Pacific Northwest
NATIONAL LABORATORY

PNNL-38601

Reliable Integration of AI Data Centers at Scale – Analysis, Modeling and Synthetic Data Generation

Technical Report

October 29, 2025

Sai Pushpak Nandanoori
Veronica Adetola
Bowen Huang
Ramij Raja Hossain

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

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

Printed in the United States of America

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

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

PNNL-38601

Reliable Integration of AI Data Centers at Scale – Analysis, Modeling and Synthetic Data Generation

Technical Report

October 29, 2025

Sai Pushpak Nandanoori
Ramij Raja Hossain
Bowen Huang
Veronica Adetola

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

Pacific Northwest National Laboratory
Richland, Washington 99352

Acknowledgments

This work was performed as part of the Pacific Northwest National Laboratory's Laboratory Directed Research and Development initiative SMARTT-2 Transition.

The authors thank Karan Kalsi, Buxin She, Soumya Kundu, and Brett Ross for their for the fruitful discussions related to this work.

Acronyms and Abbreviations

PNNL	Pacific Northwest National Laboratory
FFT	Fast Fourier Transform
GPU	Graphical Processing Unit
CPU	Central Processing Unit
LEL	Large Electric Load
LSTM	(Long Short-Term Memory)
HVAC	Heating, Ventilation and Air-Conditioning

Contents

Figures

1.0 Introduction

Data centers are rapidly becoming major energy consumers on the grid. In 2023, U.S. data centers consumed around 4.4% of the national electricity use (176 TW) and their consumption is expected to grow substantially, reaching 9% of total U.S. electricity generation by 2030 [1, 2]. Fueled by generative AI and the rapid integration of AI across sectors, this growth introduces greater uncertainty and fast, pronounced demand swings. The combination of large load and distinctive characteristics (e.g., transient swings) poses significant challenges for grid stability and reliability, with implications for national energy security.

Access to high-resolution electricity usage data from data centers is essential for understanding their dynamic behavior, demand flexibility, and grid impacts. However, confidentiality restrictions make it difficult to obtain real-world operational data. High-fidelity synthetic datasets can fill this gap, enabling researchers and grid operators to better model data center loads and develop resilient, adaptive energy systems. Creating accurate synthetic models requires: 1) Capturing complex load dynamics, which are highly variable across multiple timescales due to varying AI workloads; and 2) representing short-term fluctuations and long-term trends in data center-grid interactions to support both real-time operations and strategic planning..

To explore load characteristics of AI-driven data centers, we analyze publicly available datasets, including SURF and the MIT Supercloud. The SURF dataset provides long-term, fine-grained operational server metrics gathered from a scientific computing infrastructure over a period of nearly 8 months at 15-second intervals [3]. The MIT supercloud dataset includes parsed logs from the scheduler, compute nodes, CPU and GPU usage time series data, as well as sensor data from physical monitoring of the facility housing the cluster itself [4]. Both were released to catalyze advances in understanding data center behavior, predicting and mitigating failures, and optimizing operations.

We gathered, cleaned, processed, and classified these open datasets, which provide insights into large-scale HPC and data center operations. These datasets capture CPU/GPU utilization, memory allocation, file system logs, and power consumption at resolutions as fine as 100-milliseconds. Our comprehensive analysis highlights key characteristics of AI data center loads that can help grid operators perform stability studies, prepare for worst-case scenarios, and schedule generation aligned with load patterns.

Focusing on the MIT SuperCloud dataset, we present average daily power-demand variations, maximum ramp-up and ramp-down rates, and correlations among GPU allocation, memory allocation, GPU temperature, memory temperature, and GPU power draw. We also introduce a rudimentary model for generating synthetic power-consumption profiles that can be further refined.

The rest of the report is organized as follows. Section 2.0 and 3.0 provides a summary of SURF and MIT supercloud dataset accordingly providing the key characteristics of AI data center loads. Section 4.0 presents an LSTM based model to generate synthetic data using the available open source datasets. Finally, Section 5.0 summarizes this work.

2.0 Characterizing Large Electric Loads using the SURF Dataset

2.1 Dataset Overview

The SURF machine metrics dataset provides a broad collection of datacenter telemetry collected from the SURFsara Lisa HPC cluster over multiple months, exposed via a Prometheus-based monitoring stack [5, 6]. The public release contains hundreds of machine- and rack-level time series spanning CPU load, memory usage, network throughput, disk I/O rates, power-related metrics, and selected GPU indicators (where available). Machine and rack identifiers are pseudonymized.

A critical constraint for our Large Electric Load (LEL) analysis is that *job-level metadata is not included* in the public archive due to privacy considerations, i.e., the dataset lacks scheduler logs (job IDs, user IDs, start/stop times). Consequently, the dataset enables fine-grained system introspection but does not provide ground-truth job timelines.

2.2 Goal and Use in Our Study

Our objective is to infer the timing and characteristics of LEL episodes—periods corresponding to major jobs or power-intensive workloads—directly from system metrics. We focus on reconstructing job-like intervals and then examining their power/thermal signatures and resource footprints.

2.3 Bottleneck

- **Missing job data:** No Slurm (or equivalent) job records are available in the public SURF release.
- **Scope of metrics:** Only system metrics (CPU load, memory, network, disk I/O, power-related signals, and limited GPU indicators) are provided; there are no per-job labels.

