

A Methodology to Evaluate the Grid Reliability Impact of Oscillations Induced by Large Loads

June 2026

Shuchismita Biswas¹

Sameer Nekkhalapu¹

Antos C. Varghese¹

Brett Ross¹

Jim Follum¹

Kaustav Chatterjee¹

Rahul Chakraborty²

¹Pacific Northwest National Laboratory

²Dominion Energy



U.S. DEPARTMENT
of **ENERGY**

Prepared for the U.S. Department of Energy
Under contract DE-AC05-76RL01830

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

www.osti.gov
ph: (865) 576-8401
fax: (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>

A Methodology to Evaluate the Grid Reliability Impact of Oscillations Induced by Large Loads

June 2026

Shuchismita Biswas¹
Sameer Nekkhalapu¹
Antos C. Varghese¹
Brett Ross¹
Jim Follum¹
Kaustav Chatterjee¹
Rahul Chakraborty²

¹Pacific Northwest National Laboratory

²Dominion Energy

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

Pacific Northwest National Laboratory
Richland, Washington 99352

Contents

Abstract	vi
1.0 Background	1
2.0 Power System Oscillations	4
2.1 Oscillations in Different Frequency Bands	4
2.2 Studying Oscillation Impact in Planning Studies	5
3.0 Methodology	9
3.1 Module 1: Screening	9
3.1.1 Voltage sensitivity screening	10
3.1.2 Angle sensitivity screening	11
3.1.3 Oscillation excitability screening	11
3.2 Module 2: Simulation	13
3.2.1 Data center load representation	13
3.2.2 Selecting output channels	15
3.2.3 Cyclic load behavior	16
3.3 Module 3: Analysis	16
3.3.1 Risk metrics	18
3.3.2 Visualizing the risk metrics	19
4.0 Case Studies	22
4.1 WECC 240-bus Model	22
4.1.1 Example 1: 1 Hz mode	22
4.1.2 Example 2: 0.38 Hz mode	24
4.2 WECC Planning Model	25
4.2.1 Example 1: N-1 security constraint violation	25
4.2.2 Example 2: Inter-area propagation	29
5.0 Implications and Future Directions	33
5.1 Key Takeaways	33
5.2 Future Directions	34
Appendix A User-Defined Model for Oscillatory Load	A.1
A.1 Capabilities	A.1
A.2 Implementation	A.1
A.3 Validation	A.1
A.4 Advantages Over the Script-based Approach	A.1

Acknowledgements

This work was funded by the U.S. Department of Energy Office of Electricity. The authors gratefully acknowledge Nathaniel Horner and Sandra Jenkins of DOE for their guidance and support throughout this project, as well as Frank Tuffner of PNNL and M. Logan Pokallus of the Western Area Power Administration (WAPA) for their valuable review and insights. We also thank Pavel Etingov of PNNL for his contributions in generating some of the figures used.

Abstract

The rapid growth of hyperscale AI data centers is bringing renewed attention to the reliability risk that sustained forced oscillations pose to bulk power systems, with cyclic computational workloads emerging as a new forcing source. Unlike the broadband, stochastic disturbances from traditional industrial loads such as arc furnaces, AI training and inference facilities can inject large active power swings concentrated at specific frequencies over extended durations - characteristics that existing grid planning practices do not account for. While the North American Electric Reliability Corporation (NERC) has recognized this gap and called for system-level studies of large load interconnections, no standardized methodology exists to screen, simulate, and quantify these risks at the planning stage.

This report presents the Risk Assessment Tool for Large Load-induced Events (RATLLE), a Python-based, publicly available script suite developed at the Pacific Northwest National Laboratory to evaluate bulk power system reliability risks from data center-induced oscillations. RATLLE implements a three-module workflow: a screening module that identifies vulnerable interconnection locations and excitable system modes; a simulation module that models cyclic data center load behavior using a commercial positive sequence simulation platform; and an analysis module that computes risk metrics and generates interactive visualization dashboards. The risk metrics, formulated around simulation observables, map oscillation impacts to a three-stage severity scale spanning latent equipment fatigue through imminent cascading failure.

The methodology is demonstrated on two Western Electricity Coordinating Council (WECC) system models: a publicly available 240-bus reduced representation and a detailed 2031 Heavy Winter planning case. Case studies illustrate that even modest 50 MW forced oscillations at resonant frequencies can produce wide-area power swings, N-1 security constraint violations, and cascading generator trips through protection actions - outcomes that would not occur under normal operating conditions without oscillations present. The results underscore the need for standardized oscillation impact assessment in large load interconnection studies and provide a reproducible, extensible framework for utilities to adopt or customize within their existing planning workflows.

Figures

1	Example demand periodicity due to an AI-training workload	1
2	Natural and forced oscillations	5
3	Response to non-sinusoidal forcing inputs at resonant frequencies	7
4	RATLLE architecture	9
5	Voltage sensitivity computed for 69-138 kV buses in WECC (2031 Heavy Winter Planning Case)	10
6	Ringdown analysis process	12
7	Analyzing ringdown response due to a 0.5-second load impulse at bus 6508 in the WECC 240-bus system.	12
8	Sequence of steps in the simulation module	13
9	Data center load representation	14
10	Cyclical load patterns that can be simulated in RATLLE	16
11	Process flow in the simulation and analysis modules	17
12	Conceptual illustration mapping oscillation impact to progressive severity stages . . .	18
13	Risk metrics visualization for a scenario simulated in the WECC 240-bus model	21
14	Visualizing transmission lines with highest power swings for a scenario simulated in the WECC 240-bus model	21
15	Grid elements where oscillations higher than 20 MW are recorded when 100 MW load perturbations are introduced from the H ALLEN bus	22
16	Oscillations observed in generator power outputs when 1 Hz, 100 MW load perturbations are introduced from the H ALLEN bus.	23
17	1-Hz mode shape plot	23
18	Max. oscillation amplification vs. mode shape magnitude of source	24
19	Grid elements where > 20 MW oscillations are observed when 0.38 Hz, 100 MW load perturbations are introduced from the CANADA bus	24
20	Modal analysis of load impulse response from the California bus considered	25
21	Approximate locations of power systems elements that cross the RATLLE risk metric thresholds in Example 1	27
22	Time-series plots for the most impacted generator and transmission line in Example 1	28
23	Contingency 1 – Generator A loss leads to the tripping of Generator B due to undervoltage protection action	29
24	Contingency 2 – Line A loss leads to the tripping of Generator B due to undervoltage protection action	30
25	Modal analysis of load impulse response for Example 2 in the WECC planning model	31
26	Approximate locations of impacted elements in Example 2	31
27	Time-domain simulation plots in Example 2 after the Generator C loss contingency . .	32
A.1	Parametrized representation of different oscillatory load profiles in the UDM	A.2
A.2	UDM performance for various load characteristics for square wave profile	A.3

A.3 Zoomed-in plot of Figure A.2 A.3

A.4 Comparing UDM performance with simulations leveraging Python scripts A.5

Tables

1 Oscillations in different frequency bands 6

2 Load composition in the CMLD model used to represent data center dynamics 15

3 CMLD switching characteristics parameters representation of large loads 15

4 Risk metrics for generators 19

5 Risk metrics for other elements 20

6 Risk assessment for Example 1 26

7 Risk assessment for Example 2 32

A.1 Description of the UDM Parameters A.4

1.0 Background

The power grid is experiencing a period of rapid load growth driven by surging energy demand from data centers as a global race for artificial intelligence (AI) dominance accelerates. Hyper-scale facilities can consume hundreds of MW at a single load site, and their dynamic behavior poses unique risks to power grid operations (NERC, 2025b). Recognizing these risks, the North American Electric Reliability Corporation (NERC) Large Loads Working Group has published a series of white papers identifying: key reliability concerns associated with emerging large loads (LLs) such as AI-training facilities (NERC, 2025a); gaps in current practices and standards (NERC, 2026a); and reliability guidelines for mitigating these risks (NERC, 2026c). NERC has further issued a Level 3 Alert calling upon industry stakeholders to improve load modeling practices, conduct system-level studies that incorporate LL operational characteristics, and strengthen performance monitoring infrastructure (NERC, 2026b). A key reliability risk identified in the NERC white papers is forced oscillations induced by cyclic computational loads. While system-wide studies have been recommended to understand risks from site-specific behavior, the lack of a standardized methodology to quantify oscillation impact remains a gap (NERC, 2026a).




Figure 1: Example demand periodicity due to an AI-training workload, reproduced from Choukse et al. (2025)

In hyperscale AI training and inference facilities, synchronized computational tasks across large GPU/TPU clusters can lead to cyclic power demand characterized by periodic consumption bursts followed by cool-down periods (Figure 1) (Choukse et al., 2025). This synchronization effect is not exclusive to AI facilities; coordinated task execution across compute clusters in general-purpose data centers can similarly produce periodic demand impulses (Mishra et al., 2025). Such periodic load patterns can introduce forced oscillations into the grid and/or intermittently excite natural oscillation modes, posing reliability challenges. The frequency of the demand fluctuations determines the nature of reliability risks. High-frequency oscillations (in the ~ 5 –59 Hz range, commonly referred to as subsynchronous oscillations or SSOs) can cause power quality problems and interact with generator torsional modes, but their impact is typically confined to the vicinity of the oscillation source. On the other hand, oscillations at lower frequencies (roughly 0.1–1.5 Hz, often called the electromechanical range) can propagate across large portions of an electrical interconnection and cause protection/control misoperations. Interaction with control loops of inverters or synchronous generators is possible throughout the subsynchronous frequency range. A detailed discussion of the reliability risks associated with oscillations in different frequency ranges is provided in Chapter 2.0.

The threat of oscillatory instability is well recognized, and grid operators employ various monitoring and mitigation solutions to manage forced oscillations during operations (Chatterjee et al.,

2023; NERC, 2021). These practices, however, are premised on the assumption that forced oscillations arise from equipment malfunction and are not a feature of normal operations. Malfunctions are relatively infrequent and inherently difficult to predict; thus, it is more practical to address oscillations as they occur rather than to formulate case-specific mitigations during planning. The cyclic nature of large computational loads fundamentally challenges this assumption and makes assessing forced oscillation risk at the planning stage a critical reliability need.

