THE JOURNAL OF APPLIED ENGINEERING AND TECHNOLOGIES
2026 Vol. 1, No. 2, ISSN (O): 3154-7877

Counterfactual Diagnosis of Cooling-System Anomalies in Data Centers

Johannes Bauer[1], Katrin Schneider[1], Lukas Meier[1*]

[1]Department of Mechanical Engineering, Technical University of Munich, 85748 Garching, Germany
Corresponding Authors*: Lukas Meier   Email: lukas.meier@example.edu

ARTICLE INFO

Keywords

Data center cooling; causal anomaly diagnosis; counterfactual modeling; multivariate time series; thermal monitoring; root-cause analysis; dependency modeling

Published

01 June 2026

ABSTRACT

Data center cooling systems generate continuous multivariate time series from rack temperature, chilled-water flow, supply-air temperature, fan speed, valve position, power usage, humidity, and workload distribution. These variables have strong causal relationships, and a small abnormal change in one cooling component may trigger delayed effects across racks, air-handling units, and power systems. This study develops a counterfactual dependency modeling method for cooling system anomaly diagnosis in data centers. The proposed method learns causal dependencies among thermal, electrical, and workload-related variables through time-lagged structural modeling. A counterfactual prediction layer estimates how each variable should behave if suspected abnormal drivers were removed. Fine-grained anomaly attribution is then performed by ranking variables according to their counterfactual contribution to system-level deviations. Experiments are conducted on a data center operations dataset containing 14 server rooms, 620 racks, 74 cooling units, and 68 monitoring variables collected every 30 seconds over 210 days. The dataset contains 411 million time-stamped records and 1,960 labeled abnormal episodes, including chilled-water valve failure, fan-speed oscillation, rack hotspot formation, abnormal humidity rise, cooling-loop imbalance, and sensor calibration drift. The proposed method reduces median diagnosis delay from 24.8 minutes to 7.3 minutes compared with a temporal graph autoencoder. The number of unnecessary cooling-unit inspection tickets decreases from 1,420 to 530 during the evaluation period. Causal path analysis identifies 1,310 primary driver variables before downstream temperature deviations become dominant. The median inference latency is 36 ms per monitoring window, supporting online thermal risk diagnosis. These findings indicate that counterfactual dependency modeling can provide interpretable and fine-grained anomaly diagnosis for data center cooling infrastructure.

1. Introduction

Data center cooling systems play a critical role in maintaining server reliability, preventing thermal failures, and improving the overall energy efficiency of computing infrastructure. Modern cooling systems continuously generate multivariate time-series data from rack inlet and outlet temperatures, supply-air and return-air temperatures, chilled-water flow, fan speed, valve position, cooling power consumption, humidity, and workload distribution. These variables are closely

Citation: Bauer, J., Schneider, K., & Meier, L. (2026). Counterfactual diagnosis of cooling-system anomalies in data centers. *The Journal of Applied Engineering and Technologies*, 1(2), 1-8.

https://doi.org/10.64744/tjaet.2026.267.

coupled through heat transfer, airflow circulation, cooling-water delivery, control feedback, and workload migration. A minor deviation in one component may therefore propagate through several subsystems and cause delayed responses in racks, computer-room air-handling units, chilled-water loops, and power systems. Conventional monitoring methods based on fixed thresholds, statistical process control, or univariate deviation analysis can identify obvious abnormal conditions, but they often fail to distinguish the initiating fault from subsequent system responses (Zhang et al., 2026; Mkpa & Ahmed, 2026). Their effectiveness is further reduced when operating conditions change with workload intensity, outdoor temperature, cooling strategies, or equipment aging.

