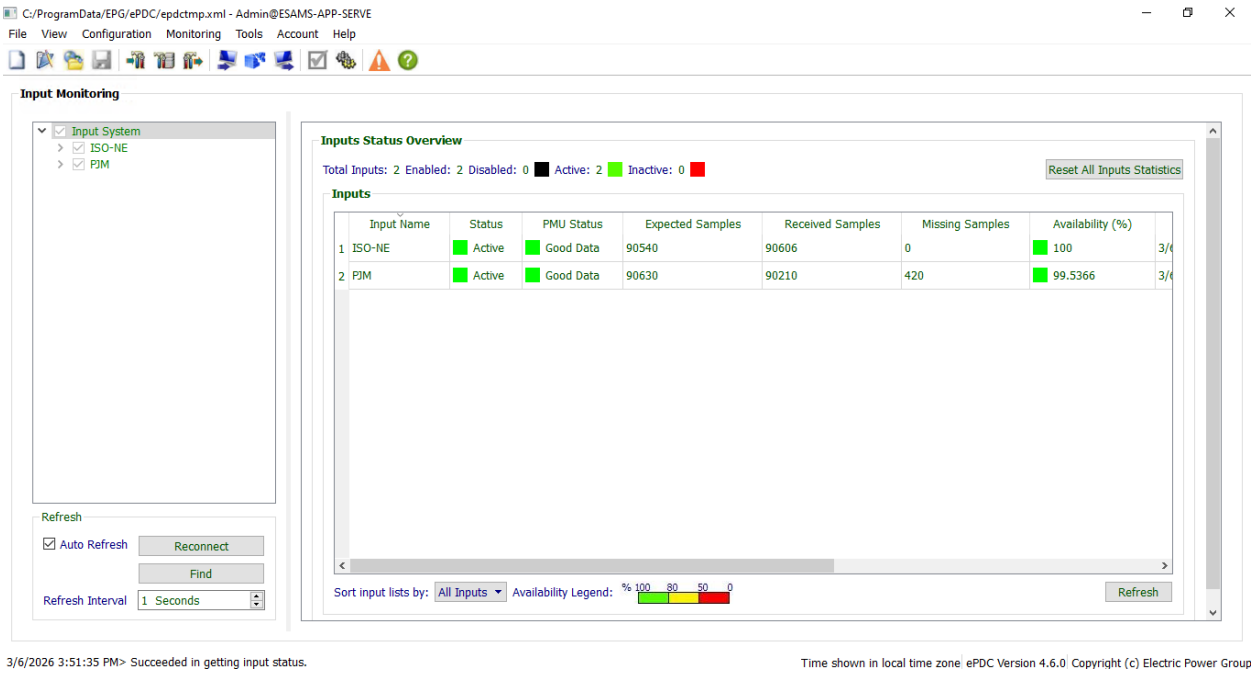

Figure 7. Input stream monitoring dashboard.

In addition to availability statistics, ePDC provided detailed monitoring of PMU communication status, synchronization flags, PMU health indicators, and measurement latency. Figure 8 illustrates the PMU status monitoring view available within the ePDC. Individual PMUs can be monitored in real time, including PMU identifier, number of phasors, synchronization status, communication status, and end-to-end latency measurements.

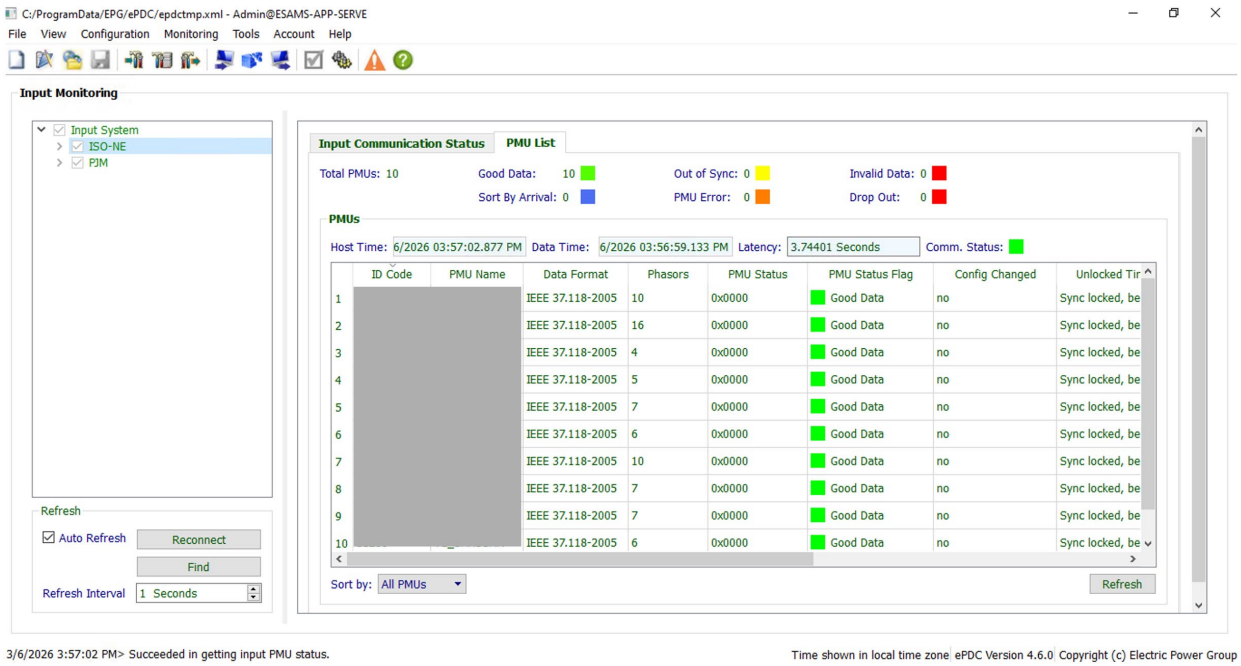

Figure 8. PMU status monitoring dashboard.