This is not to say that load demand periodicity is unfamiliar to the industry. Fluctuating loads such as arc furnaces and rolling mills have long been known to inject harmonic distortion, flicker, and power quality problems, and mitigation measures, including passive filters and static VAR compensators (SVCs), have been deployed to limit their adverse impact on the bulk power system (BPS). However, these disturbances, being broadband, stochastic, and largely confined to the local interconnection point, posed limited forced oscillation risks to the BPS in the electromechanical frequency range. Industry practice, therefore, focused on site-specific power quality mitigation. Nonetheless, the potential to affect neighboring utility systems had been acknowledged, and instances of generator shaft damage attributable to industrial load characteristics have been documented (Solanics et al., 2000; Dorfman et al., 1996). In contrast, hyperscale AI facilities can drive large active power swings with energy concentrated at specific frequencies, persisting for hours over the duration of training or inference workloads. Unlike the broadband, stochastic disturbances of arc furnaces, these deterministic and sustained oscillations are not adequately mitigated by SVCs - devices designed for reactive power compensation and voltage regulation, not for damping active power oscillations.

The experience with large pulsating loads from particle accelerators offers a closer analogue. The first known case of a large accelerator being supplied by a utility system without an intermediate on-site motor-generator set occurred in the midwestern United States in the 1970s (Pinneilo and Van Ness, 1971). Utility analysis deemed that the impact of expected load pulsations on nearby generators was acceptable, but future load growth might require mitigation measures. In the 1990s, the Long Island Lighting Company (LILCO) conducted a system impact analysis before interconnecting a synchrotron at the Brookhaven National Laboratory (Rostamkolai et al., 1994). Risks from dynamic power swings, voltage fluctuations, torsional interactions with nearby turbine-generators, and harmonic voltage distortions were studied. It was found that dynamic power swings were the limiting case, and one of the synchrotron's planned operation modes could cause 5% power swings for a nearby generator due to resonance with system modes. This was deemed unacceptable, and the load operation mode was redesigned. Additionally, a power pulsation monitoring relay (PPMR) was installed to protect the LILCO grid from problems on the synchrotron's power control system. The PPMR analyzed the frequency spectrum of the three-phase power on a dedicated 69 kV feeder supplying the load, and was designed to issue alarms or trip the supply if power swings crossed design thresholds. Different oscillation amplitude limits were prescribed for the 0.75-5 Hz and 10-40 Hz frequency ranges.

The LILCO study establishes the value of site-specific system impact analyses before interconnecting large cyclic loads, but it also illustrates a fundamental scalability problem. What was feasible for a single synchrotron in the 1990s becomes impractical for the hundreds of hyperscale data center interconnection requests that utilities now have to process, especially since each facility's workload characteristics may change over its lifetime. Each load site would require a custom study, yet most utilities lack the standardized workflows to screen, simulate, and interpret these risks systematically. The result is planning studies that are slow, inconsistent, and difficult to reproduce. NERC has explicitly acknowledged this gap, noting that "there are significant gaps regarding the defining of allowable limits for forced oscillations" (NERC, 2026a).

To address this, we present the Risk Assessment Tool for Large Load-induced Events (RATLLE)

- a publicly-available¹ Python script suite for evaluating large load-induced oscillation risks in bulk power systems developed by the Pacific Northwest National Laboratory (Biswas et al., 2026). RATLLE is suitable for studying wide-area phenomena with dynamics below ~ 10 Hz that are typically captured in phasor-domain models used by utilities. The framework guides planners progressively through a structured workflow, from system-wide screening that identifies the most vulnerable interconnection locations and frequencies, to site-specific risk assessments that quantify oscillation severity and spatial propagation. The scripts are designed to be adopted directly or customized to fit a utility's existing processes. The current implementation interfaces with PSS/E as the underlying simulation engine, with extensions to PSLF and PowerWorld under development.

This report details the methodology underlying the RATLLE framework and illustrates findings with case-studies on two system models - a 240-bus publicly available reduced representation of the Western Electricity Coordinating Council (WECC) system and a detailed WECC 2031 heavy winter planning case (Yuan et al., 2020). The present work extends our earlier work published in Biswas et al. (2026).

The key contributions of this work are summarized as follows:

- To the authors' knowledge, RATLLE is the first end-to-end methodology — spanning vulnerability screening, time-domain simulation, and risk quantification — for evaluating large load-induced oscillation impacts at the planning stage. The scripts are publicly available, and results can be independently reproduced.
- This work formulates a set of simulation-based risk metrics, organized around a three-stage severity scale. This is one of the first attempts to comprehensively quantify BPS reliability impact from sustained oscillations, for which no NERC planning standard currently prescribes limits.
- Two practical modeling approaches for cyclic computational loads are presented: a Python-scripted representation and a scalable PSS/E user-defined model. Both represent AI workload periodicity in positive-sequence planning studies.
- Case studies on a detailed WECC planning case show that 50 MW oscillations can convert an N-1 secure operating condition into one that leads to cascading trips. This motivates a proposed criterion for setting oscillation amplitude limits: permissible amplitudes should be such that no credible contingency produces failures that would not occur in the oscillation's absence.

The remainder of the report is organized as follows. Chapter 2 provides a brief overview of power system oscillations and the challenges associated with capturing their impact accurately in planning studies. Chapter 3 describes the risk assessment methodology and RATLLE script architecture, alongside outlining opportunities for further refinement. Besides detailing the tool implementation, the chapter also provides insights into design choices made during development and their limitations. Chapter 4 illustrates case studies, while Chapter 5 outlines future directions and concludes this report.

¹<https://github.com/pnnl/LL-risk-assessment>

2.0 Power System Oscillations

Understanding power system oscillations requires two complementary perspectives: a causal view that identifies the underlying mechanisms driving oscillatory behavior, and a frequency-based view that highlights the physical processes and system components most likely to be involved and impacted. From a causal standpoint, oscillations can be natural or forced (Rogers et al., 2025).

Power system dynamics can be analyzed by linearizing the system around an operating point. Like any linear system, the relationship between inputs and outputs is governed by the system's modes - the eigenvalues of the state matrix. Each mode is characterized by three properties: frequency, damping ratio, and mode shape. The damping ratio determines how quickly oscillations at a given frequency decay once excited. The mode shape describes the relative participation of different power system elements in that oscillation. Together, these properties emerge from the interactions among generators, loads, and control loops that define the system state matrix.

Natural oscillations arise from these inherent system modes. During normal operation, stochastic load and generation variations act as a broadband noise input, continuously exciting low-level oscillations that are always present under ambient conditions. Transient disturbances such as faults or generation loss act as impulse-like inputs and produce ringdown oscillations that decay according to the modal damping. In both cases, the oscillation properties - frequency, damping, and shape - are inherited from the system's modes.

Forced oscillations (FOs), by contrast, arise when a persistent periodic input is applied to the power system. Common sources include malfunctioning control equipment, cyclic industrial or computational loads, and resonant interactions with converter-based resources. Unlike natural oscillations, the frequency and characteristics of forced oscillations are determined by the forcing input rather than the system's own modes, and the oscillations persist for as long as the forcing input is active. If the forcing frequency coincides with a system mode, the resultant oscillations may be amplified due to resonance.

Figure 2 illustrates these phenomena using a simplified example. Consider a power system represented as a linear system with three modes: (0.2 Hz, 9% damping), (0.64 Hz, 3% damping), and (1.2 Hz, 5% damping). A random noise input excites all three modes simultaneously, producing small broadband oscillations in the output that reflect the superposition of the individual modal responses. A step change produces ringdown oscillations in which all three modes are excited and then decay, with the speed of decay determined by the damping ratio of each mode. In both cases, the output is a superposition of the system's modal responses, as evident from the broad peaks at mode frequencies in the output responses' power spectral density plots. The modal response is clearer in the ringdown oscillations - and hence extracting mode estimates from ambient measurements is considered a more challenging problem.

In the same system, a sinusoidal input at 0.4 Hz produces steady-state output oscillations (forced) at the same frequency once the initial transient decays. Since 0.4 Hz does not coincide with any system mode, the output amplitude remains modest - attenuated from the input amplitude. A sinusoidal input at 0.64 Hz, however, directly excites the lightly damped 3% mode, and the output amplitude grows substantially - illustrating the amplification that occurs when the forcing frequency coincides with a system mode.

2.1 Oscillations in Different Frequency Bands

The frequency of oscillations is a valuable diagnostic: it provides insight into the underlying physical phenomena, identifies the power system elements most likely to be involved, determines how far oscillations can propagate spatially, and guides the choice of modeling approach needed




Figure 2: Illustrating natural and forced oscillations with the example of a simplified power system representation. The power system is represented by a linear system with three modes - (0.2 Hz, 9% damping), (0.64 Hz, 3% damping), and (1.2 Hz, 5% damping).

to capture them accurately. Oscillation frequencies are commonly grouped into four bands, as shown in Table 1 (NERC, 2021). The typical phenomena associated with each band, and the reliability risks they introduce, differ considerably. Studying risks associated with high-frequency subsynchronous interactions requires high-fidelity electromagnetic transient (EMT) models, while positive-sequence phasor-domain studies are adequate for the slower dynamics in bands 1–3. The frequency band in which computation cycle-driven oscillations fall therefore determines which power system elements are most likely to be impacted. It is worth noting that the control loops of inverter-based resources (IBRs) span a wide frequency range, and hence interaction with IBR controls is possible across all four bands.

2.2 Studying Oscillation Impact in Planning Studies

As discussed in Chapter 1.0, there is a lack of standardized study procedures to evaluate the wide-area reliability impact of load-induced low-frequency oscillations. Prior work has explored both small-signal stability studies and explicit time-domain simulations to characterize this impact (Biswas et al., 2026; Ko and Zhu, 2026).

Small-signal stability studies can identify dominant system modes, their excitability, controllability, and damping ratios, helping planners understand which frequencies can be excited by a load at a given interconnection point. Commercial tools such as Powertech Lab's Small Signal Stability Analysis Tool (SSAT) support mode estimation through direct eigenanalysis, and methods to derive mode properties from time-domain simulations are also used (WECC, 2021).