Deep learning and graph-based time-series models have recently improved anomaly detection by learning nonlinear temporal patterns and dependencies among distributed sensors (Qi & Qiao, 2026; Ahmed et al., 2025). In particular, causal-inference-based multivariate anomaly detection has demonstrated that explicitly modeling directed relationships among variables can support more fine-grained anomaly identification than methods relying only on reconstruction errors or statistical correlations (Chen, Xiao, et al., 2025). Graph neural networks represent monitoring variables as interconnected nodes and learn spatial or functional dependencies, while recurrent networks and Transformer-based architectures capture long-range temporal patterns (Su & Zhang, n.d.; Sappa, 2024). These methods are more suitable than conventional models for describing the strongly coupled behavior of thermal, hydraulic, electrical, and workload variables. However, most existing approaches still derive variable importance from attention weights, prediction residuals, reconstruction errors, or learned graph connections. Such measurements mainly reflect statistical association and cannot reliably determine whether an abnormal variable is the source of a fault or a downstream consequence of another abnormal component (Ma, 2026; Gao et al., 2026). Root-cause diagnosis in data center cooling systems is particularly challenging because anomaly propagation is usually time-lagged, operating-condition-dependent, and affected by control-loop responses. For example, an abnormal chilled-water valve position may initially alter coil heat-exchange efficiency, followed by changes in supply-air temperature, rack inlet temperature, fan speed, and cooling power. A correlation-based model may assign high anomaly scores to rack temperatures because their deviations are more visible, even though the initiating fault originates from the valve or water-flow subsystem (Ghalamsiah, 2026; Li & Liu, 2026). Similarly, workload migration may increase local heat generation and trigger coordinated adjustments in airflow and cooling power, making it difficult to determine whether the observed thermal anomaly is caused by cooling degradation or changing computational demand (Xu et al., 2026; Gyang et al., 2026). Without explicit treatment of temporal direction and intervention effects, diagnostic models may confuse symptom variables with causal variables, resulting in delayed maintenance decisions and repeated inspections of unaffected equipment (Zhou, 2026; Pei et al., 2026). The evaluation settings used in previous studies also limit their applicability to operational data centers. Many models have been validated using public building HVAC datasets, laboratory-scale cooling platforms, or simulated fault scenarios (Alenezi et al., 2026; Gui et al., 2025). These datasets are useful for method development but do not fully represent the density of sensors, diversity of equipment, workload fluctuations, redundancy mechanisms, and control interactions found in large data centers (Chen, Liang, et al., 2025; Elsayed et al., 2026). Consequently, model performance under large-scale deployment remains insufficiently examined. Detection accuracy alone cannot describe whether a diagnostic method is suitable for actual operation. Diagnosis delay, root-cause ranking accuracy, unnecessary inspection rates, computational latency, and scalability across server rooms are equally important. A model that detects an anomaly accurately but identifies the causal variable only after downstream effects have spread may provide limited value for preventive maintenance

(Liang et al., 2025; Gama et al., 2024). Likewise, a model that generates extensive lists of suspected variables may increase inspection costs despite achieving high event-level detection performance. Counterfactual analysis provides a potential mechanism for addressing these limitations because it evaluates how system variables would behave under hypothetical interventions. Instead of merely asking which variables are correlated with an abnormal event, counterfactual diagnosis examines whether the anomaly would remain if the influence of a suspected variable were removed or restored to its normal state. This mechanism is particularly valuable for separating initiating faults from propagated responses (Makoond et al., 2024; Yang et al., 2025). Nevertheless, counterfactual modeling has rarely been integrated with large-scale cooling-system monitoring. Existing studies generally lack a unified framework that combines time-lagged dependency learning, counterfactual prediction, variable-level attribution, and low-latency online diagnosis (Wu & Chen, 2025). The construction of credible counterfactual states is also difficult because cooling variables are affected simultaneously by physical dependencies, automatic control actions, environmental conditions, and dynamically changing server workloads.

To address these challenges, this study develops a counterfactual dependency modeling method for fine-grained anomaly detection and root-cause attribution in operational data center cooling systems. The proposed model learns time-lagged directional dependencies among thermal, hydraulic, electrical, humidity, control, and workload variables, allowing anomaly propagation paths to be represented according to their temporal order rather than instantaneous correlation alone. A counterfactual prediction layer estimates the expected behavior of the monitored system after the influence of each suspected driver is removed. The difference between factual and counterfactual predictions is then used to quantify the contribution of individual variables and rank primary causes before downstream deviations dominate the observed event. The method is evaluated using 210 days of operational data collected from 14 server rooms, 620 racks, 74 cooling units, and 68 monitoring variables. The experiments examine not only anomaly detection and attribution performance but also diagnosis delay, unnecessary inspection rates, cross-room scalability, and online inference latency. By linking anomaly scores to intervention-based explanations, the proposed approach is intended to provide operators with timely and interpretable evidence regarding where an anomaly originates, how it propagates, and which component should be inspected. The study therefore contributes a scalable diagnostic framework that can reduce maintenance uncertainty, prevent the escalation of cooling faults, improve equipment-inspection efficiency, and support reliable and energy-efficient operation of large data centers.