2.4 Heuristic Job Timeline Reconstruction

Without scheduler logs, we estimate job start/end times using multi-signal change detection on resource metrics:

1. **Signal selection:** CPU load, memory usage, GPU utilization (if present), and node/cluster power proxies.
2. **Onset detection:** Identify sustained, concurrent increases in selected signals beyond adaptive baselines (e.g., moving median + robust threshold).
3. **Termination detection:** Identify sustained reversion of signals toward baseline.
4. **Merging/splitting:** Merge adjacent detections separated by short gaps; split plateaus with distinct utilization levels when justified by a clear shift in multiple metrics.

The result is a set of *inferred job intervals* that approximate job execution windows. While heuristic, this timeline allows segmentation of the power/thermal series into candidate LEL episodes for further analysis.

2.5 Findings from the SURF Data Exploration

- **LEL episodes are clearly separable** when GPU utilization and power proxies co-vary strongly, exhibiting step-like transitions and long plateaus.
- **Non-GPU LELs** (CPU/memory intensive) are also detectable via sharp, sustained rises in CPU load and memory occupancy, often coupled with increased disk and network activity.
- **Idle vs. active states** show distinct baselines; the active states often maintain elevated levels for tens of minutes to hours, consistent with large batch jobs/training runs.

2.6 Limitations

- **Resource usage $\neq$ job execution:** System events (maintenance, firmware updates, reboots) can mimic workload signatures, producing *false positives*.
- **Lack of ground truth:** Without job logs, start/end estimates carry uncertainty; overlapping jobs may be conflated or partially separated.
- **Heterogeneity across nodes:** Asynchronous activity and uneven sampling can mask weak or short LELs.

2.7 Code Artifacts and Reproducibility Notes

For documentation, our internal analysis scripts implement the pipeline above: (i) batch ingestion of raw metric CSV/TSDB exports, (ii) resampling/cleaning to uniform cadence (e.g., 0.1–1 s where available), (iii) robust thresholding with hysteresis for onset/termination, and (iv) export of inferred intervals for visualization and downstream modeling.[1]

[1] The repository includes utilities for multi-signal change-point detection, quality checks for missing columns, and figure generation. Paths and environment files are documented in the repo's README.

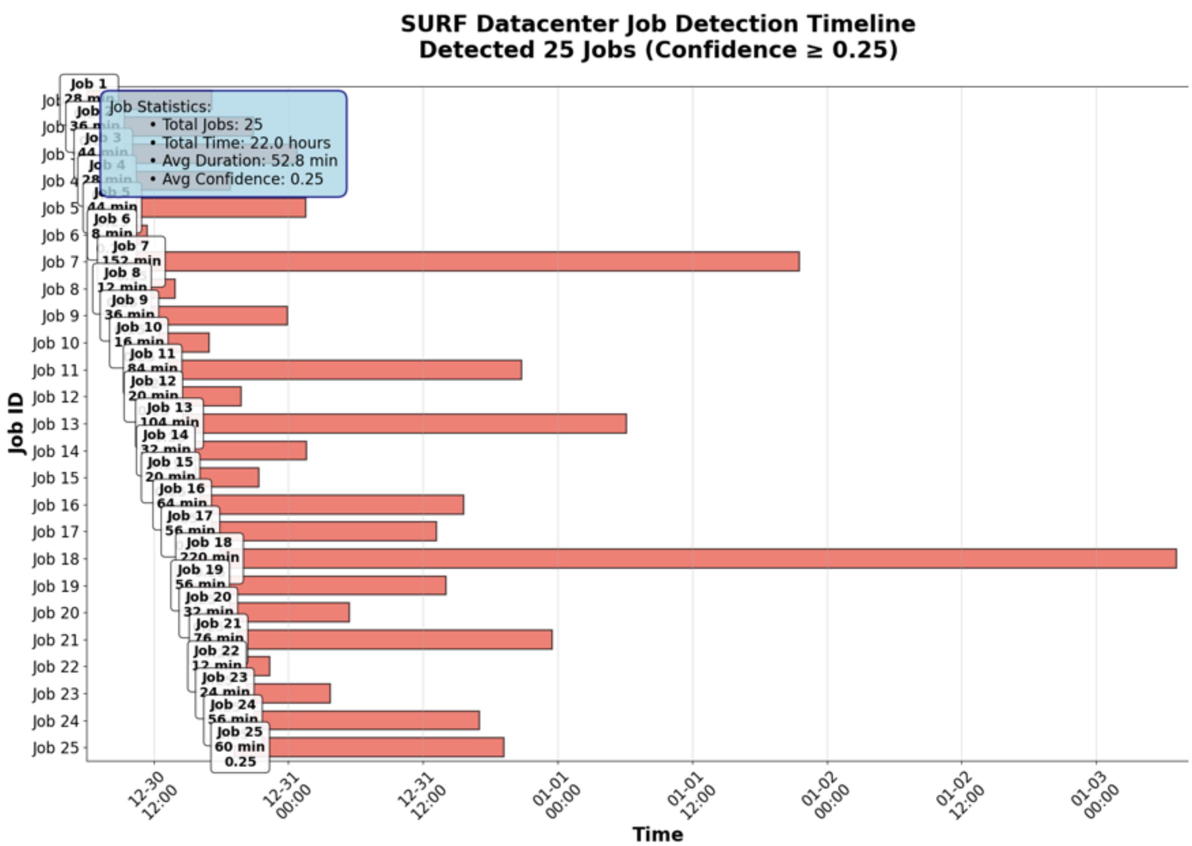

Figure 1: Estimated job detection timeline for 25 inferred jobs on SURF.