Table 1: Oscillations in different frequency bands

Band	Frequency range	Characteristics	Risk from LL-induced forced oscillations	Modeling approach
1	0.01–0.15 Hz	Typically linked to slow dynamics due to governor actions, automatic generation control (AGC), and plant controller malfunctions. Wide spatial propagation expected.	Slow demand fluctuations may cause inaccurate AGC actions or interact with frequency regulation control loops. Mistuned governor parameters can cause hydropower plants to provide negative damping in this band.	Positive sequence
2	0.15–1 Hz	Typically linked to electro-mechanical inter-area and local plant oscillations. Wide spatial propagation expected. Hydropower rough zone operations and power plant control malfunctions have been known to introduce forced oscillations in this band. If the forcing frequency coincides with an inter-area mode, oscillations can be amplified away from the source.	Resonance with local or inter-area modes can cause oscillation amplification, increasing the likelihood of protection and control misoperations.	Positive sequence
3	1–5 Hz	Typically linked to local and intra-plant modes, excitation controls, and DC circuit controls. Wide-area propagation not expected; oscillation impact is mainly visible in the vicinity of the source.	Interaction with co-located or nearby generator and inverter controls.	Positive sequence
4	>5 Hz	Typically linked to generator torsional modes, subsynchronous interactions, and fast inverter control loops. Oscillations remain local and can cause power quality problems.	Interaction with local UPS and inverter controls. If oscillation frequencies coincide with torsional modes of a nearby generator, increased equipment fatigue and shaft failure risk possible.	EMT studies required

These studies are valuable for identifying the frequencies of greatest concern where forced oscillation amplification is possible. However, they do not directly indicate which transmission lines will experience high power swings due to these oscillations. Nor do they reveal how individual generators will respond to oscillations at frequencies not coinciding with system modes without




Figure 3: Response to non-sinusoidal forcing inputs at resonant frequencies

explicit transfer function estimation. Some commercial small signal analysis platforms may also face compatibility issues if the models being analyzed include many user-defined dynamic models in non-native formats.

A further complication is that load-induced oscillations are rarely ideal sinusoids. The actual system response is therefore a superposition of responses across multiple frequencies contained in the forcing function, each interacting differently with system modes. An example with the same linear system as Section 2.1. is shown in Figure 3. It can be seen that the resonance-driven amplification for square/triangular inputs differs from a pure sinusoidal input as shown in Figure 2. Further, additional frequency content is present in the outputs compared to a sinusoidal forcing function². It might be challenging to capture this multi-frequency character using eigen-analysis alone, necessitating validation using time-domain simulations. Whether load-induced active power swings can be large enough to drive the system outside the small-signal regime and introduce nonlinear behavior is also an open question. Explicit time-domain simulations can address these limitations, but become computationally expensive for large interconnection-scale models.

The risk assessment framework presented in this report utilizes both small-signal analysis principles and time-domain simulations, using each where it is most informative. Two limitations are worth stating before describing the methodology in further detail. The limitations are inherent to planning-stage assessments of this kind, and both shape the design choices described in Chapter 3:

- Planning model limitations:** The usefulness of any simulation-based risk assessment ultimately depends on the accuracy of the underlying model. Interconnection-level planning models typically do not include detailed representations of all protection systems and remedial action schemes, and thus the direct impact of load-induced oscillations on protection and control actions cannot be studied within the simulation environment. These models span territories that different transmission operators and reliability coordinators are responsible for, and the associated modeling limitations cannot be resolved unilaterally by the planning engineer conducting the risk assessment. In the RATLLE framework, this limitation is partially addressed by deriving generic protection element settings and estimating their trip likelihood through post-processing of time-domain simulation outputs outside the simulation platform.

²By Fourier's theorem, any periodic non-sinusoidal waveform can be represented as a sum of sinusoids at the fundamental frequency and its integer multiples (harmonics), each with a specific amplitude and phase. A square wave contains only odd harmonics with amplitudes decaying as $\frac{1}{n}$, while a triangle wave also contains only odd harmonics but with amplitudes decaying faster, as $\frac{1}{n^2}$. Here, n is the harmonic number.

- *Determining quantitative risk thresholds:* Defining quantitative thresholds for oscillation risk is inherently challenging. NERC transmission planning standards do not prescribe explicit limits for oscillatory behavior beyond requiring that oscillations neither grow unboundedly nor persist beyond a set time following a disturbance — criteria designed for transient stability rather than sustained forced oscillations. Moreover, the impact of oscillations is not uniform across operating conditions; system stress reduces operating margins and can increase vulnerability to oscillation-induced disturbances. The RATLLE framework uses a three-stage severity scale that provides a structured, consistent basis for risk quantification — discussed in further detail in Chapter 3.

3.0 Methodology

Addressing the reliability risks posed by large-load-induced oscillations requires a structured, repeatable methodology that transmission planners can integrate into existing interconnection and planning study workflows. To this end, RATLLE is a Python-based script suite that automates the evaluation of bulk power system reliability risk from data center oscillations, using PSS/E as the simulation engine. The scripts are written in a modular function-based style, so that they can be easily understood, customized, and extended to suit a utility's needs. A planning case in the form of raw and dynamics files serves as the primary input, and users can pass configuration parameters to tailor the study process by editing csv files. The codebase is hosted on PNNL's public-facing GitHub repository³ with detailed documentation.




Figure 4: RATLLE architecture

Planners may approach the risk assessment problem with two questions:

- How can vulnerable locations in the system be identified efficiently, without exhaustive simulation of every candidate interconnection point?
- Once a high-risk location is identified, how can a detailed site-specific risk analysis be conducted?

RATLLE comprises three modules, designed for screening vulnerable locations, running time-series simulations, and analyzing simulation outputs (Figure 4). The screening module answers the first risk assessment question, while the simulation module tackles the second. The analysis module helps interpret the simulation outputs in terms of risk metrics. Natural decision points exist between consecutive modules - for instance, screening results can inform which studies are run by the simulation module. The rest of this chapter discusses each module in detail.

3.1 Module 1: Screening

The screening module operates in two stages. In the first stage, locations with high voltage and bus angle sensitivities to MW and MVar injections are identified. The motivation is simple - at high sensitivity locations, a persistent load cycling would produce the highest voltage and angle fluctuations. In the second stage, any system modes excitable from the high sensitivity locations are determined. Ideally, planners would conduct this analysis with several seasonal cases.

³<https://github.com/pnnl/LL-risk-assessment>




Figure 5: Voltage sensitivity computed for 69-138 kV buses in WECC (2031 Heavy Winter Planning Case)

3.1.1 Voltage sensitivity screening

The screening methodology is simple – fictitious loads representing ± 1 MW and ± 1 MVar injections are applied at each candidate bus, and a power flow solution is obtained for each perturbation. The resulting bus voltage deviations are recorded, and the voltage sensitivities $\Delta V/\Delta P$ and $\Delta V/\Delta Q$ are computed as the average deviation from the base case across the injection and withdrawal scenarios.

Large loads are typically interconnected to BPS circuits at voltages below 230 kV. By default, RATLLE limits the sensitivity screening to buses at voltage levels between 69 and 138 kV, though users can modify this range as needed. Figure 5 shows the voltage sensitivities computed for 69–138 kV buses in the WECC 2031 Heavy Winter planning case. Buses with switched shunts are excluded from the screening, as the shunts provide reactive power support. These sensitivity metrics can be supplemented with other system strength indicators such as the short circuit ratio (SCR) and its variants.

Although many sophisticated methods have been proposed in the literature to identify weak locations in a power grid, the brute-force approach adopted here was chosen for its implementation simplicity. Computation times are modest - sensitivity computations for 9,600 buses in the WECC case were completed in under 18 minutes on a 13th Gen Intel Core i7-1365U 1.80 GHz processor with 32 GB RAM.

Buses with high voltage sensitivity are identified as candidates for further screening. A 50 MW load step is applied at each candidate bus, and a dynamic simulation is run for 3 seconds following the disturbance. If the simulation remains stable, the location is deemed capable of accommodating a 50 MW step change. If the simulation diverges, the load magnitude is reduced in steps of 5 MW until a stable response is obtained, yielding a conservative estimate of the maximum allowable load step at that location. It is assumed that if a large load is ultimately sited at the bus under study, appropriate network upgrades to maintain local voltages within acceptable

limits would be undertaken prior to interconnection.

3.1.2 Angle sensitivity screening

Following a procedure analogous to the voltage sensitivity screening, generator buses are identified based on their voltage angle sensitivity to active power injection, i.e., the change in voltage angle per unit change in active power ($\Delta\theta/\Delta P$). Sensitivity is computed at all buses hosting generator units whose combined MVA rating exceeds a specified threshold, filtering out smaller units that may exhibit spurious angle sensitivity artifacts due to inaccurate or simplified representations in the planning case. Buses with large units and high angle sensitivity indicate locations vulnerable to power oscillations, and clusters of such buses delineate high-vulnerability areas in the system.

3.1.3 Oscillation excitability screening

Consider a planner conducting a detailed risk analysis for a specific interconnection location. A key question is whether certain frequencies excitable from that location could lead to resonance with system modes. The LILCO study at Brookhaven National Laboratory provides a historical precedent: LILCO first performed a frequency-domain analysis using a state-space representation of their system to identify frequencies of concern, then validated the findings through time-domain simulations (Rostamkolai et al., 1994). Building a state-space representation of a utility-scale system is difficult to automate and may need specialized commercial tools. RATLLE therefore takes a simulation-based approach: a short 0.5 second load impulse, with magnitude below the maximum allowable step determined in the previous screening stage, is applied at the candidate bus, and the subsequent ringdown response is analyzed to extract mode estimates. This simulation-based approach for analyzing system modes has also been utilized in prior art (WECC, 2021). To capture the range of conditions likely to be encountered during operations, it is recommended that this modal analysis be repeated across multiple seasonal planning cases and with relevant contingencies or local element outages represented.