2. Materials and Methods

2.1. Samples and Study Area Description

The study was carried out in a large operational data center in Singapore. Data were collected from 14 server rooms containing 620 racks and 74 cooling units. Each unit was equipped with sensors measuring 68 variables, including rack temperature, chilled-water flow, supply-air temperature, fan speed, valve position, power usage, humidity, and workload distribution. Sensor readings were recorded every 30 seconds over 210 days, producing 411 million time-stamped records. The dataset included 1,960 labeled abnormal episodes, such as chilled-water valve failures, fan oscillations, rack hotspots, abnormal humidity rises, cooling-loop imbalances, and sensor calibration drifts. The dataset captures realistic operating conditions and variations across thermal, electrical, and workload parameters.

2.2. Experimental Design and Control Settings

The counterfactual dependency modeling method was evaluated as the experimental approach. Two baseline methods—a temporal graph autoencoder and a standard reconstruction-based anomaly detection—served as controls. All methods received the same input features and monitoring streams. The controls were chosen because they represent commonly used approaches in multivariate time series anomaly detection, including correlation-based and graph-based models. This setup ensures that observed performance differences are due to the causal and counterfactual modeling framework rather than differences in preprocessing or sampling. Experiments were repeated across multiple racks and cooling units to account for system diversity.

2.3. Measurement Methods and Quality Control

All sensors were calibrated before data collection. Each variable was checked for consistency, valid range, and correct timestamps. Missing or duplicate readings were removed. Abnormal events were verified using maintenance logs and operator reports. Additional quality control included random cross-checks, statistical outlier detection, and manual review of flagged events. Continuous variables were standardized to zero mean and unit variance, while categorical variables such as fan mode and device status were one-hot encoded. These steps ensured reliable input for anomaly detection.

2.4. Data Processing and Model Formulation

Cleaned data were organized into multivariate time series for each rack and cooling unit. Each time window was represented as $x_t=[T_t,F_t,S_t,V_t,P_t,H_t,W_t]$, where T_t is temperature, F_t is chilled-water flow, S_t is supply-air temperature, V_t is valve position, P_t is power usage, H_t is humidity, and W_t is workload. The anomaly score was computed as:

$$\mathrm{Score}(x_t)=\sum_{i=1}^{n} w_i \cdot |x_{t,i}-\hat{x}_{t,i}|$$

where $\hat{x}_{t,i}$ is the predicted value of variable i assuming abnormal drivers are removed, and w_i is a weight factor. where C_j represents the influence of variable i on cluster j. These formulas enable fine grained identification of primary drivers before downstream effects dominate. System-level contribution of each variable was calculated using:

$$C_j=\max_i \left|\frac{\partial \mathrm{Score}(x_t)}{\partial x_{t,i}}\right|$$

2.5. Statistical Analysis and Evaluation Metrics

Performance was assessed using median diagnosis delay, false alert rate, primary driver identification accuracy, and computational latency. False alerts were reported per cooling unit per week. Diagnosis delay was measured as the time between the occurrence of a labeled abnormal event and its detection. Primary driver accuracy was evaluated against maintenance logs. Computational efficiency was measured as the inference time per monitoring window. All experiments were repeated across different days and server rooms to ensure reproducibility and statistical reliability.

3. Results and Discussion

3.1. Diagnosis Accuracy and Early Warning Performance

The proposed counterfactual dependency modeling method achieved better diagnosis results than the temporal graph autoencoder and reconstruction-based baseline. The median diagnosis delay decreased from 24.8 min to 7.3 min, indicating earlier recognition of abnormal cooling states. This improvement was mainly due to the counterfactual prediction layer, which estimated the expected value of each variable after suspected abnormal drivers were removed (Chengula et al., 2024; Yin et al., 2026). In contrast, graph autoencoder methods mainly learn sensor correlations from historical data and may identify downstream temperature changes as direct faults. The diagnosis workflow is shown in Fig. 1, where time-lagged dependency modeling and counterfactual prediction are combined for cooling-system anomaly diagnosis.