3.0 Overview of MIT Supercloud Dataset

3.1 Dataset Overview

The MIT Supercloud Dataset is a multi-source, anonymized collection covering ∼6 months of operations on the TX-Gaia cluster, including Slurm scheduler logs, per-job CPU and GPU time series, node-level monitoring snapshots, Lustre file-system activity, and facility (power/HVAC) sensors [4, 7].[2] Notable properties include:

- **Scheduler logs (anonymized):** submission/start/end times, resource requests/allocations, exit status.
- **Per-job time series:** CPU load, memory, I/O (10 s); GPU utilization, power, memory footprint, temperature (100 ms).
- **Node-level snapshots:** load, user/process counts, memory, Lustre RPCs (5 min).
- **Facility sensors:** total IT/HVAC power, aisle temperatures/pressure, environmental telemetry.
- **Labeled DNN jobs:** 3,425 GPU jobs mapped to model classes (ResNet, VGG, Inception, U-Net, BERT, etc.).

3.2 Power Profile Characteristics

Aggregate power shows a strong diurnal pattern (weekday peaks, off-hour troughs) with superimposed spikes during large training campaigns or synchronized batch starts. Weekly summaries reveal periods of heightened variability aligned with project milestones and experiment windows.

Figure 2 presents the daily average total power demand for the second week of August 2021, along with the corresponding daily maxima and minima. The results show no consistent pattern, power draw shifts markedly from one day to the next, driven by job submission activity. Figure 3 reports the maximum ramp-up and ramp-down rates, which also vary substantially over the week, highlighting the highly uncertain and rapidly changing nature of data center loads. This observation is again made in Fig. 4 over a period of two weeks.

3.3 Per-Job GPU/Memory Dynamics

GPU-intensive jobs commonly exhibit:

- **High, sustained GPU utilization** with large, constant GPU memory reservations (model weights/activations).
- **Thermal tracking:** GPU temperature co-varies with utilization and power draw; memory temperature closely tracks GPU temperature.
- **Burst vs. plateau patterns:** Iterative DL workloads can show bursty compute separated by I/O/synchronization valleys; mature training runs often show long plateaus.

For a representative job in the dataset, Fig. 5 shows GPU and memory utilization where we observe that the utilization vary over a wide range and the corresponding temperatures vary similarly. In the next subsection, we present the correlations between these key parameters with the power draw.

[2] GPU utilization is sampled at 100 ms; CPU/memory/I/O per-job series at 10 s; node-level snapshots at 5 min. See [4, 7] for full field dictionaries.