The automated ringdown analysis is conducted as described in Figure 6. First, the post-impulse voltage angle response is first-differenced, z-score normalized⁴, and decimated to 15 samples per second. This detrends the signal and removes high-frequency components. Next, dominant modal frequencies are identified via a zero-padded FFT⁵ – peaks indicate potential modes. If narrow peaks are visible in the frequency spectrum of the ringdown signal, then underdamped modes are likely to be present. If prominent peaks are not observed, then the location does not strongly excite any system modes.

Modal parameters are estimated by applying a nonlinear curve fitting technique whose accuracy is enhanced using information from an initial FFT analysis of the ringdown signal. This involves constructing a Hankel matrix from the ringdown signal. The matrix rank (estimated using singular value decomposition or SVD) is used to determine the number of damped sinusoids needed to represent the signal. The parameter fit itself is performed using Variable Projection (VARPRO), a separable nonlinear least-squares method in which the linear amplitude and phase coefficients are solved analytically at each iteration, reducing the nonlinear search to the natural frequency and damping ratio of each mode (Borden and Lesieutre, 2014). Fitting proceeds in two passes: the first fits the FFT-identified frequency seeds and subtracts the result to expose

⁴Z-score normalization rescales a signal by subtracting its mean and dividing by its standard deviation, yielding a zero-mean, unit-variance series that places signals of differing magnitudes on a common scale.

⁵Fast Fourier Transform




Figure 6: Ringdown analysis process

residual energy for any additional modes required by the Hankel order; the second then fits all seeds jointly against the full signal. Fitted poles are subsequently matched to their nearest FFT seed within a 15% frequency tolerance, and poles with damping ratios exceeding 30% are discarded as overdamped and physically non-oscillatory. The retained modes are ranked by spectral prominence in the observed FFT magnitude, and the top modes are reported with their frequency and damping ratio. An example analysis output is shown in Figure 7. Two dominant modes are estimated - at 0.345 Hz with $\sim 12.5\%$ damping, and at ~ 1 Hz with $\sim 7.5\%$ damping.




Figure 7: Analyzing ringdown response due to a 0.5-second load impulse at bus 6508 in the WECC 240-bus system.

If underdamped modes are observed, then oscillating loads with a matching frequency can pose severe risk to system operations. Therefore, oscillations at these frequencies need to be bounded with stringent amplitude limits to prevent aggravating state trajectories which may ultimately lead to out-of-step conditions, or other protection/control actions. Planners may choose to carry out time-series simulations at these frequencies with various amplitudes to analyze system impact. The module currently does not include automated mode-shape analysis capabilities, although future extensions are planned.

The accuracy of the results obtained from the screening module depends strongly on the underlying model's accuracy. Analysis of PMU data, where available, could help improve confidence in the observations. Typically, grid disturbances excite ringdown oscillations – studying




Figure 8: Sequence of steps in the simulation module

such events occurring in close proximity to the bus being studied could shed light on frequencies excitable from a location.

3.2 Module 2: Simulation

The simulation module automates the execution of time-domain simulations based on user-specified parameters: the oscillation shape, frequency, and amplitude; and the bus at which the load perturbation is introduced. These parameters can be informed either by findings from the screening module or by prior knowledge of the planned load’s operational characteristics — for instance, a planner may already know the interconnection location and expected workload behavior of a prospective data center. Study configurations are passed to the module through a CSV file, allowing for straightforward setup and reproducibility. Each simulation run proceeds through a sequence of steps, as illustrated in Figure 8.

3.2.1 Data center load representation

The major load components in a typical data center include: (a) servers, (b) cooling load for removing heat dissipated in server racks, and (c) auxiliary load like lighting. They are equipped with uninterruptible power supplies (UPS) to protect voltage-sensitive equipment and may also have backup generation. Dynamic models that represent data centers as standard PQ or voltage-dependent loads fail to capture behavior like load ramps and delayed reconnection after grid disturbances. To overcome this, the NERC Load Modeling Working Group (LMWG) recommends utilizing the Power Electronic Reconnecting and Ceasing (PERC1) load model that is now available in the latest versions of leading commercial positive sequence simulators (PowerWorld Corporation, 2025). PERC1 is designed to represent aggregations of power electronic loads such as data centers, electric vehicle charging stations, and cryptocurrency mining facilities. The model captures voltage-dependent real and reactive power response, frequency droop behavior, and cease-and-reconnect logic that governs load disconnection and restoration following voltage disturbances. However, the current PERC1 load characteristic does not include an option to introduce cyclic load profiles representative of AI workloads.

Studying the impact of forced oscillations in planning cases is relatively new to standard industry practice. Although prior industry recommendations, such as those arising from the NERC investigation of the 2019 Eastern Interconnection oscillation event, called for developing mature methods to model forced oscillations in commercial simulators (NERC, 2019), progress in this area remains limited. Earlier work has been largely confined to research settings, either for post-event investigation or for evaluating the performance of oscillation source localization algorithms through synthetic case studies (ENTSO-E Expert Panel, 2026; Joshi et al., 2026; Biswas et al.,




Figure 9: Data center load representation

2025). Modeling approaches have varied considerably, ranging from generator playback functions and user-defined models (UDM) that inject periodic disturbances into governor mechanical power or excitation controls (Maslennikov and Wang, 2022; NERC, 2021), to Python-scripted periodic adjustments of plant controller setpoints (Biswas et al., 2025). To represent AI workloads, RATLLE uses a decoupled load representation as shown in Figure 9. The data center is modeled as two load components: Load1 (id: ll), representing the constant component of demand, and Load2 (id: os), representing the cyclic oscillatory component.

In the planning model, the simulation module first modifies any existing load at the user-specified high-voltage (HV) bus to represent a data center. A medium-voltage (MV) bus connected to the HV bus through a step-down transformer is added. Default transformer parameters are used, as they are not expected to significantly influence oscillatory behavior, though planners with detailed knowledge of the interconnection can override these values. The HV load is moved to the MV bus and its dynamic characteristics are represented using the composite load model (CMLD) with typical data center parameters recommended by the NERC LMWG, as shown in Tables 2 and 3 (NERC Load Modeling Working Group, 2025). The PERC1 model would be a more natural fit for representing data center loads; however, the WECC planning cases available for studies were only compatible with older versions of PSS/E where PERC1 was not available. CMLD was therefore adopted as a pragmatic alternative. Since the tool is modular, the CMLD representation can be readily replaced with PERC1 or any other user-defined model at the user's discretion.

The cyclic load component is represented by a separate constant-P load (Load2) initialized to zero, whose active power is varied periodically during dynamic simulation using PSS/E's load_change_6 command. The grid effectively sees the superposition of the oscillatory load perturbation and the CMLD dynamic response. This representation is an approximation — it does not capture all the complexity of real data center behavior — but it is well-suited for planning-stage risk assessment, where the goal is to characterize the grid's sensitivity to oscillatory demand rather than to replicate device-level load dynamics with high fidelity. The approach is consistent with the reportedly observed behavior of AI training loads, where a cyclic oscillation is overlaid on a base consumption level.

Table 2: Load composition in the CMLD model used to represent data center dynamics

CMLD Component	Description	Data Center Component	Typical Parameter Setting
Motor A	3 ϕ motors with constant torque compressors	Cooling Load	$FMA = 0.09$
Motor B	3 ϕ motors with variable torque compressors (fans)	N/A	$FMB = 0$
Motor C	3 ϕ motors with variable torque compressors (pumps)	N/A	$FMC = 0$
Motor D	1 ϕ motor load	N/A	$FMD = 0$
Electronic Load	Power electronic loads	Server load & CRAH	$FEL = 0.89, PFEL = 1$
Static Load	Generic ZIP load	Miscellaneous load	$PFs = 1, P1e = 0, P2e = 0, P1c = 1, P2c = 0$

Table 3: CMLD switching characteristics parameters representation of large loads

Category	Parameter	Value
Disconnection	V_{d1} (start of electronic load drop)	0.65 pu
	V_{d2} (complete electronic load trip)	0.60 pu
	$VTR1A$ (motor A starts tripping)	0.65 pu
	$VTR2A$ (motor A fully tripped)	0.50 pu
Reconnection	$FRSTPEL$ (electronic load reconnection fraction)	0.25
	$FTR1A + FTR2A$ (motor ride-through fraction)	0.50

3.2.2 Selecting output channels

Before running dynamic simulations, RATLLE selects the output channels to be recorded. This choice involves a practical tradeoff: too few channels risk missing the impact of oscillations on elements outside the monitored set, while too many channels slow down the simulation and produce large output files that are difficult to analyze. The default channel selection logic is designed to capture both local and wide-area oscillation impact, and can be refined by the user based on observations from initial simulation runs to improve computational efficiency. By default, RATLLE records the following elements:

1. In-service loads greater than 50 MW in the source area (bus voltage magnitude, active and reactive power injections)
2. In-service buses above 100 kV in the source area (bus voltage magnitude and angle)
3. In-service generator buses with aggregated generation above 50 MW in the source area (generator terminal voltage, rotor angle, active and reactive power injections)

4. In-service transmission lines above 100 kV in the source area (bus voltage magnitudes, active and reactive power flows)
5. All in-service inter-area tie lines (bus voltage magnitudes, active and reactive power flows)

3.2.3 Cyclic load behavior

Having fixed what will be observed, the final modeling step is simulating the oscillatory computational workloads. Cyclic load behavior is simulated by periodically changing the load value using PSS/E's `load_change_6` command (v35 and above), which modifies load as a constant power characteristic. Users can specify the oscillation shape, frequency, amplitude, and start time. By default, RATLLE sets the dynamic simulation runtime to accommodate eight complete load cycles at the specified frequency. Two oscillation shapes are supported, as illustrated in Figure 10: monophasic square waves and bi-periodic pulse trains. In the bi-periodic representation, slower variations in the 0.1–1 Hz range reflect rest-to-compute transitions, while faster fluctuations in the 5–30 Hz range within compute phases arise from the rapid switching of individual processing modules. The default duty cycle parameter is set at 0.5, but can be modified by users.