PMU Data Quality and Performance Assessment

PMU performance reports were generated using ESAMS's built-in data quality monitoring application. These reports classify measurements into several categories, including:

- Good Samples
- Dropout Errors
- Invalid Data
- Transmission Errors
- Synchronization Errors
- Time Errors

Figure 9 and Figure 10 present examples of PMU performance reports generated during the demonstration showing results for a period of 1 hour. Figure 9 shows individual PMU performance over the reporting interval, while Figure 10 summarizes overall PMU performance across all monitored devices. Green segments represent valid measurements successfully received and processed by ESAMS, while other colors indicate various categories of data quality issues. This example shows that most monitored PMUs exhibited availability exceeding 99% with a few PMUs showing dropouts and invalid data periods.

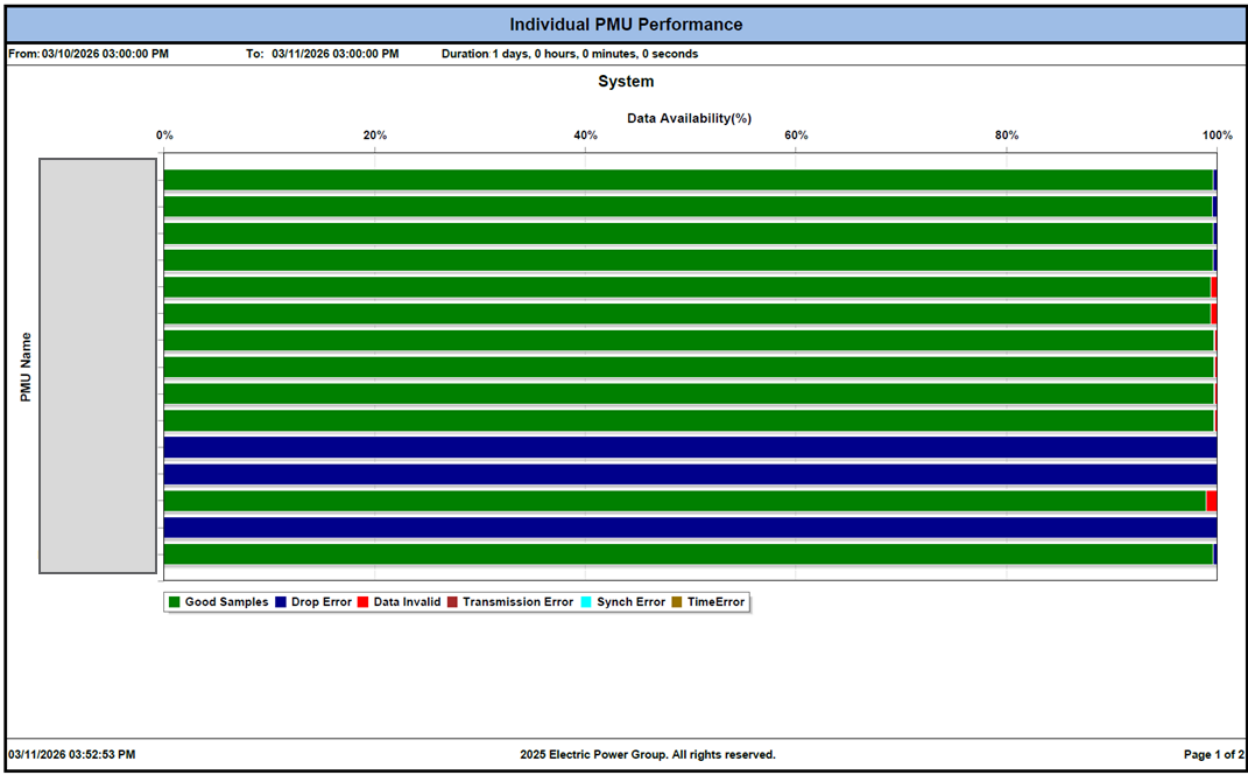

Figure 9. Individual PMU performance report.

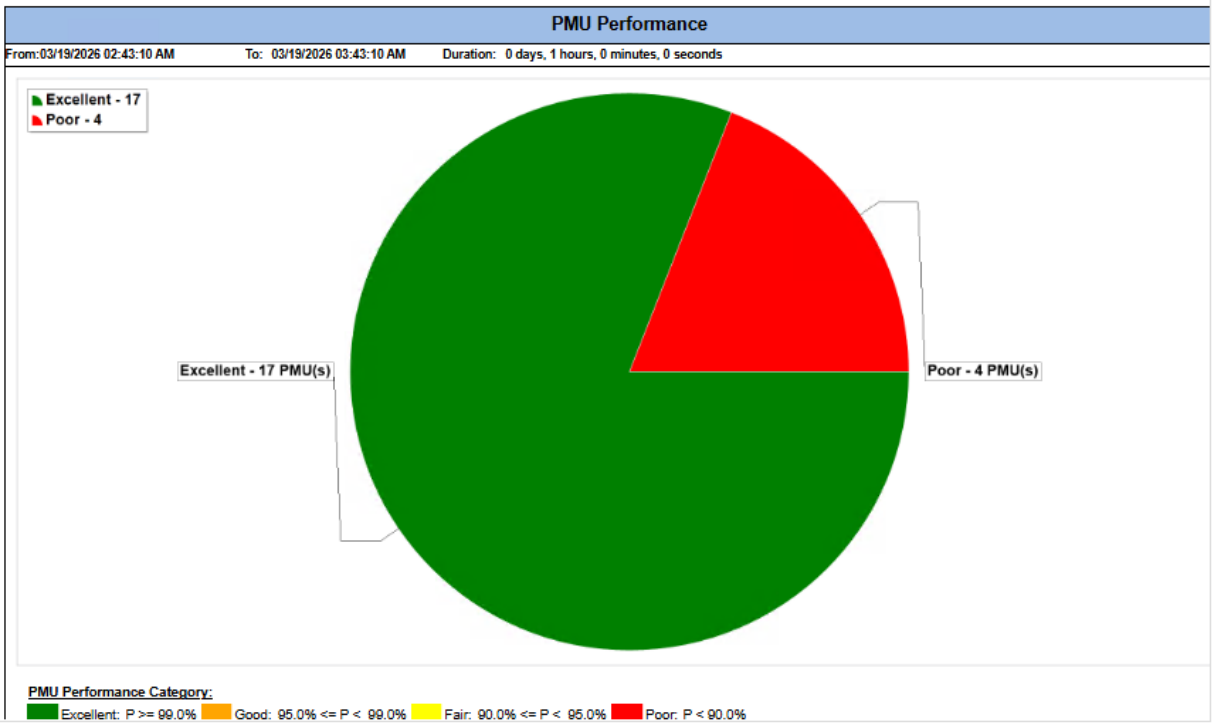

Figure 10. Overall PMU performance report.