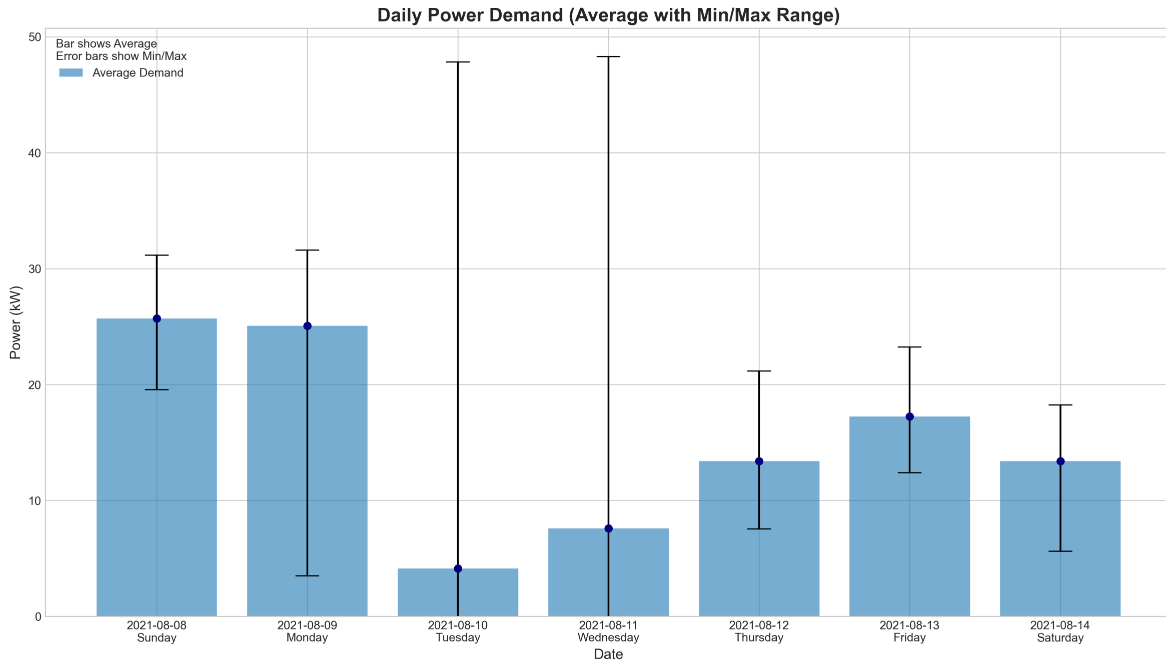

Figure 2: Daily average total power demand in the second week of Aug 2021

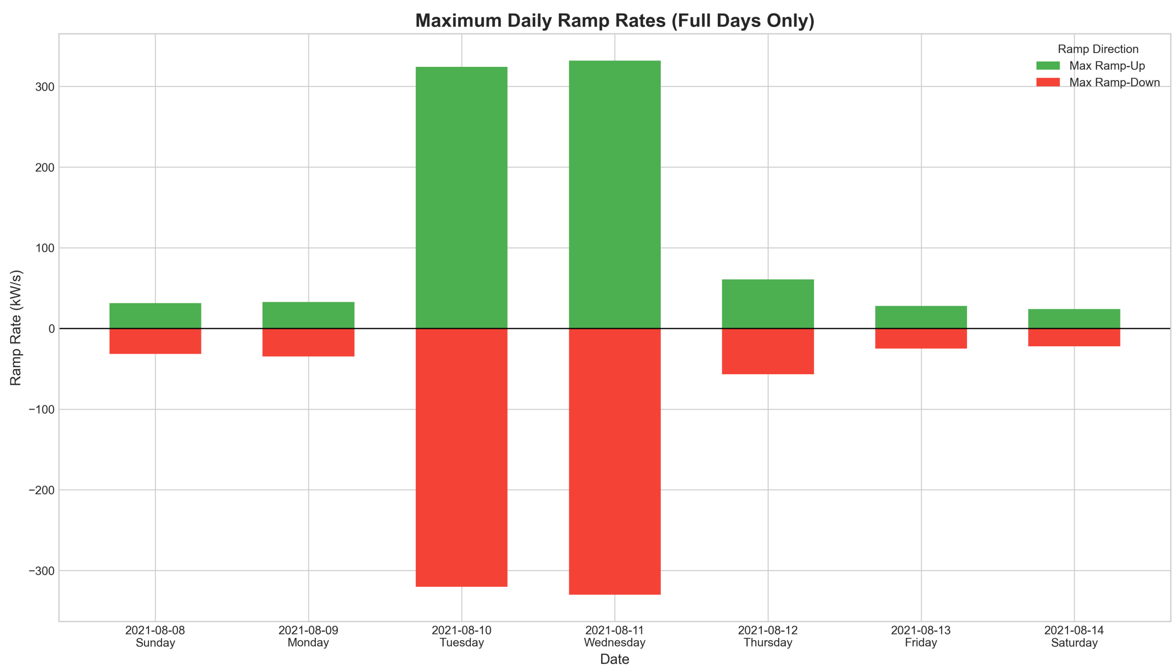

Figure 3: Daily metrics of total power demand in the second week of Aug 2021

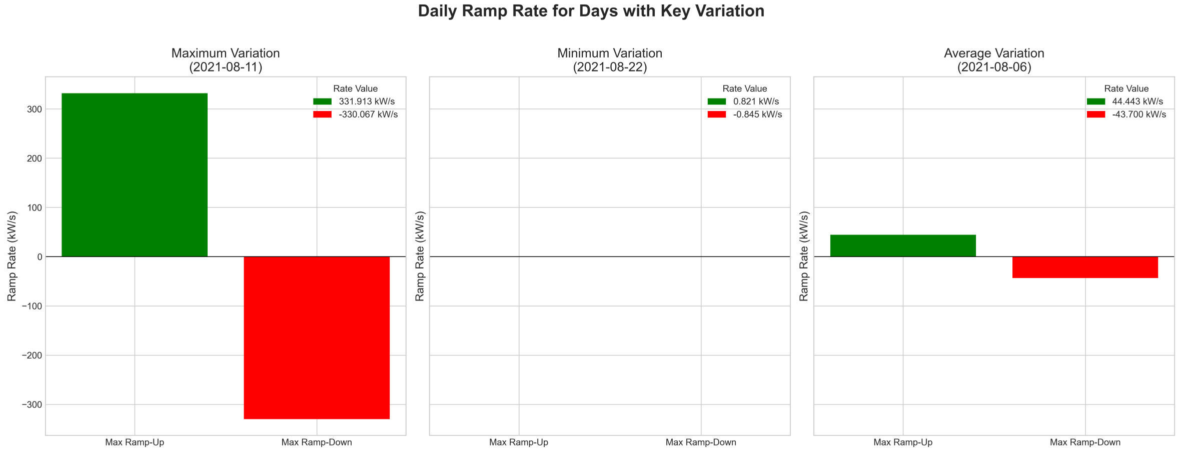

Figure 4: Weekly power demand variability

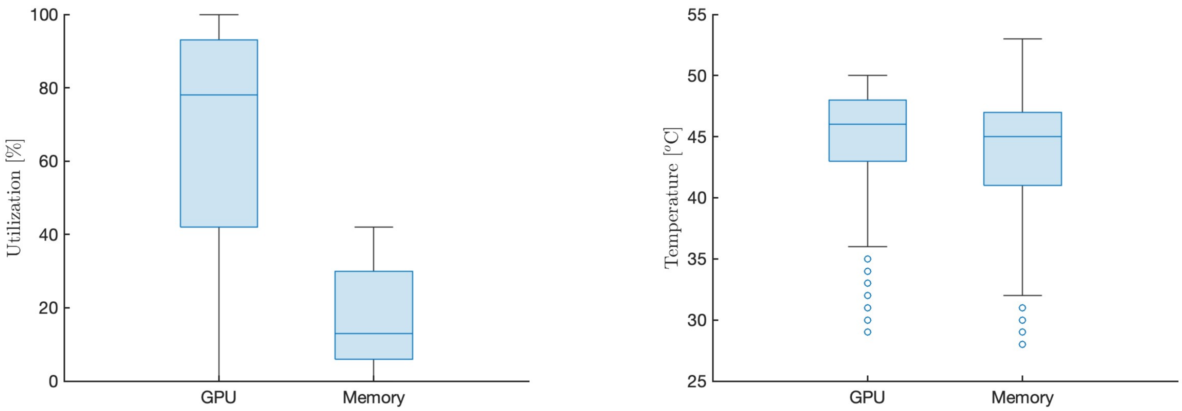

Figure 5: GPU/memory allocation, temperatures, and utilization over 150 min corresponding to a sample job

3.4 Correlation Structure

The correlation analysis between the power draw, GPU utilization, GPU temperature, Memory utilization and memory temperature reveal the following characteristics.