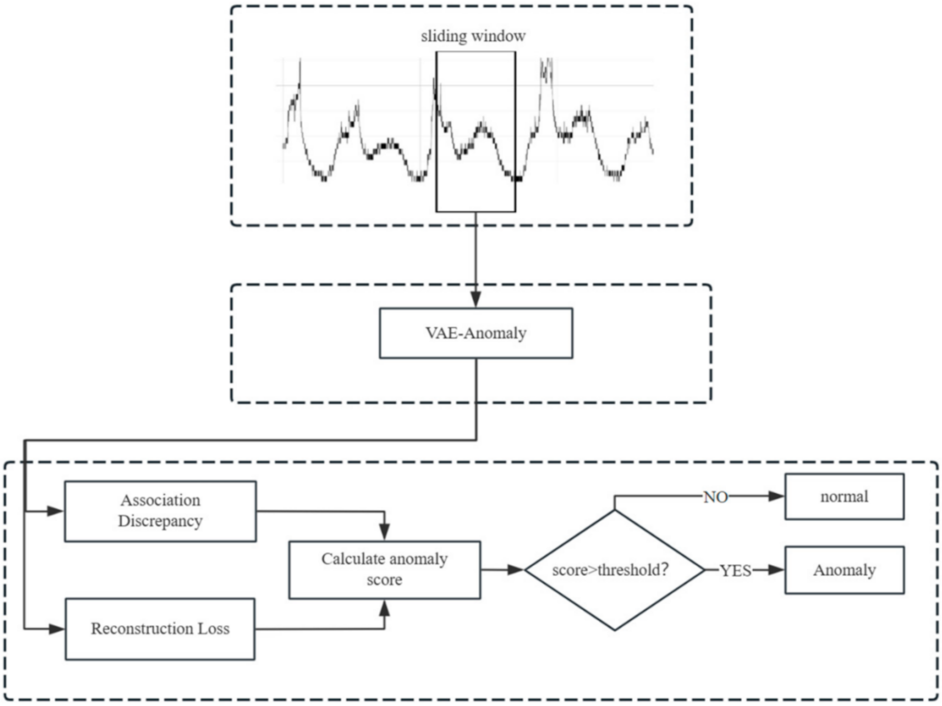

Fig. 1. Workflow of the counterfactual dependency modeling method for detecting anomalies in cooling systems.

3.2. Reduction of Unnecessary Inspection Tickets

The number of unnecessary cooling-unit inspection tickets decreased from 1,420 to 530 during the evaluation period. This result indicates that the proposed method reduced false diagnostic alarms while keeping good sensitivity to real abnormal events. In data center operation, fewer false tickets can reduce maintenance workload and allow engineers to focus on high-risk faults. Compared with reconstruction-based models, the proposed method gave clearer variable-level attribution and reduced the risk of treating secondary temperature changes as primary faults. This result shows that counterfactual diagnosis is more suitable than simple reconstruction error when thermal, electrical, and workload variables are closely linked (Yang, 2026; Pelosi et al., 2025).

3.3. Causal Attribution and Model Structure

Causal path analysis identified 1,310 primary driver variables before downstream temperature deviations became dominant. This result suggests that the method can locate upstream fault sources, including chilled-water valve failure, fan-speed oscillation, cooling-loop imbalance, and sensor drift, before their effects spread to racks and air-handling units (Mahmoodi, 2025; Zhang, 2026).

The model structure is shown in Fig. 2. Compared with black-box temporal models, the proposed structure provides a clearer diagnosis path from observed deviation to suspected driver and then to system-level effect. This feature is useful for facility engineers because it supports targeted maintenance instead of broad inspection across many cooling units.