The PMU performance reports provide visibility into the root causes of unavailable or invalid measurements, allowing investigators to determine where data issues originated. Review of the PMU performance statistics and supporting diagnostic information along with additional verification confirmed that any observed data quality issues were traced to conditions occurring upstream of the cloud environment, including:

- PMU communication interruptions
- PMU synchronization issues
- Data dropouts occurring before arrival at the cloud ePDC
- Invalid measurement flags originating from source PMUs
- Data losses occurring within source PDC infrastructures

No evidence was observed indicating that AWS networking infrastructure, the cloud ePDC, the ESAMS application layer, or cloud database services introduced additional measurement errors, synchronization problems, or communication failures.

In summary, the cloud ESAMS demonstration successfully verified that cloud-hosted ESAMS platform can maintain the same high level of data quality and availability traditionally achieved by on-premises deployments. Most importantly, the demonstration established that the cloud infrastructure does not adversely affect PMU data quality. All observed data quality deficiencies originated upstream of the cloud environment, either within source PMUs, communication networks, or local PDC systems. ESAMS faithfully preserved measurement integrity while providing enhanced visibility into data quality performance. These findings demonstrate that cloud deployment can support wide area monitoring and situational awareness without compromising the quality of the measurements on which those applications depend.

5.0 Oscillation Analysis Results

The official field demonstration of Cloud ESAMS was performed between October 1, 2025, and March 31, 2026. Several improvements to the configuration were deployed during the demonstration's first three months, so the following review of detected oscillations focuses on the final three months of the demonstration. ESAMS was configured to include oscillations of at least 2 MW/MVAR peak-to-peak amplitudes in daily reports. Each oscillation included in a daily report was reviewed by PNNL to verify correct operation. A summary of these events is presented in Table 1. Highlighted rows indicate particularly interesting events that will be discussed in detail later in this chapter.

Table 1. Summary of oscillations detected over the final three months of the field demonstration.

Event	Frequency (Hz)	Amplitude	Duration (minutes)
1	0.10	24.7 MW	1.5
2	0.33	3.7 MW	3
3	0.23	5.1 MW	4
4	0.10	5.7 MVAR	11
5	1.11	2.2 MW	21
6	0.29	5.7 MW	10
7	0.14 0.20 0.31	20.8 MW 11.5 MW 14.5 MW	2.5
8	0.40	7.4 MW	4
9	0.50	6.3 MW	4
10	1.48	21.3 MW	3
11	1.51	7.3 MW	78 (intermittent)
12	0.37	4.7 MW	26
13	0.83	5.6 MW	1.5
14	0.50	6.4 MW	5
15	0.23	6.7 MW	2
16	0.40	6.9 MW	6
17	0.73	2.2 MW	6
18	0.40	6.5 MW	6
19	2.63	3.4 MW	677 (intermittent)
20	0.47	2.2 MW	5
21	0.54	10.5 MW	2
22	0.50	4.1 MW	6
23	0.21	13.4 MW	2
24	0.50	2.8 MW	6

Of the 24 events, only 4 occurred at a frequency higher than 1 Hz. This result is expected because the oscillation detection algorithm within ESAMS is configured to detect oscillations with significant observability across a wide area, and oscillations are more prone to spreading across a wide area at low frequencies in the range of the system's inter-area dynamics. Table 1 also shows that only event 4 was observed primarily in MVAR, with the other 23 events manifesting as MW oscillations. This too is expected because MVAR oscillations tend to remain more localized. Amplitudes above 20 MW were observed for 3 events. Two additional oscillations had amplitudes above 10 MW, with the remaining 19 oscillations having amplitudes between 2 and 10 MW/MVAR. These smaller oscillations are typically not of immediate concern for power system operation, though long-lasting low-level oscillations can be symptomatic of equipment malfunctions that should be addressed. Analysis of the smaller oscillations was also useful for

validating ESAMS's performance. Four events lasted between 10 and 30 minutes, event 11 was intermittent over a 78-minute period, and event 19 was intermittent for over 11 hours.

The 24 events detected by ESAMS provided a diverse set of examples to validate performance across different frequencies, amplitudes, and durations. Across all events, the oscillation analysis components in ESAMS performed as expected. More detailed analyses for a subset of events with particularly interesting characteristics are presented in the following subsections.

5.1 Event 1: 24.7 MW Oscillation

The largest oscillation detected during the demonstration reached an amplitude of 24.7 MW at its fundamental frequency of 0.1 Hz. The oscillation was apparent in time-domain measurements, as shown in Figure 12. The oscillation persisted for only 90 seconds before disappearing without intervention by a transmission system operator. The flow of dissipating energy for this event is shown in the left plot of Figure 13. As with all results presented in this chapter, the values have been normalized based on the interface with the largest DEF value. For this event, only the interface between ISO-NE and NYISO had a substantial flow of dissipating energy. As a result, ISO-NE has a much higher net export than any other region and is identified as the source, as shown in the right plot of Figure 13.

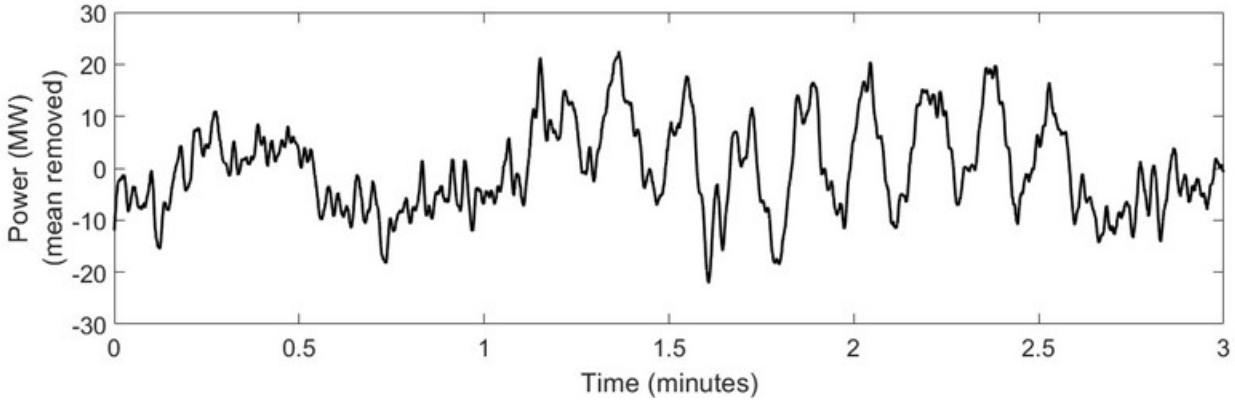

Figure 11. Active power measurements of the 24.7 MW oscillation.