1. There is always a high correlation between
   - The GPU and Memory temperatures
   - The GPU and Memory utilization
2. The correlation between
   - The GPU temperature and GPU utilization is high for most jobs
   - The Memory temperature and Memory utilization is high for most jobs
3. The correlation between the power draw and temperatures is observed to be varying across different jobs. There is no consistent trend.
4. In most jobs, the power draw is highly correlated with GPU and memory utilization and is not

a function of the job length. Observe this difference between the left and right heatmap of Figure 6.

These relationships support using thermal and power sensors as proxies for instantaneous compute intensity in LEL detection.

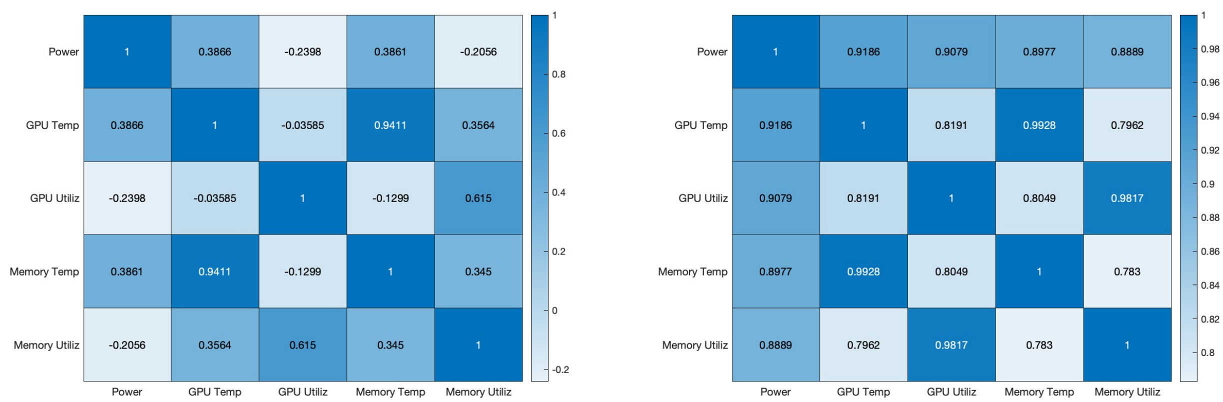

Figure 6: Correlation matrix among per-job GPU metrics (utilization, memory usage, temperatures, power): strong positive associations

3.5 Frequency-Domain Signatures

FFT analysis of GPU power draw reveals periodic components that could potentially threaten the stability of the operation of the power grid.

- **Low-frequency (diurnal/operational) components** reflecting daily load cycles and scheduled routines.
- **Mid/high-frequency components** linked to algorithmic iteration cadence and synchronization intervals in distributed training.

Figures 7 and 8 reveal the low frequency oscillations on the grid caused due to these LEL loads.

3.6 Implications for LEL Characterization

The multi-source MIT dataset enables alignment of job timelines with per-job power/thermal footprints and facility responses. Clear LEL episodes manifest as (i) high, sustained GPU/CPU utilization with elevated temperatures and power, (ii) correlated file-system load (Lustre RPC bursts during checkpoints), and (iii) facility-level power and HVAC adjustments. These signatures can inform predictive scheduling, thermal-aware job placement, and power-aware orchestration.