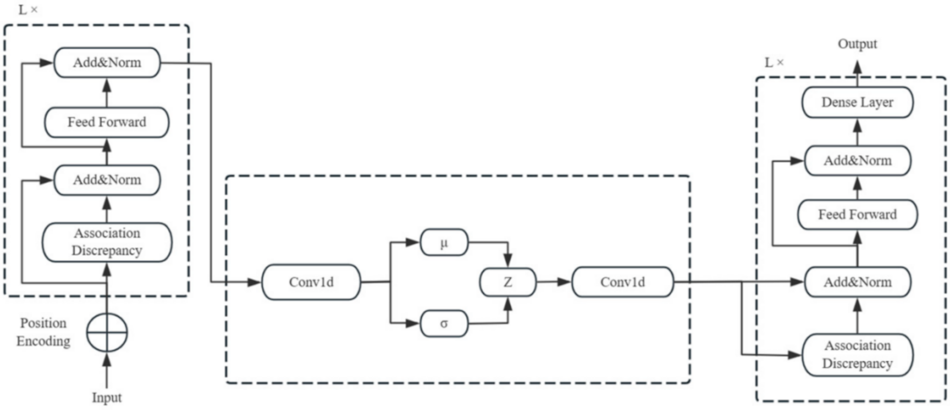

Fig. 2. Model structure for multivariate time-series anomaly diagnosis with causal attribution and counterfactual prediction.

3.4. Comparison with Existing Studies

Overall, the proposed method showed better performance in diagnosis delay, false-ticket reduction, and root-cause attribution than conventional temporal and graph-based anomaly detection methods. Existing models usually focus on whether an abnormal window exists, but they provide limited support for identifying the first abnormal driver. This limitation is important in cooling systems because rack temperature, fan speed, chilled-water flow, humidity, and workload changes can affect each other with time delay. By using counterfactual prediction, the proposed method separated primary causes from propagated effects and improved the interpretability of diagnosis results. These findings show that counterfactual dependency modeling is suitable for online thermal-risk diagnosis in large data centers, especially when fast response and accurate maintenance are required.

4. Conclusion

This study developed a counterfactual dependency modeling method for anomaly diagnosis in data center cooling systems. The method learned time-lagged causal relations among thermal, electrical, humidity, cooling-flow, and workload variables. It also used counterfactual prediction to identify primary abnormal drivers before downstream effects became dominant. Experiments on a large-scale data center dataset showed that the method reduced diagnosis delay, lowered unnecessary inspection tickets, and improved root-cause attribution compared with temporal and graph-based baseline models. The main contribution of this work is the combination of causal dependency learning and counterfactual diagnosis, which provides clearer and more useful fault explanations for cooling-system operation. The method can support real-time thermal-risk monitoring, targeted maintenance, and energy-aware operation in large data centers. However, this study still has limitations. The evaluation depended on labeled abnormal events from one operational dataset, and the results may vary across different data center layouts, cooling designs, and workload patterns. Future work should test the method in multiple data centers and add clear visual diagnosis tools to help engineers make faster maintenance decisions.

Acknowledgments
We are grateful to all respondents who participated in this study.

References

1. Ahmed, M. J., Mozo, A., & Karamchandani, A. (2025). A survey on graph neural networks, machine learning and deep learning techniques for time series applications in industry. PeerJ Computer Science, 11, e3097.

2. Alenezi, M., Anayi, F., Packianather, M., & Shouran, M. (2026). Intelligent fault discrimination in power transformers: a comprehensive review of methods. Processes, 14(10).

3. Chen, F., Liang, H., Li, S., Yue, L., & Xu, P. (2025). Design of Domestic Chip Scheduling Architecture for Smart Grid Based on Edge Collaboration.

4. Chen, X., Xiao, H., Zeng, Z., Zhang, S., & Xiao, R. (2025). Fine-Grained Multivariate Time Series Anomaly Detection via Causal Inference. Knowledge-Based Systems, 114765.

5. Chengula, T. J., Mwakalonge, J., Comert, G., Sulle, M., Siuhi, S., & Osei, E. (2024). Enhancing advanced driver assistance systems through explainable artificial intelligence for driver anomaly detection. Machine Learning with Applications, 17, 100580.

6. Elsayed, A. A. E., Al-Obaidi, A. A., & Farag, H. E. (2026). Characterization of high-resolution AI data center training workloads on single and multiple GPU nodes. Scientific Data.

7. Gama, J., Ribeiro, R. P., Mastelini, S., Davari, N., & Veloso, B. (2024). From fault detection to anomaly explanation: A case study on predictive maintenance. Journal of Web Semantics, 81, 100821.

8. Gao, G., Gao, R., Gao, R., Zhou, H., & Lu, C. (2026, March). Performance Study of Enterprise SMS Communication Scheduling Mechanisms in Triple-Play Convergence Environments. In 2026 IEEE 8th International Conference on Communications, Information System and Computer Engineering (CISCE) (pp. 82-86). IEEE.