Although the approach described above approximates cyclic AI workloads reasonably well, some limitations remain. The constant power load characteristic, while representative of compute load behavior, can introduce numerical instabilities during fault simulations. Additionally, repeated interaction with the simulation environment through Python scripting increases runtime. To address these limitations, PNNL has developed companion PSS/E user-defined models (UDMs) that embed the oscillatory load behavior directly in the dynamic simulation and can be used for standalone scenario simulations outside the RATLLE framework. The UDMs provide additional capabilities including triangular load variations and the ability to play back arbitrary load profiles, and planned work includes migrating the simulation automation in RATLLE from the API-based load-changing approach to the UDMs. A full description of the model, its parameters, and validation against the script-based approach is provided in Appendix A.




Figure 10: Cyclical load patterns that can be simulated in RATLLE

3.3 Module 3: Analysis

As illustrated in Figure 11, the analysis module processes simulation outputs to characterize the impact of a user-selected forced oscillation scenario. To quantify impact, a set of risk metrics is formulated that captures effects on equipment wear and tear as well as the likelihood of protection and control misoperations. Oscillation propagation within the load source area is assessed through the response of major generators and transmission lines, while spread outside the area is captured through tie-line flows. These observables are evaluated against risk metric




Figure 11: Process flow in the simulation and analysis modules

thresholds to identify elements of concern: the greater the number of flagged elements, the higher the risk posed by the oscillation scenario under study.

Elements that significantly exceed the risk threshold also serve as an information proxy for vulnerability — these are the locations where protection and control misoperations are likely, and where the impact of contingencies is most likely to be exacerbated by the presence of forced oscillations. For example, consider a heavily loaded transmission line carrying large MW oscillations. An N-1 contingency that reroutes additional power through that line may simultaneously increase oscillation amplitude, potentially leading to thermal overloads, relay actions, and cascading failures — even though the base system is N-1 secure under normal conditions. Prescribing oscillation amplitude limits at the load site is therefore essential to prevent such compounding failures during operations.

Assessing the impact of sustained oscillations on power grid reliability is an inherently qualitative problem, as the severity of oscillations depends on their duration, the elements affected, prevailing operating conditions, and whether alarms are triggered. A conceptual three-stage severity scale, illustrated in Figure 12, captures this range of impacts and serves as the guiding principle for the risk metrics formulated in this work. These metrics represent one of the first attempts to quantify bulk power system reliability in the context of forced oscillations, and the severity scale described herein extends naturally beyond the planning domain that is the primary focus of this work.

Stage 1 in the severity scale maps to small oscillations below control room alarm thresholds that may go undetected for extended periods. Despite their low amplitude, prolonged oscillations can degrade reliability through cumulative equipment fatigue, accelerated aging of battery energy storage systems (BESS) through micro-cycling, and operator desensitization - all of which erode the system's capacity to respond to subsequent disturbances. Stage 2 represents a transition to active system stress, where oscillation magnitudes rise to levels that may trigger control room alarms. At this stage, the system becomes increasingly vulnerable to hidden failures and device misoperations, and the risk of incorrect mitigation actions by operators or generator owners compounds the hazard. Contingencies may be exacerbated by the presence of oscillations - threatening N-1 security. Stage 3 characterizes imminent collapse, where a triggering event precipitates rapid failure propagation with insufficient time for operator intervention. Sustained oscillations frequently expose latent vulnerabilities (e.g., relay misoperations, faulty control logic) that can precipitate cascading failures (ENTSO-E Expert Panel, 2026). Such hidden failures are most consequential when operating reserves are low, with the worst case being a widespread blackout.

Stage 1 Latent failures accumulate	
Manifestation	Small recurring oscillations that do not cause alarms
Worst-case outcome	Symptom of hidden failures leading to a catastrophic event if left unaddressed
Immediate consequences	• Equipment fatigue • Accelerated BESS aging (micro-cycling) • Affects grid monitoring (e.g. mode meters) • Challenges maintaining voltage schedules • Operator desensitization
Stage 2 Early warning signs and near-misses	
Manifestation	Oscillations that cause control room alarms
Worst-case outcome	Oscillations and other system factors lead to a blackout
Immediate consequences	• Stage 1 consequences • System vulnerable to hidden failures and device misoperations • Potential for incorrect mitigation by operators/generator owners • Could be one contingency away from tipping point
Stage 3 Fast propagation — imminent collapse	
Manifestation	A triggering event precipitates fast propagation of failures
Worst-case outcome	Blackout!
Immediate consequences	Too late for operator intervention

Figure 12: Conceptual illustration mapping oscillation impact to progressive severity stages

The risk metrics in the following section operationalize this severity scale as follows. Some thresholds are deliberately set below typical control room alarm levels; a threshold crossing therefore indicates that an element has entered Stage 1 — oscillations large enough to drive cumulative wear, headroom erosion, or repeated control excursions, yet likely to escape operator attention. Crossings at levels comparable to control room alarm settings or the NERC SMWG reportability criteria (e.g., line swings above 20 MW, voltages outside operating bands) indicate Stage 2, where misoperation risk becomes active and contingencies may be exacerbated. Stage 3 is not assessed through a static threshold: it is established by overlaying credible contingencies on the active oscillation scenario and observing whether element trips or cascading failures result, as demonstrated in Chapter 4.0. Each metric in Tables 4 and 5 is accordingly tagged with the severity stage(s) it primarily indicates.

3.3.1 Risk metrics

Prior literature offers little in the way of quantified oscillation impact metrics. NERC Synchronized Measurements Working Group (2024) provides a useful benchmark through their oscillation reporting guideline that defines a reportable oscillation event as one satisfying at least two of the following criteria :

- **Duration and magnitude:** duration of at least 5 minutes; peak-to-peak amplitude of at least 20 MW or MVar on a single element; or large magnitude of at least 50 MW or MVar peak-to-peak on a single element.
- **System impact:** caused transmission system impacts such as load or generation loss, equipment damage, or cascading outages; caused a generation loss of at least 50 MW; or exhibited

large geographic or electrical spread, affecting multiple Transmission Operators or exciting an inter-area mode.

- **Unique phenomena:** oscillation causes or effects that may be informative to the industry.

The risk metrics formulated in this work to quantify oscillation severity are listed in Tables 4 and 5, categorized by whether they affect generation units or other power system elements. Each metric is presented alongside its anticipated near-term operational impact and the latent failure modes it may expose, recognizing that oscillations need not cause failures directly to activate pre-existing vulnerabilities. The formulation draws on a rigorous review of post-event investigation reports published by system operators and discussions with industry practitioners. Default threshold values are assigned to each metric, though they will ultimately depend on the risk appetite of the organization conducting the analysis. All swing amplitudes are calculated as the difference between maximum and minimum values in each oscillation cycle.

Table 4: Risk metrics for generators

Simulation observations	Severity stage	System impact	Potential hidden failures
No. of units with P swing > 5 MW or 3% Rated MVA	1	Equipment fatigue, reduced headroom, accelerated BESS aging due to microcycling	Control valve malfunction, AGC hunting, plant controller logic problems
No. of units with Q swing > 5 MVar or 3% Rated MVA	1	Equipment fatigue, reduced headroom	Excitation control failures, excitation limiters/protection misoperation
No. of units with Pmax/Pmin or Qmax/Qmin violations	2	Reduced headroom	Limit cycle logic or reset timers may lead to unpredictable system behavior
No. of Terminal Voltage Swings >0.05 p.u.	2	Challenge maintaining voltage schedule, excessive tap change operations	Inverter controller misoperation
No. of units with frequency outside 60 ± 0.036 Hz	1	Reduced frequency response availability, repeated excursions outside governor deadband	Control valve malfunction, inverter controller misoperation
No. of units with frequency swing >0.036 Hz	1	Oscillation damping mechanisms may be activated	Control malfunction, overspeed protections may act erroneously, PLL failures
No. of units with rotor angle swings >10 degrees	1	Equipment fatigue	Inverter protection misoperation, out-of-step protection misoperation

3.3.2 Visualizing the risk metrics

The RATLLE scripts automatically calculate risk metrics from simulation outputs and generate interactive, self-contained HTML dashboards. Users can adjust the metric threshold values and see how the number of impacted elements changes. The dashboard supports multi-level interrogation of simulation data without requiring additional simulations or customized data analysis scripts, and includes time-series plots for the most impacted system elements. Example dash-

Table 5: Risk metrics for other elements

Simulation observations	Severity stage	System impact	Potential hidden failures
No. of lines with P/Q swing >20 MW	2	Reduced transfer capacity, inefficient operations	Inadvertent zone 3 trips, RAS/SPS
No. of buses with voltages outside 0.95–1.05 p.u.	2	Reduced reserves	Inadequate voltage support availability
No. of buses with voltage swings >0.03 p.u.	1	Excessive tap change operations	Inverter controller misoperation

board screenshots for a scenario simulated on the WECC 240-bus model are shown in Figures 13 and 14.

The visualization dashboard helps users identify the most impacted elements and assess which contingencies are most likely to cause subsequent failures. Automating this contingency identification is an ongoing area of development for RATLLE, but at present, engineering judgment is required. Contingency events can be overlaid on an ongoing oscillation simulation, and any protection activation is discernible from the simulation log file.

As discussed earlier, planning models may not accurately represent all protection elements. The RATLLE framework approximates the likelihood of protection operations as follows:

1. Voltage and frequency deviations are analyzed to assess the likelihood of under/overvoltage and under/overfrequency protection operations.
2. Zone 3 distance relay misoperations are among the protection events most likely to precipitate cascading trips, yet Zone 3 relays are frequently absent from planning cases. RATLLE therefore post-processes simulation outputs to compute the impedance trajectory for selected lines, estimate likely Zone 3 reach settings, and determine whether the Zone 3 element would have picked up. Generic reach settings are derived from NERC loadability criteria (maximum permissible reach before load encroachment): Zone 3 relays must not operate when a circuit is loaded at 150% of its highest seasonal facility rating (Rate B), with loadability evaluated at 0.85 p.u. voltage and a 30° power factor angle (NERC System Protection and Control Working Group, 2025). If the computed impedance trajectory enters the relay characteristic, a misoperation is considered likely. The intended workflow is for the user to identify heavily loaded lines with significant power swings via the RATLLE risk visualization dashboard, then estimate Zone 3 misoperation likelihood using the distance protection subroutine.