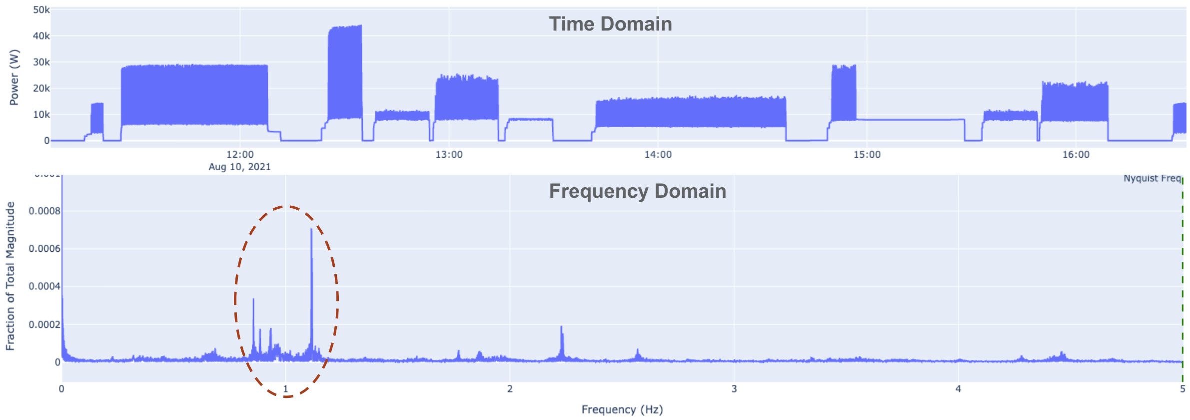

Figure 7: FFT of GPU power during peak load

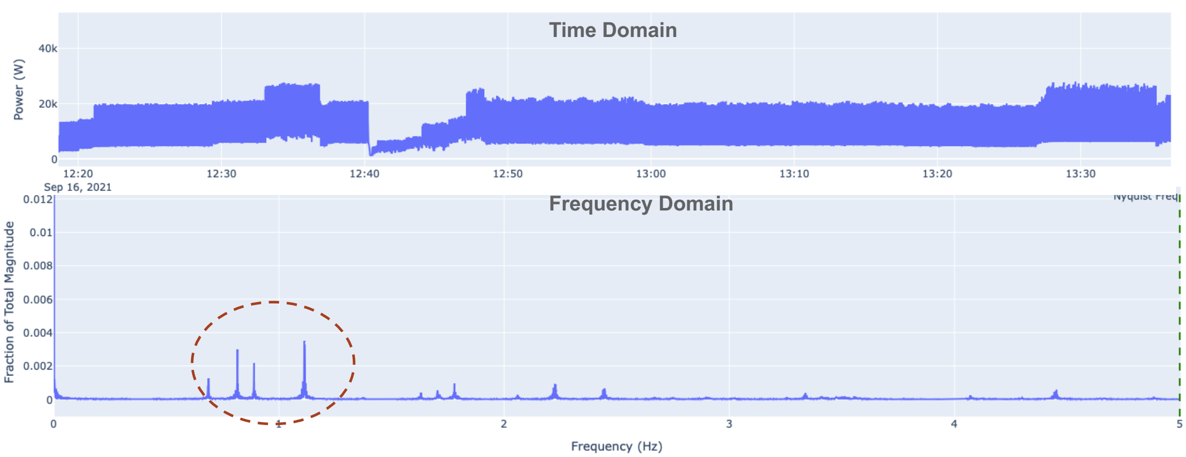

Figure 8: FFT of GPU power during medium load

4.0 Synthetic Model Development

4.1 Model Design and Training

To predict datacenter power consumption under Large Energy Load (LEL) scenarios, we developed a Long Short-Term Memory (LSTM) recurrent neural network model. The model was trained on preprocessed GPU job traces extracted from the MIT Supercloud dataset. Each training sample corresponds to a windowed segment of the power profile with aligned system metrics. Based on preliminary experiments with multiple architectures, the final model was selected using mean prediction error as the criterion.

The following design choices were made:

- **Input features:** GPU utilization, GPU memory usage, node load, and relevant resource allocation features from Slurm logs. Weakly relevant fields were pruned during preprocessing.
- **Sequence length:** Window size determined by the typical duration of GPU job traces (on the order of a few thousand samples).
- **Network architecture:** Multi-layer LSTM with hidden states sized to capture long-term dependencies, followed by fully connected dense layers for power prediction output.
- **Training strategy:** Supervised learning using mean squared error (MSE) loss. Validation sets were held out to monitor overfitting and guide hyperparameter selection.

The training was initially performed locally for prototyping, with the plan to migrate to the Deception cluster for large-scale training on thousands of jobs in the future work.

4.2 Validation Results

This section summarizes the model performance on the validation dataset. Key evaluation dimensions include:

- **Time-domain prediction:** Comparison of actual vs. predicted power traces shows the LSTM captures the main dynamics of workload-induced power fluctuations. (see Fig. 9)
- **Correlation analysis:** A strong correlation between predicted and actual power indicates the model learns the dominant load-response relationships. (see Fig. 10)
- **Frequency-domain comparison:** FFT analysis demonstrates that the predicted power preserves the key low-frequency components of the actual power profile, crucial for capturing daily/weekly load cycles. (see Fig. 11)
- **Error distribution:** The absolute error profile suggests prediction errors are concentrated during transition phases (job start/end), while steady-state loads are reproduced with lower error. (see Fig. 12)

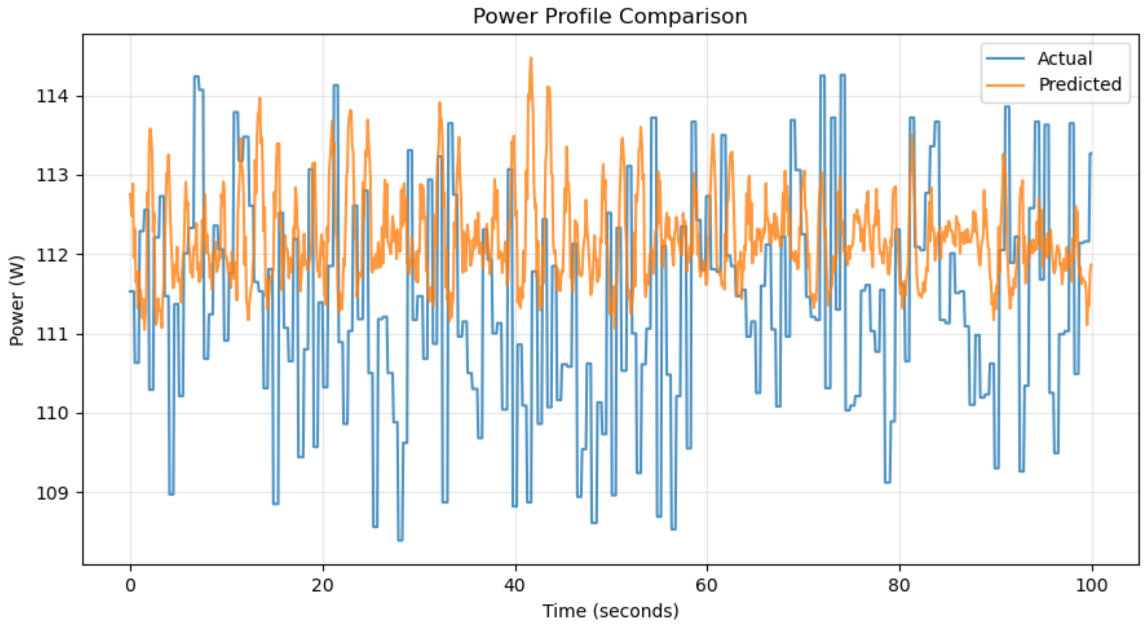

Figure 9: Actual vs. predicted power profile.