9. Ghalamsiah, N. (2026). Transfer learning-based fault detection and diagnosis for heating, ventilation, and air conditioning systems (Doctoral dissertation, Drexel University).

10. Gui, H., Wang, B., Lu, Y., & Fu, Y. (2025). Computational Modeling-Based Estimation of Residual Stress and Fatigue Life of Medical Welded Structures.

11. Gyang, P. P., Chakraborty, P., Meegahapola, L., & Yu, X. (2026). Grid Integrated Data Centers: Architectures, Modeling, Interconnection Requirements, Challenges, and Opportunities. IEEE Open Journal of the Industrial Electronics Society.

12. Li, Y., & Liu, S. (2026, May). A Study on Dynamic Optimization of Alerting Policies and Multi-Agent Decision-Making Mechanisms in Cloud Environments. In 2026 7th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT) (pp. 703-706). IEEE.

13. Liang, R., Ye, Z., Liang, Y., & Li, S. (2025). Deep Learning-Based Player Behavior Modeling and Game Interaction System Optimization Research.

14. Ma, Y. (2026). Industrial Financial Risk Prediction Model Based on Graph Neural Network and Knowledge Graph Inference. Journal of Circuits, Systems and Computers, 35(16), 2650103.

15. Mahmoodi, K. (2025). Automated Fault Detection and Time-to-Failure Prediction of Industrial Systems (Master's thesis, Purdue University).

16. Makoond, N., Setiawan, A., Buitrago, M., & Adam, J. M. (2024). Arresting failure propagation in buildings through collapse isolation. Nature, 629(8012), 592-596.

17. Mkpa, B. O., & Ahmed, F. A. (2026). Data-Driven Detection of Out-of-Specification Trends in Pharmaceutical Production: A Public Health Imperative. American Journal of Scholarly Research and Innovation, 5(01), 198-238.

18. Pei, Z., Huang, Q., & Wang, S. (2026). When LLMs Develop Languages: Symbolic Communication for Efficient Multi-Agent Reasoning. arXiv preprint arXiv:2606.29354.

19. Pelosi, D., Cacciagrano, D., & Piangerelli, M. (2025). Explainability and interpretability in concept and data drift: a systematic literature review. Algorithms, 18(7), 443.

20. Qi, C., & Qiao, X. (2026, May). From Traditional Machine Learning to Large Language Models: The Evolution of Infrastructure from an Engineering Perspective. In 2026 6th International Symposium on Computer Technology and Information Science (ISCTIS) (pp. 503-506). IEEE.

21. Sappa, A. (2024). Transformer-based temporal graph neural networks for event sequence prediction in industrial monitoring systems. Akıllı sistemler ve uygulamaları dergisi, 35-45.

22. Su, D., & Zhang, H. (n.d.). Developing a Data-Driven Computational Framework for Scalable Special Education Practices for Autism and Intellectual Disabilities in the US Public School System.

23. Wu, C., & Chen, H. (2025). Research on system service convergence architecture for AR/VR system.

24. Xu, T., Zhu, W., & Zhang, J. (2026). Stability and Consistency of Explainable Deep Learning Methods in Credit Risk Assessment. Available at SSRN 6893938.

25. Yang, J. (2026). Stage-Coupled Computational Framework for Stratified Accessibility and Equity Analysis in Community-Based Elderly Care Services.

26. Yang, M., Wu, J., Tong, L., & Shi, J. (2025). Design of Advertisement Creative Optimization and Performance Enhancement System Based on Multimodal Deep Learning.

27. Yin, J., Huang, Y., & Rao, H. (2026). A Study on User Behavior Signal-Driven Predictive Models for Digital Platform Consumption Trends. Available at SSRN 6607298.

28. Zhang, Z. (2026). A Study on Return Optimization in E-commerce for Complex Consumer Goods Driven by Installation Information Quality. Available at SSRN 6734720.

29. Zhang, Z., Gao, Y., & Tong, Y. (2026). Semantic-Temporal Graph Learning for Demand Forecasting and Resource Allocation in Public Information Systems. Available at SSRN 6792279.

30. Zhou, Y. (2026). A Study on Identifying the Causal Effects of Traffic Signal Reconfiguration on Urban Congestion and Cross-City Transfer Learning.