It must be acknowledged that since RATLLE uses the maximum permissible reach philosophy, the estimated relay reach is very likely an overestimate. Actual relay reach settings would be lower.

3. RATLLE allows users to post-process simulation outputs and assess whether remedial action schemes⁶ (RAS) may arm due to the simulated oscillations. At present, only simple threshold-based rules can be evaluated — of the form: if voltage, frequency, or active/reactive power magnitudes exceed specified thresholds for longer than a predetermined duration, the RAS is considered armed.

⁶also called Special Protection Schemes or SPS




Figure 13: Risk metrics visualization for a scenario simulated in the WECC 240-bus model




Figure 14: Visualizing transmission lines with highest power swings for a scenario simulated in the WECC 240-bus model

4.0 Case Studies

This chapter illustrates the RATLLE risk assessment pipeline using two system models: the WECC 240-bus reduced model (Yuan et al., 2020) and the WECC 2031 Heavy Winter planning case. Although findings from the 240-bus model do not directly translate to the actual WECC system, it serves as a useful medium for demonstrating the methodology. The model is publicly available, so all results presented here can be independently reproduced. Moreover, its modest size allows new users to familiarize themselves with the RATLLE script suite before applying it to larger operational models. For simulations in the WECC planning model, element identifiers are anonymized.

All simulations are carried out using the Siemens PTI PSS/E platform - version 34 and 35 for the interconnection-scale and reduced models respectively. A few Python API commands change between the two PSS/E versions.

4.1 WECC 240-bus Model

The 240-bus reduced WECC model comprises about 140 GW of generation with a realistic fuel mix. As a researcher-oriented test system, it aggregates generators and lacks operation-specific information such as line emergency ratings and protection configurations, meaning detailed generator interactions and protection assessments cannot be replicated. Case studies on this model are therefore illustrative rather than operationally representative, and serve primarily to demonstrate RATLLE's workflow before applying it to the full planning model. Some of the presented results have been reproduced from the authors' earlier work (Biswas et al., 2026).

4.1.1 Example 1: 1 Hz mode

From Bus 1302 (H ALLEN), three types of 100 MW load perturbations were simulated: (a) 0.5 Hz square wave, (b) 1 Hz square wave, and (c) a bi-periodic load pattern with the slow and fast frequencies of 0.1 and 1 Hz, respectively.




Figure 15: Grid elements where oscillations higher than 20 MW are recorded when 100 MW load perturbations are introduced from the H ALLEN bus

Figure 15 plots all elements experiencing oscillations above 20 MW, and illustrates two key trends. First, spatial propagation is wider at lower forcing frequencies. Second, amplification is highest at 1 Hz, consistent with the presence of a system mode excitable from H ALLEN at that frequency. The bi-periodic case impacts more elements than either monoperiodic case




Figure 16: Oscillations observed in generator power outputs when 1 Hz, 100 MW load perturbations are introduced from the H ALLEN bus. Mean values subtracted for easy visualization.




Figure 17: 1-Hz mode shape plot. Circles mark generator locations. Radii are proportional to mode shape magnitudes (measure of observability). Arrows show mode shape angles. Generators with similar angles swing together.

alone, likely because the 0.1 Hz component additionally excites locations in the hydro-rich Pacific Northwest. Figure 16 shows the resulting generator power swings for the 1 Hz perturbation.

The estimated shape of the 1 Hz mode is shown in Figure 17. The mode is observable primarily in the Southwest, with generators in Nevada swinging against those in Southern California, Arizona, and New Mexico. Mode shape magnitude — a measure of observability — is highest at the H ALLEN bus, and oscillation amplification is correspondingly highest when the forcing originates there. Figure 18 shows that a linear relationship exists between the mode shape magnitude at the source location and the maximum observed amplification when oscillating at modal frequency. This is a practically useful result: where mode shape or participation factor analysis has already been performed, it can serve as an initial screening criterion for identifying high-risk load cycling frequencies at candidate interconnection locations. Mode shape information is also directionally symmetric — it not only reveals which elements a given source location can impact, but equivalently, which load locations pose the greatest oscillation risk to a given site.




Figure 18: Max. oscillation amplification vs. mode shape magnitude of source




Figure 19: Grid elements where > 20 MW oscillations are observed when 0.38 Hz, 100 MW load perturbations are introduced from the CANADA bus

4.1.2 Example 2: 0.38 Hz mode

The 0.38 Hz mode in the 240-bus model is representative of the WECC North–South (N-S) mode and is the most geographically widespread mode in the system. A 100 MW load perturbation at this frequency was introduced from the CANADA bus in the British Columbia zone, which was found to have a high mode shape magnitude for the N-S mode. As shown in Figure 19, oscillations exceeding 100 MW peak-to-peak are observed on key transmission lines in the Northwest corridor connecting Canada, the Pacific Northwest, and Northern California. Power swings above 20 MW are observed in generators as far south as Arizona - a striking result given that the N-S mode is well-damped in this model.

Note that observations from the reduced WECC model cannot be directly extrapolated to detailed WECC planning cases.




Figure 20: Modal analysis of load impulse response from the California bus considered

4.2 WECC Planning Model

The 2031 Heavy Winter planning model used comprises about 160 GW load, with $\sim 12.5\%$ generation contributed by non-synchronous resources. The RATLLE pipeline was applied to the WECC model in the following manner. First, the voltage and angle sensitivity screening steps identified the 20 highest-vulnerability buses, several of which were spatially clustered. Next, load impulses were applied at these locations, and their ringdown response was analyzed to identify excitable modes. Buses from which lightly-damped modes were excitable were selected for detailed time-domain simulation. The intent was to efficiently identify interesting case studies, rather than to comprehensively characterize risk across the interconnection. This process helped identify two interesting cases that are discussed next. The case studies presented demonstrate the range of outcomes the RATLLE framework can help analyze - one in which a contingency converts comparatively local forced oscillations into cascading failure, and one in which wide-area propagation occurs but reliability risks remain contained.

4.2.1 Example 1: N-1 security constraint violation

The first example is for a location in California. Modal analysis identified two excitable frequencies from this location — at approximately 1.15 and 1.7 Hz (Figure 20), both with damping ratios below 10%. Based on this result, square wave 50-MW load perturbations were introduced at 1.15 Hz.

This scenario led to risk threshold violations on several system elements, as summarized in Table 6. The relatively modest 50-MW swings had a pronounced impact: four monitored transmission lines experienced active power swings exceeding 10% of their thermal rating (Rate A), nine generators had active power swings above 3% of their MVA ratings, and two generators had reactive power swings exceeding 3% of their rated MVA capacity. The highest active and reactive power swings recorded on any single generator were approximately 13.1 MW and 12.6 MVar (on a 134 MVA machine), respectively, while the highest line power swing reached approximately 185 MW (line normal rating 296 MVA).

Table 6: Risk assessment for Example 1

Element	Risk Metric	No. of Impacted Elements
Generator	Active power swing >5 MW	14
	Active power swing >3% of rated MVA	9
	Reactive power swing >5 MVar	1
	Reactive power swing >3% of rated MVA	2
	Terminal voltage swing >0.05 p.u.	0
	Frequency outside 60 ± 0.036 Hz	2
	Frequency swing >0.036 Hz	7
	Rotor angle swing >10 degrees	1
Lines	Active power swing >20 MW	28
	Active power swing >10% of rated capacity	4
HV Buses	Voltage excursion outside 0.9–1.1 p.u. band	4
	Voltage swings >0.05 p.u.	3

The approximate geographic locations of impacted elements are shown in Figure 21, where line widths are proportional to active power swing amplitude. While the impact is not interconnection-wide, significant oscillations are observed across California. Time-series plots for the most impacted generator (Generator A) and transmission line (Line A) are shown in Figure 22. Oscillation amplitudes are observed to grow over the simulation period.

To assess N-1 vulnerability, two contingencies were simulated two seconds after the oscillation onset. In the first, Generator A was tripped. As Figure 23 shows, this precipitated growing oscillations in the output of a nearby Generator B, which was subsequently tripped by under-voltage relay action approximately six seconds later. Generator B had not been included in the initial monitored set as its base-case output was about 40 MW, lower than the 50 MW criteria used to select monitored generator buses. The protection action was identified by analyzing the simulation log file. Closer inspection revealed that even without contingencies, Generator B was already experiencing active power oscillations exceeding 15 MW — larger than those on Generator A.

The second contingency simulated was a Line A outage, which produced similar cascading behavior. As shown in Figure 24, the line trip was followed by growing oscillations in both Generator A and B, with Generator B again tripped by undervoltage protection at approximately nine seconds. After Generator B was lost, growing oscillations ceased in Generator A's response, though forced oscillations persisted in the system.

In the base case without oscillations, neither contingency caused overloads or element trips — the system was N-1 secure. The introduction of 50 MW forced oscillations from the load bus changed this: the same contingencies now led to cascading generator loss, constituting a violation of the N-1 security criterion. In severity-scale terms, the pre-contingency oscillation condition corresponds to Stage 2 — multiple alarm-level threshold crossings — while the contingency overlay demonstrates Stage 3 conditions: a triggering event producing cascading generator loss. The potential of a N-1 security constraint violation yields a physically meaningful and operationally grounded way to set acceptable oscillation limits: load oscillations must not, under any credible contingency, precipitate element trips or cascading failures that would not have occurred in their absence.






Figure 21: Approximate locations of power systems elements that cross the RATLLE risk metric thresholds in Example 1


The present work selected contingencies by tripping the elements most impacted by the oscillation — a reasonable but ad hoc approach. A systematic methodology for identifying credible contingencies likely to produce N-1 violations under active oscillation scenarios remains an opportunity for future development of the RATLLE framework.