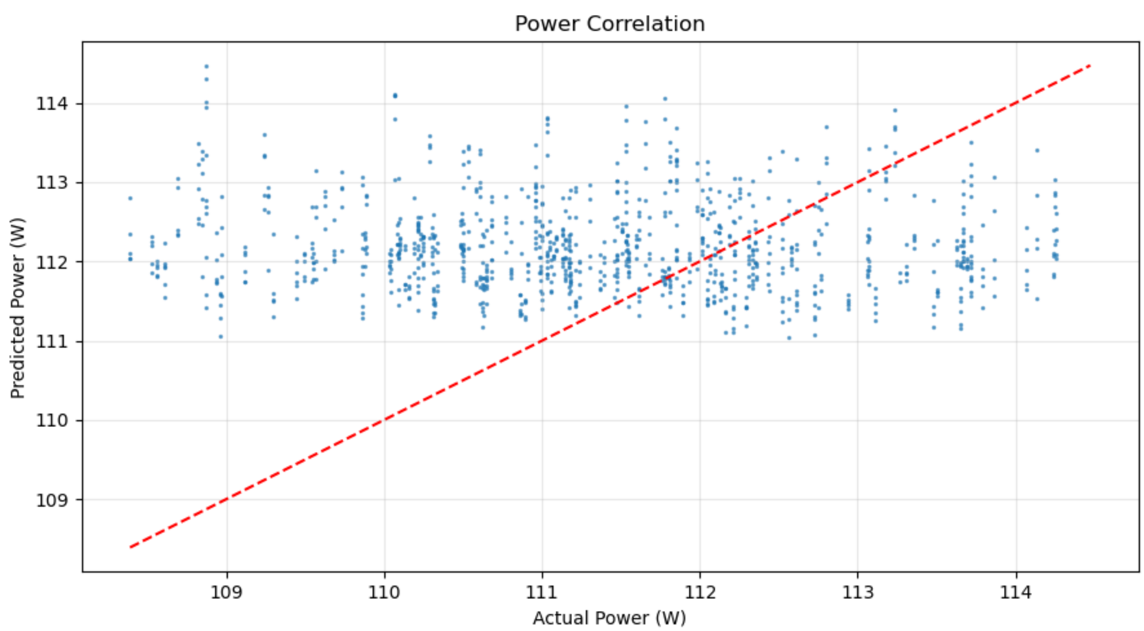

Figure 10: Power correlation between predicted and actual signals.

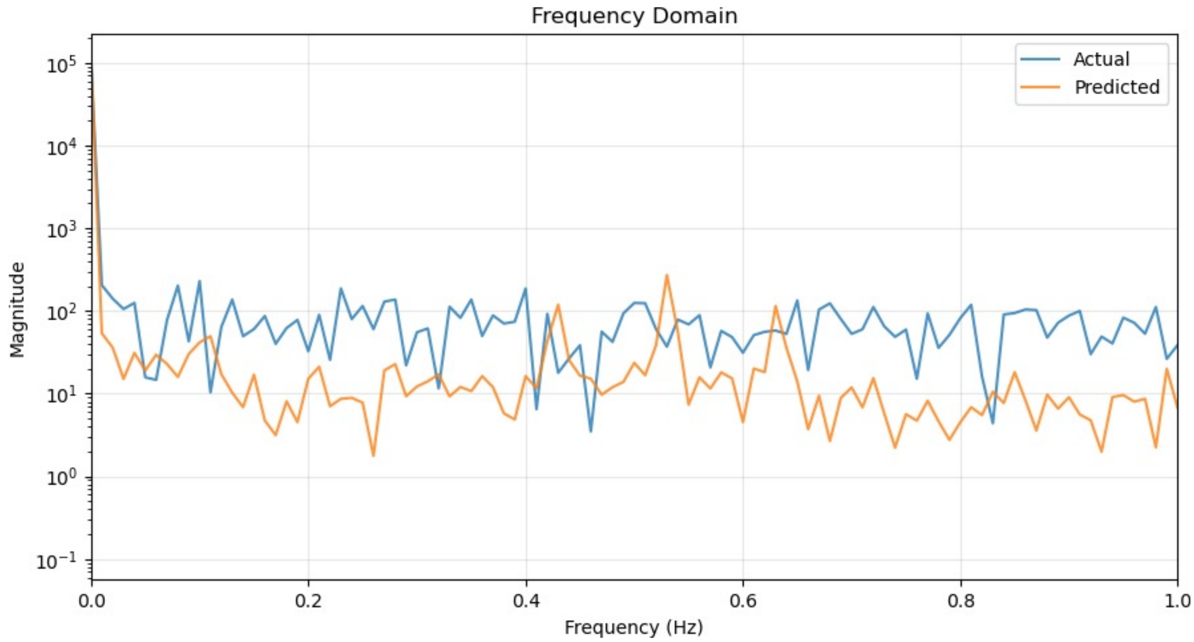

Figure 11: Frequency-domain comparison of actual vs. predicted power.