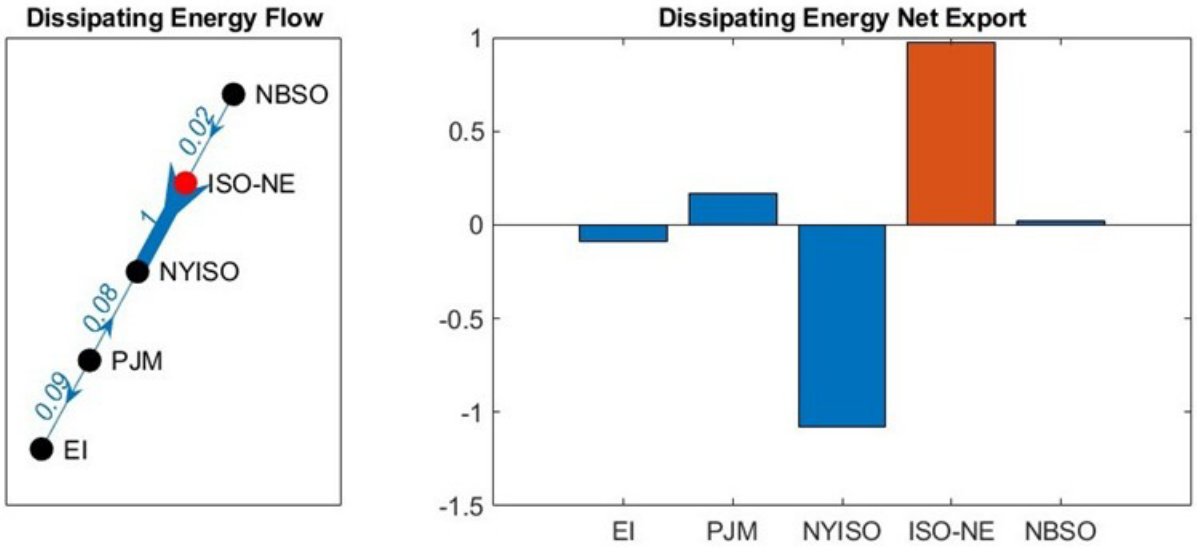

Figure 12. Regional dissipating energy flow (left) and each region's net export of dissipating energy (right) for the 24.7 MW oscillation. As the region with the highest net export, ISO-NE was identified as the source.

5.2 Event 7: Multi-Frequency Oscillation

As shown in Figure 14, the next event appeared as a series of step changes in measurements of active power. Similar to a square wave, this type of oscillation can be described as the combination of sinusoids at distinct frequencies. In this case, oscillations were detected at 0.14 Hz, 0.20 Hz, and 0.31 Hz. The largest of these was the 0.14 Hz component at 20.8 MW. Note, though, that swings greater than 70 MW were observed in the step changes. ESAMS was configured to monitor frequencies as low as 0.1 Hz, so it did not detect the oscillation's frequency components below 0.1 Hz, including the largest component at approximately 50 MW with a frequency of 0.055 Hz. An offline test confirmed that adjusting the lower end of the monitored frequency range to 0.015 Hz allowed the 0.055 Hz component to be detected. The lower limit of 0.1 Hz was retained in the online implementation because oscillations below this frequency tend to remain localized and are therefore of less interest for the ESAMS application.

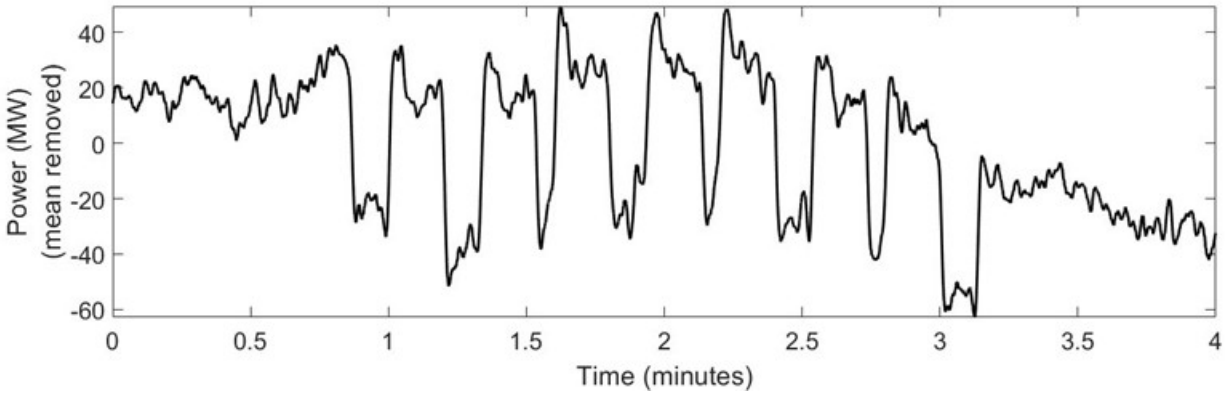

Figure 13. Active power measurements of the multi-frequency oscillation.

The detector's operation is depicted in Figure 15. The test statistic combines information from all input signals to represent the amount of evidence there is for an oscillation at each frequency, with large, narrow peaks indicating the presence of a forced oscillation. The test statistic incorporates the oscillation's amplitude, its dispersion through the power system, and the power of the signal's ambient noise[1] at the oscillation's frequency. These latter aspects explain why the peaks in the test statistic at 0.2 Hz and 0.31 Hz are larger than the peak at 0.14 Hz, despite the 0.14 Hz component of the oscillation having a larger amplitude. Detection occurs when the test statistic exceeds the threshold. The threshold is generated based on a statistical approach that maintains a consistent limit on the probability of false alarm for each frequency. The threshold varies with frequency because it also accounts for coherence among input signals, which is frequency dependent.