(a) Generator A active and reactive power output



(b) Generator A terminal voltage



(c) Line B active power flow

Figure 22: Time-series plots for the most impacted generator and transmission line in Example 1




Figure 23: Contingency 1 – Generator A loss leads to the tripping of Generator B due to under-voltage protection action

4.2.2 Example 2: Inter-area propagation

A cluster of high voltage sensitivity locations was identified in the Wyoming–Colorado region. The second example introduces forced oscillations from one of these locations. The RATLLE screening module identified an excitable frequency at approximately 0.38 Hz, coinciding with the WECC East–West A (E-W A) inter-area mode. However, the estimated damping ratio exceeded 15%, indicating a heavily damped mode. Oscillations at this frequency were nonetheless simulated, to evaluate the geographic spread expected when an inter-area mode is excited.

As in Example 1, 50-MW square-wave load perturbations were introduced at 0.38 Hz. The approximate locations of elements where at least one risk metric threshold was violated are shown in Figure 26. The spatial spread of impacted elements is considerably wider than in Example 1, with significant oscillation amplitudes visible on a transmission line between Idaho and Oregon. The geographic distribution of the impacted elements agrees with the E-W A mode shape results reported in prior work (WECC, 2021). This demonstrates that inter-area mode excitation can drive wide-area oscillation propagation even when the mode is heavily damped, though unlike Example 1, oscillation amplitudes did not grow significantly over the simulation period.

The total count of impacted elements is lower than in Example 1, as summarized in Table 7. The highest active and reactive power swings recorded at any generator were approximately 17 MW and 16.5 MVar, respectively. The highest line active power swing reached approximately 27




Figure 24: Contingency 2 – Line A loss leads to the tripping of Generator B due to undervoltage protection action

MW, with a corresponding reactive power swing of approximately 15.5 MVar.

A contingency was simulated for Example 2 by tripping the generator with the highest active power swings in the oscillation scenario (Generator C). Generator C was a ~1000 MVA unit experiencing only 15 MW swings, meaning the system response was dominated by the generation loss itself rather than any oscillation interaction. Post-contingency oscillation impacts did not vary significantly and the system remained N-1 secure — in contrast to Example 1. As shown in Figure 27, while power flow is redistributed following the trip, oscillation amplitudes are largely unchanged. The contingency analysis was not exhaustive; the result nonetheless suggests that the heavily damped mode and modest per-unit oscillation amplitudes relative to Generator C's rating limit the compounding risk that was observed in Example 1.

This example remains within Stages 1–2: thresholds are crossed across a wide geographic footprint, but the contingency overlay does not produce trips or cascades, and the immediate reliability consequence is limited to wide-area equipment stress and monitoring burden.




Figure 25: Modal analysis of load impulse response for Example 2 in the WECC planning model




Figure 26: Approximate locations of impacted elements in Example 2

Table 7: Risk assessment for Example 2

Element	Risk Metric	No. of Impacted Elements
Generator	Active power swing >5 MW	6
	Active power swing >3% of rated MVA	2
	Reactive power swing >5 MVar	11
	Reactive power swing >3% of rated MVA	1
	Terminal voltage swing >0.05 p.u.	0
	Frequency outside 60 ± 0.036 Hz	0
	Frequency swing >0.036 Hz	3
	Rotor angle swing >10 degrees	0
Lines	Active power swing >20 MW	5
	Active power swing >10% of rated capacity	2
HV Buses	Voltage excursion outside 0.9–1.1 p.u. band	4
	Voltage swings >0.05 p.u.	0




Figure 27: Time-domain simulation plots in Example 2 after the Generator C loss contingency

5.0 Implications and Future Directions

The computational workloads of hyperscale AI data centers present a forced oscillation risk that existing planning frameworks were not designed to address. This report has developed and demonstrated RATLLE, a structured, automatable methodology for evaluating that risk in bulk power systems. Chapter 2 established the physical basis for concern — the conditions under which large cyclic loads can excite natural system modes, drive wide-area oscillation propagation, and interact with protection and control systems in ways that standard transient stability analysis does not capture. Chapter 3 described the three-module RATLLE pipeline, from system-wide sensitivity screening through time-domain simulation to risk metric computation, and outlined the modeling choices and limitations inherent in applying these methods to interconnection-scale planning cases. A framework for quantifying system reliability risk due to the presence of sustained oscillations is proposed. The case studies in Chapter 4 then grounded these concepts in simulation results on both a reduced test system and a detailed WECC planning model, producing findings with direct implications for how interconnection studies should be conducted and how oscillation amplitude limits should be set. The remainder of this chapter synthesizes those findings and identifies directions for future development.

5.1 Key Takeaways

1. Forced oscillations from large loads are a credible planning-stage reliability risk.

The California case study in Section 4.2.1 illustrates that 50 MW of oscillatory demand, well within the range of plausible AI workload fluctuations, can trigger N-1 security constraint violations that would not otherwise occur. An operating condition deemed secure under standard contingency analysis can become failure-prone when sustained oscillations are superimposed. This has direct implications for planning practice: system studies involving hyperscale facilities should assess N-1 security not only under steady-state or transient conditions but also across representative forced-oscillation scenarios. There is also an opportunity to develop automated methods for identifying the contingencies most susceptible to N-1 constraint violations in the presence of sustained oscillations, so that exhaustive analyses do not need duplication.

Nonetheless, reaching consensus on the specific modifications needed in current study practices remains an open challenge. Planners may have limited visibility into facility workload characteristics at the interconnection stage; those profiles can evolve over the load's operational lifetime; and as new resources are integrated into the grid, system modes may shift themselves — potentially rendering earlier interconnection study conclusions outdated.

2. Reliability risk is frequency-dependent and location-dependent.

The screening module results and case studies together reinforce that oscillation risk is not simply a function of load magnitude. A 50 MW oscillation at a frequency coinciding with a lightly damped local mode, from a location with high mode shape magnitude, can produce larger system-wide impact than a 100 MW oscillation at an off-resonance frequency from a weaker coupling point. Utilities and hyperscale facility operators can work together to understand if modifying the facility workload characteristics can help in accommodating higher oscillation amplitudes - a historical precedent is provided by the LILCO experience with interconnecting the Brookhaven synchrotron described in Chapter 1.

3. Inter-area mode excitation warrants attention even when modes are well-damped.

The Wyoming–Colorado example in Section 4.2.2 shows that forced oscillations can drive wide-area propagation, with impacted elements visible across multiple transmission areas, even when the excited mode is well-damped. The system remained N-1 secure in this case, and os-

cillation amplitudes did not grow, suggesting the immediate reliability risk was limited. However, the geographic spread of impacted elements underscores that persistent oscillations from the interconnected facility would be observed by multiple transmission operators, requiring coordination across organizations. These results suggest that wide-area risk exposure from AI workloads warrants evaluation.

4. The N-1 security criterion provides a principled basis for setting oscillation amplitude limits.

A persistent challenge in this domain, acknowledged in NERC LLWG's white papers, is the lack of quantitative standards for permissible forced oscillation amplitudes. The case studies in this report suggest a physically grounded approach: oscillation amplitude limits at the load interconnection point should be set such that no credible contingency, applied during an active oscillation scenario, produces overloads, voltage limit violations, element trips, or cascading failures that would not occur in the oscillation's absence. This criterion avoids the need to define oscillation limits in isolation from system conditions. Translating this principle into a systematic limit-setting procedure — accounting for seasonal operating conditions, contingency selection, and uncertainty in load frequency content — is an important direction for future work.

5. Planning model limitations constrain but do not invalidate simulation-based risk assessment.

As discussed in Section 2.2, interconnection-scale planning models may not fully capture protection systems, remedial action schemes, or the dynamic characteristics of all load types. The RATLLE framework partially compensates through post-processing — estimating Zone 3 relay reach, assessing RAS arming likelihood, and flagging undervoltage and underfrequency conditions — but these approximations carry uncertainty. The appropriate response is not to forgo simulation-based assessment, but to treat its outputs as conservative indicators that flag elements and scenarios warranting closer examination, rather than definitive predictions of protection system behavior. Where PMU data is available near candidate interconnection locations, comparison against simulated ringdown responses can improve confidence in modal estimates and partially validate the planning model's dynamic representation.

6. Methodology standardization remains the central challenge.

The LILCO synchrotron study of the 1990s demonstrated that rigorous site-specific impact analysis is feasible for oscillations emerging from individual large cyclic loads. What has changed is scale: utilities are now processing hundreds of hyperscale interconnection requests from load owners who have limited visibility into the grid conditions that govern oscillation risk. The RATLLE framework is designed to address this scalability problem by providing a repeatable, automatable workflow that can be integrated into existing interconnection study processes. The scripts are publicly available and documented, and the modular architecture allows utilities to adopt components selectively or substitute their own modeling approaches where needed. Broader adoption will require not only tooling but also the development of industry consensus around study procedures, reporting requirements, and acceptable risk thresholds — work that is underway in the broad industry and that the methodology presented here is intended to inform.

5.2 Future Directions

The findings presented in this report suggest several natural extensions to the RATLLE framework that the authors aim to pursue.

RATLLE currently simulates oscillations injected from a single large load location at a time. As hyperscale AI facilities continue to proliferate, analyzing the impact of synchronized workloads across multiple facilities owned by the same operator becomes increasingly relevant. Extending

the platform to simulate such multi-site synchronized demand patterns, and characterizing their aggregate impact on the grid, is an area of active development.

The authors aim to add an automated contingency selection module for identifying which contingencies are most likely to produce N-1 violations under active oscillation conditions. This would replace the current ad hoc approach and make the risk assessment more systematic and reproducible. On the platform front, RATLLE currently uses PSS/E as the underlying simulation engine; expanding to other widely used platforms such as PSLF, PowerWorld, and TSAT will broaden applicability across the industry. Feedback from early users has also identified several visualization improvements to the risk assessment dashboard that the authors plan to incorporate.

The risk quantification framework and associated thresholds in the current work are empirically derived, and further calibration against industry experience and historical oscillation events is required to build confidence in their general applicability.

Finally, while the current focus has been on the planning stage, opportunities exist to extend these findings to the operations timeframe. The risk quantification framework and associated case studies can inform operator guidance — specifically, helping define when sustained oscillations are severe enough to warrant immediate action as opposed to a monitoring posture.