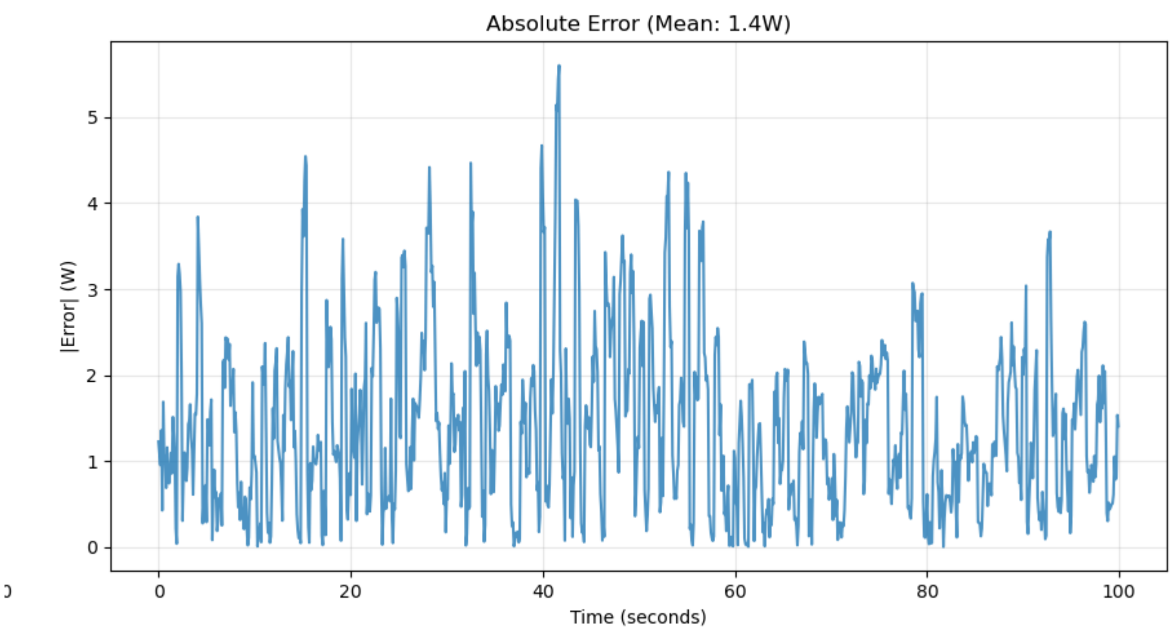

Figure 12: Absolute error distribution of predicted vs. actual power.

5.0 Conclusion

AI-driven data centers are rapidly expanding their footprint on the electric grid, bringing large, fast-changing loads that challenge stability and planning. In this report, we analyzed high-resolution, publicly available datasets (SURF and MIT SuperCloud) to characterize key operational features relevant to grid integration. Our analysis reveals pronounced day to day variability in power demand, significant ramp rates over short intervals, and strong relationships between GPU/memory utilization, thermal behavior, and power draw patterns consistent with workload-driven dynamics and scheduler activity.

We developed an initial synthetic load model that reproduces key aspects of data center power profiles, including steady-state behavior and frequency-domain characteristics. Model errors were found to be concentrated around job transitions (start/end), while steady-state loads were captured with lower error. These findings underscore the need to better encode workload-aware transitions, scheduler effects, and thermal/power management dynamics to accurately reflect short-term volatility.

This work demonstrates the value of high-fidelity synthetic datasets for grid studies, enabling assessment of worst-case ramps, stability risks, and operational strategies that align generation with data center demand patterns. The synthesis framework outlined here will be extended to fuse diverse open-source datasets, incorporate richer workload and control features, and undergo validation against real operational measurements. Future work will focus on fusing diverse open-source datasets to generate realistic scenarios with additional validation studies and extract key load characteristics for research on data center–grid interactions.

References

[1] J. Aljbour, T. Wilson, P. Patel, Powering intelligence: Analyzing artificial intelligence and data center energy consumption, EPRI White Paper no. 3002028905 (2024).

[2] A. Shehabi, A. Hubbard, A. Newkirk, N. Lei, M. A. B. Siddik, B. Holecek, J. Koomey, E. Masanet, D. Sartor, et al., 2024 united states data center energy usage report (2024).

[3] L. Versluis, M. Cetin, C. Greeven, K. Laursen, D. Podareanu, V. Codreanu, A. Uta, A. Iosup, A holistic analysis of datacenter operations: Resource usage, energy, and workload characterization–extended technical report, arXiv preprint arXiv:2107.11832 (2021).

[4] S. Samsi, M. L. Weiss, D. Bestor, B. Li, M. Jones, A. Reuther, D. Edelman, W. Arcand, C. Byun, J. Holodnack, M. Hubbell, J. Kepner, A. Klein, J. McDonald, A. Michaleas, P. Michaleas, L. Milechin, J. Mullen, C. Yee, B. Price, A. Prout, A. Rosa, A. Vanterpool, L. McEvoy, A. Cheng, D. Tiwari, V. Gadepally, The MIT supercloud dataset, in: 2021 IEEE High Performance Extreme Computing Conference (HPEC), 2021, pp. 1–8, dataset and docs at https://dcc.mit.edu. doi:10.1109/HPEC49654.2021.9622850.
URL http://arxiv.org/abs/2108.02037

[5] SURFsara / SURF, Surf machine metrics dataset (prometheus exports), https://zenodo.org/records/4459519, accessed 2025-09-25 (2021). doi:10.5281/zenodo.4459519.

[6] NLeSC, SURF, Surface: Tools and documentation for surf metrics dataset, https://github.com/sara-nl/SURFace, accessed 2025-09-25 (2021).

[7] MIT Supercloud Team, Mit supercloud dataset: Labeled subset readme, https://github.com/MIT-AI-Accelerator/MIT-Supercloud-Dataset, accessed 2025-09-25 (2021).

Pacific Northwest National Laboratory

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

1-888-375-PNNL (7665)

www.pnnl.gov