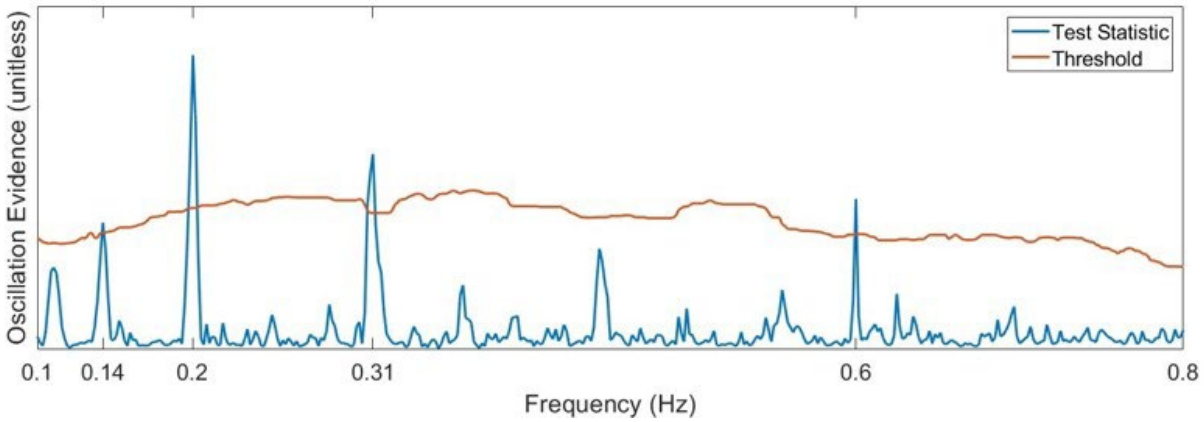

Figure 14. Detection of the multi-frequency oscillation at its 0.14 Hz, 0.2 Hz, and 0.31 Hz components. The oscillation at 0.6 Hz was unrelated and had an amplitude less than the 2 MW threshold for reporting.

The regional dissipating energy flow for the 0.31 Hz component of the oscillation is presented in Figure 16. The diagram on the left clearly shows ISO-NE exporting to both NYISO and NBSO, leading to the very large net export in the plot on the right. Results for the 0.14 Hz and 0.2 Hz components also identified ISO-NE as the clear source. The results for the 0.31 Hz component were distinct in that the flow between ISO-NE and NBSO was much larger for it than the other frequencies. This result highlights the role that inter-area dynamics play in how quickly an oscillation's amplitude dissipates as it spreads through the power system.

[1] Ambient noise refers to the process noise that results from random load changes perturbing the power system's dynamics. The power of the process noise varies with frequency based on the grid's dynamics.

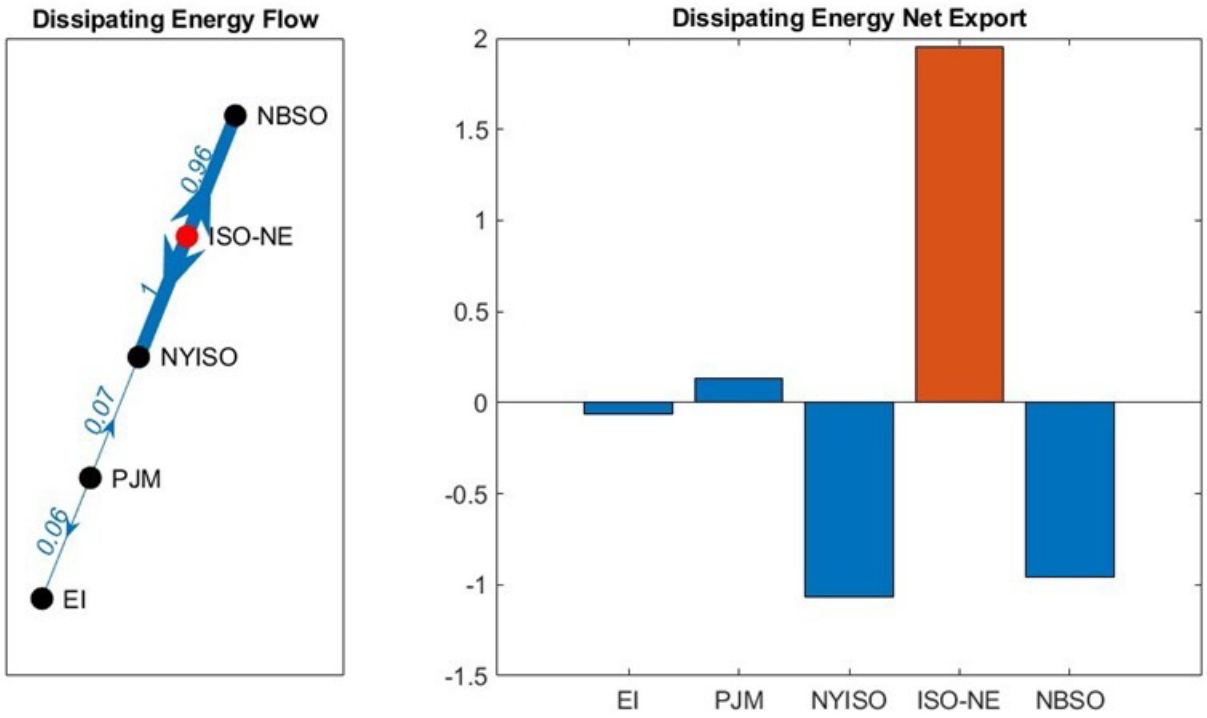

Figure 15. Regional dissipating energy flow (left) and each region's net export of dissipating energy (right) for the 0.31 Hz component of the oscillation. As the region with the highest net export, ISO-NE was identified as the source.

5.3 Events 10 and 11: Recurring Oscillations Near 1.5 Hz

Event 10 occurred at 1.48 Hz with an amplitude of 21.3 MW and a duration of 3 minutes. The next day, the oscillation occurred again at 1.51 Hz. The amplitude was lower at 7.3 MW, but it was present intermittently over a 78-minute period. Measurements from ISO-NE were unavailable at the time, so only a subset of interfaces were visible. The localization results for the 7.3 MW occurrence are presented in Figure 17. The results for the 21.3 MW occurrence were nearly identical.