References

- Biswas, S., X. Lyu, Q. Nguyen, L. He, D. Khan, M. Lu, X. Fan, and S. Kincic (2025). Modeling forced oscillations in bulk power grids with high hybrid power plant penetration. In *2025 IEEE Energy Conversion Conference Congress and Exposition (ECCE)*, pp. 1–6.
- Biswas, S., A. Varghese, K. Chatterjee, B. Ross, and J. Follum (2026). pnnl/ll-risk-assessment (33525-E). *US Department of Energy (DOE) Software*, 68. <https://github.com/pnnl/LL-risk-assessment>.
- Biswas, S., A. C. Varghese, K. Chatterjee, S. Nekkalapu, B. Ross, and J. Follum (2026). Evaluating the risk to bulk power system reliability from large load induced oscillations. In *IEEE PES T&D*.
- Borden, A. R. and B. C. Lesieutre (2014). Variable projection method for power system modal identification. *IEEE Transactions on Power Systems* 29(6), 2613–2620.
- Chatterjee, K., S. Biswas, J. D. Follum, and S. Kincic (2023). Online monitoring applications enabled by phasor measurement units: technical assistance to the power sectors of Southeast Asia. Technical Report PNNL-35221, Pacific Northwest National Laboratory (PNNL), Richland, WA (United States).
- Choukse, E., B. Warriar, S. Heath, L. Belmont, A. Zhao, H. A. Khan, B. Harry, M. Kappel, R. J. Hewett, K. Datta, et al. (2025). Power stabilization for AI training datacenters. *arXiv preprint arXiv:2508.14318*.
- Dorfman, L., P. Riccardella, D. Rosario, and T. Sherlock (1996). Retaining ring failure at Comanche unit 2: root cause analysis. Technical report, Electric Power Research Inst., Palo Alto, CA (United States).
- ENTSO-E Expert Panel (2026, March). Final report on the 28 April 2025 blackout in Spain and Portugal. Technical report, European Network of Transmission System Operators for Electricity (ENTSO-E).
- Joshi, D., M. Kc, H. Wajid, and V. Venkatasubramanian (2026). Descriptor: Dataset for power system forced oscillation responses (DPSYFOR). *IEEE Data Descriptions* 3, 19–24.
- Ko, M.-S. and H. Zhu (2026). Wide-area power system oscillations from large-scale AI workloads. *IEEE Transactions on Power Systems*.
- Maslennikov, S. and B. Wang (2022). Creation of simulated test cases for the oscillation source location contest. In *2022 IEEE Power & Energy Society General Meeting (PESGM)*, pp. 1–1. IEEE.
- Mishra, C., L. Vanfretti, J. Delaree Jr, T. Purcell, and K. D. Jones (2025). Understanding the inception of 14.7 Hz oscillations emerging from a data center. *Sustainable Energy, Grids and Networks*, 101735.
- NERC (2019, December). January 11, 2019 Eastern Interconnection oscillation disturbance. Technical report, North American Electric Reliability Corporation.

- NERC (2021, November). Recommended oscillation analysis for monitoring and mitigation reference document. Technical report, NERC Synchronized Measurements Working Group.
- NERC (2025a, July). Characteristics and risks of emerging large loads: Large loads task force white paper. Technical report.
- NERC (2025b, January). Incident review: Considering simultaneous voltage-sensitive load reductions. Technical report.
- NERC (2026a, March). Assessment of gaps in existing practices, requirements, and reliability standards for emerging large loads. Technical report.
- NERC (2026b, May). Essential action to industry: Computational load modeling, studies, instrumentation, commissioning, operations, protection, and control. Technical report.
- NERC (2026c, May). Risk mitigation for emerging large loads: Large loads working group reliability guideline. Technical report.
- NERC Load Modeling Working Group (2025). Data Center Information Collection Questionnaire. [Online]. Available: <https://www.nerc.com/comm/RSTC/LMWG/Data>
- NERC Synchronized Measurements Working Group (2024). Oscillation report document. Technical report, North American Electric Reliability Corporation.
- NERC System Protection and Control Working Group (2025, February). Determination of Practical Transmission Relaying Loadability Settings. Technical report, North American Electric Reliability Corporation.
- Perić, V. S., X. Bombois, and L. Vanfretti (2017). Optimal multisine probing signal design for power system electromechanical mode estimation. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*, pp. 3156–3164.
- Pinneilo, J. A. and J. E. Van Ness (1971). Dynamic response of a large power system to a cyclic load produced by a nuclear accelerator. *IEEE Transactions on Power Apparatus and Systems PAS-90*(4), 1856–1862.
- PowerWorld Corporation (2025). Power electronic reconnecting and ceasing (PERC1) load model. <https://www.powerworld.com/knowledge-base/power-electronic-reconnecting-and-ceasing-perc1-load-model>.
- Rogers, G. J., R. T. Elliott, D. J. Trudnowski, F. Wilches-Bernal, D. Osipov, and J. H. Chow (2025). *Power System Oscillations: An Introduction to Oscillation Analysis and Control* (Second ed.). New York: Springer.
- Rostamkolai, N., R. J. Piwko, and A. S. Matusik (1994, August). Evaluation of the impact of a large cyclic load on the LILCO power system using time simulation and frequency domain techniques. *IEEE Transactions on Power Systems* 9(3), 1411–1416.
- Solanics, P., K. Kozminski, M. Bajpai, J. Esztergalyos, E. Fennell, J. Gardell, C. Mozina, S. Patel, A. Pierce, V. Skendzic, P. Waudby, and J. Williams (2000). The impact of large steel mill loads on power generating units. *IEEE Transactions on Power Delivery* 15(1), 24–30.
- WECC (2021). Modes of inter-area power oscillations in the Western Interconnection.

Yuan, H., R. S. Biswas, J. Tan, and Y. Zhang (2020). Developing a reduced 240-bus WECC dynamic model for frequency response study of high renewable integration. In *IEEE PES T&D*, pp. 1–5.

Appendix A – User-Defined Model for Oscillatory Load

As described in Chapter 3.0, the RATLLE simulation module introduces cyclic load behavior by repeatedly modifying load values through PSS/E's 'load_change_6' command. This script-based approach treats the load variation as an externally imposed sequence of discrete events, with the simulation stopped, modified, and advanced at each step. It is well-suited to studies involving oscillations at a few locations, but scales poorly when many widely distributed loads must be represented or when complex oscillatory waveforms are required across a large system.

To address this, a user-defined model (UDM) has been developed in PSS/E that injects a forced oscillation, reproducing the cyclic profile of AI training workloads. Unlike the scripted approach, the UDM computes the load variation internally as a function of simulation time. The waveform logic therefore remains synchronized with the dynamic solver and requires no interruption of the simulation run.

A.1 Capabilities

The UDM provides the following options:

- **Load characteristic:** constant power, constant current, or constant impedance. Although data centers are physically constant-power loads, the constant-current and constant-impedance representations are useful under fault or weak-grid conditions, where constant-power loads can introduce numerical instabilities.
- **Oscillation waveform:** mono-periodic, bi-periodic, triangular, or arbitrary play-in profiles.

Input parameters are supplied through the model's DYR data and are described in Table A.1. Time-domain examples of the parameterized load profiles are shown in Figure A.1.

A.2 Implementation

The model is implemented in Fortran and compiled into a DLL using the Intel Fortran compiler within the PSS/E Environment Manager, after which it is available for dynamic simulation. It has been developed and tested in PSS/E versions 34, 35, and 36.

A.3 Validation

Figure A.2 compares the load profiles produced by the UDM for a square-wave oscillation under the three load characteristics. As the zoomed view in Figure A.3 shows, the constant-current and constant-impedance cases do not hold a flat peak amplitude, since the power drawn depends on the voltage at the load terminals — in contrast to the constant-power case. Figure A.4 compares the UDM against the Python API-based scripting approach and shows close agreement, confirming that the UDM reproduces the intended behavior.

A.4 Advantages Over the Script-based Approach

By defining the waveform type and its parameters through the model's DYR data, the oscillatory load UDM can be used to represent different load profiles in different buses without excessively





Figure A.1: Parametrized representation of different oscillatory load profiles in the UDM




Figure A.2: UDM performance for various load characteristics for square wave profile




Figure A.3: Zoomed-in plot of Figure A.2

Table A.1: Description of the UDM Parameters

Parameter	Parameter Description
<i>Flag1</i>	0 (constant power), 1 (constant current), 2 (constant impedance) for P component
<i>Flag2</i>	0 (constant power), 1 (constant current), 2 (constant impedance) for Q component
<i>Flag3</i>	0 (monoperiodic), 1 (bi-periodic), 2 (triangular)
<i>Flag4</i>	0 (Oscillation applied to only P), 1 (Oscillation applied to only Q)
<i>Amp</i>	Amplitude of the oscillation
<i>Amp1</i>	Lower Amplitude of the fast frequency cycles in the bi-periodic waveform. Used only when Flag3 =1
<i>Freq</i>	Frequency of the oscillation
<i>Fast_Freq</i>	Frequency of the fast component in the bi-periodic waveform. Used only when Flag3 =1
<i>NC</i>	Number of cycles that the oscillation persists
<i>D</i>	Duty cycle
<i>D_Fast</i>	Duty cycle of the faster cycles in the bi-periodic waveform. Used only when Flag3 = 1
<i>X</i>	Start time of the oscillation

interrupting the simulation run as the script-based approach required. This improves reproducibility and reduces the scripting overhead and runtime growth incurred when many API-based load changes are issued during a simulation. These properties make the UDM particularly suited to large-scale dynamic screening and sensitivity studies — running multiple contingencies, oscillatory profiles, and data-center penetration levels — where the scripted approach becomes cumbersome.

The model structure is extensible, and capabilities can be expanded to include additional signal types that can be useful for other use-cases. For example, including a multi-sine functionality could be useful for testing approaches related to probing-based system identification on interconnection-scale models Perić et al. (2017) .

RATLLE's current script-based approach remains useful for analyzing a small number of locations at a time. However, the UDM offers advantages in parameterization, scalability, and computational speed. As noted in Section 3.2.3, integrating the UDM as the default perturbation mechanism within the RATLLE pipeline is planned future work.




Figure A.4: Comparing UDM performance with simulations leveraging Python scripts

Pacific Northwest National Laboratory

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnnl.gov