Following the event, a transmission system operator within the EI region of Figure 17 confirmed that they were aware of event 11 and that its source was within their footprint. However, the transmission system operator was not aware of the larger short-duration occurrence the previous day. Upon further investigation, they confirmed that it originated at the same source as event 11. The oscillation was significantly larger at the source than in the measurements supplied to ESAMS.

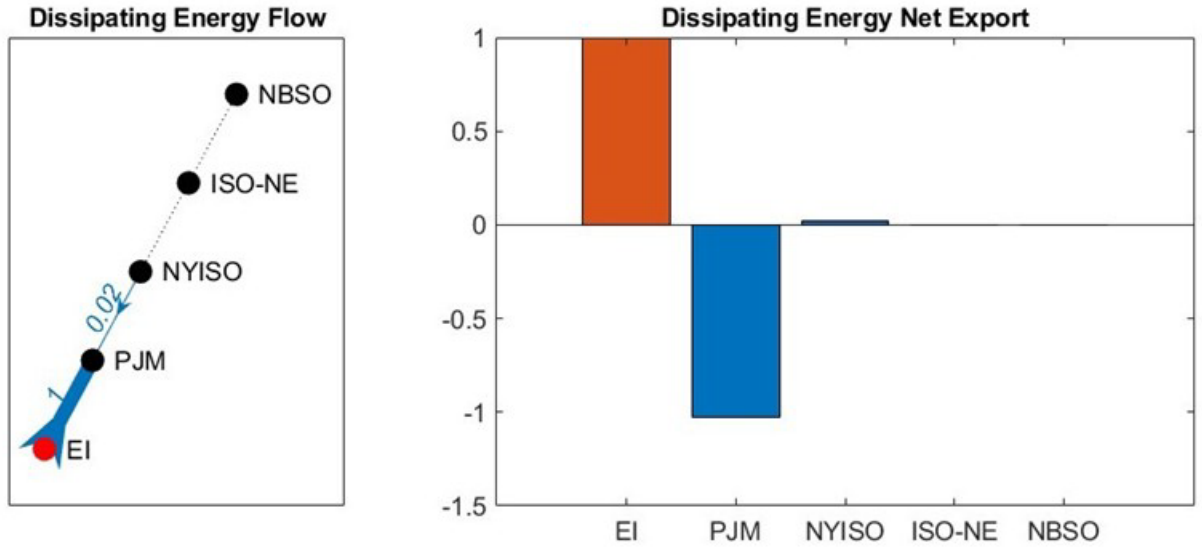

Figure 16. Regional dissipating energy flow (left) and each region's net export of dissipating energy (right) for the 7.3 MW oscillation. As the region with the highest net export, the rest of the Eastern Interconnection bordering PJM was identified as the source.

5.4 Event 13: 5.6 MW Oscillation

Event 13 was a relatively small oscillation at 5.6 MW, but it was apparent in active power measurements, as shown in Figure 18. The oscillation had a frequency of 0.83 Hz and persisted for approximately 1.5 minutes. The localization results are presented in Figure 19. The dissipating energy flowed out of NYISO to both ISO-NE and PJM, resulting in a large next export of dissipating energy that led to NYISO being identified as the source.

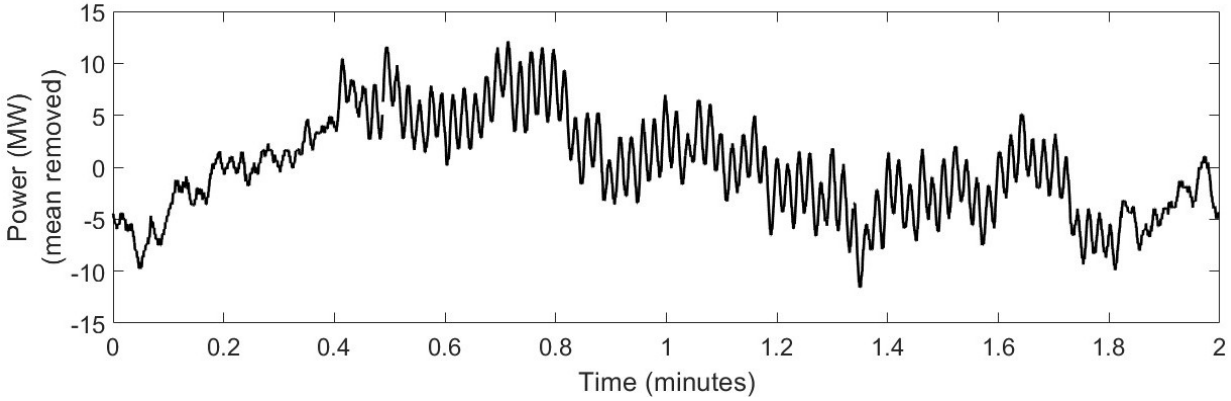

Figure 17. Active power measurements of the 5.6 MW oscillation.

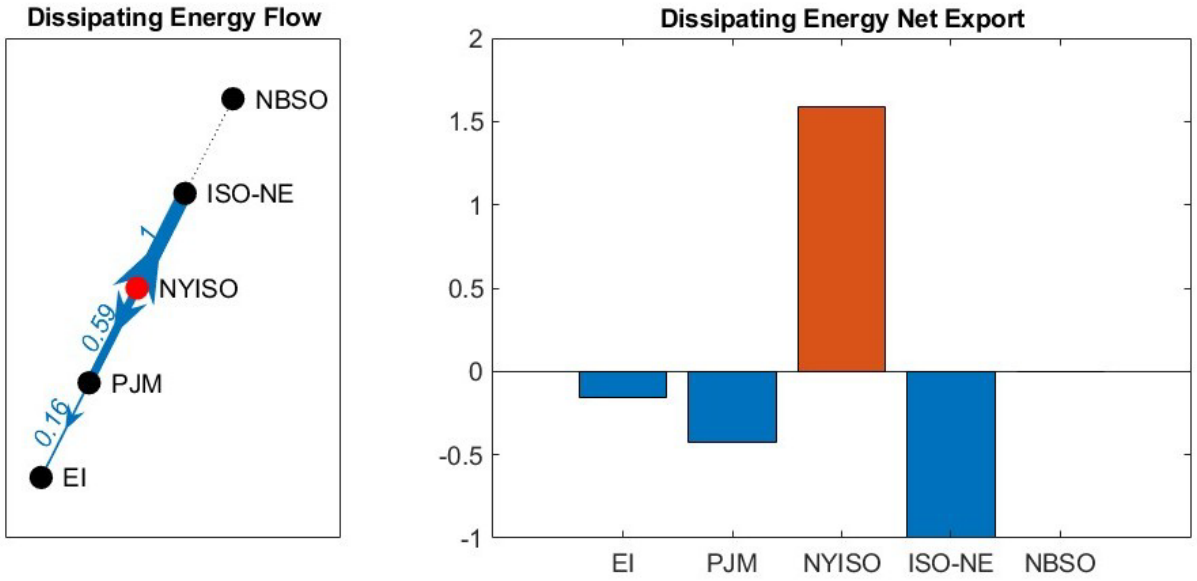

Figure 18. Regional dissipating energy flow (left) and each region's net export of dissipating energy (right) for the 5.6 MW oscillation. As the region with the highest net export, NYISO was identified as the source.

5.5 Event 19: 3.4 MW Long-Duration Oscillation Example

Event 19 was by far the longest duration oscillation that was reported. It occurred intermittently over a 677-minute period spanning two days at a frequency of 2.63 Hz and a maximum amplitude of 3.4 MW. Despite its small size, oscillations of such long duration can point to equipment problems that need to be addressed. A plot of the oscillation is provided in Figure 20, and the localization results are presented in Figure 21. The only major flow of dissipating energy was from NYISO to ISO-NE, leading NYISO to be identified as the source.

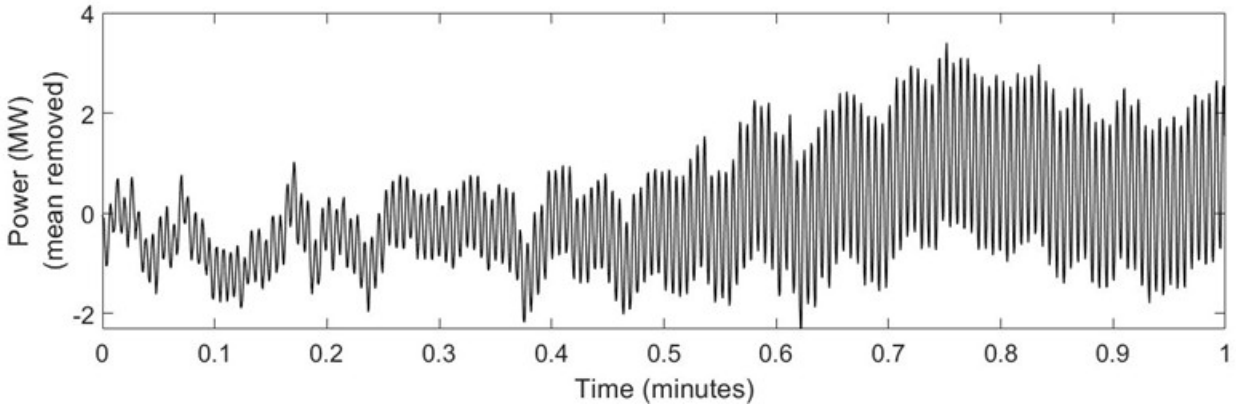

Figure 19. Active power measurements of the oscillation as it increased to its maximum amplitude of 3.4 MW.

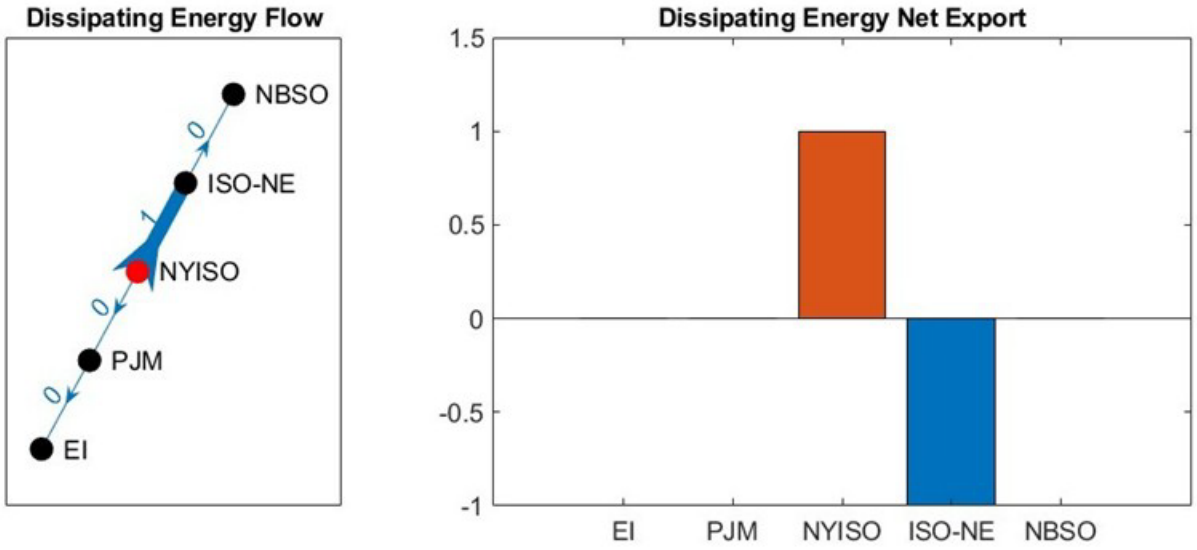

Figure 20. Regional dissipating energy flow (left) and each region's net export of dissipating energy (right) for the 3.4 MW oscillation. As the region with the highest net export, NYISO was identified as the source.

6.0 Conclusion

One of the key wide-area monitoring capabilities enabled by synchrophasor measurements is oscillation detection and source localization. Oscillations can spread from their source and impact entire interconnections, making it critical for reliability coordinators to have information about where the oscillation is originating. ESAMS has the technical capability to address this need, but the initial demonstration identified on-premises hosting at a RC as a major barrier to adoption. In this work, the project team demonstrated that transitioning ESAMS to a cloud environment is a viable approach to removing this barrier.

In preparation for the six-month field demonstration, ESAMS was deployed in an AWS cloud environment hosted by ISO-NE. Synchrophasor measurements were streamed from reliability coordinators to the cloud using two different approaches. ISO-NE streamed measurements using AWS Direct Connect, while PJM connected to the cloud using a site-to-site VPN. Measurements from these organizations provided observability from as far north as Maine to as far west and south as Indiana and Virginia. By monitoring measurements along the borders of ISO-NE and PJM, oscillation sources could be localized to five regions: NBSO, ISO-NE, NYISO, PJM, and the rest of the Eastern Interconnection.

During the final three months of the demonstration, 24 oscillation events larger than 2 MW/MVAR were detected and localized to a region. The largest of these events was nearly 25 MW and the longest persisted intermittently for over 11 hours. The successful detection of so many wide-area oscillations reflects the successful use of cloud resources to gather measurements and perform the analysis. The inbound and outbound network traffic remained at approximately 2.5 – 2.8 MB and 4.5 – 5.0 MB every 5 minutes, respectively. These values are low for typical cloud network capacity, demonstrating the scalability of the Cloud ESAMS approach. Likewise, the transit time from on-premises PDCs to the cloud environment of 150 – 200 milliseconds indicates that the system can scale to include more measurements spanning additional regions of the grid. Based on the costs observed during the demonstration, this scaling could be achieved with relatively modest cost increases.

Based on the success of the initial six-month demonstration, the Cloud ESAMS deployment has been extended to provide additional reliability coordinators with an opportunity to consider participation. Expanding to cover a broader portion of the Eastern Interconnection will allow ESAMS to provide more detailed information about each detected oscillation's amplitude and source region. The periodogram detector used in ESAMS is designed to focus on wide-area oscillations, and adding measurements from a wider region will also strengthen this focus so that local oscillations that can be addressed with an individual entity's tools are not reported.

Wider visibility also presents opportunities to expand ESAMS's capabilities to increase value. For example, continuous tracking of the system's small-signal stability, which was configured during the initial 2021-2022 field demonstration, could resume. New capabilities, such as localization of frequency excursion events could also be considered.

Over multiple field demonstrations, ESAMS has clearly shown the value of analyzing synchrophasor measurements across not only the footprint of a reliability coordinator, but across an entire interconnection. Effective monitoring at this scale introduces significant challenges, particularly for on-premises hosting at a single RC's premises. The Cloud ESAMS demonstration clearly showed the suitability, reliability, readiness, and cost effectiveness of cloud-based hosting to address this challenge and create further opportunities for interconnection-scale monitoring that can enhance the reliability of the bulk power system.

7.0 References

Agarwal, Abhijeet, John Balance, Bharat Bhargava, Jim Dyer, Ken Martin, and Jianzhong Mo. 2011. "Real Time Dynamics Monitoring System (RTDMS®) for use with SynchroPhasor technology in power systems." *2011 IEEE Power and Energy Society General Meeting.* 1-8. doi:10.1109/PES.2011.6039688.

Biswas, Shuchismita, and Jim Follum. 2024. "Interpreting Forced Oscillation Notifications from ESAMS: General Guidance for Reliability Coordinators." Tech. rep., PNNL. PNNL-37162.

Biswas, Shuchismita, Jim Follum, and Joseph H. Eto. 2023. "Confidence Assessment for Regional Forced Oscillation Source Localization: Formulation and Field Validation." *IEEE Transactions on Power Delivery* 38: 3739-3748. doi:10.1109/TPWRD.2023.3284418.

Donnelly, Matt, Dan Trudnowski, James Colwell, John Pierre, and Luke Dosiek. 2015. "RMS-Energy Filter Design for Real-Time Oscillation Detection." *2015 IEEE Power and Energy Society General Meeting.* Denver, CO. 1-5.

Eto, Joseph H., Neeraj Nayak, Simon Mo, Ken Martin, Song Xue, Horacio Silva-Saravia, Jay Chen, et al. 2022. "Eastern Interconnection Situational Awareness Monitoring System (ESAMS) Demonstration Project." Tech. rep., Lawrence Berkeley National Laboratory.

Follum, Jim. 2020. "Statistical Evaluation of New Estimators Used in Forced Oscillation Source Localization." *Hawaii International Conference on System Sciences (HICSS).*

Follum, Jim, and Francis Tuffner. 2016. "A multi-channel method for detecting periodic forced oscillations in power systems." *2016 IEEE Power and Energy Society General Meeting.* 1-5. doi:10.1109/PESGM.2016.7741382.

Follum, Jim, and John W. Pierre. 2016. "Detection of Periodic Forced Oscillations in Power Systems." *IEEE Transactions on Power Systems* 31: 2423-2433. doi:10.1109/TPWRS.2015.2456919.

Follum, Jim, Shuchismita Biswas, Tawsif Ahmad, Clifton Black, and Michael Breuhl. 2025. "A Bootstrapping Approach to Setting Thresholds for RMS-Energy Oscillation Detectors." *2025 IEEE Power and Energy Society General Meeting.* 1-5. doi:10.1109/PESGM52009.2025.11225125.

Follum, Jim, Tianzhixi Yin, and Nicholas Betzsold. 2020. "Regional Oscillation Source Localization: Implementation in the ESAMS Tool." Tech. rep., PNNL. PNNL-29612.

NERC. 2019. "Eastern Interconnection Oscillation Disturbance: January 11, 2019 Forced Oscillation Event." Tech. rep., North American Electric Reliability Corporation (NERC), https://www.nerc.com/globalassets/our-work/reports/event-reports/january_11_oscillation_event_report.pdf.

NERC. 2021. "Recommended Oscillation Analysis for Monitoring and Mitigation." Tech. rep., North American Electric Reliability Corporation (NERC), https://www.nerc.com/globalassets/who-we-are/standing-committees/rstc/smwg/oscillation_analysis_for_monitoring_and_mitigation_trd.pdf.

Pacific Northwest National Laboratory

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

1-888-375-PNNL (7665)

www.pnnl.gov